



NEXT ▶

Gigabit Ethernet: Technology and Applications

by Mark Norris

ISBN:1580535054

Artech House © 2003 (270 pages)

This resource offers a detailed understanding of how Gigabit Ethernet is used to build storage area networks, to scale up local area networks into total area networks, and to provide wireless solutions.

Table of Contents

[Gigabit Ethernet Technology and Applications](#)

[Chapter 1](#) - The Quiet Revolution

[Chapter 2](#) - Ethernet—The Story So Far

[Chapter 3](#) - Gigabit Ethernet

[Chapter 4](#) - Wireless Ethernet

[Chapter 5](#) - Total Area Networks

[Chapter 6](#) - Storage Area Networks

[Chapter 7](#) - A Changing Marketplace

[Chapter 8](#) - Managing Total Area Ethernet networks

[Chapter 9](#) - Through the Looking Glass

[Appendix A](#) - Complementary Technologies

[Appendix B](#) - Competing Technologies

[Glossary](#)

[Index](#)

[List of Figures](#)

[List of Tables](#)

Team LiB

NEXT ▶

Team LiB

◀ PREVIOUS

NEXT ▶

Back Cover

Gigabit Ethernet delivers the speed and high bandwidth that today's organizations demand from their local area network. It is being chosen over other high-speed technologies because it is a flexible and cost-effective solution that can be used in a wide range of applications. This comprehensive guide offers a clear picture of how Gigabit Ethernet works and how it can be used for a broad range of services. Practitioners learn how Gigabit Ethernet is utilized in the wide area and how it can support mobile systems.

The book also offers a detailed understanding of how Gigabit Ethernet is used to build storage area networks, to scale up local area networks into total area networks, and to provide wireless solutions. By comparing Gigabit Ethernet to competing technologies, this guide helps professionals decide when Gigabit Ethernet is the right solution for their networking needs. This complete reference also provides a thorough treatment of the Ethernet standard and its many variants.

About the Author

Mark Norris is a technical director of Norwest Communications, Suffolk, UK. He is the author of *Mobile IP Technology for M-Business* (Artech House, 2001), *Understanding Network Technology: Concepts, Terms, and Trends, Second Edition* (Artech House, 1999), and *Survival in the Software Jungle* (Artech House, 1995); the coeditor of *Systems Modeling for business Process Improvement* (Artech House, 1995); and the coauthor of *Component-Based Network System Engineering* (Artech House, 2000).

Team LiB

◀ PREVIOUS

NEXT ▶

Gigabit Ethernet Technology and Applications

Mark Norris

Artech House

<http://www.artechhouse.com>

Library of Congress Cataloging-in-Publication Data

Norris, Mark.

Gigabit Ethernet technology and applications / Mark Norris.

p. cm. - (Artech House telecommunications library)

Includes bibliographical references and index.

ISBN 1-58053-505-4 (alk. paper)

1. Ethernet (Local area network system) I. Title. II. Series.

TK5105.8.E83 .N67 2003

004.6'8-dc21 2002033227

British Library Cataloguing in Publication Data

Norris, Mark

Gigabit ethernet technology and applications. - (Artech House telecommunications library)

1. Ethernet (Local area network system) 2. Gigabit communications

I. Title

004.6'8

1-58053-505-4

Cover design by Gary Ragaglia

Copyright © 2003 ARTECH HOUSE, INC.

**685 Canton Street
Norwood, MA 02062**

All rights reserved. No part of this book may be reproduced or utilized in any form or by any means, electronic or mechanical, including photocopying, recording, or by any information storage and retrieval system, without permission in writing from the publisher.

All terms mentioned in this book that are known to be trademarks or service marks have been appropriately capitalized. Artech House 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.

International Standard Book Number: 1-58053-505-4

Library of Congress Catalog Card Number: 2002033227

10 9 8 7 6 5 4 3 2 1

About the Author

Dr. Mark Norris has more than 20 years' experience in software development, computer networking, and telecommunications systems. During this time, he has managed dozens of projects to completion, from small projects to multisite, multimillion-dollar projects. He has been published widely during the last 10 years, writing a number of books on software engineering, computing, project and technology management, and communications and network technologies. Dr. Norris lectures on network and computing issues, has contributed to references such as Encarta, is a visiting professor at the University of Ulster, and is a fellow of the IEE. He plays a mean game of squash but tends not to mix this with networking of any kind. He can be found at mnorris@iee.org.

Acknowledgments

Writing a book is a bit like water torture-it goes on for a long time and is never out of your thoughts. One of the nicest moments is when you finish, stand back, and look at the final product. It is then that

you realize just how much other people have helped. There are a few people without whom this text would never have seen the light of day (or, if it had, as an altogether poorer work).

If you do come across a particularly insightful, interesting, (or witty) passage in this book, it is probably due to Dave Piscitello, a star if ever there was one. Every inch of the text has benefited from Dave's deep understanding of networks in general and Ethernet in particular. Dave is an avid user of Ethernet technology and can be contacted (very quickly and reliably) via dave@corecom.com.

Two more heroes (or rather heroines) to cite are Tiina Ruonamaa and Julie Lancashire at Artech House. They have been a constant source of encouragement and have made sure that this manuscript saw the light of day in the time and style it should.

Finally, there are many friends and colleagues who have contributed their ideas and observations. In particular, John Atkins (who first thought of the Total Area Network), Dave Sutherland, Steve Pretty, and John Nolan (a mine of useful information).

Team LiB

◀ PREVIOUS

NEXT ▶

Chapter 1: The Quiet Revolution

Overview

The revolution will not be televised.
--Gil Scott Heron

For nearly a hundred years, the telephone network was the only communications "game" in town. Public operators provided this for the benefit of the global masses. The basic service, telephony, was stable and predictable and change was evolutionary—until about 20 years ago. Around that time, new technology allowed people to set up their own information *and* telephony networks. These operated both parallel to and in conjunction with the established, global telephone network. Shortly thereafter, deregulation and liberalization of the telecommunications industry began to fragment the telecommunications market place, which has and continues to be fiercely competitive.

For all these dramatic changes, there is still a clear distinction between the privately owned and operated local area networks (LANs) and the wide area network (WAN) products offered in the public marketplace. The issue that perpetuates this distinction is the costs of local access and long-distance connections, in particular, broadband connections (i.e., those providing transmission service in excess of one megabit per second). Until recently, only the large operators have been able to capitalize in this area, and only barely through economies of scale. But this statement is not as true as it was a few years ago.

There is a quiet revolution in telecommunications. The forces behind this revolution are a huge growth in long-distance transmission capacity, to the point of overabundance, the speed with which local area networking technology has advanced, and the ubiquitous adoption of LANs, from large enterprises to small businesses and even to individual residences. The full impact of these forces has yet to emerge. They could well combine to enable the long-heralded death of distance.

Interesting though this may be, the revolution is not complete, and the dramatic development in distance-insensitive communications is only one of a number of changes. In this book, we focus on Ethernet, which reveals, in a microcosm, how a single technological innovation can dramatically alter (and enhance) a 125-year-old industry in a scant 20 years. In this chapter, we start with a brief description of modern network technology. We will explain what a LAN does, where the WAN fits, and how the two are used to deliver applications to an end user. This chapter will paint a picture of the current communications market with particular attention to the oversupply of long-distance transmission, as well as cover the technology. In later chapters, we focus entirely on Ethernet—the technology that is at the nexus of the quiet revolution. First, though, a brief resume of the wider world of communications.

1.1 The Information Economy

Information has always been a valuable commodity. For years, people have sought to be the first to discover, apply facts and concepts, and then *share* these or the results of their application with others. On a few occasions, this has been purely through altruism on others it has been in pursuit of gain—much of the Rothschild family's fortune was made because they knew of (and capitalized on) the British victory at Waterloo before anyone else.

Access to information continues to grow in importance year after year. Information is now a vital business commodity. Its effective use is, at least, an issue of money and often one of survival. Information, shared over a distance, is now the key resource in many businesses. This point is readily illustrated. Not many years ago, the major assets of the airline business were the planes themselves. These days, even modern jets would be easier to replace, and judged of less value, than the associated flight booking systems.

The problem with information sharing is that there is so much information around to share in a timely manner. It is rather difficult to quantify this statement, as there is no one defined measure of information, but it is rather self-evident. Extreme sceptics might consider that more books have been published in the last 50 years than in the previous 500. The interesting statistics in [Figures 1.1\(a\)](#) and [1.1\(b\)](#) nicely illustrate the point, albeit in a slightly oblique way.

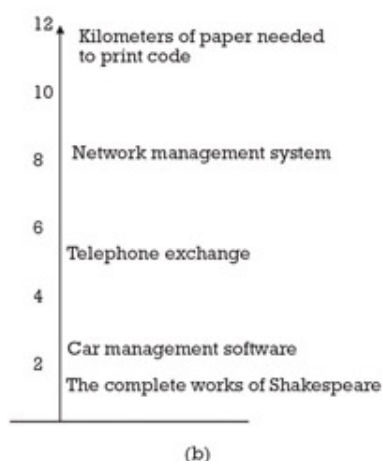



Figure 1.1: (a) The volume growth of information. (b) The amount of information that resides in software code.

[Figure 1.1\(a\)](#) is fairly self-evident. In [Figure 1.1\(b\)](#), each item is represented by a length of print—the millions of lines of code in modern software systems represents considerably more text than a set of works that has satisfied generations of readers. Moreover, consider that in less than two decades, the Internet and World Wide Web have made authors and publishers of millions of users who have

contributed a huge number of information sources.

None of this implies that people in the twenty-first century are any more brilliant or wise than their forebears, or that they value information retention more. Only that they have superior tools at hand to record it. Society seems to have increased its predilection for codifying; we record *everything* and as much about everything as possible. The net effect is that we now live in a society that creates lots of information, and one in which the competitiveness of business depends on it.

Of course, the volumes of information that have been filed away are only useful if they are available in the right place at the right time. This implies communication, which, in turn, implies a fast and efficient network.

Team LiB

◀ PREVIOUS

NEXT ▶

1.2 The Network Is King

The telephone network has been in place for many years now. In many ways, it is everything you would want of a global communications infrastructure-it is reliable, ubiquitous, and cheap. However, it was built to carry voice traffic and does not perform particularly well when used for data transmission. The ever-increasing demand to send pictures, video, and text grows, putting the telephone network's shortcomings in sharp relief. This is because the telephone network was originally designed to carry analog voice before being redesigned for digital voice. The transmission rates that were adequate for voice connections are absurdly small for data, and the consequent difficulties of reengineering a network based on time division multiplexing (TDM) are daunting.

But is not the network a secondary consideration? Should we not be more concerned with developing better on-line content, a wider range of network-based services, and more powerful applications? This seems to be what telecommunications providers are thinking. Up until a few years ago, they could only offer low-capacity connections to end users.

But now technologies such as digital subscriber loop (DSL) are being planned for introduction around the world, and these promise considerably higher bandwidth than current offerings. This bandwidth is universally touted as a way to provide Internet access, and, in particular, to sell content (that is, access to a fixed source of information) to users. The need for higher speed access is something with which few people would argue. Accessing content on the Internet today is like driving to a giant shopping mall that has only a single lane road. Latency and inaccessibility are killing content providers and business-to-consumer applications. Enterprises doing electronic business to business are struggling as well, for mostly this reason. But does the future lie in content provision, with network access no more than a basic enabler?

Throughout the history of communications, connections between people have consistently been more important (and lucrative) than access to information. For instance, the U.S. Postal Service of the 1800s was awash with newspapers that needed to be delivered. This was the dominant content in terms of volume, weighing about 20 times as much as the letters carried. But letters brought in most of the money needed to run the postal system-around 85% of the revenue. On the Internet, electronic mail is still the king, even if its volume is relatively small.

In both instances, the content may be dominant in terms of volume, but the real demand, and, hence, most of the money, is in providing communication from any one point to any other.

What does this mean for the network operators and their suppliers? In broad-brush terms, a Machiavellian approach (focusing on what people actually do, rather than ideology) would be to provide the connectivity people need, rather than try and second-guess what sort of content or "killer application" they might want (or think they want).

Put more simply, network providers should stick to what they are good at! Application services, e-business, and the like may capture the headlines and the imagination, but the operator that can offer the underlying network is likely to capture the money. On that basis, we do not stray into such exotica in this book but rather stick to the basics of communication networks.

The next question, of course, is what sort of network does an operator need to invest in? The simple and correct answer to this is that the operator should invest in the network that the customer wants. Given that the prime need is for users to establish whatever sort of connections they need, the more flexible, familiar, and, to no small degree, inexpensive, that network is to use, the better.

1.3 Technology to the Rescue

Traditionally, there are two types of networks. WANs are the elder statesmen of networks. The public-switched telephone network (PSTN) and the Internet are prime examples of WANs. These public networks have been around for years and have been engineered to provide a specific range of services (i.e., voice, electronic mail) to a large population.

On an altogether smaller geographic scale exists the LAN, which, as the name suggests, serves a small number of users over a restricted distance. The LAN tends to be considerably faster than a WAN, an advantage inherited from the original design goals of providing high transmission rates at a sacrifice of transmittable distance. Also, because LANs are usually privately owned and do not offer public service, they can be readily reconfigured to suit specific user needs.

Of course, the world of communications is not really this simple—real networks, public *and* private, tend to consist of both WAN and LAN elements, as shown in [Figure 1.2](#).




Figure 1.2: A typical network with wide area and local elements.

A user at one end of the network can access information stored locally over a LAN. If users need to get something stored on a remote database or public (Web) server, then they must use the WAN to get to the relevant site. Putting aside the details of routing and addressing and the complications of dial-in and mobile access, the point is that the end-to-end communications path over the example network in [Figure 1.2](#) is a composite one. It mixes two type of network *technology*, each of which has its own characteristics, to provide a network *service*. A very important consequence of the majority of LAN-WAN communications paths is that, from the user's point of view, the resulting network service often exhibits considerable delay. This is because of the transmission rate mismatch or the bottleneck that exists where the LAN connection (usually measured in megabits per second) connects to the WAN (which offers speeds more often measured in kilobits per second).

Of course, users should not really have to be concerned with network details. They are more interested in how long it takes to retrieve the files they want or to get the results of database queries. It should come as no surprise that users are curious or frequently dissatisfied when two seemingly similar actions offer vastly different performance. The question users, perhaps naively, ask, but one that accurately defines the network service commonly sought, is "Why do I need a LAN and a WAN? Why isn't there just one network that covers everything—a *total area network*?"

This is not an unreasonable question at all! We will come back to the concept of total area networking. Before discussing that, however, we need to look in more depth at the potential for networking over the wide area.

Over the last 10 years, we have seen huge investments in network infrastructure. A significant portion of the planet is now covered with high-capacity optical fiber. Coupled with transmission technology that makes better use of that fiber, it is now feasible to transfer huge amounts of data at very high speed

from one end of the world to the other.

Figure 1.3 shows a common arrangement where the WAN bottleneck is bypassed by linking two distant LANs with a private transmission link—in this case, a leased line to the point of presence of a long-distance carrier that provides international optical circuits. Of course, this is probably not a particularly useful option as it constrains communication to just two sites. A better option would be to connect to a public WAN that can provide switching. Public networks also bring economies of scale. So, even if ample fiber exists for all of us to own a strand, we still need to have something do the switching.




Figure 1.3: Avoiding the WAN bottleneck.

In order to understand what the current situation offers, we now look at the long-distance transmission situation in some detail.

1.3.1A Wired World

One of the very obvious legacies of the industrial revolution is the train system. In just about every country in the world, there is an extensive collection of main lines that serve the big cities and branch lines to connect the smaller outposts. The current information revolution is going much the same way, only this time it is optical fiber that is being installed as the information rails to serve future generations.

Just as we have multiple rail service providers in the world, so too do we have many optical network providers. Across Europe alone, there are at least 20 main players building high-capacity fiber-optic networks. Most of these connect financial and industrial centers and extend transatlantic capacity. Typically, the fiber networks installed are designed as rings, broken at key locations known as points of presence (PoPs). The potential of optical transmission technology combined with the flexibility of ring topology promises to be a major enabler of the information economy.

The basic reason for the increased interest in operating fiber-optic networks is that technology for transmitting information over such networks has improved rapidly during the last few years. As a result of the advances in technology, the cost of fiber networks (in terms of dollars per bit) has dropped dramatically.

The dramatic pace of change is readily illustrated by the relative bandwidth of two recent transatlantic cables. Global Crossing (previously known as Atlantic Crossing) built its first cable (AC-1) in 1998 with a capacity of 80 Gbps. Its second transatlantic cable (AC-2) entered service 3 years later with a capacity of 1.28 *terabits* per second (Tbps). Another operator, FLAG, installed its own link shortly after Global Crossing at 2.56 Tbps. The long-planned Oxygen cable started its design life at 600 Gbps and is now planning to carry 2.56 Tbps.

Fiber-optic cable typically has passbands in the infrared area of the electromagnetic spectrum. Narrowbands at around 1,330 and 1,550 nm wavelength, depending on the type of fiber, are the most

effective frequencies for carrying a light signal. Modern fibers are doped to shape the passband at these frequencies and therefore increase the width of the usable bandwidth of the fiber. It is also possible to pump the fiber using a laser with a known frequency (a process known as Raman amplification) and this increases the gain of the fiber in parts of its spectrum.

Dense wave division multiplexing (DWDM) is a technique that exploits the way in which fibers carry data. With this technique, the fiber passband is broken into individual channels ("colors"). This effectively turns the single fiber into a multifiber bundle, each of which is capable of providing a high-capacity channel. Early DWDM systems offered 10 to 20 channels. Today, terrestrial systems using DWDM are reaching upwards of 100 channels, and the technology holds the promise of being able to increase to 2,000 channels. Submarine cables tend to adopt advanced technologies such as these slightly behind their terrestrial counterparts because of the reliability issues that affect long-haul undersea cables and the difficulty of gaining access to them once laid. However, each new cable that is designed carries significantly more traffic than its predecessors.

In addition to adding more channels, the basic frequency at which each channel operates is also rising. Whereas transmission occurred at STM-4 or 620 Mbps a short while ago, systems using STM-16 (2.5 Gbps) are commonplace today, with STM-64 (10 Gbps) becoming the current de facto technical standard. Systems operating at 40 Gbps and even 100 Gbps have been deployed and can be expected to become more widespread.

A new level of technology called optical switching offers functionality in the optical domain that has hitherto been the province of the electronic domain. As this takes hold, another level of functionality and capacity can be offered from the same basic fiber.

With such a wide range of different technologies all extending the capacity of fiber, we have seen a huge rise in the availability of long-haul bandwidth. This is illustrated in [Figure 1.4](#), which shows the accelerating trend in capacity of transatlantic cables.




Figure 1.4: Transatlantic transmission capacity.

It is clear that the capacity available in a single fiber is huge. Even with today's technology, one fiber can have 16 colors, each carrying 2.4 Gbps, which equates to more than 16,000 standard 2-Mbps (E1) circuits, enough to carry 4,000 million minutes of voice traffic per month. Put another way, this is enough capacity to download the entire movie *Gladiator* in a fraction of a second. In a global context, all of the international telephony traffic leaving the United Kingdom during 1997 would fit on one fiber using current technology. With 64 wavelengths on a fiber, each carrying 80 Gbps, there would be enough capacity (over 5 Tbps) to virtually hold all of the world's telephony traffic!

While these advances are releasing significant new capacity from fibers, the basic costs of laying undersea cables are staying broadly constant, or even dropping slightly. Furthermore, landline trenching of new optical cable is practically unnecessary in all but suburban and rural areas, as the fiber already in the ground far exceeds anticipated needs for decades. A sample of cables laid over the last few years indicates laying costs dropped from around \$40,000 to \$50,000 per kilometer to \$20,000 and below, depending on the undersea conditions. [Figure 1.5](#) illustrates this.




Figure 1.5: The falling cost of long-distance transmission.

The net effect of these developments is that the unit cost of bandwidth is dropping extremely quickly. Cables coming into service today have average construction costs of around \$30,000 per STM-1, and this figure is dropping at a rate estimated at more than 50% per annum. So the consequence of the technological advances being made is that for the immediate and midterm future, cable capacities will continue to rise and, consequently, unit capacity costs will continue to drop rapidly. The favorable conditions for installing cable has prompted a dozen or more companies to lay fiber-optic cable across Western Europe alone, many more worldwide. With each believing that they would gain first mover advantage, these companies have together created a surfeit of transmission capacity.

Not surprisingly, several of these pioneering companies are going out of business because they spent a lot of money on now obsolete assumptions of how much capacity each fiber provides (and, hence, how many customers each would support). Too many players resulted in too much capacity. Few can recoup their investment in equipment or even continue operations in the resultant fire sale economy.

In the past, the cost of the long-distance part of communications paths was dominant. The most striking impact of high-capacity fiber from the consumer's point of view is that cost dynamics are now, and perhaps forever, very different from those of the recent past. Costs are now driven by the "last mile"-the link between the user (office or residence) and the network provider's PoP. Today, the local loop access, or "tails", is a significant if not dominant part of total cost, and this does not seem to be falling at the same rate as the long-distance part, probably because there are fewer (if any) competitors (even after deregulation). With the local loop seen by many operators as something of a millstone and ownership nine-tenths of the law, there is little drive to innovate.

1.3.2 The Total Area Network

The availability of affordable and plentiful long-distance capacity is likely to have significant impact on the way networks evolve over the next few years. When the customary bottleneck is removed from the long-haul part of an end-to-end connection, the differentiation between LANs and WANs begins to disappear-they both switch traffic and, when distance-related factors are removed, they look very much the same.

This may seem to be no more than a logical progression enabled by better network technology. But there are significant implications to total area networking, some driven by user needs, others by economics, and others still by the technology itself.

The first implication derives from looking at the innovation that has been evident over the last 10 years through the use of LANs in general and Ethernet technology in particular. With a shift from local to total area networking, there is likely to be a matching shift from local to global business operation. This trend is already well established but will be accelerated to a new level by technology.

A second implication is that the removal of barriers, imposed at present by disjoint networks, will make it viable to treat the network as a commodity rather than a slightly fragile resource, tended by specialists. This will encourage the use of the network and, in turn, the demand for greater connectivity. Put another way, the more users on a network, the more useful that network is and the greater the likelihood that it will expand to accommodate yet more users. This virtuous circle was first noticed by Robert Metcalfe and has been dubbed "Metcalfe's Law."

One of the biggest factors in the speed of development of the total area network will be economics. LANs are already commodity, and we now have WAN costs dropping (albeit slowly). This leaves local access, between the user and the network PoP, as the missing link. The sooner cost is removed from

this part of the total area network, the sooner it will take off. And there are plenty of prospects for Ethernet solutions here.

Team LiB

[← PREVIOUS](#) [NEXT →](#)

1.4 Whence We Came

In order to understand what sort of total area networking we might expect to evolve, we need to look at the established network. As the rest of the book covers the LAN and, in particular, the Ethernet, we provide some balance by giving a brief overview of today's leading voice and data networks.

1.4.1 The PSTN

The telephone network that people are most familiar with has evolved over more than a hundred years. It is a very large and complex system with global coverage, engineered to provide universal and nearly uninterrupted public service. It consists of a number of linked elements. The main ones are described in the following paragraphs.

An *access layer* is used to deliver the service to the end customer and is usually provided using twisted copper pairs. There is some use of optical fiber for higher volume business customers, and coaxial cable is used for TV distribution, but coverage is variable—fiber tends to be preinstalled on new sites but economics dictate whether it is laid to established sites.

There is some use of multiplexing in the access layer, but the primary interface to the switching and core transmission layers is through a main distribution frame (MDF), where physical cross connections are made between copper pairs.

With copper likely to remain the dominant form of delivery to the wall socket in domestic premises, many network operators are seeking to "sweat the asset" of their installed plant. The main technology used to speed up the "last mile" of the domestic network connection is DSL. This takes various forms, all of which deliver a bandwidth of at least 500 Kbps over the existing copper pair infrastructure.

A *core transmission layer* of the network is provided primarily by digital transmission systems based on TDM. The use of TDM allows many digital signals to be interleaved onto a high-speed trunk or bearer circuit. Standards based on TDM, notably synchronous optical network (SONET) and synchronous digital hierarchy (SDH), have developed for the efficient carriage of digital voice. These standards provide inbuilt management and monitoring capabilities and allow simple dynamic reconfiguration in response to changing demands or network failures, which yields high availability.

The data rates specified by SONET/SDH are shown in [Table 1.1](#).

Table 1.1: Data Rates for SONET/SDH

SDH Signal	SONET Signal	Bit Rate	Capacity–Bearer Circuits	Capacity–Voice (Mmins/month)
STM-0	STS-1/OC1	51.48 Mbps	630 ISDN channels	5
STM-1	STS-3/OC3	155 Mbps	63E1 or 3E3	15
STM-4	STS-12/OC12	622 Mbps	252E1 or 4E4	60
STM-16	STS-48/OC48	2.488 Gbps	1008E1 or 16E4	240
STM-64	STS-192/OC192	9.95 Gbps	4032E1 or 64E4	960

Note SONET and SDH are equivalent technologies. The former is the United States variant, the latter is the European Union version.

A *switch layer* is used to provide the various switched services available (mostly dialed phone calls,

including voice and facsimile services). Most PSTN switches are limited to a maximum bit rate of 64 Kbps (the digital equivalent of one speech channel). There are three separate major switching services: circuit-switched PSTN and integrated services digital network (ISDN), which are partially integrated, and the packet-switched public data network (PDN), which is a separate service based on the X.25 standard for packet switching. The switch layer is also used to provide cross-connect functions for some leased services.

Controlling all this are various *network management* systems—usually one for each of the different types of equipment in the network. In addition, various service management systems are in place to control the different services provided. The service provisioning process is still dominated by the need to make physical interconnections at various points in the network, which makes an automated end-to-end process for service provision and rearrangement slow to effect and difficult to achieve with any measure of consistency.

1.4.2 The Internet

Like the PSTN, the Internet is a communications network that spans the globe. Unlike the PSTN, the Internet relies on packet switching rather than circuit switching. It was originally conceived as a data network rather than a voice network and has grown through the interconnection of many local area and regional networks, rather than being built as a single entity.

In essence, the Internet works by passing data using a standard communications protocol known as the Internet protocol (IP). The IP can be used on virtually every type of computer in use today and will operate over almost any network infrastructure, local and wide area, at subkilobit to terabit transmission rates. The Internet has an associated naming and addressing scheme, known as the domain name service (DNS), that allows resources (information, services, and people) on the Internet to be easily located. It is the universal acceptance and wide availability of IP, DNS, and other key standards, such as simple mail transfer protocol/POP (SMTP/POP) for mail, transmission control protocol (TCP) for end-to-end connection (transport and reliability) control, and Hyper-Text Transfer Protocol (HTTP) for rich text transfer, that give the Internet its global reach, application extensibility, broad availability, and vast user appeal.

One of the basic functions for which the Internet was built was to allow the transfer of digital information from one computer to another and, hence, from one person to another. In addition to the electronic mail protocols (notably, SMTP and POP), two information and file transfer application protocols, file transfer protocol (FTP) and HTTP, have become the framework on which the majority of Internet applications are built. Virtually all networked PCs have a preinstalled FTP client program (it is included with any browser), and simple FTP clients can connect to FTP hosts and view directories and download them according to user choice. Typically, FTP servers require a username and password before they will allow connection. There are some FTP servers that will allow anonymous login; these require the username to be anonymous and a password that is either your IP address or your Internet e-mail address.

HTTP is the protocol that enables perhaps the best-known Internet, the World Wide Web (WWW, or simply, the Web). The Web is an extension of the global Internet and builds on its established file transfer capabilities. Like FTP, Web browsers such as Netscape Navigator and Internet Explorer use a simple file transfer protocol, known as HTTP, to connect to Web hosts and can download information in any file format. The user is provided with a simple and intuitive interface for navigating information that is distributed throughout the Internet. Access to Web pages does not usually require you to enter a username or password (although some commercial sites do require registration and/or a fee).

For all of the sophistication that has been built into the Internet, it is the IP that provides the common base. IP, like Ethernet, is a packet-switching technology, so there is a high degree of synergy between the two. Together, they are capable of supporting a broad variety of applications, including voice, static, and streaming data, and video in its many incarnations. The natural fit with IP—the basis for such a wide range of applications—is what makes Ethernet the prime candidate technology for any network that it is possible to use it in.

1.5 About This Book

The ultimate purpose of this book is to give a comprehensive account of Ethernet technology and applications. This remit covers a fairly wide range as, from modest beginnings, the Ethernet has blossomed into a broad and fast-moving communications phenomenon. Given the potential scope of the book and the limited coverage that is possible in a few hundred pages, here is brief summary of what lies ahead.

This chapter sets the groundwork by reviewing, at a high level, the current picture of communication network provision. In particular, the way in which long-haul transmission technology has advanced, and the extent to which the matching network infrastructure has grown, is explained. [Chapter 2](#) introduces the Ethernet as a technology that has taken and seems likely to hold center stage in communications networks. The origins of the Ethernet are detailed along with its evolution into the widely deployed Fast Ethernet.

[Chapter 3](#) presents the technical details of the latest and fastest versions of the Ethernet-the Gigabit Ethernet. In this chapter, we explain the workings of both the Gigabit Ethernet and the latest development of the technology at the time of publication, the 10-Gigabit Ethernet.

In [Chapter 4](#), we take a brief look at an interesting addition to the Ethernet family, the wireless Ethernet and its potential role in both the LAN and local access. In addition to explaining the workings of this variant, we also take a look at its leading competitors.

In order to emphasize the evolutionary, rather than revolutionary, way in which Ethernet technology is developing, the structure of [Chapters 2](#) through [4](#) is very similar. The focus is on the changes from one generation of Ethernet to the next and the presentation follows a consistent path.

The last part of the book is all about what sort of applications are likely to be enabled by a Gigabit Ethernet. In particular, the potential for wider area networking services are considered in [Chapter 5](#) and an illustration of the practical issues and options in building such a network is given. [Chapter 6](#) considers the way in which Ethernets can be used to build storage area networks, and [Chapter 7](#) examines the potential of Ethernet technology in offering commercial fixed and mobile data services. In this chapter, we explore how Ethernet is penetrating the WAN market, its competitive position against third generation mobile networks, and the emergence of Web service, a standard for building applications that naturally fit onto IP- and Ethernet-based networks.

[Chapter 8](#) steps back from the network technology per se to consider a vital aspect of any network-that of network and service management. The key techniques and tools needed to keep a network operational are explained. To close, [Chapter 9](#) reviews the impact of Ethernet and discusses a few of its prospects.

To complete the book, two appendixes identify technologies that surround Ethernet. The first of these appendixes deals with those that seem likely to complement Ethernet, such as Bluetooth, the second with those that appear to compete with it, most notably asynchronous transfer mode (ATM).

[Table 1.2](#) is intended to help navigate the book as a whole. Some parts contain a lot of technical detail, others are more discursive-a quick reference to the ratings given could well minimize your chances of being bogged down or bored.

Table 1.2: A Guide to This Book

	Technical Content	General Interest	Specialist Detail
Chapter 1	**	****	*
Chapter 2	****	***	****
Chapter 3	****	***	****
Chapter 4	****	***	****
Chapter 5	***	****	***
Chapter 6	**	****	***
Chapter 7	**	**	***
Chapter 8	****	**	****
Chapter 9	**	***	*
Appendix A	***	**	*
Appendix B	***	**	*
Glossary	*	*	*

To help those who prefer an occasional dip into the technical parts of the book, rather than a concerted attack, we have appended a fairly large glossary that should help you through the more challenging sections.

1.6 Summary

It was not very long ago that the U.K. Postal System had a choice of post-boxes, one marked "national", the other marked "local." To use this system effectively, you had to know how it worked-the local letters would usually arrive more quickly but would be delayed if misdirected through the national routing scheme. Similar systems are implemented in the United States and elsewhere around the world.

The current picture of information networks is somewhat similar. The onus is on the user to know how the system works in order to make the best use of it. But this situation is changing. In this chapter, we looked at some of the disruptive influences, the most dramatic of which is the huge growth in long-distance transmission capacity.

We discussed how advances in technology and investments in infrastructure have combined to yield a "wired world." There is now enough capacity along some of the world's major communication routes to carry the entire planet's voice traffic several times over, with sufficient capacity to transfer the largest of national archives in parallel. With plentiful (and, consequently, cheap) long-distance transmission capacity, we argued that at least the technology barriers to effective high-speed communication-the wide area bottleneck-is removed. This enables the local network to extend into the wide area to provide a *total* area network.

When total area networks supplant the familiar LANs and WANs we have today, it seems likely that all manner of innovation will result. However, it would be premature to speculate about this. The main conclusion of this opening chapter is that there are new opportunities for the LAN. With Ethernet being the dominant technology for local area networking, it seems sensible to assess what it has to offer and how it might be applied.

Selected Bibliography

Cook, P., *Towards Local Globalisation*, London: UCL Press, 1993.

Davies, D., C. Sandbanks, and A. Rudge, *Telecommunications After AD2000*, London: Chapman & Hall, 1993.

John, R. R., *Spreading the News: The American Postal System from Franklin to Morse*, New York: Harvard University Press, 1995.

Muller, N., *Desktop Encyclopaedia of the Internet*, Norwood, MA: Artech House, 1999.

Naisbitt, J., and P. Aburdene, *Reinventing the Corporation*, London: Futura Books, 1986.

Ohmae, K., *The Borderless World*, London: Harper Collins, 1992.

Shaw, J., *Telecommunications Deregulation and the Information Economy*, Norwood, MA: Artech House, 2001.

Tomlinson, C., *Telecommunications*, London: Addison Wesley, 2001.

Chapter 2: Ethernet-The Story So Far

Overview

Let's start at the very beginning, a very good place to start.
--Rodgers and Hammerstein

The historical roots of the Ethernet are rather interesting-and not a little exotic. It all began, according to popular belief, with the Aloha network, which was developed at the University of Hawaii. This was one of the earliest examples of a LAN and is generally considered to be the ancestor of all of the shared media networks that followed.

In 1973, researchers at Xerox's Palo Alto Research Center (PARC), principally Robert Metcalfe and David Boggs, were looking for some way to interconnect the Xerox Altos machines, early graphical personal computers developed at PARC. Their solution, which was based on ALOHA, was a 2.94-Mbps system that controlled the access to a shared network by sensing for contention and backing off until the way was clear.

This was such a successful solution that it came to be used to connect more than 100 personal workstations together over a 1-Km cable run. Metcalfe decided to call the technology "Ethernet", choosing the word "Ether" to describe the physical medium-at that time, a cable-that carries bits to all nodes in the network.

[Figure 2.1](#) shows a drawing by Metcalfe that illustrates the simplicity of the first Ethernet.

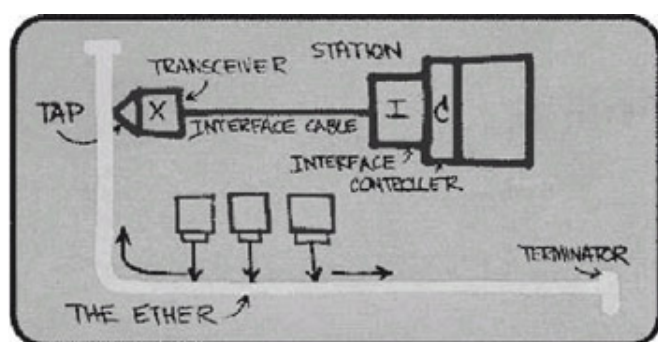


Figure 2.1: The original Ethernet "napkin diagram."

