LANs to WANs

Network Management in the 1990s

Robert P. Davidson

Nathan J. Muller

LANs to WANs: Network Management in the 1990s

Nathan J. Muller
Robert P. Davidson

Illustrated by Richard S. Stankevich

Artech House
Boston • London

Library of Congress Cataloging-in-Publication Data

Muller, Nathan J.
LANs to WANs : network management in the 1990s / Nathan J. Muller and Robert P. Davidson.
p. cm.
Includes bibliographical references and index.
ISBN 0-89006-410-5
1. Computer networks--Management. 2. Local area networks (Computer networks) 3. Integrated services digital networks.
I. Davidson, Robert P. II. Title.
TK5105.5.M86 1990 90-41621
004.6--dc20 CIP

British Library Cataloguing in Publication Data

Muller, Nathan J.
LANs to WANs : network management in the 1990s.
1. Computer systems. Networks
I. Title II. Davidson, Robert P.
004.6

ISBN 0-89006-410-5

685 Canton Street
Norwood, MA 02062

International Standard Book Number: 0-89006-410-5
Library of Congress Catalog Card Number: 90-41621

10 9 8 7 6 5 4 3 2 1

LANs to WANs:
Network Management in the 1990s

For a complete list of the *Artech House Telecommunication Library,* turn to the back of this book . . .

Preface

Throughout this book, the terms "global," "ubiquitous," and "intelligent" are used to describe communication networks of the future. Such terms convey an image of communication networks as they will be in the year 2000 and beyond. However, not only are technological trends transforming communication, but political and economic trends are as well.

On a macro level, as the world's population increases, resources are depleted, the biosphere deteriorates, and we lose our ability to maneuver in a compressed world. There is less tolerance for error; even a small mistake, a failure to respond fast enough, or a misunderstanding takes on awesome significance. Any resulting catastrophe, whether natural or of human origin, becomes the concern of the many instead of the few, if only because their frequency and intensity will make it so. Satellite technology has made us aware of how fragile our biosphere really is and of the need to preserve it. Communication networks that are global, ubiquitous, and intelligent certainly will not solve our environmental problems for us, but these networks can keep the "window of opportunity" open long enough for us to respond appropriately and quickly to prevent, contain, or remedy potentially disastrous results.

On a more mundane level, the micro level, communication networks have amply demonstrated their value to businesses, governments, and society at large. The effect of radio, television, and the telephone have made us acutely aware that we belong to a global community, with individual members sharing certain basic needs and aspirations. In fact, the availability and speed of information transfer has intertwined the economies of nations to such an extent that when stock exchanges in London and Tokyo cough, Wall Street sneezes!

If corporate America learned one lesson in the 1980s, it was that communication networks can result in competitive advantages. By drawing on the previously described macro view, we can say that businesses have been able to use their extensive communication networks to track emerging markets, respond to changes in customer demand, penetrate new markets, and support expansion op-

portunities at home and abroad. For a growing number of companies, communication networks have become the infrastructure upon which the organizational superstructure rests. Momentarily remove that block, and business comes to a halt; remove that block for a long enough time, and the entire structure collapses.

In today's aggressive business world, data communication is a strategic weapon. To know the equipment that forms the network is no longer enough; cost containment and control are also necessary for success. The ability to control costs begins with the selection of appropriate transmission services and equipment that will deliver the desired levels of performance and efficiency. Cost control does not stop with the acquisition of equipment and services; it is a continuing process that encompasses installation, systems integration, maintenance, and the ongoing management of network elements. The choice of network, whether it is a local area network (LAN) or a wide area network (WAN), is predicated on function as well as associated costs. The distinctions that once separated the public network from private networks are blurring, forcing corporations to consider the benefits of hybrid arrangements that combine the best attributes of both.

This book is intended to be a comprehensive reference for communication managers, planners, and consultants. It explores the networking landscape, present and future. As such, we describe what networks look like today, what they may look like tomorrow, and how we go from here to there. Although the focus is on technology, we weave into the discussions such often overlooked elements as integration, design, and management, and we also address operational and regulatory issues. Where appropriate, we try to provide a historical perspective to these discussions, which may be of value to those who are new to the communication field.

