Computing McGraw-Hill

LAN to WAN Interconnection

FOR THOSE WHO ARE FAMILIAR WITH ONE GEOGRAPHY AND PERPLEXED BY THE OTHER

John Enck
Mel Beckman

LAN to WAN interconnection

John Enck
Mel Beckman

McGraw-Hill
New York San Francisco Washington, D.C. Auckland Bogotá
Caracas Lisbon London Madrid Mexico City Milan
Montreal New Delhi San Juan Singapore
Sydney Tokyo Toronto

hc 4 5 6 7 8 9 10 DOC/DOC 9 9 8 7 6

Library of Congress Cataloging-in-Publication Data

Enck, John, 1956-
LAN to WAN interconnection / by John Enck, Mel Beckman.
p. cm.
Includes index.
ISBN 0-07-019614-1
1. Local area networks (Computer networks) 2. Wide area networks (Computer networks) I. Beckman, Mel, 1955- . II. Title.
TK5105.7.E53 1995
004.6'8—dc20 95-10086
CIP

Acquisitions editor: Marjorie Spencer
Editorial team: Marc Damashek, Book Editor
David M. McCandless, Managing Editor
Joanne Slike, Executive Editor
Production team: Katherine G. Brown, Director
Rhonda E. Baker, Coding
Brenda M. Plasterer, Coding
Wanda S. Ditch, Desktop Operator
Nancy K. Mickley, Proofreading
Brenda S. Wilhide, Computer Artist
Joann Woy, Indexer
Design team: Jaclyn J. Boone, Designer
Katherine Stefanski, Associate Designer

0196141
EL3

In memory of my father, Herbert,
my mother, Birdine,
and Taz, my devil dog.
John Enck

To my wife, Patricia
Mel Beckman

Acknowledgments

The authors would like to thank Duke Communications for their encouragement and support over the years; Norm Bartek for his careful and skillful review of our first draft; Marjorie Spencer, for her patience and guidance during the development of this book; and McGraw-Hill for the chance and reality of bringing this book to market.

Preface

One of the most important issues we discussed before we began writing this book was how technical the material should get. For example, did we want to discuss the properties of electrons moving through copper media? Did we want to explain the principles that allow light waves to travel through bends in fiber optic cable? Did we want to include lengthy dissertations on the differences between Manchester encoding and differential Manchester encoding?

After some brief discussions along these lines, we finally agreed that although these principles are certainly important, they aren't particularly relevant when it comes to the day-to-day reality of interconnecting networks. Having made that decision, we then agreed to avoid the use of academic terminology in favor of common terms and terminology. Our goal was simple: we wanted to create a no-nonsense book that presents the information you need to know to plan, implement, and manage Local Area Network (LAN) connections, Wide Area Network (WAN) connections, and, of course, LAN-to-WAN interconnections.

How successful were we in reaching our goal? That decision is yours to make, and we encourage you to do so.

John Enck
Mel Beckman

Introduction

This book is divided into five parts. Each part begins with an introduction that provides an overview of the technology discussed in the related chapters. The chapters following that introduction then provide detailed information on technology. The specific organization of the parts and chapters is as follows:

Part I: General overview. A brief discussion of the basic LAN/WAN problem.

Chapter 1: When LANs meet WANs. The purpose and application of Local Area Network (LAN), Metropolitan Area Network (MAN), and Wide Area Network (WAN) technology.

Part II: LAN basics. Information critical to the planning, installation, and expansion of Ethernet/IEEE 802.3 and Token Ring/IEEE 802.5 LANs.

Chapter 2: Ethernet/802.3. Detailed information on Ethernet and IEEE 802.3 LANs, including discussions on topology, construction, and the common Carrier Sense, Multiple Access with Collision Detection (CSMA/CD) access discipline. Also discusses high-speed Ethernet technology.

Chapter 3: Token Ring/802.5. Detailed information on IBM Token Ring and IEEE 802.5 LANs, including discussions on topology, construction, and the common token-passing access discipline. Also discusses the pros and cons of source routing.

