Internetworking LANs

Operation, Design, and Management

Robert P. Davidson
Nathan J. Muller

Internetworking LANs: Operation, Design, and Management

Robert P. Davidson and Nathan J. Muller
Cover and Illustrations by Linda Lee Tyke

Artech House
Boston • London

Library of Congress Cataloging-in-Publication Data

Davidson, Robert P.
Internetworking LANs : operation, design, and management / by Robert P. Davidson and Nathan J Muller ; illustrated by Linda Lee Tyke.
p. cm.
Includes bibliographical references and index.
ISBN 0-89006-598-5
1. Local area networks (Computer networks) 2. Wide area networks (Computer networks) I. Muller, Nathan J. II. Title.
TK5105.7.D38 1992 92-8204
004.6'16—dc20 CIP

685 Canton Street
Norwood, MA 02062

International Standard Book Number: 0-89006-598-5
Library of Congress Catalog Card Number: 92-8204

10 9 8 7 6 5 4 3 2

Internetworking LANs:
Operation, Design, and Management

For a complete listing of the Artech House Telecommuncinciations Library,
turn to the back of this book

To our wives,
Myrna and Linda

Preface

Networking Trends

This book covers a wide range of issues relating to the operation, design, and management of local area networks (LANs) and LAN internets. Such a book is needed if for no other reason than to chronicle the major advances in the approaches to networking that have occurred over the last decade and to bring readers up to speed on the changes and challenges to come. More important, this book describes the impact of LAN internets on the business enterprise, providing practical advice on navigating the rapidly changing communications environment for competitive advantage. In doing so, this book puts into perspective the various products, services, and technologies that play an important role in helping organizations streamline their networks for optimal efficiency and economy.

In the early 1980s, big networks were still based on hosts and terminals. From the internetworking perspective, the major concern was how to get multihost connectivity from a single terminal. The problem was addressed by *front-end processors* (FEPs) that had lines coming in from terminals and channels going out to local and remote hosts. The FEPs provided switching and other functionality between the two sides of the network. The FEPs were typically large, relatively expensive units that resided in the corporate data center alongside a mainframe computer.

A few years later, Ethernet changed this switching fabric, lessening the need for FEPs. Ethernet handled the switching and routing functions more efficiently and economically than FEPs, at least within the local area—an office floor, a building, several adjacent buildings, or a campus. Beyond this distance, the protocols that endowed LANs with their high speeds became ineffective. Nevertheless, throughout the 1980s, the number of LANs grew dramatically. LANs were no longer used solely to provide host-to-terminal connectivity but to implement distributed processing as well. This elevated networking to a new level—that of peer-to-peer computing. These concepts proved as popular as they were powerful. Soon the number of installed LANs grew into the millions. They became large enough that organizations again had to start thinking in new directions, looking for a better way to interconnect LANs, both within a building or campus and over a wide area.

To address these concerns, vendors came up with interconnect devices that were functionally similar to the FEPs that LANs replaced. Originally, bridges were used to connect local LAN segments, improving LAN performance. Later, like routers and gateways, the bridges became part of the internetwork between local and *wide area networks* (WANs). On one side, they provided LAN connections; on the other side, they provided connections to the WAN by way of T1 or other digital lines. In essence, a unique networking fabric was created using new types of internetworking devices—mostly bridges and routers—that could build large, complex internets spanning both local distribution and wide area connections.

Initially, these internetworks were arranged in tiers, using a topology that involved cascading bridges (Fig. P.1). But this imposed severe performance penalties because the arrangement forced too much traffic to traverse several bridges. With the introduction of multiport bridges, users can now hub networks (Fig. P.2) and eliminate the performance penalty associated with cascading arrangements. This precipitated the debate over when to use bridges and routers, which continues today. However, this debate may soon end, as the respective vendors contend for competitive advantage by closing the functional gap between the two types of devices.