Although comprehensive in scope, this book does not provide an exhaustive treatment of the various topics. We do not attempt to replace the many fine specialized texts that have become available in recent years. From such a comprehensive treatment, however, the reader can obtain an appreciation of the relationships that exist among various technologies, and how the deployment of these technologies affects organizational performance and translates into competitive advantage.

We recognize that many communication managers are tired of hearing about how corporate networks can be used to secure competitive advantages. A typical attitude seems to be reflected in the oft-repeated comment: "I'm too busy keeping the network up and running to care about such things." Of course, the job of most communication managers is more tactical than strategic, but to ignore totally how the corporate network can be positioned to take maximum advantage of new technologies and emerging services is short-sighted. For example, building or expanding a network without regard for industry standards can result in the early obsolescence of technology acquisitions, wasted capital investments, and overhead costs that are higher than necessary. Because capital investments and overhead

costs must be recovered over time through the prices charged for products or services, those firms with higher costs will have a difficult time competing against those with lower costs. If we assume all else to be equal, the company that can more quickly recoup its capital investment and overhead costs will have more pricing flexibility, and therefore more marketplace leverage. Surely, these matters must be on the minds of communication managers, even as they do such mundane and time-consuming tasks as interpreting alarms, initiating restoral procedures, implementing diagnostic routines, and tracking equipment moves and changes.

Part I of this book discusses internal corporate networking, as exemplified by the LAN. With the growth of distributed computing has come a corresponding growth in LANs, which have become the preferred way to interconnect computers and peripherals. Although an installed base of tens of millions of Ethernet and token-ring nodes exists, market growth, product enhancement, and innovation continue at a frenetic pace. Chapter 1 discusses LAN topologies, architectures, access methods, and standards.

We continue our discussion of LANs in Chapter 2 on operating systems, which focuses on the client-server model of distributed computing. Operating systems are composed of many modules, some of which reside in a microcomputer that acts as the server, while other modules reside in terminals, printers, and other network resources, which constitute the "clients." The models work together to provide the functions of the network, such as recognizing users, associating their identities with access privileges, and routing their requests to the appropriate server for action. Although invisible to the user, the operating system must be viewed within the context of LAN functionality because it constitutes a part of the communication among various LAN elements.

Another part of the communication between various LAN elements is the access protocol. To ensure the success of LANs, a critical factor was that the cost of the interface used to connect equipment be driven down until it was much less than that of the equipment alone. The IEEE 802 LAN standards were developed for this purpose. They may be described in terms of the seven-layer Open Systems Interconnection (OSI) reference model developed by the International Organization for Standardization (ISO). Each layer of the model has its own set of protocols. Chapter 3 describes these protocols and how they are used to allow computer systems to communicate with each other if they "speak" the same set of computer protocols at each layer.

Among the alternatives to private LANs are telephone company central office-based local area networks (CO-LANs), which are discussed in Chapter 4. While private customer on-premises LANs operate at speeds of 1 Mb/s and above, a CO-LAN has a top asynchronous speed of 19.2 kb/s and a top synchronous speed of 64 kb/s. Despite the speed limitations of CO-LANs, they may be appropriate for CENTREX users with limited data-switching needs. Instead of investing heavily in on-premises hardware and cabling, the customer premises portion of a CO-LAN

takes advantage of existing twisted-pair wiring to connect virtually all types of computers and terminals. This approach makes installation, expansion, and maintenance fairly easy.

In Part II we explore some of the technologies that are more suitable for wide area networking, beginning with high-speed diagnostic modems. Since the 1960s, analog modems have evolved from simple devices that provided point-to-point communication to the basic building blocks used to create and manage complex multipoint networks involving multiplexers, packet switches, fiber optic transmission facilities, and satellites. Despite progress toward an "all-digital" network, analog modems are still required to pass data over local access lines, most of which will not be converted to digital for many years. At high speeds, however, even slight differences in modem design can profoundly affect performance. Chapter 5 delves into such design elements as the multiplexing techniques used to derive the diagnostic channel, the quality of the filters that separate diagnostic from production data, coding-decoding schemes, and the level of circuit integration. Also discussed is the range of diagnostic features available with these modems, which can greatly facilitate network management. Over the years, the private branch exchange (PBX) has emerged to become the undisputed cornerstone of today's corporate voice network. As the networking requirements of organizations continue to change, however, new applications are being sought for the PBX, many of which may go unfulfilled. Chapter 6 explores PBX design and its affect on the integration and management of voice and data. The chapter concludes with an examination of the PBX's role in the integrated services digital network (ISDN) environment.

