



Computer Networking Essentials

An essential guide to understanding networking
theory, implementation, and interoperability



Computer Networking Essentials

Debra Littlejohn Shinder

Cisco Press

Cisco Press
201 West 103rd Street
Indianapolis, IN 46290 USA

Computer Networking Essentials

Debra Littlejohn Shinder

Copyright © 2002 Cisco Systems, Inc.

Published by:

Cisco Press

201 West 103rd Street

Indianapolis, IN 46290 USA

All rights reserved. No part of this book may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or by any information storage and retrieval system, without written permission from the publisher, except for the inclusion of brief quotations in a review.

Printed in the United States of America 3 4 5 6 7 8 9 0

Third Printing January 2002

Library of Congress Cataloging-in-Publication Number: 2001090429

ISBN: 1-58713-038-6

Trademark Acknowledgments

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

Warning and Disclaimer

This book is designed to provide information about basic networking and operating system technologies. Every effort has been made to make this book as complete and as accurate as possible, but no warranty or fitness is implied.

The information is provided on an “as is” basis. The author, Cisco Press, and Cisco Systems, Inc., shall have neither liability nor responsibility to any person or entity with respect to any loss or damages arising from the information contained in this book or from the use of the discs or programs that may accompany it.

The opinions expressed in this book belong to the author and are not necessarily those of Cisco Systems, Inc.

Feedback Information

At Cisco Press, our goal is to create in-depth technical books of the highest quality and value. Each book is crafted with care and precision, undergoing rigorous development that involves the unique expertise of members from the professional technical community.

Readers’ feedback is a natural continuation of this process. If you have any comments regarding how we could improve the quality of this book or otherwise alter it to better suit your needs, you can contact us through e-mail at feedback@ciscopress.com. Please make sure to include the book title and ISBN in your message.

We greatly appreciate your assistance.

Publisher
Editor-in-Chief
Executive Editor
Cisco Systems Management

Production Manager
Development Editor
Senior Editor
Copy Editor
Technical Editor
Reviewers

Associate Editor
Cover Designer
Composition
Indexer

John Wait
 John Kane
 Carl Lindholm
 Michael Hakkert
 Tom Geitner
 William Warren
 Patrick Kanouse
 Kitty Wilson Jarrett
 Jennifer Chisholm
 Jill Batistick
 Dr. Thomas W. Shinder
 Lynn Bloomer
 Wayne Jarvimaki
 Michael R. Hanson
 Shannon Gross
 Louisa Klucznik
 Steve Gifford
 Tim Wright

CISCO SYSTEMS



Corporate Headquarters
 Cisco Systems, Inc.
 170 West Tasman Drive
 San Jose, CA 95134-1706
 USA
<http://www.cisco.com>
 Tel: 408 526-4000
 800 553-NETS (6387)
 Fax: 408 526-4100

European Headquarters
 Cisco Systems Europe
 11 Rue Camille Desmoulins
 92782 Issy-les-Moulineaux
 Cedex 9
 France
<http://www-europe.cisco.com>
 Tel: 33 1 58 04 60 00
 Fax: 33 1 58 04 61 00

Americas Headquarters
 Cisco Systems, Inc.
 170 West Tasman Drive
 San Jose, CA 95134-1706
 USA
<http://www.cisco.com>
 Tel: 408 526-7660
 Fax: 408 527-0883

Asia Pacific Headquarters
 Cisco Systems Australia, Pty.,
 Ltd
 Level 17, 99 Walker Street
 North Sydney
 NSW 2059 Australia
<http://www.cisco.com>
 Tel: +61 2 8448 7100
 Fax: +61 2 9957 4350

Cisco Systems has more than 200 offices in the following countries. Addresses, phone numbers, and fax numbers are listed on the Cisco Web site at www.cisco.com/go/offices

Argentina • Australia • Austria • Belgium • Brazil • Bulgaria • Canada • Chile • China • Colombia • Costa Rica • Croatia • Czech Republic • Denmark • Dubai, UAE • Finland • France • Germany • Greece • Hong Kong • Hungary • India • Indonesia • Ireland • Israel • Italy • Japan • Korea • Luxembourg • Malaysia • Mexico • The Netherlands • New Zealand • Norway • Peru • Philippines • Poland • Portugal • Puerto Rico • Romania • Russia • Saudi Arabia • Scotland • Singapore • Slovakia • Slovenia • South Africa • Spain • Sweden • Switzerland • Taiwan • Thailand • Turkey • Ukraine • United Kingdom • United States • Venezuela • Vietnam • Zimbabwe

Copyright © 2001, Cisco Systems, Inc. All rights reserved. Access Registrar, AccessPath, Are You Ready, ATM Director, Browse with Me, CCDA, CCDE, CCDP, CCIE, CCNA, CCNP, CCS1, CD-PAC, *CiscoLink*, the Cisco NetWorks logo, the Cisco Powered Network logo, Cisco Systems Networking Academy, Fast Step, FireRunner, Follow Me Browsing, FormShare, GigaStack, IGX, Intelligence in the Optical Core, Internet Quotient, IP/VC, iQ Breakthrough, iQ Expertise, iQ FastTrack, iQuick Study, iQ Readiness Scorecard, The iQ Logo, Kernel Proxy, MGX, Natural Network Viewer, Network Registrar, the Networkers logo, *Packet*, PIX, Point and Click Internetworking, Policy Builder, RateMUX, ReyMaster, ReyView, ScriptShare, Secure Script, Shop with Me, SlideCast, SMARTnet, SVX, TrafficDirector, TransPath, VlanDirector, Voice LAN, Wavelength Router, Workgroup Director, and Workgroup Stack are trademarks of Cisco Systems, Inc.; Changing the Way We Work, Live, Play, and Learn, Empowering the Internet Generation, are service marks of Cisco Systems, Inc.; and Aironet, ASIST, BPX, Catalyst, Cisco, the Cisco Certified Internetwork Expert Logo, Cisco IOS, the Cisco IOS logo, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Collision Free, Enterprise/Solver, EtherChannel, EtherSwitch, FastHub, FastLink, FastPAD, IOS, IP/TV, IPX, LightStream, LightSwitch, MICA, NetRanger, Post-Routing, Pre-Routing, Registrar, StrataView Plus, Stratum, SwitchProbe, TeleRouter, are registered trademarks of Cisco Systems, Inc. or its affiliates in the U.S. and certain other countries.

All other brands, names, or trademarks mentioned in this document or Web site are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0010R)

Introduction

Computer Networking Essentials helps you understand the fundamentals of computer networking concepts and implementation and introduces you to the client and server operating systems that run on networked PCs.

Concepts covered in this book include the history of networking, networking terminology, networking theory and established standards, and implementation of local-area and wide-area networks. Special emphasis is placed on understanding network protocols and how they operate at all layers of the networking model. Emphasis also is placed on the interoperability of networks that run on multiple protocols, platforms, and operating systems.

Specialty areas such as security, remote access, virtual private networking, thin client networking, monitoring, management, and troubleshooting are covered thoroughly. Emerging technologies that are expected to impact the future of networking are also introduced.

Who Should Read This Book

This book's primary audience is professionals who are beginning training in the networking industry and those who need a review of basic concepts.

The secondary audience includes corporate training faculties and staff and members of the business world who work with information technology personnel and require a broad overview of the concepts involved in networking from the small business to the enterprise-level corporation.

A third target audience is the general user who wants to know more about how computers communicate over networks. The book's approach is designed to be user-friendly and accessible to the non-technical reader who is overwhelmed by the jargon found in vendor documentation and technical manuals.

This Book's Organization

This book is organized into four parts and includes 19 chapters, an appendix, and a glossary. The following sections describe the contents of each part of the book.

Part I: Introduction to Networking Concepts

Chapter 1, "Introduction to PC Networking," introduces you to the basic concepts of PC networking by providing a brief history of electronic communications and networking and a summary of where PC networking is today.

Chapter 2, "Categorizing Networks," discusses the categorization of networks according to physical scope, administrative model, network operating system, protocols in use, topology, and architecture.

Chapter 3, "Networking Concepts, Models, and Standards," provides an overview of binary communications and introduces two popular networking models: the Department of Defense (DoD) model on which the TCP/IP protocols are based and the Open Systems Interconnection (OSI) model, which was developed by the International Organization for Standardization (ISO). Specifications set forth by the Institute of Electrical and Electronics Engineers (IEEE) and vendor-specific models are also covered.

Chapter 4, "Networking Communications Methods," discusses signaling methods and provides an understanding of analog, digital, broadband, baseband, asynchronous, synchronous, simplex, duplex, and multiplexed signaling. Media access methods are described, including CSMA/CD, CSMA/CA, token passing, and demand priority.

Chapter 5, "LAN Links," discusses popular LAN types, including Ethernet, Token Ring, FDDI, AppleTalk, and ARCnet.

Chapter 6, "WAN Links," provides an overview of WAN connections such as PSTN, ISDN, t-carriers, Frame Relay, X.25, and CATV network, as well as high-speed connectivity solutions such as ATM, SONET, and SMDS. This

chapter also covers LAN-to-WAN connection solutions, including Internet Connection Sharing (ICS), Network Address Translation (NAT), proxy servers, and routed connections.

Part II: Networking Hardware and Software

Chapter 7, “Physical Components of the Network,” introduces students to the many types of networking media, including coax, twisted-pair cable, and fiber-optic cable, as well as to wireless technologies such as laser, infrared, radio, and satellite/microwave communications. Connectivity devices such as repeaters, hubs, bridges, routers, and switches are also discussed.

Chapter 8, “Networking Protocols and Services,” describes common LAN protocols—TCP/IP, NetBEUI, IPX/SPX—and discusses the OSI protocol suite. PPP and SLIP, which are WAN link protocols, and PPTP and L2TP, which are common tunneling protocols, are also presented.

Chapter 9, “The Widest Area Network: The Global Internet,” discusses the evolution of the Internet, the protocols used for Internet communications—HTTP, FTP, NNTP, SMTP, and POP—and the TCP/IP protocol suite.

Chapter 10, “Network Operating Systems,” discusses general network administration practices and then looks at the specifics of common server operating systems, including Windows NT, Windows 2000, NetWare, UNIX, and Linux.

Chapter 11, “Directory Services,” describes the Directory Services Protocol (DAP) and the Lightweight Directory Access Protocol (LDAP), as well as the X.500 standards developed by the ISO to promote directory services compatibility and interoperability. Novell’s NDS, Microsoft’s Active Directory, and Banyan VINES’ StreetTalk directory services are covered in some depth.

Chapter 12, “Desktop Operating Systems,” looks at the client side of the client/server network and discusses the advantages and disadvantages of common desktop clients, such as DOS, Windows, Linux, Macintosh, and OS/2, and how each can be integrated into popular NOS environments.

Chapter 13, “Hybrid Networks,” provides information about interoperability solutions and protocol gateways that allow PCs running different operating systems, protocols, and platforms to communicate with one another. This chapter also looks at PC-to-mainframe communications using Systems Network Architecture (SNA) solutions.

Part III: Network Specialty Areas

Chapter 14, “Protecting the Network,” addresses security issues and provides an overview of basic cryptography concepts, public and private key encryption, certificate services, firewalls and proxies, and internal security measures such as “smart cards” and advanced authentication technologies. It also provides guidance for developing security policies for your network. The second half of the chapter discusses disaster recovery plans, including implementation of disk fault tolerance (or RAID), regular scheduled backups, and server clustering.

Chapter 15, “Remote Access,” discusses methods of connecting to a server from a remote location using remote connectivity devices such as modems, ISDN terminal adapters, and customer premises equipment (CPE) for dedicated lines. Dial-in server configuration and special security considerations are also covered.

Chapter 16, “Virtual Private Networking,” provides an overview of VPN concepts and discusses the tunneling protocols used to provide VPN security.

Chapter 17, “Thin Client Networking,” discusses Network Computers, Net PCs, and Windows-based terminals. Windows terminal services, Citrix Metaframe, web-based computing, the X Window system and Java virtual machines—and the role each plays in thin client networking—are also discussed.

Chapter 18, “Monitoring, Management, and Troubleshooting Tools,” presents an introduction to the TCP/IP utilities and other tools built into the various operating systems. This chapter also examines commercial products such as Sniffer Pro, LANalyzer, Microsoft’s Systems Management Server, Novell’s ManageWise, and IBM’s Tivoli.

Part IV: The Future of Networking

Chapter 19, “Tomorrow’s Technologies,” takes a look into the future of PC networking. It discusses ways of overcoming the current limits of IP, including the new version of IP—IPv6. The goal of universal connectivity is addressed, and more exotic possibilities such as artificial intelligence, quantum computing, and cybernetic life forms are presented as possible components of tomorrow’s networks.

This Book’s Features

This book contains several elements that help you learn about operating systems and networking:

- **Figures, listings, and tables**—This book contains figures, listings, and tables that help to explain concepts, commands, and procedural sequences. Diagrams illustrate network layouts and processes, and screenshots assist students in visualization configuration procedures. In addition, listings and tables provide summaries and comparisons of features and characteristics.
- **Author’s notes, tips, sidebars, and cautions**—These elements are included to provide you with extra information on a subject. You will probably find these asides to be very beneficial in real-world implementations.
- **Chapter summaries**—At the end of each chapter is a summary of the concepts covered in the chapter, which provides a synopsis of the chapter and can serve as a study aid.
- **Further Reading**—Each chapter includes a list of resources for additional information about the topics covered in the chapter, including website URLs and books and articles that cover the topic in more detail.
- **Review questions**—After the Further Reading section in each chapter are 10 review questions that serve as an end-of-chapter assessment. The questions are designed to reinforce the concepts introduced in the chapter and to help students evaluate their understanding before moving on to the next chapter.

The conventions used to present command syntax in this book are the same conventions used in the *Cisco IOS Command Reference*, as follows:

- **Boldface** indicates commands and keywords that are entered literally as shown. In examples (not syntax), boldface indicates user input (for example, a **show** command).
- *Italics* indicates arguments for which you supply values.
- Square brackets [] indicate optional elements.
- Vertical bars (|) separate alternative, mutually exclusive elements.
- Braces and vertical bars within square brackets—for example, [x {y | z}]—indicate a required choice within an optional element. You do not need to enter what is in the brackets, but if you do, you have some required choices in the braces.

Introduction to PC Networking

Welcome to the world of personal computer (PC) networking. In this world, it is no longer enough to simply have and use PCs; today it is imperative that you also “get connected.” The real power and usability of PCs becomes apparent only when they are linked so that they can communicate with one another. From the simple two-computer home or small office local-area network (LAN) to the ever-growing global Internet, networking is *the future* of computing, and that future is here today.

In many areas of the United States, the demand for trained networking professionals far exceeds the supply. According to projections of the U.S. Department of Labor, computer networking as an occupation has a bright future. Businesses and individuals are buying PCs, and those computers are linking within LANs and wide-area networks (WANs) at an astonishing pace. We literally are “networking the world.”

Because network communications is quickly becoming a part of our lives, even those not directly involved in the information technology (IT) industry should know something about the basics of networking. Just as it would be difficult to function in today’s world if you knew nothing about a telephone and its features, in the not-too-distant future, knowing how to “get on the network” will be a requirement for many individuals, both at work and at home.

A Brief History of PC Networking

The desire to communicate with others is a driving force among human beings, and the sophisticated means we have developed to communicate sets us apart from other species. From the moment it became possible to link two computers and get them to talk to one another, the concept of the Internet was inevitable.

In the early days of computing, computers were enormous machines that filled entire rooms—sometimes entire city blocks—and cost hundreds of thousands of dollars. Although these expensive behemoths had less processing power and memory than today’s tiny handheld computers, they were state-of-the-art technology in the 1950s and 1960s. In a world in which human beings who were slow and prone to error had done calculations manually, the capabilities of the computer were amazing.

At the midpoint of the twentieth century, computers were still rare, exotic, mysterious machines owned only by large companies, governmental bodies, and educational institutions. For the most part, computers were standalone systems, isolated from one another.

In the 1940s, Thomas Watson, the chairman of IBM, said that a market existed in the world for approximately five computers. Even as recently as 1977, Ken Olson, president of Digital Equipment Corporation, said, "There is no reason anyone would want a computer in their home" (ISBC [International Small Business Consortium], www.isbc.com/isbc/business/wisdom.cfm). Of course, both have been proven not just wrong, but *very* wrong. However, no one would have predicted, even a decade ago, that PCs would proliferate as they have or that computer networking would become a mainstream topic.

The First Communications Networks

By the mid-1900s, electronic communications had been around for over a century and was being implemented in both Europe and the United States. These early networks took many forms and sent only coded signals. They later became capable of sending voice across the wire.

This section provides a rough time line of how the first networks were developed.

Telegraph Cables

In the early 1800s, the French developed the first optical telegraph network, which sent information at the blazing speed of 20 characters per second, and Samuel Morse demonstrated the electrical telegraph, which spurred the development of networked communications in the United States.

The Telephone Network: Circuit-Switching Technology

In the late 1800s, a vast telephone network began to be built. Technology leaders of the day, however, were no more farsighted than those of the early computer age. In 1876 an internal memo at Western Union stated that "This 'telephone' has too many shortcomings to be seriously considered as a means of communication. The device is inherently of no value to us" (www.isbc.com/isbc/business/wisdom.cfm).

Despite that attitude, there were more than 50,000 telephone lines in the U.S. by 1880, and by 1960, telephone lines covered urban areas, and the telephone network became a global communications network.

A telephone system uses circuit-switching technology, in which a circuit, or virtual pathway, is established when one telephone connects to another on a network. This works

well for voice transmission because the sounds being transferred over the wire flow at a relatively constant rate.

In a circuit-switched network, a connection is established, as shown in Figure 1-1. All signals are passed over this circuit for the duration of the session. If you disconnect and reconnect, a different circuit can be used, as represented by the dotted line.

Figure 1-1 *In a circuit-switched network, a connection is established, as represented by the solid line.*



The technology works less well for transfer of computer data, which has a tendency to be sent in bursts; that is, periods of high activity are interspersed with intervals of low activity or inactivity.

Packet-Switching Technology

During the 1960s, the U.S. government became interested in developing a computer network that would enable systems at military installations and major educational institutions to communicate with one another. Because this was during the middle of the Cold War, they wanted the network to have robustness, reliability, and redundancy so that the network would survive a nuclear war.

Researchers working at the Massachusetts Institute of Technology (MIT), the RAND Institute, and the National Physical Laboratory (NPL) in England invented a new technology called *packet switching*, which worked better for bursty transmissions than did the traditional circuit-switching technologies. Their work created a foundation for the communications technology used on the Internet today.

In a packet-switched network, as shown in Figure 1-2, a connection is not established for the entire transmission. Instead, each individual packet of data can take a different path.

Communications from different sources can share the same line, rather than the line being dedicated to one end-to-end communication for the duration of a session, as is the case with circuit switching.

Figure 1-2 *Networked computers share data, software, and hardware resources.*



Circuit Switching Versus Packet Switching

The terms *circuit switching* and *packet switching* sound alike but have different meanings.

The public telephone system, sometimes referred to as POTS (plain old telephone service), is a switched-circuit communications network. When you place a telephone call in this type of network, only one physical path from your telephone to the one you're dialing is used for the duration of that call. This pathway, or *circuit*, is maintained for your exclusive use, until you end the connection by hanging up your telephone.

Note, however, that if you call the same friend at the same number tomorrow, and do so at the same location from which you placed today's call, the path is not necessarily the same. That's why the circuit is referred to as *switched*. It also explains why you can get a clear connection one day and noise and static on another.

With a packet-switching network, no dedicated pathway or circuit is established. Packet switching is sometimes referred to as a *connectionless* technology because of the lack of a dedicated pathway. If you transfer data, such as a word processing file, from your computer to another using a packet-switched network, each individual *packet* (that is, each small chunk of data) can take a different route. Although it all arrives at the same destination, it doesn't all travel the same path to get there. Internet traffic generally uses packet-switching technology.

The difference between circuit and packet switching can be compared to the different ways in which a large group of people traveling from Dallas to San Francisco can reach their destination. For example, circuit switching is similar to loading the entire group on a bus, a train, or an airplane. The route is plotted out, and the whole group travels over that same route.

Packet switching is like having each person travel in an automobile. The group is broken down into individual components as the data communication is broken into packets. Some travelers can take interstate highways, and others can use back roads. Some can drive straight through, and others can take a more roundabout path. Eventually, they all end up at the same destination. The group is put back together, just as packets are reassembled at the endpoint of the communication.

The ARPANet

The first packet-switched computer network was conceived in the late 1960s, under the auspices of the U.S. Department of Defense (DoD). It was christened the ARPANet (for Advanced Research Projects Agency network). The ARPANet's first *node*, or connection point, was installed at the University of California at Los Angeles in 1969. In just three years, the network spread across the United States, and two years after that, it spread to Europe.

As the network grew, it split into two parts. The military called its part of the internetwork *Milnet*, and ARPANet continued to be used to describe the part of the network that connected research and university sites. In the 1980s the Defense Data Network (a separate military network) and NSFNet (a network of scientific and academic sites funded by the National Science Foundation) replaced ARPANet. Eventually this WAN grew into what we today call the Internet.

Yesterday's Networks

Computer networking didn't begin on such a large scale as the ARPANet project; that is, the LAN came before the WAN. As computers became less expensive and more powerful, businesses of all sizes more commonly used them. Although the first machines were useful for only very limited types of data processing, as software development flourished, new programs enabled users to do much more than just collect and sort data.

With early mainframe systems, for instance, multiple users could access the same stored data from *terminals*, which were stations with input and output devices (for example, keyboards and monitors). These stations had no computing power of their own; they were points from which the mainframe computer could be accessed.

Using mainframes worked well in many respects, but they had several disadvantages when compared to smaller computers (then called microcomputers). Expense was one

disadvantage; large mainframe systems cost far more than the so-called “personal” computers designed to sit on a desktop and function independently.

Another disadvantage of mainframes was the *single point of failure* concept. With mainframe computing, if the computer was down, it was down for everyone. Nobody could access data, and nobody who depended on the computer could get any work done. The use of individual PCs, on the other hand, circumvented this problem.

PCs were full-fledged computers that ran programs and performed tasks entirely on their own. They also provided some measure of *fault tolerance*, which is the capability of a system to continue to function and ensure data integrity when failures occur. If one employee’s computer crashed, it didn’t affect the capability of the rest of the employees, who had their own PCs, to continue working. In fact, if an employee had saved data to a floppy disk, he or she could move to a functioning machine and continue working.

These factors contributed to the increased popularity of PCs as a computing solution for small and large businesses (and everything in between). However, once everyone had a PC on the desktop, companies were faced with a dilemma: How could workers share information as they had with the old mainframe computing model? The solution was networking.

Disadvantages of Standalone Systems

In the early days of desktop PCs, networking hardware and software were not readily available, and many businesses used the machines as standalone systems. If all users needed to print documents on occasion, there were three possible ways to provide that ability:

- A printer could be attached to each machine. This was a costly solution because it necessitated buying multiple printers, even though it was unlikely that they all would be in use at the same time.
- The file to be printed could be saved to floppy disk and transferred to a machine that had an attached printer. This was a less-expensive option, but it was an inconvenience both to the person who had to go begging for a printer and to the person with the printer, whose work was interrupted while someone else used his or her machine to print.
- A printer could be moved from one workstation to another, depending on who needed to print. This was a somewhat cumbersome solution; nonetheless, it was widely implemented, using rolling printer carts that were wheeled around the office. Each move required that cables be disconnected and reconnected, and sometimes, a move involved software reconfiguration as well.


High cost, inconvenience, and extra work are the primary disadvantages of standalone, or non-networked, solutions.

What Is a Network, Anyway?

The *American Heritage Dictionary* defines a *network* as “a system of lines or channels that cross or interconnect.” Earlier we mentioned the telegraph and telephone networks, and of course, we’ve all heard references to the television networks. Using the dictionary definition, we can call even the state highway system, or the railways that crisscross the country, a network.

That being said, what is a computer network? Simply, it is two or more devices linked for the purpose of sharing information, resources, or both. The link can be through cable (coaxial, twisted-pair, or fiber optics, as you’ll learn later in this chapter), or it can be a wireless connection that uses radio signals, laser or infrared technology, or satellite transmission. The information and resources shared can be data files, application programs, printers, modems, or other hardware devices. See Figure 1-3 for an illustration.

Figure 1-3 *This time line shows significant events in PC networking history.*



Why Network Computers?

If the advantage of PCs were each user having an independent computer, why would we want to turn around and link them again? We link them because networked PCs give us, in many ways, the best of both worlds. Each user has independent processing power, but still can enjoy all the benefits of sharing. On the other hand, a company sees a significant cost savings when expensive, occasionally used peripherals are shared over the network. For example, an expensive color laser printer might be used only for special projects, yet many

different members of the organization will need to print to it from time to time. With network access, it's easy for them to do so.

Benefits of Getting Connected

Many business owners and managers state that the primary reason for networking their PCs was the need to share printers, as described in the earlier example. Of course, once the systems were linked, people discovered the usefulness of being able to share much more than printers.

The cost involved in linking computers in a LAN—the network interface cards (NICs) for the computers, the cabling or wireless media, the hubs and other connectivity devices—often pays for itself many times over by reducing expenditures and lost production time.

Sharing Output Devices

As discussed, printers and other output devices can be shared on a network, saving time, money, and a great deal of aggravation. Items that can be shared include plotters, which are devices used to draw diagrams, and charts. They also include line-based graphics devices that use pens or electrostatic charges and toner. Fax machines, which can be either input or output devices, also are easy to share.

Sharing Input Devices

You can share scanners, digital cameras, and other input devices across the network. Because these devices, even more so than printers, are generally used on an occasional basis and are often relatively expensive, it makes sense to configure them for multiple users on the network.

Sharing Storage Devices

Networked computers can share the use of hard disks and floppy and CD-ROM drives. With this type of sharing, you can save files to the disk of another computer across the network if you run out of hard disk space on your computer. In addition, if your computer doesn't have a CD-ROM drive installed, you can access the shared CD drive of another computer. This ability to share also occurs with Zip and Jaz drives, magneto-optical drives, tape drives, and just about any other type of storage device that can be connected to a PC.

Sharing Modems and Internet Connections

Another important feature of networking is the ability of networked PCs to share modems, ISDN lines, cable modems, and DSL adapters. With the proper software—such as proxy or

Network Address Translation (NAT) software, which we discuss in detail in Chapter 9, “The Widest Area Network: The Global Internet”—an entire LAN can connect to the Internet through one phone line and a single ISP account.

Sharing Data and Applications

Hardware devices are not the only, or even the most important, resources that can be shared on a network. Data files and application programs also can be made available to multiple users. This sharing results in the efficient use of disk space and easier collaboration on multiuser projects. For example, if several managers need to access and revise a spreadsheet containing a department’s budget, the file can be stored in a central location. After each manager makes the desired changes, the file can be saved to the network location so that the updated version is available for the next manager.

Application programs, such as word processing programs, can be installed to a network server. Users can connect to the share and run the application on their own machines, without using space on their local hard disks for the program files.

Be aware that software vendors’ licensing agreements can require that you purchase additional licenses for each workstation that uses a network application, even though only one copy is actually installed and all users are accessing that same copy.

The Birth of the Internet

As mentioned previously, back in the 1960s, usable networking technologies became available, and in the early 1970s, the ARPAnet was created by a collaborative effort between the U.S. government (primarily the DoD) and several large universities.

The Role of the DoD

As the Cold War between the United States and the Soviet Union intensified in the 1960s, the DoD recognized the need to establish communications links between major U.S. military installations. The primary motivation was to maintain communications if a nuclear war resulted in mass destruction and breakdown of traditional communications channels. Major universities, such as the University of California and MIT, were already involved in networking projects too.

The DoD funded research sites throughout the United States, and in 1968, ARPA contracted with BBN, a private company, to build a network based on the packet-switching technology that had been developed for better transmission of computer data.

The 1970s: The Growth Spurt Begins

When the ARPAnet project began, no one anticipated that the network would grow to the extent it did. Throughout the 1970s, more nodes were added, both domestically and abroad.

The 1980s: More Is Better

In 1983, the ARPAnet network was split, and 68 of the 113 existing nodes were taken by Milnet, which was integrated with the Defense Data Network. The Defense Data Network had been created the previous year.

The Domain Name System (DNS) was introduced in 1984, providing a way to map “friendly” host names to IP addresses that was much more efficient and convenient than previous methods. We discuss these previous methods in Chapter 8, “Networking Protocols and Services.” In 1984, there were more than 1000 host computers on the network.

During the last half of the 1980s, the networking picked up considerably. For instance, the NSF created supercomputer centers at Princeton, in Pittsburgh, at the University of California at San Diego, at the University of Illinois at Urbana-Champaign, and at Cornell. The Internet Engineering Task Force (IETF) also came into being during this time. By 1987, there were 10,000 hosts on the network, and by 1989, that number increased to over 100,000.

The 1990s: The Net Becomes Big Business

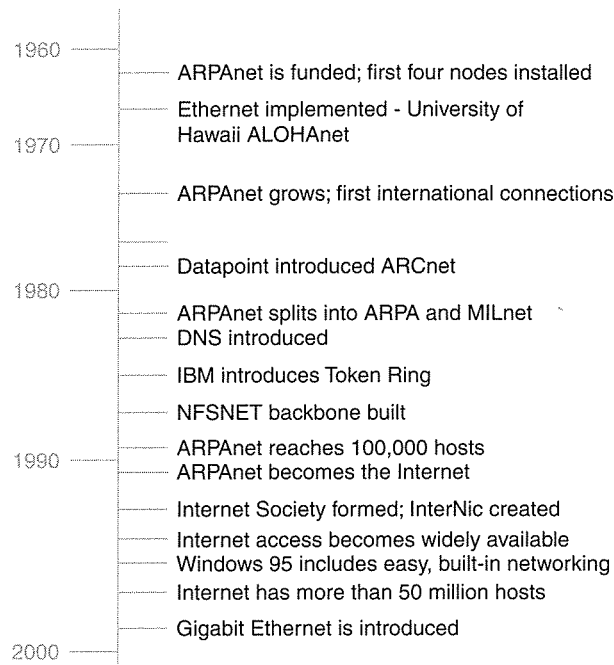
The phenomenal growth rate of the 1980s was nothing compared to what came in the 1990s. ARPAnet ceased to exist, and the Internet was “invented,” with the U.S. government getting involved in pushing the development of the so-called information superhighway. The NSFnet backbone was upgraded to T3 speed (that is, 44.736 Mbps), and in 1991 it sent more than 1 trillion bytes per month. The Internet Society (ISOC) was formed, and in 1992 more than 1 million hosts existed on the Internet.

The 1990s was the decade that the Internet went commercial. As more and more college students and faculty, individual home users, and companies of all sizes got connected, the business world recognized the opportunity to reach a large and expanding affluent market. By 1995, online advertising had caught on, online banking had arrived, and you could even order a pizza over the Internet.

The last half of the last decade of the century ushered in new major developments almost on a daily basis. Streaming audio and video, “push” technologies, and Java and ActiveX scripting took advantage of higher-performance connectivity available at lower and lower prices. Domain names became big business, with particularly desirable names selling for upwards of \$1 million. In December 1999, almost 1 billion sites existed on the World Wide Web, with well over 50 million host computers participating in this great linking.

Figure 1-4 shows a time line of significant events in PC networking history.

Figure 1-4 *The global network began in the 1960s and continues to grow today.*



The Cost of Technology: More and More for Less and Less

As computer and networking technology have advanced over the past few decades, the cost of that increasingly sophisticated technology has fallen dramatically. Those falling prices are at least partially responsible for the rising popularity of connectivity solutions in the business world and in personal lives.

In the 1970s and 1980s, a PC that was considered state of the art for the time cost several thousand dollars. Online services existed, but with fees of \$25 or more *per hour* of access, only big businesses and the wealthy could afford them. PC veterans still can remember the announcement of Prodigy's "bargain rates" of only \$9.95 an hour for online access. This was at blazing speeds of 1200 or 2400 baud.

Today, of course, for under \$1000, you can buy a computer system capable of doing much more, and doing it better and faster, than the \$500,000 mainframe of 20 years ago. Internet access at speeds equivalent to T1 is available through DSL or cable modem for \$30 to \$40 per month, and the price is falling all the time. Basic Internet access at 56 kbps can be had

for much less—even for free, if you can tolerate a bit of advertising taking up space on your screen.

PC Networking Today

As we enter the 2000s, we are on our way to networking the world (and beyond). The Net is beginning to permeate almost every area of our lives. We have computers at work, computers at home, and portable computers that we carry with us on airplanes and to the beach.

As always, where there are multiple computers, networking usually follows. Indeed, a primary function of many of these computers is to connect to the Internet. In this section, we look at some of the ways in which networked computers are changing our lives.

Home Computing

Home PCs are commonplace, and many of these PCs are being marketed specifically as e-machines, ready to connect to the Internet. Many households have multiple computers, and where two or more computers exist under the same roof, the desire to link them is sure to rear its head sooner or later. There is a new, booming market in home LAN technology, which uses wireless solutions or the house's telephone or electrical wiring in place of traditional Ethernet cabling.

Web Presence and E-commerce

Businesses of all sizes and types are finding that having a Web site is beneficial—or even essential—to their advertising strategies. Even small companies are turning to e-commerce (that is, selling their products or services directly over the Internet) as a cost-effective solution. Large corporations are running *enterprise networks* (that is, large multisite networks) that connect offices around the world to their own internal intranets as well as to the Internet.

High-Performance Business Solutions

As high-speed connectivity over fiber-optic and other fast media becomes commonplace, live videoconferencing is becoming a viable replacement for face-to-face meetings. Executives are staying in close touch with their offices even while traveling, thanks to the availability of remote dial-in access and virtual private networking. Employees at all levels of the organization are telecommuting, enjoying more flexibility in their work schedules while the company benefits from savings in facilities and onsite equipment. In addition, transfer of large data files can now be accomplished quickly and efficiently. Even traditional

low-bandwidth activities such as scheduling and e-mail access are improved by emerging and affordable high-performance technologies.