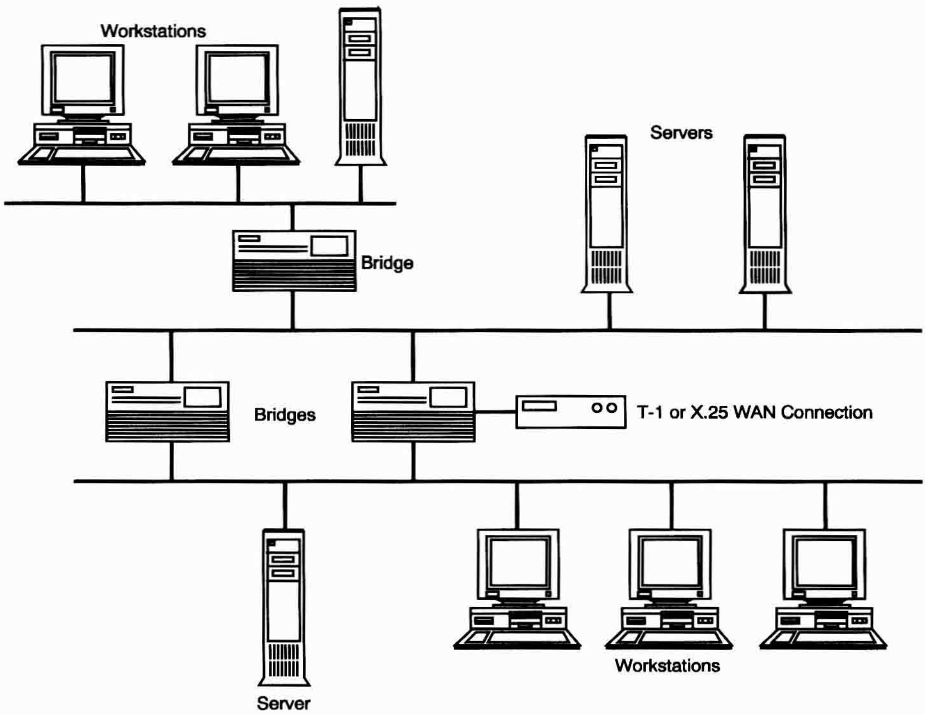

Figure P.1 Network of cascading bridges.

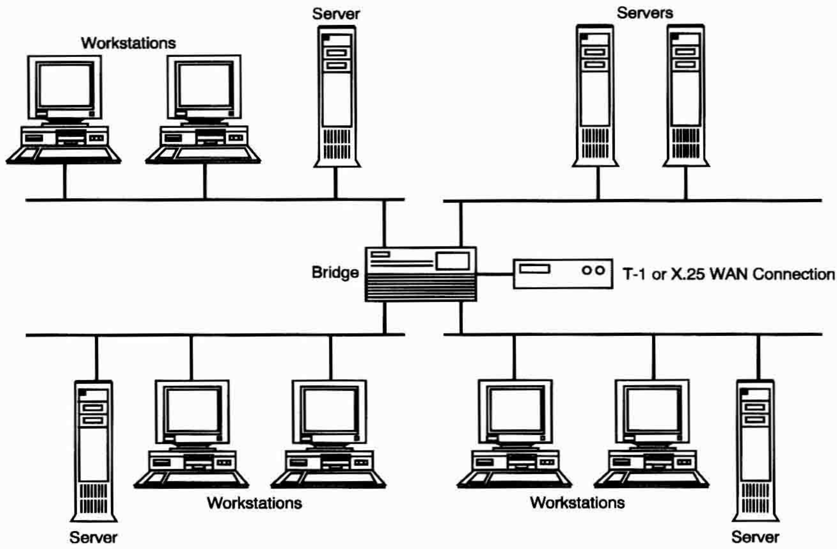

Figure P.2 Hub network with multiport bridge.

Of equal importance is that the switching and routing functions of internetworking have migrated back to individual boxes. Before distributed computing can scale effectively to global WANs, technology that returns the switching and routing functions to the fabric of the public network must emerge and be effectively implemented. Much of the present debate over the intelligent public network centers around the needs of this emerging computing environment. Worldwide, carriers are looking at multibillion-dollar investments in their infrastructure to upgrade the relatively dumb, asynchronous network elements into manageable, synchronous elements.