The curious clock rate for the first Ethernet was no more than a matter of convenience. A 2.94-MHz timing signal was derived from the system clock on the Altos machines, the Xerox PCs of the day. Very soon, the main proponents of this new technology, DEC, Intel, and Xerox (collectively referred to as DIX), came up with the now familiar 10-Mbps technology, and the first of the commercial Ethernets was specified by DIX to operate at this speed.

By 1985, Ethernet technology was embodied in a set of standards from the Institute of Electrical and Electronic Engineers (IEEE), known as IEEE 802.3. Originally, two types of coaxial cables were used, called thick Ethernet and thin Ethernet. Later, copper unshielded twisted pair (UTP), which was widely used for telephone connections, was added to the repertoire.

In the 1980s, 10 Mbps was a lot of bandwidth. As computing technology improved and distributed, networked applications have grown in sophistication and popularity, and network bandwidth requirements have increased quite dramatically. The simple and elegant design of the Ethernet allowed it to respond to this growth in demand. In 1995, the IEEE adopted the 802.3u Fast Ethernet standard. This was significant because it specified an operating speed of 100 Mbps but could readily interoperate with the original 10-Mbps Ethernet.

An important principle, that of backward compatibility, had been established, and this principle has been maintained throughout Ethernet's history. The Gigabit Ethernet specification appeared less than 5 years after the Fast Ethernet standard was published, and the 10-Gigabit version has already been

specified. The track record for Ethernet technology indicates that it can keep pace with the most demanding requirements.

Before looking at Ethernet's technology, it is worth dwelling on the reasons for its preeminence. Back in the 1980s, there were a number of technologies that did much the same as Ethernet. Some adopted the Ethernet's backing off approach for resolving access contention to a shared medium, others used a token to indicate which station was allowed to use the network. Most of these technologies worked well and, as is always the case with competing technologies, each could claim some advantage over the competition.

The beauty of the Ethernet, and probably its key differentiator, was and still remains its simplicity. Simplicity has enabled Ethernet to evolve while maintaining compatibility with earlier installations. This, in turn, has kept the whole-life cost of deploying Ethernet attractively small. As the importance of networks grew at the end of the twentieth century, the buying public chose their winner.

A final point to make is that Ethernet appeared at just the right time. A famous prediction by Robert Metcalfe, now known as "Metcalfe's Law", was that *the value of a network expands exponentially as the number of users increases*. As the speed of Ethernet has increased, first to 10 Mbps, then to 100 Mbps, and more recently to 1 Gbps and beyond, more users have been accommodated, and the network has become a more central part of their world. Plentiful bandwidth at commodity pricing has also enabled application innovation, and Ethernet, more than any other technology, has provided bandwidth when demanded, if not "on demand." At the same time that the Ethernet has become more important, the cost of the technology has plummeted (mostly thanks to another law, from Gordon Moore of Intel, who predicted the cost of processing power will be reduced by half every 18 months).

Affordability, ubiquity, and compatibility have made Ethernet the dominant network technology in LANs. With the advent of Gigabit Ethernet, it is now making substantial inroads into metro area networks (MANs) and even WANs. Compounding this expansion into the MAN and WAN is the fact that the volume of data communications traffic has grown to the point where it outweighs circuit-switched voice traffic four to one. The telecommunications infrastructure, purposely built to carry voice, is clearly due for an overhaul when 80% of all network traffic is data.

2.1 Basic Concepts

As we have already indicated, Ethernet is admirably simple—it consists of little more than a communication medium (i.e., some form of cable, a specific wavelength of light, or particular radio frequency), a protocol that determines how you access the medium, an arbitration method to accommodate multiple, simultaneous users, and a format for the data transmitted over the medium. Of course, there is a considerable amount of practical detail that needs to be attached to these bare bones, but for now, let us look at these three fundamentals.

2.1.1 Packet Format

All the data carried over an Ethernet is carried in a packet that conforms to the format shown in [Figure 2.2](#). For all intents and purposes, this packet format defines Ethernet. It has persisted since the early days and provides commonality across all the various Ethernet flavors.




Figure 2.2: The basic Ethernet packet.

The packet preamble is normally generated by the Ethernet hardware, which also appends the frame check sequence or checksum, a redundant series of bits that ensures data integrity during transmission. Software is responsible for setting the destination and source addresses, as well as the data that is carried in the frame's payload.

At the start of every frame, there is the *preamble*, which is a series of alternating ones and zeroes that can be used by the Ethernet receiver to acquire bit synchronization. Next comes the start of *frame delimiter*, a series of alternating ones and zeroes that ends with two consecutive ones and is used to acquire byte alignment.

The address information follows. The first part of this is the *destination Ethernet address*, which indicates the address of the intended receiver. If this field is set to all ones, then the message is broadcast to all attached stations. The next piece of address information is the *source Ethernet address*. This is the globally unique Ethernet address of the sending station (that is, the unique identity of the PC, workstation, server, or whatever other device initiates communication).

There are two interpretations for the next field—it can denote either *message length* or *message type field*. The reason for this duality is that the IEEE standard for Ethernet is slightly different from the original, proprietary specification from Xerox. The latter did not need a length field because all of the vendor protocols that used it (XNS, DECnet, IPX, and IP) had their own length fields. However, the IEEE committee needed a standard that did not depend on the good behavior of other protocols, so they replaced the two-byte type field with a two-byte length field.

The reason the two definitions of this field can coexist is that Xerox had not assigned any upper layer protocol-type values below the decimal value of 1,500. Since the maximum length of an Ethernet frame is 1,500 bytes, all possible length values can be indicated without any conflict or overlap. Hence, any Ethernet frame with a type/length field less than 1,500 is in IEEE format (with length defined as a value between 64 and 1,500). Any frame in which the field's value is greater than 1,500 must follow the Xerox format (with predefined type values such as 0x800 for IP packets or 0x600 for DECnet).

The actual information sent over the network follows next in the *data field*. This part of the frame is

where an IP packet (or any other type of data) would be carried. The data field can be up to 1,500 bytes in length-and this sets the upper limit on the amount of data that can be transported inside any one frame. There is also a minimum frame size, so the data field is padded up to 46 bytes if needed.

There is good reason for specifying maximum and minimum frame sizes. If the frame is too long, it can block other users from getting fair access-other users continuously detect a potential collision and therefore back off from sending. If the frame length is too short, the last bit of it can leave the sender before the first bit has arrived at the recipient, thereby making it difficult to test when the network is free for use. With reference to [Figure 2.2](#), the shortest Ethernet frame is $6 + 6 + 2 + 46 + 4 = 64$ bytes and the longest frame is $6 + 6 + 2 + 1,500 + 4 = 1,518$ bytes.

The last field is a *frame check sequence*. This is a 32-bit cyclic redundancy check that operates on the whole frame (except for the preamble and itself). It serves to let the receiver of the frame test know whether any errors have occurred in its transmission.

2.1.2 Network Protocol

A *network protocol* is a standard that allows computers to communicate with each other across some form of link. The protocol needs to define how the computers interact with the network and then how they find the device they want to communicate with. A good protocol should also define how to handle damaged transmissions. IPX, TCP/IP, DECnet, AppleTalk, LAT, X.25, Netware/IPX, and NetBEUI are all well-known examples of network protocols.

We start by looking at how the computer gets on to the network in the first place. The precursor to Ethernet-ALOHA-had the simplest possible scheme for gaining access to a shared medium. When a terminal had something to send, it simply sent it, regardless of what anyone else was doing. If another computer tried to send at the same time, the two messages would collide, both would be corrupted, and everyone would have to back off for a while before starting again. It was only when you received an acknowledgement to your message that you could be sure that it had arrived safely and there had been no collision. Of course, you might have had to wait for the acknowledgment as it, too, could have suffered a collision. This was not an efficient way to control access, so a more sophisticated mechanism was sought.

Enter slotted ALOHA. Now transmissions were constrained to occur in frames, rather than at any time. This meant that collisions only occurred when two terminals chose to transmit in the same frame (rather than when there was any overlap at all between the two transmissions). This simple modification led to a significant improvement in throughput, allowing up to 35% of the channel capacity to be used, compared with a previous peak utilization of less than 20%.

Ethernet built on the lessons of ALOHA by enforcing the following regime:

- Terminals always listen before they send.
- They do not send when someone else is doing so.
- When they do send, they limit the amount of data that is transmitted.
- If they find themselves sending at the same time as someone else, they back off.

These are much the same set of rules that govern a (polite) telephone conference. The main addition to the ALOHA practice is to listen before sending. If the medium is free, the terminal can transmit immediately. If it is busy, it backs off for a random time before trying again.

Collisions still occur, of course, as it takes a finite time for frames to travel from one point of connection to another. Hence, two terminals may start to send, both believing the medium to be free, only to realize that someone else has already started to send at the same time.

The basic physics of transmission speed have an impact on the range of frame sizes that are specified for use on an Ethernet. If two stations are at opposite ends of the network-as shown in [Figure 2.3](#)-they need to know when the other one is transmitting. If the frame length is too short, then this is compromised. The listening station is unaware that the medium is busy, so it would feel free to send. A collision would then happen close by this station. Meanwhile, the station that had already sent the frame would be blissfully ignorant of what was going on at the other end of the network and would

have to wait to find out if there had been a collision.




Figure 2.3: Collisions on an Ethernet bus.

In order to remove the indeterminacy caused by frames in transit, a minimum frame size is needed. This is equal to twice the propagation time from one end of the network to the other. With multiple sections joined by repeaters, there can be up to 2.5 km between the two end stations. Hence, with signals traveling at a speed of about 2×10^8 m/sec, the "there and back" propagation time is just over 50 ms. At a bit rate of 10 Mbps, this equates to a minimum frame size of 64 bytes.

The mechanism for controlling access described above is known as *collision sense multiple access/collision detect* (CSMA/CD). It has proved a simple yet effective way of admitting traffic on to the network. Adjusting the back-off time after a collision and the maximum frame size that can be sent both help to ensure fair access and increase the achievable throughput of data across the network. In practice, CSMA/CD Ethernet networks could cope perfectly well with traffic loads that previously used up to 70% of available bandwidth.

Now that we have a solution to the main problem of accessing shared media, the issue of getting the right connection remains. This is where the address resolution protocol (ARP) for IP comes in handy.

All stations attached to an Ethernet have a *physical address*—a 48-bit (or 6-byte) identifier known as the hardware or media access control (MAC) address. When user A wants to communicate with user B at another station, it must know both the MAC address and the destination station's higher level protocol address (i.e., its IP address). To obtain station B's IP address, user A's computer composes a packet known as an ARP request, encapsulates this within an Ethernet frame, and "broadcasts" this to all stations on the Ethernet.

Typically, the ARP request contains the destination IP address of the intended host, since this is the most commonly recognized form of identification these days. However, ARP variants have been developed for nearly every protocol operated over Ethernet, including Novell/Netware, AppleTalk, and DECnet.

The hardware location of the addressee is not known to start with, so the hardware address field in the broadcast message is set to zero. The frame containing the ARP request is broadcast to all devices connected to the Ethernet by setting the destination Ethernet address to all ones (FF FF FF FF FF FF). All stations on an Ethernet must listen for and process frames having this *broadcast address* as the destination MAC address.

When the broadcast frame is received, the station (B) whose upper level address (i.e., its IP address) matches the address encoded in the ARP request is the only one that responds to the broadcasting station. It returns an ARP reply message containing its hardware (MAC) address, which is inserted in the field that was previously set to zero.

When computer A receives and processes the ARP response message, it will maintain association of the logical IP address of the intended recipient with its physical MAC address in what is called an ARP cache. Once this step is complete, IP packets can be sent between the two without further concern; a logical path now exists between the sender and receiver.

2.1.3 Communication Media

As it became popular, four main cable types were used to carry Ethernet frames. These were thick coaxial cables (known as 10Base-5), thin coaxial cables (10Base-2), twisted wire pair (10Base-T), as

well as over-radio frequencies and fiber optic (the initial specification for which was known as FOIRL). The IEEE 802.3 standard set rules for each cable type for the maximum permissible segment length and the maximum number of attachments on a single segment. Over the years, fiber has increased in popularity, primarily due to its high capacity, so we expand on fiber-based Ethernet in the next chapter. For now, here is an overview of the media used to support the earlier networks.

10Base-5, or thick Ethernet, was the original wiring type. The transceivers for *10Base-5* were fitted with nonintrusive connectors (commonly known as vampire taps). These taps pierced the cable sheathing to connect devices to the network. The cable was several times thicker than that of thin Ethernet, giving it the following advantages and disadvantages:

- Maximum segment length of 500m;
- Maximum number of repeaters between any two nodes on the network of four;
- Maximum of 100 nodes on any one thick Ethernet segment;
- Maximum of three thick Ethernet segments between any two end nodes;
- Need to terminate both ends of a coax segment with 50-Ω terminators;
- Attached devices placed at 2.5-m intervals along the segment.

The main attributes of thick Ethernet, or thickwire, were its range and the number of connections it enabled. In its heyday, it was generally used to create the backbone that joins a number of smaller network segments into one large network. Thickwire made an excellent backbone because it supported many nodes in a bus layout and the segment could be quite long. Thickwire was commonly run from workgroup to workgroup, where smaller, shorter-range networks could then be attached to the backbone using less expensive cabling. Over time, thickwire was replaced by less expensive thin wire—the expense of thickwire, coupled along with the expense of vampire taps, made it an increasingly unattractive option.

10Base-2, or thinnet, uses cable that resembles coaxial aerial cable externally but makes connections via T-connectors. As with thickwire, the following restrictions impact network design and installation:

- Maximum segment length is 180m;
- Maximum number of repeaters between any two nodes on the network is four;
- Maximum of 30 nodes can be placed on any one thin Ethernet segment;
- A maximum of three active thin Ethernet segments can exist in a series;
- Maximum length of any single-series path must not exceed three thin coax segments and two inter-repeater links;
- Both ends of a coax segment must be terminated with 50-Ω terminators;
- Devices should be placed at 0.5-m intervals along the segment.

Thin coax offers the advantages of thicknet's bus topology, with reduced cost and easier installation. However, it can support only 30 nodes per segment, and each node must be at least 1.5m apart.

10Base-T (also known as UTP) has rapidly grown in popularity since the early 1990s and would seem to be the wave of the future as far as local area cabling goes. The cable consists of four pairs of wire similar to telephone cable, both in appearance and end-connector type. It comes in a variety of grades, with level one being the lowest quality and level five being the highest. Unlike the previous two types of cabling, both of which present a shared bus to all devices, *10Base-T* is used to connect devices into a single point known as a hub. The advantage of this star layout is that a faulty connector affects only one device, rather than the whole network. *10Base-T*'s characteristics include the following:

- The maximum length is 100m.
- Each cable will be a point-to-point connection.
- Each cable will be a single segment (i.e., one cable per device).

- Each cable must be terminated with RJ45 connectors.
- The cable must have a resistance between 85Ω and 100Ω.
- The cable should have two twists per foot (this is to minimize interference).

With 10Base-T, level one and two cabling can only really be used for voice and low-speed transmissions (less than 5 Mbps). Level three can be used for data speeds up to 16 Mbps, while level four can handle speeds up to 20 Mbps. The highest specification of 10Base-T cable, level five, can handle speeds up to 100 Mbps, so it can also be used to carry Fast Ethernet.

FOIRL, or 10Base-FL optical fiber, is similar to twisted pair. It can handle 100-Mbps transmission speeds but is not affected by electrical emissions or electro-magnetic interference. The major advantage of fiber-optic cable is its 2-km maximum length. The disadvantage is the higher cost of cable and equipment.

For the sake of completeness, it is worth mentioning home phoneline networking (HPNA). This accommodates the simultaneous use of "any" category twisted pair in arbitrary tree topologies by telephony, DSL, and Ethernet.

2.1.4 Topology

As we have already inferred, Ethernet is deployed in two basic topologies, bus and star. The physical [topology](#) defines how a node (which could be any connected device, such as a computer, printer, or server) is connected to the network.

Abus topology consists of nodes connected together by a single long cable. Each node taps into the bus and directly communicates with all other nodes on the bus. The major advantage of this topology is its easy expansion, by adding extra taps, and the lack of a central point of failure. The major disadvantage is that any break in the cable will cause all nodes on the cable to lose their network connection. A *star topology* links exactly two nodes together on the network. A device called a hub is used as a collection point, where many of the connections come together. The major advantage is any single break only disables one host, and the major disadvantage is the added cost of a hub. In practice, the reliability of most hubs is very high and they are much less likely to fail than a heavily used strand of coaxial cable. Hence, the star topology gained favor with users seeking a highly dependable LAN.

As well as having a physical topology, Ethernets also have a logical topology that is inferred from the way signals flow over the set of media segments that make up the whole system.

Multiple Ethernet segments can be linked together to form a larger network using a signal amplifying and retiming device called a repeater. Through the use of repeaters, multiple segments can grow to form a *nonrooted branching tree*. This means that each media segment is an individual branch of the complete system. Even though the media segments may be physically connected in a star pattern, with multiple segments attached to a repeater, the logical topology is still that of a single Ethernet channel that carries signals to all stations.

A tree is a formal name for systems like this, and a typical network design actually ends up looking more like a complex concatenation of network segments of both topologies. On media segments that support multiple connections, such as coaxial Ethernet, you may install a repeater and a link to another segment at any point on the segment. Other types of segments known as link segments can only have one connection at each end.

A system of linked segments that may grow in any direction, and does not have a specific root segment, is termed *nonrooted*. Most importantly, segments must never be connected in a loop. Every segment in the system must have two ends, since the Ethernet system will not operate correctly in the presence of looped paths.

If we have several media segments linked with repeaters and connecting to stations, a signal sent from any one station travels over that station's segment and is repeated onto all other segments. This way, it is heard by all of the other stations over the single Ethernet channel.

The physical topology may include both bus cables and a star cable layout. In [Figure 2.4](#), we have three segments connected to a single repeater, laid out in the star physical topology.

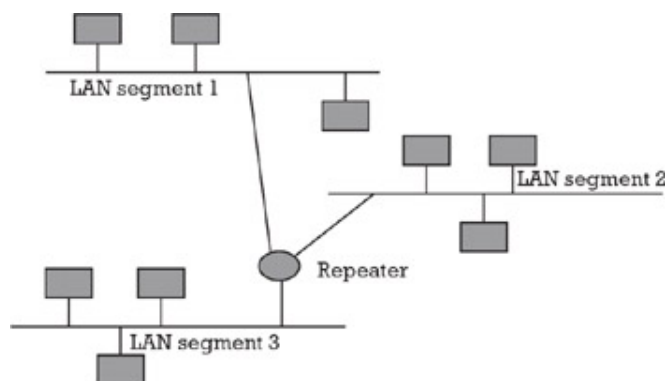


Figure 2.4: A typical Ethernet system.

The point is that no matter how the media segments are physically connected together, there is one signal channel delivering frames over those segments to all stations on a given Ethernet system.

2.1.5 Network Components

We have introduced quite a few terms by stealth in this chapter, so now is a good time for a few definitions. Here are the main physical elements that are used to build a real Ethernet.

Hub-a central point where multiple cables come together. A hub usually allows 8, 16, or 64 nodes to connect to each other. If any single connection becomes disconnected or is having problems, the hub can partition it (that is, remove it from the network) and allow all other nodes to continue to communicate. Ethernet hubs are necessary in star topologies such as 10Base-T.

A multiport twisted-pair hub allows several separate segments to be joined into one network. One end of the point-to-point link is attached to the hub and the other is attached to the computer. If the hub is attached to a backbone, then all computers at the end of the twisted-pair segments can communicate with all the hosts on the backbone. The number and type of hubs in any one collision domain is limited by the Ethernet rules. These repeater rules are discussed in more detail later.

An important fact to note about hubs is that they only allow users to share one Ethernet. A network of hubs/repeaters is termed a *shared Ethernet*, meaning that all stations of the network contend for transmission of data onto a single network or *collision domain*. This means that individual members of a shared network will only get a percentage of the available network bandwidth. The number and type of hubs in any one collision domain for 10 Mbps Ethernet is limited by the rules given in [Table 2.1](#).

Table 2.1: Planning Rules for Different Types of Ethernet

Network Type	Max Nodes Per Segment	Max Distance Per Segment
10Base-T	2	100m
10Base-2	30	185m
10Base-5	100	500m

Transceivers-also known as media attachment units (MAUs) and used to connect nodes to the Ethernet medium. For coaxial cable, these would be "vampire taps" or *BNC connectors* that directly access the bus. With a star topology, the transceiver allows the attachment of 10Base-T cable on one side, and the connection of a device via a 15-pin D-shell connector, known as an *application user interface (AUI)*, on the other. The user would connect the AUI connection to the computer and the twisted pair to the network media.

Repeaters-used to connect two or more Ethernet segments of any given media type. They can be

used to extend a segment beyond its maximum length or maximum number of nodes by restoring signal quality and timing. Repeaters can also be used to connect segments consisting of different media types together into one larger segment. It is worth noting that a repeater counts as a node on every segment to which it is attached. Repeater delay is very significant in practice, so much so that there are two types-type 1 (which is slower) and type 2 (which is faster).

Bridges-connect separate Ethernets together. They map the Ethernet addresses of the nodes residing on each separate network segment and then allow only the necessary traffic to pass through the bridge. A bridge can also filter out certain traffic and prevent it from passing through. When an Ethernet frame (as distinguished from a packet-the information received from a higher layer, carried in the frame's payload) is received by the bridge, the bridge determines the destination and source segments. If the segments are the same, the frame is dropped (or, more euphemistically, filtered) and if the segments are different, the frame is forwarded to the proper segment. In addition, bridges prevent all bad or misaligned frames from spreading by not forwarding them.

Bridges are called "store-and-forward" devices because they look at the whole Ethernet frame before making their filtering or forwarding decisions. Filtering frames, and regenerating forwarded frames, enables bridging technology to split a network into separate collision domains. This allows for greater distances and more repeaters to be used in the total network design.

Most bridges are self-learning; they determine the user Ethernet addresses on the segment by building a table as frames are passed through the network. This self-learning capability, however, dramatically raises the potential of network loops in networks that have many bridges. A loop presents conflicting information on which segment a specific address is located and forces the device to forward all traffic. The spanning tree algorithm is a software standard (found in the IEEE 802.1d specification) for describing how switches and bridges can communicate without creating network loops.

Switch-a bridge that can connect more than two segments together. The idea behind a switch is that it removes all unnecessary traffic from each segment by only forwarding the traffic that is relevant to that segment. This provides better performance on the network as a whole. Switches provide the same "hub" functionality in a hub-and-spoke topology as its predecessor, the nonswitching hub hardware component.

Routers-work in a manner similar to switches and bridges in that they filter out network traffic. Rather than using frame addresses, they do this by filtering a specific protocol. For instance, an IP router can divide a network into various subnets so that only traffic destined for particular IP addresses can pass between segments. The price paid for this type of intelligent forwarding and filtering is usually calculated in network speed, because protocol filtering usually takes more time than packet filtering.

An important function that routers perform when used on an Ethernet is the segmentation of collision and broadcast domains. In doing this, the router operates on the Ethernet frames, rather than the packets they carry.

Closely related to the router is a hybrid piece of equipment, the brouter (shorthand for bridge/router), which can operate in both bridge and routing mode, even doing so simultaneously for different protocols (i.e., routing IP packets and bridging Ethernet frames).

Network interface cards-commonly referred to as NICs, are used to connect a computer to a network. The NIC provides a physical connection between the networking cable and the computer's internal bus. Different computers have different bus architectures (e.g., PCI bus master slots are most commonly found on Pentium PCs). NICs come in three basic varieties: 8 bit, 16 bit, and 32 bit. The larger the number of bits that can be transferred to the NIC, the faster the NIC can transfer data to the network cable. The NIC is a very low-cost item (\$20-\$30 dollars for 10/100 Mbps). Firewire and Universal Serial Bus (USB) are emerging as successors to PCI.

Most NIC adapters comply with plug-n-play specifications. On these systems, the NIC is automatically configured without user intervention, while on non-plug-n-play systems, configuration is done manually through a setup program and/or switches.

2.2 Ethernet in Action

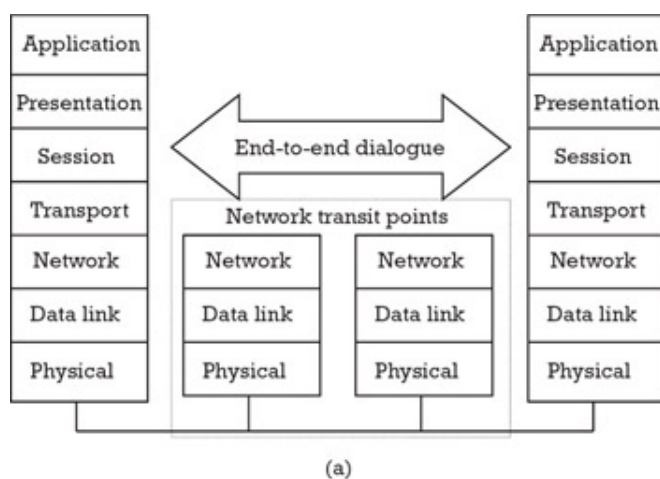
By now, we have quite a reasonable picture of the Ethernet. We have seen the components used to deploy it, we know how it can be configured, and we learned what sort of media can be used. We even know how the media is accessed and how information is exchanged.

One more aspect of the Ethernet needs to be covered before the picture is complete. This section describes the logical structure of communication used with the Ethernet.

2.2.1 It Always Ends in Tiers

Virtually all communication systems segment their activities in some way. This is convenient, as it allows low-level aspects, such as coding, to be treated separately from higher-level concerns, such as end-to-end session management. It also provides a reference for equipment suppliers, which helps them produce equipment that interoperates with others.

Perhaps the best-known example of segmenting communication activities is the International Standards Organization (ISO) 7-layer model and the IP architecture. Both define (as you might expect) the separate concerns in setting up a communication link. With the ISO model, seven separate layers are defined. The first three are all about the transmission of data from A to B while, at levels four and above, the concern is with end-to-end communication. The Internet architecture is somewhat simpler with only four layers. These architectures are illustrated in [Figures 2.5\(a\)](#) and [2.5\(b\)](#), respectively.



Applications	FTP, SMTP, Telnet, HTTP File transfer, Web browsing
Transport	TCP, UDP End-to-end delivery
Internet	IP Point-to-point packet routing
Network access	Ethernet Basic connection to the Internet

(b)

Figure 2.5: (a) The ISO layered models of communications, and (b) the Internet layered models of communications.

In concept, there is an interface between layers, with the higher ones passing information down to the

lower ones until the physical layer is reached, whereupon actual communication takes place. At the same time that information is flowing down our segmented communicating "stack", communication takes place between peer levels in different stacks. Hence, one application will "talk" to another and both of them will be supported by services provided by the other layers, as shown by the arrow between two communicating entities in [Figure 2.5\(a\)](#).

The layers within the ISO model are:

Physical-concerned with transmitting bits over a communication channel. The main issue is to ensure that, when one stack sends a "1", it is received as a "1" and not a "0." Examples of specifications that sit within this layer are RS232-C, for electrical data signal characteristics.

Data link-which builds up a frame for transmission via the physical layer. The main concern is to ensure that error-free frames can be transmitted between any two points in the network. The ISO 3309 high-level data link control (HDLC) standard is an example of data link protocol standards.

Network-which is concerned with getting an end-to-end connection, so issues such as routing and congestion control all lie here. The IP that carries much of the Internet traffic would fit here.

Transport- which is the first real end-to-end layer and assures end-to-end transmission. Two (exemplary) transport layer protocols are the Internet's TCP and user datagram protocol (UDP). The former is connection oriented, the latter connectionless.

Session-aims to ensure that different machines can establish communication; for instance, allowing two computers to complete a file transfer. The ITU-T X215 standard defines a connection-oriented session layer protocol.

Presentation-as the name suggests, is responsible for encoding data from a computer's internal format so that it is suitable for transmission. Hence, it has to deal with compression and decompression.

Application-the top of the stack, which contains the communication applications that use the services of the lower layers, an example being the ITU-T X400 message handling system. This is not the application itself, but the bits within it that provide the basic communication primitives.

The ISO model is meant to be technology independent. It states how the communication systems should be structured, not how each part should be built. In principle, this leaves suppliers to provide components that fulfill a specific function or sit in commonly recognized parts of a communicating system. Hence, the model is widely applicable but rather conceptual.

The Internet architecture, which readily maps to the ISO model, is more specific and directly reflects the way in which the Internet is structured. In this instance, each layer has the following purpose:

4. The *application layer* defines the application software, its processes, and the protocol it uses to convey its data to the communications protocol stack. In the case of e-mail, the protocol it uses is SMTP. E-mail applications wrap up messages with start and end markers and attach header information about where the mail is from and to whom the mail is to be sent. This is passed to the layer below to be sent on its way, much like putting a letter in an envelope, writing an address on the front, affixing the proper postage, and dropping it into a postbox.
3. The *transport layer* wraps up the application layer message in its own data that defines the application that is sending it and the application to receive it. These are known as the source and destination ports. It will also add data to specify the overall length of the message and a number, the checksum, to use to check if any of the data it is carrying has been corrupted. This is the layer in which both TCP and UDP reside. UDP is generally used to convey small messages of a request-response nature within single packets, and TCP is used to convey larger messages within a byte stream where some delivery assurances are needed.

The reasoning behind this difference is that, for small messages, the overhead of creating connections and ensuring reliable delivery is greater than the work of retransmitting the entire message. To this end, TCP will attach further information to the message passed from the application to ensure that the reliable connection is maintained during the transmission and that the segments of the byte stream all arrive at their destination. Continuing the earlier post office analogy, this is like having a set of numbered boxes being dispatched from a shipping office that

looks after the customer's accounts (and, hence, ensures that they get all of their boxes when they should).

2. The *Internet layer* provides the most important function of the TCP/IP stack. It structures the data into packets, known as datagrams, moves the datagrams between the network access layer and the transport layer, routes the datagrams from source to destination addresses, and performs any necessary fragmentation and reassembly of datagrams. The Internet layer wraps up the transport layer data in its own data, which includes the length of each datagram and the source and destination addresses (the IP addresses) that specify the network and host of the source and destination. This layer is akin to the internal workings of a post office-how letters are bundled into bags, how delivery trucks are routed and scheduled, which commercial subcontractors are used, and so on).
1. The *network access layer* is perhaps the least discussed of all the layers since the protocols within it are generally specific to a particular hardware technology for the delivery of data. Therefore, there are many protocols, one or more for each physical network implementation. The role of the network access layer is to ensure the correct transmission of IP datagrams across the physical medium and map IP addresses to the physical addresses used by the network. In the "real world", this layer provides the planes, boats, and trains that get packages from A to B.

As with the ISO model, received packets are "peeled" as they ascend the stack, each layer removing the layer-specific data they put there by the layer below. The concept of wrapping up data, layer by layer, in this way is referred to as encapsulation.

By the time that information reaches an Ethernet, it is almost fully clothed in all of its layer wrappings. The last step is to add the last layer or so. In the models mentioned, Ethernet can be positioned in the two lowest layers. We now look at this in some detail.

2.2.2 The Ethernet Protocol Stack

The ISO model is quite comprehensive. It has to be, because it aims to describe a wide range of communicating systems. The situation with Ethernet is considerably more straightforward and is restricted to layers 1 and 2, as previously defined. This is illustrated in [Figure 2.6](#).

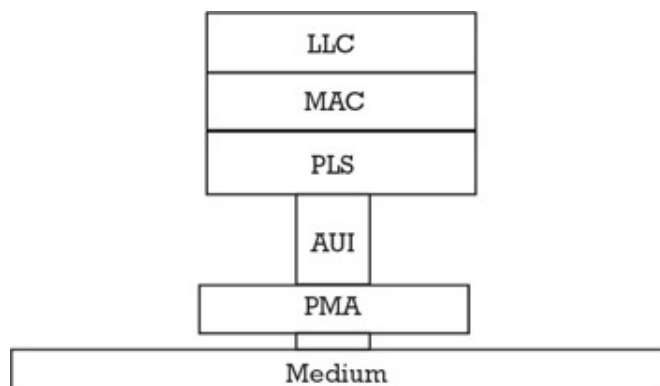


Figure 2.6: The Ethernet protocol stack.

In the diagram, the physical layer is split into two parts, with an interface between them. The physical medium attachment (PMA), as the name suggests, denotes the connector used to hook a device onto the LAN. The physical signaling sublayer (PLS) reports the state of the medium (i.e., idle, busy) to the MAC. The interface that sits between the PMA and the PLS is known as the attachment unit interface (AUI). It usually takes the form of a 15-pin D connector.

It should be noted that the data link activities defined by ISO are split into two areas here. The MAC defines how the physical medium is accessed, and the logical link control (LLC) provides a consistent interface to higher level protocols (e.g., IP) and the underlying network.

2.2.2.1 LLC

One of the key functions of the LLC is to make a broadcast network appear to the network layer as a set of point-to-point links. In doing this, it provides the following three distinct classes of service:

- Acknowledged connectionless service, which offers a mechanism where the delivery of a frame is acknowledged without a connection being set up. The main use of this is with real-time applications that require acknowledgement but cannot tolerate the delay caused by setting up a connection.
- Connection-oriented service, where a virtual circuit is established between two end points. This allows users on the LAN to set up and clear a point-to-point connection.
- Unacknowledged connectionless service, which is a datagram service that supports only the sending and receiving of frames. Because it is simple, it is the easiest service class to implement and, because end-to-end error and flow control are often provided by higher layers, it is useful in practice.

In providing these three classes of service, the LLC renders the MAC and the physical implementation of the LAN transparent to the higher layers of protocol. As a result, the underlying network can evolve (e.g., get faster, use different media) without requiring any changes to the applications that users access over it.

The LLC uses a two-byte address and a one-byte control field to provide the classes of service mentioned (see [Figure 2.7](#)).

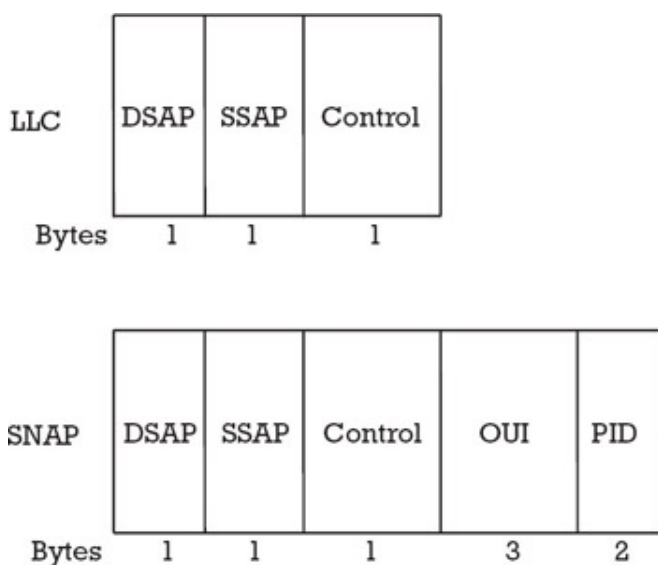


Figure 2.7: The LLC and SNAP formats.

The first of the address bytes is a destination service access point and the second a source service access point. The limited signaling capacity allocated to the LLC has led to an alternative, known as the Sub-Network Access Protocol (SNAP). With SNAP, the three LLC bytes are set to predefined values and five new bytes are inserted. The first three of these [the organizationally unique identifiers (OUI)] are used to identify an organization (e.g., Novell), the following two (the PID), identify a specific protocol (e.g., Netware/IPX).

2.2.2.2 MAC

The other part of the data link layer is provided by the MAC. The MAC handles the following:

- Data encapsulation from higher levels;
- Frame transmission;
- Frame reception;
- Data decapsulation to lower levels.