Online Learning

Public schools are getting wired, and online learning is becoming a reality as major colleges and universities offer credit courses that can be completed either partially or wholly over the Internet. In addition, hardware and software vendors, led by big players such as Microsoft Corporation and Cisco Systems, are partnering with commercial and nonprofit organizations to bring networked computers and Internet access within the reach of almost everyone who wants it.

Tomorrow's Networks

At the end of this book, we look at emerging technologies in the networking field. Some of these are still on the drawing board, and others have been tested and found feasible in the lab or in limited scope in the field.

As political entities and large corporations get behind the push to develop new, better, and faster means of bringing networked communications into the daily lives of people at all socioeconomic levels, exciting new developments in both software and hardware are being announced on a regular basis. The way we work and play is being transformed, and optimists predict that international and cultural barriers will melt as global connectivity makes the world a smaller—or at least more accessible—place. Many people seek to establish shared technology centers that provide access and training to groups that traditionally have not had the opportunity to benefit from computer technology and connectivity.

Speculation about where networking is headed is an attempt to predict the unpredictable. Twenty years ago, many of the technologies we now take for granted were unthinkable to anyone but science fiction writers. It's likely that twenty years in the future, computing and networking will have gone in drastically new directions that we can't begin to imagine today. There are, however, some interesting possibilities on the horizon, as discussed in the following sections.

'Smart' Appliances and Homes

Although not yet widely available, the concept of smart appliances and homes is already a reality. You can buy kitchen appliances today that have embedded microprocessors, or miniature computers, that control temperature, cooking time, and so forth. The next step—and it's not a very big one—is connecting those tiny computers to a network so that you can issue commands remotely. It's likely that in the not-too-distant future, you will get online

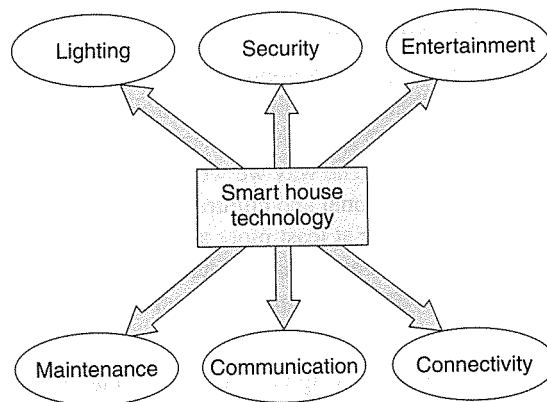
at the office, access your oven, and tell it what time to start preheating in preparation for your arrival home.

This concept logically leads to the next: the integrated smart house, which is computer-controlled and wired to a central network. The prime example of this is the multimillion-dollar estate of Microsoft chairman Bill Gates. This estate features computerized climate control and customized music “zones” that follow a person from room to room.

Numerous companies can make your present home smart by wiring it for Ethernet connectivity and by integrating your home computers with your phone, security, home theater, heating and air conditioning, and lighting systems. You can control these systems with the click of a mouse or even through voice command. Figure 1-5 illustrates some features currently available in smart house technology.

We discuss the house of the future a bit more in Chapter 19, “Tomorrow’s Technologies.”

Figure 1-5 *Smart house technology brings the network of the future home today.*



Phoning Home

Telephone technology has been quietly advancing over the past few decades. The developments in this area, however, have been marketed and integrated into our lives so smoothly and gradually that their impact has not received the sort of attention it might have otherwise.

Providers of mobile phones have built cellular- and satellite-based networks that now offer sophisticated services, such as wireless Internet access, at affordable prices. The technology is improving all the time.

Regular telephone companies (referred to within the industry as *telcos*), spurred by competition from cable TV companies, are implementing high-speed services for data

transfer, such as Asynchronous Digital Subscriber Line (ADSL), at a cost low enough to market to home users.

Telephony applications are integrating the PC and the telephone in a variety of ways that allow for automatic answering, message storage, and message retrieval through the computer.

NOTE

Telephony applications are those that combine telecommunications technologies with computing technologies. This includes implementations in accordance with the Telephony Applications Programmers Interface (TAPI) specifications and the standards of the Internet and Telecoms Convergence Consortium (ITC).

Internet phone programs using voice over IP enable you to bypass the telephone lines entirely (if you have an Internet connection through cable, wireless, or some other medium) to make long distance calls from your computer without paying long distance charges.

New operating systems, such as the latest version of Microsoft Windows, include support for IP telephony, which blends voice, video, and data transmission over TCP/IP connections with improved quality of service.

We take a closer look at personal communications systems of the future in Chapter 19.

The Wired Workplace

It seems that in any business facility you enter, there is a computer (or several) on every desktop, all of them linked to an internal LAN, the external Internet, and a few remote private networks for good measure. We can only wonder how the work environment of the future will be *more* network-centric than it is today.

The office of tomorrow will be even more reliant on network technology. As business environments are wired with fiber-optic and other high-performance media, increasing bandwidth and real-time video will make teleconferencing an attractive alternative to face-to-face meetings. It's likely that more and more employees will work from home. They will not, however, be isolated with their own little individual projects. Instead, they will share documents over the network to foster a more team-oriented environment.

As full-featured computers shrink in size while growing in capability, it will become easier to "take the office on the road." Handheld systems that today integrate e-mail, Web access, calendaring, and task management will no doubt in the future provide for on-the-go video on demand, voice communications, and notification services—combining the functions now performed by separate devices. No longer will you need to carry a palm computer,

mobile phone, and pager; one machine, small enough to wear on your belt or slip into a purse, will do it all.

This universal connectivity and accessibility have the potential to increase the productivity of businesses of all types.

Schools of the Future

The future of education will be impacted by developments in networking technology. In fact, it is already commonplace for libraries from elementary to university level to include one or more Internet-connected computers for students.

The most dramatic effect of technology on schools is the availability of information on a scale that was never before possible. As more public and private schools get wired, the way in which students do research for papers will change. In fact, the very nature of those assignments can change as well, to require the inclusion of multimedia material.

New learning methodologies that involve more of the senses are likely to become more popular. At the lower age levels, this can mean educational interactive games that engage the attention of students and contribute to the development of hand-eye coordination.

At the college and university levels, offsite learning seems to be the trend. Students can attend class through computer, downloading assigned reading material, submitting completed assignments through e-mail, and participating in class discussions through live chat. There are already a small percentage of courses offered in this way, and we expect that in the next decade, this will become a standard way to gain college credit or even complete degrees.

Networking Health Care

The advancement of computer and networking technology will make an enormous impact in medicine and health care. The most obvious benefit of networking to the healthcare industry is the capability of physicians to share patient records and diagnostic and treatment information.

Other ambitious and exciting developments are already taking place or are expected to become reality in the near future. For instance, long-distance surgery—in which the physician performing the operation from a remote location controls the robotic device that is actually operating on the patient—has been successful in experimental situations.

A related concept is the tele-examination, in which the physician conducts a preliminary physical exam over two-way video, perhaps aided by a proxy or stand-in such as a physician's assistant or nurse practitioner who is actually with the patient. The idea of telemedicine is especially attractive in rural areas where physicians are in short supply.

Technology and the Law

Legal research and courtroom procedures have become more efficient because of networking and the subsequent improvement in record keeping and information sharing abilities.

Law enforcement, in particular, has put the new technologies to work. Many police agencies now have mobile data terminals (MDTs) installed in squad cars, allowing officers to quickly and directly access criminal history files, license records, and departmental references and resources. MDTs also make it possible for officers to communicate with one another over the network in a fashion that is more secure than broadcasting over open radio channels. Departments are also making use of global positioning system (GPS) technology, combined with computer-aided dispatch, to provide accurate information to dispatchers and supervisors at the station about where each unit is located at any given time. This gives command staff more control for better deployment of officers.

Intergalactic Networking?

Now that Internet connectivity has permeated almost every corner of the globe, the next logical step is to network the final frontier: outer space. The idea is not as futuristic as it sounds; in fact, computer communications are already sending signals to satellites orbiting high over the earth, and there are projects underway that broadcast signals into deep space in hopes of catching the attention of life, if it exists, on other planets.

Computers on Earth already control the activities of unmanned space missions and play a huge role in the journeys of manned shuttles. Networked communications make interplanetary exploration feasible.

A Brief Overview of Networking Terminology

When you visit another country, the first task to tackle is learning the language. The same is true of entering a new area of study or a new career field.

Computer networking, like most professions, has its own jargon, such as technical terms, abbreviations, and acronyms, that can, at first glance, look as foreign to the uninitiated as does the alphabet of a country halfway around the world.

Without a good grasp of the terminology, you will have difficulty understanding the concepts and processes in this book. This section gives you a head start on deciphering some of the tech talk in this and other introductory guides to networking and network operating systems.

This is not intended to be a comprehensive glossary of networking terms, but a quick reference that defines and briefly discusses some of the most important and most basic words, phrases, and acronyms that enable you to navigate through the next few chapters.

Each definition is expanded on in the chapters that follow. Please refer to the glossary for a more comprehensive list of definitions.

Concept-Related Terminology

networking model—A networking model is a graphical representation of the processes involved in network communications. The popular models represent these processes as layers or levels; thus, they are called *layered models*. The most commonly referenced are the Open System Interconnection (OSI) seven-layer model, the four-layer DoD model (sometimes called the TCP/IP networking model), and the Microsoft Windows networking model. We discuss each in detail in Chapter 2, “Categorizing Networks.”

client/server networking—In computer networking terms, a *client* computer is one that sends a request to another computer for access to its data or resources. The computer that responds to that request and shares its data or resources over the network is called the *server*. In a *peer-to-peer network* (also called a *workgroup*), all computers on the network act as both clients and servers. In a *server-based network* (sometimes called a client/server network, and in Microsoft Windows networking, called a *domain*), there is a dedicated server computer running special server software, which performs user authentication/security functions. We discuss both peer-to-peer and server-based networking in Chapter 2.

Network Hardware-Related Terminology

NIC—It is pronounced “nick” and refers to the network interface card, also called the network adapter card (but for some reason never called a NAC), or just the network interface. This card typically goes into an ISA, PCI, or PCMCIA (PC card) slot in a computer and connects to the network *medium*, which in turn is connected to other computers on the network.

media—Media are the means by which signals are sent from one computer to another by cable or *wireless* means.

wireless media—Wireless media, such as the radio, laser, infrared, and satellite/microwave technologies, carry signals from one computer to another without a permanent tangible physical connection (cable).

coax—Coaxial cable, or coax, is similar to cable TV cable, which is copper-cored cable surrounded by a heavy shielding that is used to connect computers in a network. Either thin or thick coax can be used.

twisted-pair—Twisted-pair is a type of cabling, also used for telephone communications, that consists of pairs of copper wires twisted inside an outer jacket. There are two basic types: UTP (unshielded twisted pair) and STP (shielded twisted pair). UTP is the most commonly used cabling in modern Ethernet networks. It comes in different category ratings depending on whether it is considered voice or data grade and the transmission speed it

supports. “Cat 5” refers to a category 5 rating, which can be used for voice (telephone) or data and which supports speeds up to 100 Mbps.

plenum—The plenum in a building is the space between a false ceiling and the floor above, through which cabling can be run. *Plenum-grade* cable, often called plenum cable, refers to cable with an outer jacket made of Teflon or other material that complies with fire and building codes for installation in the plenum area.

PVC—In the context of network hardware and cabling, PVC stands for polyvinyl chloride, the material out of which the jacket on non-plenum-grade cable is made. It is less expensive than plenum-grade materials but does not meet most safety codes for installation in the ceiling because it gives off a poisonous gas when burned.

fiber optics—Often shortened to just *fiber*, fiber optics refers to cabling that has a core made of strands of glass or plastic (instead of copper), through which light pulses carry signals. Fiber has many advantages over copper in terms of transmission speed and signal integrity over distance; however, it is more expensive and more difficult to work with.

connectivity devices—This term refers to several different device types, all of which are used to connect cable segments, connect two or more smaller networks (or subnets) into a larger network, or divide a large network into smaller ones. The term encompasses repeaters, hubs, switches, bridges, routers, and brouters. Each is discussed in detail in Chapter 7, “Physical Components of the Network.”

Software-Related Terminology

protocol—A network protocol is a set of rules by which computers communicate. Protocols are sometimes compared to languages, but a better analogy is that the protocol is like the syntax of a language, which is the order in which processes occur. There are many different types of computer protocols. A *protocol stack* refers to two or more protocols working together. The term *protocol suite* describes a set of several protocols that perform different functions related to different aspects of the communication process.

NOS—*NOS*, which stands for network operating system, usually refers to server software, such as Windows NT, Windows 2000 Server, Novell NetWare, and UNIX. The term sometimes refers to the networking components of a client operating system such as Windows 95 or the Macintosh OS.

client operating system—Also referred to as the desktop operating system, *client operating system* refers to the operating system software that runs on the network’s workstations, which access the server and/or log onto the network as clients.

hybrid network—A hybrid network (also called a multivendor network) is one in which the software products of different vendors interoperate, especially in regard to the server operating systems. For example, a network that has Windows NT domain controllers, NetWare file servers, and a UNIX Web server is a hybrid network.

Design and Topology Terminology

LAN—A local-area network (LAN) is a network that is confined to a limited geographic area. This can be a room, a floor, a building, or even an entire campus.

WAN—A wide-area network (WAN) is made up of interconnected LANs. It spans wide geographic areas by using WAN links such as telephone lines or satellite technology to connect computers in different cities, countries, or even different continents.

MAN—A MAN (metropolitan-area network) is a network that is between the LAN and the WAN in size. This is a network that covers roughly the area of a large city or metropolitan area.

physical topology—This refers to the layout or physical shape of the network, whether the computers are arranged so that cabling goes from one to another in a linear fashion (linear bus topology), the last connects back to the first to form a ring (ring topology), the systems “meet in the middle” by connecting to a central hub (star topology), or multiple redundant connections make pathways (mesh topology). The characteristics of each are discussed in Chapter 3, “Networking Concepts, Models, and Standards.”

logical topology—The logical topology is the path that signals take from one computer to another. This can correspond to the physical topology. For instance, a network can be a physical *star*, in which each computer connects to a central hub, but inside the hub, the data can travel in a circle, making it a *logical ring*. The difference between physical and logical topologies is discussed in Chapter 3.

Measurement-Related Terminology

bit—The smallest unit of data in a computer. A bit equals 1 or 0, and it is the binary format in which data is processed by computers.

byte—A byte is a unit of measure used to describe the size of a data file, the amount of space on a disk or other storage medium, or the amount of data being sent over a network. 1 byte generally equals 8 bits of data.

KB (kilobyte)—A kilobyte is approximately 1000 bytes (actually, it’s 1024 bytes). It can be abbreviated as “K.”

KBps (kilobytes per second)—This is a standard measurement of the amount of data transferred over a network connection.

kbps (kilobits per second)—This is a standard measurement of the amount of data transferred over a network connection.

MB (megabyte)—A megabyte is approximately 1 million bytes (actually 1,048,576). A megabyte is sometimes referred to as a “meg.”

MBps (megabytes per second)—This is a standard measurement of the amount of data transferred over a network connection.

Mbps (megabits per second)—This is a standard measurement of the amount of data transferred over a network connection.

Hz (Hertz)—A unit of frequency. It is the rate of change in the state or cycle in a sound wave, alternating current, or other cyclical waveform. It has one cycle per second and is used to describe the speed of a computer's microprocessor.

MHz (megahertz)—One million cycles per second. This is a common measurement of the speed of a processing chip such as a computer's microprocessor.

GHz (gigahertz)—One thousand million, or 1 billion (1,000,000,000), cycles per second. This is a common measurement of the speed of a processing chip such as a computer's microprocessor.

NOTE

A common error is confusing KB with kb and MB with Mb. Remember to do the proper calculations when comparing transmission speeds that are measured in KB with those measured in kb. For example, modem software usually shows your connection speed in *kilobits* per second (for example, 45 kbps). However, popular browsers display file-download speeds in *kilobytes* per second, meaning with a 45 kbps connection, your download speed would be a maximum of 5.76 KBps. In practice, you cannot reach this download speed because of other factors consuming bandwidth at the same time. We discuss data transfer rates in more detail in Chapter 7.

NOTE

PC processors are getting faster all the time. The microprocessors used on PCs in the 1980s typically ran under 10 MHz (the original IBM PC was 4.77 MHz). As the year 2000 began, PC processors approached the speed of 1 GHz.

What This Book Covers and What It Doesn't

This book provides an overview of networking fundamentals and popular server and client operating systems in use on networks today. Because the scope is broad, we are not able to go into the depth or detail on individual topics in a manner possible in more specialized books.

Throughout the book, we provide a resource list at the end of each chapter to point you toward sources of more detailed information on each of the topics introduced.

Networking Certifications

Because of the shortage of qualified professionals and the high demand for personnel in the networking industry, certification has become a popular means of measuring basic knowledge and qualifications, especially for entry-level positions. Many organizations offer certification examinations to test your grasp of networking technologies.

Vendor-Specific Certifications

Some certification programs are *vendor-specific*; exam candidates are tested on their abilities with particular hardware or software products and are expected to know the “party line” and answer exam questions in keeping with the particular vendor’s philosophy and focus. Many of these certifications, such as the Cisco CCIE and the Microsoft MCSE, are well respected in the industry.

NOTE

Vendor-specific certifications are useful for demonstrating specific capabilities with a particular company’s products and in many instances are desired or even required by employers.

Some of the most popular vendor-specific certification programs include the following:

- Cisco Certified Network Associate (CCNA), Cisco Certified Network Professional (CCNP), and Cisco Certified Internetwork Expert (CCIE)
- Microsoft Certified Professional (MCP) and Microsoft Certified Systems Engineer (MCSE)
- Novell Certified NetWare Administrator (CNA) and Novell Certified NetWare Engineer (CNE)

Other companies, such as Sun, Lotus, IBM, and RedHat and other Linux vendors, also offer certification exams for their networking software. There are also vendor-specific hardware certifications offered by IBM, Compaq Computers, Digital Equipment Corporation, and others.

Non-Vendor-Specific Certifications

Non-vendor-specific certification programs attempt to measure general knowledge and skills applicable to the networking products of a wide range of vendors.

The most popular non-vendor-specific networking skills certification is Network+, which is offered by the Computing Technology Industry Association (CompTIA). This association also developed the vendor-neutral A+ PC hardware technician’s exam.

Non-vendor-specific certifications are useful for demonstrating a broad base of knowledge and skills pertaining to generic networking concepts, practices, and terminology.

This Book and Certification

This book was designed to give you a broad overview of the essential elements of PC networking. It can serve as an introductory guide for those new to the IT industry and those who plan to seek vendor-specific certification such as the CCNA, MCSE, or CNA.

It can also be used as a study guide, in conjunction with other preparatory material, for the Network+ exam. We have covered all topics included in the exam objectives.

NOTE

The Network+ certification was supported and sponsored by such companies as Microsoft, Novell, IBM, Lotus, and many more.

Although the new Windows 2000 MCSE certification track does not include an exam devoted exclusively to networking essentials, much of the material in this book is useful in studying for the Windows 2000 core examination 70-216, *Implementing and Administering a Microsoft Windows 2000 Network Infrastructure*.

Summary

This chapter introduced you to the world of computer networking and how standalone systems began to be linked into networks. You learned about early LANs and the development of the ARPAnet, which was the joint DoD/university project that became today's global Internet.

This chapter touched on some of the ways in which technologies are affecting our lives, and it provided an overview of some of the concepts that are explored in depth in later chapters of this book. You learned some common networking terms and about the role of technical certifications in the networking industry.

In upcoming chapters, you will build on this information as you learn about the models and standards on which today's networks are built. We will go under the hoods of small LANs and complex WANs, and you will become familiar with the signaling methods, architectures, hardware, protocol operating systems, and services that provide the foundation of modern networking.

You will learn about hot topics such as security and troubleshooting, and we will discuss specialty areas such as remote access, virtual private networking, and thin client networking. We will look at emerging technologies that promise that the networks of the future will be even more fascinating, and more practically useful, than the networks of today.

Further Reading

An excellent Web-based resource for definitions of networking terms and acronyms is www.whatis.com.

A good reference for additional information about the history of networking is www.silkroad.com/net-history.html.

For more information about CompTIA and the Network+ certification program, see CompTIA's Web site at www.comptia.com.

For more information about vendor-specific certification training and exams, see the following Web sites:

- Cisco: www.cisco.com/warp/public/10/wwtraining/certprog
- Microsoft: www.microsoft.com/train_cert
- Novell: education.novell.com/certinfo

Review Questions

The following questions test your knowledge on the material covered in this chapter. Be sure to read each question carefully and select the *best* correct answer or answers.

- 1 What was the early implementation of networking technology developed by the French in the early 1800s?
 - a The telephone network
 - b The optical telegraph network
 - c The Ethernet network
 - d The ARPAnet
- 2 What is the name of the technology that the telephone network uses?
 - a Packet switching
 - b Layer 2 switching
 - c Layer 3 switching

- d** Circuit switching
- 3** What is the technology that works best for bursty data transmissions?
 - a** Packet switching
 - b** Analog transmission
 - c** Circuit switching
 - d** Switchboard technology
- 4** Which of the following is a disadvantage of mainframe-based networks? (Select all that apply.)
 - a** The mainframe hardware is more expensive than PC hardware.
 - b** Mainframes are incapable of processing the large amounts of data that are processed by PC servers.
 - c** Mainframes represent a single point of failure.
 - d** Mainframe terminals are less secure than networked PCs.
- 5** Which of the following can be shared across a computer network? (Select all that apply.)
 - a** Data
 - b** Applications
 - c** Printers
 - d** Modems
- 6** Which of the following is classified as an input device? (Select all that apply.)
 - a** Plotter
 - b** Fax machine
 - c** Digital camera
 - d** Printer
- 7** What is the name of the method introduced in the 1980s to provide a means for mapping friendly host names to IP addresses?
 - a** DoD
 - b** DHCP

- c DSL
 - d DNS
- 8 What are the applications that combine telecommunications and computer technologies called?
- a Computel technologies
 - b Telephony technologies
 - c TPI technologies
 - d ITC technologies
- 9 What type of computer sends a request to another computer for access to its data or resources?
- a Server
 - b Workstation
 - c Client
 - d Terminal
- 10 What is the set of rules by which computers communicate?
- a Protocol
 - b Media type
 - c Byte
 - d Topology



Categorizing Networks

Technology professionals break networks into categories based on characteristics required to administer or troubleshoot a particular network. You can break down types of networks according to physical properties or characteristics of the software running on them. For example, categorization can be based on the following:

- Physical scope
- Administrative method
- Network operating system
- Networking protocols
- Topology
- Architecture

We look at each classification method in the following sections.

Categorizing Networks by Physical Scope

One network categorization method is based on the physical scope of the network, which includes the geographic area it spans and, to a lesser extent, the size of the network. Using this method, you can generally place a network into one of three categories:


- Local-area network (LAN)
- Metropolitan-area network (MAN)
- Wide-area network (WAN)

These categorizations correlate somewhat with the network size, which is the number of computers and users (a LAN tends to be smaller than a MAN, which is in turn smaller than a WAN). They also correlate to some degree with financial resources (a WAN is generally much more expensive to set up and maintain than a LAN), but the most important determining factor is the geographic area covered by the network.

Characteristics of a LAN

The *American Heritage Dictionary* defines the word *local* as “of, related to, or characteristic of a particular place, rather than a larger area.” Similarly, the term *LAN* describes a network that spans a limited area; the computers that belong to the network are physically close to one another. However, LANs can vary drastically in their numbers of computers and users. For instance, a LAN could consist of two PCs sitting a few feet apart in a home or office, or it could include hundreds of computers spanning several floors of a skyscraper, or even in some cases, multiple office buildings in close proximity. Figure 2-1 shows a graphical representation of the layout of a simple LAN.


Figure 2-1 A LAN is limited to a specific geographic area.



NOTE Network architecture and cable type can limit the number of computers a LAN can contain. You'll learn about these limitations later in this chapter, in the section "Categorizing Networks by Architecture."

Large LANs can be divided into *workgroups* for easier management. In this context, a workgroup consists of users who share the same resources, such as files, printers, and applications. For example, within a company LAN, you can create workgroups for different departments, such as Sales, Finance, and Human Resources. Figure 2-2 shows how a LAN can be divided into workgroups.

Figure 2-2 LANs can be divided into workgroups.

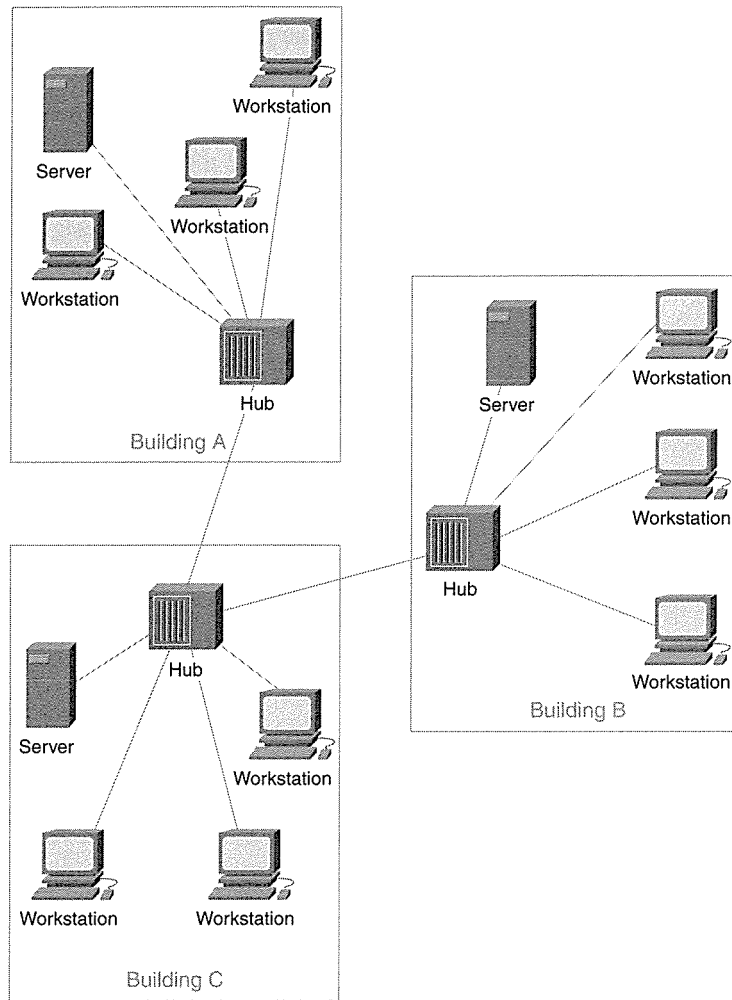


LANs are the basis of larger networks (MANs and WANs), which are created by networking two or more LANs together.

Characteristics of a MAN

A MAN consists of two or more LANs networked together within the confines of a space roughly corresponding to a metropolitan area, hence the name. The typical MAN is a high-performance, public network. See Figure 2-3 for an illustration of a MAN.

Figure 2-3 A MAN covers a wider area than a LAN, but it is more geographically limited than a WAN.



The term *MAN* is used less frequently to define networks than are the terms *LAN* and *WAN* because it is less frequently implemented. Most networks are contained within a building or campus and thus fall into the category of *LAN*, or they span a greater distance, with nodes in different cities, states, or even countries, and thus qualify as a *WAN*. The maximum distance defining a *MAN* is approximately 50 miles, or 80 kilometers.

Characteristics of a WAN

A WAN is a network that spans a large geographic area. The best and most familiar example of a WAN is the Internet. However, a WAN can also be a private network. For example, a company with offices in many countries can have a corporate WAN connecting locations through telephone lines, satellites, or other technologies. The WAN generally consists of many interconnected LANs.

NOTE

When networks are networked to one another, the result is referred to as an *internetwork*, or *internet*. When you see the word internet beginning with a lowercase letter, it refers to any network of networks. When the first letter is capitalized, the word pertains to the global public network of networks we call the Internet.


Related terms that evolved from the term internet are intranet and extranet. An *intranet* is a private network within an enterprise that uses the same protocols (such as TCP, HTTP, and FTP) and technologies used on the Internet. An *extranet* also uses Internet technologies, but is made accessible across remote links to a business's customers, employees, vendors, and partners.

Although WANs can use private links to connect networks, they often use public transports such as the public telephone system. Consequently, transmission speed is often slower than on a LAN; typical speed over analog phone lines using a top-of-the-line modem is at or under 50 kbps. Even the high-speed WAN links, such as T1, cable modem, and digital subscriber lines (DSL) top out between 1 and 6 Mbps. On the other hand, the slowest Ethernet LAN links sends at 10 Mbps.

Another characteristic of the WAN is that a connection to it may not be permanently wired, as with a cabled LAN. Instead, WAN links are often, but not always, dialed "on demand." Many WAN links are in fact dedicated, always-on connections, but temporary connections are more common on the WAN than on the LAN.

In summary, WANs can use either private or public transports and can consist of permanently dedicated or dial-up connections. WAN links are usually slow when compared to LAN links. WANs are categorized as either distributed or centralized. A *distributed WAN*, such as the Internet, has no central point of control. A *centralized WAN*, on the other hand, is based on a central server or a centralized site (such as company headquarters) to which all other computers are connected. Figure 2-4 shows a centralized WAN.

Figure 2-4 A centralized WAN is built around a “master” computer or site to which others connect.



WANs are *routed* networks, which means that in order for messages to get from one LAN to another, the packets must go through a *gateway*. The gateway is a router, or computer that is configured to perform routing functions. We discuss how routing works in much more detail in Chapter 9, “The Widest Area Network: The Global Internet.”

Categorizing Networks by Administrative Method

You can categorize networks based on administrative method, that is, how and by whom shared resources are managed. A network can be organized as follows:

- As a peer-to-peer workgroup in which each computer functions as both client and server and each user administers the resources on his or her computer
- As a client/server, or server-based, network in which administration is centralized on a computer running special network operating system (NOS) server software that authenticates username and password information to enable authorized users to log on and access resources

Which of these is better? It depends on the situation. Each has advantages and disadvantages, which are summarized in Table 2-1.

Table 2-1 *Advantages and Disadvantages of Peer-to-Peer and Client/Server Networks*

Advantages of a Peer-to-Peer Network	Advantages of a Client/Server Network
Less expensive to implement.	Provides for better security.
Does not require NOS server software.	Easier to administer when the network is large because administration is centralized.
Does not require a dedicated network administrator.	All data can be backed up on one central location.
Disadvantages of a Peer-to-Peer Network	Disadvantages of a Client/Server Network
Does not scale well to large networks; administration becomes unmanageable.	Requires expensive NOS software such as NT or Windows 2000 Server or Novell NetWare.
Each user must be trained to perform administrative tasks.	Requires expensive, more powerful hardware for the server machine.
Less secure.	Requires a professional administrator.
All machines sharing the resources negatively impact performance.	Has a single point of failure if there is only one server; users' data can be unavailable if the server is down.

As you can see, choosing the method of administrative organization depends on factors such as the number of computers and users, the security requirements, the hardware, the personnel, and the available budget.

Servers and Clients

Before we look more closely at the characteristics of each type of network and how each is implemented, let's review the terms used to describe the roles of computers on a network.

A *server* is a computer that makes its resources (data, software, or attached peripherals such as printers) available for access by other computers on the network. On the other hand, a *client* is a computer that accesses the resources of a server.

TIP

Remember the two this way: Servers give (that is, share their resources) and clients take (that is, access the resources of servers).

It's easy to become confused by these terms; although we sometimes use the word server to describe a dedicated server computer that is usually a powerful computer running NOS software that enables centralized management of the network, in reality, *any* computer that shares its resources is acting as a server.

Servers Sharing “Shares”

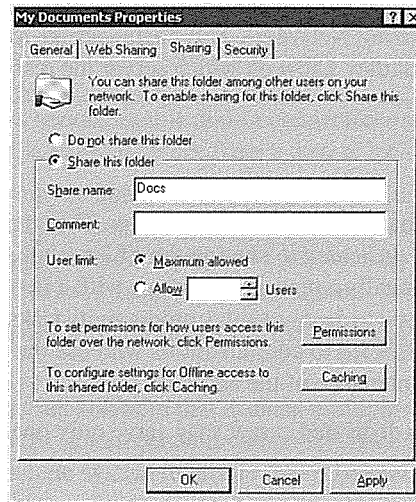
The operating systems that we think of as client or desktop operating systems—Windows 95/98, NT Workstation, and Windows 2000 Professional, for example—can and do function as servers when you create *shares* on them to enable other computers to use their resources.

NOTE Some operating system vendors use the term *share* to describe the resources that have been made available to other computers, while others vendors stay with the more traditional terminology *shared resource*.

To enable resources to be accessed by others, you must specifically designate the resource as shared and give it a *share name* that identifies it on the network. This name does not necessarily have to be the same as the name of the resource itself. For example, if you had a folder called Administration that you wanted to share over the network, you could name the share “Admin.” Admin is the name others would see when they browse the network. They would also use this name when they connect to your share.

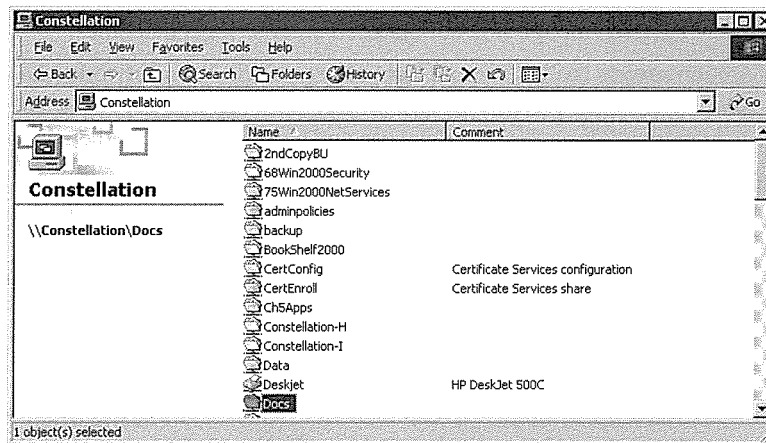
Figure 2-5 shows the dialog box used in Windows 2000 to create a share for the folder named My Documents.