Chapter 7 discusses multiplexers. The economic benefits of multiplexing are readily appreciated by most network managers. In optimizing the network through multiplexing, the number of local voice-grade leased lines can be reduced, while high-capacity digital facilities may be utilized to full advantage. The resulting efficiencies can greatly reduce operating costs. When equipped with the requisite management functionality, as well as support for fractional T1 and the ISDN primary rate interface (PRI), today's T1 multiplexers provide users with unprecedented flexibility to design and maintain fault-tolerant networks.

Just as analog private line data transmission requires modems, digital data services and high-capacity T1 service require channel service units (CSUs) and data service units (DSUs) to interface data terminal equipment with the network. Chapter 8 discusses the regulatory issues, functionality, and diagnostic capabilities associated with these devices, as well as their complementary relationship to the network.

The proliferation of T1 links has created the need to simplify their administration and control, facilitate the ability to access and test individual circuits, switch channels from one link to another, and reconfigure circuits to achieve optimal efficiency and cost savings. Digital cross-connect systems (DCSs), the subject

of Chapter 9, perform these tasks and more. The DCS is associated with customer premises management systems that enable users to control their own circuit configurations. This chapter also contrasts AT&T's customer-controlled reconfiguration (CCR) service with its Bandwidth Management Service (BMS).

Very small aperture terminals (VSATs) have become a reliable transmission technology for today's private networks. VSAT systems integrate transmission and switching functions to provide preassigned links and on-demand assignment for point-to-point and broadcast networks. Advances in technology have combined to bring about smaller, more powerful and economical satellite dishes that in many cases offer considerable efficiency, convenience, and flexibility over terrestrial private leased lines, including fiber optic transmission systems. Much of this contest is related to cost. Chapter 10 compares the various types of VSATs, as well as their components. Also discussed are the various applications of VSAT technology, configuration possibilities, protocols used, and routing and management functions.

The demise of packet switching has been widely predicted since the 1980s, yet it remains surprisingly healthy. While the latest developments in T1 networking and its derivatives, fractional T1 and switched 56 kb/s service, dominated trade press coverage of the industry throughout 1989, the packet-switching domain has quietly expanded. Chapter 11 describes packet transmission technologies and how they are being developed and deployed, particularly "fast packet" and "frame relay," which form an integral part of broadband ISDN. Despite the pervasiveness of the X.25 standard, some confusion remains about Recommendation X.25 and ancillary standards. The chapter delves into these and other issues.

Instead of locking themselves into a single network-wide architecture, users may now select from among several architectural building blocks, fitting them to the needs of particular segments of the network. Chapter 12 deals with the "nuts and bolts" of assembling a hybrid network consisting of statistical and time division multiplexers, as well as X.25. This hybrid approach provides greater flexibility in configuring the network for expansion, and more effectively positions the network to accommodate new communication technologies such as ISDN, which is the ultimate in hybrid networking, offering a blend of circuit switching and packet switching. Also included in this chapter are several case studies.

With individual LAN islands continuing to proliferate among multiple corporate locations, the need to interconnect them becomes increasingly apparent. Often, the connectivity vehicle is the WAN. There are several types of devices available to meet the varying interconnection needs of LAN users: bridges, routers, and gateways. Part III on internetworking contains separate chapters for each type of device, exploring protocols, applications, and performance issues.

Part IV, on the network support infrastructure, starts with Chapter 16 on network design. The increasing frequency of changes in telecommunication rates and services virtually requires that network design be a continuous process. The power of computers can be harnessed to help managers find the proper blend of

facilities and equipment, and the most appropriate topology to support organizational objectives efficiently and economically. However, network design requires more than having the right tools and tariff databases; it requires this very complex task to be approached in a systematic way. This chapter focuses on some of the considerations that apply to the network design process.