Part III: Interconnection tools. Examination of the fundamental tools that interconnect multiple LANs in a local or wide area environment.

Chapter 4: Bridges & routers. An examination of bridging and routing technology and applications. Also includes definitions of the major bridging and routing protocols and techniques, e.g., the Routing Information Protocol (RIP), the Open Shortest Path First (OSPF) protocol, and Data Link Switching (DLSw).

Chapter 5: Gateways & hubs. An examination of gateway and hub functions and applications. Also includes discussions on protocol encapsulation, how gateways operate in Transmission Control Protocol/Internet Protocol (TCP/IP) networks, and "smart" versus "dumb" hubs.

Part IV: Specialized high-speed interconnections. Explorations of high-speed interconnections that can be used to implement LANs, MANs, or interconnect LANs/MANs to create a WAN.

Chapter 6: Fiber links & FDDI. An exploration of the networking applications for fiber optic links. Includes discussion of how fiber optic links operate, and of the Fiber Distributed Data Interface (FDDI).

Chapter 7: Frame & cell relay. An exploration of frame relay technology, cell relay technology, and the Asynchronous Transfer Mode (ATM) architecture.

Part V: Traditional WAN interconnections. Information on conventional wide area connections commonly used for low-speed network connections (less than 100 Mbps).

Chapter 8: Point-to-point links. A look at switched and leased analog phone links, digital phone links, and T1, fractional T1, and T3 services.

Chapter 9: Multi-point links. A look at X.25 packet switching networks, Integrated Services Digital Network (ISDN), and Switched Multimegabit Digital Services (SMDS).

Contents

3 Token Ring/802.5 51

Part III: Interconnection tools

4 Bridges & routers 77

5 Gateways & hubs 101

Part IV: High-speed interconnections

6 Fiber links & FDDI 127

1

General overview

IF YOU THINK ABOUT IT, NETWORKING TECHNOLOGY HAS become so intertwined with computer technology that it is often difficult to separate the two. For example, users stationed at terminals depend on the link between their terminals and the main computer just as much as they depend on the computer itself. Similarly, PC users who access files and printers on a common server rely on the network that interconnects them with the server. And in many large organizations, terminals, host computers, PCs, and PC servers rely on network connections to perform day-to-day business transactions—even though these devices may be scattered across the face of the globe.

Like computers, network technology comes in all shapes and sizes. Simple network connections accommodate the physical interface between a PC and a modem, or a PC and a printer. More intricate connections handle the attachment of multiple terminals to a common host, or facilitate the interconnection of PCs and PC servers. As you progress farther up the scale of network technology, you run into complex connections that function as the building blocks for large (possibly world-wide) networks.

All of these types of connections—from the lowly PC serial cable up to world-spanning fiber links—are important components in the grand scheme of networking. Unfortunately, they cover so much technical ground that it is difficult to get your arms around networking as a whole. Consider, for instance, how vastly different the electrical current technology used to carry information in most Local Area Networks (LANs) is from the light-wave technology used to carry information over fiber optic links.

Even the application of networking technology varies from one situation to another. The networking solution that works well for interconnecting PCs in a single office building may fall flat on its face when the network expands into multiple buildings located miles apart. The wide area strategy linking remote terminals back to a central computer may crumble when the terminals are replaced by PC LANs. Any network solution appropriate for a specific computing environment is subject to change if that environment changes.

As complicated as networking is, it is not rocket science. You do not need a degree in electrical engineering or quantum physics to understand and appreciate its subtleties. If you have a good working understanding of one aspect of networking, you'll find it relatively easy to learn other aspects. For example, if you've been working with LANs for several years, you'll find that Wide Area Network (WAN) technology has many familiar elements. Or if you've been working with wide area links for years, you'll find similarities in how local area links operate.