What is driving this change is the mismatch of LAN-WAN technology and the complexity it creates. LANs, connectionless services that carry traffic in frames, are appreciated for their bandwidth-on-demand capabilities as well as for their high speed. Nodes are logically connected to any other node on the network. This peer-to-peer connectivity is a function of the frame-oriented nature of LAN traffic. The frames themselves contain the addressing information that allows any station on the LAN to communicate with any other station.

When connected to WANs, however, the situation changes. The standard topology involves the interconnection of LANs by bridges or routers. Under current WAN schemes, these devices are often connected to circuit switches deployed throughout the public network. These switches are designed to support voice traffic over point-to-point lines. The lines provide bandwidth in fixed increments, usually 64-Kb/s channels; they make no provision for elasticity based on the requirements of the applications. This inflexibility prevents users from allocating bandwidth on demand and forces lines to be dedicated to specific applications. When the application is not using the WAN, its allocated bandwidth goes unused, even if other applications need it.

For this and other reasons, users experience significant performance degradations when going from the LAN to the WAN. The growth of distributed databases and client-server applications only worsens this situation by creating more traffic that must move from high-performance LANs across the relatively slow WAN. Simply increasing the speed of the WAN will not help the situation much because WAN technologies were not designed to carry data traffic—at least not at the speeds to which LAN users are accustomed.[1]

To handle traffic, the WAN opens a circuit—usually a point-to-point line—and sends the data over it. Even though the WAN establishes a circuit that can be used for data transmission, the WAN contains no inherent routing capabilities. Rather, the intelligent interconnection devices at each end of the circuit determine the routes. Data that must be sent to multiple destinations travel to each different destination by way of a different circuit. Consequently, the routing or switching function resides at the customer premises: in the bridges and routers that connect the LAN to the WAN. This is the source of problems in scalability and complexity that are inherent in today's WANs.

Because the switching function of today's WANs is in the bridges and routers, the responsibility of managing the WAN routing function lies with corporate *management information system* (MIS) professionals and LAN managers. Essentially, they must manage their own enterprises' WAN. But until they obtain the tools that make LAN-to-WAN internets easier to implement, maintain, and manage, they will not be able to meet the global connectivity needs of the enterprise.

Frame relay and IEEE-based 802.6 metropolitan area networks (MANs) represent two solutions to the problems posed by conventional public data services. In fact, frame relay and 802.6 standards are both intended to change the networking fabric of WANs by allowing them to accommodate the frame-oriented nature of LAN data communications. While frame relay grew out of the ISDN standards and is a streamlined version of the X.25 standard for packet data networks, 802.6 constitutes a radically new approach to wide area networking. A connectionless service, 802.6 provides a logical networking fabric that allows the WAN to act like a LAN. The 802.6 standard provides the same datagram services over a WAN that LAN protocols provide over a LAN. An addressing scheme based on telephone numbers allows the 802.6 protocol to route packets to multiple nodes on the network by way of a single connection to the network. In the same way that nodes on a LAN are logically connected to every other node on the network by way of one LAN con-

[1] WANs are referred to as both public and private. In actuality, many private networks connect to the public network by dedicated or switched circuits. It is the analog- or voice-compatible elements within the public network that degrade data transport. For LANs, even data-compatible WANs that are packet-switched employ relatively slow X.25 systems.

nection, all networks connected by way of an 802.6 MAN will be logically connected to every other network on the MAN by way of one line.

802.6 MANs provide the necessary transition from narrowband to broadband services, creating new types of WANs. Because of its high bandwidth, ability to scale, and frame-switching orientation, 802.6 will facilitate the building of national and global internets designed to handle the bursty nature of LAN traffic. Links among MANs will be provided by way of another emerging backbone technology, the *synchronous optical network* (SONET), which is capable of gigabit speeds, using intelligent network elements.[2] However, the complexity and cost of 802.6 switches and networks virtually assures that 802.6 networks and the links between them will be built and maintained by the *Regional Bell Operating Companies* (RBOCs) and interexchange carriers. Users who would rather not build and manage their own wide area internets will welcome 802.6-based services as an attractive alternative.