The MAC does not know, or care, about the physical layer being used. Neither does it care about the operating speed of the network. It serves to get frames on and off the network and can operate either full duplex (send and receive at the same time) or half duplex (send and receive in turn). The latter was the original mode of operation, but when there is no contention for the medium (as is the case with a 10Base-T star network), there is no need to restrict transmission. Full-duplex operation achieves higher throughput, as there is no collision penalty, and it offers greater reach, as there is no collision domain, so the medium, not the protocol used, is the limiter.

One of the key aspects of the MAC is that it holds the addresses of all stations connected to the network. The MAC address, illustrated in [Figure 2.8](#), is divided into four parts. The first two bits indicate whether the frame is unicast (zero) or multicast (one) and is universally (zero) or locally (one) managed. The third field is predefined by the IEEE and is called an OUI—a unique 22-bit OUI is assigned to each organization that wishes to build Ethernet. The organization, in turn, completes the 48-bit addresses using its assigned OUI. The full 48-bit address is known as the physical address, hardware address, or MAC address. Each MAC address is globally unique to a device (and so usually comes as a hard-coded identifier).




Figure 2.8: The structure of the MAC address.

Having a unique 48-bit address preassigned to each Ethernet interface when it is manufactured vastly simplifies the network's setup and operation. For one thing, preassigned addresses keep you from getting involved in administering the addresses for different groups using the network. And if you have ever tried to get different work groups at a large site to cooperate and voluntarily obey the same set of rules, you can appreciate what an advantage this can be.

As each Ethernet frame is sent onto the shared medium, all of the attached devices look at the first 48-bit field of the frame, which contains the destination address. The interfaces compare the destination address of the frame with their own address. The Ethernet interface with the same address as the destination address in the frame (already determined using ARP, as described in [Section 2.1](#)) will read in the entire frame and deliver it to the networking software running on that computer. All other network interfaces will stop reading the frame when they discover that the destination address does not match their own address.

The separation of concerns evident in the early Ethernet structure has served it well over the years. It has provided the basis for preserving some key aspects (such as the frame structure) while changing the technology and increasing the speed of operation.

2.3 Fast Ethernet

The popularity and large installed base of 10-Mbps Ethernets made it a natural springboard for faster networking technologies and, as we indicated in [Section 2.2](#), the separation of concerns in the Ethernet structure enable it to evolve with minimal disruption. However, the course of standards does not always run true and, in the course of devising the Fast Ethernet, two competing standards, 100Base-T and 100VG-AnyLAN were created, but only one, 100Base-T, endured and grew to the ubiquity of its ancestor, 10Base-T. A bit of history follows.

The first standard is known as 100Base-T and is supported by companies such as 3Com, Intel, and Synoptics. It is purely an extension of the original 802.3 standard, retaining the CSMA/CD (listen before you send) protocol as its communication technology. The second, competing technology, known as 100VG-AnyLAN, was put forward by companies including IBM and Hewlett Packard. The technology behind AnyLAN was completely new; instead of using CSMA/CD, it uses an unrelated demand-priority protocol. We now explain how each works before comparing their performance.

2.3.1 100Base-T

This version of Fast Ethernet gained initial standards approval in the mid-1990s, and the major vendors involved in the product have formed a committee known as the Fast Ethernet Alliance. Its transmission is limited to cable runs of 250m, before a device such as a bridge or router is needed to regenerate the signal. This meets the needs of most businesses but could prove to be a limiting factor in some cases, hence the interest in an alternative.

The main components of the 100Base-T standard are as follows:

Physical layer-The main alternatives for media that can be used are 100Base-TX (unshielded twisted pair), 100Base-T4 (shielded twisted pair), and 100Base-FX (fiber optic).

With 100Base-TX, the medium is two pairs of high-quality category 5 unshielded, balanced twisted-pair cable. The coding scheme specified is 4B5B, which is more complex than the Manchester encoding used for 10-Mbps Ethernet but has fewer interference problems. This standard has become popular due to its close compatibility with the 10Base-T Ethernet standard. Its two-pair configuration is the same as 10Base-T and the RJ45 connectors, and its 100-m segment length is also common. This means that the upgrade path from 10Base-T to 100Base-TX is an easy one that usually requires little more than the installation of higher quality cable. Because 100Base-TX uses exactly the same protocol as its predecessors, it supports full- and half-duplex mode.

The 100Base-T4 specification is intended for networks operating over four pairs of lower grade category UTP cabling. It achieves its speed by dividing a 100-Mbps data stream into three 33-Mbps streams. These three streams are sent over three of the four pairs, with the remaining one reserved for collision detection. No data is sent on the fourth pair; instead, the hub uses it to signal a workstation when a collision occurs. Splitting the data stream across the wires helps ensure the signal integrity, but full-duplex operation cannot be supported. The coding scheme for 100Base-T4 is known as 8B6T.

Optical fiber is used with 100Base-FX. Two fibers are used—one for transmission and one for reception. The same coding scheme is used as with 100Base-TX. The major advantage of this variant of Fast Ethernet is its long range, which makes it ideal for constructing large networks.

In addition to these three, there is now an option to use two pairs of category 3 UTP (100Base-T2). This was not part of the original specification and has not been a widely deployed practice. That said, practically speaking, today's NICs are entirely insensitive to what the cabling is, provided the distance is suitable for the medium. Home offices in particular can and do use cheap category 3 UTP, as the required distances here are often less than 10m.

One important addition to the physical layer (but only for 100Base-TX) is autonegotiation. This is a mechanism that seizes control of a link when a connection is established between two devices, determines the network capability of each device, and optimizes the dialogue between them. Autonegotiation works by using a feature built into 10Base-T called link integrity test. This was originally used as a media check when no frames were being sent but has been converted so that autonegotiation-enabled devices can exchange setup parameters.

Media independent interface (MII)-a new sublayer in the stack that is located above the physical layer. It defines a standard interface between the MAC layer (below) and any of the three physical layers. It performs essentially the same function as the AUI in the 10Base-T system (i.e., to decouple the MAC from the underlying medium). The main difference between the MII and the AUI is that data passes through the latter one bit at a time but through the former as 4-bit parallel bytelets. The reason for this is simple physics-driving a bit stream with a 100-MHz clock would severely limit range. The 4-bit option allows the clock speed to be reduced to a more forgiving 25 MHz.

MAC-as already described, now located above the new MII layer and based on the CSMA/CD protocol previously used in standard 10-Mbps Ethernet.

Figure 2.9 depicts the protocol layer arrangement for 100Base-T. The figure includes the 10Base-T stack for comparison.

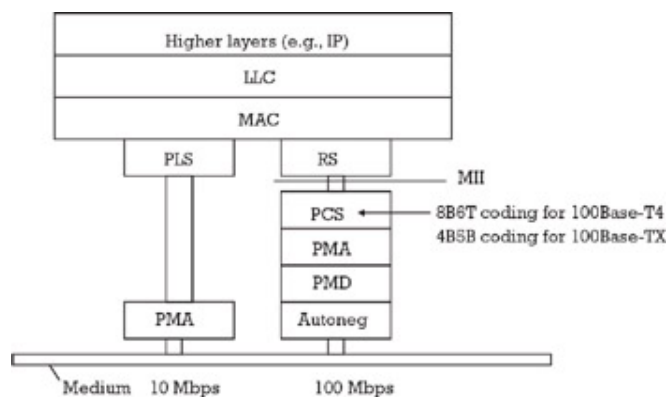


Figure 2.9: The Fast Ethernet layers.

The addition of the MII is not the only change with Fast Ethernet. There is also considerable extension to the structure of the physical layer, including the following sublayers:

- A physical coding sublayer (PCS) is introduced to provide a uniform interface to the reconciliation sublayer (RS) for all of the possible physical media.
- A physical medium dependent (PMD) sublayer is added to map the physical medium to the PCS. This layer defines the physical layer signaling used for various media.
- The RS maps status signals from the physical layer (PHY) into PLS primitives so that they can be understood by any MAC sublayer.

Despite the apparent wholesale change in the PHY layer, there is still compatibility between basic and Fast Ethernet, with the RS ensuring that the same MAC is used in both cases.

2.3.2100VG-AnyLAN

This variety of Fast Ethernet runs on four separate pairs of copper wire, distributing data equally among each pair. It allows for cable runs of up to 4,000m, without the need for devices to regenerate the signal. It has been given its own standards specification by the IEEE, called 802.12, and the major vendors involved have grouped together to form the 100VGAnyLAN Forum.

100VG-AnyLAN supports both Ethernet and token-ring frames, making it very flexible. However, the method of transmission it uses for these frames differs greatly from any introduced earlier. Rather than using CSMA/CD or token-passing methods, it uses a deterministic, demand-priority access protocol. A device sends a control tone to the hub on the network when it wishes to transmit frames, and the hub polls all connected devices and then grants access based on priority.

Supporters of 100VG-AnyLAN claim that it is a better protocol for LANs that transmit great amounts of time-sensitive data, such as fullmotion video, and other multimedia files. It could be argued that 100VG-AnyLAN is not actually Ethernet at all, due to the fact that it uses a different transmission protocol. Network purists claim that as CSMA/CD is the whole basis of Ethernet, 100VG does not

qualify to call itself Ethernet. That said, it is widely viewed as a viable addition to the clan, so we will now take a brief look at how it works.

Demand priority. Instead of having each controller check for a busy network, 100VG-AnyLAN uses a demand priority scheme. Demand priority works like a traffic signal; the hub logic determines which controller has access to the network. The hub polls each controller to determine if that controller has data to transmit and then allows transmission in port order.

For example, if there is a request waiting on port one and port three, all requests are of equal priority. The hub begins by servicing port one. Next, the hub checks to make sure that no new requests have come in for port two, the next in line. Assuming no requests have come in, the hub proceeds to service the request on port three.

If a request came in from ports two and four while the request at port three was being serviced, port four would be the next one serviced, and then the hub would start back at port one. This approach allows equal access to network media but priorities are also allowed and the network manager can set these from the hub.

In theory, every workstation could be set to high, which negates the purpose of priorities, but in a ideal network, only ports using applications such as video conferencing or multimedia would be set to high. These types of applications require frequent response in real time, so they need the priority set to high.

If port one has a normal priority request and ports two and three have high priority, the hub bypasses port one and services ports two and three first.

Of course, if after servicing port three, high-priority requests continue to come from port two and three, station one again is bypassed. Normal requests will continue to be bypassed until that request has been waiting for a few hundred milliseconds (depending upon the configuration of the hub timer).

Once the timer has expired:

- The hub completes the processing of the current high-priority request.
- The hub changes the priority of the request on port one from normal to high.
- The hub again polls all ports and determines, using the port order, which port should be processed next.

In the example previously mentioned, the hub would finish the request from port two, complete the outstanding request from port three, and then service the request waiting at port one. This handling of priorities guarantees that a maximum settable delay will not be exceeded.

Ideally, the network management software should be used to ensure that a large number of users are not all configured to send frames marked as high priority. Indeed, the network manager can configure each hub to treat frames from a particular port as normal, regardless of request priority assigned locally by the driver.

Hub layers. Up to three hierarchical (cascaded) layers of hubs are allowed in a 100VG-AnyLAN network. The root hub (first-level hub) controls the order in which requests should be serviced. When the root hub receives a request from another hub, it passes control of the network to the second-level hub. The second-level hub services its requests in port order. The root hub continues to process all requests from the second-level hub before it continues servicing the third port on its own hub.

2.3.3 Comparison of Fast Ethernet Alternatives

The two options described in [Section 2.3.2](#) take very different approaches and therefore have quite different characteristics. Here we list the main advantages and disadvantages of each. This is not intended as any sort of critique of either—the real aim is to point out the consequent operating idiosyncrasies that derive from design choices that have been taken in each case.

100Base-T is an inexpensive technology. It fits well with existing LAN technology and leverages existing expertise. Moreover, familiarity with Ethernet should enable users to incorporate the new technology easily into their existing networks.

However, there is a smaller network radius with 100Base-T networks than with 100VG-AnyLAN if repeater hubs are used instead of the (more expensive) switching hubs. Also, this is a shared bandwidth, nonpriority solution-so many end stations contend for the 100 Mbps on an equal (nonpriority) basis.

100VG-AnyLAN offers an easier migration path from token ring networks than 100Base-T technology because 100VG-AnyLAN is able to use token-ring (802.5) frame formats and is deterministic-the maximum time to send a frame on a network can be calculated.

However, it does not afford the same degree of backward compatibility with existing Ethernets that 100Base-T does, and the total throughput of 100VG-AnyLAN cannot exceed 100 Mbps. With 100Base-T and efficient switching hubs, it is possible to reach higher throughputs.

One final issue that should not be overlooked is enterprise administration-having two different technologies in an enterprise means you must have expertise in both, manage both using disparate management information bases (MIBs), and have instrumentation to analyze both. In addition, you cannot have one stock of common equipment. The upshot of this is that compatibility with the installed base makes 100Base-T the market choice.

Team LiB

[◀ PREVIOUS](#) [NEXT ▶](#)

2.4 Design

There are a number of design rules that both Ethernets and Fast Ethernets must follow if they are to function correctly. The maximum number of nodes, number of repeaters, and upper limit on segment distances are defined by the electrical and mechanical design properties of each type of Ethernet and Fast Ethernet media.

A network using repeaters, for instance, needs to function within the timing constraints that apply to an Ethernet. Although electrical signals on the media travel near the speed of light, it still takes a finite time for the signal to travel from one end of a large Ethernet to another. As a rule of thumb, the Ethernet standard assumes it will take roughly 50 ms for a signal to reach its destination.

Ethernet is subject to the "5-4-3" rule of repeater placement: the network can only have five segments connected; it can only use four repeaters; and of the five segments, only three can have users attached to them; the other two must be inter-repeater links.

If the design of the network violates these repeater and placement rules, then timing guidelines will not be met and a sending station will spuriously resend frames. This can lead to lost frames and excessive resent frames, which can slow network performance and create trouble for applications.

Fast Ethernet has modified repeater rules, since the minimum frame size takes less time to transmit than regular Ethernet, and the length of the network links allows for a smaller number of repeaters. In Fast Ethernet networks, there are two classes of repeaters. Class I repeaters have a latency of 0.7 ms or less and are limited to one repeater per network. Class II repeaters have a latency of 0.46 ms or less and are limited to two repeaters per network. [Table 2.2](#) lists the distance (diameter) characteristics for these types of Fast Ethernet repeater combinations.

Table 2.2: Distance Limits for Fast Ethernet

Fast Ethernet	Copper	Fiber
No repeaters	100m	412m [□]
One class I repeater	200m	272m
One class II repeater	200m	272m
Two class II repeaters	205m	228m

[□]In full-duplex mode, this is 2 km.

Both 100Base-TX and 100Base-T4 support a maximum of 100m of UTP from the hub to the workstation. A total of 205m of cable is allowed for the local end station to local hub to remote hub to remote end station connection. This is referred to as the network diameter.

To increase the network diameter, the use of switching hubs at key central locations is recommended. These devices join two or more separate networks, allowing network design criteria to be restored. Switches allow network designers to build large networks that function well. Each network connected via one of these devices is referred to as a separate collision domain in the overall network.

A switching hub allows 100m of cable between stations and hubs because it stores and forwards the frames, rather than repeating them along the wire. This store and forward works as if the frame came from the switching hub itself, although none of the frame information is changed (the originator's address is still stored, rather than the switching hub address).

If networks running different variants of Fast Ethernet are being joined, they must communicate through a hub that handles both types of 100Base-T or have bridges and/or routers that connect the two types of networks.

2.5 Standards

Throughout this chapter, we have been a bit lax with terminology, using Ethernet as a generic term instead of being more specific about which flavor of that technology we meant. To paraphrase Saki, "a little bit of inaccuracy saves a load of explanation," but before going on, we do need to get more specific, so this section defines all of the standards we need to know about.

At the beginning of the chapter, we explained that the early development of Ethernet was done by Xerox research, and the name "Ethernet" was a registered trademark of Xerox Corporation. The original technology was refined, and a second generation (called Ethernet II) was widely used during the 1980s. Ethernet from this period is often called DIX after its corporate sponsors Digital, Intel, and Xerox. As the holder of the trademark, Xerox established and published the standards.

Obviously, no technology could become an international standard for all sorts of equipment if a single U.S. corporation controlled the rules, so the IEEE was assigned the task of developing formal international standards. This remit was broader than just Ethernet—it covered all flavors of LAN technology.

As a result, the IEEE formed the 802 committee to look at Ethernet, token ring, fiber optic, and any other LAN technology. The objective of the project was not just to standardize each LAN individually, but also to establish rules that would be global to all types of LANs so that data could easily move from Ethernet to token ring and so on. The IEEE view of the network layers was, therefore, a little different than that shown in [Figures 2.6](#) and [2.9](#).

The way in which the IEEE 802 committee organized itself to meet its objectives is shown in [Figure 2.10](#). In the diagram, a set of options sits below a common LLC layer. Each of the options was developed under the auspices of a different part of the IEEE 802 committee—IEEE 802.3 looked after CSMA/CD, IEEE 802.4 after token bus, and IEEE 802.5 after token ring. Relevant to all of these was another body, IEEE 802.1, which looked at the management of networks.

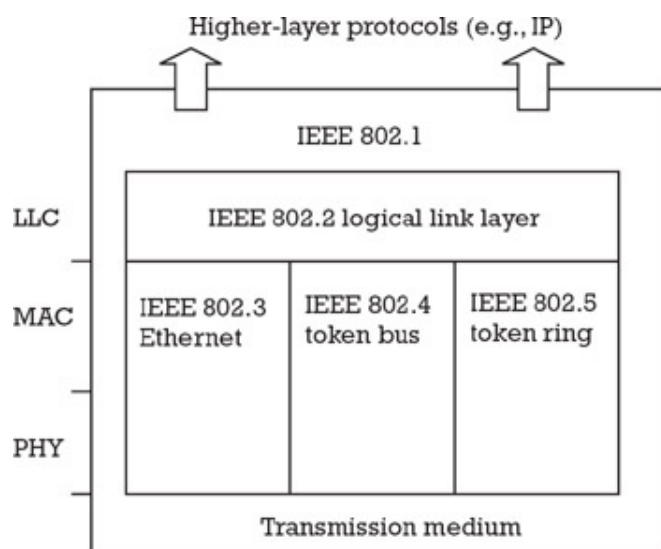


Figure 2.10: The IEEE view of the Ethernet layers.

The more rounded view of the IEEE 802 committee created some conflict with the existing practice under the old Xerox DIX system (e.g., the DIX type field was defined as a length field by the IEEE, for reasons explained earlier). However, one of the precepts of the IEEE standards was that old DIX messages and new IEEE 802 messages would have to coexist on the same LAN. And so, a little of the history of the Ethernet was perpetuated through the standardization process.

As it got down to business, the IEEE 802.3 group refined the specification for electrical connection to the Ethernet. All of the hardware vendors immediately adopted this specification and, to this day, all cards and other devices conform to this standard. However, the refinements required a change to the network architecture of all existing Ethernet users. Apple had to change its Ethertalk, and did so when

converting from Phase 1 to Phase 2 Appletalk. DEC had to change its DECnet. Novell added IEEE 802 as an option to its IPX, but supports both DIX and 802 message formats at the same time.

The TCP/IP protocol used by the Internet refused to change. The Internet Engineering Task Force (IETF) manages Internet standards, and they decided to stick with the old DIX message format indefinitely. This produced a deadlock between two standards organizations that has not yet been resolved.

IBM waited until the IEEE 802 committee released its standards and then it rigorously implemented the IEEE 802 rules for everything except TCP/IP, where the IETF rules take precedence. This means that NetBEUI (the format for NetBios on the LAN) and SNA obey the IEEE 802 conventions.

So, in some ways, we actually suffer from a surfeit of standards. The old DIX rules for message formats persist for some uses (Internet, DEC-net, some Novell). The newer IEEE 802 rules apply to other traffic (SNA, NetBEUI). Given this situation, it is difficult to talk both accurately and briefly about "Ethernet." For the purposes of communication, we have chosen to be a little permissive with the word. A key point, though, is that it pays to be somewhat more strict when the time comes to put theory in practice.

Notwithstanding that caveat, the IEEE 802.3 committee has produced (and continues to produce) a comprehensive set of specifications that covers all aspects of Ethernet. [Table 2.3](#) lists some of the main offerings.

Table 2.3: Main Specifications in the IEEE 802.3 Family

Reference	Contents	Description
802.3	The original standard for the Ethernet	10Base-T, 10Base-2, 10Base-5,
802.3u	Specification for Fast Ethernet	100Base-TX, 100Base FX, 100Base-4
802.3x	Flow control	
802.3z	Gigabit Ethernet	1,000Base-SX, 1,000Base-LX, 1,000Base-CX
802.3ab	Copper Gigabit Ethernet	
802.3ac	Frame tagging for virtual LAN Support	
802.3ad	Link aggregation	
802.3ae	10-Gigabit Ethernet	

It is interesting to note that the rate with which the IEEE 802.3 committee has produced new specifications has increased over recent years. The sound base, laid down back in the late 1970s and early 1980s, has enabled Ethernet technology to keep pace with the today's demands.

2.6 Summary

For years, Ethernet has been the networking technology of choice for most organizations, and it is the most popular physical layer technology for local area networking in use today. In fact, it is currently estimated that there are close to 100 million Ethernet users worldwide. Ethernet has become popular because it strikes a good balance between speed, cost, and ease of installation. These strong points, combined with wide acceptance in the computer marketplace and the ability to support virtually all popular network protocols, make it an ideal networking technology for most computer users today.

As Ethernet's popularity has risen, so have the demands placed on most Ethernet-based LANs. Powerful workstations and sophisticated networked applications are placing ever-increasing demands on the networks that connect them. These new demands are exceeding the capacity of the well-established 10Base-T technology.

Ethernet cannot afford to stand still-vendors need to offer ever-higher transmission speeds if they are to maintain their preeminent position. The Fast Ethernet standard (IEEE 802.3u) was the first major evolutionary step. This standard has raised the Ethernet speed limit from 10 to 100 Mbps, with only minimal changes to the existing cable structure.

This chapter explained the basic operating principle of Ethernet and illustrated its evolution from the early days through to Fast Ethernet. We introduced all of the component parts of an Ethernet system and showed how it operates.

Selected Bibliography

Metcalfe,R., and D.Boggs,"*Ethernet: Distributed Packet Switching Technology for Local Computer Networks*",*ACM*,Vol 9,No. 5,July 1976 (the original paper on the Ethernet).

Shotwell,R. (ed.) *Ethernet Sourcebook*,North-Holland,1985 (probably provides the most comprehensive treatment of basic Ethernet technology).

<http://www.iguest.com/~nmuller/fast.html> (contains a good introduction to 100Base-T: Fast Ethernet technology, from the Fast Ethernet Alliance).

<http://www.lantronix.com/htmlfiles/mrktg/catalog/et.htm> (contains an in-depth introduction to LANs from Lantronix).

Chapter 3: Gigabit Ethernets

Overview

There is more to life than increasing its speed.

—*Mahatma Gandhi*

One of the basic laws of network design is that you always underestimate the demand for increased capacity. Only a few years ago, the idea of an Ethernet running at gigabit speeds would have seemed a little excessive. But today's data-intensive applications, an increasing number of network users, and new methods of information delivery are driving an ever-increasing demand for more bandwidth.

Most existing networks struggle to cope with the demands of their users. The latest computers and servers are capable of processing a data throughput of up to a gigabit per second. But even when they connect through a Fast Ethernet, they find they are faced with a transfer rate of only a tenth of that required. So a bottleneck occurs. This bottleneck stimulates interest in networks that offer the sort of capacity that matches the capabilities of the devices that connect to them.

When a user strays from the LAN, the speed disparity between his application and the network is accentuated. Data warehousing, high-quality document transferring, and videoconferencing all demand more bandwidth than most local access and mobile services offer. Currently, LAN-like bandwidth remains expensive, and service areas are severely limited. One of the primary goals of modern data communications is not only to facilitate fast exchange between computing systems but also to do it over initially reasonable, and ultimately great, distances.

This chapter introduces Gigabit Ethernet. It explains the modifications to the basic Ethernet technology needed to achieve such high speeds. The details Gigabit Ethernet's physical and media access options are explained. In doing this, we point out the fundamental similarities between Gigabit Ethernet and its lower speed forebears.

The second part of this chapter deals with the latest step in the Ethernet evolution—the 10-Gigabit Ethernet. Again, the basic technical details are explained and the commonalities with, and advances beyond, antecedents pointed out. To complete our technical story, we give an overview of the Fiber Channel, one of the main technologies used in conjunction with Gigabit Ethernet.

3.1 Gigabit Ethernet

The first move toward a Gigabit Ethernet can be traced to the time that the IEEE 802.3 standard committee created the IEEE 802.3z Task Force. This committee was given the task to develop the standard that would address the need for a high-speed technology for both backbone networks and local access (that is, the connection from a user's site to a PoP on a network).

Like all other Ethernet technologies, the IEEE 802.3z Gigabit Ethernet standard is an extension of the base IEEE 802.3 standard. Gigabit Ethernet has a lot in common with its forebears with respect to MAC layer characteristics and framing, but has a physical layer and data link layer that enables it to operate at a considerably higher speed.

The specific design objectives for the Gigabit Ethernet specification given to the IEEE 802.3z task force included the following:

- Should offer 10 times the bandwidth of Fast Ethernet-1,000 Mbps; Must use the IEEE 802.3 Ethernet frame format;
- Should employ the same half-duplex and full-duplex MAC operation schemes as its predecessors;
- Should be backward compatible with 10 Mbps and 100 Mbps Ethernet technologies;
- Should support all existing network protocols used with the Ethernet family.

This list should come as no surprise (especially the first point!). All of the objectives put before the IEEE 802.3z task force are consistent with established Ethernet principles. In those instances where the demands imposed by a move to a much higher speed cannot be made using an established mechanism, the approach taken has been to specify some alternative way of operating. A case in point is the *carrier extension field* of the frame format, an extra field defined so that it can operate at gigabit speeds. This will be discussed in more detail later. First, we explain the physical and media access arrangements for Gigabit Ethernet.

3.1.1 Physical Layer

The physical layer of Gigabit Ethernet describes the physical properties of the communication media, as well as the electrical properties and interpretation of the exchanged signals. In order to achieve 1,000 Mbps throughput, a modified version of the ANSI X3.230 Fiber Channel standard physical layer is added to the established technologies already used with the Ethernet family.

There were several good reasons to adopt Fiber Channel. It was a proven technology, so reusing it would reduce the Gigabit Ethernet standard development time. It was also commercially mature, so products were already available (and reasonably priced).

The Fiber Channel technology (explained in more detail later in this chapter) uses long wavelength lasers to transmit data over fiber-optic cable. The following three types of media are specified in the IEEE 802.3z standard (collectively known as 1000Base-X standard):

- 1000Base-SX: Short wavelength laser transmitted over multimode fiber.
- 1000Base-LX: Long wavelength laser transmitted over both single mode and multimode fiber.
- 1000Base-CX: Short-haul copper shielded twisted pair (STP), sometimes known as twinax cable.

3.1.1.1 1000Base-X

ANSI Fiber Channel has been widely developed as an interconnection technology for the connection of workstations, supercomputers, storage devices, and peripherals. It has a four-layer architecture. The lowest two layers, FC-0 (interface and media) and FC-1 (encode/decode), are used in Gigabit Ethernet.

Approximate cabling distances that can be supported are given in [Table 3.1](#).

Table 3.1: Maximum Reach for Various Media Types

Type of Cable	Max Distance
Single-mode fiber	3km
Multimode fiber-850-nm laser	300m
Multimode fiber-1,300-nm laser	550m
Short-haul copper	25m

3.1.1.21000Base-T

The fourth media type specified for Gigabit Ethernet is known as 1000Base-T. This defines the use of long-haul copper UTP, which should allow a range of between 25m and 100m using four pairs of category 5 UTP. A separate committee from the 1000-Base-X standards, the IEEE 802.3ab task force, developed the 1000Base-T standard.

3.1.2MAC

The MAC layer of Gigabit Ethernet contains all of the capabilities that exist in other Ethernet technologies, along with a number of additional features and functions that older Ethernet technologies do not have. The main examples of the new features specific to Gigabit Ethernet operation are carrier extension and frame bursting.

The IEEE 802.3z MAC operation can be in either half-or full-duplex mode. A half-duplex channel can receive and transmit, but not at the same time. With full-duplex transmission, it is possible to transmit and receive data at the same time. In full-duplex mode, the Gigabit Ethernet MAC uses the IEEE 802.3x full-duplex specification, which includes the IEEE 802.3x flow control. In half-duplex mode, the Gigabit Ethernet MAC uses the long-established CSMA/CD access method.

The use of full-duplex transmission in Gigabit Ethernet increases the overall bandwidth from 1 to 2 Gbps for point-to-point links. It also increases the maximum transmission distances for the particular media. CSMA/CD is not used for media access in full-duplex Gigabit Ethernet because the use of full-duplex operation eliminates collisions on the wire. Full-duplex operation is best suited to backbone transmission and as access to high-speed servers.

In half-duplex mode, Gigabit Ethernet employs an enhanced version of the CSMA/CD protocol. As explained in [Chapter 2](#), when a frame is transmitted onto a network, it stands a chance of colliding with another frame transmitted at exactly the same time by a different station on the same network. The CSMA/CD access method requires that attached stations listen for traffic on the network and transmit frames only when there is no other traffic. Collisions can still occur when two end stations listen for traffic, hear none, and then transmit simultaneously. In this situation, both transmissions are damaged, and the stations must retransmit at some later time (see [Figure 3.1](#)). CSMA/CD allows the MAC layer to stop transmitting and retransmits the frame when the transmission medium is clear.

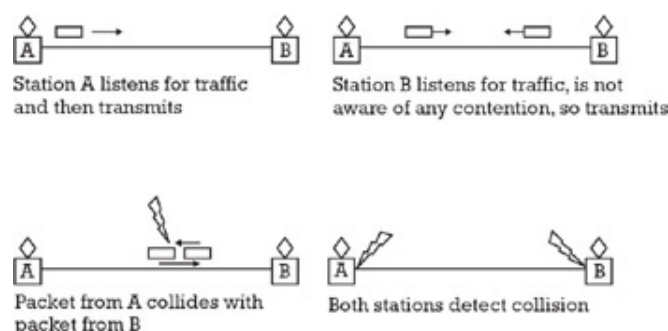


Figure 3.1: Collisions on a link.

The network must wait for a collision to abate before engaging in any further transactions. However,

the use of CSMA/CD access method in Gigabit Ethernet creates a problem, since CSMA/CD is very sensitive to frame length. As we have already seen, the minimum frame transmission time must be longer than the maximum round-trip propagation time of the LAN. If this is so, it ensures that a station will still be transmitting a frame when it is informed of a collision. The minimum time to detect a collision is the time it takes for the signal to propagate from one end of the station to the other. This minimum time is called the slot time.

A more useful metric to look at is the slot size, which is the number of bytes transmitted in one slot time (see [Figure 3.2](#)). Since the frame transmission time is inversely proportional to the data rate, increasing the speed decreases the frame transmission time. In the case of Gigabit Ethernet, increasing the speed by 100 from original Ethernet decreases the frame transmission time by a factor of 100.




Figure 3.2: Slot sizes.

The maximum cable length permitted in Ethernet is 2.5 km (with a maximum of four repeaters on any path). If the same frame sizes and cable lengths are maintained, then a station may transmit a frame too fast and not detect a collision at the other end of the cable. One of the following two remedies must be considered:

1. Keep the maximum cable length and increase the slot time (and, therefore, the minimum frame size).
2. Keep the slot time the same and decrease the maximum cable length.

In Fast Ethernet, the maximum cable length is reduced to only 100m, leaving the minimum frame size and slot time intact. With Gigabit Ethernet 10 times faster than Fast Ethernet, to maintain the same slot size, the maximum cable length would have to be reduced to about 10m, which is not very useful. So the first of our two options is the more sensible-Gigabit Ethernet uses an increased slot size of 512 bytes. However, in order to meet a basic remit of the IEEE 802.3z task force to maintain compatibility with Ethernet, the minimum frame size cannot be increased from 64 to 512 bytes; rather, the *carrier event* is extended through a technique known as carrier extension. This technique, illustrated in [Figure 3.3](#), is used to solve the timing problem associated with the use of CSMA/CD at gigabit rates.

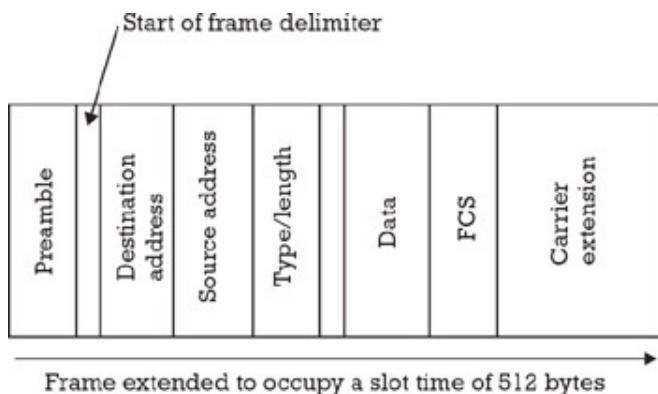


Figure 3.3: Carrier extension.

With carrier extension, the slot size is increased from 64 to 512 bytes. Hence, any frame that is smaller than 512 bytes must be padded. Padding is achieved using special extension symbols that cannot

occur in the payload. These carrier extended frames only exist in transit across the medium. As far as both sender and receiver are concerned, the frames are normal Ethernet packets. Thus, the frame check sequence (FCS) is calculated only on the original information (without extension symbols) and the extension symbols are removed before the FCS is checked by the receiver. In this way, even the LLC layer is not aware that carrier extension is being used.

Carrier extension is an admirably simple solution and provides a way of maintaining the IEEE 802.3 minimum and maximum frame sizes with meaningful cabling distances. But it does have a downside—it is very bandwidth inefficient. This is especially true with small packet sizes, where up to 448 bytes of padding may have to be added. This results in low throughput. In fact, for a large number of small packets, the throughput is only marginally better than Fast Ethernet.

And so an additional feature, known as frame bursting, has also been incorporated into the enhanced CSMA/CD scheme used by Gigabit Ethernet.

Frame bursting allows stations to send a number of short frames so that the full available bandwidth is used. When a station has a number of packets to transmit, the first packet is padded to the slot time, if necessary, using carrier extension. Subsequent packets are then transmitted back to back, with the minimum interpacket gap (IPG) until a burst timer (which counts up to 1,500 bytes) expires. This technique, illustrated in [Figure 3.4](#), substantially increases the throughput and allows users to fully exploit the native speed of Gigabit Ethernet.




Figure 3.4: Frame bursting.

When the MAC is ready to send a frame, it checks a *burst timer*. If this timer is not running, the medium is not busy, which means this must be the first frame being sent. The MAC starts the timer as it sends the first frame. If the frame is less than 512 bytes, it will include a carrier extension. Once this frame has been sent, the MAC signals the PHY layer to send a carrier extension instead of leaving the quiet time of the interframe gap. This allows the transmitting station to hold the medium until the next frame is sent, which happens right after the carrier extension is sent. When this next frame is sent, the burst timer is already running. This means that the medium must be clear to send. With no chance of a collision occurring, the packet can be sent without carrier extension. Packets continue to be sent until the burst timer expires.

The carrier extension symbol is placed between packets simply to delineate the frames. If this were not done, a receiving station (which may be a 10 or 100-Mbps Ethernet installation) would not understand the incoming data. The inclusion of the carrier extension symbol indicates where the interframe gap should be and so restores normal service.