Today's networks consist of myriad intelligent elements. Not only must diverse elements be seamlessly interconnected on the same network, users must be provided with transparent access to every other element on it. The selection, installation, and maintenance of this equipment requires a broad range of expertise that typically is not found within a single organization, and therefore many companies must rely on systems integrators. Chapter 17 discusses the role of systems integrators in the network support infrastructure, as well as the strengths and weaknesses of the various types of firms specializing in such services.

The plethora of equipment types that must be interconnected on the same network to achieve the organization's business objectives makes the task of management difficult because each type of device uses unique control protocols and management routines. The key issue is how to manage these diverse elements with the objective of providing full network connectivity, correct functionality, and full flexibility to the network's users. Chapter 18, on network management, sorts out the different types of proprietary management systems, discusses the capabilities of each, and illustrates how they are being unified by such vendors as AT&T, DEC, and IBM. We will also discuss the OSI reference model, and how it may encourage the development of expert systems for network management.

Computer and communication systems have increased in functionality and complexity, in keeping with the demands of users for more processing power, speed, and connectivity. This has catapulted maintenance and support services to a key status in the networking infrastructure. Whether problems are revealed through alarms, diagnostics, or users experiencing trouble with their equipment, the need for timely and qualified maintenance and support services is critical. Chapter 19, on maintenance and support services, describes the roles of hardware vendors, third-party maintenance firms, and users in providing such services, as well as the ramifications of renewed competition in the service industry. Also discussed are some of the service options available and the key issues related to service provision.

T1 private networks, stimulated by post-divestiture (of AT&T) tariff relationships, have grown enormously, resulting in cost-effective private "dedicated networks." However, T1 users continue to seek improved network survivability because, as more companies are realizing, the ability to remain competitive rests on the quality of their networks. Chapter 20 on network restoral compares various restoral techniques used on T1 networks, and explores the relationship between private T1 and carrier-based rerouting alternatives such as AT&T's DACS/CCR and evolving ISDN PRI services. Also included is a discussion of restoral options offered by hardware vendors, especially those available through T1 multiplexers.

One of the most overlooked elements of the network support infrastructure is regulatory participation by MIS and telecommunication professionals. Chapter 21 on regulatory involvement explains what is at stake when users fail to make their views known to state and federal regulatory agencies, either individually or through involvement with trade associations. In addition to a discussion of how specific regulatory actions affect the cost and characteristics of networks, this chapter provides advice on how to influence regulatory decisions in their formative stages.

Part V, on the future of networking, starts with a discussion of the latest innovation in T1 technology: fractional T1. It entails the provision and use of incremental bandwidth between 56 kb/s and 1.544 Mb/s without the need to lease an entire T1 facility. Since its introduction in 1988, fractional T1 has been a topic of intense interest among those who operate corporate networks. After all, being able to order—and pay for—only enough bandwidth to support existing applications is a novel concept. Chapter 22 discusses the ramifications of fractional T1, providing extensive cost-benefit analyses and a discussion of hardware requirements.

Another widely publicized topic is ISDN, the subject of Chapter 23. An important milestone in the continuing evolution of telecommunication networks worldwide, the long-term objective of ISDN is to implement a common standard for the provision of carrier services in such a way as to minimize operating costs, while maximizing the performance and reliability of the network. ISDN is also intended as a network platform from which new information services can be launched in a more efficient and expeditious manner. This chapter provides a comprehensive explanation of ISDN, separating the myths from the realities as well.

For those with requirements for both fractional T1 and ISDN, technological developments currently under way make feasible the combination of fractional T1 for dedicated circuits with primary rate ISDN for switched circuits over the same T1 facility, thereby implementing "fractional ISDN." Chapter 24 provides an in-depth look at this emerging service, illustrating its benefits for both large and small users, including lower access costs and bandwidth partitioning for enhanced management and control.

Open network architecture (ONA) refers to the overall design of a carrier's basic network facilities that allows traditional carriers to offer enhanced services without the requirement of establishing a structurally separate subsidiary. At the same time, ONA allows independent operators to offer new services over the public telephone network. As described in Chapter 25, this will result in lower costs for users as well as virtually instantaneous "cutover" of new services from a variety of providers.