Figure 2-5 Creating a share in Windows 2000.



Notice that the folder name appears in the title bar, but the share is named Docs. When another computer on the network consults the *browse list*, which is the list of available shared resources, the share appears as shown in Figure 2-6.

Figure 2-6 The share named Docs appears in the browse list for the server named Constellation.



NOTE Some NOSs allow for shared resources to be “published” to the directory (for example, shared folders and printers can be published to the Active Directory in Windows 2000 and to the NDS tree in Novell NetWare). This enables users to locate the shared resources without requiring that they know on which server the resource is physically located.

Dedicated Servers

Any computer that shares its resources is a server. The term is often used, however, to refer to machines that are *dedicated* to sharing their files, applications, or peripherals. A *dedicated server* is usually a powerful machine with a fast processor and a large amount of memory that is not used to do day-to-day productivity tasks; it is dedicated to being a server. Access is usually restricted to administrators who use the server to perform management, monitoring, and maintenance of the network.

In a large network, a dedicated server may serve only one function, as in the following situations:

- File servers are servers on which data files are stored. Users save their application data to a hard disk on the server instead of saving to the local hard disks on their workstation machines. Files on the server are easier to locate because they are all in a central location; they also are easier to back up.
- Print servers are machines that control one or several printers to which users can send documents across the network to be printed.
- Application servers are computers on which network applications are installed. Users can run the application (such as a word processing or database program) across the network, even though it is not installed on their local machines.
- Logon servers (called *domain controllers* on Windows networks) hold a *security database*, which contains information comprising user accounts. The server checks user credentials against the database and controls access to the network and its resources.
- Web servers run Web server software on top of the operating system in combinations such as Microsoft Internet Information Server on NT and Windows 2000, Apache on Linux and UNIX servers, and Netscape’s Enterprise Server on various platforms. Many operating systems come with built-in Web server software. Web server software often includes both File Transfer Protocol (FTP) and Network News Transfer Protocol (NNTP) server software.
- Mail servers provide mailboxes to collect e-mail sent to users of the network. This e-mail then can be downloaded to the mail client software on the users’ individual machines. Examples of mail server software include Microsoft Exchange for Windows networks and sendmail, a popular UNIX mail daemon.

NOTE

A *daemon*, in UNIX terminology, is a program that runs continuously and handles periodic requests for service.

- Remote access servers enable dial-in connections so that other computers can access the server or the entire network from a distant location over the telephone lines.
- Terminal servers run software that enables client applications to be run on the server so that “thin client” computers (low-powered, inexpensive systems) can function as terminals rather than as independent systems. The server provides a multisession environment and runs the application programs being used on the clients.
- Telephony servers provide answering machine and voice mail services and also route calls.
- Cluster servers run software that enables multiple servers to be joined in *clusters*, which are groups of independent computer systems, also known as *nodes*, working together as a single system to ensure that mission-critical applications and resources remain available to clients.
- Proxy servers act as intermediaries between workstation users and the Internet to ensure security and provide administrative control and caching services.
- Fax servers provide a central point on the network to send and receive facsimile messages and to distribute incoming faxes to the appropriate users.
- BOOTP servers use the bootstrap protocol to enable client computers to boot an operating system and receive IP configuration information over the network.
- DHCP servers assign IP addresses and TCP/IP configuration information to computers configured to be DHCP clients so that administrators do not have to manually assign a unique IP address to each client machine on the network.
- Name resolution servers provide mapping of friendly network names, which enable users to identify computers without having to remember numerical identifiers. These names map to IP addresses, which are used by the TCP/IP protocol suite to locate computers on the network. Name resolution servers include *Domain Name System (DNS) servers*, which map hierarchically structured host names to IP addresses, and *NetBIOS Name servers* (such as Microsoft’s WINS server) that map flat NetBIOS names to IP addresses.

NOTE

We discuss BOOTP, DHCP, and name resolution services in more detail in Chapter 8, “Networking Protocols and Services,” in the section “The TCP/IP Suite.”

Clients, Workstations, Hosts, and Nodes

As you have learned, a network client is a computer or other network device that requests access to network resources. A client is usually, but not always, a computer. Note that a printer or other network device that can request resources is also technically a client.

The term *client* can also refer to software programs that access the resources of server programs. For example, the e-mail program that runs on your desktop computer and that sends a request to download your new messages from a server computer running e-mail server software is called a *mail client*.

A *client operating system* is commonly installed on desktop computers, or *workstations*, which are either networked as peers in a workgroup or logged on to a logon authentication server to access the network. Examples of client operating systems are Windows 95/98 and Windows NT Workstation.

The term *workstation* is used in many ways. It refers to the NT client operating system, and it is sometimes used to mean any desktop computer running any client operating system. Workstation also has a secondary meaning: a powerful computer used to run resource-intensive application software. Such a computer is a graphics workstation or a computer-aided design (CAD) workstation.

Another term you hear used to refer to computers on a network is *hosts*. This term is often used in reference to TCP/IP-based networks and can include any network device that is assigned an IP address.

NOTE When used as a verb, *host* generally refers to providing a service to another device or computer. For instance, a Web server hosts users' Web pages.

A *node* is a connection point on a network. In some contexts, the word refers to any computer or other network device. In other cases, it indicates a redistribution point, or device, that is programmed or engineered to recognize and process transmissions to other nodes.

Characteristics of Peer-to-Peer Networks

The peer-to-peer structure is appropriate for small networks, where strict security is not required. Most networking books recommend a maximum of ten computers for this type of administrative model.

Implementing a small peer-to-peer network is inexpensive and relatively easy with modern operating systems such as Windows 95/98, NT Workstation, Windows 2000, and various versions of Linux, all of which have networking components built in.

If you want to participate on a peer-to-peer network, you configure the computer to join a *workgroup* in the networking configuration properties in Windows 9x, Windows NT, and Windows 2000. All computers that share their resources with one another must have the same workgroup name entered in this field.

Administration in Peer-to-Peer Networks

Administration over users and resources is *decentralized* in a peer-to-peer network. Every computer on the network can act as both client and server; that is, every computer can share its own resources with the rest of the network and access the resources of others.

The user of each computer is responsible for administration of that computer's resources (creating user accounts, creating shares, and assigning permissions to access those shares). Each user is responsible also for backing up the data on that computer. Unfortunately, locating a resource can be difficult in a peer-to-peer network that has more than a few computers.

Security in Peer-to-Peer Networks

A peer-to-peer network has no central database where user account information is kept. All security is local. User accounts and passwords must be created and maintained on each individual machine. This is both inconvenient and inherently less secure than the centralized security model used in client/server networks.

For example, consider a workgroup consisting of four NT Workstation machines belonging to Mary, Joe, John, and Jane. For Mary to access files stored on Joe's machine, Joe must create a user account for her on his computer. If she also wants to access a shared printer on Jane's computer, Jane must create a user account for Mary, and so on.

If Joe, Jane, and John assign different usernames and passwords to these accounts, Mary can be in the position of having to remember to use the name *maryjohnson* and the password *tree* when she accesses resources on Joe's machine, username *maryj* and password *Shelton* when she accesses resources on Jane's machine, and username *mjohnson* and password *flute* when she accesses resources located on John's computer. You can see how unmanageable all those accounts can get as the network grows.

Share-Level versus User-Level Security

In a Windows for Workgroups or Windows 95/98 workgroup, security is implemented at the *share level* rather than at the *user level*. Table 2-2 outlines the differences between share- and user-level security.

Table 2-2 *Characteristics of Share- and User-Level Security*


Share-Level Security	User-Level Security
Used by Windows for Workgroups 3.11, Windows 95, and Windows 98.	Used by Windows NT and Windows 2000.
A password is assigned to each shared resource.	A password is assigned to each individual user.
To access a resource across the network, the user must enter the password for that resource.	To access a resource across the network, the user's account must have permissions assigned to access that resource.
Users must remember multiple passwords.	Users must remember only one password.

In our small workgroup example from before, if Mary, Joe, Jane, and John are running Windows 95 on their computers, they need not create user accounts for each person who will access their resources. Instead they assign a password to every shared resource. This means if Mary wants Jane to be able to access the Documents folder, but doesn't want Joe or John to have access to it, she must give the folder a password (for example, "health") and tell only Jane the password. Jane is prompted to enter the password when she tries to access the folder.

If Mary has another folder, named Pix, that she wants John and Jane to access, but not Joe, she must assign that folder a different password (let's say "crate") and give that password only to John and Jane.

Suppose Mary has 20 shared folders on her hard disk. She has to remember 20 different passwords that grant access to those folders, and the other users have to remember multiple passwords *and* keep track of which password goes with which folder. (And you thought the multiple user account scenario with the NT workgroup was bad!) Figure 2-7 illustrates this situation.

Figure 2-7 With share-level security, the number of passwords that must be remembered can quickly get out of hand.



Because this situation becomes intolerable as the network grows, most users in Windows 9x workgroups soon resort to the easiest solution—not securing resources at all. Obviously, leaving resources unsecured does not work in a network environment that contains highly sensitive data to which access must be restricted. It is better to use a server-based network, as described in the following section.

Characteristics of Server-Based Networks

The defining characteristic of a server-based (also called a client/server network) network is *centralized control*. In such a network, at least one machine runs an NOS such as Windows NT, Windows 2000 Server, or NetWare. User accounts are created on the server, and the network administrator can control the entire network from this central location.

Server-based networking solves the problem of performance degradation that occurs on workgroup computers when other users are accessing a system's resource while a user is working at the station. Generally, performance and throughput are better on a server-based network.

Server operating systems generally provide for additional sophisticated services. For example, a Windows NT/2000 server can function as a remote access (dial-in) server to which multiple users simultaneously can connect over phone lines. Note that Windows 9x and NT Workstation allow one incoming connection at a time. The server also has other services built in, such as the capability to provide name resolution services, to allocate IP addresses as a DHCP server, and to host Websites with add-on or built-in Web server software.

Administration in Server-Based Networks

The server-based networking model simplifies administration, especially for networks with many computers and a large number of shared resources. Shared files are stored on the server so that they can be easily located and backed up.

One important characteristic of server-based computing, which can be seen as either a drawback or a blessing, is the requirement for a dedicated professional network administrator. Although this can increase the cost of maintaining the network, having all network operations under the control of a trained administrator, rather than having end-users manage their own resources, can actually reduce overall costs in the long run.

Security in Server-Based Networks

Server-based security is inherently more secure than that of peer-to-peer networks. To log on to the network, each user must have a valid user account and password created on the server. This enables access to resources throughout the network, and there is no need for multiple user accounts on different machines, as is the case with an NT Workstation workgroup. There is also no need to remember passwords for different shared resources, as is the case with Windows 9x workgroups, because server-based networks rely on *user authentication* and *permissions* to control access.

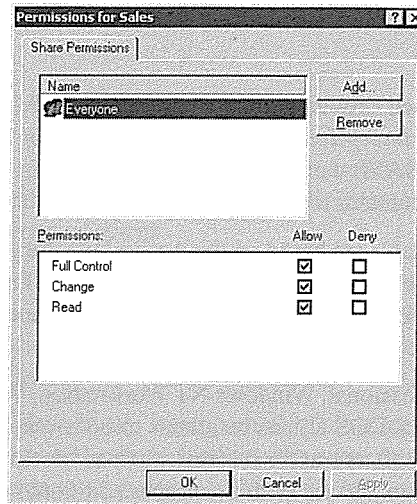
The network administrator can assign permissions for each shared resource to individual users or groups. Those permissions can enable different levels of access for different users. For example, let's say there is a shared folder on the server called Sales. We want Mary, who is a secretary in the Sales department, to be able to read the documents in the folder, but she should not be able to make changes to them. We want Jane, who is vice president in charge of Sales, to be able to read, change, delete, and otherwise fully control the documents in the Sales folder. To accomplish this in a Windows 2000 network, we can edit the share's properties and assign read permissions to Mary's user account and full control permissions to Jane's.

NOTE

A better way to manage permissions, instead of assigning them to individual user accounts, is to organize the user accounts into *security groups* and assign permissions to the groups. We discuss use of groups in more detail in Chapter 10, "Network Operating Systems," in the section "General Network Administration."

When we assign permissions for the Sales folder to Mary and Jane (or groups to which they belong), we must remove the default share permission, as shown in Figure 2-8.

Figure 2-8 In Windows networks, the default share permission gives Full Control to Everyone.



As you can see in Figure 2-8, when you create a new share, the Everyone group is automatically given Full Control. This group includes all users who have a valid account to log on to the network. Luckily, you can easily delete this permission to restrict access to only those who should have it.

NOTE

This is a simplified explanation to illustrate the difference between server-based and workgroup-based *share permissions*. Windows 2000 can implement several different layers of security, depending on the file system and protocols used. The permissions assigned to the share make up only one of the locks on the door to a shared resource.

Categorizing Networks by NOS

Networks are sometimes categorized based on the NOS that is installed on the servers and that is used to control the network. Examples of network types based on the server operating system include the following:

- Windows (Windows NT and Windows 2000)
- NetWare
- UNIX

Many networks combine two or more server types on the network; these are often referred to as *hybrid networks*.

Windows Networks

Windows server-based networks are called *domains*. In Windows NT 4.0 domains, a master computer called a *primary domain controller* holds the only read/write copy of the security accounts manager (SAM) database. Since the release of Windows 2000, Microsoft has referred to NT 4.0 domains as *downlevel domains*.

Windows 2000 domains are based on the Active Directory, a copy of which is held on each domain controller and which contains security account information and objects representing network resources. A network can have multiple domain controllers; all can read and write to the directory database.

Windows 2000 Professional, Windows NT Workstation, Windows 95 and 98, Windows for Workgroups, and MS-DOS all can be clients to both Windows NT and Windows 2000 servers. Non-Microsoft operating systems, such as Macintosh and Linux, also can access resources on Windows servers with the proper additional software installed.

NOTE

Each popular network operating system is discussed in detail in Chapter 10.

NetWare Networks

Novell NetWare is a popular NOS that provides logon security and functions as a file and print server. Windows desktop operating systems can access NetWare servers if they have the proper client software installed. Windows 9x, NT, and 2000 all include client services for NetWare. Novell also makes an add-on client program, Client32, that can be installed on Windows 32-bit operating systems to provide greater functionality.

NetWare versions 4.x and 5.x provide directory services through NetWare Directory Services (NDS), a hierarchical database similar to (and which preceded and perhaps inspired) Microsoft's Active Directory. Older versions of NetWare (3.x and later) are still in use on some networks and use a database called the *Bindery* to organize network objects.

UNIX Networks

UNIX was the NOS originally used by most hosts on the ARPAnet, the predecessor to the Internet. UNIX was developed by Bell Labs in 1969 and comes in many flavors because of its open code distribution. It is a powerful NOS, but most UNIX implementations are text-based and are relatively difficult to learn.

Linux is a variation on UNIX that recently has become popular both as a server and as a desktop operating system. Like its older brother UNIX, Linux is an open standard, and many different companies market their own versions. Popular versions include RedHat, Caldera, and Corel.

Hybrid Networks

Most medium to large networks today can be considered hybrid networks. They run software made by different vendors, use multiple protocols, and can even combine the domain and workgroup concepts.

A Microsoft network in which clients log on to a Windows NT domain controller may also have a NetWare file server that those same clients access and a UNIX machine that is used to provide Web hosting services. The PCs can even connect to an IBM AS/400 mainframe to access certain applications and records.

Most vendors provide software interoperability tools, either included with the operating system or available as add-ons, to facilitate their integration into this type of multivendor environment. For instance, Microsoft's Windows NT and 2000 Server products include Gateway Services for NetWare and Services for Macintosh.

Some popular interoperability programs include the following:

- **Client Services for NetWare (CSNW) and Gateway Services for NetWare (GSNW)**—CSNW enables individual Microsoft client computers to directly access NetWare servers. GSNW enables a Microsoft server's clients to access the NetWare server's resources by going through the gateway software installed on the NT or Windows 2000 server.
- **File and print services for NetWare**—This enables clients to a NetWare server to access resources on a Windows server.
- **Services for Macintosh**—This enables Macintosh computers to access files and printers on a Microsoft network.
- **Systems Network Architecture (SNA)**—SNA enables PC networks to connect to IBM mainframes.
- **SAMBA**—SAMBA is a set of utilities that enables Microsoft computers to access files and print services on UNIX servers.

These interoperability solutions are discussed in greater detail in Chapter 13, "Hybrid Networks."

Categorizing Networks by Protocol

Sometimes networks are categorized according to the protocols they use for communications. The network protocols are the rules of order followed by the linked computers in establishing and maintaining communication over the network. The three most popular LAN protocols are NetBEUI, IPX/SPX, and TCP/IP, as described in the following sections. Other LAN protocols include AppleTalk and the OSI protocol suite.

NetBEUI Networks

A small, simple LAN using Microsoft operating systems can communicate using the NetBEUI protocol. NetBEUI (which stands for NetBIOS Extended User Interface) is based on the NetBIOS (Network Basic Input/Output System) protocols developed by IBM for workgroups. Note that a NetBEUI network cannot be routed because the protocol is nonroutable. This means if your network is divided into subnets, you have to use a different LAN protocol for computers to be able to communicate with computers in a different subnet.

You will learn more about why and how networks are subnetted in Chapter 8 in the section “IP Subnetting and Supernetting.”

Advantages of NetBEUI include its simplicity and low resource overhead. It is fast and requires no complicated configuration information to set up.

IPX/SPX Networks

Novell uses the Internet Package Exchange/Sequenced Packet Exchange (IPX/SPX) protocol stack as its LAN protocol, and it is required for NetWare networks before version 5.0.

NOTE

NetWare version 5.0 is the first version of NetWare that supports running on “pure IP” (the Internet Protocol of the TCP/IP protocol stack) and does not require IPX/SPX.

IPX/SPX is usually associated with NetWare networks, but is not limited to that purpose. A workgroup or domain of Microsoft computers can use the IPX/SPX protocol as well. Microsoft includes its own implementation of IPX/SPX-compatible protocols, called NWLink, in the Windows 9x, NT, and 2000 operating systems. A Microsoft client must have NWLink or IPX/SPX installed to connect to a NetWare server running NetWare 4.x or below.

IPX/SPX requires minimal configuration (more than NetBEUI, but less than TCP/IP) and offers faster performance than TCP/IP. IPX/SPX is sometimes run on internal Microsoft networks that are connected to the Internet for security purposes. We discuss this option in more detail in Chapter 13 in the section “Security Aspects of Using Multiple Protocols.”

TCP/IP Networks

Despite the fact that it is the slowest and most difficult to configure of the popular LAN protocols, TCP/IP is the most widely used. There are several good reasons for this:

- TCP/IP uses a flexible addressing scheme that is extremely routable, even over the largest networks.
- Almost all operating systems and platforms can use TCP/IP.
- A huge number of utilities and tools are available, some of which are included with the protocol suite and some of which are add-on programs for monitoring and managing TCP/IP.
- TCP/IP is *the* protocol of the global Internet. A system must run TCP/IP to connect to the Internet.

Most enterprise level networks run on TCP/IP, and it is imperative that network administrators be familiar with its protocols. We cover the individual components of the TCP/IP suite in Chapter 8.

Other LAN Protocols Used in Networks

Most LANs use one of the three protocols, but you might encounter a few other protocols and protocol suites in your study of PC networking.

AppleTalk

AppleTalk is the set of protocols developed by Apple for networking its Macintosh machines. The AppleTalk suite includes the following protocols:

- **LocalTalk**—Used for connecting Macintosh computers in small workgroups. It is relatively slow (230.4 kbps) and supports only 32 devices.
- **EtherTalk**—Used to connect Macintosh workgroups to Ethernet networks.
- **TokenTalk**—Used to connect Macintosh workgroups to Token Ring networks.

AppleTalk networks use AppleTalk Address Resolution Protocol (AARP) to map AppleTalk addresses to Ethernet and Token Ring physical Media Access Control (MAC) addresses.

The OSI Protocol Suite

The Open System Interconnection (OSI) protocol suite was intended to be a replacement for TCP/IP, which was expected to be phased out. The OSI suite was developed by the International Organization for Standardization (ISO) to provide an improved set of protocols for less confusion and easier standardization of networking products among multiple vendors.

NOTE

The International Organization for Standardization is also called the ISO. No, that's not a typo—"ISO" is not an acronym, but a word derived from the Greek *isos*, meaning "equal."

Phasing out TCP/IP and using OSI instead sounded good in theory, but TCP/IP proved to be a little harder to kill off than anticipated. The U.S. government got behind the OSI suite, and in the late 1980s the Department of Defense decreed that by August 1990, all its computer communications would use OSI protocols. But it didn't happen. Networks didn't make the change, and TCP/IP still reigns as king of the Internet. Its position seems strong, especially with its capability to evolve (demonstrated by the planned transition to IPv6) to meet the challenges of future growth.

Categorizing Networks by Topology

Networks are sometimes categorized based on the physical or logical *topology* of the network. The *physical topology* refers to the shape of the network—the way in which the cable is laid out. The *logical topology* refers to the path the signals travel as they make their way from one point on the network to another.

The physical and logical topologies can be the same; in a network physically shaped as a linear bus (that is, in a straight line), the data travels in a straight line from one computer to the next. A network also can have physical and logical topologies that are not the same. The cable segments can connect all computers to a central hub in a star shape, but inside the hub, the connections can be wired so that the signal travels around in a circle from one port to the next, creating a logical ring.

The following are the most popular LAN topologies for networks:

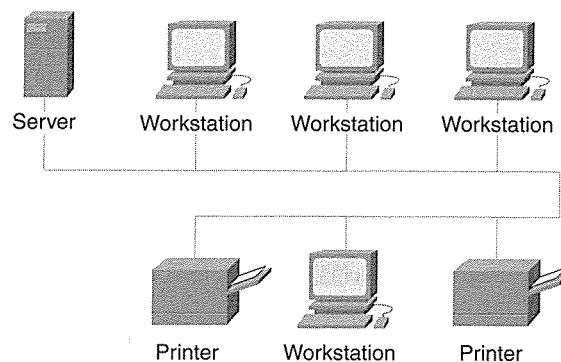
- Linear bus
- Ring
- Star bus
- Mesh
- Hybrid

They are described in more detail in the following sections.

Linear Bus Networks

A linear bus (sometimes called just a *bus*), as the name implies, is a network that is laid out in a straight line. The line doesn't actually have to be physically straight; rather, the cable proceeds from one computer to the next, and then to the next, and so on. Figure 2-9 illustrates a linear bus.

Figure 2-9 Computers in a linear bus network are connected in a line from one to the next.



Because it has a beginning and an end, a linear bus network requires *termination* at each end. Failure to terminate both ends of the cable results in *signal bounce*, which can disrupt or prevent communications on the network. One end of a linear bus—but *not both*—should be grounded.

NOTE

To end a bus, a device called a terminator is attached to the “empty” side of the NIC T-connector on the first and last computers on the linear cable.

Bus networks usually use thick or thin coax cable and the Ethernet 10Base2 or 10Base5 architecture. You learn more about the characteristics of different cable types in Chapter 7, “Physical Components of the Network.”

Communications on a Bus Network

On a bus network, when a computer sends a message, that message goes to every computer on the bus. Each network interface card (NIC) examines the headers of the message to determine whether the message is addressed to that computer. If it is not, the message is discarded.

Advantages of Bus Networks

The bus topology is very simple and easy to set up. It is relatively inexpensive because it uses less cable than other topologies. The bus is especially suitable for small, temporary networks, such as in a classroom that might be used for only a few days or weeks.

Disadvantages of Bus Networks

A bus is known as a *passive* topology because the computers do not regenerate the signal and pass it on as they do in a ring. This makes the network vulnerable to *attenuation*, which is the loss of signal strength over distance. *Repeaters* can be used to address this problem. We discuss repeaters in more detail in Chapter 7.

Another disadvantage of the bus is that if there is a break in the cable (or one user decides to unplug his or her computer from the network), the line is broken. This means not only that computers on opposite sides of the break cannot communicate, but also that two new ends are not terminated and the resulting signal bounce can bring the entire network down.

Ring Networks

If you connect the last computer in a bus back to the first, you have a *ring* topology. In a ring, every computer is connected to two other computers, and the signal can travel around and around the circle (see Figure 2-10). Because a ring has no endpoint, termination is not necessary (or possible).


A physical ring network generally uses coax cable, as with the bus. A Token Ring network, which is a logical ring, uses STP cable (IBM type) and complies with IEEE 802.5 specifications.

Communications on a Ring Network

On a ring network, the signal travels in one direction. Each computer receives the signal from its *upstream neighbor* and sends it to its *downstream neighbor*. The ring is considered an *active topology* because each computer regenerates the signal before passing it on to the next.

The ring topology is most commonly associated with the Token Ring architecture. In this implementation, the ring is generally a *logical* ring, the circle being wired inside the Token Ring hub, which is called a multistation access unit (MSAU).

Figure 2-10 In a ring network, computers are connected in a circle, with the last connected back to the first.



Advantages of Ring Networks

The ring is relatively easy to troubleshoot, and like the bus, is simple to set up. A physical ring requires more cable than a bus and less than a star topology.

Disadvantages of Ring Networks

The ring suffers from some of the same drawbacks as the bus. If the circle stays unbroken, it is a reliable topology. If a break or disconnection of the cable occurs anywhere on the network, however, it brings all network communications to a screeching halt.


Another disadvantage of the ring is the difficulty in adding more computers to the network. Because the cabling runs in a closed circle, it is necessary to break the ring at some point to insert the new computers. This means the network is out of commission while you make the additions.

Star Bus Networks

The *star* (also called a *star bus*) is one of the most popular LAN topologies. It is implemented by connecting each computer to a central hub, as shown in Figure 2-11.

The hub can be *active*, *passive*, or *intelligent*. A passive hub is just a connection point. It does not require electrical power. An active hub (the most common) is actually a *repeater* with multiple ports; it boosts the signal before passing it to the other computers. An intelligent hub is an active hub with diagnostic capabilities. It has a processing chip built in. We discuss hubs in more detail in Chapter 7.

Figure 2-11 *The star topology connects all computers to a central hub.*



The star topology is generally used with unshielded twisted-pair (UTP) cabling and the Ethernet 10BaseT or 100BaseT architecture.

Communications on a Star Network

On a typical star network, the signal is passed from the sending computer's NIC to the hub, boosted (that is, amplified), and sent back out all ports. On a star, in a manner similar to a bus, all computers receive the message, but only the computer whose address matches the destination address in the message's header pays attention.

Advantages of Star Networks

The star topology has two big advantages over the bus and ring. First, it is far more *fault tolerant*; that is, if one computer becomes disconnected or there is a break in its cable, only *that* computer is affected, and the rest of the network can communicate normally. Second, it offers ease of reconfiguration. Adding more computers to the network, or removing computers, is as simple as plugging in or unplugging a cable. Troubleshooting physical layer problems in a star network is also easy, especially with an intelligent hub that provides diagnostic information. You learn more about the physical and other layers defined by the standard networking models in Chapter 3, "Networking Concepts, Models, and Standards."


Disadvantages of Star Networks

Despite the star's benefits, it does have a few disadvantages, primarily related to cost. First, it uses more cable than the linear bus or ring because there must be a separate length of cable going all the way from the hub to each computer. Another source of additional expense is the hub itself, which must be purchased in addition to the cable. The good news is that UTP cable is relatively inexpensive, and there is no need for terminators on a star network.

Mesh Networks

The *mesh* is a topology you won't see as often as the three discussed so far. In a mesh network, every computer has a direct connection to every other computer on the network, as shown in Figure 2-12.

Figure 2-12 *In a mesh network, every computer is connected to every other computer.*



These redundant connections make the mesh the most fault-tolerant of all topologies. If one pathway from the sending to the destination computer is down, the signal can take another path.

Unfortunately, this advantage is offset by both the high cost of the huge amount of cable required to implement a mesh and the complexity of the network if more than a few computers are involved. The number of connections increases exponentially as each new computer is added. It is no coincidence that “mesh” sounds a lot like “mess”—that’s exactly what you have as a mesh network grows.


Hybrid Topologies

The word *hybrid* is used in a couple of different ways in reference to network topology. In Chapter 13, we examine the use of the term *hybrid* to describe networks that run multiple protocols, operating systems, or computing platforms. In this chapter, the word is used in reference to a topology that combines elements of two or more of the standard topologies (for example, hybrid mesh, star bus, or ring bus).

Hybrid Mesh Networks

Because the mesh topology quickly grows complex and unmanageable, many networks are based on a semi-mesh topology, in which there are redundant connections between some of the computers, but not all; this type of network is often referred to as a *hybrid mesh*. The redundant connections should be made between the computers that have the greatest need for a fault tolerant connection. Figure 2-13 shows an example.

Figure 2-13 A hybrid mesh provides for redundant connections between some computers, but not all.




The hybrid mesh provides much of the benefit of the mesh at less expense, and it is easier to set up and manage.

Combined Topologies

Hybrid is also used to refer to networks that use multiple topologies. Many networks combine one or more topologies. For instance, you can have several hubs with computers connected to each in a star configuration, and then network the hubs in a linear bus. Many hubs include a BNC connector for thin coax cable, along with several RJ-45 ports for the UTP connections, for this purpose.

In this type of connection (see Figure 2-14), the coax cable connecting the hubs is called the *backbone*. The backbone is the part of the network that connects all smaller parts, or *segments*. Several segments can be connected to one backbone to create a larger network.

Figure 2-14 *The coax cable connecting the hubs is called the backbone of the network.*



Categorizing Networks by Architecture

Yet another way to describe a network is by its architecture. Generally speaking, the *network architecture* includes a set of specifications that take into account its physical and logical topologies, the type of cable used, distance limitations, media access methods, packet size and headers, and other factors. You might see these specifications referred to as the *data link layer protocols*.

The most popular current LAN architectures are Ethernet and Token Ring. Other common LAN architectures are AppleTalk and ARCnet. All are described in the following sections.

Ethernet Networks

Ethernet, developed in the 1960s and refined by Xerox, Digital, and Intel to form the basis of the IEEE 802.3 specifications, is today's most popular networking architecture. We examine the role of IEEE specifications in Chapter 3.

Ethernet networks are configured as a physical bus or star, and use the carrier sense multiple access collision detect (CSMA/CD) method of media access. Standard Ethernet was limited to 10 Mbps, but *Fast Ethernet*, which runs at 100 Mbps, is now commonplace, and *Gigabit Ethernet*—capable of speeds exceeding 1 Gbps—is emerging as a new and exciting technology.

There are different subcategories of Ethernet networks, depending on the cable type used. They include the following:

- 10Base5
- 10Base2
- 10BaseT
- 100BaseT
- 1000BaseT
- 100BaseVG-AnyLAN
- 10BaseFL
- 100BaseFL

In the next sections, we discuss characteristics of each type of Ethernet network.

NOTE

Specifications about maximum and minimum segment lengths are given in meters in most books because it is standard practice in the scientific community to use metric measurements rather than the U.S. measurement system. We follow that convention in this book. However, popular certification exams have been known to mix measurement types. For example, one question can give a cable length in meters, while another gives the measurement in feet, so it is important to know how to make the conversion: One meter equals 3.28 feet.

10Base5 Ethernet Networks

Referred to as standard Ethernet, although it is no longer as popular as some of the other types, 10Base5 networks use thick coaxial cable (approximately one-half inch in diameter) and are also called *thicknet networks*. Thick coaxial cable is also called *RG-8* or *RG-11*.

The 10 in 10Base5 refers to its maximum speed: 10 Mbps. The 5 refers to maximum segment length, which is 500 meters. Thicknet networks use the bus topology.

Thicknet is more difficult to work with than other Ethernet cabling for a couple of reasons. First, it is less flexible because of the larger diameter. Second, connections are made using a device called a *vampire tap*, which requires that you drill a small hole into the cable to attach the connector. 10Base5 networks use *external transceivers*. A transceiver is a device that generates and receives the data signals; in other Ethernet architectures, the transceiver is built into the network card. 10Base5 technology uses DIX connectors with AUI cable to connect the transceiver to the NIC.

NOTE We discuss NICs and the various connector types in detail in Chapter 7.


Thicknet often is used in conjunction with other Ethernet cable types to provide a *backbone* linking thinnet segments or 10BaseT hubs.

10Base2 Ethernet Networks

A popular coax-based network type is 10Base2 Ethernet, which uses thinner (approximately one-quarter inch in diameter), less expensive, and more flexible cable than 10Base5. The 2 in this case is a bit misleading; it is actually a rounded up approximation of the maximum segment length, which is 185 meters. Like 10Base5, these thinnet networks are physically structured as a linear bus, which requires termination at each end.

10Base2 networks are easier to set up and work with at the physical level than thicknet networks. Twist-and-push connectors (called *BNC connectors*) are used to connect the cable to a *T-connector* on the network card, as shown in Figure 2-15, and the transceiver is built into the network card.

Figure 2-15 Thinnet connects to the NIC's T-connector with a BNC twistlock device on the end of the cable.



Thin coaxial cable for a 10Base2 network is 50-ohm RG-58A/U or RG-58C/U. (The first is most often used; the latter is a military specification.)

WARNING Although RG-58 cable is similar in appearance to RG-59, which is used for cable television installations, don't confuse the two. RG-59 is 75-ohm cable and should *not* be used for computer networking. RG-58U is a similar cable type that does not meet IEEE specifications and should be avoided.

Thinnet cable is sometimes called *cheapernet* because it costs less to implement than thick coax cable. Thinnet is generally more expensive per foot than UTP, but less cable is required for a coax network because of the bus configuration.

UTP Ethernet Networks

The most popular media for new LAN installations is UTP cable, which is used in 10BaseT (T for *twisted*) networks. Most of us are already familiar with twisted-pair cabling, which is used for telephone cabling.

UTP comes in different grades, identified as numbered *categories* that follow the pattern of Cat 1, Cat 2, and so on. Table 2-3 shows the UTP categories and the use and characteristics of each.