Of course, with frames arriving at gigabit speed, there is a greater chance that the buffers set up to smooth packet arrival distribution will overflow. Therefore, Gigabit Ethernet needs to provide some form of flow control to avoid traffic congestion and overloading. This is similar to that specified for Fast Ethernet, but there are some significant differences. Perhaps the most important is that the flow control mechanism is specified for all physical media—including fiber. With Gigabit Ethernet, both flow control and autonegotiation are part of the PCS.

3.1.3 The Gigabit Ethernet Protocol Stack

The layers that make up the physical and media access layers of the Gigabit Ethernet are much the same as those we have already seen. The protocol architecture is shown in [Figure 3.5](#).




Figure 3.5: Gigabit Ethernet protocol stack.

In the diagram, the gigabit media independent interface (GMII) is the interface between the MAC and physical layers. GMII allows any of the available physical layers to be used with the same MAC layer. It is an extension of the MII used in Fast Ethernet and uses the same management interface as MII.

The GMII supports 10-, 100-, and 1,000-Mbps data rates. It provides separate 8-bit wide receive and transmit data paths, so it can support both full- and half-duplex modes of operation. The GMII also provides two media status signals. The first indicates the presence of the carrier and the second the absence of a collision.

The RS maps these status signals into PLS primitives so that they can be understood by any MAC sublayer. With the GMII, it is possible to connect various media types such as STP and UTP, as well as single-mode and multimode optical fiber, while using the same MAC controller. This arrangement is similar to the more familiar 10/100 Mbps NICs, where you can purchase NICs with multiple media connectors (BNC/ AUI/RJ45).

The GMII sits above three main sublayers—the PCS, the PMA, and the PMD. The purpose of each of these is explained in the following sections.

3.1.3.1 PCS

This sublayer of the GMII provides a uniform interface to the RS for all of the possible physical media. It uses an 8B/10B coding scheme just like Fiber Channel. In this type of coding, 10-bit "code groups" represent groups of 8 bits. Some code groups represent 8-bit data symbols and others are control symbols. An example of a control signal is the base-link code word used for flow control. These code words supplant the link pulses used with Fast Ethernet and allow flow control to be exercised over fiber links, the dominant medium for use with Gigabit Ethernet.

The PCS sublayer also generates the carrier sense and collision detect indications for half-duplex operation. It manages the autonegotiation process by which a NIC communicates with the network to determine the network speed (10, 100, or 1,000 Mbps) and mode of operation (half duplex or full duplex).

3.1.3.2 PMA

This sublayer provides a medium-independent means for the PCS to support various serial bit-oriented physical media. This layer takes the 10-bit code groups sent at 125 MHz by the PCS and renders them into serial format before transmission. In the reverse direction, it deserializes bits received from the medium back into code groups for delivery to the PCS.

3.1.3.3 PMD

This sublayer maps the physical medium to the PCS. This layer defines the physical layer signaling used for various media. The medium dependent interface (MDI), a component of the PMD, is the actual physical layer interface. This layer defines the actual physical attachment, such as connectors, for different media types.

3.1.4 Signal Encoding

So far, we have said very little about the way in which information sent over an Ethernet is encoded. With the very high operational speeds of Gigabit Ethernet, it is necessary to look quite closely at the options for signal coding. There are many such schemes available. Here, we look at some of those that would be appropriate for use with Ethernets.

In general, encoding schemes are often labeled with a very shorthand notation, such as "8B/10B." This label denotes a scheme that encodes 8 data bits into 10 code bits (or, put another way, "8 bits into 10 baud"). A code is considered to be more bandwidth-efficient when more data bits are encoded into fewer code bits, but this encoding usually requires more sophisticated algorithms, and therefore more complex encoders and decoders.

The 8B/10B coding scheme originally developed for Fiber Channel is particularly well suited for use with Gigabit Ethernet networks. Its main advantages are that it offers good error detection, along with reliable synchronization of bits and clock recovery. Both of these attributes are key in very high-speed networks.

Gigabit Ethernet employs the 8B/10B coding system in order to properly interpret the data at the receiver. In this system, the 8 data bits received from GMII are sent in 10 bits as a code group prior to transmission. The extra two bits included in the code require a signal transmission rate of 1.25 baud for every bit (that is, 10 baud divided by 8 bits) to transmit a net 1 Gbps of user data. The extra two bits are included to contain transfer control information such as start of packet, end of packet, and idle.

At the receiving end of a link, all symbols are valid if they contain five 1's and five 0's. If this is not the case, then some form of transmission error has occurred. The consistency of the 8B/10B format permits the generation of "DC-balanced" signal, so there is no net voltage on the transmission link. It also allows the receiver to perform bit synchronization easily and ensures the incoming bit stream has frequent transitions for performing clock recovery. IBM invented the 8B/10B coding scheme for low-cost devices and implementations. This means that, as well as having desirable technical features, the encoding/decoding algorithm is fairly simple and can be implemented in inexpensive hardware.

The main drawback of 8B/10B is the 25% overhead and consequent loss of bandwidth efficiency. This means that a line rate of 1.25 Gbps is needed to implement a usable transmission rate of 1 Gbps; for a serial implementation, this can be a significant disadvantage.

3.1.4.1 Scrambling

Scrambled encoding has its place in WAN applications such as SONET and SDH. It has virtually no overhead and is therefore more bandwidth efficient than 8B/10B. This allows a lower line rate to be used and can result in extended reach. The implementation of this code is also simple and can be done in hardware. However, its maximum run-length is nondeterministic, it has no guarantee of DC balance, and no built-in special characters. The cost of devices with this encoding scheme is typically higher, which could make this technique less attractive for use on a LAN unless an extended reach is really needed.

It could be argued that using a scrambled encoding scheme would enable information to make a seamless transition from a LAN to a WAN. However, there is a lack of physical compatibility between WAN technology and the IEEE 802.3z standard and, in practice, a compromise has to be made. Typically, a bridge would be inserted somewhere between LAN and WAN in order to handle the difference of line encoding techniques and other attributes. Given this, it is unlikely that scrambled encoding will become an end-to-end standard-8B/10B will likely be used in the LAN/MAN world, while scrambled encoding will likely be specified in WANs.

3.1.4.2 MB810

This coding technique has similar advantages and disadvantages as those in 8B/10B. The key benefit claimed by proponents of this code is that it offers a 50% improvement of bandwidth efficiency over the 8B/10B code. If this claim can be realized, the required line rate can be reduced in half, so only 0.5 Gbaud is needed for a serial implementation of 1-Gigabit Ethernet. This benefit sounds very promising. However, until the technologies supporting the MB810 encoding are proven, a leading role for this coding technique is not expected.

3.1.4.3 PAM-5

This coding technique was adopted in the 1000-Base-T standard. It employs multilevel amplitude signaling to increase the number of bits per baud. In twisted-pair lines, the PAM-5 encoding can achieve 2 bits per baud with a 3-dB coding gain, yielding a significantly lower line rate. For Fiber Channels, however, more work needs to be done to study the impact of the signal-to-noise ratio and nonlinearity penalty on multilevel amplitude signaling. The PAM-5 signaling may not offer very good reach when used with fiber, and the current devices may not support this signaling. Until these concerns are assuaged, the PAM-5 encoding is unlikely to compete with the 8B/10B and scrambled coding schemes.

3.1.4.4 16B/18B

The benefits of this code are similar to those of the 8B/10B, but with fewer overheads-12.5%, compared with 25% overhead for the 8B/10B. However, it is incompatible with the physical layer of the IEEE 802.3z standard and the SONET/SDH transmission standards.

3.1.4.5 Forward Error Correction

The forward error correction (FEC) code has extra bits appended to a frame that allow the correction of corrupted data at the receiving end, with no need for retransmission. This effectively lowers the bit error rate (BER) by a great amount but does incur the expense of some additional overhead; for example, with 6% overhead, an FEC code can achieve an effective 10^{-14} BER from an input with an actual 10^{-4} BER. The most popular FEC codes are known as broadcast channel (BCH) codes. These can be used in conjunction with 8B/10B or scrambled encoding. The FEC technique is useful in applications, such as long-haul networks, where coding advantage is strongly needed.

3.210-Gigabit Ethernet

The newest member of the Ethernet family is the 10-Gigabit Ethernet. Another task force, IEEE 802.3ae, has created the specification for this variant, in much the same way that Gigabit Ethernet was specified. Needless to say, there is considerable commonality with all of the previous Ethernet specifications. The overall structure is substantially the same as before, and changes are only introduced where needed to cope with the higher speed of operations.

The 10-Gigabit Ethernet is specified as a full-duplex, fiber-only technology. It does not need the carrier-sense multiple access with CSMA/CD that was defined for the half-duplex Ethernet technologies. Neither does it include copper as a medium-the reach would be so limited as to be impractical. In every other respect, though, 10-Gigabit Ethernet remains true to the original Ethernet model. A PHY device, which corresponds to layer 1 of the ISO model (see [Figure 2.5](#)), is still used to connect the media (optical in this instance) to the MAC layer, which corresponds to ISO layer 2.

As with Gigabit Ethernet, the 10-Gigabit Ethernet architecture further divides the PHY (layer 1) into a PMD and PCS.

Optical transceivers, for example, are PMDs. The PCS is made up of coding and a serializer or multiplexing functions. With the IEEE 802.3ae specification, two PHY types are defined-one for LANs and another for WANs. The various architectural components of 10-Gigabit Ethernet are shown in [Figure 3.6](#).

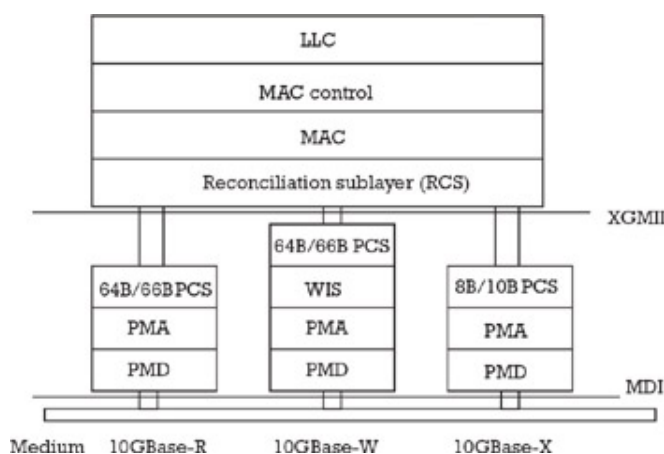


Figure 3.6: 10-Gigabit Ethernet protocol stack.

The WAN interface sublayer is introduced into the central option in [Figure 3.6](#) to provide the WAN PHY. The WAN PHY has an extended feature set added onto the functions of a LAN PHY. The two LAN PHYs are solely distinguished by the PCS.

3.2.1 Physical Layer Architecture

One of the first things that anyone familiar with Ethernet would notice about the 10-gigabit variant is that there are two options for implementing the physical layer-a serial solution and a parallel solution.

The serial solution uses one high-speed (10 Gbps) PCS/PMA/PMD circuit block, and the parallel solution uses multiple PCS/PMA/PMD circuit blocks, each of which operates at lower speed. The two solutions have different advantages and disadvantages. These are discussed in the following sections, as we introduce the working details in each instance.

3.2.1.1 Serial Implementation

In the serial implementation, illustrated in [Figure 3.7](#), one physical channel operates at 10 Gbps.




Figure 3.7: Serial physical layer implementation.

The operation is straightforward. For transmission, the reconciliation module passes the signals, corresponding to the MAC data, word-by-word to the PCS module. The PCS module then encodes the signals with a predefined coding technique and passes the encoded signal to the PMA module. The PMA module then serializes the encoded signals and passes the stream to the PMD module. The PMD module transmits the signal stream over the fiber at 10 Gbps. For receiving, the process is reversed.

The main attribute of the serial architecture is its simplicity-the transmitting and receiving operations are both straightforward. It does not require a complicated multiplexing/demultiplexing like the one needed in the parallel implementation. Thus, the timing jitter requirement can be more relaxed. Furthermore, serial only requires one fiber channel and one set of laser equipment, so the cost of implementation is minimized.

The downside to the serial architecture is the need for expensive high-speed logic circuits and technology. In order to reduce the transmission rate, one of the higher rate coding techniques introduced in [Section 3.1](#) (e.g., PAM-5) could be used. In such a case, only one lower cost laser unit may be needed. At publication, practical experience has yet to confirm the viability of this alternative.

3.2.1.2 Parallel Implementation

In the parallel implementation, illustrated in [Figure 3.8](#), there are multiple physical channels. Each subchannel may be implemented by using parallel cable using multiplexing.

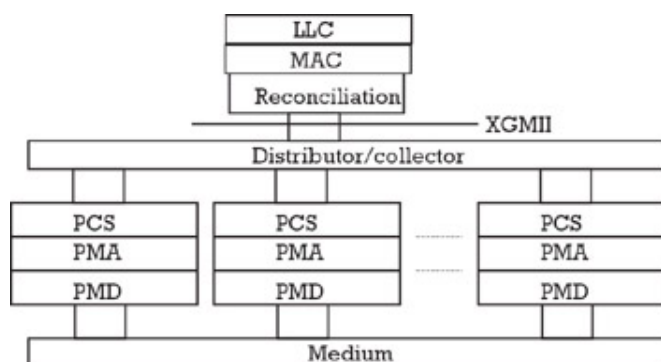


Figure 3.8: Parallel physical layer implementation.

For transmission, the distributor multiplexes the data (frames and idles) accepted from the MAC layer into a number of separate streams using a straightforward round-robin scheme. Each stream is given to each PCS module. Each PCS module encodes the received stream and passes it to each PMA module for serialization. After serialization, each PMD module transmits each serialized data stream at a fractional bit rate, as determined by the level of parallelism being used. For receiving, the reverse process is performed.

The main advantage of the parallel implementation is that the operating rate in the PCS/PMA modules

is reduced, which enables cheaper devices such as standard CMOS technology to be used. The disadvantages are the need for a distributor/collector module that may be sensitive to timing jitters and the use of multiple sets of logic circuits and laser equipment. Two techniques are used to achieve multiple channels: parallel cabling and wavelength division multiplexing (WDM).

For the parallel solution to achieve a total rate of 10 Gbps, 10×1 (or 10×1.25 Gbps to accommodate the coding scheme) fibers may be used. The availability of a 1-Gigabit Ethernet solution makes this a viable upgrade path. However, the cost of this solution is 10 times that of the 1-Gigabit Ethernet. This violates the economy of scale one expects from 10-Gigabit Ethernet, making it a somewhat less-than-attractive solution. However, as 10-Gbps equipment becomes more widely available, only four parallel fibers or one single fiber would be needed. This technique will be very attractive for short-haul (less than 200m) applications, where the cost of expensive optical multiplexing equipment may outweigh the cost of fiber cabling. However, the parallel cabling solution does not apply to the existing infrastructure.

3.2.1.3PMDs

In view of the extended range of applications for the 10-Gigabit Ethernet, the IEEE 802.3ae task force has specified the following number of PMDs:

- A 850-nm serial PMD specified to achieve a 65-m objective over multimode fiber.
- A 1,310-nm serial PMD to meet a 2-km and 10-km single-mode fiber objective.
- A 1,550-nm serial PMD to meet (or even exceed) a 40-km single-mode fiber objective.

The third bullet item—a PMD that works over a 40-km range—reflects the successful deployment of Gigabit Ethernet solutions in metropolitan and long-distance applications.

In addition, the task force has selected the following two versions of the WDM PMD:

- A 1,310-nm version over single-mode fiber with a target distance of 10 km.
- A 1,310-nm PMD to meet a target of 300m over installed multimode fiber.

These last two PMDs are included to meet the 10-Gigabit Ethernet aim of supporting LANs, MANs, and WANS. [Table 3.2](#) summarizes the various PMD options specified by the 802.3ae task force.

Table 3.2: Target Distances Specified by the IEEE 802.3ae Task Force

Optical Source and Fiber Type	Transmission	Target Distance
850-nm multimode	Serial	65m
1,310-nm multimode	WDM	300m
1,310-nm single mode	WDM	10 km
1,310-nm single mode	Serial	10 km
1,550-nm single mode	Serial	40 km

3.2.1.4Physical Layer

Both of the physical layers specified for 10-Gigabit Ethernet—the LAN PHY and the WAN PHY—will operate over common PMDs and, therefore, will support the same distances. It is only the PCS that distinguishes the two PHYs.

The 10-gigabit LAN PHY is intended to support all existing Gigabit Ethernet applications at 10 times the bandwidth in a more cost-effective manner—we will discuss this in more detail later. Over time, it is anticipated that the LAN PHY will be used in pure optical switching environments that extend over all WANs. However, for compatibility with existing WAN technology, the 10-Gigabit Ethernet WAN PHY supports connections to SONET/SDH circuit-switched telephony transmission equipment.

The WAN PHY differs from the LAN PHY by including a simplified SONET/SDH framer in the WAN

interface sublayer (WIS). Because the line rate of SONET OC-192 (and its SDH equivalent, STM-64) is very close to 10 Gbps, it is relatively simple to implement a MAC that can operate with a LAN PHY at 10 Gbps or with a WAN PHY transmission rate of approximately 9.29 Gbps. (See [Section 3.2.2](#))

To keep the implementation cost of the WAN PHY at an acceptable level, the IEEE 802.3ae task force wisely rejected full conformance to SONET/SDH. Instead, the WAN PHY is defined to be a cost-effective mechanism that uses common Ethernet PMDs to access SONET/SDH. The benefit of this is that it enables attachment of packet-based IP and Ethernet routers and switches to the SONET/SDH infrastructure. Hence, an Ethernet user can transmit information across the WAN backbone over SONET/SDH.

It is also important to note that Ethernet remains an asynchronous link protocol. As in every Ethernet network, 10-Gigabit Ethernet's timing and synchronization must be maintained within each character in the bit stream of data. The receiving hub, switch, or router may retime and resynchronize the data. In contrast, synchronous protocols, such as SONET and SDH, require that each device share the same system clock to avoid timing drift between transmission and reception equipment and subsequent increases in network errors where timed delivery is critical.

The WAN PHY allows the attachment of data equipment, such as switches or routers, to a SONET/SDH network and, thus, accommodates simple extension of Ethernet links over these networks. Two routers will behave as though they are directly attached to each other over a single Ethernet link. Since no bridges or store-and-forward buffer devices are required between them, all of the IP traffic management systems for differentiated services and quality of service (QoS) [e.g., differentiated services (DiffServ), multiprotocol label switching (MPLS)] operate over the extended 10-Gigabit Ethernet link connecting the two routers. In order to simplify management of extended 10-Gigabit Ethernet links, the WAN PHY provides most of the SONET/SDH management information. This allows a network manager to view the Ethernet WAN PHY links as though it was a SONET/SDH link. It is therefore possible to do end-to-end performance monitoring and fault isolation on the entire network, including the 10-Gigabit Ethernet WAN link, from the SONET/SDH management station. The SONET/SDH management information is provided by the WIS, which also includes the SONET/SDH framer. The WIS operates between the PCS and serial PMD layers common to the LAN PHY.

3.2.1.5 Chip Interface

One of the technical innovations introduced by the IEEE 802.3ae task force is an interface called the XAUI (pronounced "Zowie"). The "AU" portion of the name is taken from the Ethernet attachment unit interface. The "X" part is meant to represent the Roman numeral for 10 and implies 10 Gbps. The XAUI is designed as an interface extender, and the interface, which it extends, is the XGMII, the 10-gigabit MII. Its place in the overall 10-Gigabit Ethernet picture is shown in [Figure 3.9](#).

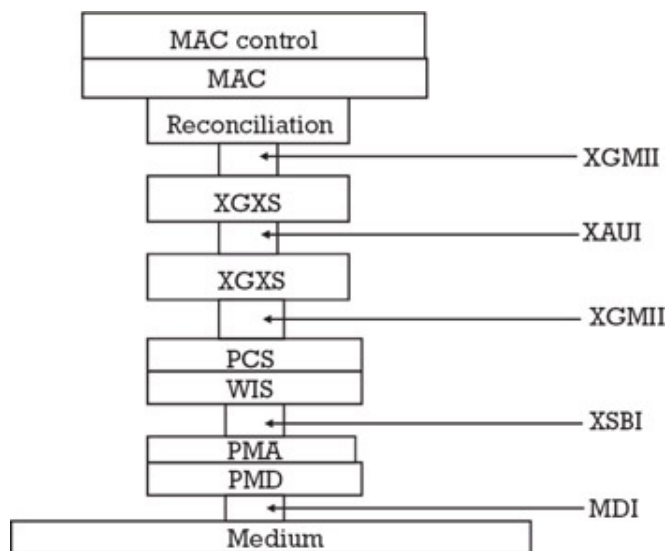


Figure 3.9: Where the XAUI fits.

For completeness, [Figure 3.9](#) also shows where other interface points fit (the MDI and a newly

introduced break point to allow access to the WIS).

The XGMII is a 74-signal wide interface (with 32-bit data paths for each transmit and receive) that may be used to attach the Ethernet MAC to its PHY. The XAUI may be used in place of, or to extend, the XGMII in chip-to-chip applications typical of most Ethernet MAC to PHY interconnects. The XAUI is a low pin count, self-clocked serial bus that is a direct descendent of the Gigabit Ethernet 1000Base-X PHY.

The XAUI interface speed is 2.5 times that of 1000Base-X. This means that, by using four serial lanes, the 4-bit XAUI interface is capable of supporting the tenfold increase in data throughput required by 10-Gigabit Ethernet. The XAUI employs the same robust 8B/10B encoding used with 1000Base-X to provide a high level of signal integrity through the copper media that is typical for chip-to-chip printed circuit board traces. Additional benefits of the XAUI technology include its inherently low electromagnetic interference (EMI), compensation for multibit bus skew-allowing significantly longer distance chip-to-chip links-error detection, and fault isolation capabilities.

The XAUI interface is becoming widely available, and its equivalence to the technology employed in other key industry standards, notably 10-Gigabit Fiber Channel, assures the lowest possible cost.

3.2.2MAC

The MAC layer of 10-Gigabit Ethernet is very similar to the MAC layer of previous Ethernet technologies. It uses the same Ethernet address and frame formats, but it does not support the half-duplex mode. It supports data rates of less than 10 Gbps using a pacing mechanism for rate adaptation and flow controls. The following sections consider these features one at a time.

3.2.2.1Full-Duplex Only

In previous Ethernet standards, there were two modes of operation: half-duplex and full-duplex modes. The half-duplex mode has been defined since the original version of Ethernet. In this mode, data are transmitted using the popular CSMA/CD protocol on a shared medium. Its simplicity contributed to the early success of the Ethernet standard. This mode of operation is so famous that many wrongly associate the CSMA/CD protocol with standard Ethernet operation. Efficiency and distance limitation are the main disadvantages of the half-duplex mode. In this mode, the link distance is limited by the minimum MAC frame size. This restriction reduces the efficiency drastically for high-rate transmission. For instance, the carrier extension technique is used to ensure the minimum frame size of 512 bytes in Gigabit Ethernet to achieve a reasonable link distance.

At the transmission rate of 10 Gbps, the half-duplex mode is not an attractive option and no market would realistically exist for the half-duplex operation at this rate of transmission, as most of the links at 10 Gbps are point-to-point over optical fibers. In this case, the full-duplex operation would be the preferred option, so the standard for 10-Gigabit Ethernet specifies only the full-duplex operation.

In full-duplex operation, there is no contention. The MAC layer entity can transmit whenever it wants, provided that its peer is ready to receive. The distance of the link is limited by the characteristic of the physical medium and devices, power budgets, and modulation. In this case, a desired topology can be achieved with the use of switches or distributed buffers.

3.2.2.2MAC Frame Format

One of the primary targets during the developing 10-Gigabit Ethernet standard was to use the same MAC frame format as specified in preceding Ethernet standards. This allows a seamless integration of the 10-Gigabit Ethernet with existing Ethernet networks. There is no need for fragmentation and reassembling or address translation and this, in its turn, implies faster switching. Since only full-duplex operation is used, the link distance does not affect the MAC frame size. The minimum MAC frame size will be made equal to 64 octets, as specified in previous Ethernet standards. Carrier extension is not needed.

3.2.2.3Data Rate

Although contrary to many people's idea of the obvious, the specification of the data rate for 10-Gigabit Ethernet was no simple task. Most of the LAN-oriented IT community wanted the data rate to be 10

Gbps so that a 10-Gigabit Ethernet switch can support exactly 10 1-Gigabit Ethernet ports. The telecommunications community, on the other hand, wanted the rate to be 9.584640 Gbps so that it equates exactly to the established OC-192 standard used for transmission systems.

The solution is to support both rates. This is done by specifying the data rate at 10 Gbps and then using the pacing mechanism described in [Section 3.2.2.4](#) to accommodate the (slightly) slower data rates. What might be an issue in this solution is that it can require a device with a large buffer to bridge the two rates. It should be noted that if the data rate is specified at 9.584640 Gbps, it is not possible to support the 10-Gbps data rate.

3.2.2.4 Pacing Mechanism

The pacing mechanism allows the MAC layer to support transmission rates, for instance, of 1 Gbps or 10 Gbps on LANs, as well as a rate of 9.584640 Gbps for transmission across a wide area. To achieve this, the MAC layer entity has the ability to pause data transmission for an appropriate period of time to provide a flow control or rate adaptation. There are two techniques for pacing mechanism. The first is the word-by-word hold mechanism, and the second is the interframe Gap (IFG) stretch technique. In the word-by-word technique, the MAC layer entity pauses to send a 32-bit word of data for a prespecified period of time upon request from the physical layer. In the IFG stretch technique, the interframe gap is extended for a predefined period of time with or without a request from the physical layer. The main disadvantage of the IFG stretch technique is that a large data buffer is required because the algorithm operates between frames. The word-by-word hold mechanism is preferred due to its main advantages of being able to support any of the encoding techniques. In addition, it does not need a large data buffer to hold multiple MAC frames, and so the buffer size is independent of link speed.

3.2.3 The 10-Gigabit Media Independent Interface

The 10-Gigabit MII (10GMII) provides the interface between the MAC layer and the physical layer. It allows the MAC layer to support various physical layer variations.

The TX_word_hold line is provided to support word-oriented pacing mechanism. The 32-bit data paths are provided for transmit and receive functions, each with four control bits (one per byte). The control bit is set to "1" for delimiters and special characters and "0" for data. Delimiters and special characters are determined from the 8-bit data value when the control bit is set to "0." The delimiter and special characters include the following:

- *IDLE*, which is signaled during the interpacket gap and when there is no data to send;
- *SOP*, which is signaled at the start of each packet;
- *EOP*, which is signaled at the end of each packet;
- *ERROR*, which is signaled when an error is detected in the received signal or when an error needs to be put to the transmitted signal.

These delimiters and special characters enable a proper synchronization for multiplexing and demultiplexing operations. It should be noted that the interface could also be scaled in speed and width. For example, an 8-bit data path with 1 control bit may be used at four times faster than clock rate. In this way, the delimiter and special characters remain unchanged. Thus, it can support both serial and parallel implementations of the PCS.

3.3 Fiber Channel

Fiber Channel has been designed as a universal standard. It is currently specified to support a wide range of popular applications, including the Small Computer Systems Interface (SCSI), ATM adaptation layer (AAL), and IP, as well as IEEE 802.3. It aims to combine the best of both channel and network data communication.

A channel is a direct or switched point-to-point connection. One benefit of a channel is that it is mostly hardware-intensive for speed and efficiency of data transport. This is compared to the higher overhead of network data transfer, which is slower because it is software-intensive. The benefit of networks is that they can handle a greater range of tasks because they operate in an environment of unanticipated connections.

The downside of network data transfer is minimized by providing a way to transfer data from a buffer at the source device to a buffer at the destination device. Fiber Channel does not need to know what the data is or how it is formatted. What the individual protocols do with the data before or after they are placed in the buffer is independent of the function of Fiber Channel. All a Fiber Channel port has to do is to manage a simple point-to-point connection between itself and the fabric.

Fiber Channel is currently specified at rates up to and including 2 Gbps, although higher data rates are planned.

Fiber Channel will allow simultaneous transmission of different protocols over a single optical-fiber pair and it can allow a number of existing services, such as network, point-to-point, and peripheral interfaces, to be accessed over a single medium using the same hardware connection.

The structure of Fiber Channel is defined as a multilayered stack of functional levels, not unlike those used to represent network protocols, although not mapping directly to OSI layers. The layers of the Fiber Channel standard define the physical media and transmission rates, encoding scheme, framing protocol and flow control, common service, and the upper level applications interfaces.

FC-0, the lowest layer, specifies the physical link in the system, including the media, transmitters, receivers, and connectors that can be used. This also includes electrical and optical characteristics, transmission rates, and other physical components of the standard. In line with its design aims, the Fiber Channel physical level has been designed so that it can be used with a large number of technologies and meets the widest range of system requirements. Indeed, an end-to-end route can use different link technologies for increased performance and decreased cost, while at the same time, systems integrators can tailor an installation to meet the specific needs of their customers.

FC-0 also specifies the Open Fiber Control (OFC) system, which is a safety system used to control the optical power level of laser data links in which an open fiber condition occurs. This feature is required because the optical power levels in this kind of system exceed the limits defined by laser safety standards. Whenever this open fiber condition occurs in the link from the sending port, the receiver port detects it and pulses its laser at a low duty cycle within the laser safety requirements. The receiver at the other port detects the pulsing signal and itself sends pulses within the specified laser safety range. If the open fiber condition is restored, receivers from both ports receive the pulsing signals, and this results in a double handshaking procedure that restores normal transmission after a couple of seconds.

FC-1 defines all transmission protocols, including serial encoding and decoding rules, special characters, and error control. Every 8 bits of data are encoded into a 10-bit transmission character, as explained in [Section 3.2](#). A transmission word is composed of four contiguous transmission characters. The transmission code is DC-balanced, and the transmission character is used to ensure that clock recovery is possible by having enough transitions present in the serial bit stream.

Character conversion is accomplished by taking an unencoded information byte that is made up of 8 information bits logically labeled and a control character. The first 5 bits are converted into a decimal representation, xx (5 bits have $2^5 - 1 = 31$ values, so no more than two decimal digits are required for this). The remaining bits are similarly converted into a single-digit decimal value, y . The control character Z is designated to be D for data-type or K for special-type. The resulting combination of this information forms a name in the form Zxy that represents a valid transmission character. After

transmission, the D-type transmission characters are decoded into one of 256 8-bit combinations. Any K-type transmission characters are used for protocol management functions. All other codes besides these types are invalid.

The physical layer uses a running disparity (RD). This is a binary number that is calculated based on the number of 0's and 1's in the two sub-blocks. The first sub-block is the first 6 bits of the transmission character and the second sub-block is the last 4 bits of the transmission character. A new RD is calculated at the transmitter and receiver, and if the RD values are not the same, a disparity violation condition is indicated.

Team LiB

[← PREVIOUS](#) [NEXT →](#)

3.4 Summary

The rate at which Ethernet has moved forward over the last few years is quite remarkable. It is only a few years ago that 10 Mbps was the state of the art in local-area networking. We saw in [Chapter 2](#) how the Fast Ethernet evolved. In this chapter, we described two more major advances—first the Gigabit Ethernet and then the 10-Gigabit Ethernet.

In some ways, this progression has been very straightforward. A common structure has persisted throughout, with the changes for each variant confined to those areas where a higher speed of operation makes it necessary. The uniformity of Ethernet technology in moving from 10 Mbps to 10 Gbps remains a testimony to the effective separation of concerns in its architecture.

Although focused on the way in which Gigabit Ethernets are implemented, this chapter has not ignored the wider range of applications that are now within its reach. With data rates comparable to WAN transmission technology and a reach that comfortably covers metropolitan areas, the Ethernet has evolved from a local area solution into a *total* area solution. And, having laid out the technical basics in this chapter, we can now move on to consider what part the Ethernet might come to play in the total area network.

Team LiB

◀ PREVIOUS

NEXT ▶

Selected Bibliography

Most of the reference material for Gigabit Ethernet can be found on the WWW. Some of the most useful links are to the various alliances that look after the development of a particular Ethernet type, notably the following:

<http://www.gigabit-ethernet.org>, the Gigabit Ethernet Alliance.

<http://www.10gea.org>, the 10-Gigabit Ethernet Alliance.

Both sites contain good and up-to-date information on their respective topics, along with links to related areas.

Team LiB

◀ PREVIOUS

NEXT ▶

Chapter 4: Wireless Ethernet

Overview

Each slight variation, if useful, is preserved by natural selection.
--Charles Darwin

Before moving on from the technology of the Ethernet to consider what sort of networks and services can be built on it, we describe one more variant on the basic theme: wireless networking. In doing this, we buck the trend established over the last couple of chapters by moving down in speed, rather than up. The reason for this is that our focus in this chapter is on the wireless version of the Ethernet. As we will see, what this member of the family, specified in the IEEE 802.11 standard, lacks in speed, it more than makes up for in mobility.

After many years of being something of a forgotten country cousin, the wireless Ethernet has at last established a significant niche in the LAN market. By and large, it acts as an adjunct to traditional wired LANs, satisfying mobility, relocation, and ad hoc networking requirements, as well as providing a way to cover locations that are difficult to wire and where wiring would be a lost investment (e.g., in a temporary or short-term leased office). It is also used to quickly establish LANs and is increasingly important in business continuity planning, as well as becoming the preferred alternative for home and small-office networks.

As the demand for mobile services grows, so might the importance of the wireless Ethernet. In [Chapter 7](#), we consider the market prospects for 802.11, with particular regard to its position against the mobile service offered by the third generation mobile network. Before that, though, we concentrate on the technical aspects.

The physical and media access arrangements characterize this variant of Ethernet. In line with the principles we laid out in earlier chapters, the LLC, frame format, and addressing scheme are identical to those used with all the other 802 LAN types. So, despite a very different medium, a wireless Ethernet will interwork quite happily with a wired LAN through a device known as a media converter.

Before taking a brief look at the way in which wireless LANs operate, we give some specifics details of the PHY and MAC layers it uses.

4.1 The Wireless LAN

As the name suggests, a wireless LAN uses a wireless transmission medium—a predefined range of frequencies in the radio spectrum. Until relatively recently, few organizations used wireless LANs because they cost too much, their data rates were too low, they posed occupational safety problems because of concerns about the health effects of electro-magnetic radiation, and significantly, the radio spectrum used required a license. Today, however, these problems have largely diminished; the radio spectrum used today is unlicensed, and there is renewed interest in the use of wireless LANs.

There is nothing new or particularly unique about today's wireless Ethernet. As we mentioned in [Chapter 2](#), the first wireless Ethernet products appeared back in the late 1980s. These tended to be marketed as special-purpose substitutes for traditional wired LANs. The idea was to use wireless links to avoid the cost of installing cables and so ease the task of relocating or otherwise modifying the installed network structure.

Over the years, however, organizations began to have second thoughts about this substitution strategy. With LANs becoming more and more popular, architects started to design new buildings that were extensively prewired to accommodate internal communication networks. Also, as data transmission technology advanced, it became possible to use inexpensive twisted-pair cabling for local networks—in particular, category 3 and category 5 UTP introduced in previous chapters. So there was little motivation to replace or supplant wired network with wireless connections.

The higher speed of operation offered by the latest incarnation of wireless LAN technology—IEEE 802.11—combined with hype surrounding third generation mobile data services has sparked renewed interest in mobility. With transmission around 10 Mbps, the wireless network offers comparable transmission rates to its wired counterparts and can readily reestablish itself in environments where the motivation to avoid wiring is high. For instance, buildings with large open areas, such as manufacturing plants, stock exchange trading floors, and warehouses, are all good candidates. In these instances, wired networks are awkward to install because of the limited scope for cable placement. In addition, older buildings often have insufficient twisted-pair cabling and are protected from development, so the drilling holes for new wiring is not an option. In short, there are plenty of places where the latest wireless networks look to be a good investment.