And what if you haven't been working with either LAN or WAN technology? Don't worry—we make no assumptions about your skill level. In fact, we cover a broad spectrum of networking technology that includes LANs, WANs, and the overlap between the two. We hope that this approach will save wear and tear on two of your most highly treasured networking tools: your time and patience.

When LANs meet WANs

1

NETWORKING HAS BEEN A KEY ELEMENT IN DATA PROCESSing since it moved from a batch environment to an interactive environment in the early 1960s. With that transition, the computer input/output environment changed from card readers, paper tape readers, and line printers to remote terminals and dot matrix printers. Programmers, operators, and eventually users learned to interact directly with the computer—punched cards and paper tape strips were no longer required.

The creation of an interactive user environment was quickly followed by the introduction of computer-to-computer links to allow computers to exchange information with one another. At first these links required the computers to be physically close to one other, but over time the technology matured to the point where computers could communicate with each other regardless of distance. Once the interconnection technology was established, the focus turned to speed—how to move information over these computer-to-computer links faster and more efficiently.

Taken together, these two developments—interactive users and computer-to-computer links—formed the foundation for today's towering empire of Local Area Network (LAN) and Wide Area Network (WAN) technology.

Despite this common foundation, LAN and WAN technology evolved separately. WAN technology was developed to address the requirements of a centralized computer network, while LAN technology was advanced to handle the issues associated with a distributed computer network. This distinction can best be seen by looking at the network approaches used by two major computer companies—IBM's centralized network architecture and Digital Equipment's distributed network architecture.

The need for wide area networking

Under IBM's vision for computing in the 1960s and 1970s, a large central computer (i.e., a mainframe) provided processing and mass storage services to handle both interactive and batch applications. This central computer was designed to provide application services to hundreds (and eventually thousands) of remote terminals and printers. Given this scope of operation, it became clear to IBM that a single central computer could not efficiently handle both the application load and the communications load.

IBM's solution to keeping communications overhead off the central computer was to develop a hierarchical network structure that distributed the communications functions over a number of devices, as shown in Fig. 1-1.

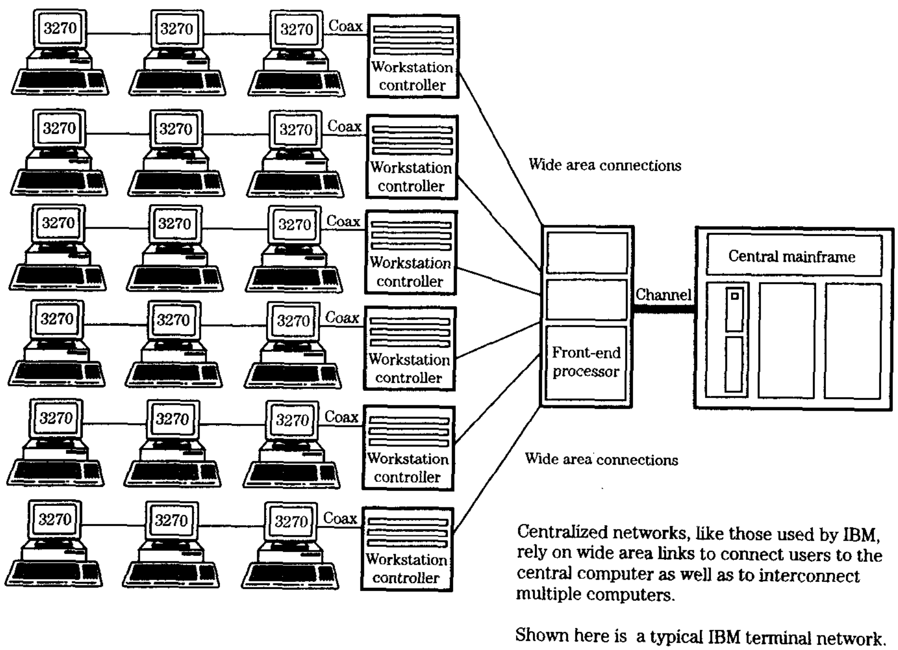

■ **1-1** *Centralized network implementation*

Under this design:

- ☐ User terminals and printers attach to workstation controllers using coaxial connections. A single workstation controller handles a number of terminals and printers.
- ☐ Each workstation controller is responsible for delivering messages sent from the central computer to its attached terminals and printers, and for collecting information entered

into terminals and transmitting it to the central computer. Workstation controllers can interface with either the central computer, or more likely, with a communications device called a "front-end processor" (FEP).