Table 2-3 Available Categories of UTP Cable

UTP Category	Maximum Transmission Speed	Characteristics and Uses
Cat 1	Voice only	Used in old telephone installations
Cat 2	4 Mbps	Not recommended for data transmission
Cat 3	16 Mbps	Lowest recognized data grade; used for most telephone wiring
Cat 4	20 Mbps	Suitable for networking 10 Mbps Ethernet networks
Cat 5	100 Mbps–1 Gbps	Most popular grade for LAN networking; used for Fast Ethernet (100 Mbps)
Cat 5 Enhanced (Cat 5e)*	155 Mbps	Used for Fast Ethernet and 155 Mbps Asynchronous Transfer Mode (ATM)
Cat 6 & 7*	1 Gbps and up	Used for new Gigabit Ethernet technologies

* You often do not see Cat 6, Cat 7, and enhanced Cat 5 mentioned in networking texts because these new cable types have specifications that have been only recently established.

As you can see, UTP supports higher transmission speeds than coax. UTP is also highly flexible and easy to install. It uses RJ connectors, which are the modular plug types used for telephones. Although telephone cables usually use the smaller RJ-11 connector, most Ethernet cables connect with the slightly larger RJ-45 connector.

In a UTP-based network, each computer has a length of cable connecting it to a central *hub* in a star topology. We discuss various types of hubs in Chapter 7.

10BaseT Networks

The 10BaseT specification is popular for LANs of all sizes. It can run on Cat 3 cable, which is already installed in many buildings for telephone communications. New 10BaseT networks are usually set up using Cat 5 or 5e cable so that it is easy to upgrade to 100 Mbps later.

100BaseT Networks

The 100BaseT classification refers to Ethernet networks running at 100 Mbps over Cat 5 or 5e cable. These networks use the same topology and access methods as 10BaseT. Indeed,

the only differences are the requirement for the higher-grade cable and the fact that the network cards and hubs must support the 100-Mbps transmission speed.

Many NICs and hubs are made to support both 10- and 100-Mbps transmission speeds, which makes it easy to upgrade incrementally. In addition, with the proper hardware, you can run part of the network at 100 Mbps while other parts still run at 10 Mbps.

1000BaseT (Gigabit Ethernet) Networks

The initial standards for very high speed Ethernet, most commonly called Gigabit Ethernet, were established by the IEEE in 1996 and published as the 802.3z specifications. These standards provide for 1000-Mbps (1-Gbps) transmission, using the 802.3 Ethernet frame format and the CSMA/CD access method.

NOTE

The IEEE 802.3ab standard sets specifications for the operation, testing, and usage requirements for Gigabit Ethernet for distances of up to 100 m, using four pairs of Cat 5 copper cabling. This includes most of the cabling already installed in buildings for 10BaseT and 100BaseT networks.

A big advantage of Gigabit Ethernet is its interoperability and backward compatibility with its predecessors. Gigabit Ethernet makes an excellent backbone technology in conjunction with 10BaseT and 100BaseT LANs.

NICs and hubs that support the 1000 BaseT standards are available, although at a cost that is several times that of 100-Mbps Ethernet components.

100BaseVG-AnyLAN Networks

Hewlett Packard developed 100BaseVG-AnyLAN technology as a fast, reliable networking architecture that uses a special type of hub that functions as an intelligent central controller. This hub manages network access by continually performing a rapid round-robin scan of network port requests to check for service requests from the nodes that are attached to them. The hub receives the incoming data packet and directs it only to the port with a matching destination address. This matching provides inherent network data security.

Hubs can be linked together, and each hub can be configured to support either 802.3 Ethernet or 802.5 Token Ring frame formats. However, all hubs on the same network segment must be configured to use the same frame format.

100BaseVG-AnyLAN networks often are placed into the Ethernet category, but they use a different media access method, called *demand priority*, that is defined in IEEE specification 802.12.

10BaseFL and 100BaseFL Networks

The FL in 10BaseFL stands for *fiber link*, and these networks use baseband signaling over fiber optic cable. You learn more about signaling methods in Chapter 4, “Networking Communications Methods.”

Fiber-optic cabling uses pulses of light instead of electrical signals to represent the 0s and 1s of binary communication used by computers. A big advantage of fiber optics over copper cabling is its resistance to interference and *attenuation* (that is, loss of signal strength over distance). A cable segment under the FL specifications can be 2000 meters in length, which is 4 times that of 10Base5, over 10 times that of 10Base2, and 20 times the limit for 10BaseT!

Token Ring Networks

Token Ring is a networking architecture that was developed by IBM in the 1980s. It was designed to overcome some of the problems inherent in a *contention*-type network such as Ethernet, where the computers on the network compete or contend for the chance to broadcast. We discuss this more in Chapter 4 in the section “Media Access Methods.”

In a Token Ring network, which uses a logical ring topology, a signal called a *token* is passed around the circle and a computer cannot broadcast until the token gets around to it. This means that, unlike Ethernet networks, Token Ring networks do not experience *data collisions*. A data collision occurs when two computers send at the same time.

Although logically a ring, Token Ring networks are physically laid out as star topologies. Computers connect to a central hub or concentrator called an MSAU. The ring is actually inside the hub, where the wiring connects the ports in a continuous circle and the data travels in a circular path.

The standards for Token Ring networks are defined in IEEE 802.5. IBM cable types, primarily shielded twisted-pair (STP), are used. Token Ring network cards and other components are generally more expensive than those made for Ethernet networks. However, Token Ring is a highly reliable architecture that is still in use in many LANs.

Older Token Ring components supported only 4-Mbps transmissions, but newer implementations can transfer data at 16 Mbps.

NOTE Efforts are under way to develop an effective high-speed Token Ring technology that runs Token Ring over FDDI or ATM. For more information, see the white papers provided by the High-Speed Token Ring Alliance at www.hstra.com/whitepapers.html.

Advantages of Token Ring include the following:

- As a Token Ring network gets overloaded, its performance degrades gracefully. That is, the network gradually gets slower but does not fail suddenly, as can happen with Ethernet.
- Token Ring uses an *active topology*, in which each computer in the circle regenerates the signal as the signal travels around the ring. Ethernet networks require *repeaters* to boost the signal as distance increases.

AppleTalk Networks

Apple Computers developed the AppleTalk protocol suite to network its Macintosh machines for file and print sharing in a workgroup environment. AppleShare is a suite of application layer protocols that provide this functionality. AppleShare components, which are built into the Macintosh operating system, include the following:

- **AppleShare File Server**—Enables users to access the computer's resources
- **AppleShare Print Server**—Provides for sharing of printers

AppleShare PC is a service that runs on DOS computers to enable them to access files on an AppleShare file server or print to a shared AppleShare print server.

AppleTalk networks are often referred to as *LocalTalk* networks. The LocalTalk Link Access Protocol (LLAP) is the basic protocol used, and it supports dynamic addressing.

AppleTalk networks can be divided into groups, called *zones*, and serve a purpose similar to dividing a large network into workgroups. Users only see those shared resources that are in the zone to which he or she belongs. Routers use the Zone Information Protocol (ZIP) to communicate between zones.


ARCnet Networks

Attached Resource Computer Network (ARCnet) is an older LAN architecture that has, for the most part, been replaced by Ethernet and Token Ring. ARCnet uses a token-passing access method, but implements the network topology as a bus or star instead of as a ring.

In an ARCnet network, the token is passed in numerical order according to the *node address*, which is an eight-digit binary number set on the ARCnet NIC using DIP switches

or jumpers. If the cards are not numbered carefully so that the signal travels in a logical order—or if computers are later moved to different locations without the addresses being changed—you can end up with the signal traveling greater distances than necessary. Figure 2-16 illustrates this problem.

Figure 2-16 *If the addresses on ARCnet NICs are not set properly, the path taken by the token can be very inefficient.*

**NOTE**

ARCnet networks are slow compared to most LAN technologies: 2.5 Mbps for standard ARCnet, although a newer standard called *ARCnet Plus* improves on this considerably, with a top speed of 20 Mbps.

ARCnet can use coax, UTP, or even fiber-optic cable, but it is most commonly associated with RG-62/U 90-ohm coax cable. One of the drawbacks of ARCnet, and a reason it is no longer a popular choice for local-area networking, is its very proprietary nature. It is, however, a stable and reliable architecture.

Summary

In this chapter, we have looked at many different ways in which networks can be categorized, including the following:

- The physical scope of the network
- The administrative model used to manage the network
- The NOS run on the servers that control the network
- The protocols used by the network to communicate
- The topology or layout of the network
- The architecture, which encompasses the standards and specifications by which the network operates

There are, of course, other ways in which networks can be categorized. This book focuses on networking PCs, but in the bigger picture, networks could also be categorized as PC networks and mainframe networks.

Many of the categories we have discussed can be broken down further; for instance, UNIX networks can be broken down into HP-UX networks, SUN Solaris networks, AIX networks, SCO networks, and so on. Microsoft networks can be divided into Windows 2000 networks, Windows NT 4.0 networks, and Windows NT 3.51 networks—not to mention Windows 9x workgroups.

As we examined each method of categorization, you might have noticed that certain architectures are associated with certain topologies, or that particular NOSs use specific administrative models, or that some protocols are used more often in networks of a particular physical scope. In later chapters, we build on these observations.

You will learn about the different signaling methods used by different technologies and the different cable types associated with different media access methods. Chapter 3 provides a look at the popular models used to describe the complex process in which all these components engage. We also discuss the specifications and standards on which networks are based.

Further Reading

A good resource for information on Gigabit Ethernet is the Gigabit Ethernet Consortium homepage: www.iol.unh.edu/consortiums/ge/main.html.

GEC2000 is an annual conference addressing Gigabit Ethernet, most recently held March 27–29, 2000, in San Jose, California. See the GEC2000 Web site, at www.etherconference.com/English/Main_Splash.html.

An excellent resource for information about Hewlett Packard's 100BaseVG-AnyLAN is available at www.100vg.com.

For general information on how LANs operate, see www.net-engineer.com/lanbasics1.html.


Review Questions

The following questions test your knowledge on the material covered in this chapter. Be sure to read each question carefully and select the *best* correct answer or answers.

- 1 Which of the following is a category used when classifying networks according to the physical scope of the network?
 - a 10Base2
 - b WAN
 - c IPX/SPX
 - d Ring
- 2 Which of the following is true of the typical LAN? (Select all that apply.)
 - a A LAN is a large network comprised of many small networks.
 - b A LAN typically spans a dispersed geographic area, with nodes in different cities, states, or countries.
 - c A LAN is typically confined to a limited area, with the computers in close physical proximity.
 - d LANs can be divided into workgroups for easier management.
- 3 What is a group of networks that are networked to each other called?
 - a An internet
 - b An intranet
 - c An extranet
 - d A workgroup
- 4 Which of the following is true of a client/server (or server-based) network but not true of a peer-to-peer network? (Select all that apply.)
 - a Client/server networks are generally less expensive to implement.
 - b Client/server networks provide for centralized data storage, making backups easier.
 - c Client/server networks generally require a dedicated network administrator.
 - d Client/server networks offer better security.

- 5 Which of the following terms is used to identify a logon authentication server in a Microsoft client/server network?
- a Daemon
 - b Security Accounts Manager
 - c Domain controller
 - d Cluster server
- 6 Which of the following is true of security in a peer-to-peer (workgroup) environment? (Select all that apply.)
- a Security is centralized and administered by a network administrator.
 - b Security is stronger and more difficult to defeat than in a client/server environment.
 - c Each machine in the workgroup maintains its own local security database.
 - d Security is implemented at the share level rather than at the user level.
- 7 Which of the following is a characteristic of Windows 2000 domains that is a change from the Windows NT domain model? (Select all that apply.)
- a Windows 2000 domains are based on a directory service called the Active Directory.
 - b Windows 2000 domains are based on a directory service called the Bindery.
 - c Windows 2000 domains are referred to as downlevel domains.
 - d Windows 2000 domains can contain multiple domain controllers, all of which can read and write to the directory database.
- 8 Which of the following is true of TCP/IP networks?
- a TCP/IP networks are typically medium- to large-sized networks that contain multiple subnets.
 - b The global Internet is a TCP/IP network.
 - c Few tools and utilities are available for TCP/IP networks.
 - d TCP/IP networks are easier to administer than NetBEUI networks.

- 9 What is the name for the network topology that has a beginning and an end and that requires termination to prevent signal bounce?
- a Star
 - b Hybrid mesh
 - c Linear bus
 - d Token Ring
- 10 Which of the following network architectures uses unshielded twisted-pair cabling?
- a 10Base2
 - b 10BaseT
 - c 100BaseFL
 - d Token Ring



LAN Links

In the first four chapters of this book, you learned about the purpose and structure of networks, networking models and standards, how signals are sent across network media, and the methods used for controlling access to those media.

In this chapter, we put together some of these concepts as we discuss how they all fit into the specifications that make up a particular type of network link (also referred to as the networking *architecture*).

Generally, we can divide link types into two categories:

- Those used to connect local-area networks (LANs) or metropolitan-area networks (MANs); that is, networks of limited geographic scope
- Those used to connect wide-area networks (WANs), which span cities, countries, or even continents

In this chapter, we look at the characteristics of the first link type, and how to best make choices about which is easiest, most cost effective, and otherwise the optimum solution for a given situation. The second type is covered in Chapter 6, “WAN Links.”

There are several link types used for connecting computers over a LAN or MAN. Because these networks span relatively short distances, cable types can be used that would not work with WANS because of their susceptibility to attenuation.

You are already familiar with most of the LAN architectures from earlier chapters, but this chapter outlines the specifications and standards for each, including:

- Ethernet
- Token Ring
- Fiber Distributed Data Interface (FDDI)
- AppleTalk
- ARCnet

Ethernet

The Ethernet architecture is the most popular type of LAN link today. It is based on the 802.3 standard, which specifies a network that implements the CSMA/CD access control method using baseband transmission over coaxial or twisted-pair cable laid out in a bus topology (that is, a linear or star bus). Standard transfer rates are 10 Mbps or 100 Mbps, but new standards provide for *Gigabit Ethernet*, capable of attaining speeds up to 1 Gbps over fiber-optic cable or other high-speed media.

Ethernet originated with the ALOHA WAN at the University of Hawaii, which used the CSMA/CD access control method. In the 1970s the Palo Alto Research Center (PARC), owned by Xerox, developed a 2.95-Mbps Ethernet design. Soon afterward, Xerox, Intel, and Digital collaborated on a standard for Ethernet that transmitted at 10 Mbps.

Today there are numerous Ethernet topologies in use, including the following widely implemented forms:

- 10Base2 (thinnet)
- 10Base5 (thicknet)
- 10BaseT (UTP)
- 100BaseT (Fast Ethernet)
- 100BaseFX (Ethernet over fiber-optic cable)
- 1000BaseT (Gigabit Ethernet)


10Base2

The 10Base2 Ethernet topology implements the CSMA/CD access method in a linear bus layout as specified in IEEE standard 802.3, using thin coax cable. This topology is also referred to as *thinnet*.

The 10 in the name refers to the transfer speed of thinnet, which is 10 Mbps. Ethernet is a baseband transmission network. The 2 is a little trickier to figure out; it indicates a rough approximation of the maximum cable segment length for this topology, which is 185 meters (rounding that up to 200 gives you the 2).

A 10Base2 network is limited to 30 *nodes* (that is, computers or other network devices) per 185-meter segment. Thinnet uses barrel connectors and T-connectors (both of which are called *BNC* connectors) and BNC-type terminators on each end of the bus. Figure 5-1 illustrates how a barrel connector is used to extend the length of the cable.

Figure 5-1 A BNC barrel connector can be used to extend the length of a piece of thin coax cable.



NOTE

There is some disagreement in the networking industry about the origins of the abbreviation BNC. You might see it defined in some texts as *British Naval Connector*. Others define it as *Bayonet Nut Connector*. Many sources attribute the abbreviation to the names of those who designed it: Neil and Concelman. Among those who advocate the last definition, there is further contention over whether the name is spelled *Neil* or *Neill*. Research indicates that Paul Neill, an employee of Bell Labs, created a connector he called an “N” connector, which was used by the British Navy. Carl Concelman later adapted the design to create a bayonet-mount connector, which he called the “C” connector. Finally, Neill and Concelman got together and created a smaller version, the BNC connector.

Advantages of 10Base2

The 10Base2 topology is easy to install and configure. For this reason, it is appropriate for small temporary networks (such as classrooms that must be set up and torn down regularly). 10Base2 uses less cable than 10BaseT, and the cable it uses costs less per foot than that which is used with 10Base5. It does not require extra components such as hubs or external transceivers. Thus, 10Base2 is a relatively inexpensive network to implement.

Disadvantages of 10Base2

Unfortunately, 10Base2 is not appropriate for large networks because of the limitations on cable segment length and the number of nodes per segment. In addition, its transfer speed, 10 Mbps, is relatively slow in today’s bandwidth-intensive world.

NOTE


Network specifications mandate a maximum cable segment length because a signal loses strength over distance (attenuation). Minimum length specifications are also important because if the cable segment between two computers is too short, the signals of the two devices will overlap, resulting in an abnormal voltage on the cable that is called a *collision*.

10Base5

10Base5 is sometimes referred to as *standard Ethernet*. It is also called *thicknet*. Thicknet cable is approximately one-half inch in diameter, which is about twice the thickness of thinnet. Thicknet uses the linear bus topology, but because of the thicker coax cable used, it is capable of transmitting for a greater distance (500 meters) without signal loss. Thicknet specifications mandate a minimum segment length of 2.5 meters.


Setting up a network with thicknet is more complex than doing so with thinnet. 10Base5 uses external transceivers, which provide for transmission and reception. They are attached to the network card through a *transceiver cable*, using AUI (DIX) 15-pin connectors, and to the coax cable through a *vampire tap*. See Figure 5-2 for an illustration of this connection.

Figure 5-2 A NIC attaches to thicknet through an external transceiver and a vampire tap.



10Base5 is often used as a *backbone*, with the thicknet used to span long distances, such as connecting repeaters on different floors, to which thinnet is then attached and the computers on each floor are connected to the thin coax. Figure 5-3 shows how this works.

Figure 5-3 *Thick coax can be used as a backbone in combination with thinnet segments.*



Like other Ethernet networks, 10Base5 uses the CSMA/CD media access method and baseband transmission. A length of cable can be extended using an N-series barrel connector, and the bus is terminated with N-series terminators.

Advantages of 10Base5

Thicknet offers two big advantages over thinnet:

- Capability to span a greater distance between repeaters
- Capacity for more nodes (computers) on each segment

An advantage of thick (and thin) coax over UTP is the partial protection against EMI that the thick outer jacket affords.

Disadvantages of 10Base5

Thick coax does have some disadvantages. Its inflexibility makes it difficult to work with, and adding stations requires drilling into the cable to make the tap.

Thicknet is also relatively expensive. Not only does the thick cable cost more than thin coax or UTP, but also you must purchase the additional equipment (for example, transceiver and transceiver cable) for each computer on the network.

10BaseT

Currently, 10BaseT is one of the most popular Ethernet implementations. You could say it is the star of the PC networking world—especially because it uses a star bus topology.

NOTE Inside the Ethernet hub, the signaling system is a bus, as with coax Ethernet networks.

In fact, you will probably hear the term *Ethernet cable* used to describe the unshielded twisted-pair (UTP) cabling generally used in this architecture (shielded twisted-pair [STP] also can be used). 10BaseT and its cousin, 100BaseX, make for networks that are easy to set up and expand.

Advantages of 10BaseT

Networks based on the 10BaseT specifications are relatively inexpensive. Although a hub is required if you are connecting more than two computers, small hubs are available at a low cost, and 10BaseT network cards are inexpensive and widely available.

NOTE 10BaseT specifications require a hub. However, if you wish to connect only two computers (for example, for a home network), and you want to use UTP rather than thinnet, you can do so using a *crossover cable*. This is a type of cable in which the wire pairs are cross-connected, and it is used also to connect two hubs to each other if the hubs do not have uplink ports.

Twisted-pair cabling, especially the UTP mostly commonly used, is thin, flexible, and easier to work with than coax. It uses modular RJ-45 plugs and jacks, so it is literally a “snap” to connect the cable to the NIC or hub.

Another big advantage of 10BaseT is upgradability. Although by definition a 10BaseT network runs at 10 Mbps, by using Category 5 or above cable and 10/100-Mbps dual-speed NICs, you can easily upgrade to 100 Mbps by simply replacing the hubs.

Disadvantages of 10BaseT

The maximum length for a 10BaseT segment (without repeaters) is only 100 meters (about 328 feet). The UTP used in such a network is more vulnerable to EMI and attenuation than other cable types. Finally, the extra cost of a hub may make this solution slightly more expensive than a thin coax network.

100BaseX

The high-bandwidth demands of many modern applications, such as live video conferencing and streaming audio, have created a need for speed, and many networks require more throughput than is possible with 10-Mbps Ethernet. This is where 100BaseX, also called *Fast Ethernet*, comes into play.

100BaseX comes in several different flavors. It can be implemented over 4-pair Cat 3, 4, or 5 UTP (100BaseT), over 2-pair Cat 5 UTP or STP (100BaseTX), or as Ethernet over 2-strand fiber-optic cable (100BaseFX).

Advantages of 100BaseX

Regardless of the implementation, the big advantage of 100BaseX is high-speed performance. At 100 Mbps, transfer rates are 10 times that of 10Base2, 10Base5, and 10BaseT.

Because it uses twisted-pair cabling, 100BaseX also shares the same advantages enjoyed by 10BaseT: low cost, flexibility, and ease of implementation and expansion.

Disadvantages of 100BaseX

100BaseX shares the disadvantages of 10BaseT, which are inherent to twisted-pair cabling, such as susceptibility to EMI and attenuation. 100-Mbps NICs and hubs are generally somewhat more expensive than those designed for 10-Mbps networks, but prices have dropped as 100BaseX has gained in popularity.

Fiber-optic cable remains an expensive cabling option, not so much because of the cost of the cable itself, but because of the training and expertise required to install it. We discuss cabling more in Chapter 6.

1000BaseT

If 100BaseX is known as *Fast Ethernet*, this new addition to the Ethernet family, 1000BaseT, must be considered a speed demon. Its common nickname is *Gigabit Ethernet*. Although not yet in widespread implementation in production networks, this architecture supports data transfer rates of 1 Gbps, which is almost seven times faster than a T-1 line. Gigabit Ethernet is, for the most part, a LAN architecture, although its implementation over fiber-optic cable makes it suitable for MANs as well.

Advantages of 1000BaseT

The greatest advantage of 1000BaseT is, of course, performance. At 1 Gbps, it is 10 times as fast as Fast Ethernet and 100 times as fast as standard Ethernet. This makes it possible to implement bandwidth-intensive applications, such as live video, throughout an intranet.

Disadvantages of 1000BaseT

The only disadvantages associated with 1000BaseT are those common to all UTP networks, as detailed in the sections on 10BaseT and 100BaseT.

The Structure of an Ethernet Frame


The package into which data is broken for transmission is called a *frame* in Ethernet terminology. The frame contains header information, the data being transmitted, and trailer information.

Ethernet Frame Types

There are four different frame types (Ethernet_802.2, Ethernet_802.3, Ethernet II, and Ethernet_SNAP) associated with Ethernet communications. The two frame types generally used on a TCP/IP network are Ethernet_II and Ethernet_SNAP (Subnetwork Access Protocol). The frame type is defined by the structure of the frame (that is, the number of bytes allocated to each field in the header, data field, and trailer field). For example, the Ethernet_II header is simpler, containing only four fields. The Ethernet_SNAP header contains nine fields.

Figure 5-4 illustrates the structure of an Ethernet frame (Ethernet II type).

Figure 5-4 An Ethernet II frame contains four header fields and a trailer.



The first header field, which is simply an indicator of the beginning of the frame, is the *preamble*. The next two fields contain the destination and source addresses (that is, the receiving and sending computers), and the fourth field is used to identify the protocol operating at the network layer (typically IP or IPX). The data follows, and can be from 46 to 1500 bytes in length. Ethernet uses an additional 18 bytes for construction of the frame itself. Finally, the trailer contains a cyclic redundancy check (CRC), which is a calculation used to verify that the frame that is received matches what was sent.

Checking for Errors Using CRC

The CRC is a standard way of checking for errors in the data after it has been transmitted over a network. Here's how it works: The computer that is transmitting the data treats the data as one long polynomial. A polynomial is an algebraic expression consisting of one or more summed terms, each term consisting of a constant multiplier and one or more variables raised to integral powers. The transmitting computer divides that polynomial by a predefined 16- or 32-bit polynomial. The quotient from that division is the CRC, and it is appended to the block of data before transmission. Then the computer on the receiving side uses the same predefined 16- or 32-bit polynomial, which it applies to the incoming data. The receiving system compares its result with the result that was appended to the data by the sending system. If the two are the same, the receiving computer concludes that the data has been received successfully. If they are different, the sending computer is notified so that it can resend that block of data.

Rules of Ethernet Engagement

To ensure that the network operates properly, it is important to comply with the rules of engagement laid forth in the specifications for each implementation. These are restrictions and limitations imposed by the characteristics of the media. Note that it may be possible to break or bend these rules and still have functional network communication; however, the network would not be in compliance with standards, and you could experience connectivity problems from pushing beyond the limits.

We discuss two important rules in this section:

- The 5-4-3 rule
- 10BaseT node capacity limitations

The 5-4-3 Rule

One well-known limitation on coax-based Ethernet networks is referred to as the 5-4-3 rule. The numbers represent the following limitations:

- **5**—The maximum number of cable segments allowed in a thinnet or thicknet network is five (each segment can be up to 185 meters for thinnet or 500 meters for thicknet).
- **4**—The maximum number of repeaters that can be used to connect the segments is four.
- **3**—The maximum number of segments that can be *populated* (that is, contain nodes) is three. The other two segments are for purposes of extending the distance only.

NOTE

When applying the 5-4-3 rule to a 10Base5 (thicknet) network, the length of the thick coax cable is used to measure the distance. The transceiver cable does not count in this measurement.


5-4-3 and 10BaseT

Does the 5-4-3 rule apply to twisted-pair Ethernet networks? Many books lead you to believe it doesn't. However, the limitation on the number of hubs that can be uplinked is really an application of this same rule.

Applying the 5-4-3 rule to a UTP star-bus network means that a populated hub, for purposes of this rule, counts as a segment, and the maximum number of cascaded populated hubs allowed in a twisted-pair network is three.

If you think about the fact that the wiring inside each hub is a bus with computers connected to it, just as in a cable segment between repeaters on a linear bus network, this makes sense, as shown in Figure 5-5. This is the reason the topology is called a *star bus*.

Figure 5-5 Each hub to which computers are connected represents a populated segment.



5-4-3 and Expansion of the Network

At this point, you might be wondering how large Ethernet networks can ever be created within the parameters of this rule. Although it is possible to have hundreds of computers connected through twisted-pair or coax cabling, how do those networks circumvent the 5-4-3 limitations?

To understand how Ethernet networks can be expanded, we must first really understand the rule itself. The 5-4-3 rule is often stated as “no more than five segments, four repeaters, and three populated segments per *network*.” In this case, a “network” is defined as a *collision domain*, that is, a network within which there will be a collision if two computers transmit at the same time.

The reality of the collision domain is the key to connecting more devices than the 5-4-3 rule would seem to allow. We can create many separate collision domains on a network by dividing the network by using connectivity devices such as switches and routers. In Chapter 7, we look in detail at how these devices work.

10BaseT Node Capacity

The number of nodes per segment on a twisted-pair star-wired bus is always exactly two: the computer or device at one end of the cable and the hub at the other. The maximum number of nodes in a hub-based network should not exceed 100 because of the exponential increase of collisions after 100 nodes. This limit can be extended with connectivity devices

such as bridges and switches. The theoretical total number of network devices on a 10BaseT network is 1024.

Summary of Ethernet Characteristics

Table 5-1 is a brief summary of specifications for the most widely used Ethernet implementations.

Table 5-1 *Ethernet Specifications Summary*

	10Base2	10Base5	10BaseT	100BaseX
Cable Type	Thin coax RG-58 A/U	Thick coax RG-8 or RG-11	UTP Cat 3, 4, 5, and 5e	UTP Cat 3, 4, 5, and 5e
Connector Type	BNC connector	AUI/DIX (to transceiver)	RJ-45 modular	RJ-45 modular
Maximum Segment Length	185 meters (607 ft)	500 meters (1640 ft)	100 meters (328 ft)	100 meters (328 ft)
Maximum Network Length	925 meters (3035 ft)	2500 meters (8200 ft)	Star bus topology	Star bus topology
Nodes per Segment	30	100	2 (1024 per network)*	2 (1024 per network)*
Transfer Rate	10 Mbps	10 Mbps	10 Mbps	100 Mbps

* Because a hub is used as a central connection point with UTP-based networks, each segment of cable has only the computer or network device at one end and the hub at the other as nodes on that segment.

Ethernet is a popular LAN architecture that works with almost all popular network server and client operating systems, including Windows for Workgroups, Windows 9x, Windows NT, Windows 2000, Novell NetWare, AppleTalk/AppleShare, and UNIX/Linux.

Token Ring

Token Ring was originally developed by IBM and was designed to be a reliable network architecture based on the token-passing access control method. It is often integrated with IBM mainframe systems such as the AS/400, and it was intended to be used with PCs, minicomputers, and mainframes. It works well with Systems Network Architecture (SNA), the IBM architecture used for connecting to mainframe networks.

The Token Ring standards are provided in IEEE 802.5.

Token Ring Topology

Token Ring's topology can be confusing at first. It is a prime example of an architecture whose physical topology is different from its logical topology. The Token Ring topology is referred to as a *star-wired ring* because the outer appearance of the network design is a star, with computers connecting to a central hub, called a multistation access unit (MSAU). Inside the device, however, the wiring forms a circular data path, creating a logical ring.

Token Ring is so named because of its logical topology and its media access control method: token passing. The transfer rate for Token Ring can be either 4 Mbps or 16 Mbps.

Token Ring is a baseband architecture using digital signaling. In that way it resembles Ethernet, but the communication process is quite different in many respects. Token Ring is an *active* topology: As the signal travels around the circle to each network card, it is regenerated before being sent on its way.

The Token Ring Communication Process

In an Ethernet network, all computers are created physically equal. At the software level, some may act as servers and control network accounts and access, but the servers communicate physically on the network in exactly the same way as the clients.

The Monitor of the Ring

In a Token Ring network, the first computer that comes online becomes the "hall monitor" and must keep track of how many times each frame circles the ring, and it has the responsibility of ensuring that only one token is out on the network at a time.

The monitor computer periodically sends a signal called a *beacon*, which circulates around the ring as does any other signal. Each computer on the network looks for the beacon. If a computer does not receive the beacon from its nearest active upstream neighbor (NAUN) when expected, it puts a message on the network that notifies the monitoring computer of the beacon that was not received, along with its own address and that of the NAUN that failed to send when expected. In many cases, this caused an automatic reconfiguration that restores the communications.

Data Transfer

A Token Ring network uses a *token* (that is, a special signal) to control access to the cable. A token is initially generated when the first computer on the network comes online. As you learned in Chapter 4, when a computer wants to transmit, it waits for and then takes control of the token.

The token can travel in either direction around the ring, but only in one direction at a time. The hardware configuration determines the direction of travel.


Token Ring Hardware Components

Token Ring networks use special hardware. One reason Token Ring is less widely implemented than Ethernet is the higher cost of the Token Ring components.

The MSAU

The Token Ring hub is called a MSAU or smart MSAU (SMSAU). The distinguishing characteristic of an MSAU is the ring configuration of the wiring inside. Multiple MSAUs can be joined together, as long as the “out” port of one is connected to the “in” port of the next, maintaining the integrity of the ring topology. The data must travel in a continuous circle.

Figure 5-6 *Multiple MSAUs can be connected, as long as the ring is maintained by ensuring that the ring cable goes in and out properly.*



The Cabling and Network Card

Token Ring uses IBM type 1, 2, or 3 cabling. IBM labels its twisted-pair cabling as “types” rather than as “categories.” See Table 5-2 for complete descriptions.

Table 5-2 *IBM Cable Types*

IBM Cable Type	Description
Type 1	STP; two pairs of wires with outer shield
Type 2	STP; two pairs of wires for data and four pairs of wires for voice
Type 3	UTP; uses RJ-45 or RJ-11 connectors; four wire pairs
Type 5	Fiber-optic cable
Type 6	STP; used as data patch cable
Type 9	Plenum-grade STP cable

NOTE IBM cable uses a special IBM connector, called a type A connector, that is not compatible with standard BNC connectors.

Token Ring network cards, such as Ethernet UTP cards, come in different speeds, designed to run at either 4 Mbps or 16 Mbps. If the network runs at 4 Mbps, you can use a 16-Mbps card (although it will run only at the slower speed).

WARNING You *cannot* use a card designed to run at 4 Mbps on a 16-Mbps Token Ring network. The computer will not be able to communicate at all.

Other Components

A variety of other hardware components are used with Token Ring. Media interface connectors (MICs), for instance, are used to connect type 1 and type 2 IBM cable. Media filters are used to connect a Token Ring NIC to a modular jack and to reduce noise on the line.

Repeaters can be used to regenerate the digital signal, as in Ethernet networks, and extend the length of the network.

Advantages of Token Ring

Token Ring is a highly reliable architecture, and its token-passing scheme eliminates data collisions. Additionally, the MSAUs can detect the failure of a network card and automatically disconnect it from the ring, thus enabling the token to continue around the ring. Hence, the network is not brought down by the failure of one computer.

Another advantage is the capability to easily interoperate with PC and mainframe networks.

Disadvantages of Token Ring

The major disadvantages of the Token Ring architecture are its higher cost when compared to 10BaseT or 100BaseT and its slower relative speeds.


The Structure of the Token Ring Frame

Token Ring networks use three different types of frames:

- **Data Frame**—Carries information being sent from one computer to another
- **Token Frame**—Circulates on the network until it is captured by a computer that is ready to send information
- **Management Frame**—Transmits error or other management information

The Token Ring data frame structure is a bit more complex than that of the Ethernet II data frame. Figure 5-7 illustrates a Token Ring data frame, which is sometimes also called an LLC (logical link control) frame.