4.2 Physical Layer

As before, the physical layer is the part of the Ethernet specification that defines the actual transmission details. On this occasion, there is no tangible medium—no wires or cables that you can touch—so we are left with only radio spectrum and free-space light signals to carry our frames.

The IEEE issued the 802.11 physical layer specification in three stages. The first part, which was issued in 1997, was called, simply, IEEE 802.11. This comprised a MAC layer and three PHY specifications—all of which were expected to operate at data rates of 1 and 2 Mbps. These three were the following:

- A direct-sequence spread-spectrum (DS-SS) system, operating in the 2.4-GHz industrial, scientific, and medical (ISM) band;
- A frequency hopping spread-spectrum (FHSS) system, operating in the 2.4-GHz ISM band;
- A diffuse infrared signal, operating at a wavelength between 850 and 950 nm.

The infrared option never really gained much market support and has quietly disappeared from view because it requires an unobstructed line-of-sight path and also because the available data rates are limited. So, the IEEE 802.11 specification quickly became a radio-based variant.

Both radio options use spread-spectrum approaches, which have their roots in military applications and require a much wider bandwidth than is actually necessary to support a given data rate. The idea behind using the wider bandwidth is to minimize interference and reduce the error rate during transmission.

Most of the early 802.11 networks used the *frequency hopping* (FH) scheme, which is simpler to implement and consequently less expensive. Networks that used the *direct sequence* (DS) scheme could be more effective, but all the early 802.11 products had data rates of at most 2 Mbps, which limited their usefulness.

By the end of the 1990s, the IEEE issued the second and third stages of the PHY definitions. These two—known as IEEE 802.11a and IEEE 802.11b—were released at roughly the same time and have both been actively developed over the last few years.

IEEE 802.11a operates in the 5-GHz band at data rates up to 54 Mbps. IEEE 802.11b operates in the 2.4-GHz band at 5.5 and 11 Mbps. Because the latter is easier to implement, it has yielded products first and has thus been the first to appear in the marketplace.

The various layers in the IEEE 802.11 architecture are depicted in [Figure 4.1](#). This shows all five of the physical medium options that have finally emerged.

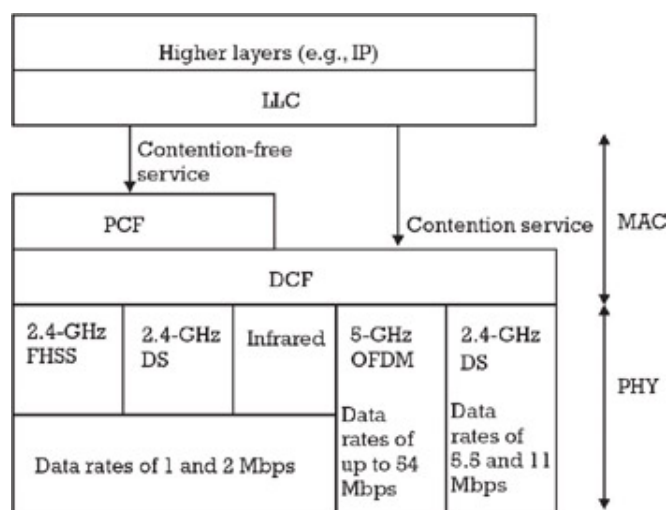


Figure 4.1: The IEEE 802.11 protocol layers.

We now give a little more detail on the latest PHY layers and the technology that supports them all. We next consider the options that have been introduced into the MAC layer for wireless operation.

4.2.1 IEEE 802.11b

The IEEE 802.11b PHY has been specified to provide data rates of 5.5 and 11 Mbps. In order to achieve higher data rates, it uses the DS spread-spectrum technique. This means that IEEE 802.11b can interwork with 1-Mbps and 2-Mbps IEEE 802.11 systems that use DS, but it will not work with IEEE 802.11 FH systems.

To extend range and guard against interference, IEEE 802.11b uses dynamic rate shifting. This allows the actual data rates used to be adjusted to compensate for the condition of the radio channel. Simply put, when the radio signal is strong, you get a high data rate, and when the radio signal is weak, you get the best data rate the stations can sustain. Ideally, the full 11 Mbps would be used, but if a transmitter is beyond the optimal range or there is substantial interference, it will fall back to 5.5 Mbps, 2 Mbps, and so on. The full rate is restored if radio conditions improve. Rate shifting is a PHY feature and is transparent to both the user and the upper layers of the protocol stack.

The IEEE 802.11b specification quickly led to product offerings, including chip sets, PC cards, access points, and systems. Apple Computer was the first to field IEEE 802.11b products, offering the AirPort wireless network option to users of its iBook portable computer. Other companies, including Cisco, 3Com, and Dell, quickly followed with their own wireless LAN offerings. At the time of publication, IEEE 802.11b PC cards and access points available from commodity-focused manufacturers like Linksys and NetGear make wireless LANs easily affordable for home and small businesses.

Although all these products are based on the same standard, it is still a concern of many users that products from different vendors may not interoperate successfully. Recognizing this concern, the Wireless Ethernet Compatibility Alliance (WECA) created a test suite to certify interoperability for IEEE 802.11b products. Interoperability tests have been going on since early 2000, and a considerable array of (competitively priced) products has now achieved certification.

4.2.2 IEEE 802.11a

Although IEEE 802.11b is successful to some degree, the data rate is still too low for applications that need a truly high-speed LAN. IEEE 802.11a targets this specific need. Unlike the other IEEE 802.11 standards, IEEE 802.11a uses the 5-GHz band, also an unlicensed band, and replaces the spread-spectrum scheme with the faster *orthogonal frequency-division multiplexing*, thankfully abbreviated to OFDM. This scheme, which is also called *multicarrier modulation*, uses up to 52 carrier signals at different frequencies, sending some of the bits on each channel. Possible data rates that can be enabled with OFDM are 6, 9, 12, 18, 24, 36, 48, and 54 Mbps.

We return briefly to the two original radio options within IEEE 802.11 for the physical layer. Both were specified to operate in the unlicensed radio band around 2.4 GHz and to use spread-spectrum techniques. Before looking in more detail at these two physical layer options, it is worth dwelling on the radio frequency at which both were originally intended to operate. The frequency band around 2.4 GHz is known as the *ISM band*. This frequency band is recognized by international regulatory agencies, such as the Federal Communications Commission (FCC) in the United States, European Telecommunications Standards Institute (ETSI) in Europe, and MKK in Japan, for unlicensed operation. Although unlicensed, and therefore free from cost and planning constraints, there are some important regulations that govern equipment operating in the ISM band. Most notable are the limits on transmission power that are imposed to minimize interference. One of the key drivers in the design of the IEEE 802.11 radio PHY has been the limited transmission power (as low as 100 mW) that is available.

4.2.3 Spread-Spectrum Techniques

As previously stated, both options for the IEEE 802.11 radio PHY are spread-spectrum techniques. Each option, FH and DS, is explained in detail later. First, a few words on their common root.

As is implied in the name, *spread-spectrum techniques* use a broader frequency range than is needed

for a given bit rate and modulation technique. The frequency distribution of the transmission signal is introduced by coding at the transmitter. This coding is independent of the data being sent and is carried out before modulation of the data takes place.

The performance of a spread-spectrum system is not much different from an ordinary, narrowband system in terms of immunity to noise. A spread-spectrum system is considerably less prone to detection or interception. So the complexity introduced by adding a coding stage to signal transmission pays off in terms of security. Because transmissions over a spread spectrum are difficult to detect and intercept, military applications drove the development of early spread-spectrum techniques.

4.2.4 Frequency Hopping Techniques

The code introduced in this instance relates to the sequence of frequency changes applied to the signal carrier. FH systems use conventional modulation but change the carrier frequency at a given rate. A receiver that does not know the code cannot follow the frequency hops, so it misses most, if not all, of the data being sent.

With frequency hopping, the 2.4-GHz band is divided into 75 1-MHz channels. Once the sender and receiver have agreed on a hop pattern, they start to jump from one channel to another as their dialogue unfolds.

It is possible to have a frequency hopping system where the hop rate is faster than the bit rate of the data being transmitted. However, this is quite complex and tends to be expensive to implement in hardware. The more popular option is a slow FH system. In such systems, the data to be transmitted is split into packets, and each packet is sent using a different carrier frequency.

One of the main attributes of FH is that it is fairly immune to interference. If there are overlapping signals, they only have an effect on each other when their carrier signals coincide, which is limited by the continual resetting of the FH carrier. Another attraction of FH is that it is relatively simple.

FH's main limitation is that it is restricted to speeds of no more than 2 Mbps. FCC regulations restrict subchannel bandwidth to 1 MHz, forcing FH systems to spread transmissions across the entire available 2.4-GHz band. FH systems must hop frequently, and this incurs a significant hopping overhead.

4.2.5 DS Systems

With DS systems, the bit sequence transmitted is combined with a higher rate bit sequence (known as the chip sequence) to generate a coded signal at the chip sequence rate. This sequence, which is always a multiple of the bit rate of the raw data, is then used to modulate the carrier.

The selection of the chip sequence has a significant effect on the properties of a DS system. A short code-that is, a bit pattern that repeats frequently-makes it easy to acquire synchronization at the receiver. A long code is more difficult to decipher and therefore protects the information being sent.

Like FH, DS divides the 2.4-GHz band into channels-14 in this instance-and these channels are separated by 25 MHz in order to minimize interference. DS systems have a low power density, which minimizes the probability of signal interception, as well as the amount of interference generated. Also, because the signal bandwidth is much narrower than the transmission bandwidth, interference is less of a problem. Most of its power will be dissipated in the transmission bandwidth.

4.3 MAC

Every LAN consists of a set of devices that shares network transmission capacity, so there must be some way of controlling access to the specified transmission medium. Ideally, this mechanism will ensure that all devices use the available capacity in an orderly and efficient fashion. This is one of a number of responsibilities that fall to the MAC protocol. Other key MAC functions are to ensure that all LAN stations cooperate; that only one station transmits at a time; and that the data a station sends is formed into standard format frames.

The transmission of information over an unprotected medium that is open to forces of nature poses an interesting challenge of having to cope with an unreliable medium. Up to this point, the Ethernet MAC has evolved along one path and the transmission aspects (when copper wire or fiber are used) have not been an issue. When the available medium is radio spectrum, it is necessary to carefully reevaluate not only how the MAC operates but also what it needs to do. So, before explaining what the result is, let us consider the requirements.

4.3.1 Features Required for Wireless Operation

The IEEE 802.11 study group identified certain key issues that had to be taken into account in the design of their MAC protocol. The main items on the list are the following:

- *Throughput*. Radio spectrum is a limited resource, so it is important that the MAC makes the best use of it. The target here would be wired LANs, where the user traffic can use up to 70-80% of the total bandwidth.
- *Transparency*. The various PHY layers that can be used with IEEE 802.11 have different propagation characteristics, and these need to be accommodated by the MAC.
- *Delay*. This is an important consideration with unpredictable media, so steps should be taken to bound delay.
- *Security*. Because the medium is open to both interference and eavesdropping, a successful network security scheme and MAC should minimize this.
- *Fairness*. The fading characteristic of a radio signal makes it likely that one station may, at times, have to operate at reduced power. The MAC design should ensure that the disadvantaged station is given equal access opportunity.
- *Power consumption*. One of the essential features of a wireless device is its portability. The MAC design should minimize device activity (and, hence, power consumption).
- *Deployment optimization*. A practical wireless LAN may have to support several hundred stations over an area of several 100m². The MAC design should be capable of handling this footprint (and the consequent propagation delays), as well as the required number of attached stations.
- *Robustness*. If two wireless LANs are operating in close proximity, it is likely that there will be significant interference. It is, therefore, important for the MAC protocol to be resilient enough to carry on with normal operation in the face of this interference.
- *Roaming support*. As wireless LANs become more popular, it is likely that a user will want to move from one LAN to another. If this were the case, the continuity of connection would have to be supported by a hand-off mechanism. Roaming may be supported by higher level protocol (such as mobile IP) but the MAC would also have to play its part.
- *Broadcast capability*. This is a natural form of Ethernet communication so should be supported in its wireless variants.
- *Priority marking*. A characteristic of wireless LANs is that the down-link traffic is usually considerably higher than the uplink. To expedite the delivery of information, some form of priority marking is needed.
- *Connectivity*. Most applications require interconnection with stations on a wired backbone LAN.

Wireless LANs easily satisfy this requirement by using control modules that connect to both types of LANs.

- *Collocation.* As wireless LANs become more popular, multiple installations are likely to operate in close proximity. Consequently, a device assigned to one LAN may be able to transmit or receive without authorization on a nearby LAN. To prevent this, the wireless LAN scheme must use addressing and access control techniques.
- *Manageability.* The MAC protocol provision for addressing and network management should let organizations dynamically and automatically add, delete, or relocate end systems without disrupting other network users.

One final requirement, driven more by commercial good sense rather than operational necessity is for license-free operation. The experience of third generation license auctions is that operators find it difficult to meet the cost of service provision when burdened with a heavy up-front barrier to entry.

All of these issues pertain equally well to any wireless MAC. However, it is the IEEE 802.11 variety that currently concerns us, so we now move on to explain how this works.

4.3.2 The Structure of IEEE 802.11 MAC

In order to address all of these requirements, the resulting MAC layer turns out to be rather complex. Unlike the previous MACs that we have examined, the IEEE 802.11 variant specifies a distributed coordination function with a rudimentary priority scheme, in which all stations cooperate for medium access. We return to the finer points of the wireless Ethernet MAC shortly. Before that, we need to examine the layout of an 802.11 LAN.

IEEE 802.11 defines two pieces of equipment. The first is a wireless station, which is usually a computer equipped with a NIC. The second is an access point (AP), which acts as a bridge between wireless and wired networks and is where the wireless station connects to the wireline network. An AP usually consists of a radio transceiver, a wired network interface, and bridging software that conforms to the Ethernet bridging standard defined in IEEE 802.1d.

The smallest building block in the IEEE 802.11 specification is a *basic service set* (BSS). This consists of at least one access point connected to the wired network infrastructure and a set of wireless stations. Most wireless LANs would be built up from several of these. If more than one BSS is deployed, they have to be connected together using a distribution system. These single and multiple BSS options are shown in [Figure 4.2](#).




Figure 4.2: Options for connectivity.

A BSS, shown on the left in [Figure 4.2](#), is defined to be a group of stations that are under the direct control of a single coordination function. In other words, they all execute the same MAC protocol and compete for access to the same, shared medium. This coordination function can be a distributed coordination function (DCF) or a point coordination function (PCF), both of which we will define and explain. The area served by the BSS is known as the *basic service area* (BSA). This is analogous to

the cell in a familiar mobile telephone network.

All of the stations within a BSS can communicate directly with each other (assuming that there are no problems with the transmission medium). There is no support for mobility-this option simply enables ad hoc communication between a number of stations over a defined geography.

In contrast with this ad hoc network, the arrangement shown on the right of [Figure 4.2](#) provides users with extended services and range. These conjoined networks are established by connecting each constituent BSS into the distribution system via an AP. The AP acts rather like the base station that serves several cells in a mobile phone network. The network formed by the collection of BSS is known as an *extended service set* (ESS), and this looks like one big BSS to the LLC of each station.

The distribution system that connects the various BSSs together can be thought of as a backbone network that supports the operation of the MAC across the whole of the ESS. This backbone is technology independent and can be provided with any of the other IEEE 802 networks-a Fast Ethernet or another wireless LAN-or even a fiber distribution network. In addition to connecting the separate BSS together, the distribution system allows a portal into other networks (such as the Internet or an organization's virtual private network) to be provided, as shown in [Figure 4.3](#).




Figure 4.3: A wireless LAN as part of a larger network.

The portal logic is implemented in one of the devices introduced in [Chapter 2](#). If the connection is to an IEEE 802.X network, it acts like a bridge and provides range extension and transfer between different frame formats. If the connection is to another network altogether, a router is used.

4.3.3 Coordination Functions

The IEEE 802.11 specification defines one basic coordination schemes plus two further optional ones. These are the following:

- The DCF, which fills the lower part of the MAC layer and uses a standard Ethernet style of contention algorithm that provides access for all traffic. Ordinary, asynchronous traffic uses this function.
- DCF with handshaking, which adds request to send (RTS) and clear to send (CTS) signals to the DCF. These familiar signals are used to control the dialogue between stations.
- The PCF, which is used for distributed, time-bounded services. PCF allows priority access to the medium to be granted to a particular station.

We now look at each of these options in a little more detail.

4.3.3.1 DCF

This is the fundamental access method used to support asynchronous data transfer on a best-effort basis. It is based on the CSMA/CD protocol introduced in [Chapter 2](#), with one difference-collision avoidance is implemented, rather than just collision detection.

CSMA with collision avoidance works like this. A station wishing to transmit listens and, if no activity is detected, waits an additional, randomly selected period of time. It then transmits, assuming the medium is still free, and runs a "waiting for acknowledgement (ACK)" timer. If the destination station receives the frame intact, it issues an ACK frame. When this ACK reaches the sender, the process is complete. If the sender does not receive an ACK frame before the expiry of the "waiting for ACK" timer, it concludes that either the original frame or the ACK was lost in a collision. The sender repeats the transmission process.

The extra caution built into CSMA/CA makes it very effective in an environment where the medium is not always reliable. The overhead it introduces does slow performance, though, and wireless LAN performance can suffer, especially in cases where collisions occur frequently. In cases where collisions are few, degradation is less marked.

The DCF is the only one of the coordination functions used in the BSS. In an ESS, it may be used on its own or in tandem with the PCF.

The DCF sits directly on top of the PHY layer to support contention services. The DCF contention service allows fair access to the medium for all stations. When a station has some data to send, it first must contend for access to the medium and then, once it has sent the first frame of that data, it must contend again to send subsequent frames.

4.3.3.2 DCF with Handshaking

Two handshaking signals have been added to the basic DCF in order to solve the "hidden node" problem. Hidden nodes occur in radio systems that rely on the physical sensing of the carrier. In such instances, any one receiver can hear (i.e., is in radio range of) two different transmitters, but these two transmitters cannot hear each other. Both transmitters can send to the receiver but neither will back off when there is contention because they are not aware of each other (cannot hear the other's signal, so cannot discern a collision).

Just as in a RS232 cable link, an RTS is issued by a hopeful transmitter to check whether he might continue what he intends. A CTS frame gives the requesting station permission to send a frame. At the same time, it tells all stations within range not to start sending for a specified time known as the *net allocation vector* (NAV).

Because of the signaling overhead that it introduces, handshaking is not used when the packets sent are small or if there is a low probability of collision. A station can choose to implement handshaking none of the time, all of the time, or when the frames to be transmitted exceed a threshold.

4.3.3.3 PCF

The PCF sits on top of the DCF, in the upper portion of the MAC layer. It is connection-oriented and allows frames to be transferred without contention. The PCF uses a facility called the point coordinator (in acronym overloading so fashionable in technology, this is referred to as a PC). The PC polls the stations on the network so that they can send without having to contend for the channel. The polling sequence is not preordained but is left to the implementer.

When the PCF and DCF work together, they share the time that is available for accessing the medium; a portion is allocated for contention-free access, the rest is used for contention-based traffic. A special frame, known as a beacon frame, is used to delineate PCF and DCF modes.

4.3.4 Other MAC Functions

In order to address the list of requirements given at the start of this section, the IEEE 802.11 MAC has been designed to provide the following functions.

4.3.4.1 Power Management

The MAC layer supports power conservation to extend the battery life of portable devices, as well as control access to the medium. There are two modes of power use—continuous aware mode and power-save polling mode. With the former, the radio is always on (and drawing power). In the latter, the radio is asleep, and the access point queues any data that is intended for it. Periodically, the radio

wakes up in response to signals from the access point. These signals tell the station whether there is data queued for reception. If there are, the station receives the data and returns to its quiescent state.

4.3.4.2 Association

When a station enters the range of one or more access points, it chooses one to associate with (i.e., it joins a BSS). This choice is usually based on signal strength and observed error rates. Periodically, the station will check to see if there is a better access point that it should connect to. If there is, it reassociates. The IEEE 802.11 MAC is responsible for handling both the initial association and the subsequent reassociation.

4.3.4.3 Security

Access control and encryption facilities, known as wireless equivalent privacy (WEP), are included in the IEEE 802.11 MAC. Together, they are intended to provide the same level of security on a wireless LAN that you get on a wired LAN.

For access control, a wireless LAN service area identifier is programmed into each AP, and this identifier must be known by any station that wishes to associate with that AP. In addition, it is possible to establish a table of MAC addresses, called an access control list, in the AP. This restricts association to stations that have an address on the list-if a MAC is on the access control list, the AP will not bridge packets or forward them to any other station it is serving.

Although affording some level of security, spread-spectrum transmission does not provide strong protection. If a transmission is intercepted, data sent in clear can readily be extracted. For this reason, a 40-bit key is provided for data encryption. This conforms to the RC4 algorithm from RSA data security. All data sent and received is encrypted with this key as long as the station is associated with the AP. Furthermore, the AP can issue an encrypted challenge to any station that tries to associate with it. In order to gain access to the network, the station must use its key to return the correct response to the challenge.

4.3.4.4 Fragmentation

It is sometimes desirable to break large frames down into a number of smaller ones-a process known as fragmentation. This is useful when the medium is congested, as smaller frames are less likely to be corrupted. In such circumstances, fragmentation can reduce the need for retransmission and thus improve overall network performance, although fragmentation over multiple hops is traditionally proven to be inefficient. The MAC layer is responsible for reassembling fragments received so that the process is transparent to higher level protocols.

4.4 Typical Configurations

In most cases, an organization will already have a wired LAN to connect its servers, printers, and workstations.

A typical case in point would be Norwest Nuts & Bolts, a small manufacturing plant that needed to install a network to link its office and factory floor operations. This company had an office area that was physically separate from the factory floor but needed to be logically linked via the LAN. The office area was already fitted with a wired LAN but the factory did not lend itself to any sort of fixed network—errant forklift truck and delivery lorries had disabled previous attempts. Given this, linking a wireless LAN into the wired LAN on the same premises was the only viable option—and one that delivered exactly what was needed within a demanding deadline.

This kind of application, or LAN extension, can take one of several forms.

Single and multiple cells. In [Figure 4.4](#), we show the basic single and multiple cell configurations. The former is so named because all the wireless end systems are within range of a single control module. In the multiple-cell wireless LAN, the wired LAN connects multiple control modules. Each control module supports the stations within its transmission range.

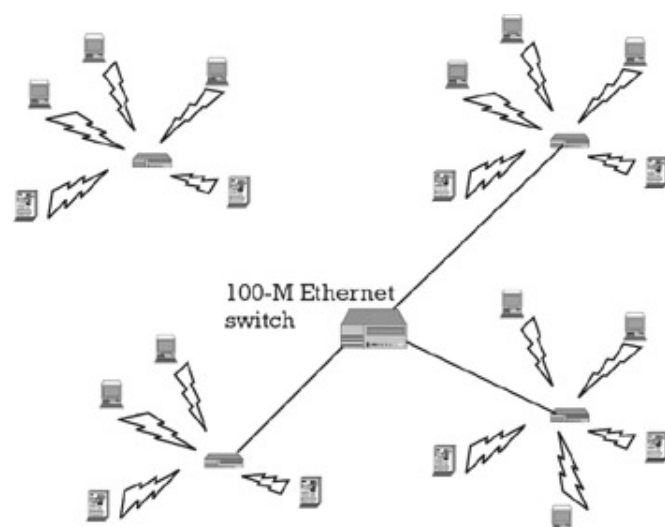


Figure 4.4: Two basic deployment options.

Nomadic access. In this configuration, the wireless LAN links a LAN hub and a mobile station equipped with an antenna, such as a laptop or notepad computer. This would suit a high-tech firm with multiple buildings in a small campus. Employees would routinely bring laptops to conference rooms or to other employees' offices to collaborate. Nomadic access is needed to accommodate such mobility across campus. It allows, for example, an employee returning from a trip to transfer data from a personal portable computer to an office server.

Nomadic access is also useful in a fragmented environment where people do not have a fixed base, such as a campus or a business operating from a cluster of buildings. With nomadic access, users can move around with their portable computers and access the servers on a wired LAN from various locations.

Ad hoc network. This network is set up temporarily to meet some immediate need—it has no centralized server and just connects people together. This does not mean that the ad hoc network is of limited use—indeed, it is a very handy configuration.

It is all too common these days for a dozen people to show up at a meeting with laptops. Once they get into the conference room, all they find is a tangled mass of RF45 cables and an 8-port hub smothering all the surface area of the conference table. Once the initial cries of "what a mess" have died down, there are several options. Everyone can decide to play musical chairs and timeshare whichever connections can be persuaded to work. Or they could wait while the local network guru runs

down the hall for an additional hub, only to discover that no one really knows how to cascade hubs. Or they could just get bad tempered and abandon the meeting.

An altogether better alternative would be to just plop an AP in the room, free it of the wired clutter, and let any number of people "link up" (and keep a few extra wireless PC cards handy, just in case someone does not have one). Thus, in meetings, a group of employees, each with a laptop or palmtop computer, can link their computers in a network that lasts just as long as the meeting.

Team LiB

[◀ PREVIOUS](#) [NEXT ▶](#)

4.5 Open Issues

Despite becoming firmly established as a flexible and effective networking solution, there are still some issues with wireless Ethernet that have yet to be resolved. Perhaps the most important to be aware of is its potential for interference with other systems that operate in the 2.4-GHz band, most notably Bluetooth and HomeRF. With so many devices that use the same portion of the spectrum as 802.11b, including baby monitors and even garage door openers, some frequency planning is needed before installing a wireless LAN.

Although still an open issue, there is a coexistence study group, IEEE 802.15, examining the various aspects of interference and, so far, the prospects are encouraging. The aim of the group is to provide mechanisms that let devices of all radio persuasions exchange information and cooperate to minimize mutual interference.

As more IEEE 802.11-compliant products become available, security is likely to become a major issue. Wireless LANs are uniquely vulnerable to both eavesdropping and unauthorized transmission simply due to the fact that radio spectrum is used rather than cabling. As explained earlier, the IEEE 802.11 standard has provided some ways to address these concerns, and an increasing number of equipment suppliers are routinely implementing the security portion of the standard as part of their offerings. However, when compared with the security measures used to build virtual private networks, the IEEE 802.11 security provisions are fairly weak. In mitigation of this are strong security measures such as IPSec that can be applied at higher layers of the networking protocol.

Although listed here as open issues, it would be unfair to cast these as debilitating or even significantly inhibiting problems with wireless LANs. In truth, these are areas where there is not yet the same level of practice or experience as there is with wired LANs. And the history of Ethernet is that practical solutions soon appear.

4.6 Summary

As the desire and sometimes need for mobility increases, so will the appeal of wireless networks grow. One of the more recent additions to the Ethernet family, the IEEE 802.11 standard, offers a comprehensive, scalable, compatible, and flexible solution. It has taken nearly 20 years to get wireless LANs to a viable state but now that the technology is demonstrably stable, their application is exploding and we are seeing a rapidly growing number of wireless LANs in offices and factories.

In this chapter, we explained the workings of the wireless LAN. In particular, we focused on the PHY and MAC layers that distinguish a wireless LAN from its wired cousins. The various options for deploying a wireless LAN and current implementations (notably 802.11b and 802.11a) are explained. Along with this, the way in which a network comprises wireline and wireless components is illustrated.

The fact that the popularity of wireless LANs has rocketed over the last couple of years indicates that there is considerable mileage in this technology. As the demand for flexible networking rises, it is likely that many more IEEE 802.11 solutions will be deployed in coming years.

Having dealt with the technology of Ethernet over the last few chapters, we now go on to look at its application. In the case of wireless Ethernet, there is an interesting and, as yet, unanswered question of how it will fare against third generation mobile networks that aim to provide data services.

Selected Bibliography

Frankel,S.,*Demystifying the IPSec Puzzle*,Norwood, MA:Artech House,2001.

GeierJ.,*Wireless LANs*,New York:McMillan Technical,1999.

O'HaraR., and A.Petrick,*IEEE 802.11 Handbook*,New York:IEEE Press,1999.

Santamaria,A., and F.Lopez-Hernandez *Wireless LAN Standards and Applications*,Norwood, MA:Artech House,2001.

In addition to these books, which provide excellent technical reference for all aspects of wireless LANs, the following industry-hosted Web sites provide current information:

<http://www.wlana.com>, Wireless LAN Association.

<http://www.wirelessthemet.org>, Wireless Ethernet Compatibility Alliance.

<http://www.grouper.iee.org/groups/802/11/index.html>, IEEE 802.11 Wireless LAN Working Group.

Chapter 5: Total Area Networks

Overview

Invention is the talent of youth and the judgment of age.
--Jonathan Swift

One of the enduring truths of prediction is that people overestimate what will happen in the next couple of years but underestimate what will happen over the next 10. This is as true of the Ethernet as it is of any other major phenomenon. For all the high expectations for Ethernet when it emerged as the dominant LAN technology in the 1990s, few would have seen it as a viable candidate to be the one and only network technology. Yet this is precisely what is happening. The capacity, range, and flexibility of the Ethernet has evolved to such a degree that it can now be used to build the *total area network*-the homogeneous medium that supports communication in the office, across town, and beyond.

In keeping with the current state of telecommunications, we look at the prospects for Ethernet in both the WAN and LAN in this chapter. In particular, the prospect of high-speed Ethernet technology that can support an organization spread over tens of kilometers with no need for a WAN connection is discussed.

In the wide area, the de jure technologies are SONET and ATM. Both are well established and have been designed to operate at the levels of resilience required for public networks. The pros and cons of Gigabit Ethernet against these incumbents are considered, and an outline business case is presented.

In the local areas, Ethernet is (by its very nature) already well established. There are, however, many local-access connections that do not use Ethernet. These are the local loops that connect a phone user to the nearest exchange. Here, too, there are prospects for Ethernet, and the relevant flavor-EtherLoop-is examined in the second part of the chapter.

The overall intent of this chapter is to show that Ethernet has the flexibility to provide a uniform network technology that can be used from one end of a link right through to the other.

5.1 Ethernet in the Wide Area

For many years, networks were something of a mystery to the common man. They were provided by national operators, such as AT&T, NTT, and BT, which had invested in a range of different technologies to build an all-purpose infrastructure.

As LANs became more ubiquitous and a commodity, more and more people became network-aware and, as the speed of the LAN increased, these people wanted to extend the application of their local network. In this section, we look at the prospect of Ethernet as a metropolitan and wide-area networking technology. First, we consider the current situation.

5.1.1 Traditional Setup

As explained in [Chapter 1](#), the core of the modern telecommunications network is typically comprised of high-capacity optical fiber. These fibers provide an enormous amount of raw bandwidth—more than enough to carry every telephone call in the world, several times over.

But there is more to communications than raw capacity. In order to be effectively deployed, the fiber highway needs to be structured in some way so that it can carry that traffic where it needs to go. This is where SONET comes into play. SONET, and SDH, its equivalent in Europe, is a multiplexing transmission carrier system in which lower bit rate channels are interleaved into a higher level, fixed-length, frame structure and transmitted in a hierarchy of successive levels.

The levels commonly deployed are STM-1 at 155 Mbps, STM-4 at 622 Mbps, STM-64 at 2.4 Gbps. [Table 5.1](#) relates the SDH and SONET signal names to the capacity available and traffic that can be carried.

Table 5.1: Capacity and Traffic Offered by Current Transmission Technology

SDH Signal	SONET Signal	Bit Rate	Capacity-Bearer Circuits	Capacity-Voice (Mmins/month)
STM-0	STS-1/OC1	51.48 Mbps	630 ISDN channels	5
STM-1	STS-3/OC3	155 Mbps	63E1 or 3E3	15
STM-4	STS-12/OC12	622 Mbps	252E1 or 4E4	60
STM-16	STS-48/OC48	2.488 Gbps	1008E1 or 16E4	240
STM-64	STS-192/OC192	9.95 Gbps	4032E1 or 64E4	960

A single OC192 connection can carry more than 32,000 uncompressed simultaneous voice calls, and data and voice can be carried in the same signal. There are literally hundreds of OC192 connections in the combined core networks of the various global and national carriers.

In addition to its high capacity, SONET is a very resilient and flexible transmission technology. It is usually deployed over fiber rings, can carry a wide variety of traffic, and has inbuilt mechanisms that allow recovery from breaks in the underlying network.

The resilience, flexibility, and capacity of SONET make it very popular option for implementing high-capacity connections with bounded metrics (e.g., latency, loss, information rate). It has rapidly replaced its predecessor [known as the plesiochronous digital hierarchy (PDH)] and is the technology of choice for long-distance voice transmission networks.

SONET has proven to be an excellent transport for a voice-centric network, but the mix of traffic being

carried is changing dramatically and voice traffic is being overtaken by data. Unlike voice traffic, data traffic tends to be bursty, and it is best served by a transmission technique that does not assume that all traffic looks the same (like 64 Kbps connection-oriented voice). Carriage of data traffic does not play to the strengths of SONET, which offers a series of time slots for individual streams of predictable traffic rather than a less rigid format that can cope with variable rate traffic.

SONET's lack of efficiency in a data-centric world is highlighted by the lack of native interfaces for the leading packet technologies-IP and ATM. In both instances, additional, expensive equipment is required to make a connection. DS3/OC-x/ATM interfaces-and router/switch products capable of driving them-are hideously expensive compared even to Gigabit Ethernet NICs.

Another limitation of SONET is the equipment is speed-specific. That means that if a company wants to upgrade to a higher speed, for instance from 2.4 Gbps to 10 Gbps, all equipment must be replaced. Such "forklift" upgrades are disruptive and very expensive.

One final point on SONET is that it is a very complex technology, both to understand and to implement. It demands high-caliber engineers to deploy it and requires a good deal of expertise to keep it going-when you must reconfigure a SONET ring, it is rarely something that can be accomplished in real time. Needless to say, this all incurs more expense.

So much for the current state of transmission in the core network of the stereotypical telecommunications operator. If we move our focus a little closer to the end user and look at the technology that is used to deliver bandwidth from the core into the end organization, we would typically find ATM.

ATM is a flexible switching technology that can be configured to support a wide range of different services. It is capable of establishing a link or *virtual path* between users and, through its "adaptation" layer, it can adjust the link to suit different traffic types: For instance, a constant bit rate can be specified to carry voice, an available bit rate for e-mail and file transfer.

ATM uses a defined frame or *cell* structure to carry all of its data (in this case, a standard 53-byte packet with 5 bytes for control information and 48 bytes of payload). When first introduced, ATM offered up to 155 Mbps bandwidth, which is 50% more than Fast Ethernet. It was ideal for the emerging high-bandwidth applications of the time, most notably multimedia. The assured QoS that could be offered, combined with scalability, made ATM an attractive option for both public data services and corporate networks.

As its popularity grew through the 1990s, ATM extended its range of applications into the local area. The emergence of LAN emulation (LANE) allowed the same technology to be used in both the wide and local areas, a way of establishing a high-speed, end-to-end link with predictable and manageable characteristics. It seemed to many people (author included) that ATM had all of the credentials to be crowned as *the* technology for total area networking. But prediction only serves to make our descendants laugh. As ATM plodded forward one step at a time, Ethernet was evolving in leaps and bounds.