- ☐ The purpose of a FEP is to further isolate the central computer from the communications network. The FEP manages the traffic to and from workstation controllers and then funnels it into the central computer over a specialized high speed connection called a "channel" attachment. This is the same kind of attachment used to connect disk drives and tape drives to the central computer.

As you can see, IBM's hierarchical design frees the central computer from handling the communications environment, because communications functions are distributed across the FEP and workstation controllers. IBM continued to refine its network design, adding capabilities like FEP-to-FEP communications so multiple central computers could communicate with one another, and could share terminal and printer resources as well. In 1974, IBM turned their network design into a formal architecture, called the Systems Network Architecture (SNA).

IBM's centralized network design depends on wide area links to work effectively. For example, if a group of users located in Boston needs to communicate with a central computer located in New York, a wide area link is required to connect a workstation controller in Boston with the FEP in New York. Or if a central computer in Los Angeles needs to exchange information with a sister company's central computer located in Dallas, then a wide area link needs to be put in place between the two FEPs.

In short, when you introduce significant geographical distances into a centralized network, that network's dependency on wide area technology increases.

Now IBM was not the only company with a network design dependent on wide area technology—many other major computer manufacturers followed IBM's lead and introduced centralized network designs. By the time the 1980s rolled around, the majority of the computer industry was clamoring for wide area technology to improve the speed and efficiency of their centralized network designs.

Wide area network evolution

The first major milestone for wide area technology was the creation of a modulation–demodulation device, better known as a "modem." A modem translates digital information sent from a computer into an analog format suitable for transmission over a standard telephone line. A second modem on the other end of the line then translates the analog information back into digital information for the other computer. Because telephone lines are available over a wide geographical area, they were (and still are) perceived as an ideal means of transmitting terminal, printer, and other computer information over long distances.

In the beginning, modems supported very modest transmission speeds—for example, 300 bits per second (bps). Over time, modem technology improved and transmission speeds gradually increased. Thus the transmission rate of 300 bps was superseded by 1200 bps, then by 2400 bps, 4800 bps, 9800 bps, and so on, until today's current speed of 28.8 kbps (28,800 bps).

The way that modems interacted with the phone lines also changed over time. In the beginning, modems could work over standard voice phone lines or dedicated leased lines. When standard voice lines were originally used, someone (some person) had to manually dial up the other modem to establish the connection. This manual intervention was viewed as less than ideal, so the Automatic Calling Unit (ACU) was developed to handle the dialing automatically. The ACU technology was then surpassed by the Hayes modem command set, which remains the de facto standard for modems today.

Leased phone lines avoid the issue of dialing because each end of the line is permanently assigned to a modem. Leased lines, however, brought their own set of technology issues. For example, leased lines can be configured for two wires or four wires, and can even be configured so that a single line has multiple connections, allowing a pool of modems to share the line. The dedicated nature of leased lines is very attractive for centralized network implementations—wide area links can be available 24 hours a day, 7 days a week, with no chance of accidental hang-ups.

Unfortunately, the analog technology used in voice telephone lines has its limitations. The most important limitation is that because the computer information is sent in an analog format, any "noise" or disturbances on the line have a significant negative impact on the transmissions. So the long distance phone carriers created and offered another option for computer networking—digital links.

Unlike analog phone links, digital links do not need to have the computer's digital information translated into an analog format—instead, the digital information is sent across the wide area link in