Commercial implementations of 802.6 are referred to by Bellcore as *switched multimegabit data services* (SMDSs). Initially, SMDS will offer only connectionless services to meet the urgent need for LAN interconnection. Circuit switching may be added later to support voice and video as user acceptance of the service grows. With SMDS, the user can create a private data network over a public service. Through address filtering, users will be able to specify to the carrier those addresses from which they want to receive data; all others will be locked out.

Access to SMDS will be by way of T1 or T3 lines. Large users will prefer T3 connections: at 44.736 Mb/s, real-time metropolitan area networking can be achieved; whereas the use of T1 connections entails speed adaptation that could degrade internetworking performance. Once SMDS becomes available and all of the enterprise's LANs are connected it, the switching function will revert to the fabric of the network and will no longer have to reside at the customer's premises other than for local switching. The responsibility of managing the network will largely be lifted from the customer and fall to the carrier. Since the network is highly scalable, customers will be able to build large global internetworks without having to increase their management burden. Already, carriers are aligning with bridge-and-router vendors to offer services and equipment as part of an overall package that includes management, service, and support.

Because SMDS constitutes an entirely new approach to networking, it cannot be simply overlaid onto the existing telephone network. Rather, it is a separate network requiring new high-speed switches and standardized interfaces for access. Implementing this new network will not be cheap or easy; consequently, SMDS will be phased in as the RBOCs amass the resources to build and support SMDS networks and the multiprotocol environments they will enable. This presents some interesting management challenges that affect both carriers and their customers.

---

[2] SONET is not switched, so the 802.6 links may need to be supported by cell-switching technologies such as the *asynchronous transfer mode* (ATM).

Management systems designed by vendors of *customer premises equipment* (CPE) must be able to exchange information with the carrier's management systems. The two "open" approaches to network management are the *simple network management protocol* (SNMP) developed for the *transmission control protocol/internet protocol* (TCP/IP) environment and the *common management information protocol* (CMIP) developed for *open system interconnection* (OSI) compliant networks. These standards offer data collection and some rudimentary control, but neither does much to explain what the collected data mean. Still needed are management tools that allow users to identify and solve problems quickly, or allow the user to be notified of carrier-related problems. Beyond this, carrier management systems must assume responsibility for problem correction, using a new class of intelligent networking elements.

Bridge, router, and intelligent hub vendors are working with the carriers to provide such solutions. These solutions will offer real-time information about the health of the device or subnetwork at a glance. Performance data will be used to create a baseline for measuring future performance and for setting user-defined performance thresholds that will allow ports on a hub, for example, to be automatically enabled or disabled. Trouble tickets will be automatically generated and sent to the carrier or to the manager of a remote LAN. Real-time network performance information presented in more intelligible format will reduce the need for around-the-clock troubleshooting by a staff of experts.

Most routers and bridges use the *high-level data link control* (HDLC) protocol to communicate over point-to-point lines. Equipment vendors are looking to frame relay, a modification of HDLC, to provide a multiplexed interface among bridges and routers and T1 nodal processors that compose the WAN. As a multiplexed WAN interface, frame relay, like 802.6, does not require a dedicated connection to each device on the WAN or for each application that needs to use the WAN.

Frame relay pushes the limits of today's WANs, permitting the full bandwidth of a T1 or T3 line to carry LAN traffic. Frame relay not only pushes the limits of WANs but allows users to upgrade their existing equipment at only an incremental cost to access the service. With traditional T1 multiplexer vendors supporting frame relay and building the next generation of cell relay switches, users will have the means to build their own high-performance networks that not only surpass what is offered by conventional circuit-switched *time-division multiplexers* (TDMs) but also provide an economical alternative to 802.6 carrier-based services.

Unlike 802.6, frame relay is a connection-oriented service. This means network managers must preassign addresses to specific LANs on the WAN. When a device supporting frame relay communicates over the WAN, the network decodes the destination from the address field, and the frame can be segmented into smaller packets for transmission. The frame is reassembled at the destination and submitted to the receiving device.

Modifications to the HDLC addressing scheme allow frame relay to support multicast services. Every connected network will be logically adjacent and will appear to be only one hop from every other network connected to the WAN. Because the interface is multiplexed, frame relay can provide bandwidth on demand. Data from multiple sources can be interleaved on the same available bandwidth.