Figure 5-7 *The Token Ring data frame is more complex than the standard Ethernet frame.*



The *start delimiter* in a Token Ring frame serves the same purpose as the preamble on an Ethernet frame. The other fields function as follows:

- **Access control field**—Signifies whether the frame is a token or a data frame, and its priority
- **Frame control field**—Contains access control information
- **Destination address**—The address(es) of the computer(s) to which the data is being sent
- **Source address**—The address of the sending computer
- **Data**—The actual information being sent (for example, the contents of an e-mail message)
- **Frame check sequence (FCS)**—CRC error-checking bits
- **End delimiter**—Signifies the end of the frame
- **Frame status field**—Indicates whether the address was recognized and the frame copied (marked by the receiving computer before being sent back around the ring)

NOTE The token frame (sometimes just called the “token”) circulates on the network until it is captured by a computer that is ready to transmit. A token frame has only three fields: the start delimiter, the access control field, and the end delimiter.

Rules of the Ring

Specifications for Token Ring place some limitations on its implementation:

- Distance between MSAUs is limited to 365 meters for type 3 cable and to 730 meters for type 1 or 2 cable.
- Maximum length for a cable segment depends on the cable type and varies from 45 to 200 meters.
- Maximum length for type 6 patch cable is 46 meters.
- Minimum segment length is 2.5 meters.
- Maximum number of computers per segment, per IBM specifications, is 72 with unshielded cable and 260 with shielded.
- Maximum number of segments (that is, MSAUs) that can be connected is 33 (again, according to IBM specifications).

Summary of Token Ring Characteristics

Token Ring is a reliable technology, but less popular than Ethernet because of the relatively higher cost and slower transfer rates. However, specifications are being developed for 100-Mbps Token Ring and Gigabit Token Ring, so this architecture should not be counted out just yet.

FDDI


FDDI is a type of Token Ring network. Its implementation and topology differ from IBM’s Token Ring LAN architecture, which is governed by IEEE 802.5. FDDI is often used for MANs or larger LANs, such as those connecting several buildings in an office complex or campus.

How FDDI Works

As its name implies, FDDI runs on fiber-optic cable, and thus combines high-speed performance with the advantages of the token-passing ring topology. FDDI runs at 100 Mbps, and its topology is a *dual ring*. The outer ring is called the *primary ring* and the inner is the *secondary ring*.

Normally, traffic flows only on the primary ring, but if it fails, the data automatically flows onto the secondary ring, in the opposite direction. The network is said to be in a *wrapped state* when this occurs. This provides fault tolerance for the link. Figure 5-8 illustrates how FDDI works.

Figure 5-8 FDDI uses a dual-ring fault-tolerant topology.



Computers on a FDDI network are divided into two classes:

- **Class A**—Computers connected to the cables of both rings
- **Class B**—Computers connected to only one ring

Another difference between FDDI and 802.5 Token Ring is the allowed number of frames on the network. A FDDI network enables multiple frames to circulate on the ring simultaneously. More than one computer can transmit at the same time. This is called *shared network technology*.

Like 802.5 Token Ring, FDDI uses beaconing to detect and isolate problems within the ring.

FDDI Specifications

A FDDI dual ring supports a maximum of 500 nodes per ring. The total distance of each length of the cable ring is 100 kilometers, or 62 miles. A repeater is needed every 2 kilometers, which is why FDDI is not considered to be a WAN link.

NOTE These specifications refer to FDDI implemented over fiber-optic cable. It is also possible to use the FDDI technology with copper cabling. This is called *Copper Distributed Data Interface (CDDI)*. The maximum distances for CDDI are considerably lower than those for FDDI.

FDDI's ring topology can be implemented as a physical ring or as a star-wired logical ring by using a hub.

Advantages of FDDI

FDDI combines the advantages of token passing on the ring topology with the high speed of fiber-optic transmission. Its dual ring topology provides redundancy and fault tolerance. The fiber-optic cable is not susceptible to EMI and noise, and it is more secure than copper wiring. It can send data for greater distances between repeaters than can Ethernet and traditional Token Ring.

Disadvantages of FDDI

As always, high speed and reliability come with a price. FDDI is relatively expensive to implement, and its distance limitations, though less restrictive than those of other LAN links, make it unsuitable for true WAN communications.

The FDDI Frame Structure

The FDDI frame structure is similar to that of Token Ring; however, FDDI uses a timed token protocol, while Token Ring uses a priority/reservation token access method. The maximum size for a FDDI frame is 4500 bytes.

The FDDI token is a special frame consisting of three octets. A computer that has data to transmit waits for the token frame, takes possession of it, transmits one or more data frames, and then releases the token.

Summary of FDDI Characteristics

FDDI is a good choice for medium-sized networks—MANs and large LANs. It is good for networks that require high bandwidth, such as engineering, graphics, and video applications.

AppleTalk

AppleTalk networking is designed, not surprisingly, to connect Apple computers, and its software components are built into the Macintosh operating systems. AppleTalk is the name of the architecture. It comes in two varieties: Phase 1 (included in early versions of the Mac OS) and Phase 2 (the current release).

LocalTalk is the term used to refer to the cables, the hardware, and the data link layer protocol used in AppleTalk networks.

AppleTalk Specifications

AppleTalk/LocalTalk networks use the CSMA/CA media access control method. STP cabling is most commonly used, but it is also possible to use UTP or fiber-optic cable. The network topology is a bus or tree, and a LocalTalk network (using Apple components) is limited to 32 nodes.

AppleShare is the protocol that provides file and print sharing on an AppleTalk network.

NOTE

Vendors besides Apple also offer cabling and connection components that enable more computers on an AppleTalk network.

How AppleTalk Works

AppleTalk networks use an addressing scheme in which each computer that comes on the network first looks for a stored address (one that it used in a previous session). If it finds one, it uses that; if not, it assigns itself an address, chosen at random from the range of allocated addresses. Then it broadcasts the address to determine if another computer is using it. If so, it repeats the process. If not, it stores the address to be used again the next time it comes on the network.

AppleTalk was designed for small networks. However, separate networks can be connected to one another. Each subnetwork is called a *zone* and has a zone name to identify it. Zones function somewhat like workgroups. You can access resources in a different zone by clicking the zone name.

AppleTalk networks can also be connected to networks using other architectures, such as Ethernet or Token Ring. Apple provides EtherTalk and TokenTalk, which are cards that enable Macintosh computers to connect to networks operating under the 802.3 and 802.5 specifications, respectively.

Advantages of AppleTalk

Because Apple included AppleTalk in the Macintosh operating system, it is easy to implement and configure. Setting up a small workgroup of Macintosh computers is simple and inexpensive.

Disadvantages of AppleTalk

AppleTalk is not suitable for very large networks. It is very slow compared to other LAN links (230.4 Kbps), and thus not suitable for bandwidth-intensive applications.

LocalTalk Data Transmission

LocalTalk is the data link layer protocol originally used by AppleTalk. Macintosh computers using LocalTalk are linked to one another through their printer ports.

NOTE

Although AppleTalk networks originally used LocalTalk cabling, AppleTalk runs also over Ethernet. The latter has become the popular choice because LocalTalk is slow (230.4 kbps) and was designed for small, low-traffic workgroups.

Summary of AppleTalk Characteristics

Despite performance limitations, the AppleTalk protocol suite provides a quick, easy way to connect Macintosh computers in small workgroups for sharing of files, printers, and other resources.

ARCnet

Attached Resource Computer Network (ARCnet) is an older LAN technology that uses a special token-passing access control method over a star-bus topology.


Some books associate ARCnet with the 802.4 token bus specifications, but actually ARCnet was developed before the 802 standards and does not map completely to this specification. It was designed for small workgroups and is a very simple and inexpensive architecture.

How ARCnet Works

The token-passing method used by ARCnet differs from that of Token Ring. On an ARCnet network, the token does not move along a path based on physical location. It instead goes from one computer to another based on that computer's position in the assigned numbering

order. This makes ARCnet an inefficient transmission medium in that the data may travel a much longer path than necessary to reach its destination. Figure 5-9 illustrates how ARCnet works.

Figure 5-9 ARCnet's token-passing method can result in data taking an inefficient route to its destination.



ARCnet Specifications

ARCnet uses a hub to connect computers in a star bus configuration, and it generally uses 93-ohm coax cable specified as RG-62 A/U. However, ARCnet can run also on twisted-pair or fiber-optic cabling.

Using the more common coax cable, the maximum segment length (computer to hub) is 610 meters. Using UTP, the maximum segment length drops to 244 meters.

NOTE ARCnet can use a linear bus topology, instead of a star, with coax cable, in which case, the maximum segment length is 305 meters.

Standard ARCnet suffers from slow performance; the transfer rate is only 2.5 Mbps. A later implementation, ARCnet Plus, can transmit at 20 Mbps.

Advantages of ARCnet

ARCnet is inexpensive and works adequately for small networks. ARCnet Plus offers a higher data transmission rate than standard Ethernet.


Disadvantages of ARCnet

ARCnet is an older technology that is not as popular as Ethernet and Token Ring. Its networking components are not as widely available, and it is more difficult to find technicians who are experienced in working with ARCnet technology. Even in its faster variety, it is not nearly as fast or efficient as 100BaseT, and it is not suitable for high-bandwidth applications.

The ARCnet Packet Structure

A unit of data transmitted across the network is called a *packet* on an ARCnet network. This equates to the Ethernet or Token Ring *frame*. The standard ARCnet packet is very simple in structure, as shown in Figure 5-10.

Figure 5-10 *The ARCnet packet contains addressing information and data.*



As you can see, the header and trailer fields contain the addresses of the source and destination computers, with the data between the addresses. A standard ARCnet packet is limited to 506 bytes of data, and ARCnet Plus packets can contain up to 4096 bytes.

Summary of ARCnet Characteristics

ARCnet is a simple technology that is being replaced, in many environments, by Ethernet or Token Ring. Although ARCnet is relatively inexpensive and provides adequate performance for small, low-bandwidth networks, its limited transmission speed and inefficient token-passing method have led to a decline in popularity.

Comparing Networking Architectures

Table 5-3 summarizes the characteristics of the networking architectures.

Table 5-3 *Networking Architectures Compared*

Networking Architecture	Cabling Type(s)	Transfer Speed	Access Method	Topology
Ethernet 10Base2	Thin coaxial	10 Mbps	CSMA/CD	Linear bus
Ethernet 10Base5	Thick coaxial	10 Mbps	CSMA/CD	Linear bus
Ethernet 10BaseT	UTP Cat 3–5	10 Mbps	CSMA/CD	Star
Ethernet 100BaseT	UTP Cat 5	100 Mbps	CSMA/CD	Star
Ethernet 100BaseFX	Fiber-optic	100 Mbps	CSMA/CD	Star
Ethernet 1000BaseT	UTP Cat7	1 Gbps	CSMA/CD	Star
Token Ring	STP or UTP	4 Mbps or 16 Mbps	Token passing	Physical star, logical ring
FDDI	Fiber optic	100 Mbps	Token passing	Dual ring
AppleTalk	STP or UTP	LocalTalk: 230.4 Mbps Ethernet: 10 Mbps	CDMA/CA	LocalTalk: Linear bus Ethernet: Star

Summary

In this chapter, we examined the characteristics, specifications, and implementations of several popular LAN architectures.

You learned that Ethernet is one of the most popular LAN links and that it is based on the CSMA/CD access method and governed by IEEE 802.3 standards. We discussed the many varieties of Ethernet, including:

- 10Base2 (thinnet)
- 10Base5 (thicknet)
- 10BaseT, 100BaseT, and 1000BaseT

We also discussed the Token Ring architecture governed by IEEE 802.5 and how communication occurs on a Token Ring network. Next, we looked at FDDI, a token-passing architecture that runs over fiber-optic cable and that is appropriate for large LANs and MANs. Then we took a look at AppleTalk and the components that make up the

architecture: LocalTalk, AppleShare, EtherTalk, and TokenTalk. Finally, we discussed ARCnet, an older technology that uses a unique token-passing scheme in which data follows a logical, rather than physical, path.

In the next chapter, we will expand our horizons a bit and look beyond the LAN to the concepts and practices of wide area networking. We will discuss popular WAN links and how they are implemented, and we will examine how you can connect your LAN to a WAN (including the widest area network of all, the Internet).

Further Reading

For more information on LAN links, see the following:

- Charles Spurgeon's Ethernet (802.3) Web site at wwwhost.ots.utexas.edu/ethernet/ethernet-home.html.
- The Token Ring Consortium at www.iol.unh.edu/consortiums/tokenring/main.html.
- The Interoperability Lab's FDDI tutorial at www.iol.unh.edu/training/fddi/htmls/index.html.
- Core Competence AppleTalk whitepaper at www.corecom.com/html/appletalk.html.
- The ARCNET trade association Web site at www.arcnet.com/home.html.

Review Questions

The following questions test your knowledge of the material covered in this chapter. Be sure to read each question carefully and select the *best* correct answer or answers.

- 1 IEEE 802.3 defines specifications for which of the following network architectures?
 - a Token Ring
 - b Ethernet
 - c ARCnet
 - d FDDI
- 2 Which media access method is used by Ethernet networks?
 - a CSMA/CD
 - b CSMA/CA
 - c Token Passing
 - d Demand priority

- 3 10Base2 Ethernet specifications limit the number of nodes on a segment of cable to what number?
 - a 30
 - b 100
 - c 185
 - d 1024
- 4 Which of the following cable types uses a BNC connector?
 - a UTP
 - b thick coax
 - c thin coax
 - d Fiber optic
- 5 The term “standard Ethernet” is used to refer to which of the following?
 - a 10BaseT networks
 - b 100BaseT networks
 - c 10Base2 networks
 - d 10Base5 networks
- 6 What is the specified maximum cable distance for a 10Base5 network?
 - a 100 meters
 - b 500 meters
 - c 100 feet
 - d 500 feet
- 7 Which of the following is true of the Token Ring architecture? (Select all that apply.)
 - a Token Ring is a passive topology.
 - b Token Ring networks run at 4 or 16 Mbps.
 - c Token Ring hubs are called multistation access units.
 - d Token Ring is a broadband architecture.

- 8 What is the name of the circulating signal sent by the monitor computer in a Token Ring network, whose purpose is to detect when a station fails so that the network can reconfigure itself to restore communications?
- a token
 - b data frame
 - c beacon
 - d NAUN
- 9 Which of the following is the data link layer protocol originally designed for AppleTalk networks?
- a Ethernet
 - b LocalTalk
 - c EtherTalk
 - d TokenTalk
- 10 What is the transfer speed for standard ARCnet?
- a 2.5 Mbps
 - b 25 Mbps
 - c 16 Mbps
 - d 10 Mbps



WAN Links

The technologies, media, and equipment that work well for the short distances spanned by a LAN or MAN are generally not suitable for long-distance wide-area networks (WANs). In today's very mobile world, high-performance, cost-effective WAN technologies are a necessity for many reasons:

- Executives and other employees need access to their corporate networks while on the road or at home.
- Companies with branch offices in widely dispersed geographic locations need network connectivity between locations.
- Organizations want to share information with other organizations physically separated by long distances.
- Commercial, governmental, and educational bodies and individuals need access to the resources available on the global Internet.

It is obviously impossible to string Ethernet cable from the home office in Denver to the branch office in Houston. Even if cabling distance limitations did not apply, this would not be a viable solution for connecting international sites.

WANs require a whole new set of technologies and rules of implementation. In this chapter, we discuss the concept of networking over long distances and the technologies commonly used to connect computers that are located in different states, countries, or even different continents. These range from the Public Switched Telephone Network (PSTN) already in place in most of the world to modern high-tech solutions such as satellite communications technologies that enable us to "talk" to computers in space.

Wide-area networking presents many challenges not encountered in implementing a network that is confined to one geographic area. A WAN is *not* just a really big LAN. Rather, it is a collection of many separate LANs, connected by links that are different in many ways from LAN links. WANs that span international boundaries require consideration of even more factors, including time zones and language differences.

Designing a WAN is a complex task. Choosing the appropriate technology involves analyzing the purpose(s) the WAN will serve, the number of users, the bandwidth requirements, and the patterns of use. We can categorize these considerations as follows:

- WAN hardware
- WAN topologies
- Network switching types
- New and emerging WAN technologies
- LAN/WAN connectivity

We look at each issue in the sections that follow.

WAN Hardware

The hardware necessary to implement a WAN link can be as simple and inexpensive as a telephone line and a modem at each end. On the other hand, it can be complex and costly. In general, equipment cost and complexity increases with increased speed and reliability.

In the following sections, we discuss common WAN devices, including modems, ISDN and digital subscriber line (DSL) terminal adapters, and customer premises equipment (CPE) used with dedicated links such as T-carrier connections and X.25.

Modems

To establish a network connection (to an Internet service provider or to a dial-up server on a private network) over public telephone lines, you use a device called a *modem*.


TIP The word “modem” is derived from the actions it performs; a modem *modulates* and *demodulates* a signal. In other words, it converts the sending computer’s digital signal to analog for transmission over the analog line and then converts it back to digital for processing by the receiving computer.

Modems come in two physical types: *internal* and *external*. Each has advantages and disadvantages, and configuration is slightly different depending on the type. Either way, modems are *serial* devices, which means bits are sent one at a time. This can be contrasted with *parallel* devices, such as printers, to which multiple bits can be sent simultaneously. A serial transmission is analogous to a group of people marching in a straight line, and a parallel transmission is like having the same group marching in rows of three across.

Internal Modems

One advantage of the internal modem is compactness. It is a circuit board card that fits in an ISA (Industry Standard Architecture) or PCI (Peripheral Component Interconnect) slot inside the computer, as shown in Figure 6-1. This means that you don't have to find room for an extra device on your desk. In addition, you are not required to buy a serial cable, which you might be forced to do if you use an external model that doesn't include one in the box.

Figure 6-1 *An internal modem is a circuit board that fits inside the computer.*



Internal Modem Configuration Parameters

Internal modems are traditionally more difficult to configure than external modems. You must set the IRQ, the input/output addresses, and the virtual com ports to ensure that they don't conflict with the settings of some other device in your computer. Let's look more closely at each setting and how it is used:

- **Interrupt Request (IRQ)**—This is an assigned location that designates where the system expects the device to interrupt it when the device sends a signal. Signals from different devices that go to the processor on the same interrupt line would interfere with each other, so a separate IRQ must be assigned to each device.
- **Input/Output (I/O) address**—This is the location where data sent from the device is stored before it is processed by the CPU. As with the IRQ, if multiple devices attempt to use the same I/O address, one or both devices might not work properly.
- **Virtual com port**—This is a logical port number, by which the operating system identifies a serial port. You must set each serial device to use a different com port.

All popular operating systems provide a means by which you can view how resources are being used, and which ones are not in use, so that you can choose free resources to assign to your new device.

Changing the Internal Modem Settings

Internal modems generally provide a way to change the configuration settings. Depending on the manufacturer and model, you can change IRQ, I/O, and com port settings with the following:

- **Dip switches**—These are small switches on the circuit board that can be moved to a different position. The position of the switch designates which setting is to be used.
- **Jumpers**—Pairs of metal pins built into the circuit board, these represent an electrical contact point. Jumpers are configured by placing a small plug on the pins to complete the circuit. The instructions that come with your internal modem tell you how the jumpers should be set to use a specific IRQ, I/O address, or com port.
- **Software**—Some modems do not have physical switches or jumpers, but do come with a software program that is run to change the configuration.

Plug and Play

Many modern modems support *Plug and Play (PnP)* technology, which enables the operating system to detect the device, install the necessary software drivers, detect what resources are free on the computer, and assign those resources to the device automatically. Little or no intervention is required from the user.

PnP is great—when it works and when you are aware of a few caveats. If you buy a modem or other device that is advertised as Plug and Play, it is automatically configured *only* if the following is true:

- Your computer's BIOS (Basic Input/Output System) supports PnP.
- You are running a PnP operating system.

Both criteria must be met. Computer motherboards produced after 1995 usually support PnP. Operating systems that support PnP include Windows 95, 98, ME, and 2000.

NOTE


Windows NT is *not* a PnP operating system. However, it does have limited PnP functionality and detects some modem types.

External Modems

External modems have a couple advantages over the internal variety:

- Most external modems provide status lights, which indicate when the modem is powered on, connected, or transferring data. See Figure 6-2.
- External modems are generally easier to install and configure. There are no switches or jumpers to set, and you don't have to open the computer case.

Figure 6-2 External modems provide status indicator lights.



External modems require power cords to plug into an electrical outlet, but internal modems run off the computer's power. A serial cable connects the modem to one of the serial ports on the back of the computer.

Serial Port Considerations

To use an external modem, you need a free *serial port*. Most computers have two built-in serial ports, labeled COM 1 and COM 2, with connectors on the back of the computer.

Many devices, such as scanners, digital cameras, and serial pointing devices, also use serial ports. If your computer does not have a free serial port, you have a few options:

- Use an internal modem.
- Install an expansion card in your computer, which enables you to add serial port connections.
- If your computer's motherboard supports universal serial bus (USB), you can chain multiple serial devices, such as modems, off a single serial port. You might have to add a card to provide a USB connector, and you need a USB modem.

UART Chips

Serial ports use a chip called a UART (Universal Asynchronous Receiver/Transmitter) to handle serial communications. This chip comes in different types, and the type used determines how fast data can be transferred over that serial port.

The first PCs had 8250 UART chips. The top speed for this chip is 9600 bps, which means that even if you attach a high speed (56 kbps) modem to one of these ports, your speed would be limited by the UART.

Modern computers have UART chips in the 16450 or 16550 series. These serial ports can support transfer speeds of up to 115,200 bps.

16650 and 16750 UART chips are also available as add-on “enhanced serial port” cards. Internal modems have their own UART chips built into the card, so the speed of the computer’s com port is irrelevant.

NOTE If you have a high-speed modem and a modern computer, but are able to connect only at low speeds, check the com port configuration settings. Some operating systems set the com ports to 9600 bps by default; you need to change this setting to realize the port’s full capacity.

Modem Drivers

Drivers are software programs that act as a liaison between the hardware device and the operating system. Driver software is usually supplied by the modem manufacturer with the device, or it can be downloaded from the manufacturer’s Web site.

You must install the correct driver software for your device because if operating system code included support for all hardware devices that could possibly be used with it, the operating system would require significantly more disk space—much of it wasted on driver software that would never be used.

Modem Configuration

In addition to installing the driver software that enables the operating system to recognize the modem, and setting the IRQ, I/O address, and com port that the modem will use, you have to configure the modem to dial and maintain a connection. Modern operating systems have built-in support for dialup networking. You might have to install the remote access services if the modem was not present when the operating system was installed.

Modem Banks

A computer can be configured as a *dialup server* (also called a *remote access server*) to enable other computers to dial into it and connect to it over the phone lines. Computers running powerful server software can support many incoming remote access connections simultaneously; for instance, Windows NT Server supports up to 256 connections.

How can you connect 256 modems to a remote server? When you have many simultaneous dial-in connections (for example, when the server belongs to a company with many telecommuters who need to connect to the corporate network from home), you can use a *modem bank*. Modem banks are also called *modem nests* or *modem pools*.

A modem bank enables you to use a group of modems (usually mounted together in a rack) with a single server, and host multiple remote connections. The rack of modem cards is controlled by an interface that connects to the server, to a router, or directly to the local network. Of course, you need a phone line for each separate connection.

ISDN and DSL Adapters

The device used to connect a computer to an Integrated Services Digital Network (ISDN) or DSL telephone line is often referred to as a modem. It is more accurately called a *terminal adapter* because it does not modulate and demodulate signals because ISDN lines are digital, unlike the analog PSTN lines.

ISDN Adapters

ISDN adapters, such as modems, come in both internal and external varieties. They are configured similarly to modems, but the typical 128-kbps ISDN service consists of two data channels that each run at 64 kbps. The two channels are commonly used in a *multilink* configuration to provide the 128-kbps bandwidth. We discuss ISDN technology later in this chapter in the section, "ISDN."

The two data channels have separate telephone numbers in most cases. ISDN adapters are configured with information about the *service profile identifier (SPID)* for each channel, which consists of the telephone number, a two-digit sharing terminal identifier, and a two-digit terminal identifier (TID). Some modern ISDN adapters support automated SPID selection and do not require you to enter this information.

DSL Adapters

Both ends of a DSL connection require a device called an *endpoint* (and often referred to as a *DSL modem*), which connects to an Ethernet NIC installed in the computer. In some cases, the endpoint/modem is external. In others, the endpoint and the NIC are placed together on the same card.

Customer Premises Equipment

Customer premises equipment (CPE) is a general term that encompasses several different devices. The customer's site requires this hardware to process incoming transmissions from WAN links such as T-carrier lines, X.25 connections, and Frame Relay links.

Common types of CPE include the following:

- A channel service unit/digital service unit (CSU/DSU), used with circuit-switched connections such as a T-1 line. The CSU receives and transmits signals to and from the WAN line. The DSU manages line control, timing errors, and signal regeneration.
- A packet assembler/disassembler (PAD), used with packet-switched connections such as X.25. The PAD is an asynchronous device that enables multiple terminals to share a network line. Users dial into PADs through modems.

WAN Topologies


We discussed LAN topologies in Chapter 2, “Categorizing Networks,” and some of those same concepts apply to WANs. In the context of WANs, however, the *topology* describes the arrangement of the transmission facilities.

The simplest WAN topology is a simple point-to-point connection. The WAN, like the LAN, also can use traditional networking topologies such as a ring or star.

The Point-to-Point WAN

A point-to-point WAN is similar to the LAN topology referred to as a linear bus. A remote access link, which can be anything from a 56-kbps dial-up modem connection to a dedicated T-1 line, connects each point on the WAN to the next. See Figure 6-3 for an illustration of this.

Figure 6-3 A point-to-point WAN directly connects two endpoints.



This is a relatively inexpensive way to connect a small number of WAN sites. However, it is not fault tolerant. For example, in Figure 6-3, if the equipment at the Dallas office fails, San Francisco and Boston cannot communicate with one another. Limited scalability (the capability to “grow gracefully,” that is, to continue to function efficiently as the network grows larger) is another disadvantage. If you add another point to the WAN at Nashville, between Dallas and Boston, you increase the number of hops required for Boston to communicate with Dallas or San Francisco.


NOTE In wide-area networking, a *hop* is defined as the trip from one router to the next. The *hop count* is the number of routers the packet passes through from source to destination.

The point-to-point link works best for small WANs with only two or three locations.

The WAN Ring

A ring is constructed by establishing a point-to-point connection from Point A to Point B, from Point B to Point C, and from Point C back to Point A, as shown in Figure 6-4.

Figure 6-4 A WAN can use a ring topology.



The ring topology provides redundancy. In the example shown in Figure 6-4, if the line between Dallas and San Francisco goes down, data can still be transferred between the two cities by going through Boston.


The ring topology is more expensive to implement than the single point-to-point topology, and it suffers from the same scalability problem as the point-to-point topology.

The ring topology works well for WANs that connect only a few locations and that need the reliability offered by the redundant pathways.

The WAN Star

When a WAN is laid out in a star configuration, a device called a *concentrator router* is used; it serves as a central point to which all network routers are connected. In Figure 6-5, for example, the concentrator router is located at the Dallas headquarters.

Figure 6-5 A WAN arranged in a star topology is scalable.



The star topology is more scalable than the ring, and in a star, it is easier to add locations to the WAN.

The disadvantage of the star is its single point of failure. In the case shown in Figure 6-5, this point of failure is the Dallas concentrator router. If this device fails, communications cease among all points on the network.

Full- and Partial-Mesh WANs

A mesh topology, where there are multiple connections between points, provides the most fault-tolerant and reliable WAN. Unfortunately, it is also the most expensive to implement, and it becomes cumbersome if there are a large number of sites.


A full mesh topology requires that every site in the network be connected to every other site. With a partial mesh, a smaller number of redundant connections exist. This provides reliability approaching that of the full mesh, at significantly lower cost.

Refer to Figures 2-10 and 2-11 in Chapter 2 for an illustration of mesh and partial (hybrid) mesh topologies.

Multitiered WANs

A multitiered WAN is similar to the star in that it uses concentrator routers, but is more reliable because it links two or more of these concentrators with other locations “cascaded” off the main routers. See Figure 6-6 for an illustration.

Figure 6-6 A multitiered WAN offers more reliability than a simple star.



The multitiered WAN is scalable because additional locations (and even additional tiers) can be added to the network relatively easily. This topology is used for large, fast-growing networks.

Traffic flow can become a problem on large, multitiered WANs. Flow patterns should be carefully analyzed to ensure the most efficient placement of equipment for best performance.

Network Switching Types

Data can travel to a distant destination over several different types of lines, using one of two popular switching technologies:

- Circuit switching
- Packet switching

In the next sections, we examine how each technology works.

Circuit Switching Versus Packet Switching

Many people confuse circuit switching and packet switching. An example of a circuit-switched network is the telephone network. When you place a call to a branch office in Boston, a circuit is established for the duration of the call. The signal uses that circuit, or pathway, until you break the connection by hanging up. If you call again tomorrow, a different pathway might be taken.

An example of a packet-switched network is the Internet. When you send an e-mail message to the Boston office, it is broken down into chunks called *packets*. Each packet might take a different pathway to reach the destination computer; they are reassembled at the other end.

Circuit-Switching Networks

The first switching type we examine is the *circuit-switched* network.

Dialup Versus Dedicated Connections

Circuit-switched networks include both dialup and dedicated leased lines.

A dialup connection is a *temporary* connection, established for the duration of the session. Dialup connections can, however, be implemented as “always-on” connections with which you dial up the remote server and then do not end the connection. With a dialup connection, it is possible to hang up and dial a different location if you choose to do so. For example, you can end the connection to one ISP and then dial in and connect to another.

A dedicated connection is one that goes only from one specific point to another (for example, from your business office to your ISP).

Circuit-switching technology has a long history and is older than packet switching. Circuit switching is most appropriate when data must be transmitted in real time (as with a

telephone conversation). Circuit-switched networks are connection-oriented networks because a connection is established before transmission begins.

We briefly examine the following circuit-switched technologies in the following subsections of this chapter:

- PSTN
- ISDN
- DSL
- Leased lines
- Digital data service (DDS)
- T-carriers
- Switched 56

PSTN

The most common type of dialup WAN link is made by using the PSTN—the analog phone lines that are installed in most residences and businesses.

PSTN has two big advantages:

- It is available almost all over the world.
- It is inexpensive.

A dialup connection using ordinary phone lines is easy to implement. Besides a modem, no special equipment is required, and analog modems are readily available, simple to configure, and inexpensive.

The telephone system was not originally created with data transfer in mind; it was designed to transmit voice. High speed was not an issue, so there is an inherent limit to the attainable transfer rate.

Line quality is also a factor. Even with top-of-the-line 56-kbps modems, many telephone lines are capable of providing no more than 40 kbps–45 kbps.

ISDN

ISDN was designed to eventually replace POTS and provide a reliable digital connection suitable for both voice and data. Although that hasn't happened, and the recent advent of faster technologies at lower cost means it probably won't, ISDN still offers some advantages.

The characteristics of ISDN include the following:

- As its name implies, ISDN is a digital link. Because it does not have to convert data from digital to analog format and back, performance and reliability are high.
- It is more readily available than some of its newer competitors, such as DSL.
- Although it is a dialup technology, ISDN can be used as an always-on link (that is, dedicated ISDN).
- ISDN service is more expensive than analog service (PSTN) and requires specialized equipment, both at the telephone company central office (that is, at the digital switch) and at the customer's premise (that is, at the ISDN terminal adapter).

An ISDN circuit is made up of one or more channels that carry data (called bearer channels, or B channels) and a control channel (called the Delta channel, or D channel).

Each B channel provides 64 kbps of bandwidth, and B channels can be aggregated by using *inverse multiplexing*. This enables you to combine the bandwidth of multiple channels to create one high-speed connection. The D channel provides either 16 or 64 kbps, depending on the interface implementation.

ISDN is offered by most telephone companies in two standard access interfaces:

- **Basic Rate ISDN (BRI)**—This interface consists of two 64-kbps B channels (for an aggregate usable bandwidth of 128 kbps) and one 16-kbps D channel.
- **Primary Rate ISDN (PRI)**—This consists of 23 64-kbps B channels (for an aggregate bandwidth of 1.472 Mbps) and one 64-kbps D channel.

BRI is often implemented for residential or small business high-speed data transfer, and PRI is commonly used for digital voice transmission in conjunction with private branch exchange (PBX) telephone systems.

NOTE

A *PBX* is a private telephone network operated within an organization. Internal users share outside lines, and calling within the organization requires dialing only a four-digit extension. A traditional PBX required a switchboard operator, who answered all incoming calls and then routed each to the appropriate extension. Modern equipment automates this process.

DSL

DSL is a relatively new technology, offered by telephone companies as an add-on service over existing copper wires. DSL offers several advantages over other WAN link types.

The following list contains characteristics of DSL:

- DSL offers speeds up to and exceeding those of T-1, at a fraction of the cost. In many areas, DSL service costs less than ISDN.
- DSL is an always-on technology. There is no need to dial up each time you wish to connect.
- Both voice and data can be transmitted over the same line simultaneously.
- At present, availability is limited. The telephone company central office (CO) that is servicing the location must have DSL equipment installed, and for most “flavors” of DSL, you must be within a specified number of feet from the CO to get DSL service.

DSL comes in several varieties:

- **ADSL (Asymmetric DSL)**—This is the most common implementation. Speeds vary from 384 kbps to 6 Mbps (or more) downstream, typically combined with a lower upstream speed.
- **SDSL (Symmetric DSL)**—This provides the same speed for downloads and uploads.
- **HDSL (High Data Rate DSL)**—This variety typically provides bandwidth of 768 kbps in both directions.
- **VDSL (Very High Data Rate DSL)**—This is capable of bandwidths between 13 Mbps and 52 Mbps.
- **IDSL**—This is DSL over ISDN lines. It has a top speed of 144 kbps, but is available in areas that don’t qualify for other DSL implementations.

The generic term for DSL, encompassing all implementations, is *xDSL*.