As described in Chapter 26, ONA is an important step toward the realization of another emerging concept called the "intelligent network." The intelligent network is industry jargon intended to convey an image of how the public network

may look and function in the future. It provides the means with which carriers can create, and uniformly introduce and support, myriad services and features via a common architectural platform that uses ISDN and Signaling System Number 7 (SS7). This intelligence will be derived from sophisticated software embedded at strategic locations within the public network. By accessing this intelligence, with or without ISDN, users will be able to engineer their own services and to customize features without telephone company involvement. The intelligent network will far surpass the capabilities offered by ISDN, but the two complement each other in that ISDN is the ideal transport medium for accessing the intelligent networks.

Perhaps no other change affecting data communication has been as dramatic as the tremendous amount of fiber cable that has been deployed since the 1980s. Yet this bandwidth capacity has not been fully exploited for digital transmission. Chapter 27 discusses an emerging technology, SONET (synchronous optical network), which will eventually allow the full potential of fiber optics to be realized. In addition to specifying a standard fiber optic interface, current SONET standards specify transmission rates that start at 51.84 Mb/s and reach to 2.488 Gb/s, with provision to go to 13 Gb/s. The long-term ramifications of these standards are far-reaching: SONET will eventually replace the asynchronous networks of today; with the same fiber cable that supports asynchronous networks, transmission capacity can be increased a thousandfold by using end-to-end SONET equipment. The chapter also explores the relationship of T3 to SONET, and how the choice will affect corporate users.

Chapter 28 discusses the technologies that go into the development of metropolitan area networks (MANs), which are essentially huge LANs that can encompass an entire city, providing data transport at fiber optic speeds. The development of MANs is being driven by the increasing need among LAN users for high-speed data services, unconstrained by geographical limitations. Although teleports and fiber carriers have expressed interest in the MAN concept, most of its development activity focuses around telephone companies, which recognize that future account control hinges on their ability to incorporate such a service into their customers' embedded base of information processing and networking systems. Moreover, MANs provide an infrastructure within which telephone companies can better position themselves for the provision of multiple services by deployment of broadband ISDN.

Part VI concludes the book with a discussion of networking for competitive advantage, a topic that will assume added significance in the 1990s as businesses worldwide begin to appreciate the strategic role of communication networks in tracking and penetrating emerging markets, responding to changes in customer demand, and supporting expansion opportunities at home and abroad.

Nathan J. Muller
Robert P. Davidson

Acknowledgments

This book integrates a broad range of technical concepts, operational issues, and management concerns into a single easy-to-read reference that can be used to help organizations plan and implement networks in the 1990s. We drew freely on the experience and expertise of our colleagues at General DataComm, who have assisted us with advice, ideas, and critiques.

We are especially grateful for the assistance of Jack O'Neil, who contributed generously to the chapters on PBXs, multiplexers, digital cross-connect systems and services, network design, network restoral, and fractional T1. He also provided much of the tariff analyses sprinkled throughout the book.

Mike McLoughlin also provided generous assistance. Many of the ideas included in the chapter on fractional ISDN originated with him. He also contributed to the chapter on SONET. Bill Victoria provided insights that were incorporated into the chapters on network management and, with Dom DiSario and Max Arbo, high-speed diagnostic modems. Hugh Goldberg played a key role in the development of the chapter on channel service and digital service units, and provided much of the technical input about regulatory involvement.

Many of the ideas in the chapter on network design are the result of collaboration with Dick Bachmann and Jim Senyszyn. Tom Bednarczyk assisted in the development of the chapter on systems integration as well as that on maintenance and support services. Timothy J. Smith and John Galloway assisted in the development of the chapter on multiplexers, and provided information for the chapters on network management and ISDN.

Special thanks go to Charles P. Johnson, chairman and CEO of General DataComm, and Ross Belson, GDC's president and COO, for maintaining the kind of work environment that makes projects such as this possible. In this regard, we would also like to thank Fred Cronin, Bob Smith, and Dan Young.

We also wish to acknowledge the valuable contribution of Vinton G. Cerf, chairman of the Internet Activities Board and series editor for the Artech House telecommunication library.

The opinions expressed in this book do not necessarily reflect the views of General DataComm, Inc.

Contents