The availability of frame relay and 802.6-based MANs (including SMDS) portends sweeping changes in the way LANs are interconnected over the WAN, determining, among other things, the scalability of today's LANs. While 802.6 represents a completely new approach to networking, frame relay is only a modification of the existing network elements to make them more suitable for high-speed data communications. While frame relay will soon match the SMDS initial T3 rate, SMDS itself will migrate to ultra-high-speed SONET OC-3 and OC-12 rates, 155 Mb/s and 622 Mb/s respectively, satisfying even higher bandwidth imaging, videoconferencing, and multimedia applications. Eventually, frame relay networks could feed information onto these higher rate SMDS highways. So the two are not necessarily competing services; both will be around for quite some time.

These are only a few of the high-impact topics that are fully explored in this book. We start with a discussion of micro-to-mainframe links in Chapter 1. Since the 1960s, mainframe computers have become a fixture in most large organizations, yet for many years the absence of real-time access to mainframe information became widely recognized as a barrier to timely and effective decision-making. With microcomputers beginning to migrate into corporate offices in the late 1970s, enterprising vendors started developing ways to link them to the mainframe with different combinations of hardware and software. The result was an immediate improvement in the speed at which mainframe information could be accessed and distributed.

As organizations became more reliant on information and its movement among different workgroups and departments, the stage was set for the acceptance of local-area networking. Chapter 2 traces the origins of Ethernet and describes the topologies and principles of operation of the popular LAN types.

Corporate computing has largely moved from the traditional centralized mainframe toward distributed networks of people and resources sharing information. Chapter 3 discusses the distributed computing environment as exemplified in the emerging client-server architecture, contrasting it with the centralized mainframe environment and, to a lesser extent, with another form of distributed computing known as peer-to-peer. The tight coupling of people and resources implied under the distributed computing concept makes for better and faster decision-making, which, in turn, affects corporate profitability and competitiveness.

Chapter 4 discusses a topic that is not often covered in literature about LANs: *interprocess communication* (IPC). IPC is fundamental to an understanding of enterprise management and may be implemented within a single device or between remote devices, underlying the communication for both LANs and WANs. The IPC facilities of operating systems allow arbitrary threads and processes to exchange data and to synchronize execution. With these facilities extended over a network, true distributed applications can be built out of simpler, reusable building blocks. In the process, the need to develop separate modules for each service environment is eliminated, which decreases development costs. This capability has already proven itself in environments that use multilayered communications protocol suites, such as TCP/IP, and will become equally valuable in emerging OSI networks.

With LANs becoming more complicated at the physical and management levels, the need has arisen for devices that not only facilitate the management of sprawling LANs by providing real-time status information and statistical reports, but prevent relatively simple problems from bringing down the entire network. Chapter 5 describes how this functionality is achieved through a relatively new device called the "intelligent hub." These devices evolved from relatively simple wiring hubs to become indispensable components of LANs and internetworks.

Approaches for interconnecting LANs over the WAN are addressed in Chapter 6. The devices that feed traffic across the WAN fall into the categories of repeaters, bridges, routers, and gateways. These devices all interconnect networks, and, to varying degrees, their functions overlap. So the choice of the most appropriate interconnection device may not always be clear. This chapter contrasts the functions of each device within the framework of the international OSI reference model.

Chapter 7 offers case studies in LAN internetworking, drawing upon the previous chapter's discussion. The case studies come from information-intensive financial organizations, putting interconnectivity solutions into focus against real-world business issues.

Chapter 8 describes a suite of nonproprietary high-speed networking protocols known as TCP/IP, which traces its origins to the Federal Government's *Defense Advanced Research Projects Agency* (DARPA), the research and development arm of the *Department of Defense* (DoD). TCP/IP is valuable in environments where communicating, translating, and requesting data from diverse computer systems is done on a fairly routine basis. Many corporations have adopted TCP/IP and its related network management protocol, SNMP, because they could not wait for industry consensus on the much-heralded OSI standards.