ADSL, currently the most popular DSL implementation, generally provides a fast downstream transfer rate (typically 1.5 Mbps) and a slower upstream rate. This is based on the theory that most users of the Internet primarily access e-mail and surf the Web, which are download-intensive tasks. The lower upload rates do not work as well, however, if you wish to host a Web or FTP server, or engage in other upload-intensive tasks.

ADSL typically uses frequency-division multiplexing (FDM) to split the bandwidth and create multiple channels. Some ADSL implementations use a different method, called *echo cancellation*. It is more efficient, but also more complex and more costly.

Table 6-1 summarizes currently available DSL implementations.

Table 6-1 *Comparison of DSL Implementations*

DSL Type	Average Speeds	Advantages	Disadvantages
ADSL	384 kbps to 6 Mbps (downstream)	Relatively inexpensive; more widely implemented than other types.	Can be installed only within 17,500 ft of a telephone company CO; upstream speed is usually much slower.
SDSL	Up to 3 Mbps	Offers the same data rate upstream and downstream.	Generally more expensive and less widely available than ADSL.
IDSL	144 kbps	Can be installed in many locations where other DSL types are not available because of distance.	More expensive than ADSL; considerably slower speed.
HDSL	768 kbps up and downstream	Faster than IDSL and some implementations of ADSL.	Not widely available.
VDSL	13 Mbps to 52 Mbps	Extremely high speed for live audio and video.	Not widely available; most expensive DSL type.

Leased Lines

For WAN links that require guaranteed high performance and reliability, an option is to lease lines from the telephone company for private use. A leased line provides a permanent connection from one point to another (for example, from one branch office to another, or from your company LAN to your ISP).

DDS

DDS was one of the first digital services made available to the public. DDS provided a 56-kbps transfer rate. It lost the popularity contest to T-carrier technology because a T-1 line typically provides more bandwidth per dollar.

T-carriers

T-carriers are dedicated digital circuits that are typically leased by large companies to provide high-speed data, voice, audio, and video over a highly reliable point-to-point connection.

T-carrier circuits are typically established over copper wires, but they can also run over fiber-optic cable, coaxial cable, and even wireless technologies.

A CSU/DSU is used at each end of the connection to encode the data to be sent over the T-carrier.

Although prices for T-1 lines have fallen dramatically over the last decade, it is still an expensive option. A T-1 line typically costs 10 to 20 times that of DSL service for comparable speed (1.5 Mbps).

Why would anyone pay for T-1 when low-cost high-speed options are available? Availability itself is one reason; DSL has only recently become widespread, and even in areas where it is offered, many businesses and residences are not within the distance limitations required to order the service.

Another reason to pay extra for T-1 is guaranteed bandwidth. This is called the *committed information rate (CIR)*. With 1.5-Mbps DSL service, the telephone company sells you a service that has a maximum transfer rate of 1.5 Mbps. Your line might or might not actually perform at that speed at a given time. When you lease a T-1 connection, the telephone company guarantees the data rate of 1.544 Mbps. This can be an important consideration for corporate enterprises that depend on network performance.

The T in T-carrier refers to transmission channel; the signal itself is more accurately referred to as the data signal (DS) rate. You hear both DS-1 and T-1 used to refer to the same line type. Table 6-2 lists common T-carrier implementations.

Table 6-2 *Common T-Carrier Implementations*

Carrier Designation	Data Signal Rate	Data Transfer Speed
T-1	DS-1	1.544 Mbps
T-2	DS-2	6.312 Mbps
T-3	DS-3	44.736 Mbps
T-4	DS-4	274.760 Mbps

T-carrier lines consist of multiple 64 kbps channels. It is possible to lease just part of a T-1 line (in 64 kbps increments) if you don't need the entire 1.544 Mbps bandwidth; this is called *fractional T-1*.

Switched 56

Switched 56 is an enhanced version of PSTN. It is a digital switched-circuit connection that transfers data over one 56 kbps channel. Switched 56 is less expensive—but also much slower—than a T-1 line. It is a dialup technology, thus it can be appropriate in cases in which a dedicated connection is unnecessary.

Unlike PSTN, which is theoretically capable of data rates of 56 kbps but in practice generally attains speeds no higher than 50 kbps (and often falls far short of that), Switched 56 can provide a reliable full 56 kbps connection. Because it is a digital connection, error rates are lower than on a regular PSTN (analog) line.

The connection is called “switched” because individual 56 kbps channels are switched out of a T-1 circuit and sent to a specific user location.

The popularity of Switched 56 has suffered as other low-cost, high-bandwidth options (such as ISDN and DSL) have become more widely available.

Packet-Switching Networks

Packet-switching networks are networks in which data packets can take different routes to reach the same destination. At the receiving end, the packets are put back together in the correct order. Packet-switched networks are often depicted as a cloud because the exact route of travel of the data is unknown.

Packet-switching technologies include the following:

- X.25
- Frame Relay
- Asynchronous Transfer Mode (ATM)

NOTE Although you often hear the term *X.25 network*, the technically correct term is *Public Switched Data Network (PSDN)*. X.25 is the protocol that is used for communication between the data terminal equipment and the network.

We discuss each in more detail in the following sections.

X.25

X.25 was one of the first packet-switching networks and was designed to work with IBM mainframes, such as the IBM 360, and use analog transmission.

X.25 was originally called the ARPAnet 1822 protocol; the name X.25 came from the specifications for the protocols used by this technology, established by the International Telegraph and Telephone Consultative Committee (CCITT) in 1976.

NOTE The CCITT changed its name to the International Telecommunications Union (ITU) in 1993.

PSDN technology operates at the first three layers of the OSI model. The technology called X.25 is actually made up of several protocols:

- PSDN uses a protocol called X.21 at the physical layer (a variation of the X.21 physical layer protocol, X.21bis, is used in the United States).
- A protocol named Link Access Procedure Balanced (LAPB) is used at the data link layer.
- At the network layer, the Packet Layer Protocol (PLP) is used to assemble frames from the data link layer into packets.

The primary objective in designing the X.25 protocol was reliability. At the time it was designed in the 1970s, both the computers in use and the telephone lines were prone to error. Thus, the PSDN running on X.25 included redundant error-checking to compensate for these problems. The result was a highly reliable means of data transfer, but performance was slowed by the extra error-checking activity. The PSDN usually transfers at 64 kbps or below.

There are still public switched data networks in use today. To make a WAN connection over a PSDN, you can do one of the following:

- Dial into a packet assembler/disassembler (PAD) with an asynchronous modem
- Make a synchronous connection using the X.32 protocol
- Use an X.25 smart card to connect directly to the PSDN

Frame Relay

Frame Relay is a newer packet switching technology, which was designed to be used over digital lines and which grew out of X.25. Frame relay is a variation on and improvement to the X.25 technology, developed by the CCITT. Frame Relay uses only the first two OSI layers rather than the first three (as X.25 does). It was developed to take advantage of modern computers and telephone lines, which are far more reliable than those in use when X.25 originated. It has become a popular option for WAN links, and it generally offers a higher-performance, more cost-effective solution than does X.25.

Frame Relay operates only at the two lowest levels of the OSI model, the physical and data link layers. Frame Relay uses less overhead than X.25, and thus, it is faster. Frame Relay can run at T-1 and T-3 speeds (from 1.5 Mbps to almost 45 Mbps).

Frame relay is called a *fast packet* technology.

A typical Frame Relay implementation uses a *permanent virtual circuit (PVC)* to provide an always-on connection. Because service providers generally charge fees based on usage (referred to as *bandwidth on demand*), you can avoid the cost of a dedicated leased line.

Frame Relay has high performance because it does not include the extensive error checking and correction of X.25. Frames with errors are discarded, and it is up to the endpoints

(communication computers) to detect the missing packet and request retransmission. Because transmission is digital, there are relatively few errors to contend with, and Frame Relay works well in a WAN environment over T-1 lines.

ATM

ATM is a popular packet-switching technology that was designed to support high-speed applications such as streaming audio and video. An important concept for ATM networks is quality of service (QoS), which is a way to control the allocation of network bandwidth to specific applications to provide guaranteed bandwidth where it is most important.

ATM is hardware based, which means that all equipment on the network must be designed to work with ATM. The advantage is that this results in high speeds for processing and switching. Standard ATM transfer rates are 25 Mbps, 155.520 Mbps, and 622.080 Mbps, and ATM is capable of speeds of 10 Gbps. Unfortunately, that performance comes at a high price. ATM is expensive to implement because all network hardware must support ATM, and network interface cards (NICs), hubs, and other ATM-compatible equipment is costly.

ATM is a modern digital technology that breaks data into 53-byte fixed-length units called *cells*. Five bytes are used for the ATM header, which contains addressing information.

ATM can be used for both LANs and WANs, and it uses multiplexing to transfer voice, data, and video simultaneously over the network. Cell switching is a function of the ATM hardware (unlike Frame Relay, in which it is a software function).

You can connect to an ATM network through a direct connection or an on-demand connection. The connection between the two endpoints is a *virtual circuit*; it can be either a PVC or a switched virtual circuit (SVC). With either a PVC or an SVC, ATM uses predefined circuits instead of establishing the virtual circuits at the time of connection, as X.25 and Frame Relay do. This saves a great deal of time and is another factor in ATM's high speed.

Many networking experts predict that in the future, ATM will become the technology of choice for both LANs and WANs.

Emerging WAN Technologies

New, faster, and more efficient WAN technologies are being developed all the time. Many of these interoperate with one another to provide support for the high-bandwidth applications in use today and those expected to be in demand in the future.

In the following sections, we look at new high-speed technologies, including OC-SONET, Broadband ISDN, CATV, and SMDS.

OC-SONET

OC stands for *optical carrier*, and SONET stands for *Synchronous Optical Network*. SONET is a physical layer protocol that provides for high-speed transmission using fiber-optic media. SONET is capable of rates of almost 20 Gbps, and ATM can run over SONET to achieve very high data transfer speeds.

NOTE You might see the term Synchronous Digital Hierarchy (SDH) used to refer to the SONET technology outside the United States.

The SONET signal rate is measured by OC standards. Table 6-3 illustrates the available transmission rates (called optical carrier levels).

Table 6-3 *OC Signal Transmission Rates*

OC Level	Signal Transmission Rate
OC-1 (base rate)	51.84 Mbps
OC-3	155.52 Mbps
OC-12	622.08 Mbps
OC-48	2.488 Gbps

SONET is used as the physical basis for another technology, broadband ISDN, which is discussed in the next section.

Broadband ISDN

Broadband ISDN (BISDN) is an emerging technology designed to use fiber-optic cable and radio waves to transmit data at high speeds over SONET, FDDI (the Fiber Distributed Data Interface), and Frame Relay.

Broadband technologies, which can send multiple channels of data, video, and voice over the same medium, are growing in popularity as Internet connectivity and other high-bandwidth network usage increases. Other broadband technologies include DSL and cable modem.

CATV

Cable TV (CATV) companies saw a great opportunity: They already had a vast infrastructure of coaxial cable in most major cities and many rural areas, and this cable

could be used not only to transmit television signals, but also to transmit computer data. Numerous cable providers now offer Internet access accounts.

Cable is not a general WAN technology. It was originally designed to enable you to communicate only with the cable company/service provider's server (in the form of receiving incoming television channels). Although customers were all connected to the same network through the coax cable running through their neighborhood, the network was not designed to enable them to communicate with one another. In fact, the network was not designed to enable its users to send data at all—only to receive it. Cable modem changes all that.

Cable Internet access requires a cable modem that connects both to the incoming coax cable and to a NIC in the user's PC (typically this is a 10BaseT Ethernet NIC).

In this scenario, the cable company is also the user's ISP. There is no option to separate the provision of the physical line from the access service as there is with access over telephone lines. In other words, you cannot lease the line from the cable company and use it to connect to some other ISP's server.

On the other hand, when you pay a telco for the use of a phone line (whether an analog PSTN line or a dedicated T-1 line), you can purchase an Internet account from any ISP you choose (including the phone company that provides the line). Although this same method is technically possible with cable, the cable companies have packaged the two services together and the terms of their service contracts require that you use the cable company as your ISP.

Cable infrastructure can support either one-way or two-way transmissions. *One-way cable* provides only downstream transmission over the coax. Uploading must be done over a regular analog phone line that also plugs into the cable modem. With one-way cable, upload speeds are limited to standard rates attainable over PSTN, which is less than 56 kbps. Download speeds vary from 364 kbps to 1.5 Mbps.

Two-way cable provides both uploads and downloads over the coax. Nonetheless, many cable companies limit the upstream speed to 128 kbps to discourage customers from running servers (which is often prohibited by the CATV terms of service contract).

Cable is an always-on technology, but one-way cable still requires you to dial up to establish a connection. A big advantage of CATV is its low cost; however, in many areas, users experience reliability problems. Because cable is a "shared-bandwidth" technology (that is, the entire bandwidth of the cable is divided between all users on that cable segment at a given time), performance might degrade as more users in the neighborhood are added to the network. There are also security issues that, at this time, make CATV more viable for residential use than for business.

SMDS

Switched Multimegabit Data Service (SMDS) is a new packet-switching technology that is designed especially for WAN links that experience a lot of “bursty” traffic. (*Bursty* refers to transmission that comes in “bursts” rather than in a constant, even stream.)

SMDS is connectionless; that is, there is no requirement that a connection or circuit be established before transmitting the data. It uses relatively large packets, up to 7168 bytes in length. SMDS addresses, which are ten-digit numbers (such as a telephone number), are used to identify the SMDS subnetwork. SMDS links are connected to an SMDS switch on the telephone company’s backbone network, typically by multiple OC-3 SONET links.

SMDS was designed as a public network to provide services similar to those of a LAN, except that it spans a metropolitan area. Data transfer speeds typically range from 1.544 Mbps to 45 Mbps. It is scalable and can be used in conjunction with ATM. However, SMDS is not as widely available as Frame Relay and other services, and SMDS equipment may be more difficult to find.

Wireless WAN

In many cases, it is impossible—or at least inconvenient or expensive—to run a wired link to connect WAN sites. Wireless solutions are especially appropriate when it is important that data be communicated in real time, or when users are on the move. Wireless works best for communicating small amounts of data.

The wireless technologies used for WANs include the following:

- **Radio frequency (RF) technologies**—Specialized Mobile Radio (SMR) provides data rates of 1200 bps to 19,200 bps. Enhanced SMR (ESMR) is the digital implementation.
- **Satellite technologies**—This provides both circuit-switched and packet-switched services at speeds of 4800 to 9600 bps.
- **Microwave technologies**—This technology uses cellular techniques over microwave frequencies to provide higher speed and capacity (wireless broadband).
- **Cellular technologies**—This provides a circuit-switched connection over analog or digital cellular links.
- **Packet data network technologies**—This technology provides a packet-switched WAN with no call setup involved.

Compared to wired links, wireless communications are often more costly and relatively slow. For example, analog cellular systems typically provide no more than 14,400 bps transfer rates, while digital cellular offers up to 64 kbps.

LAN/WAN Connectivity

In today's wired world, local connectivity often is not enough. No LAN is an island, or at least, fewer and fewer of them are islands as it becomes vital to business interests that a LAN be able to communicate with the outside. This can mean connecting the LAN to a corporate WAN, the global Internet, or both.

There are several ways to connect your LAN to the outside world, depending upon your budget and needs. Of course, the most obvious way to provide network users with access to other networks is to equip each PC with a modem and phone line. In this manner, each user can establish a dialup connection to an ISP or other remote server. However, this solution has many drawbacks:

- It becomes prohibitively expensive as the number of users increases. Not only must you purchase hardware (the modem) for every computer, you must also pay for a separate telephone line *and* if users are to connect to the Internet, a separate ISP account for each.
- Allowing users to dial out using a modem can create serious security risks if the nature of the data on your network is confidential. The company has little control over which networks the user connects to and the audience to which the company's data might be exposed.
- A high degree of user sophistication is required, which means a significant expense for training users to configure and manage their own dialup connections.

There is a better way. In fact, several alternatives offer advantages over the old-fashioned way of connecting LAN users to a WAN. Each has advantages and disadvantages, and which is best depends on your particular situation.

In the following sections, we briefly discuss the following LAN/WAN connectivity options:

- Translated connections
- Proxy servers
- Routed connections

Translated Connections

One of the most cost-effective ways to connect all computers on a small LAN to the Internet or to another WAN link is through address translation. *Address translation* enables all computers to access the WAN through a single host computer, using only one telephone line and ISP account (or other WAN link) and only one registered public IP address.

How Address Translation Works

A computer running address translation software sits between the public WAN and the private LAN. It has interfaces to both networks. This computer has a private IP address used for communications with other computers on the LAN and a public IP address (which can be assigned through Dynamic Host Configuration Protocol [DHCP] from an ISP's server at the time the WAN connection is established). We refer to this computer as the *address translation host*.

NOTE DHCP is a service that automatically assigns IP addresses and other TCP/IP settings to computers that are configured to use DHCP. You will learn more about DHCP in Chapter 8, "Networking Protocols and Services."


Address translation works by mapping the private IP address of each computer on the LAN that sends data "outside" to a port number on the host computer. This information is added to the IP header of the packet, which is then sent out over the WAN with the IP address of the host computer (the one that has the physical connection to the WAN) as the source address.

When a computer on the local network opens a Web browser and sends a request to view a URL, for example, the host computer assigns a port number to that request, which identifies the original sending computer. Then the host sends the request out to the ISP's Web server. When the page is returned to the host computer (whose IP address is listed in the header as the source of the request), the host consults its address translation table, matches up the packets with the computer that originally sent the request, and then forwards the Web page to that computer.

The information in the address translation table includes the following:

- The original source and destination IP addresses (identifying the sending computer within the network and the computer outside the network to which the data is sent)
- The original source and destination port numbers (identifying the application making or receiving the request; for example, HTTP requests for Web pages are normally sent to TCP port 80)
- Sequence numbers (identifying the order in which the packets are sent)
- A timestamp

Network address translation (NAT) is the common term for which standards have been developed and published as RFC 1631. Not all address translation technologies comply with these standards. Figure 6-7 illustrates the steps in the address translation process.

Figure 6-7 *The NAT process involves translating private addresses to a public address.*

- 1 The user at the client computer (IP address 192.168.1.9) opens a Web browser application and enters the URL `www.tacteam.net` into the address box. The browser software sends an HTTP request to the IP address associated with the `www.tacteam.net` “friendly name.”
- 2 The NAT host on the client’s LAN maps the request from 192.168.1.9 for `www.tacteam.net` to a port number in the mapping table. The table contains the original source and destination IP addresses and original source and destination TCP/UDP port numbers.
- 3 The NAT host changes the header so that to the outside network, the packet appears to originate not from 192.168.1.9, but from the public IP address assigned to the NAT host’s external network adapter by the ISP.
- 4 The NAT host sends the request for `www.tacteam.net` to the ISP server. Domain Name System (DNS) maps the name to the IP address of the server on which the `www.tacteam.net` homepage is stored.
- 5 The request is received by the `www.tacteam.net` server and the page is returned to the public IP address used by the NAT host.
- 6 The NAT host consults its address translation table to determine whether the page should be sent to the client at 192.168.1.9 (and the TCP/UDP port number to which it should be sent).

NAT Software

Some operating systems, such as Windows 98/2000 and current versions of Linux, have built-in support for address translation. In Linux it is referred to as *IP masquerading*. Windows 98 and 2000 Professional call the feature *Internet Connection Sharing (ICS)*, although it can be used to share a connection to a private remote network or to a VPN as well. Windows 2000 Server supports the more flexible and robust version of address translation, NAT. If an operating system does not include address translation, you can use an add-on NAT program to provide the same functionality. Examples of these programs include:

- **Sygate, from Sybergen Software**—www.sygate.com
- **NAT32, from A.C.T. Software**—www.nat32.com

Another, more sophisticated (and slightly more difficult to configure) type of software that you can use to share a connection is a *proxy*. We discuss this option in the next section.

Proxy Servers

A proxy server does more than provide a shared connection to the WAN, although it does serve this purpose. A proxy server acts as an intermediary, separating the LAN from the outside network, and it can provide protection by filtering incoming and outgoing packets. It also enhances Web performance by *caching* often-requested Web pages.

Proxies use an address translation method, but do not necessarily comply with the NAT specifications in RFC 1631.

How Proxies Work

The proxy server receives requests for Internet resources from proxy clients, similar to the way in which NAT works. The proxy server checks its filter settings (which are configured by the administrator). If the request meets filter requirements, the server looks first in its cache of stored pages. If the requested page is there, the proxy server returns the page to the requesting client. There is no need to send the request on to the ISP server. If the page is not there, the proxy server requests the page from the ISP server, receives it, and returns it to the client.

As with NAT, the internal clients that access the Internet through a proxy server are invisible to the Internet; all outside communication is done by the proxy server.

Proxy Software

Proxy software typically provides more protection and performance enhancement than NAT. However, it might also be more expensive, and it is typically more difficult to set up

because the Internet applications on all the client machines (such as the Web browser) must be individually configured to use a proxy server.

To use NAT, you need only set the client's TCP/IP configuration to obtain an IP address through DHCP, and other necessary information can be automatically distributed to the clients by the DHCP server.

Numerous proxy server applications are available for popular operating systems, and they include the following:

- **Rideway, from DGL**—dgl.com/rideway
- **Winproxy, from Osis**—www.winproxy.com/toc
- **Microsoft Proxy Server, from Microsoft**—www.microsoft.com
- **Squid for UNIX, from SCO**—www.sco.com

Some proxy programs, such as Microsoft Proxy Server, run only on a network operating system (NOS) such as Windows NT or Windows 2000 Server. Others, such as Rideway and Winproxy, can be used on desktop operating systems such as Windows 95/98, Windows NT Workstation, and Windows 2000 Professional.

Software that combines proxy and NAT technologies are sometimes referred to as *transparent proxies*.

Routed Connections

A *routed connection* is another way to provide the computers on the LAN with access to a WAN. A routed connection enables each computer to participate directly on the Internet, unlike NAT and proxy connections, where the individual computers must go through an intermediary.

Configuring a routed connection requires extensive knowledge of TCP/IP addressing, and you must purchase and configure a *router*. Additionally, every computer on the LAN that connects to the outside network must have a “legal” registered public IP address.

NOTE See Chapter 8 for more information on TCP/IP and IP addressing and Chapter 9, “The Widest Area Network: The Global Internet,” for more information on routing.

Why Use a Routed Connection?

A big advantage of the NAT and proxy solutions is the capability to connect the small LAN to a public WAN by using only a single IP address. However, this might not be the best solution in some situations.

Because of the way address translation works, protocols that do not store the addressing information in the IP header do not work with NAT. In some cases, *NAT editors* can be added to make modifications to the IP packet so that NAT will work. In other cases (for example, when packets are authenticated and encrypted using IP Security [IPSec]), address translation is not possible.

Configuring a Routed Connection

A routed connection requires either a dedicated routing device such as a router or the use of a computer running an operating system that enables IP forwarding (in the latter case, the computer acts as the router).

Computers that use TCP/IP to communicate must have the following properties configured:

- An IP address that is valid for the network on which they will communicate
- A subnet mask that designates what part of the IP address identifies the computer and what part identifies the network

Computers participating on a *routed* network must also have a *default gateway* configured. This is the address of the router, which has two network connections: one to the LAN and one to the outside network.

NOTE

The address for a DNS server must be entered if you want to use “friendly” host names (for example, URLs such as www.tacteam.net) instead of IP addresses. You learn more about DNS in Chapter 8.

To set up a routed connection to the Internet, the TCP/IP protocol on the router is configured with an IP address, a subnet mask, and a DNS server address obtained from the ISP, and a static default route is configured to use the Internet interface.

The computers on the LAN that connect to the Internet are likewise configured with IP addresses, a subnet mask, and a DNS server address obtained from the ISP. They are also configured to use the IP address of the router on its LAN interface as their default gateway address.

Summary

In this chapter, you have learned about wide-area networking and the established and emerging technologies that can be used to connect computers in distant locations.

We discussed WAN hardware, and then we moved on to WAN topologies and discussed the advantages and disadvantages of WAN configurations such as the point-to-point connection, the ring, the star, the full or partial mesh, and the multitiered WAN.

We then delved into the differences between circuit-switched and packet-switched networks, and we discussed the characteristics and technologies associated with PSTN, ISDN, xDSL, DDS, T-carriers, X.25, Frame Relay, ATM, OC-SONET, BISDN, CATV, SMDS, and wireless WAN technologies.

We next discussed how to connect a LAN to a WAN. Specifically, we learned several ways to provide all computers on a local network with access to the Internet or another outside network. We learned about translated connections that use NAT, which is sometimes called connection sharing or IP masquerading. Then we discussed how to set up a routed connection for those circumstances in which address translation is undesirable or impossible.

This chapter wraps up Part I of this book. You should now have a grasp of basic networking concepts. In Part II, we look at the hardware and software that makes the network run. Chapter 7 introduces you to the components of the physical network.

Further Reading

An excellent resource for information on various WAN technologies is WANsites, at www.networkcomputing.com/wansites/default.html.

Thorough, clear explanations of the different WAN technologies can be found at the High Performance Networking Unleashed Web site, at www.officewizard.com/books/network.

An excellent, detailed discussion of NAT is available online at www.suse.de/~mha/linux-ip-nat/diplom.

Review Questions

The following questions test your knowledge of the material covered in this chapter. Be sure to read each question carefully and select the *best* correct answer or answers.

- 1 Which of the following hardware settings commonly must be configured on an internal modem?(Select all that apply.)
 - a I/O address
 - b IP address
 - c IRQ
 - d Virtual com port

- 2 The CSU/DSU and the PAD are examples of which of the following?
- a ISDN terminal adapters
 - b Customer premises equipment
 - c Terminal identifiers
 - d Concentrator routers
- 3 PSTN, ISDN, DSL, DDS, and T-carrier links are all examples of what type of network?
- a Circuit-switched networks
 - b Packet-switched networks
 - c Switched 56 networks
 - d LANs
- 4 Which of the following are true of ISDN? (Select all that apply.)
- a It is an analog link.
 - b It is generally more expensive than PSTN.
 - c It requires special equipment at the CO and the customer's premise.
 - d It uses a circuit consisting of only one channel.
 - e It can transfer both voice and data.
- 5 Which WAN topology is the most scalable?
- a Point-to-point.
 - b Ring.
 - c Star.
 - d All of the above are equally scalable.
- 6 What was the first packet switching technology that was based on the ARPAnet 1822 protocol?
- a Frame Relay
 - b X.25
 - c ATM
 - d DSL

- 7 Which of the following are characteristics of Frame Relay that distinguish it from X.25? (Select all that apply.)
- a Frame Relay offers high performance.
 - b Frame Relay uses packet switching.
 - c Frame Relay uses digital signaling.
 - d Frame Relay does not include extensive error checking.
- 8 Which of the following describes ATM? (Select all that apply.)
- a It uses variable length packets.
 - b It uses 53 byte units of data called cells.
 - c It can transfer video, voice, and data simultaneously.
 - d It uses predefined circuits.
 - e It is less expensive than other WAN technologies.
- 9 Which of the following technologies is capable of transmission rates of up to 2.488 Gbps?
- a ADSL
 - b T-1
 - c ISDN BRI
 - d OC-SONET
- 10 Which of the following is true of NAT? (Select all that apply.)
- a NAT requires that each computer on the internal network have a registered public IP address.
 - b NAT is compatible with all applications.
 - c NAT is incompatible with technologies that encrypt IP data.
 - d NAT uses a table to map private internal IP addresses to one or more external public addresses.



Protecting the Network

In Part I, “Introduction to Networking Concepts,” you learned the technical concepts involved in sharing data across a computer network. In Part II, “Networking Hardware and Software,” we discussed the details of how the hardware components, the server and client operating systems, and the network protocols work together to make that data sharing possible.

We haven’t yet examined, however, an important aspect of sharing computers’ resources across a network: maintaining the integrity of the shared data. There are two broad issues involved in protecting our data from loss or misuse:

- Network security
- Disaster protection and recovery

The more that businesses become dependent on their computer networks and the information that they contain, the more vulnerable the businesses become when something goes wrong. Murphy’s infamous law assures us that sooner or later, something *will* go wrong. The “something” can involve deliberate or accidental breaches of network security, up to and including corporate espionage. On the other hand, it can take the form of a hard disk crash, flood, or fire that destroys the server.

In this chapter, we examine data protection issues and discuss ways to reduce the risk of data loss without giving up the many benefits of having computers networked and connected to the outside world. You will learn that in some cases, networking even works in your favor by making data protection and recovery easier than it would be on a standalone system.

Network Security

Secure networking is a hot topic in the information technology world. Well-publicized intrusions into governmental and big business networks, widespread attacks of computer viruses, and high-profile criminal cases involving computer hackers are constantly in the news. From the administrators of multinational enterprise networks to home computer users with dialup Internet accounts, almost everyone who is “connected” is also concerned, to some degree, about the possibility of unauthorized access.

Security means different things to different people. Although the word (according to the *American Heritage Dictionary*) is synonymous with “guarantee” and “warranty,” in the context of networking, security is never absolute. The only completely secure system is the one to which *no one* has access. This is obviously unworkable. *Computer security* is defined by the *Microsoft Press Computer and Internet Dictionary* as “the steps taken to protect a computer and the information it contains.” No guarantees are implied in this definition.

Because the entire purpose of computer networks is to share resources, the trade-off between security and accessibility is always a delicate balancing act. The more secure your network, the less accessible it is, and the more accessible it is, the less secure it is.

Security issues can make the network administrator’s relationship with network users an adversarial one. Users generally prefer more accessibility, and administrators like to err on the side of more security.

In the following sections, we discuss how you can assess the security needs of a particular network, how to assess existing and potential threats to security, and how to implement the appropriate security measures. We also take a look at how security components work and at some advanced identification and authentication technologies.

Assessing Security Needs

When it comes to a computer network, how much security is enough? The answer depends on the organization. The first step in developing a viable plan for protecting your network’s data is assessing its security needs. Factors to be considered include the following:

- The type of business in which the company engages
- The type of data stored on the network
- The management philosophy of the organization

Let’s take a look at these individually and discuss why each is important.

Type of Business

Some businesses, such as law or medicine, by their very nature generate confidential data. The privacy of a patient’s medical records and attorney-client communications are protected by law. If sensitive documents are stored on your network, it is imperative that a high level of security be maintained. To do otherwise puts the organization at risk of civil liability and even criminal charges.

Other organization types that often produce sensitive data include the following:

- Law enforcement agencies, courts, and other governmental bodies
- Educational institutions that store student records on a network
- Hospitals, mental health facilities, and substance abuse facilities

- Companies that contract with the Department of Defense (DoD) or that perform other national security-related work
- Any organization that gathers data under a guarantee of confidentiality
- Any organization that produces a product or provides a service in a highly competitive industry or field of research
- Any organization whose network is connected to the Internet

Type of Data

Regardless of the type of business, certain types of data are considered to be private and should be protected. These types include the following:

- Payroll records and employees' personal information
- Accounting and tax information
- Trade secrets such as original code, plans and diagrams, recipes, and business strategies

If these types of information are stored on your network, you should implement a security plan to protect them.

Management Philosophy

If the data on the network is not subject to privacy laws, the security level might be dependent on the business owners' or managers' personal philosophies about how open (or closed) they want the network to be.

In some organizations, everyone is considered to be part of one big, happy family. Accessibility and ease of use enjoy a higher priority than privacy and security. Other organizations operate on a "need to know" principle; management prefers that information be accessible only to those whose jobs require it. Neither policy is right or wrong; the network administrators simply need to know and must be willing to implement network security in keeping with the organization's management style.

NOTE

Distributed Denial of Service (DDoS) attacks work by compromising intermediary computers and surreptitiously installing software on them that will be used as part of the attack platform. Thus, your company's computers can unknowingly participate in the attack and can incur liability as a result. For this reason, it is important for all networks that are connected to the Internet to implement security policies, regardless of the sensitivity of their own data.

Assessing Security Threats

After you have decided that your business type, data type, and management philosophy require the implementation of security measures, you should assess the likely sources of threat to your data's integrity.

It is easy for organizations to underestimate, overestimate, or completely overlook the risks that make their networks vulnerable. However, there are many different types of threats to network security, and they all can be classified in one of two broad categories:

- External threats
- Internal threats

The following sections describe these types of threats.

External Threats

Once upon a time, external threats were not a serious concern for most company LANs. The network was generally self-contained, and for an intruder to penetrate from the outside, he or she would have to dial in to a modem somewhere on the network or tap into the cabling. Now that most LANs are connected to the Internet, all that has changed. When your network can access the outside world, those outsiders also can access your network.

The motives of external intruders vary. Common motives include revenge (such as dissatisfied customers, disgruntled former employees, and angry competitors), recreation (those who hack into networks “just for fun” or to prove their technical skills), and remuneration. In this last case, the intruder is being paid to invade the network or does so for personal gain; for example, the hacker might attempt to transfer funds to his own bank account or erase records of his debts.

External security breaches can take many forms, including the following:

- Unauthorized use of passwords and keys
- Denial of Service (DoS) attacks
- IP spoofing
- Computer viruses and worms
- Trojan horse programs

We discuss each infiltration method separately in the following sections.

Unauthorized Use of Passwords and Keys

A *password* is a sequence of alphabetic, numeric, or symbol characters used to verify that a user is really the person authorized to use a particular account to access a system or

network. A *key* is a number or cipher used by the system to verify the integrity of a communication.

Passwords and keys are security measures designed to keep unauthorized persons out, but they are effective only if they are kept secret. If someone knows your user account name and password and has a physical connection to the computer or network, that person can access anything to which you have access permissions. Obtaining a password is often the first step in hacking into a system.

The term *hacker* has different connotations, depending on who is using it. In the early days of computing, it meant a good programmer, but our modern news media has popularized a more negative meaning. Today it usually implies someone who breaks into computer systems, often with the intent to steal or destroy data.

Hackers are often represented as evil geniuses who use sophisticated, esoteric techniques to gain access to computers and networks. The truth is a little less exciting—in a high percentages of incidents, the intrusion is not based on high-level programming skills at all. Hackers often access the network in the same way as authorized users do—by entering a valid user account name and password. In this instance, they are often referred to as *crackers* because they “crack” or obtain the password and enter the network much as a safecracker opens a vault.