5.1.2 The Ethernet Alternative

For all of the virtues of SONET and ATM, they suffer two major drawbacks-they are complex and expensive. Ethernet, by contrast, has always been a more affordable option and, even if it were not, would be preferred by most people because it is easier to install and operate. With the speed and range of Ethernet increasing so dramatically over the last few years, so have its claims to the crown that was fleetingly awarded to ATM.

Although important, cost and complexity are not the only factors that will drive the future path of total area networking-there are a host of engineering considerations, such as compatibility and interworking. This is where Ethernet scores well. One of Ethernet's greatest strengths is that it has remained Ethernet through all its advances. Upgrading from one version to another is virtually painless, as is the integration of different technologies (including wireless). All of the applications that work on Ethernet will work just the same on Gigabit Ethernet.

The practical elegance of Gigabit Ethernet has persuaded many people to use it in place of ATM in their campus data networks. Gigabit Ethernet is beginning to make inroads as a transport for MANs.

Figure 5.1 illustrates how a MAN can be built using Gigabit Ethernet switches connected with the ubiquitous and plentiful fiber links that grace so many of our major cities.

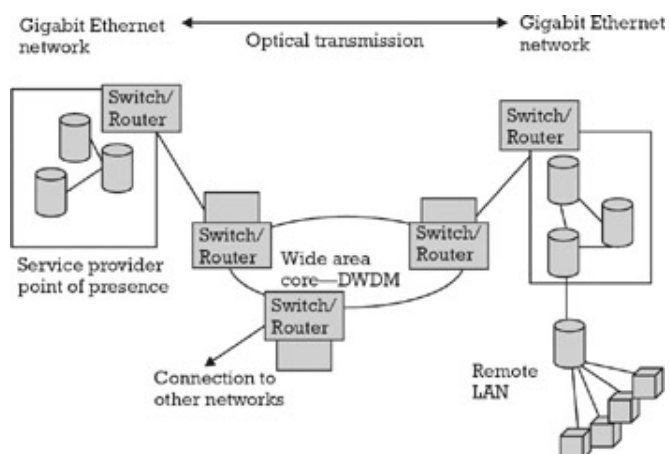


Figure 5.1: A MAN built with Gigabit Ethernet.

A number of innovative service providers, such as Yipes and Telseon, are buying dark fiber (i.e., unlit cable that they put their own transmission systems on). With these basic units, they are building fiber rings and running Ethernet over them. They then sell bandwidth (usually in 1-Mbps increments) to service providers, who, in turn, sell these services to end users. Exodus Communications, as well as GlobalCenter, the Web-hosting arm of Global Crossing Ltd., also offer Gigabit Ethernet services. It has to be said that many of these operators are currently struggling, but this has far more to do with the downturn across the entire telecommunications sector, not the technical merits of the solution offered.

Although in its early days, the Gigabit Ethernet market has already grown rapidly. There will inevitably be a period of consolidation as some of the new players succeed and others do not. Table 5.2 illustrates the range of providers already established.

Table 5.2: Gigabit Ethernet MAN Operators in Europe

Operator	Target Market	Services Offered	Main Equipment Suppliers
B2 (Sweden)	Residential users in multitenant units	10-Mbps Internet access	Cisco
COLT (Europe)	Commercial multitenant units	Gigabit Ethernet for data centres	ONI Systems
e.Biscom (Germany and Italy)	Corporations, SMEs, SOHO, and residential users	Always-on, 10-Mbps bidirectional IP connection for telephony, Web access, and DVD-quality video-on-demand (VOD)	Cisco
IntelliSpace (United Kingdom and United States)	Corporations	High-speed Internet access and data services. In-building networks offering services scaling from 64 Kbps to 10 Gbps	Extreme Networks, Juniper, Riverstone Networks
Utfors (Scandinavia)	SMEs and residential users	All IP, Ethernet-based multiservice carrier network	Sycamore Networks

Sphera Optical Networks (Europe and United States)	Carriers, ISPs, and data centers on metropolitan area rings	Optical-level transport from 155 Mbps to 10 Gbps in the metropolitan core	ONI Systems
--	---	---	-------------

(From: Analysys, 2001.)

As is clear from [Table 5.2](#), operators are already using Gigabit Ethernet to offer a wider range of services than those traditionally available from a metropolitan area data carrier. e.Biscom, for example, is moving far beyond standard voice services with its always-on, bidirectional, high-speed IP connections. The company aims to offer customers a full portfolio of services, both commodity and value-added, from voice and Internet connectivity to VOD.

Among the possible customers for Gigabit Ethernet technology are fiber-wired buildings, such as multitenant units, application service providers, companies with campuses containing buildings without servers, and companies interested in using this technology for remote backup and/or disaster recovery.

So Gigabit Ethernet seems to be ready to succeed. It is backed by the industry in the form of the Gigabit Ethernet Alliance, and it has a proven track record of being flexible, usable, and inexpensive. Furthermore, the fastest competing ATM products available now operate at a top speed of 622 Mbps. At 1,000 Mbps, Gigabit Ethernet is now almost twice as fast.

Of course, there are some areas in which Gigabit Ethernet might be perceived as falling short of ATM, most notably in its assurance of QoS. This means that the latter may still be better suited for applications such as voice and video, because it can condition a connection to minimize delay (which suits this sort of application). That said, QoS does not have to be built in to Ethernet, as it can be accommodated at layer 3 with techniques developed for QoS control in IP networks such as MPLS and DiffServ.

The aim of this section was not so much to compare Ethernet with established technologies but more to check its credentials as a WAN contender. It is evident that each of the technologies has desirable features and advantages. It could be said that seemingly divergent technologies are actually converging. ATM was touted to be the seamless and scaleable networking solution that extended into the local area. Ethernet, which was for a long time restricted to LANs alone, evolved into a scalable technology. What happens from here on is guesswork but it is worth noting that at the time of publication, an OC48 SONET port running at 2.4 Gbps costs around \$30,000, while a 1-Gbps Ethernet port that cost around \$1,200 a year ago is now available for under \$100.

5.1.3 The Quest for QoS

If Gigabit Ethernet is to become a dominant WAN technology, one of the key issues that must be resolved is how it assures QoS. The only option native to Ethernet for supplying any sort of acceptable levels of QoS is that which relies on overprovisioning of capacity. This approach has been tenable so far because transmission capacity has risen so fast, but as the size of a network grows, so does the need for traffic management. What worked in the fastest of local networks may not in a network shared over a wide area.

Of course, reserving extra capacity does offer some measure of protection against the problems of delay, jitter, packet loss, and throughput. However, the fact remains that guaranteed bandwidth between end points does not itself guarantee low delay or low jitter. If Gigabit Ethernet operators are to be able to move on to offering guaranteed service level agreements (SLAs), rather than limited assurance on a best-efforts basis, the other options must be explored.

There is no easy solution to assuring QoS in a packet network. The same issue faces the designers of IP-based networks. They also have to devise schemes for applying a connection-oriented feature in a connectionless network. However, there are candidate solutions, both in the short and long term.

One example of a stopgap measure is Extreme Networks' policy-based QoS monitor that it offers with its Alpine WAN switches. This relies on communication between the switches that control the traffic they admit. Extreme's QoS solution certainly works, but it is limited to an all-Extreme network since the

technology is proprietary. Companies deploying their own MANs using high-speed Gigabit Ethernet are likely to be forced, in the short term, to go with a single vendor's switches and take advantage of proprietary technology in order to be able to ensure timely delivery of high-priority packets.

In the longer term, there are a number of standards that promise to add QoS control to the Gigabit Ethernet repertoire. A number of alternatives are developing fast to address the QoS issue. The following are two main routes to solving the problem:

- Incorporate technology standards that refine Gigabit Ethernet itself, such as resilient packet ring (RPR) and IEEE 802.1p/Q.
- Use complementary QoS technologies such as MPLS and the Diff-Serv model.

The following discusses each of these options:

- RPR is a protocol designed to address the issue of network uptime. It uses the existing Ethernet MAC layer (rather than requiring a new layer) and uses TDM to allocate capacity to streams on a ring. The protocol, which is in the process of becoming a standard, is compatible with both DiffServ and MPLS. RPR can also accommodate the IEEE 802.1p/Q standards, which add traffic management capabilities to Gigabit Ethernet.
- IEEE 802.1p/Q allows switches to reorder packets based on priority level, adding traffic prioritization capabilities by marking packets as belonging to one of three data prioritization levels. IEEE 802.1Q, meanwhile, provides a platform that enables Ethernet to handle time-sensitive applications such as voice.
- With MPLS, a number of label switch routes are defined and established across a set of routers in a network by populating a label forwarding table on each router. On ingress to an MPLS network, the destination address of the incoming packet is evaluated and a label added to the packet that indicates the next router and the priority of the packet. All packets to the same next router and with the same QoS requirements are assigned the same label. The output queue of the ingress router and the links between routers can be engineered so that set amounts of bandwidth are allocated to certain labels and, hence, QoS-type packets. By this mechanism, the packet travels from router to router with its label being swapped but always going via a specific route, engineered for its particular QoS requirements. At the egress router, the last label is stripped off and the packet is transmitted to its destination. MPLS can be used with Ethernet to tackle the delay, jitter, and packet-loss aspects of QoS. Ethernet and MPLS, in combination, promise to offer the best of both worlds—layer 1 and 2 connection enabled by Ethernet and connection-oriented capabilities provided by MPLS at layer 3. MPLS and the DiffServ model have also made possible the delivery of multiple service levels over an IP backbone.
- DiffServ is a traffic prioritization mechanism that operates at layer 3 only and, when used with an appropriate network management system, can allow operators to offer their customers a range of next-generation applications, including real-time applications, voice over IP (VoIP), and videoconferencing. Despite operating separately from the underlying Ethernet, there are mappings between DiffServ type of service (TOS) bits and 802.1p/Q.

In truth, none of these options could be deemed as mature solutions, although MPLS is becoming more common among Internet service providers (ISPs). They each offer some distinct capability and, together, they promise to add the traffic handling capabilities that Gigabit Ethernet will need to become the technology for total area networking.

Currently, support for delay-sensitive services, most notably, voice, is limited. VoIP appears to be the dominant technology in providing support for voice services; Gigabit Ethernet operators Cogent and Telseon, for example, currently offer their customers a VoIP service via TalkingNets, an Application Service Provider (ASP) that provides voice services. However, VoIP technologies can still provide a less-than-perfect user experience. A large part of this is due to increased end-to-end delays in VoIP transmissions as compared to their more traditional counterparts—delays that may be discernible to the end user. VoIP solutions are also still a long way from offering the many features currently offered by traditional operators. However, companies such as Ufors believe that they are already on the way to solving the problem.

5.1.4 The Road Ahead

While Gigabit Ethernet is going through growing pains, 10-Gbps Ethernet development for the WAN is in the embryonic stage. Certain customers, such as Qwest, have asked vendors like Cisco to scale such deployments using early versions of 10 Gbps, and vendors are already implementing these solutions with 10-Gbps Ethernet framing over fiber.

Cisco demonstrated an early 10-Gbps Ethernet link several years ago at the May 2000 Network World Interop trade show in Las Vegas. It now plans to ship a single module for its switches that will support 10-Gbps Ethernet with different transceivers that will connect to single and multimode fiber. To make this happen, Cisco formed an alliance with Metromedia Fiber Network (MFN) that will give customers access to MFN's fiber networks that it is building in major metropolitan areas. To highlight how important this technology is to Cisco, it has established a new Metropolitan Services Business Unit and has acquired Qeyton Systems, a metropolitan dense wave division multiplexing (MDWDM) company that makes switches designed for MAN environments.

Some vendors have already begun to talk about 100-Gbps Ethernet as the next standard to be developed in order to ensure that service providers and enterprises have networks that can continue to scale. Another possible scaling approach would exploit DWDM (which combines multiple colors of light, each of which carries a separate communications channel, into one optical fiber). There have been successful tests of DWDM combining 100 colors at 10 Gbps each to produce a 1 Tbps pipe.

The logical end point here would be the elimination of the need for SONET. It would be replaced with DWDM, leaving Gigabit Ethernet as the ubiquitous networking technology.

5.1.5 The Business Case

Virtually all of the discussion so far has been about technology. A few hints have been dropped about the cost-effectiveness of Ethernet but we have not really considered the business side of the argument. This section does just that. We now present the outline of a business case for a small U.K.-based networking company that is considering the strategic deployment of Ethernet technology. The section takes the form of a response to the key questions that would need to be addressed before an investor would consider signing on the dotted line.

Why Gigabit Ethernet from a market point of view? According to Nortel, 95% of all LAN nodes are Ethernet. Nortel's observation is confirmed by a Network News/Black Box survey of 425 U.K. IT network and telecommunications staff that shows that Ethernet has entered the corporate sector as a mainstream technology.

Furthermore, a number of high-profile users (major airlines and utilities) have replaced their ATM LAN networks with Gigabit Ethernet. According to Ovum, by 2005, more than 1.8 million broadband corporate lines in the United Kingdom alone will be based on fiber. Given the dominance of Ethernet technology on the LAN, corporations deploying fiber connections on a WAN will prefer a native Ethernet end-to-end service as the most cost-effective option. The ease with which long-established Ethernet installations can be incorporated with the latest versions, as shown in [Figure 5.2](#), strengthens this case.

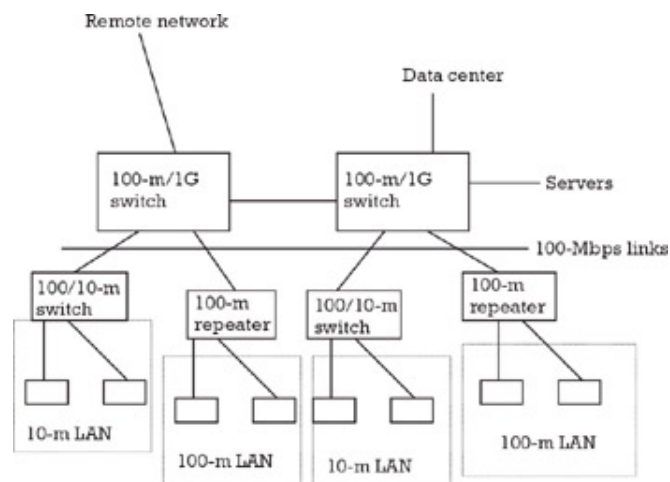


Figure 5.2: The compatibility across the Ethernet generations.

Given the statistics mentioned, it is not surprising to learn that 95% of all IP traffic generated from the growth in e-commerce and the outsourcing of key enterprise services (such as extranets, back-office automation, content/Web hosting, telecommuting, and business-critical application hosting) originates on Ethernet. This traffic might come from any part of [Figure 5.2](#)—a server farmer, a Web site, a remote site, or a local workgroup. A 1999 Ovum report ("IP, ATM and MPLS-Strategies for Broadband Networking") states that by 2004, 4 Tbps of installed IP capacity is expected within the United Kingdom as enterprises continue to build out intranets and extranets and add Gigabit Ethernet to LAN networks.

What are the customer benefits? With the development and commercialization of carrier class switches from vendors such as Riverstone and Extreme Networks and a massive deployment of metro fiber in cities, a new breed of Ethernet metro service providers has emerged in the United States. The likes of Telseon have rolled out Gigabit Ethernet services and regional area IP networks to aggregate and transport native Ethernet. These Ethernet metro service providers have proven attractive because they offer a clear-cut value proposition for customers who want to buy Ethernet instead of SONET-based services.

The simple reason for this is that Ethernet-based services take cost and complexity out of a customer's network. A switch sits in the basement of the customer's premises, aggregating lower speed traffic onto a Gigabit Ethernet pipe and adding QoS and service classes. The customer interface is a port on a LAN switch. Direct LAN-to-LAN service between business locations means there is no costly and complex multiplexing or protocol conversion at the customer end, and no incremental (and formidable) expense in ATM network interfaces for both the provider and customer.

Ethernet services also reduce one of the major hidden costs of enterprise networks, as the same people managing the corporate LAN can support extended virtual LANs and the end user understands the solution. IP over Ethernet is something every LAN administrator knows. With no additional training, the WAN administration is taken care of!

In addition, as Ethernet connections are not constrained by SDH hierarchy, customers can make soft changes to bandwidth by 1-Mbps increments through end-user controlled Web front-end interfaces. While some of this granularity exists for SDH-based services, expensive, big-step changes from 2 Mbps to 45 Mbps to 155 Mbps cannot be avoided.

What advantages accrue to the operator? From a carrier and competitive point of view, the unique proposition of the Gigabit Ethernet MAN and WAN is to allow an operator to eliminate expensive and inflexible packet over SONET (POS) across the network

Many layers of network elements and protocol complexity make services costly to provide. By stripping away layers such as SDH and ATM and deploying layer 3 Gigabit Ethernet switches at both the metro and backbone networks, an operator can provide a range of IP services and rapid provisioning that its customers require at a fraction of the price and complexity of legacy networks.

In [Table 5.3](#) is an approximate (but representative) cost comparison that shows that Ethernet is a relatively cheap network medium in terms of equipment costs, provisioning, and maintenance charges

compared to other technologies.

Table 5.3: Cost Advantages of Gigabit Ethernet

	Electronics/Optics \$/Mbps	Bandwidth Management and Provisioning	Annual Maintenance Upgrades	Bandwidth on Demand
IP/SDH	\$6-\$35	\$5K	\$750-\$3,750	Very difficult
IP/Ethernet	\$1-\$3	\$1K	\$150-\$450	Yes
Gigabit Ethernet advantage	6:1-12:1	5:1	5:1-8:1	Yes

(From: Yipes, Dell'Oro, Yankee, Extreme Networks and Juniper Networks.)

Ethernet networks also have the added benefit of optimizing network utilization. Ethernet active flow control (specified in IEEE 802.3x) has the ability to prevent the loss of data and allow transmission links to operate at very high utilization levels, whereas standard SDH transmission links (when carrying a mix of subrate tributaries) often operate at somewhere around 50% as the designed utilization.

Are there enhanced service creation opportunities for the operator? Unlike traditional SDH or POS services, Ethernet based-services can be provided for and billed on diverse bit rates (e.g., 1 Mbps, 10 Mbps, 100 Mbps, and up to 1 Gbps) with no engineering changes involved. This granularity allows customers to get the bandwidth they need. At the same time, it frees up excess capacity in the metro network for other customers.

In addition to bandwidth, Ethernet-based services allow an operator to provide and bill for QoS. Service providers can deploy Gigabit Ethernet products that offer different service classes based on either the customer or the application and collect the appropriate billing information based on the QoS parameters.

Hardware-based accounting (the ability to collect detailed customer usage) enabled by Ethernet switches allows an operator to offer tailor-made services and innovative billing plans based on actual usage patterns.

There are also other opportunities. While pure-play metro Ethernet service providers such as Yipes have used Gigabit Ethernet switches in metropolitan networks, they still deploy expensive high-speed routers as the backbone technology to connect metro islands. It is entirely feasible to optimize the pure-play metro Ethernet service by deploying some of the more recent technology. For instance, an operator could do one or more of the following in order to differentiate itself and/or reduce its cost base:

1. Use an MPLS core using Ethernet switches over a long-haul DWDM network to connect metro islands.
2. Design the network to feature a high number of low-cost layer 2 switches in metro networks (i.e., no layer 3 routers). The traffic from these could be backhauled to an MPLS core.
3. Have a limited number of Internet gateways at the edge of the network for value-added services such as fire walling, DSL aggregation, and Web caching.
4. Have an Internet peering cluster at the edge with a few mid-range low-cost routers to do Internet routing.

In tandem, these steps could achieve a truly low-cost end-to-end network and provide integrated service offerings that meet the needs of a diversified customer base by taking advantage of the latest technologies. Of course, a full business case needs to show a reasonable return on investment, a positive cash flow in the foreseeable future, and so on. The step, though, is to make the right strategic choice.

Team LiB

◀ PREVIOUS

NEXT ▶

5.2 Ethernet Over the Last Mile

As far as most telecommunications providers are concerned, the link between the user and the network is the last mile. It is the most remote and distant point from the core of their centralized network. For many providers, it is also the oldest and fussiest technology. Designed decades ago for analog voice traffic, it is a snarly mess of poorly (if at all) shielded copper pairs that is one of the most expensive maintenance items for providers. Of course, the user sees things very differently: This connection to the network is the *first* mile for voice, facsimile, Internet, and ideally, cable and broadcast television and entertainment. Whichever viewpoint we take, the fact is that there is a lot of (usually copper) wire between telephone exchanges and end users. From here on, we refer to this tangled morass of copper as the local loop, and we will now look at some of the ways in which it can be exploited.

5.2.1 Traditional Setup

For many years, the local loop carried analog telephony signals. This situation persisted until the 1980s, when it was realized that the digital techniques used in the core of most public networks (i.e., digital exchanges and transmission techniques) could be extended out to the user. The thinking was that an all-digital network would support a range of voice and data services, thereby leveraging the value of the installed base.

Worldwide, there are several hundred million twisted-pair subscriber loops already installed. These loops are used to carry voice traffic from the subscriber's premise to the local exchanges or central offices of the network provider, which takes this voice traffic and delivers it to another subscriber.

For a variety of reasons, mostly historical, voice traffic uses only the "bottom" 4 kHz of the analog bandwidth available in the copper wire. This is an extremely small portion of the actual bandwidth available. The actual amount of bandwidth varies with the quality of the loop, distance, and the gauge of the wire, but in general, most subscriber copper loops can reliably use about 1 MHz of bandwidth. If this unused bandwidth can be made available, it can provide a substantial resource for the telephone companies.

There have been a number of attempts to make use of this latent capacity, most notably ISDN and DSL. Both have sought to make better use of the existing local loop by replacing traditional analog transmission techniques with digital ones.

ISDN has been designed to allow a number of services to be carried together on the same telephone wire. It can be considered an extension of the PSTN, the key difference being that the analog transmission over the local loop is replaced by a digital echo canceling scheme. This means that the ISDN can readily carry any form of data, such as voice, video, and computer files, without the need for any sort of analog-to-digital conversion.

The initial motivation behind ISDN was to replace the analog telephone network with a less noisy, digital one. It was, therefore, designed around the same notion that already existed in the PSTN, with two separate channels operating at 64 Kbps. This number springs from the fact that basic, analog voice transmission requires 8K samples per second, each of which is encoded as 8 bits.

In the United Kingdom and Europe, ISDN is offered in two forms, ISDN2 and ISDN30, where the number suffixes denote the number of 64-Kbps channels that are provided. ISDN2, also known as basic rate access, provides two 64-K (B or bearer) channels and a single 16-K signaling (D or delta) channel. ISDN30, also called primary rate access, provides 30 B channels along with a D channel.

In the United States, primary rate access is based around 24 B channels, with one D channel. In both cases, basic rate is intended for home use, and primary rate is meant for businesses.

The problems with ISDN are that it took more than 25 years to deploy, is still not as ubiquitous as the PSTN, and is expensive. Coupled with the fact that the bandwidth needs for data and multimedia services leapfrogged ISDN's meager offering by nearly two orders of magnitude, the service looks lame to many people.

DSL defines how a pair of modems—one located at the local telephone exchange and the other at the

customer site-can be used to deliver high-speed signals over their established twisted-pair copper connection. There are several varieties of DSL, including the following:

- *Asymmetrical DSL (ADSL)* allocates the available bandwidth in an asymmetric spectrum so that ore data is delivered downstream (toward the user) and is then returned to the exchange in the upstream channel. ADSL is well suited for high-speed Internet/intranet access, video-on-demand, and telecommuter applications. ADSL speeds range from T1 (1.544 Mbps) and E1 (2.048 Mbps) to 6 Mbps and beyond downstream. Upstream return channel speeds range from 64 Kbps to 384 Kbps to 640 Kbps. ADSL transmissions operate at distances up to 5 km (between the customer and the local exchange or serving central office switching system) via a single copper twisted pair. The practical issue here is that the copper pair has to be nonimpaired, and bandwidth diminishes sharply as distance from the serving switch increases.
- *High-speed DSL (HDSL)* is a symmetric technology with speeds of 1.5 or 2.0 Mbps (upstream and downstream). Its main purpose is to replace traditional T1/E1 leased circuits. As per the standards bodies, HDSL is a two-wire implementation with an operating range somewhat more limited than that of ADSL-much over what 3-km telephone companies need to install signal repeaters to extend the service. Because HDSL is a two-wire implementation, it is deployed primarily for PBX network connections, digital-loop carrier systems, interexchange PoPs, Internet servers, and private data networks.
- *Symmetric DSL (SDSL)* is similar to HDSL in that it delivers 1.5 Mbps or 2.0 Mbps (or submultiples), but it does so using a single line, downstream toward the user and upstream. The use of a single line further limits SDSL's operating range; 10,000 feet is the practical limit for SDSL applications. Because of its symmetrical nature, it is well suited for videoconferencing applications or remote LAN access.
- *Very high-speed DSL (VDSL)* is asymmetric. Its operating range is limited from 1,000 to 4,500 feet but supports very fast transmission via single twisted-pair copper. Data can travel at rates up to 51.84 Mbps from 330 to 1,000 feet, with rates of up to 1.6 Mbps on the upstream return path. VDSL is positioned as the eventual modem of choice for fiber-based full-service networks. The extra bandwidth allows telephone companies to deliver high-definition television (HDTV) programming using VDSL technology.
- *Rate-adaptive DSL (RADSL)* automatically adjusts to copper quality degradation or can be manually adjusted to run at different speeds up to ADSL rates.
- *ISDN-based DSL (IDSL)* "inverse multiplexes" two ISDN 64-Kbps B channels using 2B1Q coding into one 128-Kbps channel.

Both ISDN and DSL turn the local network into a digital one and, in doing so, offer higher speed access to the end customer. By and large, though, both of these local network enhancement schemes have experienced no more than moderate success. ISDN and DSL have proved slow to deploy and have experienced patchy acceptance, part due to performance, part to cost, but mostly due to a weak regulatory environment.

5.2.2 EtherLoop

EtherLoop is a technology designed to provide high-speed data access to the home over the standard twisted pair of wires that constitute the local loop. Although created by a Nortel spin-off known as Elastic Networks, it is not a proprietary technology and there is a range of EtherLoop suppliers and products.

EtherLoop uses the basic concepts of DSL technologies and combines them with standard packet delivery system algorithms to provide a high-speed solution that overcomes many of the DSL limitations, without sacrificing speed or data quality. Through a combination of advanced signal modulation, burst delivery technology, and use of the Ethernet packet-data protocol, EtherLoop provides a solution that is simple to install, robust, and efficient in power consumption.

Because EtherLoop is Ethernet compliant, it is easily adaptable to established Ethernet systems. It allows for simultaneous voice and high-speed data communications at speeds ranging from 125 Kbps up to 6 Mbps with a range of distances up to 21,000 feet. Similar to Ethernet, it transmits data packets

in bursts. Between bursts, it looks for problems and interference in the lines and, if any are encountered, it finds an alternate path.

Since the local loop is essentially a point-to-point connection (between the local exchange and the subscriber premises), it is possible to define one end of an EtherLoop as a "server" and the other as a "client." If we say that the client only speaks when the server allows, this effectively eliminates all collisions. Also, because the subscriber loop is generally of lower quality in terms of BER than a private LAN, EtherLoop has to manage frame-error checking and retransmission, using the Ethernet checksum. This is typically not done on LANs, since the BER is low. It is, however, an appropriate action to take on a noisy copper telephone loop.

EtherLoop also takes advantage of the flexible symmetry available in a half-duplex Ethernet-type communication link. The amount of time spent transmitting in each direction is directly proportional to the amount of traffic being offered in each direction. If the user is downloading a large file, then most of the time is spent transmitting in the downstream direction. In this case, nearly all of the bandwidth is going downstream, with only a small amount of bandwidth being used to transmit acknowledgements upstream. If the user is uploading a large file, then nearly all of the bandwidth is being used in the upstream direction. If the user is engaged in a videoconference, then the bandwidth is split symmetrically. These are only a few examples of how the bandwidth can be split. Any combination is possible, as EtherLoop does not impose a fixed symmetry scenario on the user, but instead allows the user traffic to set the symmetry dynamically.

EtherLoop contains some technology that is unique. When the transmitter is silent, the quality of the signal is monitored. Cross talk and interference can be measured and the device can constantly change internal frequencies to reduce cross talk and avoid interference. This is used as a continuous rate adaptation technique. This rate adaptation allows the modem to immediately adapt to any noise it receives or generates and generally improves the quality of the modem's transmissions while reducing its interference with the other lines in the cable.

The combination of DSL signaling technology, burst mode delivery, and Ethernet frame technologies gives EtherLoop its competitive advantages. DSL technology allows the use of the existing telephone infrastructure, and Ethernet burst technology reduces interference and power consumption and provides a low-cost interface.

In summary, EtherLoop takes the best features of DSL technology, then uses burst technology to reduce the interference problems and the Ethernet packet-data model to reduce the cost of connecting to a data network.

5.2.3 Structure and Operation of EtherLoop

An EtherLoop configuration is relatively straightforward. Each subscriber has a modem that connects over the local loop to an EtherLoop multiplexer shelf product, which is EtherLoop ready. The voice and data channels are separated, and the voice channel is passed on to the switch. The data channel is passed on to an Ethernet hub, or switch, which then connects to any standard data network. Depending on the needs of the customer, multiple networks can be attached. For example, some users may wish to use the public Internet, some may wish to use an operator's regional broadband network, and some may wish to connect to private corporate networks.

This discrimination of service is typically performed by an Ethernet switch, which intelligently routes the data to the appropriate remote resource (either the Internet, or one of potentially several private networks), based on the destination. This is only one of many possible network architectures, as Ethernet networking is very well understood and there are a variety of ways in which data can be delivered in secure and flexible ways, including the following:

- Residential access. This is the most straightforward option. The end user will typically have only one device connected to the Ether-Loop link, which simplifies the overall architecture. The local exchange end provides Ethernet switching to deliver data to one of several routing resources, which may be connections to the Internet, to private Intranets, or to an ATM/frame relay transport network. The home user will typically be assigned one or more static IP addresses (addresses that do not change each time the user connects to the network), or alternatively have an IP address provided at boot-up via the dynamic host configuration protocol (DHCP).

- Small office. This differs somewhat from the residential access in that there will usually be a more extensive LAN in the subscriber's office than in the average residential user's home. In this case, the local exchange and customer modems must take on the additional responsibilities of bridging Ethernet traffic, keeping local traffic on the subscriber premises' side of the link. This capability is already built into the EtherLoop modems and, in this configuration, an external modem is more appropriate.

The small office option does require some additional resources [typically either virtual LAN (VLAN) or permanent virtual circuit (PVC)-enabled Ethernet switches] to correctly route traffic to the corporate network, or tunneling protocol software to encrypt packets that are delivered over a public network to the corporate access point. Each of these secure communications strategies has its advantages and disadvantages. Future EtherLoop products that will enhance this market include EtherLoop/Ethernet switches, VLAN switches, EtherLoop-based routers (with an internal EtherLoop NIC as opposed to an internal 10Base-T 10/100 NIC, external EtherLoop modem, and PPOE software), and ATM/frame relay interfaces.

5.2.4 Coding

EtherLoop is designed to use a range of frequencies, from roughly 30 KHz up to about 2.5 MHz for high-quality subscriber loops, although this is divided up into 12 overlapping frequency spectra, only one of which is active at any point in time. The lowest spectrum has a total frequency range of 62.5 KHz and the highest has a frequency range of 1.667 MHz.

Historically speaking, one hertz was equivalent to one symbol per second, which would give EtherLoop a theoretical maximum symbol rate of 1.667 megasymbols per second. Using standard modulation techniques that give one bit per symbol, this would translate into 1.667 Mbps. The bad news is that not all of the frequency spectrum can be used effectively because of cross talk, loop quality, and so forth. Hence, the maximum symbol rate is reduced. On the plus side, equalization technology can improve the bit rate considerably and at shorter distances (less than about a mile) that have less noise, allowing the symbol rate to reach the 1.667 megasymbol per second limit.

Furthermore, there are a variety of advanced signaling techniques that "squeeze" more than one bit out of one symbol. Several of these techniques have been used within analog modems to bring the maximum speed of a standard modem up to 56 Kbps. EtherLoop uses two related signal modulations techniques: quadrature phase shift keying (QPSK) and quadrature amplitude modulation (QAM) to extend the amount of information transferred.

5.2.5 Architecture

EtherLoop technology was purposely designed to follow the format of the Ethernet IEEE 802.3 standard as closely as possible. This has a dual purpose. First, it allows EtherLoop to appear like Ethernet natively, which reduces the product costs of performing a protocol conversion. Second, it allows EtherLoop devices to fit into an existing network, in almost any location where an Ethernet device currently operates. There are some limitations, however. EtherLoop is only appropriate in a point-to-point link, it provides a lower peak bandwidth on longer/lower quality loops, and it will always be more expensive than Ethernet.

These limitations are not as severe as they might first appear. Much of the existing Ethernet now in place uses the 10Base-T substandard, which requires point-to-point connections between a network device and a hub, so it is a perfect match for local loop. Peak bandwidth is important, but so is a combination of technical improvement and fundamental issues of distance and wire quality. Even conservative estimates of EtherLoop's capabilities indicate that it should reach 10 Mbps at 3,000 feet on a majority of existing wiring while simultaneously supporting an analog voice channel. The expense of EtherLoop is only a factor when compared with standard Ethernet-benchmarking against DSL is more favorable.

5.2.6 EtherLoop in Action

The Pinehurst Resort Country Club, site of the 1999 U.S. Open Championship, wanted to install a state-of-the-art communications network. This took the form of an EtherLoop installation in 30 guestrooms and two meeting rooms. Elastic Networks was the equipment supplier for this network,

and Sprint served as the ISP.

To connect to the Internet, Pinehurst guests now simply plug their laptop computer into the modem in the hotel room. The end-user solution does not require special set-ups, configurations, or protocol changes. The system uses existing telephone lines to provide simple, reliable Internet access for every hotel unit.

EtherLoop technology was an ideal solution for Pinehurst because it made it very easy for guests to get on-line and did not require any rewiring to a historic property-Pinehurst's primary building is nearly a century old and has been designated a National Historic Landmark.

In addition to guest convenience, EtherLoop technology offers hotel property management compelling benefits, including the reduction of PBX congestion, fast connectivity, and a greatly enhanced range of services.

Perhaps the most compelling aspect of EtherLoop technology in this application was its simplicity. Other competitive solutions would have proved much more difficult to install than EtherLoop, and less capable options would have resulted in wasted time setting up a computer or waiting for Internet access.

Team LiB

[← PREVIOUS](#) [NEXT →](#)

5.3 Summary

Gigabit Ethernet is already making inroads in MANs, where new technology has made it possible to transmit over long distances and therefore provide a homogeneous Ethernet solution in the wide area. Several service providers are now offering Gigabit Ethernet service in major metropolitan areas. Switching vendors, particularly Extreme Networks and Cisco, have been actively developing gigabit interfaces for WAN switches, as well as the QoS features required for Gigabit Ethernet bandwidth to be a commercial proposition for carrying both voice and data.

The key driver of Gigabit Ethernet in the WAN is price. The difference between established SONET transmission installations and Gigabit Ethernet is now so great that the former is being questioned, despite its reputation for high resilience. This advantage is likely to harden as 10-Gbps switches become commodity items with a consequent sharp drop in pricing. The fact that more than 50 vendors have joined a 10-Gbps Ethernet Alliance is ample evidence that 10-Gbps Ethernet has considerable momentum and will not be derailed.

And it is not just in the wide area that Ethernet is increasingly looking to be the network technology of choice. The version of Ethernet that operates over point-to-point wires-EtherLoop-allows the same basic information format to be used at all points in the network.