The OSI reference model offers a useful framework for understanding the functional aspects of today's complex networks. Chapter 9 describes the model in terms of the functionality found in each layer. Issues of OSI conformance and product interoperability are also covered, along with product evaluation criteria.

The emergence and rapid growth of client-server, processed-image, and *computer-aided design/computer-aided manufacturing* (CAD/CAM) applications will require new communications methodologies and wide area networking protocols. These applications generate large amounts of data locally and then require that high-speed bursts of data be transferred between dispersed corporate locations. Chapter 10 discusses how fast-packet technologies, including frame relay, handle these applications.

Chapter 11 examines carrier-provided alternatives for LAN interconnection, highlighting the IEEE's 802.6 standard for *metropolitan-area networks* (MANs) and Bellcore's service implementation, SMDS.

Many high-speed services are already supported over fiber-optic facilities. Aside from the obvious speed advantages offered by optical fiber, the cost of new installation is now low enough to rival that of copper wire. Chapter 12 describes the *fiber distributed data interface* (FDDI) and its self-healing functionality for LANs and MANs. This chapter also discusses the emerging SONET, highlighting its advantages over T3, integral network management and control functions, and the types of equipment that go into building a totally new wide area networking infrastructure capable of providing gigabit transmission.

From FDDI and SONET, we go to link test and measurement, the subject of Chapter 13. Isolating problems on today's distributed enterprise-wide network is consuming an increasing share of valuable organizational time and resources. This chapter describes the varied functionality of diagnostic equipment that includes complex protocol and line testers, specific service testers, and the old-standby breakout box.

The continued expansion of LANs, from the department level to the building and campus level to the worldwide enterprise level, has raised LAN management needs to the highest level of sophistication. Chapter 14 describes the different types of LAN management systems and the features and functionality associated with each. To illustrate, the products of IBM, Banyan, and Novell are profiled.

Chapter 15 describes the IEEE's 10Base-T standard, which provides 10-Mb/s Ethernet performance and functionality over ubiquitously available *unshielded twisted-pair* (UTP) wiring. Because 10Base-T specifies a star configuration, installation entails the use of some components not always found on traditional Ethernets. This chapter describes these components and discusses the significance of 10Base-T for LAN management and control.

Chapter 16 discusses the various approaches to WAN management, from proprietary management systems to the "open" integrated network management system that spans the traditionally separate domains of LANs and WANs.

With networks becoming more data intensive, precautions must be taken to safeguard one of the organization's most valuable assets: information. Chapter 17 describes procedures, equipment, and circuitry that must be in place to maximize the reliability and availability of LANs and MANs.

In today's distributed computing environments, with so much information flowing from LANs onto wide area internets, it is imperative that procedures, equipment, and circuitry be in place to ensure uninterrupted information flow. Chapter 18 describes the methods available to do this, including carrier-provided redundancy and protection services, customer-controlled reconfiguration services, bandwidth on demand by way of ISDN, and simple dial backup. While some methods rely solely on available carrier services, others rely on the built-in capabilities of CPE. This chapter also discusses hybrid restoral arrangements consisting of integrated management systems, carrier services, and CPE capabilities.

Today's, LANs, micro-to-host links, gateways, bridges, and other data communications equipment and facilities provide microcomputer users with access to an abundance of corporate data, and this access will intensify as a result of the growing trend toward distributed computing. Chapter 19 discusses the problem of maintaining network security and elaborates on possible solutions.

Finally, Chapter 20 discusses some important technology implementation issues that corporations must deal with in the 1990s: downsizing, outsourcing, and network integration. While network integration has always been an integral part of technology implementation when communications links are involved, downsizing and outsourcing are relatively new concepts to most network managers. Nevertheless, all three strategies are oriented toward the deployment of technology in ways that will enhance the competitive position of the enterprise. This chapter explores some of the technology and management problems that companies often face as they continue their quest to streamline business operations to achieve the efficiencies and economies that will ensure their long-term survival.

ACKNOWLEDGMENT

We wish to single out Vinton G. Cerf, chairman of the Internet Activities Board and series editor for the Artech House telecommunications library, for participating in the creation of this book with his insightful critiques and technical advice.

Robert P. Davidson

Nathan J. Muller

Contents