NOTE

Gaining unauthorized access by using someone else’s credentials is a form of *impersonation*. This type of security breach includes IP spoofing and other methods of representing yourself as someone (or something, such as a system device) that you are not.

Hackers obtain the required passwords in many ways. Frequently, no technical expertise is used—the hacker simply uses observation skills or powers of persuasion to obtain the credentials from the user to whom the credentials belong.

Many users create passwords that are easy to guess (such as their birth dates or their spouses’ first names). Others create—or are assigned by administrators—passwords that are less intuitive. Because this makes the passwords more difficult to remember, users write them down, often leaving them in an easily accessed desk drawer or even on a sticky note attached to the computer monitor.

Hackers can also obtain passwords by simply asking for them by posing as network technicians, company administrators, or others to whom the naive user feels safe in confiding his or her password. This is sometimes referred to as a *social engineering attack*.

The *brute-force attack* is another way to obtain passwords. This means trying every possible password until you find one that works. Hackers write scripts and programs, called password crackers, that run through lists of common words and alphanumeric combinations

automatically so that the intruder doesn't have to sit there and manually enter password after password.

If these methods fail, technically savvy hackers can obtain passwords by intercepting the data packets that contain them by using a network "packet sniffer." This technique is why it is important that passwords be encrypted, rather than sent across the network as clear text. We talk about sniffer software in detail in Chapter 18, "Monitoring, Management, and Troubleshooting Tools."

Password security is an important part of establishing good overall security policies for your network. We discuss it in more detail later in this chapter in the "Password Security" section.

DoS Attacks

DoS attacks can be launched in several ways. Regardless of the method, they are designed to interrupt normal operations of the machine that is the focus of the attack. These types of attacks are also sometimes called *nuke attacks*. In February 2000, several major Web sites (including Amazon.com and Yahoo!) were temporarily shut down in this way.

Some DoS attacks exploit bugs in particular computer operating systems or applications, while others are aimed at the network itself. There are often patches (software provided by operating system or application program vendors to repair a problem with the program) available from the software vendor to plug the "holes" that enable attacks.

DoS attacks don't cause the computer to crash but are designed to interrupt or prevent connection to the network. They work by deluging the network with useless packets or by emulating a network problem that causes the computer to disconnect.

Some common forms of DoS attacks include the following:

- Ping/Internet Control Message Protocol (ICMP) flood
- Smurf attack
- Ping of Death
- SYN attacks

The following sections take a brief look at each.

Ping/ICMP Flood An ICMP flood is exactly what it sounds like—a "flood" of ICMP packets that overwhelms the system. ICMP is a message and error-checking protocol used to transmit information over the Internet. The **ping** command is commonly used to send ICMP packets for the purpose of verifying that a specific computer (which is identified by IP address or host name, the latter of which is ultimately resolved to an IP address) is on the network. Ping works by sending a message called an ICMP Echo Request and then waiting for an ICMP Echo Reply from the "pinged" computer.

When a flood of packets is sent continuously to an IP address, it can become too much for the server to handle and can cause the server to slow down and eventually disconnect because of ping timeout.

Smurf Attack A smurf attack is an ICMP flood that affects an entire service provider or an entire network segment. The ICMP messages are sent to a broadcast address, which causes all computers on that subnet to respond. When an ISP is smurfed, all connections are slowed, and all users are eventually disconnected.


Smurf attacks are a type of DoS attack. After an attacker gains access to a network, the attacker sends a broadcast into that network using an address in the target network as the source. Then all the devices in the compromised network send ICMP replies to the target address. You could have hundreds of hosts, each sending thousands of bits of ICMP echo requests into the target network. The traffic generated by this process can easily overwhelm the low-bandwidth connections used across WAN links between some ISPs and networks. The target network is affected, and often, the compromised middle network also suffers from all this traffic.

Ping of Death The Ping of Death is a slightly more sophisticated attack that takes advantage of the maximum transmission unit (MTU) limitations of a network. The MTU depends on the media and network architecture. If a packet is sent that exceeds the MTU, it must be broken into smaller chunks and then reassembled at the destination end.

The IP packet in which the ICMP Echo Request is encapsulated is limited to 65,535 octets (an *octet* is eight bits of data). A knowledgeable attacker can send a packet that exceeds the number of octets that are allowed in the data field of the Echo Request. When the destination computer tries to reassemble this packet, it crashes.

SYN Attacks An attacker can use the TCP synchronization sequence to disrupt communications. A process called the TCP three-way handshake is used to establish a session through TCP. See Figure 14-1.

Figure 14-1 TCP uses a three-way handshake to establish a communications session.



Here's how the three-way handshake works:

- 1 The client transmits a synchronization request (SYN) segment. This is a sequence number generated by the client.
- 2 The server sends an acknowledgement (ACK), which is the client's original sequence number plus 1. The server's SYN is a number generated independently by the server.
- 3 The client adds 1 to the server's SYN and returns it as an ACK. After both computers have acknowledged each other's communications, the connection is established.

A SYN attacker starts a large number of session requests (usually using a "spoofed" IP address, as described in the next section). The receiving computer puts these requests in a queue to wait for the completion of the process. By filling the queue and keeping it full, the attacker prevents other session requests from being established. Thus, legitimate users are unable to connect to the server.

IP Spoofing

IP spoofing involves altering the packet headers of messages being sent. This makes them appear as if they came from an IP address other than the actual originating address. Although spoofing is not in itself a form of attack, it is a method of gaining unauthorized access to a computer or network to launch an attack, to steal data, or to destroy data.

Computer Viruses and Worms

A *computer virus* is a program that can replicate and spread from one computer to another by copying its code to other files stored on the system without the user's consent or knowledge. Just as with the biological variety, some viruses are deadly and others are just annoyances. For example, the benign type might display a message on the user's screen. The more malevolent ones damage or destroy data or erase operating system files so that the computer can't be booted.

Viruses have received a great deal of media attention in the past, when programs such as Melissa, CIH (Chernobyl), Michelangelo, and the Iloveyou virus invaded thousands of computer systems all over the world.

A *worm* is a form of malicious virus that replicates itself and damages files on a computer. Worms are often disseminated as e-mail attachments, as executable files, as documents containing macros, or as HTML pages containing scripts.

Trojan Horse Programs

A *Trojan horse* is a program that presents itself as another program to obtain information. For example, there is a Trojan horse that emulates the system login screen. When a user types in his or her account name and password, the information is stored or transmitted to

the originator of the Trojan horse. The username and password can then be used to gain access to the system.

Internal Threats

Many network security policies focus on the Internet and external threats almost exclusively. This is usually a mistake. Just as retail companies find that their own employees commit as much (or more) theft of merchandise as “external” shoplifters, network administrators must not discount the risk of internal security breaches. Many instances of data theft, misuse, or destruction are “inside jobs.”

There are several motives for internal security breaches, including the following:

- Corporate espionage
- Internal politics
- Disgruntled employees (including ex-employees)
- Accidental breaches

Next, we look briefly at the characteristics of each and how to guard against them.

Corporate Espionage

Corporate espionage is the most sophisticated type of internal security threat. Theft of trade secrets is big business, and companies can become overnight successes or failures because of it.

Employees can be approached by competing companies and offered lucrative rewards for delivering an organization’s secret information. In other instances, employees of other companies can procure a job with your company to infiltrate and gather data to take back to the competitor, and they get to draw a paycheck from both companies simultaneously.

Finally, there are freelance corporate spies who take assignments on a contract basis. In highly competitive industries, they can steal the data on their own and then auction it to the highest bidder. “Data kidnappers” can even hold your confidential data hostage, promising not to release it to another company in exchange for a “ransom” payment.

Corporate spies are often intelligent, highly skilled, technically sophisticated individuals. They are usually well financed and can evade detection until it’s too late. If your business is part of a field in which corporate espionage is common (fields include the technology industry, oil and energy, research medicine, engineering, and others in which success hinges on being first to market with innovative products or services), your network can be vulnerable.

Security measures designed to thwart these professional spies must be of the highest level. You might need to call in consultants who specialize in protecting corporate networks from such infiltration.

Internal Politics

Another internal risk is the competitive employee who will do anything to get ahead and beat out his colleagues for the next promotion. This person might attempt to access and sabotage the work of those he views as personal competition. Others might attempt to harm the reputations of fellow workers by searching their e-mail or files for embarrassing messages or personal information. A particularly unscrupulous version of this ambitious person will even plant incriminating items on the victim's computer or send embarrassing e-mail that appears to originate from the victim.

Although this type of misuse of the network can pose a less serious threat to the company's important data than does espionage, it can cause a multitude of problems within the organization. Because many of these perpetrators are not highly skilled technicians, good basic network security policies (such as strong passwords and security auditing) can prevent or expose their actions.

Disgruntled Employees

A particularly destructive type of security threat is the employee or ex-employee who has a grudge against the company and wants to harm it. These people might destroy crucial data or disrupt vital network communications "just for fun" or out of a misplaced sense of "justice."

The company can be especially vulnerable to this threat when a technically knowledgeable employee is fired. It is not uncommon for an angry person to avenge his or her termination by erasing hard disks, by damaging computer hardware, or by releasing viruses onto the company network.

Your security policy should address this issue. Common sense dictates that terminated employees' user accounts be immediately deactivated and, in some instances, that their physical access to the company's computers be curtailed. The latter can be accomplished by assigning someone to escort them when they pack their personal belongings or when they otherwise reenter the company offices.

Accidental Breaches

In many cases, internal security breaches are not deliberate acts but are caused by the technical ignorance or lack of training on the part of the employees. Network administrators are well acquainted with the user who destroys necessary operating system files while

trying to “fix” a minor problem or who decides to “free up some disk space” by deleting all the application program files.

Accidental breaches are why operating systems with strong security features are appropriate for the business world. By implementing file-level permissions through the System Policies feature of Windows NT or the Group Policies feature of Windows 2000, users can be prevented from deleting or moving crucial system files.

Rebellious Users

Internal security breaches can also be the result of users who disagree with security policies that they believe are too restrictive. While not “accidental,” these breaches are not designed to cause harm; instead, they are intended to enable the user to do something he or she can’t otherwise do. For example, if security controls don’t prevent users from installing application software or hardware drivers, a rebellious user who doesn’t like your Internet access policies can connect an external modem, plug in a phone line, install the drivers, and dial in to his or her own ISP. Another user can install a remote-control application such as PCAnywhere, which can open up that computer—and the internal network—to outside users who also have PCAnywhere installed on their computers.

Your response to rebellious users will depend on the company’s policies and the degree of security required on the network. Implementing tighter controls might be appropriate; in other situations, it might be appropriate to evaluate whether your security policies really are unnecessarily stringent and whether you should allow users to have more legitimate access.

Implementing Security Measures

Protecting the network’s data often requires that a combination of security methods be applied. In this section, we discuss how user accounts, groups, and permissions can serve as the first line of defense. Then we look at how sensitive data can be encrypted, either in files on the disk or in packets that travel across the network.

Because e-mail is one of the most widely used network applications and because it is the form of network communication most likely to contain confidential information, we focus on the technologies available to secure e-mail messages. After you learn about how security measures are implemented, we discuss how these measures actually work.

Operating System Security Levels

Some operating systems are inherently more secure than others. Older operating systems, such as MS-DOS, do not enable user accounts to access the operating system (although an account is required to access the network). Anyone can sit down at the DOS machine, boot into the operating system, and then run its programs or access its files on the hard disk.

More modern operating systems, such as Windows 95, enable you to create computer user accounts. However, the passwords might be kept in an insecure text file or accessed by booting the computer from an MS-DOS floppy. By default, if you select Cancel when prompted for a username and password, the operating system still boots to full usability.

High-security operating systems such as Windows NT, Windows 2000, and Linux *require* that you enter a valid username and password to boot into the operating system and use it. They also store passwords in an encrypted form that can't be easily accessed.

The U.S. government provides specifications for rating computer security levels. These are discussed in the "Developing Security Policies" section of this chapter.

Users, Groups, and Permissions

Modern operating systems enable multiple users to access the computer and the network by creating separate user accounts and allowing a different password to be assigned to each. The combined username and password, when entered at a login prompt, results in the following:

- The user is given access to the operating system and network.
- The user can read and write to those shares to which his or her account has been given permissions.
- The user can exercise whatever rights (such as the right to shut down the computer or to install programs) have been assigned to his or her account.
- The user's preferences, such as desktop icons and wallpaper, are loaded.

In some computing environments in which there is no sensitive data stored on computers and the attitude toward security is casual, separate user accounts might not be created. All users can use the same account to log in. This means that every user has the same access permissions. Not only does this create an open, unsecured environment, but it also means users cannot customize their desktops and other settings.

In most business organizations, it is better to create separate accounts for different users. If unauthorized access is not an issue, passwords can be left blank.

Password Security

An important part of any network security plan is ensuring that user passwords are as impenetrable as possible. If users are allowed to select their own passwords, rules should be set according to the security needs of the network.

Some examples of password policies include the following:

- Passwords should never be words or numbers easily guessed because of their association with the user (for example, the user's Social Security number or a dog's name).
- Because many brute-force attack programs use lists of common words, passwords should not be words that are listed in the dictionary. A good policy is to combine letters and numbers (for example, "heavy238").
- In most operating systems, passwords are case-sensitive, so a password that contains randomly capitalized letters is difficult to guess (for example, "mYdoGspOt").
- A password *should* be easy for its user to remember. If it's not, he or she will probably end up writing it down.
- The more characters a password has, the harder it is for someone to guess it; thus, "visualize" is a better password than "see." This factor must be balanced against the user's ability to remember a long password.
- In a high-security environment, users should be required to change their passwords periodically. The new password should not be similar to the old one. For example, the user shouldn't change a password from "panda3" to "panda6." Again, it's important to maintain balance, because users who are required to change their passwords too often will not be able to remember them, and they will write the passwords down.
- Most network operating systems enable the administrator to set criteria such as minimum password length, password history (which keeps a list of the user's previous passwords and prevents reuse), and password expiration (which forces the user to change the password at specified intervals). There are also password security programs available that enable administrators to set parameters for passwords. For example, a parameter might prevent passwords that are out of the dictionary.

NOTE

Some security-conscious operating systems (such as Windows NT and Windows 2000) will not enable you to leave passwords blank by default; an administrator has to change the password policy settings to allow this.

Access Control Policies

Modern network operating systems (NOSs) enable administrators to control access to resources on a granular basis. The marysmith user account can be granted permission to view the contents of file1.doc but not allowed to make changes. At the same time, she might have permission to view, change, delete, or even set the access permissions on file2.doc, but no permissions to access file3.doc.

Each individual user can be given exactly the permissions needed. Using high-security file systems, permissions can be set not only on resources accessed across the network, but also on those same resources when accessed from the local machine. Local and network permissions do not have to be the same. Mary could be given full control over file4.doc when she accesses it while logged in to the machine on which it is stored, but limited to reading it if she accesses it from a remote machine across the network.

It is important to be aware of the default permissions for the NOS in use. On a Windows NT or Windows 2000 server, for example, when you share a resource, it is fully accessible to everyone on the network unless you explicitly change the permissions. On a NetWare server, the opposite is true; the share is accessible to no one until you explicitly change its permissions.

Neither method is right or wrong, but if you don't know how your NOS permissions work, you could end up allowing access to resources that you hadn't intended *or* prohibiting access to users who should have it.

In a security-sensitive organization, policies should be established to govern who should have access to which resources. Generally, access should be granted on a need-to-know basis. If the user needs access to the resource to perform his or her job duties, permissions should be granted. Otherwise, access should be denied.

Using Security Groups

Security groups are an administrative aid supported by many NOSs to make it easier to assign permissions on a large network. Groups are created, and those groups are given permission to resources. Then the appropriate user accounts are placed into the group, instantly giving the users all the permissions assigned to the group. This is easier than assigning permissions to individual users.

For example, if everyone in your sales department needs access to several folders (we'll call them salestats, salesbudget, and salesmemos), you could assign access permissions for each folder to all 20 users in the department. However, an easier method is to create a security group called sales, give it permission to all three folders, and put the 20 user accounts into it. Then if you create a new folder, salescalendars, that needs to be accessed by the entire department, all you have to do is assign permissions to the sales group. Otherwise, you would have to assign permissions for the new folder to 20 different accounts.

If you don't want a user to have access to a certain resource, don't assume that removing that permission from his user account is enough. You must also ensure that the user is not a member of any group that has permissions to access the resource. In addition, if the operating system (for example, Windows NT or Windows 2000) enables you to specifically deny access, you must explicitly do so.

File Encryption

Encryption involves converting data into a form that can't be easily understood by others. The technical aspects of encryption are discussed in more detail in the "How Security Components Work" section later in this chapter.

File encryption is a means of encrypting data stored on a computer's disk so that it is unreadable to anyone but the creator of the data. Some operating systems, such as Windows 2000, include a file-encryption function. For those that don't supply such a function (for example, Windows 9x and Windows NT), third-party encryption programs are available.

When documents are encrypted on the disk, only a user who has the correct key can view them. If others attempt access, either the file won't open at all or it appears as scrambled, meaningless characters. Note that sensitive data should be protected by both access permissions and encryption.

IP Security


File encryption protects data stored on a disk, but it does not offer security for data as it is sent over a network. The IP Security (IPSec) protocol was developed to remedy this shortcoming. IPSec secures data at the packet level. Because it works at the network layer of the OSI reference model, applications are not aware of it. Cisco Systems includes support for IPSec in its routers, and Windows 2000 includes IPSec in its TCP/IP stack.

IPSec uses two protocols:

- **Authentication Header (AH)**—Enables verification of the sender's identity
- **Encapsulating Security Payload (ESP)**—Ensures the confidentiality of the data itself

These two protocols can be used separately or together.

IPSec can operate in two modes: transport mode and tunnel mode. See Figure 14-2. Transport mode provides *end-to-end security*; that is, the encryption is in place from the source computer to the destination computer. Tunnel mode protects the data from the exit point of one network to the entry point of another.

Figure 14-2 Packet headers differ, depending on whether IPsec is used in transport mode or tunnel mode.

Secure Sockets Layer (SSL)

SSL is another means of securing communications on the network. The disadvantage of SSL is that it operates at the application layer; thus, it must be supported by the user application.

SSL was developed by Netscape to provide security for its Web browser. It uses *public* and *private key encryption*. These are discussed in the “How Security Components Work” section later in this chapter.


E-mail Security

Many computer users enjoy a false sense of security about network communications in general and e-mail messages in particular. Users assume that the messages they compose and send over the local network or the Internet are read only by the recipient to whom the messages are addressed. E-mail users behave as if they have the same expectation of privacy when sending e-mail as they do when sending a letter through the postal service. A more accurate expectation would be to assume that the e-mail is like a postcard that can be read by anyone who handles it during its journey from sender to recipient.

E-mail messages are very easy to intercept (see Figure 14-3). They often travel through dozens of nodes (servers) on their way from sender to recipient. Even if a message is sent to someone within the local network, a copy of it is stored on at least three machines: the sender’s computer, the recipient’s computer, and the internal mail server. E-mail sent over

the Internet can pass through several servers as well. Unless it is encrypted or digitally signed, the message can be easily read, copied, or altered at any point along the way.

Figure 14-3 *E-mail messages can be intercepted at the many points through which they travel as they move across the Internet.*



WARNING Deleting the message on the user's machine does *not* guarantee that there is not still a copy of it on one or more servers.

Many software products have been developed to address this problem. These products are designed to accomplish several goals:

- The message cannot be read by unauthorized parties.
- The message cannot be altered between the time it leaves the sender and the time it is opened by the recipient.
- The person identified as the sender of the message is actually who he or she purports to be.

An important factor in e-mail protection software is ease of use. If users must perform complex or time-consuming tasks to send or read secured e-mail, they are likely to abandon the security software altogether.

Popular e-mail protection programs include the following:

- Pretty Good Privacy (PGP)
- Kerberos
- Baltimore Mail Secure
- MailMarshal from Softek

Most e-mail protection software uses public key encryption (discussed in “Public/Private Key Encryption” later in this chapter) and digital signatures to provide data confidentiality and identity authentication.

How Security Components Work

Computer security technologies are based on the science of *cryptology*, which is the study of “secret writings” or *ciphers*. Encryption uses a code or key to scramble and then unscramble (or *decipher*) a message to return it to its original form. The task of analyzing and deciphering encrypted messages is called *cryptanalysis*.

Basic Cryptography Concepts

Data is encrypted using an *algorithm* or cipher. The encrypted data is referred to as *ciphertext*. To understand encryption, consider the “secret codes” used by children (and sometimes adults as well) to protect the privacy of their communications. A very simple encryption scheme assigns a number to each letter of the alphabet. If we start with A as 1, B as 2, and so forth, we can send an “encoded” message such as the following:

4-9-14-14-5-18 9-19 18-5-1-4-25

The message is easily “decoded,” or *decrypted*, back to the original:

Dinner is ready

Unfortunately, this cipher is easy to crack and thus not very secure. Computer encryption is more complex. Generally a *key* is used, which is a variable of some sort that is combined with the data to be encrypted. The method of combining the key with the data is called the *algorithm*. Because computers process binary information, these calculations are applied to each bit or group of bits (see Figure 14-4).

Figure 14-4 Computerized encryption techniques use keys and algorithms.



The longer the encryption key, the more difficult it is to break the code. In the field of encryption, 40- and 56-bit encryption is called *standard encryption*, and 128-bit encryption is called *strong encryption*.

Laws Governing Encryption Technologies

Many countries have legal restrictions on the export or import of encryption technologies. The U.S. classifies encryption as “munitions” and prohibits export of strong encryption software. Cryptology up to 40 bits can be exported, and cryptology up to 56 bits can be exported under some circumstances. 128-bit encryption technologies cannot be exported. There are no restrictions on importing encryption items.

The laws in other countries range from Argentina, which has neither import nor export controls and no restrictions on the use of cryptography, to Russia, where both import and export are regulated, and the use of unauthorized encryption within the country is prohibited.

Two popular types of encryption are secret key encryption and public/private (or public) key encryption. Let’s look at how each works.

Secret Key Encryption

Secret key encryption is often referred to as *symmetric encryption* because one key is used to both encrypt and decrypt the data. The sender and recipient agree to use a common key, or encryption algorithm, which is a *shared secret*. For example, if Ted and Carol wish to exchange private messages, they decide on the following code, or key, to represent each letter of the alphabet:

- 1 Convert the letter to a number using a scheme in which A = 6, B = 7, C = 8, and so forth.
- 2 Multiply the resulting number by 4.

This key, although by no means impossible to break, is quite a bit more secure than our first example. Now to send the message “Dinner is ready,” Carol goes through the following process to arrive at the final encrypted message:

$$\begin{aligned} D &= 9 \\ 9 \times 4 &= 36 \\ I &= 14 \\ 14 \times 4 &= 56 \end{aligned}$$

Eventually, she arrives at the final encrypted message:

36-56-76-76-40-92 56-96 92-40-24-36-120

Ted has the same key and knows that to decrypt the message, he need only divide each number by 4 and then apply the same A = 6, B = 7, and C = 8 table to restore the original message.

Popular secret key encryption algorithms are Data Encryption Standard (DES) and 3DES (pronounced “triple DES”). Another is RC (Rivest Cipher)-4, created by Ron Rivest. This person, along with Adi Shamir and Leonard Adleman, developed the popular RSA public key encryption scheme, which is discussed in the next section of this chapter.

There are three inherent problems with secret key encryption:

- Generating the secret keys
- Exchanging the keys between authorized parties without having them fall into the hands of unauthorized parties
- Dealing with the complexity involved in securing communications to many different parties

It is prudent to change the keys regularly to avoid compromising security. This means additional keys must be generated. There also must be a way to get the key to the party who is authorized to decipher the message. In our previous example, if Carol e-mailed the key to Ted without encryption, she would defeat the purpose of securing the communication. However, if she e-mailed it to him encrypted, he would not be able to read it, because he doesn't know the key.

Mechanisms have been developed to generate keys and exchange them securely. One example is the Diffie-Hellman algorithm, which enables two parties to create a secret known only to them, despite the fact that they are communicating on an unsecured network.

The third issue in our previous bulleted list is a bit more problematic. Carol and Ted can happily exchange (somewhat) secure messages for years using their secret key, but what happens if Carol wants to send a secure message to Bob? If she uses the same key, Ted would be able to decipher it. In addition, once Bob has the key, he can decipher Carol's messages to Ted as well.

To remedy the situation, Carol and Bob could come up with an entirely different key. For example, messages to Bob could be encrypted by converting the letter to a number using the numbering scheme A = 12, B = 13, and so on, and then adding 15 to the result. However, now Carol has to remember two keys, and if she wants to send secure messages to a large number of people, this technique quickly becomes unmanageable. A simpler solution is to use a different type of encryption: public/private key encryption.

Public/Private Key Encryption

Although often referred to as public key encryption for brevity, the more accurate term is *public/private key encryption*, because this type of encryption uses two keys, one of which is published and widely available, and the other of which is private and known only to the

user. Both keys are required to complete the secure communication. This type of encryption is also referred to as *asymmetric encryption*. With this type of encryption, each user has both a public and a private key, called a *key pair*. Here's how it works:

- 1 Carol and Ted exchange their public keys. It doesn't matter that this is done in an insecure manner, because the messages cannot be deciphered with just the public key.
- 2 Carol wants to send a message to Ted, so she encrypts the message using Ted's public key. A public key is associated with only one private key. To decrypt a message that was encrypted using a public key, the associated private key is required. (The reverse also applies—to decrypt a message that was encrypted using a private key, the associated public key is required.)
- 3 Ted, using his private key, can decrypt the message because it was encrypted using his public key. Notice that only Ted's keys, public and private, were used in this encryption process.

If Carol had encrypted the message using her private key, anyone could decrypt the message using her public key, which is available to all.

Both keys of the same key pair must be used for this encryption to work, and there is no need for anyone to know anyone else's private key. A good way to understand this type of encryption is to think of the two pieces of information required to enter a home protected by a digital combination lock. If someone wants to enter the house, he or she must know both the street address and the number sequence to enter into the locking device. The address is public information that is published in the telephone directory. It is available to anyone, just as the user's public encryption key is available to anyone. The lock combination is analogous to the user's private key; it is known only to the owner of the house. Both keys are unique to that particular home, but one is made known to the public while the other is kept secret.

Authentication

The examples we have discussed thus far deal with protecting the *confidentiality* of the data. A separate issue is the *authentication* of the sender's identity.

Using the public/private key method, you can easily secure the data itself. However, because the public key *is* public and available to anyone, there is no way for Ted to know for sure that a message he receives, encrypted with his public key, really is from Carol. For the identity of the sender to be authenticated, Carol would have to encrypt the message with her *private* key. Ted can then decrypt it with her public key, confident that she was the actual sender because she is the only person who knows her private key.

Another way to ensure the authentication of the sender is to use *digital signatures*, which we discuss next.

Digital Signatures

Digital signatures consist of encrypted signing information appended to a document. This information verifies both the identity of the sender and the integrity of the document itself. Digital signatures don't encrypt the data. They only ensure that it has not been altered and that the sender is authentic.

Public key algorithms are used to create and verify digital signatures and *hash algorithms*. We discuss hash algorithms next.

Hash Algorithms

A *hash* is the result of a one-way mathematical calculation (the *hash algorithm*) that creates a *message digest*. The algorithm is called "one-way" because you cannot reverse-engineer the result to discover the original message.

The hash verifies the authenticity of the message in the following manner:

- 1 The sender hashes the message with a key that is a shared secret; that is, the key is known to both the sender and the intended recipient. The hash produces a numerical result (let's say it's 0010110010100001).
- 2 The message and the result, or *message digest*, are sent to the recipient.
- 3 To confirm that the message has not been altered, the recipient applies the same key, and he or she should get the same numerical result (0010110010100001). If the content of the message has been changed, the result of the hash does not match.

Popular hashing algorithms include the following:

- **Secure Hash Algorithm (SHA)**—Developed by the National Institute of Standards and Technology (NIST) and the National Security Agency (NSA)
- **Message Digest 5 (MD5)**—Developed by Ron Rivest at MIT

Generally, hash algorithms produce a value of at least 128 bits, making it extremely difficult or impossible to produce the same result with a different set of input data.

Digital Certificates

Digital certificates are messages that contain the digital signature of a trusted third party, or *certificate authority*. The third party warrants that a particular public key actually belongs to a specified person. Certificates are used to ensure the authenticity of messages that travel across unsecured public networks such as the Internet.

A user who has the private key that is associated with a particular public key requests a certificate from a certificate authority. The certificate authority has the responsibility for verifying that a specific public key belongs to a specific user. Certificates are valid for a specified period of time, and the certificate authority can revoke them.

Organizations can set up their own certification authorities as part of their network's *public key infrastructure (PKI)*. A PKI is a system of verifying and authenticating the identity of parties engaging in electronic communications.

Digital certificates are issued by public certificate authorities such as the following:

- Verisign
- GTE Cybertrust
- Keywitness
- TradeWave


Digital certificates can be thought of as electronic ID cards. Just as a government-issued driver's license, passport, or other identification document can be presented to establish your identity in the physical world, digital certificates perform that function in the world of network communications.

Kerberos

Kerberos is an authentication protocol that is an Internet standard for verifying the identity of a user or computer system. It was developed at MIT and named for the three-headed dog of Greek mythology that guarded the gates of Hades. Kerberos security relies on three factors: the client, the server, and a trusted authority called the *key distribution center (KDC)*. The KDC maintains a database that keeps track of the participating entities.

Kerberos is based on the concept of tickets, which are encrypted messages that are used to request service from a server. Authentication is performed using symmetric encryption (also referred to as secret key encryption). See Figure 14-5.

Figure 14-5 *The Kerberos authentication process depends on tickets to verify identity and to grant access to servers.*



Here's how the Kerberos process works:

- 1 The first step involves a request from the client to an authentication server for a ticket and *session key*. This is an encryption key based on the requesting user's password, and it is combined with a random value representing the service being requested. The ticket returned to the client is called the *ticket-granting ticket* (TGT). This verifies the identity of the client.
- 2 The TGT is sent by the client to a ticket-granting server (which can be the same physical machine as the authentication server). Based on the TGT, the ticket-granting server returns a ticket that is presented to the file server whose resources the client wants to access.
- 3 The ticket is valid for a specified amount of time so that the client can make additional requests of the server without going through the process again.

Kerberos has been in use on UNIX networks for many years. Microsoft included Kerberos support in Windows 2000, and the Active Directory domain trust relationships are based on Kerberos authentication.

Advanced Identification and Authentication Technologies

Identification and authentication are important security concerns as corporate networks become larger and must support thousands of users. Establishing the identity of a user attempting to log in to the network or access a computer's resources is the foundation of a strong security plan.

Regardless of the sophistication of password encryption technologies, the problem with the username/password method is the possibility of an unauthorized party discovering and using the credentials of a legitimate user. Advanced identification and authentication technologies such as the following attempt to prevent this by basing identification on factors that can't be faked:

- Smart cards
- Biometrics, including fingerprint recognition, retinal scan and iris recognition, and voiceprint verification

Although these technologies are not yet commonly encountered in the workplace outside of government agencies, over the next decade, we can expect advanced authentication methods to become commonplace, and it will be increasingly important for network administrators to be familiar with their implementation. The following sections describe each technology.

Smart Cards

Smart cards are authentication devices that resemble credit cards and that store information such as public/private keys and passwords. The storage technology can be as simple as a magnetic strip or as complicated as an integrated circuit that is embedded in the card and that functions like a tiny computer.

NOTE

When smart cards are discussed in the context of network security, the term refers to an identification and authentication device that stores credentials that enable a user to access a system. The term "smart card" is also used in a more generic sense to indicate any plastic card that stores info on a chip or magnetic strip. These cards include credit and debit cards, medical information cards, and so forth.

Some operating systems, such as Windows 2000, have smart card support built into them. Without an operating system that has built-in support, you might need separate software before you can use smart card technologies.

A smart card is used to access the computer through a piece of hardware known as a *smart card reader*. The smart card provides extra security because a user must have possession of the physical card. Smart cards can be used in place of—or, for added security, in addition to—physically entered user credentials.

In a high-security environment, the card alone does not enable access; the user must insert the card into the reader and enter the correct login credentials before access is granted. This means a potential intruder who finds out a user's password still can't log in without the card;

likewise, someone who finds or steals the card cannot gain access without knowing the correct password.

A variation on the smart card is the *I-key* from Rainbow Technologies. It is a small device that fits on a keyboard and that communicates through a USB port instead of requiring reader hardware.

Biometrics

It is still possible that an intruder could gain knowledge of the password and possession of the corresponding smart card. To counteract this possibility, fingerprint recognition, retinal scans, and voice verification technologies take security a step further. These technologies are known as *biometrics*, which is the field of biological statistics.

Biometrics work because the statistical probability of two people having the same fingerprints (or retinal patterns or voice patterns) is so small that these technologies are admissible in court in many countries as evidence in criminal proceedings. Biometric identification technologies can provide a positive identification of a user based on biological characteristics.

In the following sections, we examine how various biometric methods work.

Fingerprint Recognition

A fingerprint recognition system often operates in conjunction with card-based security. The users' fingerprints are taken and entered into the database, and when a user wants to access the computer, he must place the finger (or more often, the thumb, although some systems use an entire handprint) onto a reader screen. The input from the reader is compared with the prints in the database, and access is granted or denied.

Various software and hardware systems are available for implementing fingerprint authentication. These systems include VeriPrint and U.are.U, both of which are offered by FingerSec Corporation.

Retinal Scan and Iris Recognition

Another means of identifying persons based on their unique biological characteristics is the *retinal scan*. In this technology, a low-intensity light source is used to scan the patterns of the retina, which is the delicate membrane that lines the inner eyeball.

Also based on the characteristics of the human eye is *iris recognition*. This technology involves computer analysis of the patterns found in the iris of the eye, which is the pigmented round membrane that gives the eye its color. These patterns are said to be even more unique than fingerprints or DNA. Even identical twins have distinct iris patterns.

Voiceprint Verification

Voiceprint verification operates on the principle that human voice patterns are unique to each individual and can be used to identify the speaker.