This chapter has discussed both Ethernet's expansion into the wide area and its sideways shuffle into the local loop. These two developments mean that Ethernet now has the capability to provide all of the telecommunication services that were once provided over disparate network technologies. And, because of the inherent compatibility across the whole Ethernet family, there is, after years of diversity, the prospect of a total area network.

Chapter 6: Storage Area Networks

Overview

One change always leaves the way prepared for the introduction of another.
--Niccolo Machiavelli

Most information that is carried over networks is valuable. It must be stored in a safe place and protected from disasters, big and small, that can and do afflict networks. Once upon a time, this meant installing a large amount of core memory in a secure location. This was fine most of the time, but all could be lost if a company experienced a fire, flood, power cut, or breakdown. An alarming number of companies that lose their data store suffer financial ruin within a few years. Many organizations have come to recognize this and therefore attach some importance to information storage and disaster recovery.

In recent years, we have seen the introduction of systems that distribute stored data across several devices and/or locations. By adding redundant data to the core information, a complete restoration can be made in the event that one (or more) devices fail.

Furthermore, storage across a number of physically separate sites offers a degree of protection from catastrophic but localized events, as well as the ability to recover (or continue operations) when an outage does hit home. Central to the development of high-capacity distributed storage systems is, of course, a high-speed network. The latest advances in Ethernet technology provide the ideal base for storage area networks (SANs).

As always, there is a price to pay. In this instance, greater availability, integrity, and resilience of data is bought with extra complexity. When many data images are kept in different devices and at different locations, careful thought has to be applied to maintaining that data. Organizations can all too easily replace the problem of vulnerability to equipment failure with the equally bad problem of installing a solution so complex that no one can use it. It is not enough to attach a set of big disks to the total area networks introduced in [Chapter 5](#). Something else must be added to manage the information that is held; for example, some software that enables users to keep versions and variants up to date, manage the configuration of related data, and so on. This chapter considers the evolution of SANs before going on to consider some of the practical issues in their use. In addition to covering the technical side of the coin, we also look into the economics of using a high-capacity distributed data store.

First, though, a little history.

6.1 The Growth of Network Storage

In the earliest days of computing, and for nearly two decades thereafter, virtually all information tended to be stored on the computer that used it. As little as 20 years ago, the most common way of taking a permanent copy of some data for backup was to write to a disk and put that disk in a fireproof safe or bank vault.

As LANs became popular, the need to use separate storage media declined. The network allowed users to place copies of their files first onto storage devices of server systems and later onto discrete storage devices that were connected to their networks. For many years, this was an ideal solution.

As companies rely more and more on e-business, on-line transaction processing, and databases, the amount of information that must be managed and stored has grown. It has been estimated that the proportion of the world's information that is being stored electronically has risen from 1% to 10% in just the last 5 years. The sheer volume of data that must be kept safe these days can intimidate even the most seasoned of network administrators.

While servers do a good job of storing data, their capacity is limited, and they can become a bottleneck if too many users try to access the same information. To counter the problem, many companies have come to rely on peripheral storage devices such as tape libraries, disk arrays, and even optical storage systems. These storage devices are effective for backing up data and certainly offer large amounts of storage capacity.

But as server farms increase in size, and as companies rely more heavily on data-intensive applications such as multimedia, the traditional storage model (i.e., just add another data silo if you need more space) does not work as well as it used to. The reason for this is that access to an ever-growing group of peripheral devices can be slow, and it might not always be possible for every user to easily and transparently access each storage device. Moreover, the problem of knowing exactly what you have stored, where it is located, and how you distinguish between and verify the current version from prior versions, has become very important. There are many instances in which it is vital to keep accurate records, ranging from software development (where a solution only works when the right set of files are used) to a discovery process in a law-suit (where a complete, documented audit trail needs to be kept). SANs are more attractive in these instances because they provide a firm base to which information management can be added. A number of companies, such as Vixel and Brocade, now provide tools that allow a distributed information base to be monitored and controlled.

Given this situation, a number of vendors from all walks of the industry have devised the concept of the SAN. These provide more options for network storage, including much faster access than the established network attached storage (NAS) model, as well as the flexibility to create separate networks to handle large volumes of data.

Before going into further detail on SANs, it is worth briefly looking at how other methods of adding storage to the network work and why the need arose to develop something beyond them.

6.2 Options for Network Storage

The most basic way of getting storage devices on the network is to hang disk arrays or other storage devices off of servers. This is typically done using an interface such as the SCSI to make the connection. SCSI is a ubiquitous and relatively high-speed interface that was developed more than 15 years ago.

The SCSI is a general-purpose input/output (I/O) interface used to connect storage-related devices such as tape and optical drives, as well as printers, scanners, and other external recording media. SCSI has undergone a number of changes over the years, primarily aimed at increasing the data transfer speed that the interface supports. SCSI was initially designed to handle up to 5 Mbps of data. In its current incarnation (known as Ultra3 SCSI), SCSI supports a considerably higher throughput rate-up to 160 Mbps. Ultra160 SCSI, a subset implementation of the Ultra3 specification, is gaining popularity among SCSI vendors as well as server and workstation manufacturers.

While SCSI has been a workhorse over the years for connecting peripherals at quite acceptable speeds, distance limitations have kept this particular interface from evolving more rapidly. The basic standards impose a length limit of about 5m or 6m between connected devices. While this limitation does not really affect the connection of storage devices directly to a server, it does severely restrict the placement of disk arrays and tape libraries at remote points on the network.

This is where the concept of NAS comes in. NAS is quite straightforward: disk arrays and other storage devices connect to the network through a traditional LAN interface, notably Ethernet. Storage devices would be attached to network hubs in much the same way as servers and other network devices. A typical setup for NAS is shown in [Figure 6.1](#).




Figure 6.1: The configuration for NAS.

NAS is effective in making storage resources readily available, and it helps alleviate the bottlenecks commonly associated with access to storage devices. It does, however, have a few drawbacks.

First, network bandwidth places throughput limitations on the storage devices. Most NAS servers are placed on well-established 10-Mbps or 100-Mbps Ethernet LANs, considerably slower than Ultra3's I/O transfer rates. Even if the network were to be running at gigabit speeds, most NAS vendors only offer interfaces up to Fast Ethernet.

Another limitation to NAS is that there is a lack of cohesion among storage devices. It is easy enough to add disk arrays and tape drives on to a LAN, but managing these devices can prove more challenging, as they are separate entities and are not logically tied together. The result of this is that data consistency/accuracy is not practically achieved with today's products. Hence, NAS is a useful addition to the storage toolkit and has its place as a viable storage architecture, but when it comes to the very large scale, something more is needed.

6.3 SANs

SCSI-based storage and NAS-based configurations are both important ways of bringing storage to the network, but they are best used in situations where the data stored does not have to be closely controlled. Large enterprises seeking to store and manage large amounts of information in a high-performance environment must look at another option. This is where the SAN fits.

In a SAN environment, storage devices such as disk arrays are connected to many kinds of servers via a high-speed interconnection, such as Fibre Channel. This setup allows for any-to-any communication among all devices that are connected on the SAN. It also provides resilience by providing alternative paths between the servers and storage devices. This means that if a particular server is slow or completely unavailable, another server on the SAN can provide access for a particular storage device—the switch shown in [Figure 6.2](#) can be used to select a default route and a preferred backup. In much the same way, the SAN makes it possible to mirror data, making multiple copies available. Perhaps most important is the fact that a SAN can be managed as a single entity; change control and file management can be consistent across all of the information stored.




Figure 6.2: A SAN.

The high-speed interconnection that links servers and storage devices essentially creates a separate network for data storage. This overlay network uses the same transmission technology as, and is connected to, the LAN, with all of its traditional functions. But it acts as an independent entity.

There are a number of advantages to SANs and the segmented environments they create within a network. To start with, they accommodate the addition of bandwidth without interruption to the operation of the main LAN. Also, they make it easier to conduct on-line backups without users suffering the consequences of a bandwidth-hungry application taking preference over their work. And, when more storage is needed, additional drives do not need to be connected to a specific server. They can simply be added to the storage network and, because of the topology used, be accessed from any point.

One further (and very important) reason why SANs have become prominent so quickly is that all the devices can be centrally managed. The desirability of managing a network as a single entity, rather than on a per-device basis, is explained in detail in [Chapter 7](#). The essential advantage of enabling the administration of dozens (or even hundreds of servers and storage devices) from one place is that it minimizes the cost of administration.

The interconnection technology used in most current SAN implementations is Fibre Channel (introduced in [Chapter 3](#)). It has increasingly been used as an alternative to SCSI in creating high-speed links among network devices.

6.3.1 Elements of a SAN

A number of components are now available for building SANs based on Fibre Channel. These allow a SAN to grow from an initial, and most likely relatively simple, configuration to a more complex facility. They also allow for the inclusion of other storage devices on to the network, such as SCSI drives.

A typical starting setup would be to have a group of server systems and storage devices connected by

Fibre Channel adapters to a network. As the storage network grows, Fibre Channel hubs can be added, and as the growth continues, Fibre Channel switches can be incorporated.

To enable this evolution, Fibre Channel can support several configurations, including the point-to-point and switched topologies indicated here. In most SAN environments, it is the Fibre Channel Arbitrated Loop (FCAL) that is preferred. This is ideally suited to the creation of an external, high-speed storage network, due to its inherent ability to deliver any-to-any connectivity between the storage devices and the servers (see [Figure 6.2](#)).

The FCAL configuration, as shown in [Figure 6.2](#), consists of several components, including servers, storage devices, and a Fibre Channel switch or hub. The loop operates the same way as a token ring-a station contends for access and, once granted, has a point-to-point connection. The switch adds resilience and ensures that each device operates as it should. The setup shown in [Figure 6.2](#) provides the basic network flexibility, but more features and functionality can be added.

Another component that might be found in an arbitrated loop setup is a Fibre Channel-to-SCSI bridge. As the name implies, this allows SCSI-based devices to connect into the Fibre Channel-based storage network. This not only preserves the usefulness of SCSI devices but does so in a way that several SCSI devices can connect to a server through a single I/O port on the server. This is accomplished through the use of a Fibre Channel host bus adapter (HBA). The HBA is actually a special-purpose Fibre Channel port. The Fibre Channel-to-SCSI bridge can multiplex several SCSI devices through one HBA.

The FCAL not only provides a high-speed interconnection among storage devices but also strong reliability. In practice, one, or even several, devices can be removed from the loop without any interruption to the data flow. Packets sent over an FCAL are error-checked, and, if need be, they can be resent if any are lost or corrupted.

6.3.2 Extending the SAN

Now that we have all the component parts we need for our SAN, we can look at their key features. To scale efficiently, SANs must be capable of interconnecting large numbers of servers with large amounts of storage without disrupting any other applications being used. The highest levels of scalability are possible only through networks of modular components that provide switching functions and services for attached servers and storage. This means that the ports on the SAN switches need to be capable of growing in increments as the network itself grows.

Adoption of Fibre Channel as the SAN transmission technology enables the creation of a "core-to-edge" network in which edge switches attach to devices at whatever supported speed the devices require, and traffic from the edge switches is routed through higher speed core switches. The core-to-edge network enables a very high level of scalability at incremental cost. In addition to scaling a SAN over a local area (such as a data center), the Fibre Channel option allows connections to be extended across a metropolitan network or long-distance link.

Fibre Channel has, in many ways, been the trigger to making SANs a reality, and future developments on the interface will likely bring more features, most notably, higher bandwidths.

6.4 Evaluating the SAN

As SANs begin to play a more prominent role in the network applications infrastructure of many organizations, it is important to consider just how they interact with the overall environment. As part of this process, organizations need to ask the following crucial questions:

- *Is it a proven solution?* A SAN solution should be field-proven in a wide range of industries and application environments and supported by a variety of server and storage vendors. It must also deliver the throughput and latency performance required by server-to-storage connections.
- *Is it scalable?* A SAN configuration should be able to scale independent of storage and server capacity-without disrupting applications' ability to access data. Likewise, a SAN must provide excellent configuration flexibility and extended distance connectivity.
- *Is it manageable at a reasonable cost?* A SAN provides a centralized point of management for many storage functions, and the consolidation of independent storage devices helps reduce management expenses. However, a SAN also represents an additional entity to be managed and therefore must require a minimum of direct management or ideally be manageable by existing tools that manage applications, storage, or servers.
- *Is it highly available?* Today's most demanding enterprise applications require continuous availability, with downtime limited to seconds or minutes per year. Built-in redundancy and the high-availability characteristics of a networked storage approach are critical.
- *Is it secure?* As SANs grow in size, or are used to support multiple customer environments, they must provide auditable security mechanisms that prevent unauthorized access to data. This requires specific mechanisms to secure access to SAN management functions, reliable authentication of devices, and prevention of network intrusion.
- *Is it heterogeneous?* No single storage or server vendor can meet the entire range of application requirements for all possible business needs. A SAN should be designed with interoperability and support for the widest range of devices, and it should be based on open industry standards.
- *Can it accommodate future technologies?* A SAN should be based on an architecture that can grow and adapt to new requirements and technologies. In addition, the SAN should be controlled by a set of management services that can be extended to accommodate new storage protocols.
- *Is file administration enhanced by the introduction of the SAN solution?* Can the management features of the SAN you choose assist you in file administration features such as configuration and change control, appropriate document retention and archiving, and the removal of superseded or obsolete documentation?

Ultimately, the benefit of any technology should outweigh its cost. It can be difficult to put a direct value on some of the advantages that a SAN has to offer (the level of security, for instance), but attributes such as availability can be used to derive a direct measure of benefit.

One of the obvious (but often overlooked) questions that should be explored when talking about SANs is how much benefit they actually offer. It is all very well to say that it is important to safeguard data, but to what extent does this justify the cost of installation?

One of the most basic (and appealing) features of a SAN is that it offers a higher degree of availability than a less-resilient storage solution (i.e., one with fewer alternative connections). If we say that the SAN improves the availability of network data from 99.9% to 99.99%, this equates to an extra 8 hours per year. The business benefit that accrues from this obviously depends on the number of people who need to use the stored data, the number of transactions they carry out, and the value of revenue generated by those people.

This does not tell the whole story. In addition to the gains from less-frequent downtime, there are also likely to be gains from faster response time. As before, the user profile will determine just how much a business is impacted by improved availability.

One final factor that must be taken into account when evaluating the business case for a SAN is

recovery after a failure. The extent to which business is lost when there is an outage is difficult to quantify but is, intuitively, a major consideration.

Using this approach, it is not unreasonable to build a business case that justifies a SAN (or not, as the case may be). A similar approach can be used to analyze the merits of installing management software once the SAN is established. As previously explained, the capital cost of the solution is offset against the savings on administration and the speedier resolution of network faults.

Team LiB

[◀ PREVIOUS](#) [NEXT ▶](#)

6.5 SAN Management

The management of a SAN involves looking after multiple layers that interact with each other and all of which need to be operational for the whole thing to work.

At the lowest level, the concern is interconnection and transport, and this is dependent on the hardware that supports higher level features. The mechanism predominantly used for this layer is the simple network management protocol (SNMP), which can be used to define agents that can report status and respond to commands.

At the highest level, an enterprisewide management platform is needed to oversee a variety of functions and is fed status information from multiple networking and storage infrastructures. The interfaces between different management layers may be as straightforward as event logs or SNMP traps, or as sophisticated as application programming interfaces (APIs) and common information models (CIM).

Throughout the hierarchy of management layers, the common charter of all management functions is to maintain data availability. For all enterprise networks, access to data is as essential as dial tone. Loss of data access disrupts communication, delays business transactions, and, ultimately, results in lost revenue. For SANs, an efficient high-speed transport has little value if data become corrupted as they are recorded to disk due to a faulty algorithm on the storage side. Likewise, a fault-tolerant, high-capacity storage array cannot fulfill its function if data is corrupted as it is transferred across the SAN due to poor signal quality.

Effective management of the SAN therefore requires visibility to all of the constituent aspects of data transport and storage. In light of this, partnerships between storage management vendors and SAN interconnect vendors are driving integration of management functions for more comprehensive SAN management applications.

The scope of the SAN and the sort of support arrangements already in place fundamentally determine just how much management is needed. For instance, a fairly small installation of somewhere between 5 and 20 storage nodes can get by with minimal support. The administrator may simply want to be notified when a problem occurs so that he can then call a customer service number of the solutions provider to replace or repair the failed unit. If the SAN interconnect provides some sort of autorecovery functionality, it can remain operational while the remedial action is taken. If the hub or switch also supports advanced diagnostic tools, these will be valuable for customer service for further diagnosing problems.

Large enterprise networks, however, often have their own in-house experts who are trained to deal with more extensive and complex network configurations. These experts would, most likely, use Fibre Channel analyzers and Fibre Channel protocol testers as their basic tools. In this instance, the greater the management capability, the better. Products that provide autorecovery capabilities and advanced diagnostics via management software offer the means to quickly identify more complex problems and ensure higher availability of the system. The trend in Fibre Channel hub and switch design is, therefore, toward increased SAN health monitoring capabilities and more sophisticated diagnostic tools.

Small and medium SAN installations may not have the requirements, resources, or budget to install these larger management frameworks, such as Tivoli or CA Unicenter. Tape backups may be performed through operating system utilities instead of more sophisticated backup managers from Veritas or Legato.

As SANs become more complex, though, the combination of hubs and switches in a single environment is more easily managed if there is a graphical interface from which an operator can solicit the status of both types of devices. Thus, a large SAN can be more readily managed if the operator can monitor both loop hubs and switches from a single screen and get a consistent look and feel for the various parts of the data transport layer.

Integration to upper level storage and systems management platforms gives the administrator additional flexibility. A commercial storage management platform from suppliers such as Veritas or Legato can provide the tools for data organization, tape backup, setting the configuration of disk

arrays, and so on. In addition, these tools can also be configured so that the hub or switch device manager software is launched from the main program. This gives the operator a complete view of the SAN-both the network it uses and the data it transports. Hence, the consistency of data across an organization can be monitored in the same way that an individual might care for the files on her local machine. This makes it possible to exercise an organizationwide information policy (covering file naming, versioning, approval, archiving, and retention).

Team LiB

[◀ PREVIOUS](#) [NEXT ▶](#)

6.6 SAN Evolution

While the SAN architecture does appear to be the next logical step in the evolution of network storage, there are a few points that need to be more adequately addressed before they become more widely used.

One crucial component that must in place for a large SAN is software that administers and controls all of the devices that are connected to the network. While a SAN configuration inherently makes management easier than in the case of NAS, most companies will install some sort of management (often customized) application to keep their SAN in order. Put another way, the tools to manage the technology are lagging behind the technology itself.

A related issue is that established NAS is often installed alongside a new SAN. Both do a useful job, but they do not interact—information kept on the NAS cannot readily be backed up over the SAN, so there is a divide in the storage process. This not only introduces considerable frustration for users but is also a potential cause of more serious data consistency problems. For this reason, a number of vendors have developed NAS/SAN gateways. With these devices, SAN-based arrays can be attached to a diskless NAS head. Software residing on the NAS head then distributes SAN or NAS data over the appropriate network, depending on the size of the data and the best path. Hence, the divide in storage technology is closed and a consistent data storage strategy can be followed.

In a relatively small SAN implementation, customized software can be written to ensure communication among all devices. But as SANs grow, and as more vendors enter this space, simply writing customized pieces of software will not be a sustainable solution. There must be a standard way for components from different vendors to interact within the context of a SAN. It is fairly typical for standards (as well as management systems) to lag behind the introduction of new equipment. In some ways this can be a blessing, as it allows the market to choose its preference before deciding on the favorite option.

The problem comes when a market with many competing proprietary solutions appears. Vendors in the storage, and specifically the SAN, market have realized the importance of this issue. Progress is being made both through vendor-neutral organizations and traditional standards bodies (see the "[Selected Bibliography](#)" at the end of this chapter). Two of the leading consortia in this area are the Storage Networking Industry Alliance (SNIA), an industry group of more than 75 companies working on storage networking, and the Fibre Channel Industry Association (FCIA), which develops equipment conforming to the American National Standards Institute (ANSI) Fibre Channel standard.

As we have already seen, there are a number of different components to a SAN. From the very name, you need some network and some storage devices. And, as previously indicated, a SAN must be properly managed if it is to be useful. Traditional storage and disk manufacturers, such as IBM, have a significant presence in the market, but there are also specialist SAN providers such as McData and Brocade. SANs may require a bit more thought and planning than simply adding one storage device to one server, but as companies wrestle with reams and reams of information on their networks, a high-speed alternative that is always available looks to be more and more appealing.

Before moving on, there is one further point to discuss regarding SAN evolution, and that is the way that SCSI, mentioned earlier in the chapter, has evolved. The latest incarnation of this standard is iSCSI, which is an IP-based storage standard that keeps the familiar SCSI controller intact. The deployment of iSCSI allows the creation of networks of storage devices that can be reliably accessed by hosts and servers using the ubiquitous TCP/IP protocols. With iSCSI, the TCP connection acts like a storage bus in a block storage (SCSI) I/O operation.

One of the attractions (quite apart from familiarity) of an IP-based storage network is that tunnels can be used to connect geographically distributed Fibre Channel *anywhere* they can find adequate bandwidth to satisfy storage application requirements. This could be realized using the evolving Fibre Channel over IP (FCIP) protocol, but there is also a proposed protocol (iFCP) for connecting Fibre Channel storage devices or SANs using IP instead of Fibre Channel switches. This means that organizations in metro areas where Gigabit Ethernet service is offered can look to IP storage as a means of satisfying business continuity requirements. While technology constraints currently hamper such deployment, the Gigabit services explained in [Chapter 5](#) will eventually extend this reach nationally and internationally. It would be wrong to suggest here that IP will make Fibre Channel

redundant. IP may be the long-haul preference, but both technologies have their place and are likely to be with us for the near future.

Team LiB

[◀ PREVIOUS](#) [NEXT ▶](#)

6.7 Devising a Storage Strategy

Before closing this chapter, we return, briefly, to consider the wider question of where SANs fit. There is little doubt that the amount of data that must be stored in electronic format is rising rapidly. Virtually everyone has some tale of woe relating to lost files or difficulties in archiving. The question that must be asked at this point is what you should do about it. In other words, What sort of storage strategy is needed?

A comprehensive storage strategy should address all aspects of storing and retrieving data. The main elements in the strategy should include the following:

- Connectivity method that gives computing elements access to data;
- Security and integrity arrangement for that data;
- Performance demands for access to and archiving of data;
- Likely scale of future storage requirements (and the expected rate of growth);
- Plans for developing relationships with storage vendors and management software providers;
- Proposals regarding which technologies to adopt;
- Plans for how stored data should be structured and what sort of configuration records should be kept.

Most of these issues have been addressed to some extent in this chapter. Some issues are wider than storage alone—they apply to all network-based elements and applications. So, although important, a storage strategy is only part of the overall picture and must be viewed in context with any other network, management, and application plans. The closer the SAN implementation follows established network standards (such as Fibre Channel), the easier it will be to keep in step with mainstream market offerings.

6.8 Summary

Most of the schematic representations of a LAN show a storage device as one of a number of attached devices. It is common practice for a working team to use a server as a central store for all their files. But this picture is changing. The growing amount of information that must be stored and the emergence of the total area network highlight the deficiencies in the idea of NAS.

The way in which storage techniques have evolved is explained in this chapter. The current stage in this evolution is SANs, which aim to overcome the speed and management limitations of NAS. Currently, the target users of SANs are organizations with large amounts of data and a diverse user community—corporate networks and e-business operators, for instance. The relative simplicity and component approach in SAN construction means that they will likely be a widely deployed option.

The objective of SAN suppliers and developers is to give servers the same sort of access to peripherally held data as is now possible for local data—and fast enough for streaming video. With a SAN, each peripheral device is connected to more than one server using a flexible ring/mesh network. Each connection within this network is a high-speed, typically Fibre Channel, link, operating at around 1 Gbps, with the range of each point-to-point connection around a few tens of kilometers.

Because of the high-speed connections used and the very resilient network configuration, SAN-connected devices can use any available server, and servers can have their data backed up on one of a number of separate (but still networked) peripherals. However, the true effectiveness of a SAN is not measured by improved bandwidth, reduced delay, and better availability, but by the software that allows an organization to extend and improve control and coordination of stored information. Hence, if Ethernet is the basis of the total area network, it would seem that SANs are the basis for total area information management.

Selected Bibliography

Clark, T., *Designing Storage Area Networks: A Practical Reference for Implementing Fibre Channel SANs*, London, England: Addison Wesley Longman, 1999.

<http://www.enterprisestorageforum.com>, (One of a number of useful portals that lead to information on the how to build and use SANs.)

<http://www.fibrechannel.com>, (The Fibre Channel Industry Association, for tutorials, white papers, and frequently asked questions on Fibre Channel and how it relates to SANs.)

<http://www.snia.org>, (The Storage Networking Industry Alliance, for useful introductory material on SANs.)

Chapter 7: A Changing Marketplace

Overview

There is nothing more difficult than to lead in the introduction of a new order.
—Niccolo Machiavelli

In the coming years, the entire communications market, and the environment in which it operates, is likely to be quite different from today. For any organization that trades on information and relies on the technology that supports it, an important part of being successful will be backing the winning technology. Investment in the wrong technology not only costs money, it can also stop a business from changing as quickly as it needs to, thereby putting it at a disadvantage against its competitors.

We have explained how flexible Ethernet technology can be, to the extent that it appears capable of almost anything. Born as a very high-speed local network, it has also emerged as a delivery method over long distances and wireless/mobile, as well as fixed-line media. Ethernet has even been demonstrated to deliver information using barbed wire as its transmission medium (pray that we never find a practical application for this configuration!). Wonderful though this is, having the ability to do something does not imply that Ethernet will actually succeed in all competitive markets it seems capable of satisfying. From a business perspective, the capabilities of a technology are of academic interest—it is market viability that matters.

In this chapter, we consider the extent to which Ethernet is likely to undermine the established technologies in those areas that now find Ethernet among the competitors. We look at two very different areas where Ethernet promises to play a major role. The first is as an alternative way of building the established enterprise or virtual private networks (VPN) that large organizations use to support their internal operations. In this case, traditional operators have deployed a number of technologies that, when linked together, offer a solution. Ethernet, as the new kid on this block, promises to make significant inroads into the market by solving many of the problems that encumber [heterogeneous](#) private networks.

The second area where many see Ethernet as a key technology is as a provider of mobile data services. In this instance, there is no real established market [if we discount the less-than-successful wireless access protocol (WAP) service], so the discussion is based more on expectation than any sort of facts and figures. Nonetheless, there are plenty of signs that can be used to arrive at a believable forecast of things yet to come; Ethernet-based wireless LANs have already provided several orders of magnitude more popular than WAP.

To complete the chapter, we consider how applications that capitalize on the strengths of Ethernet might be built. In particular, the constituent elements and overall structure of emerging Web services are explained.

7.1 The EtherCLEC Play

Modern organizations need to function faster, better, and more efficiently to stay competitive. Increasingly, this entails marshaling resources that are distributed across a country or even across the world. For instance, the European Airbus consortium manages the design and construction of planes despite being distributed across four countries. Many finance and communication companies deal with transactions that involve tens of countries.

Paper, fax, and telephone are no longer sufficient. There is a clear need for computer and communications systems that allow people to work together at a distance as easily as if they were in the same office. The flexibility and suitability of these systems is a vital differentiator in today's business world.

So, it is more and more in the interests of any high-tech organization to ensure that the computing and communications services they acquire and subscribe to satisfy their operational needs. Organizations should provide a uniform and flexible backbone that eliminates the distance element from information exchange and retrieval.

Regrettably, the practicalities do not quite live up to the expectation. Building a network to serve an organization or business can be something of a messy business, one that usually involves the bolting together of various network services from one or more operators.

The type of network that we are concerned with here is known as an enterprise network or a VPN. It takes a considerable amount of skill to assemble one of these networks. The end-to-end service expected from a VPN invariably relies on a number of suppliers—for example, tail circuits that connect branch offices from Colt, a national backbone from Interoute, switches from Nortel, multiplexing equipment from Kentrox, management software from Metrica, and international links from Equant. The list could well be extended. The simple fact is that the traditional VPN is built from many component parts, each with its particular characteristics.

As well as being difficult to specify and complex to operate, a VPN based on heterogeneity is not cheap. The recurring (monthly) charges on the set of managed local, national, and international circuits that comprise the network is usually very high; a short (5-km) link of 34 Mbps across London costs several tens of thousands of dollars a year. Furthermore, if you want to get high levels of availability, most suppliers will be happy to offer a premium service-at a price!

With all the difficulties and expenses in obtaining high-quality networks, customers are always searching for something new and better (that is, simpler and less expensive) on the data networking front. Ethernet technology clearly has something to offer. The extent to which it is capable of providing a seamless total area network has already been covered in [Chapter 5](#). So, our focus now is on the commercial value proposition that Ethernet brings to the market. We start this by considering the following generic high-level design objectives in any enterprise network.

Connectivity. A VPN must obviously interconnect a number of business locations. Organizations must consider many factors, such as intersite connectivity (topology), network and individual link capacity, availability, and redundancy, before they can determine which of potentially several solutions best satisfy "service criteria." For example, is connectivity required between all locations or just a few? Is the same capacity required at all locations, and is it required at all times or intermittently? Do we require path diversity into the network to assure persistent availability in the event a link fails?

Scalability. Ideally, the network should scale incrementally, without incurring considerable delivery delay or investment in new equipment. In today's economy, it is simply not acceptable to wait months (or pay) for a DS3 circuit when a bandwidth limit is reached.

Scalability must be considered not only in terms of bandwidth, but also in terms of the growth at specific locations and expansion to new locations. It should also be possible to independently expand different parts of the network. The whole network should scale without having to dispose of existing hardware (known as the forklift upgrade option) or infrastructure (having to add new circuits when capacity requirements change).

Cost. Enterprise networks usually come with a very high cost of ownership. Although central to the

effective working of many organizations, there is always pressure to reduce the level of spend on communications. It is not difficult to understand why. Many items of network equipment (switches, transmission kit, servers) are expensive to buy and require frequent upgrade and/or replacement. Communications services (bandwidth, managed applications) are a continual drain and tend to be even more expensive when anything out of the ordinary is required.

Flexibility. Change is inevitable. The VPN must be flexible in adapting to new locations, accommodating shifts in business (and, hence, in bandwidth requirements) from one office to another, and in taking on new requirements.

It is not enough simply to allow for change; the reconfiguration of an enterprise network is often required at short notice, as the business demands. Long lead times will stifle the ability of businesses to react quickly and capitalize on new opportunities.

QoS. The value of a network in conducting the business of an enterprise is strongly related to its performance-its availability, latency, and throughput. Many operators offer an SLA for their network that contains a guaranteed level of performance that, if not met, results in some form of rebate to the customer.

Management. The more complex a network is, the more difficult it becomes to support. Any network that is reliant on many different technologies, protocols, and interfaces quickly becomes complex. This, in turn, requires highly skilled network administrators and engineers. One of the easiest and most direct ways to reduce network support expenses is to simplify the network. The most commonly applied method of simplifying a network is to apply, as practically as possible, the "reduce to one" theory: one switch vendor, one operating system, one cabling medium, one carrier ... and even one transmission technology.

7.1.1 The Problem with Today's Enterprise Networks

While designing and building a LAN is typically fairly straightforward, connecting LANs between different locations certainly is not. The complexity of this task traditionally arises from a number of factors, including the following main culprits:

- Limited options for obtaining connection between sites from a public service provider;
- Lack of (granular) scalability in these options;
- Complexity of configuring and administering the various layers of protocols required;
- Recurring costs of the connections themselves.

Many of the established communications service providers (in particular, the incumbent operators such as BT, France Telecom, and so on) have infrastructures based on voice-centric technologies and, thus, have traditionally offered data services that are largely an overlay on a network fundamentally designed around voice delivery.

Some of these technologies-most notably frame relay, ATM, and private lines-have been adapted to meet the needs of the data world. But each of them actually consists of a number of protocols, encapsulations, compression schemes, and physical interface requirements. All have a very limited range of scalability and, because of the mix of technologies used, tend to offer service characteristics that do not complement LAN service characteristics. For instance, the end-to-end delay offered by most enterprise network suppliers for a national megabit network is on the order of 100 ms. A comparable figure using Ethernet would be around 20 ms.

In addition to this, the configuration and administration of the devices that use these services is typically intertwined with the mix of technologies that combine to provide the end-to-end solution. Finding the optimal combination is equivalent to solving a multivariant problem with numerous boundary conditions and large discontinuities-a task that is not trivial and one that has to be revised whenever some requirement or business assumption changes.

So the challenge of the enterprise network designer is to meet the high-level network objectives in the face of restrictions imposed by available technology. As we have seen, this is not a trivial task. It typically requires the expertise of many networking professionals and technical consultants, careful

evaluation and predictions about future business growth, and many trade-offs based on best-estimate assumptions.

Today, we have Ethernet until the first mile. For an end-to-end solution, the VPN designer has to buy expensive switches/routers, with expensive adapters that work with established public data network technologies, notably frame relay and ATM. Bandwidth and latency mismatches are virtually guaranteed, and the overall solution requires management of not one but several transmission services, so expertise is needed for all of them. Once installed, any changes to the network take time and careful planning. They also take money-sometimes an amount out of proportion to the incremental benefit received. There is plenty of motivation to seek an alternative-as anyone who designs or pays for a VPN will undoubtedly confirm!

7.1.2A Better Option

Life would be a lot easier if the WAN protocols, links, and associated complexity could be eliminated entirely. A network built entirely from one technology-Ethernet-could meet all of the objectives for an enterprise network listed earlier. With links scalable from 1 Mbps up to 1 Gbps with no change in hardware or protocols, simpler traffic engineering and performance evaluation, and low cost, it is easy to see the attraction.

With Ethernet, many of the problems highlighted in [Section 7.1.1](#) disappear; the scaling from 10/100/1,000 is, in some cases, achievable with autosensing, and where it is not, the incremental cost from one bandwidth rate to another is a fraction of the cost of the ATM alternative. The burden on the designer is also reduced, as much of the traffic engineering can be done in software rather than redesign and reconfiguration.

The past few years have seen an abundance of relatively inexpensive Ethernet switching and routing hardware come to market. This equipment is fully wire speed and nonblocking, providing zero packet drop performance in nonoversubscribed applications. The first few examples of Ethernet as the technology for total area networking have been demonstrated by suppliers such as Yipes and XO Communications, both of whom offer seamless LANs, MANs, and WANs.

It has to be said that the picture has yet to fully develop. Despite being acclaimed for their offerings, the so-called EtherCLECs have yet to show that they have a sustainable business model, but the economic conditions at the time of publication were not favorable for *any* telecommunications company. However, they have posted the writing on the wall for established operators, and it is only a matter of time until their business is undermined by a more flexible option.

7.2 Rabbit's Revenge

In the fixed network environment, there is a very clear distinction between the elegance and simplicity of Ethernet and the awkward complexities of building a network from a range of inherently mismatched technologies. Things are not so clear when it comes to the mobile market. Instead of having the contrast between something slick and new against something old and well worn, we have two unproven newcomers.

The popular contender to deliver the sort of connectivity that supports high-speed mobile data is third generation mobile, or 3G [also known as universal mobile telecommunication service (UMTS)]. A huge amount of resources have been and continue to be poured into the provision of third generation services. As one would expect with a mainstream telecommunications technology, much of these resources are focused on agreeing standards, making sure the network is scaleable, and attending to operational and support issues. 3G, when it arrives, will be of industrial strength-available for mass distribution in a public market.