NOTE

Voice verification is sometimes confused with *voice recognition*. The former is used to verify the identity of the speaker. The latter does not identify the speaker, but recognizes the words spoken and, in the case of voice recognition computer software, translates them into written form on the screen.

Generally, the user is required to record a password or phrase, which is stored in the database. To log in, the user speaks the same word or phrase, and the patterns are compared. Such factors as pitch, tone, and cadence are considered. Voiceprint technology is not considered to be as accurate as retinal scans, iris scans, and fingerprint recognition. Voice verification is commonly used in situations in which identity needs to be authenticated over a phone line.

The Future of Biometrics

Some parts of the world, such as Australia, South Africa, South America, and Europe, have moved quickly to adopt biometric recognition and authentication technologies. The United States has been slower to implement them, perhaps because of privacy issues or because of market factors.

As technologies become more accurate and less intrusive and problems such as identity theft gain more attention, it is likely that the use of biometrics will grow and become commonplace in the networks of the future.

Developing Security Policies

New security technologies are emerging constantly. To develop a good security plan for a network, you must be aware of what is available. You also must be able to determine the level of security that is necessary or desirable for your situation.

NOTE Biometrics specialists make a distinction between recognition and authentication. Authentication is regarded as a voluntary activity in which the user provides a name or identifying number, along with a biometric input such as a fingerprint, usually for the purpose of gaining access to a system. The biometric is compared with a stored biometric know to belong to the user. The comparison is used to verify the user's identity, and if verification succeeds, access is granted. In recognition, the system gets only the biometric input (voice sample, retinal scan, or fingerprint) and then must search its database and find a match. For example, recognition is used when forensics experts attempt to identify a fingerprint left at a crime scene. To make the identification, experts might search for a match in a database such as the FBI's Automated Fingerprint Identification System (AFIS).

Security policies should be the product of a team effort. Input should be solicited from technical personnel, management, and representative users. To be successful, a security policy must have the support of the company's managers and users. Budgetary and philosophical considerations must be weighed against the sensitivity of the data on the network and the ramifications if it were compromised.

Security policies should be in writing and should be reviewed and revised periodically as circumstances change. The first step in creating policies is to perform a detailed security analysis. You should determine what security measures are currently in place, whether they are accomplishing their purpose, and which of them should be removed, retained, or replaced.

In the following sections, we address the following important factors that you should consider when developing policies:

- Acceptable use policies
- Termination policies
- Government security ratings
- Security auditing and intrusion detection
- Firewalls and proxies
- Security through multiple protocols
- Physical security

Acceptable Use Policies

The network security policy should include an acceptable use policy that defines how users can legitimately access and use the resources on the network. This should include items such as the encryption of messages and files, Web site access, the downloading of files from

the Internet, bandwidth usage, the installation of programs and games, e-mail policies, and other end-user issues.

Termination Policies

A network security policy should address the procedures for ensuring continued integrity of the network data when an employee—especially one in a technical position—leaves the company either voluntarily or through termination of employment.

It is important that all company property be accounted for and that employees turn in smart cards and other access devices. Accounts of terminated employees should be immediately disabled. If the employee had access to sensitive data, he or she should not be allowed to take personal floppy disks, zip drives, or other data storage media without having that media checked to ensure that there has been no unauthorized copying of confidential company files.

Government Security Ratings

The U.S. government provides criteria for rating security implementations. These specifications are defined in *Department of Defense Trusted Computer System Evaluation Criteria*, which is referred to as TCSEC and sometimes called the “orange book.” TCSEC is published by the National Computer Security Center (NCSC). It is used in conjunction with *Trusted Network Interpretation of the TCSEC*, which is referred to as TNI and sometimes called the “red book.”

The TNI applies evaluation criteria for networks. Ratings start with A, which is the highest security rating, and go to D, which is the lowest rating. The C rating is divided into two subratings, C1 and C2; however, C1 is no longer used as a certification. A C2 rating, which is a higher rating than a C1, is sought by many businesses to obtain government contracts. A C2 rating requires that the operating system be able to track when and by whom data is accessed. A C2 operating system must have the capability to control users’ access to objects, provide for unique identification of users, and include a means to audit security-related events.

If your organization requires that its computer systems have a C2 rating, you should ensure that they meet all the criteria in TCSEC and TNI. You even might have to have the security configuration certified and accredited by administrative authorities to qualify as C2-compliant for government contract work.

NOTE Other countries have similar security rating systems. These systems include the Canadian Trusted Computer Product Evaluation Criteria (CTCPEC) in Canada, the Australian Information Security Evaluation Programme (AISEP) in Australia, and the Information Technology Security Evaluation Criteria (ITSEC) in Western Europe.

Table 14-1 lists various operating systems and the security ratings they have received.

Table 14-1 *NSA Security Ratings for Various Operating Systems*

Operating System	Vendor	NSA Certification
UNIX XTS-200 and 300	Wang Government Svcs	Orange Book B3
UNIX Trusted Xenix 3 and 4	Trusted Information Systems	Orange Book B2
UNIX HP-UX 8.04 and 9.0.9	Hewlett Packard	Orange Book B1
UNIX UNICOS 8.0.2	Cray Research	Red Book B1
UNIX RS/6000	IBM	Orange Book C2
Windows NT 3.5/SP3 and Windows NT 4.0	Microsoft	Orange Book C2
NetWare 4 and 4.11	Novell	Red Book C2

Merely installing the specified operating system does not guarantee that you meet the criteria of the security certification shown. There are other specifications, such as network connectivity and the enabled operating system features, that must be considered. In actuality, it is not the operating system itself that receives a security rating, but the entire hardware and software configuration.

Security Auditing and Intrusion Detection

One requirement for the C2 security rating is the capability to audit security events and the activities performed by individual users. *Auditing* is the process of tracking the activities of users and the system. For example, auditing includes monitoring which files are accessed, when, and by whom. Auditing can include information on who has logged in to or out of the system, who has accessed objects, and who has exercised user rights.

Operating systems such as Windows NT and Windows 2000 have auditing built into them, and auditing can be configured on a granular basis. Security events are not only tracked, but also recorded in a log file for easy review. Only administrators should have access to the security logs.

Passive Detection

Security auditing is referred to as a *passive* form of intruder detection. Although events affecting security are logged to a file, an administrator must suspect intrusion and check the file to learn of the breach. Examples of passive detection programs are Tripwire (for UNIX) and the built-in security-auditing feature of Windows NT and Windows 2000.

Active Detection


An environment in which network data is sensitive should not rely on passive intrusion detection alone. A security plan should include one or more forms of *active detection*. In active detection, software continuously scans the network for signs of intrusion, and some programs even alert the administrator and disconnect the suspicious session. SATAN (the ominous name stands for Security Administrator's Tool for Analyzing Networks) and NetRanger are examples of such software packages.

An updated version of SATAN has been released. It carries a more pleasant name—Security Administrator's Integrated Network Tool (SAINT). SAINT can detect additional vulnerabilities, and it is available in an easy-to-use version called WebSaint, which enables administrators to check for system vulnerabilities over the Internet.

Firewalls and Proxies

Strong perimeter security is another important consideration in establishing security policies. *Firewalls* and *proxies* can be used to create a barrier between the local internal network and the connection to the outside world. This area can be set up in its own subnet, and it is sometimes referred to as the *demilitarized zone* (DMZ) or *screened subnet*. See Figure 14-6.

Figure 14-6 Firewalls and proxies can be configured in a DMZ between the internal and external networks.



A firewall can be hardware- or software-based. It provides a means of filtering incoming and outgoing packets, and it determines (based on policies set by network administrators) whether to allow the packets through to the destination address. The firewall is typically located at the network's *gateway*, which is the point at which the network connects to another network.

Three basic types of filtering are performed by firewalls:


- Packet filtering
- Circuit filtering
- Application filtering

Packet filtering filters data packets based on the information in the IP, TCP/UDP, and ICMP headers. With packet filtering, you can enable or block specific IP addresses or port numbers. *Circuit filtering* is based on the connection at hand. If a packet isn't part of an established connection, it won't be allowed through the firewall. Finally, *application filtering* filters according to protocols used for specific IP applications. For instance, Java applets or Visual Basic scripts could be blocked.

Hardware-based firewalls are sometimes called "black boxes." They are dedicated computers that run a proprietary operating system (or, in some cases, UNIX). They function *only* as a firewall, and thus are faster and more stable than a computer that is running firewall software while providing other computer functions.


Proxy servers operate as “middlemen” in the network by performing functions similar to firewalls. Computers on the internal network communicate with the proxy, which then communicates “on their behalf” with computers on the external network (see Figure 14-7).

Figure 14-7 A proxy server acts as a “middle-man” between the internal and external networks.



Proxy servers also provide services such as *reverse proxy* and *reverse hosting*. Reverse proxy enables a proxy server to redirect external HTTP requests to a single designated machine. This enables secure access to an internal Web server without exposing the server to the external network. Reverse hosting enables the proxy server to redirect HTTP requests to more than one Web server by mapping several servers to one logical address (see Figure 14-8). Proxies also provide the caching of Web pages to improve Web performance.

Figure 14-8 Reverse hosting enables a proxy server to redirect HTTP requests to multiple Web servers.



Routers and Layer 3 switches can be configured with *access control lists (ACLs)* that restrict specific machines from using the router or that enable communication in one direction only. In this case, the router provides firewall protection.

Routers with configured ACLs can be thought of as a “first line of defense” for your network. They are often used to protect software firewalls from tampering.

NOTE The ACLs used by the router enable or deny access based on factors such as IP addresses. This means a spoofed IP address can defeat the router’s firewall protection.

Security Through Multiple Protocols

A network security policy should also address which protocols are allowed or required. You can provide security to the internal network by using a protocol stack other than TCP/IP, either on the LAN itself or in the DMZ subnet that separates it from the external network.

If you use IPX/SPX as the network/transport protocol on the local network instead of using TCP/IP, IP-based attacks from the Internet cannot penetrate the LAN. You simply configure a router to use TCP/IP on the side that connects to the Internet and IPX/SPX on the side that connects to the local network. If you need TCP/IP for communication between internal computers, you can use it on the LAN, but you should set up a DMZ that uses IPX/SPX between the LAN and the Internet.

Physical Security

A network security policy should control *physical access* to the network components. This is an important but often overlooked factor in developing an effective security plan. You should assess the degree of access that employees, contractors, clients, and the public have to the workstations, servers, cable, other media, routers, switches, and other physical components.

In a high-security environment, servers and connectivity devices should be kept behind locked doors. Workstations that are in nonsecured areas should have software controls that prevent access to sensitive network data. Cable should be protected by heavy conduits, not exposed on the floor or ceiling. Twisted-pair and coax are especially easy to tap into, but it is also possible to intercept data over fiber optics by using an optical splitter. Because the signal level is affected when a tap is made in this way, an observant network administrator can detect the tap with an optical time domain reflectometer (OTDR). However, experts can insert a tap that cannot be easily detected by the OTDR.

In many organizations, great pains are taken to restrict employees’ and the public’s access to the computer equipment—and then it is all left wide open every night when the janitorial staff comes in to clean. If security is an issue on your network, cleaning and maintenance personnel should be under the supervision of an authorized person when they are in the rooms containing network equipment.

Disaster Protection and Recovery

Intrusion is not the only threat to the network and its data. Hardware failure, natural disasters, and technical mistakes can also result in a devastating loss of important files. This is why disaster protection and recovery measures are essential for any production network.

Protecting against and recovering from catastrophic data loss involves several lines of defense:

- Power backup
- Data backup
- Disk fault tolerance
- Server fault tolerance (clustering)

We briefly review the basic principles of each and how each can be implemented on a network.

Power Backup

Numerous companies have had the experience of losing precious data when a summer thunderstorm caused the lights to flicker and the computers to reboot or when an electric company glitch created a sudden surge of high voltage. These incidents are all the more frustrating because they are preventable.

A great deal of data is lost or damaged every year because of electrical power surges and failures. In this section, we look at how you can protect the network from the dangers of power fluctuations with surge protectors, uninterruptible power supplies (UPSs), and power generators.

Surge Protectors

A power spike or surge can damage computer hardware and render data unreadable. A good surge protector and suppressor are the absolute minimum protection that you should provide for each important computer on the network, including all servers and mission-critical workstations.

NOTE

Surge protectors are the lowest-cost power protection devices, but their effectiveness is limited to increases in power voltage. They do not protect the systems if the power voltage decreases or is turned off. Many surge protectors cannot withstand multiple power surges—after experiencing a surge and doing its job, a surge protector might need to be replaced. Unfortunately, it is not always possible to detect that this has occurred, so you might believe that you have surge protection when in fact you no longer do.

UPSs

A form of power protection that is better, but more expensive, than a surge protector is a UPS. A UPS is a type of battery backup that provides a limited amount of stored electricity on which the systems can continue operating after the power fails.

A UPS is *not* designed to provide power to continue using the computer systems. Its purpose is to enable uninterrupted power for several minutes, which gives you time to close files and programs and shut the computers down “gracefully.” The typical UPS provides power for 5 to 20 minutes after it kicks over to battery mode.

The UPS is plugged into the wall outlet, and the computer is plugged into the UPS. Many UPSs have multiple outlets so that several systems can be run off one UPS. The device is continually charging in normal mode. When the electricity fails, the UPS senses this and often can be configured to notify users that it is not receiving power. Notification can take the form of an audio signal (such as a beep), or software can be used to send a message to administrative accounts on the network. The same software can be configured to automatically start shutdown of the attached computer when the UPS goes to battery mode.

All crucial systems on the network should be attached to UPS devices. The cost of a good UPS system is a fraction of the cost of a new computer, and it is well worth the expense when compared to the cost of losing your data (which can be priceless).

NOTE

Although UPS literature often uses the terminology “battery mode,” in reality the computer equipment is always running off the UPS battery. It’s just that the battery is always being charged as long as the power from the wall outlet is active. This is why there is no interruption to the electricity flowing to the computers when the UPS “shifts modes.”

Generators

The next step up in power protection is the *power generator*. This device actually makes electricity by using an engine powered by gasoline, kerosene, or other fuel. A generator enables you to continue using your electrical equipment (including computers) for the duration of a power outage.

Generators are expensive, and their expense is usually not warranted except in situations in which the power is expected to be out for a long period of time or the equipment that needs the power is used for life support or other emergency purposes.

Data Backup

Despite your best efforts to prevent it, eventually it will happen: A hard disk will crash, a fire or flood will damage the server, or a malicious virus will format the drive and render

your files unreadable. Simply put, data will be lost. However, if you have implemented a regular, thorough data backup program, it's not gone forever.

Devising a backup plan involves answering the following questions:

- What files should you back up?
- When should you back them up?
- How should you back them up?
- Why should you back them up?

The answer to the fourth question should be self-evident; your backup plan can mean the difference between catastrophic loss of data, time, and money and the minor inconvenience of spending a few hours restoring files to their original state. Let's look at the answers to the other three questions and see how they fit into your plan.

What to Back Up

The first step is to decide what should be backed up and to assign priorities to the files to be backed up. The ideal situation, of course, is to back up everything. However, this might not always be feasible because of time constraints and limits on the capacity of backup media.

It's not always important to back up the operating system and application files, because these can be reinstalled from the installation disks. Original data is more important to back up; original data includes word processing documents, spreadsheets, and other data created in various user applications. Creative work such as graphic art or original writing compositions should have high priority, because they can be impossible to re-create exactly. Of course, mission-critical data, such as financial information on which managers depend to do their jobs, goes at the top of the "what to back up" list.

Your policies about what to back up should be in writing. Data that must be backed up should be stored in a central location, such as a specified drive on the server. If important data is scattered across the network on individual hard disks, it is too difficult to ensure that nothing is missed in the backup cycle.

When to Back Up

You should construct a schedule for periodic, *regular* backup of important data. How often to back up depends on how much data you can afford to lose: a day's worth or a week's worth.

Backups can be scheduled to take place after business hours so that you do not affect network performance. Most good backup software, as discussed in the next section, enables you to schedule automated backups so that no one has to be present to start the backup.

Most backup schedules include different types of backup. There are three basic types (although some backup utilities give you additional choices):

- **Full backup**—All data on the specified drives is backed up, regardless of whether or when it was backed up before and whether it has changed since the last backup.
- **Differential backup**—All files that have changed since the last full backup are backed up.
- **Incremental backup**—All files that have changed since the last backup of any type (not merely since the last full backup) are backed up.

The full backup requires the most time and space on the backup media, and thus, it is the most expensive. It is also the simplest method, and if you perform a full backup every night, you are ensured that all backed-up data is up to date. It is also the easiest to restore.

The differential backup saves time. It is done in conjunction with full backups. For instance, you can do a full backup once a week or once a month, and then do a differential backup every night between the full backups. To restore the data, you have to restore two tapes (the most recent full backup and the most recent differential) to ensure that all data is up to date. This method can be desirable if you have a very large amount of data to be backed up and little time in the evenings to perform the backup. The differentials take far less time. In addition, the full backups can be performed on the weekend when there is more time available.

The incremental backup is the fastest, but restoration is more complex and takes longer than with the other methods. You must restore the most recent full backup and then you restore the incremental backup for *every* day since the full backup.

NOTE

How does the backup software “know” which data has changed since the last backup? An *archive bit* is set on the file. This bit is a file attribute that is similar to the bit that marks a file as a hidden file or as a read-only file. This bit is removed (that is, cleared) when a file is backed up during a full backup or an incremental backup. The bit is *not* removed after a differential backup.

How to Back Up

The “how” of backing up includes many decisions:

- What backup medium should be used?
- What backup software should be used?
- Who is responsible for performing the backup?
- Will the backup be done manually, or will it be automated?

For many years, tape was the preferred medium, and some backup software (such as the backup utility built into Windows NT) does not support any other media. Other backup programs enable you to choose from a variety of media, including removable disks such as Zip and Jaz, CD-R and CD-RW, DVD, and magneto optical (MO).

Most operating systems include a backup utility. Windows 9x, Windows NT, and Windows 2000 include Windows Backup; NetWare includes SBackup; and UNIX has the *tar* archiving program built in. There are many third-party backup programs available that offer a fuller feature set than those built into the operating systems. ARCserve (Cheyenne), Backup Exec (Seagate), and Norton Backup (Symantec) are popular Windows backup programs.

It is important that the responsibility for backing up critical data be specifically assigned. Ensure that the backup operator has the appropriate permissions to back up everyone's data. It is not necessary to give a person permission to read files to enable him or her to back up and even restore those files. You should also have a "backup backup operator"—someone who is trained to take over the duties if the primary backup operator is out.

Whether the backup is performed manually or automatically, it is vital that the integrity of the backup be checked periodically. Don't wait until a real disaster forces you to restore the data to find out that your tape drive hasn't been writing anything to the tapes. Do a test restore when you *don't* need it.

Finally, you should consider making multiple backups of data that is especially critical. Store at least one of these copies offsite. It does no good to have a current set of backup tapes in the desk drawer by the server if a fire, flood, or tornado destroys everything in the room. Offsite backups can be taken home with a trusted employee, stored in a bank safe deposit box, transferred daily to a branch office, or uploaded to a remote server over the phone lines or across the Internet.

Disk Fault Tolerance

Another way to protect your data from hard disk failure is to implement *disk fault tolerance*. (Fault tolerance refers to the capability of a system to recover after a failure.) Fault tolerance involves combining multiple physical hard disks into a *fault-tolerant set*, which can take on one of several configurations.

Disk fault tolerance is also called *redundant array of independent (or inexpensive) disks (RAID)*. Commonly used fault-tolerant configurations include the following:

- **Disk mirroring (RAID level 1)**—Disk mirroring requires two physical hard disks, preferably of the same size. All the data on one disk is mirrored to the second. There is an exact duplicate copy of all files and structures on the second disk. If one disk fails, the other can take over. This can be automatic, or you might have to "tell" the operating system where to find the new disk. For example, in Windows NT, this is done by editing the boot.ini file.

- **Disk duplexing (RAID level 1)**—Disk duplexing works exactly like mirroring, except that the two physical disks are attached to separate disk controllers. This adds a layer of fault tolerance, because if a controller fails, the other disk can still function, because it is on a different controller.
- **Disk striping with a parity drive (RAID level 3)**—Disk striping with a parity drive involves writing the data in stripes across multiple drives and writing parity data to another drive that is reserved for that purpose. This requires a minimum of three physical disks (two across which the data is striped and one for the parity information). If a data disk fails, the data can be regenerated using the parity information.
- **Disk striping with parity stripes (RAID level 5)**—Disk striping with parity stripes works in a manner similar to RAID 3, except that the parity (like the data) is written in stripes alternating across the disks. No one disk is designated for parity. Again, at least three disks are required, and if any one disk fails, the data can be regenerated.

NOTE

What is parity? The Microsoft Press Computer Dictionary defines it as “an error-checking procedure in which the number of 1s must always be the same—either even or odd—for each group of bits.” In the context of disk fault tolerance, you can reconstruct the missing data on the failed drive by combining the parity information with the data on the drives that are still functioning.

Other, less commonly used fault-tolerant RAID levels include the following:

- **RAID 2**—This is similar to RAID 3 in that data is striped across multiple drives and one drive is reserved for parity information. The only difference is that the data is striped in bits; in RAID 3, it is striped in bytes.
- **RAID 4**—This also is similar to RAID 2 and 3, except that the data is striped in blocks instead of in bits or bytes.

RAID 0 is used quite often to increase the performance of reading and writing data, but it is *not* a fault-tolerant method. RAID 0 stripes data across multiple disks in blocks, but it does not include parity information; thus, if one of the disks fails, there is no means of recovery.


RAID can be implemented as either a software solution or a hardware solution. Although hardware-based RAID is faster and generally more reliable, it is more expensive. Some server operating systems, such as Windows NT, and Windows 2000 Server products have built-in support for software RAID.

Clustering Technologies

The foundation of all disaster protection methods is *redundancy*. Power backup involves having redundant sources of electricity: the wall outlet and the UPS battery. Data backup involves creating redundant copies of important files on backup media. Disk fault tolerance involves writing the data to or across redundant disks. Clustering is the ultimate form of redundancy.

Clustering generally means grouping servers in *clusters*. With clusters, a group of servers is seen as one server on the network. If one of the servers goes down, another in the cluster takes over its duties. This transition is transparent to users. See Figure 14-9.

Figure 14-9 Server clustering is an expensive, but highly effective, way of providing fault tolerance.



Clustering support is built into operating systems such as Windows 2000 Advanced Server, and clustering software is available to configure servers that are running other operating systems. Clustering provides fault tolerance, along with other advantages such as load balancing.

Summary

No single security measure will protect your network. Effective security is accomplished through the use of a security plan that incorporates multiple security practices, policies, and technologies.

In this chapter, we discussed various ways of protecting the network and the data that resides on and travels across it. First we examined the very broad subject of network security, and then we discussed disaster protection and recovery methods.

You learned how to assess the security needs of your network, and you learned about external security breaches such as unauthorized use of passwords and keys, DoS attacks

(including Ping floods, Smurf attacks, the Ping of Death, and SYN attacks), IP spoofing, computer viruses and worms, and malicious code such as Trojan horse programs.

We then discussed internal threats, which can be attributed to corporate espionage, internal politics, disgruntled employees or ex-employees, and accidental breaches of security.

We provided an overview of cryptography concepts and how text can be “encoded” and “decoded” to thwart unauthorized interceptors. You learned to differentiate between symmetric secret key encryption and asymmetric public/private key encryption. We also discussed authentication, which is the process of determining that a message is really from the purported sender.

You learned about government security ratings and how to find information on the current ratings of popular operating systems, and we went on to discuss security auditing and intruder detection. Next we discussed firewalls and proxies and how you can create a DMZ between your internal network and the Internet (or other external network) to protect the resources on the LAN.

We also discussed methods of using different protocols on the internal and external networks to decrease the risk of intrusion from the outside. We wrapped up the first half of the chapter with a review of the importance of physical security for network components.

The second half of the chapter dealt with disaster protection and recovery, and you learned how to use multiple lines of defense against data loss, including power backup, data backup, disk fault tolerance, and clustering.

Now that you know how to protect your network to some extent from the outside world, you will learn in the next chapter how to get connected to that outside world by using phone lines.

Remote access is growing daily in popularity because more employees are telecommuting, accessing the company network after-hours, or staying connected to the corporate network when they're on the road. Chapter 15, “Remote Access,” teaches you about remote access devices and how to set up remote access clients and servers. We also discuss the special security considerations that come with allowing dial-in connections to your network.

Further Reading

For more information on U.S. laws regulating the export of encryption technologies, see www.bxa.doc.gov/Encryption.

An excellent overview of IPSec is located on Cisco's Web site at www.cisco.com/warp/public/cc/techno/protocol/ipsecur/prodlit/ipsec_ov.htm.

A useful book that covers many aspects of creating a secure network is *Designing Network Security*, by Merike Kaeo. It is published by Cisco Press.

A good resource for information about e-mail security is at www.emailtoday.com/emailtoday/dir/email_security.htm.

Several helpful links to Web sites that address certificate services and other encryption technologies can be found at www.security-online.com/info/certificates.html.

A simplified explanation of Kerberos can be accessed online at www.isi.edu/gost/brian/security/kerberos.html.

For more information about developing network security policies, see www.homex.s-one.net.sg/member/itsecurity/netsec1.htm.

NCSC/NSA Security certifications are available on the Web at www.radium.ncsc.mil/tppep/epl.

Review Questions

The following questions test your knowledge of the material covered in this chapter. Be sure to read each question carefully and select the *best* correct answer or answers.

- 1 What is the name for the number or cipher used by the system to verify the integrity of a communication?
 - a Password
 - b Cryptography
 - c Key
 - d Authentication
- 2 What is the process of trying every possible combination of letters and numbers to “crack” a password?
 - a Decryption
 - b Brute-force attack
 - c Trojan horse attack
 - d Nuke attack
- 3 What is the name for sending a deluge of ICMP messages in an attempt to overwhelm or shut down a server or network?
 - a A SYN attack
 - b A worm
 - c A Ping of Death
 - d A Ping flood

- 4 Which of the following are common internal threats to network security? (Select all that apply.)
- a Corporate espionage
 - b Accidental breaches
 - c IP spoofing
 - d Proxies
- 5 Which of the following is true of a good password policy in a high-security environment?
- a It enables users to create and use the passwords that are easiest for them to remember.
 - b It requires that passwords be a minimum length, such as six characters.
 - c It prohibits using numbers or symbols in the password.
 - d It prohibits users from changing their passwords.
- 6 Which of the following describes the default permissions on a newly created share on a Windows 2000 server?
- a By default, no one has permission to access the share.
 - b By default, everyone has read-only permission to the share.
 - c By default, everyone has full control permission to the share.
 - d By default, members of the administrators group have read-only access to the share.
- 7 Which of the following best describes file encryption?
- a File encryption protects data stored on the hard disk.
 - b File encryption protects data sent across the network.
 - c File encryption protects data both when it is on the disk and when it is sent across the network.
 - d File encryption is available only through third-party software.
- 8 In the context of network security, what is ESP?
- a Extra-Sensitive Packets
 - b Extended Security Protocol
 - c Encapsulating Security Payload
 - d External Security Properties

- 9** Which of the following encryption methods uses the same key to encrypt and decrypt the data?
- a** Asymmetric encryption
 - b** Public key encryption
 - c** Public/private key encryption
 - d** Secret key encryption
- 10** What is the one-way mathematical calculation that creates a message digest, which is used to verify the authenticity of messages?
- a** A digital certificate
 - b** A hash algorithm
 - c** A ticket-granting ticket
 - d** A retinal scan

Internet Corporation for Assigned Names and Numbers. *See* ICANN.

Internet Engineering Task Force. *See* IETF.

Internet Information Server. *See* IIS.

Internet protocol. Any protocol that is part of the TCP/IP protocol stack. *See* IP and TCP/IP.

Internet service provider. *See* ISP.

Internet Society. *See* ISOC.

internet. Short for internetwork. Not to be confused with the Internet. *See* internetwork.

Internet. The largest global internetwork, connecting tens of thousands of networks worldwide and having a culture that focuses on research and standardization based on real-life use. Many leading-edge network technologies come from the Internet community. The Internet evolved in part from ARPAnet. At one time, it was called the DARPA Internet, not to be confused with the general term internet.

Internetwork Operating System. *See* IOS.

Internetwork Packet Exchange. *See* IPX.

internetwork. A collection of networks interconnected by routers and other devices that functions (generally) as a single network.

internetworking. The industry devoted to connecting networks together. The term can refer to products, procedures, and technologies.

InterNIC. An organization that serves the Internet community by supplying user assistance, documentation, training, registration service for Internet domain names, network addresses, and other services. Formerly called NIC.

interoperability. The capability of computing equipment manufactured by different vendors to communicate with one another successfully over a network.

intranet. An internal network that is to be accessed by users who have access to an organization's internal LAN.

IOS (Internetwork Operating System). *See* Cisco IOS software.

JPEG or JPG. A compressible graphic image file format, often used for images embedded in web pages because it is supported by web browsers and because it downloads relatively quickly because of the small file size.

K

Kb (kilobit). Approximately 1000 bits.

KB (kilobyte). Approximately 1000 bytes.

KBps (kilobytes per second). A rate of transfer speed.

Kbps (kilobits per second). A rate of transfer speed.

Kerberos. A security method used for authentication that relies on an encrypted “ticket,” so that the user’s password is not required to be sent across the network. Kerberos is an Internet standard and is the means of authentication used by Windows 2000’s Active Directory.

kernel. The core of the computer’s operating system that provides the basic services for other parts of the operating system. Compare with shell.

kilobit. *See* Kb.

kilobits per second. *See* Kbps.

kilobyte. *See* KB.

kilobytes per second. *See* KBps.

L

LAN (local area network). A high-speed, low-error data network covering a relatively small geographic area (up to a few thousand meters). LANs connect workstations, peripherals, terminals, and other devices in a single building or other geographically limited area. LAN standards specify cabling and signaling at the physical and data link layers of the OSI reference model. Ethernet, FDDI, and Token Ring are widely used LAN technologies. Compare with MAN and WAN. *See* VLAN.

virtual private network. *See* VPN.

VINES (Virtual Integrated Network Service). A NOS developed and marketed by Banyan Systems.

virus. A program that propagates itself, and in some cases, it is intended to cause damage to other programs or files.

VLAN (virtual LAN). A group of devices on a LAN that are configured (using management software) so that they can communicate as if they were attached to the same wire, when in fact they are located on several different LAN segments. Because VLANs are based on logical instead of physical connections, they are extremely flexible.

Voice over IP. *See* VoIP.

VoIP (Voice over IP). A method for sending voice over a LAN, a WAN, or the Internet using TCP/IP packets. VoIP is also called Voice over Internet Protocol.

volume. A means of organizing the space on a disk.

VPN (virtual private network). A connection between a remote client and a private network, such as a company LAN, that uses the public Internet as a conduit. A VPN is established using tunneling and encryption protocols to enable data to travel across the public network to the private network in a secure manner.

W

WAN (wide area network). A data communications network that serves users across a broad geographic area and often uses transmission devices provided by common carriers. Frame Relay, SMDS, and X.25 are examples of WAN technologies. Compare with LAN and MAN.

WAN link. A WAN communications channel consisting of a circuit or transmission path and all related equipment between a sender and a receiver.

waveform signaling. A signaling method in which the signal follows a continuously variable wave pattern, which an infinite number of values along the wave. Compare with discrete state signaling.

web server. A server on which web pages are stored, which can be accessed through software called a web browser using the HTTP protocol.



Computer Networking Essentials

- Excellent preparation for generic or vendor-specific networking certification programs
- Written in a user-friendly manner accessible to readers who may need an introduction to networking terms
- Introduction to two popular real-world networking models: the Department of Defense (DoD) model and the Open System Interconnection (OSI) model
- Includes latest technologies such as laser, infrared, and satellite/microwave communications
- Coverage of common server operating systems, including Windows NT®, Windows® 2000, NetWare, UNIX, and Linux
- Security coverage provides basic cryptography concepts, public and private key encryption, firewalls and proxies, and internal security measures

This book is part of the Cisco Press Networking Technologies Series, which offers networking professionals valuable information for constructing efficient networks, understanding new technologies, and building successful careers.

Category: Networking
Covers: Networking Fundamentals

Computer Networking Essentials is a comprehensive introduction to the world of networking. This book carefully integrates hardware and software topics, incorporating a variety of popular vendor-prescribed introductory networking knowledge into an all-in-one guidebook. *Computer Networking Essentials* combines essential background materials such as the history of networking, networking theory, and client/server operating systems that run on networked PCs with practical information on networking terminology, established standards, and implementation of LANs and WANs. Special emphasis is placed on understanding network protocols and how they operate at all layers of the networking model. Care is also taken to explain the interoperability of networks that run on multiple protocols, platforms, and operating systems.

Several specialty areas of networking are explored, including security, remote access, virtual private networking (VPN), and network monitoring and troubleshooting. You also learn about the future of networking through a chapter dedicated to emerging technologies.

Computer Networking Essentials is written for entry-level network technicians and students in formal academic settings, training courses, or self-study environments. Presented at a level that requires little computer experience, *Computer Networking Essentials* is an excellent preparation for CompTIA's Network+ program, Novell's Networking Fundamentals, Cisco's CCNA® certification, and similar programs.

Debra Littlejohn Shinder, MCSE, MCT, is an instructor and trainer at Eastfield College in Dallas County, Texas. She is also a technology consultant for TACteam: Training and Consulting, which she co-owns with her husband, Tom Shinder. Debra authored *Troubleshooting Windows 2000 TCP/IP* and has contributed to more than 10 Windows books.

Technical Reviewer Dr. Thomas W. Shinder has co-authored several books with his wife, Debra. He teaches at Eastfield College in Mesquite, Texas, and he has provided training and consulting services for major Dallas area firms. Dr. Shinder attended the University of California at Berkeley and the University of Illinois Medical School, and practiced neurology for several years before making a career change to work with computers.

\$49.95 USA / \$74.95 CAN

ciscopress.com

ISBN 1-58713-038-6



X0019GNDZJ

Computer Networking Essent... (Cisco Press Core Series)
New