But despite the glossy media front, the prospect of 3G cellular providing *the* universal high-speed data service is being widely questioned by analysts and technology experts. Alternatives, most notably the 802.11b wireless LANs, are being suggested. In this section, we examine the characteristics of both technologies and suggest how they might complement one another. [Table 7.1](#) summarizes the basic capabilities of technologies that seem the most likely contenders to provide the mobile Internet.

Table 7.1: Wireless Technology Ranges When Deployed in Urban Areas, Providing Indoor and Outdoor Coverage

	Max Speed	Range
3G Urban Cell	384 Kbps	500m
Bluetooth 2.0	2 Mbps	50m
IEEE 802.11b	Up to 11 Mbps	100+m

The focus of the rest of this chapter is on 3G and IEEE 802.11b, as they provide network connectivity, as opposed to Bluetooth, which is, essentially, a wireless cable substitute.

The technical challenge of building a next generation cellular network is proving difficult for even the most advanced wireless carriers-nearly all of the major telcos have delayed their plans for third generation services. These setbacks, combined with the huge cost burden of 3G licenses, have allowed other technologies to put the squeeze on 3G. In particular, enhanced data capability on existing mobile networks has the advantage of lower deployment costs, and wireless LANs are being designed to interconnect portable high-speed devices over limited (but still very useful) distances.

In many ways, the existing GSM/General Packet Radio Service (GPRS) network promises to provide what many users want, which is good voice service, a modem speed data connection, and Short Message Service (SMS). Wireless LANs, with their high speed (around 10 Gbps) and proposed metropolitan coverage in some markets (up to 40-km radius with multiple base stations, each with a practical outdoor range of around 100m), will satisfy another large part of the user population. In particular, they appeal to those who want a fast connection to local services. Several airports, conference centers, and hotels have adopted wireless solutions for exactly this purpose.

It is a little more difficult to anticipate what 3G will deliver, but few really expect it to be the panacea that was once hoped for. There are good reasons for this. The first is that 3G is a shared transmission media. This reduces the speed available to each user ranges from a theoretical 1 Mbps to a more realistic few hundred kilobits per second. Furthermore, battery technology will set a fairly low limit on the time that a 3G-enabled mobile phone can be active (and will probably cause it to be rather too hot for comfort). And some might say that no one in their right mind would want to monkey with a telephone keypad for mouse and text input, anyway!

By comparison, 802.11 in a MAN deployment scenario meets-and exceeds-most users' data service

requirements. Bandwidth in nearly every deployment is better than 3G's paltry few hundred kilobits per second. NICs are readily available and inexpensive. They draw less power from a more capable (laptop) battery arrangement. People get to use all their keyboarding skills with a familiar user interface, and you can do voice *and* video. The story is one of technology adapting to users, not users adapting to technology.

So, if 3G is not going to be the panacea for a new world of mobile applications, what is likely to happen? Taking history as our guide, the probable way forward will be for several technologies, some existing, some new, to come together in some complementary manner, to meet near-term market needs, and for market needs themselves to evolve to take advantage of the developments possible under such a scenario.

Looking back through time, it is always easy to see why particular developments happened when they did, as alternative scenarios are neatly eliminated from consideration by the course of history. However, looking to the future, it is very difficult to predict what will succeed. There are many credible options, and it generally takes two or three random developments (often in unconnected fields) to come together in a fortuitous way to make one option take off and another crash.

Perhaps the best way to see the road ahead for mobile data is to look again at the need-connection, not content. Without an intuitive, inexpensive, *broadband* (i.e., megabits per second) connection, few service innovations and fewer compelling or "killer" applications will emerge into the market.

And when it comes to connection, there are two sure-fire winners in the data market-Ethernets and IP networks. Both have become popular choices because they are simple to adopt and affordable. Their universal acceptance makes them a first choice for any network.

The unique feature of mobile business mentioned (fleeting) already-is that it is largely dependent on and driven by user location. This means that it is, predominantly, a localized service and that rapid fulfillment is key. Delay is critical, bounded by the amount of time the mobile user remains within reach of a local service distribution point (a wireless LAN AP, for example). Users will not stay in one place long enough to wait for a centrally dispatched FedEx van to deliver what they want.

This means that a significant number of services will be relatively short range, with all transactions completed between the user and the nearest provider. Of course, the full range of m-business communication via the Internet is also necessary and so the network must cover the wide area (e.g., for the user who wants to access information on a remote server). But the local area is the fundamental building block.

Wireless LAN technology (notably the IEEE 802.11b standard that has now moved ahead of its erstwhile rivals, such as HiperLAN) is the current favorite to satisfy the demand for mobile data services. But is this just the next example of technology hype or is there a real business case for it? The bulk of press coverage promoting the 802.11 standard has tended to concentrate on the technology itself. While many reports quote headline grabbing potential revenue figures, there is rarely any attempt to outline how this income might actually be realized (see, for example, the INT Media report cited in the "[Selected Bibliography](#)" at the end of this chapter).

We can gain some practical insight here by looking at Finland's Jippii subsidiary, Wireless Network Solutions, as it is aiming to provide widearea coverage using IEEE 802.11-based technology. In fact, Jippii already operates a high-speed Internet access solution, called Freedom, in more than 100 locations in 50 Finnish cities, with plans to rollout the service overseas. Service is currently offered in business districts, business centers, hotels, airports, restaurants, and multitenant office buildings. The first extensive urban area has been covered in Seinäjoki, with plans to cover other residential areas such as south of Espoo.

If this type of solution is going to be a success anywhere, then Finland (a country with a seemingly insatiable appetite for all things mobile) is a safe bet. While accounts of Jippii's Freedom solution are strong on technical capabilities, there is no mention of the number of users it has attracted or, quite crucially, how they are billed.

The key issue here is that 802.11 was developed as an alternative to traditional cabled LANs and only defines the air-interface between a wireless modem and a base station hub. When used in the environment for which it was designed (i.e., private LANs with known users), the solution is perfectly adequate, but introduce public users and a whole raft of new issues need to be addressed. [Figure 7.1](#)

illustrates some of the support facilities that would need to be provided.

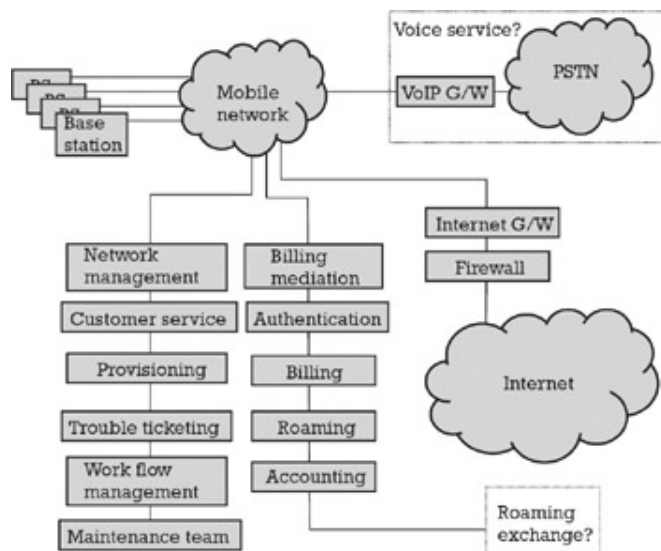


Figure 7.1: The range of operational support systems that need to be provided around a network.

In [Figure 7.1](#), users attached to the base station (top left) are part of a public access system. For any sort of commercial service, they must be provisioned, managed on an ongoing basis and, most importantly, billed. The facilities to carry out these tasks are part-and-parcel of cellular networks, but (initially proprietary) implementations would need to be developed for wireless LAN solutions, in just the same way that they have evolved for cellular networks. Articles comparing implementation costs often (conveniently) overlook the cost of "infrastructure." The importance and complexity of operational support systems (OSS) should not be underestimated—a fairly detailed treatment of OSS is given in [Chapter 8](#).

So, for all their appeal, the extent to which wireless LANs could actually supplant cellular networks (at least in the short term) may be governed by the balance of user requirements between speed and the more traditional telecommunications network features. They will doubtless provide part of the mobile data scene—the emerging model for 802.11 service is similar to that for the small-tier ISPs of the 1980s, with infrastructure concentrated in specific cities. For subsequent expansion, there are practical issues to be addressed, as we have indicated, and a lot of design effort needs to be expended along the way.

7.2.1 New Order

Given the brief comparative assessment of two candidate technologies, let us start to paint a picture of everyday life in the near future. Given the rapid uptake of personal digital assistant (PDA) devices such as the iPaq, it will probably not be long before most people will be carrying around a handheld device that is triple homed. Such a device will be able to talk 3G to a mobile operator, IEEE 802.11b for a local network provider, and Bluetooth when it needs to provide a cable replacement.

In this scenario, high-speed local connections will be routinely used for visual communications and other content-rich applications. Remote connections would be provided through an Internet gateway on the wireless LAN. Since these usually require a user to sit down, the connections will be made inside the home or office or near some (well-appointed) conurbation—all of which is well suited to the use of a wireless LAN. This may be a slightly restrictive view but just imagine someone playing multiuser Quake on his Gameboy while driving. On second thought, please don't!

When people move around, they would probably use a wide area connection, linking to one of the core 3G (or, perhaps, GPRS) networks. This parallels today's broadband connected home user or teleworker, who grudgingly uses a V.90 modem while traveling in a less civilized part of the world than what he is accustomed to. Like the modem connection, 3G/GPRS will prove good enough to pick up a specific (and usually temporary) piece of information that you cannot wait for. Some data, like travel alternatives or weather conditions at an intended destination, is needed "on the move." But most of the

time, the choice would be to use high-speed LAN access, which provides the necessary bandwidth (and carries 80% of local traffic).

How do these fit together? The full answer to this has yet to emerge, as many of the operational issues raised in [Section 7.2](#) (e.g., end-to-end billing, fulfillment, assurance) have yet to be resolved. At the networking level, both technologies peer at the IP level, which can give a specific user seamless access to both central and remote services, as shown in [Figure 7.2](#).

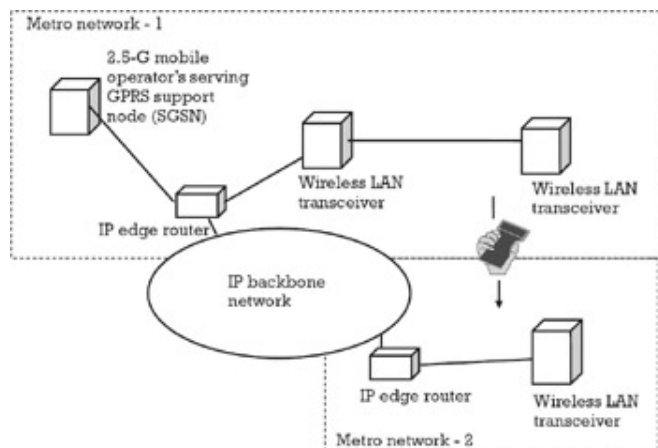


Figure 7.2: Roaming across a wireless LAN network.

In this illustration, it would be the mobile IP protocol that provides the "glue" that manages the end-to-end connection as the user roams from one place to another.

7.2.2 Reality Bytes

In theory, theory and practice are just the same. In practice, they are not. There are many practical issues in the deployment of any network. This section considers a few of the areas where unwary providers could easily find themselves mugged by reality.

Comparing technologies often helps to highlight the potential implementation costs of providing the network coverage needed to support the latest services, most notably m-business. The high bandwidth that IEEE 802.11 offers and the requirement to reduce the risk of interference in an unlicensed spectrum band will realistically limit the coverage from a base station to a few hundred meters. This is a comparable range to technologies such as DECT and CT2 in Europe and Japan's Personal Handyphone System (PHS). The CT2 system failed in the United Kingdom because the technology and service had already been superseded by cellular, but PHS did enjoy a brief period of success in the Japanese market. However, in order to provide coverage to greater Tokyo, 90,000 base stations were required. This historical perspective gives some indication of the sort of numbers needed to provide the citywide coverage needed to host useful services.

The question that needs to be addressed is whether this sort of coverage is practical? Keith Woolcock, in his report "Barbarians at the Gate-Wireless LAN Storms 3G Citadel," states that the "DM 97 billion raised in the German 3G auction could have bought 60 million wireless LAN base stations." (See the [Selected Bibliography](#).) This is more than enough to adorn every few square meters of the German countryside with an 802.11 AP. It would appear that the economics are not out of line with real life. But there is more to providing a service than simply installing enough connection points.

A full service requires both local and wide area connectivity, so the issue of backhaul provision must also be addressed. This means that each base station site must be provided with a link to a backbone network, which ultimately supports the connectivity to the Internet. In order to provide a reasonable grade of service (commensurate with that planned for 3G services), each connection must be of the order of 2 Mbps, bearing in mind that this capacity must be shared between all simultaneous users of that base site. This all adds to the cost of deployment but, of course, this may be defrayed if the EtherCLEC or other bandwidth providers take significant hold. There are a number of emerging suppliers (power and rail companies seeking to exploit their rights of way) who offer backhaul at ever decreasing cost (as explained in [Chapter 1](#)).

If the operator wishes to take full advantage of the full capability of the IEEE 802.11 standard, then a backhaul pipe of high capacity will be required—one or more 34-Mbps connections. These are currently very expensive items (the cost of renting such a circuit is several tens of thousands of dollars in a major city), but competition is easing this.

In addition to the creeping costs of deploying a solution, there are other issues that potential operators must consider. For instance, the frequency band in which IEEE 802.11 operates is unlicensed. This means that there are no licence fees to recoup, but it also implies that the air interface is open to multiple uses, including a range of domestic appliances such as microwave ovens and, critically, Bluetooth, to which the IEEE802.11 experts are currently trying to develop resilience from such sources of interference. The open question in this area is whether operators will risk offering (and indeed, whether people are prepared to pay for) a commercial service over which they ultimately have no control of the QoS?

Finally, whatever mix of technologies finally takes off, the difficulty of providing operational support should not be underestimated. This would be particularly true in a hybrid network, where customers expect services (and associated charges) to be end to end, not partitioned on the basis of underlying technology. A possible solution here would be for customer care and billing to be handled through the standard mechanisms that have been set up for cellular networks, with the wireless LAN coverage appearing as just another network cell. This has the advantage over the independent operator model, as that requires a third party (i.e., the local operator) to bill for local access on top of any service charge/subscription the user is already paying.

If there is one certainty in the future of mobile services it is that individuals and businesses will opt for the most economic way of delivering an adequate solution; pragmatism will drive progress, not ideology, and there will be a mix of solutions offered. If this means different technology for local and long-distance traffic or separate handheld devices and mobile phones, so be it. A marriage between IEEE 802.11 and 3G is by no means the only option. It is quite likely that some will opt for devices that support Etherphone for carrying voice over the IP (and Ethernet for data).

A large part of this section has discussed the practical considerations and issues that need to be addressed. Some key questions have been posed but only a few answers suggested. Real progress will be determined through compromise, barter, and good design.

As Ira Brodsky, president of Datacomm Research observes, "Public wireless LAN operators must join forces with 3G mobile phone carriers to achieve necessary coverage and service bundling. Likewise, third generation mobile phone operators need public wireless LANs to offload heavy indoor traffic from their lower speed, wide area networks."

7.3 Web Services

A recurring theme through this chapter has been that Ethernet technology promises to improve (i.e., lead to faster and cheaper) Internet access. We have already seen how a user might connect to a server over a wireless LAN to access their favorite Web site or the mail stored on their enterprise network. It seems that many technical journals indicate that this is exciting enough, or that there are sufficient business motivations to justify the installation of a new networking infrastructure. But, in and of itself, the thought of a different delivery mechanism to access what is already there does not really have much appeal when seen from end users' point of view. They are not concerned with network costs—simply with the service they get.

The real appeal of Ethernet technology in the wide area and for mobile access is that it extends the immediacy of the LAN into new areas. This does more than provide an alternative way of delivering what can be sent over dial-up and mobile networks. It means that broadband applications that involve the interactions between many parties can be supported.

The idea of cooperative communication over a network is something that has enjoyed a fair amount of attention in the applications area of late. Indeed, the whole idea of Web services is predicated on having the sort of infrastructure described earlier in support. Hence, to close this chapter, we take a brief look at Web services as the prime candidate for exploiting the twenty-first century Ethernet.

7.3.1 Some Basic Definitions

In essence, Web services provide a technical framework that defines how messages are passed between devices and information systems, along with a way to describe, publish, and locate the services that these systems offer to each other. It is assumed that there is an underlying mechanism for getting these messages to their destination but, on the strength of the last six chapters, we can take that as a given.

The prime reason for being interested in Web services is that they provide a widely accepted and effective mechanism for deploying software components, through what is emerging as an universally understood user interface. The reason that this is of interest is that this allows service solutions to be constructed by federation between suppliers for delivery to a wide range of customers.

At a technical level, these components encapsulate content, data, and business processes, and they support an interface that is accessed via HTTP (or its secure version HTTPS) using I/O data in XML format. By providing this common and nonproprietary user interface definition, Web services can be used to impose a standard set of rules that separate suppliers and consumers can use to set up direct peer-to-peer communications—in much the same way that a LAN supports workgroups. The interface that one user presents to another can be secured and restricted so that only mutually agreed information is transmitted across the boundary to authenticated recipients.

The interface that an organization supplies for use by another organization, or perhaps as a public interface for use by any user, can be considered a "[service](#)." There is an implied contract that if the requestor makes a valid call to that interface, there will be a predictable outcome supplied by the provider. Either some data will be returned, an item purchased, or an alarm sounded. Whatever the specified outcome, the interface is standard and the outcome is predictable. The key to this is that both parties understand in advance the nature of the contract, the mechanics of the interface, and the structure and meaning of the data passed across it.

Because the provider of the interface is in control of the service being provided through it, they can expose as much or as little information or functionality that they feel appropriate. Since the communication is peer-to-peer, neither player has to trust a neutral third party to mediate.

7.3.2 Federation

People today often find that they have to manage a great deal of personal information that is held in electronic format. Personal organizers, mobile phones, portable computers, network-based mail, calendars, and task and notification services all need access to common information that needs to be kept synchronized. Add on top of this information about preferred modes of communication, the priority

given to specific calls, service subscriptions, passwords, preferences about television and music tastes, Internet sites, and interests and hobbies, and the list becomes potentially very long. People tend to either manage down this information by not recording it, locating it on just one specialist device, or periodically synchronizing various subsets between devices or services.

What would be the impact if all this information could be managed as a single, logical information repository and accessed by devices, services, friends, employers, and so on under a set of simple rules defined by the user? All a person's gadgets and services would have the information required. A service provider or retailer would know the billing and delivery address without having to ask over and over again. The phone would know to ring if called by the owner's spouse but to route other calls to a voice-mail box. The list goes on.

While federating this information between devices may be possible, it may be more practical to store it with a trusted identity management service provider. This trust relationship has to be very strong, since the company is being given responsibility for valuable or at least personal information. It would not only have to be trusted to secure the information but also to ensure that shared access to the information is only given to entities that have been permitted access by the information owner.

Microsoft has provided a great example of this kind of profile management service in its Microsoft MyServices.Net service, which works in tandem with an identity and authentication service in the form of Microsoft Passport.

Microsoft Passport is an existing service used by a number of sites and which already manages 160 million user identities. It provides a single, centralized authentication mechanism. Once you have proven your identity to Passport, that identity can be supplied securely to other sites without you having to provide a username and password for each one. Passport acts as the "gatekeeper" who verifies your identity before allowing you access to the places beyond the gate. Passport is not cast as an open Web service-it is a Microsoft proprietary product. Although basically an authentication service, Passport can also hold "electronic wallet" data such as credit card payment details.

Microsoft's vision for managing user information is well advanced and includes proposals for a range of services, as summarized in [Table 7.2](#).

Table 7.2: The Range of Web Services Proposed by Microsoft

Service	Description
MyAddress	Electronic and geographic address for an identity
MyProfile	Name, nickname, special dates, picture
MyContacts	Electronic relationships/address book
MyLocation	Electronic and geographical location and rendezvous
MyNotifications	Notification subscription, management, and routing
MyInbox	Inbox items like e-mail and voice mail, including existing mail systems
MyCalendar	Time and task management
MyDocuments	Raw document storage
MyApplication Settings	Application settings
MyFavorite Web sites	Favorite URLs and other Web identifiers
MyWallet	Receipts payment instruments, coupons, and other transaction records
MyDevices	Device settings, capabilities
MyServices	Services provided for an identity
MyUsage	Usage report for above services

These services fall into the following categories:

- *Personalization information.* This is information about the users themselves, their physical location, and contact details. The services in this category are: myAddress, myProfile, myLocation.
- *Configuration information.* This is information about the hardware and software configuration of users' devices. This is akin to the current Windows "roaming profile" information. The services in this category are: myNotifications, myApplicationSettings, myFavoriteWebsites, myDevices, myServices. The myNotifications service is particularly interesting because Microsoft sees the use of "instant messaging" technology as a key part of its .NET strategy.
- *Application services.* These are "value-added" services that provide users with real Web-based applications such as document storage, calendar management, and inbox management.

The dependencies and mutual support that these services give to each other will, we believe, lead to an upward spiral of evolving personal services, creating a rich and highly active end-user experience. This will create and stimulate a market of service end-points in the form of new applications and devices, especially portable devices.

7.3.3 Software Components

Since a Web service is regarded as a remotely accessible interface onto the software component that supports it, it is worth taking a little time to walk through the idea of software components. With this background, it becomes easy to see that many of the advantages and constraints of using software components will also apply to Web services.

A well-designed software component is an independent piece of software that provides functionality that is generally useful across a range of applications. It implements one or more known functions and hides the implementation of those functions behind one or more interfaces. Ideally, it should be reusable.

Each interface on the component supports a set of semantically related operations that allow the user to access some aspect of the component's data or behavior. For example, one interface onto a component may contain management operations, another the operations to be employed by end users, while a third may be used to provide communication between peer components. The interface can be thought as defining a contract between the component and its users. Taking the gestalt view that a thing is what it does, the interface thus defines the role of the component in the environment or system that it operates.

Each operation on the interface has a name and a list of parameters of specific types, which collectively is known as the operation's signature. Interfaces supporting the same collection of operations with the same signature are, from the user's perspective, virtually identical (in truth, they are observationally equivalent). What may well differ is the quality of the service provided by the actual implementation in terms of quantitative attributes such as the speed of response, availability, and reliability.

Once we have components, we can begin to distribute them. As with other distributed software, the definition of interfaces is key. They are central in providing flexible and reusable Web service components. The interface defines the service contract between a user and the Web service, and all internal function is hidden. Thus, Web services offering the same interface may transparently be substituted for each other. This allows different vendors to provide standard functionality through a common interface but implemented by their own proprietary technology.

Web services use eXtensible Markup Language (XML) for the message, and they generally reuse the same HTTP transport mechanism. This allows Web services to be transported across the Internet or any other IP-based network over the same infrastructure as Web pages; only the language is different.

Having agreed on the language and the way this will be transmitted over the wires, the next challenge is to agree what the messages are for and what will happen when one is received. To do this, a service provider will design an application, define what it does, and then document how it is to be used in a language called Web Services Description Language (WSDL). This service definition is placed in a service registry, which can be private or public, and which can be interrogated by potential users using another standard called Universal Description, Discovery, and Integration (UDDI).

The service provider and service requestor, having respectively published and found the service, must create one application that will deliver the service and one that can call the service. However, neither party is concerned with how the other achieves this, provided that when it comes to actually calling the service, the service provider and service user abide by a common standard for using the service, which is called the Simple Object Access Protocol (SOAP). SOAP simply defines how XML messages are requested and should be processed when they arrive.

The various mechanisms described above work together to enable Web services. The basic mode of operation is that a service is first published (using WSDL) so that it can be found (with UDDI) and, assuming that it is what the recipient wanted, it can be bound (using SOAP). Hence, we have a full set of utilities that get a service from its creator to its consumer.

It would be wrong to suggest that this is the only way to provide Web services. There are other candidate utilities for each of the "publish, find, and bind" activities (e.g., ebXML could be used instead of UDDI). It is too early to say if a single, unified set of standards will be agreed for Web services, or even standards for interoperability between competing frameworks. Certainly, the pressure from buyers is for interoperability, but this may not align with vendors' desire for product differentiation and, more kindly, the genuine belief that one way is better than another for a specific application.

If the essence of Web services had to be summed up in one word, that word would be federation. This fits very well with the idea of a total area network based on Ethernet technology. If there is a basis for developing applications that take advantage of rapid peer-to-peer connectivity, then all the elements of a new age of communications are in place.

Team LiB

[◀ PREVIOUS](#) | [NEXT ▶](#)

Team LiB

◀ PREVIOUS

NEXT ▶

7.4 Summary

This chapter focused on the likely prospects for Ethernet in some of the markets in which it promises to shine. The first of these is enterprise networking, where there is already evidence that the technology described in earlier chapters can be deployed to provide an attractive and highly competitive service. In this instance, there seems little doubt that Ethernet-based solutions will, in time, overtake the less-flexible options that currently dominate the market.

A consideration of the Ethernet (in its IEEE 802.11b guise) alternative in the mobile network arena leads to a slightly different conclusion. In this instance, it is argued that Ethernet will probably sit alongside the much-hyped third generation mobile technology. Both have their strengths and limitations, so it seems logical that they will coexist, each capitalizing on its best features.

To close the chapter, we explained Web services. The term "Web service" refers to a way of building an application so that it can readily be delivered over a distributed network such as the Internet. This fits very well with the Ethernet, as it provides an underlying network well suited to such applications—one capable of delivering fast peer-to-peer connections. The basic principles and building blocks used to provide Web service are explained.

Team LiB

◀ PREVIOUS

NEXT ▶

Selected Bibliography

INT Media Research, "802.11 Wireless LAN Security: Usage, Expectations, and Strategies," June 2002, <http://www.intmediaresearch.internet.com/item/0,,2340629,00.html>.

McCue, J., "The Mobile Internet Manifesto: 4G Networks of the World, Unite!" *Economist Conference on the Wireless Internet*, Stockholm, Sweden, May 2001.

Norris, M., *Mobile IP Technology for M-Business*, Norwood, MA: Artech House, 2001.

Norris, M., and N. Winton, *Energize the Network-Distributed Computing Explained*, London: Addison-Wesley, 1997.

Norris, M., R. Davis, and A. Pengelly, *Component Based Network System Engineering*, Norwood, MA: Artech House, 2000.

Norris, M., and S. Pretty, *Designing the Total Area Network*, Chichester, England: John Wiley & Sons, 1999.

Sproull, L., and S. Kiesler, "Computers, Networks, and Work," *Scientific American*, Issue 3, 1991, p. 265.

Williamson, J., "The LAN/WAN Maze," *Global Telephony*, September 1994, pp. 20-31.

Woolcock, K., et al., "Barbarians at the Gate-Wireless LAN Storms 3G Citadel," *Nomura Equity Research*, 2001.

Chapter 8: Managing Total Area Ethernetworks

Overview

The whole problem with the world is that fools and fanatics are always so certain of themselves, with wiser people so full of doubts.

--Bertrand Russell

Effective management is vital for any network to work properly. However capable a technology, it benefits the user little if it cannot readily be tested, configured, maintained, and upgraded. To paraphrase a recent advertisement for car tires, power without control is nothing.

The traditional role of the LAN systems administrator is to look after the configuration of the network (including access rights and user profiles). If there appear to be problems with the network equipment, the administrator can often readily walk to and inspect the installation to find any piece of equipment that appears to be malfunctioning.

This is not the case when a local network grows to cover the campus, metropolitan, or wide area. A different approach must be taken to manage the network. Some degree of remote monitoring must be introduced when users and their equipment can be sited all over the place.

Managing geographically dispersed networks is something telephone operators have dealt with for many years. They established the principles and procedures-operational practices-needed to keep a large-scale, public network under control. As Ethernet installations grow to cover increasingly wide areas and support ever more users, these principles and procedures will become as relevant to total area networks as they have become to telephone networks.

For all its importance, it would be wrong to suggest that network management is anywhere near to being on a scientific footing. Like any other branch of management, it depends as much on judgment, expertise, and experience as it does on theory and logic. Somewhat regrettably, the management of large-scale networks has been a rather neglected area of study over the years. There are many more papers, books, and reports on new technologies than there are guides to their management.

Given the relative paucity of available information, this chapter presents a fairly detailed walkthrough of key management activities. It then goes on to introduce how these can be automated, as is essential for any real network. We preface this with an explanation of the various tasks that have to be carried out to keep a network operational.

The concept of OSS is also introduced in this chapter. These are software systems used to monitor network events, alter the network's configuration, and bill for its use. There are many OSS packages now available in the marketplace, and their assembly into a complete and coherent management solution is a specialist skill (or, some would say, a black art). Part of this chapter is set aside for an explanation of the basic steps in building a viable management solution.

To close, we take a brief look at the management solution that has been implemented by one of the wide area Gigabit Ethernet providers. This clearly shows how the principles introduced in the rest of the chapter can be realized in the real world.

8.1 Network and Service Management Models

The perceived effectiveness of a network depends on how available it is and how well it performs day after day, month after month, year after year. This means continuous monitoring of the equipment in the network and of the various features that it supports. The former is usually called network management, the latter is the essence of service management.

A key point that should be made is that the real problem addressed in service and network management is the control of complexity. Modern networks are simply too diverse and disparate to be controlled by anything other than well-thought-out principles and supported by powerful tools and automation. This is not something tackled as an afterthought. Management capability has to be built in as an integral part of the original network design.

The best way to deal with any complex problem is to break it down. The first step in separating key concerns in the management of a complex communications network is to look at the discrete links in the end-to-end chain that connects a user to the routers, circuits, servers, and switches that carry traffic and provide the service. We now introduce some of the models that have gained acceptance in the communications industry. These models are useful in that they provide a common point of reference between suppliers, operators, and customers. No common management strategy can be developed unless the participants can agree on the scope of management activity. Most users have an intuitive view of network management, but intuition cannot be implemented, far less shared.

8.1.1 ISO FCAPS

The first real blueprint for communications management did not appear until the early 1990s, when the ISO introduced FCAPS into its network model. FCAPS relates, quite simply, to the key management concerns of fault, configuration, accounting, performance, and security. Each of these management concerns was defined, as follows:

- Fault management-the task of monitoring network conditions to detect out-of-tolerance behavior and notifying control points when such conditions occur;
- Configuration management-the task of controlling the route that information takes through the network, the relationship between network elements, and how specific system parameters are chosen;
- Performance management-the task of monitoring the network's traffic, its remaining capacity, the rate of flow, the characteristics of network delay, and other factors relating to the network's connection, congestion control, and traffic flow capabilities. This also includes capacity management and planning functions.
- Accounting management-the task of collecting and distributing information needed to allocate network costs;
- Security management-the task of controlling access to the network and network-resident services, including the service of network management itself.

ISO stated that all networks needed to consider each of these management aspects, but need not necessarily implement all of them. There was no real prescription on how to deal with the various management concerns-just definitions. The acceptance of FCAPS as the key management aspects was a start but more detail was needed for practitioners, and this came from the Telemanagement Forum (TMF). The TMF [originally known as the Network Management Forum (NMF)] was body set up by network operators and equipment suppliers to implement practical management solutions across the communications industry.

8.1.2 The TMF

The TMF built on the ISO work by agreeing to an overall management hierarchy for communication networks, as shown in [Figure 8.1](#). This placed the user's needs firmly at the top and the management systems and network elements that fulfill those needs in a supporting role. Although simple, [Figure 8.1](#)

does carry an important message-that management design does not have to follow the specific layout of the network itself. It is often necessary to correlate inputs from several sources in order to relate what the user sees (or does not see) to the network that provides the service.




Figure 8.1: The TMF management hierarchy for communication networks.

In our management pyramid, a business is supported, successively, by a set of services, networks, subnetworks, and network elements. Having broken the management space up in this way, the TMF started to add information at each of the levels, as shown in [Figure 8.2](#).

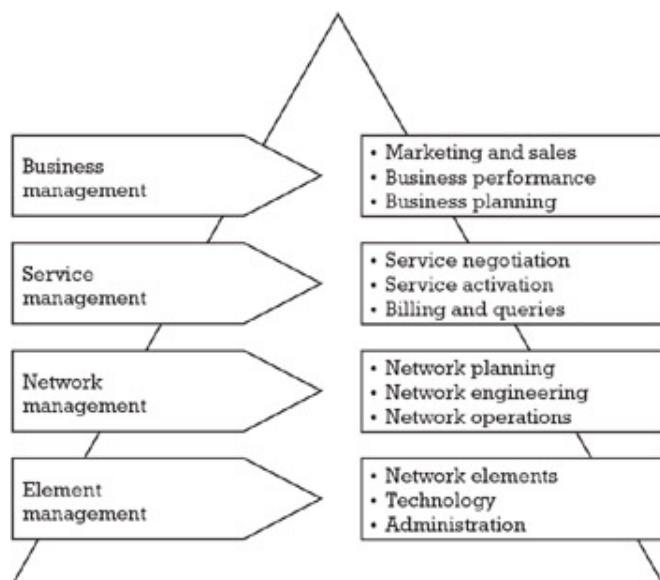


Figure 8.2: Key concerns at each level of the pyramid.

Of course, more detail is needed to make this divide-and-conquer approach practicable. This is where the definitions in the telecommunications operations map (TOM) come into play. The TOM defines the specific activities that lie within the service and network management layers. The exact definition of these activities and how they link with each other is the first step in the management design for any network.

8.1.2.1 The TOM

The structure of the TOM is shown in [Figure 8.3](#). Associated with each of the boxes in the diagram is a

wealth of supporting documentation that details the function of that box, defines its inputs and outputs, and shows where it fits in the end-to-end picture. We are now at a level of detail where we can begin to structure a practical management design, but there is a lot more that must be done before a viable solution is ready. Issues such as QoS and system performance must be addressed so that appropriate measures of availability, target response times, and so on can be established.

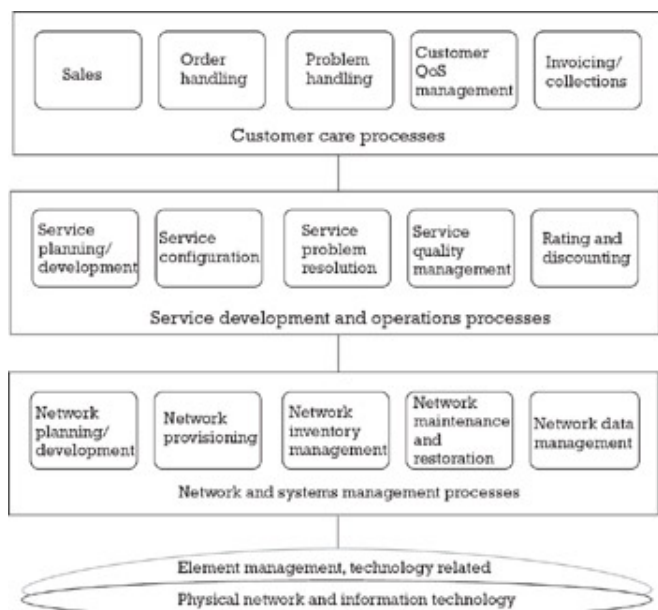


Figure 8.3: The TOM.

The TOM framework (along with some design and operational targets) provides something solid to design from and to check that design against. A specification for management tools can now be couched in terms of standard processes that must be supported. The trick now is to relate the standard processes to useful activities that must be carried out when running a network.

From an operational point of view, management is all about sets of processes that achieve a specific aim. This is the approach that the ISO took when it issued FCAPS and, for completeness, it is an aspect that has been added to the TOM with fulfillment, assurance, billing (FAB). Fulfillment (taking orders and providing service), assurance (locating, tracking, and fixing faults), and billing (allocating costs and issuing bills)- the paths through the activities defined in the TOM-are illustrated in [Figure 8.4](#).

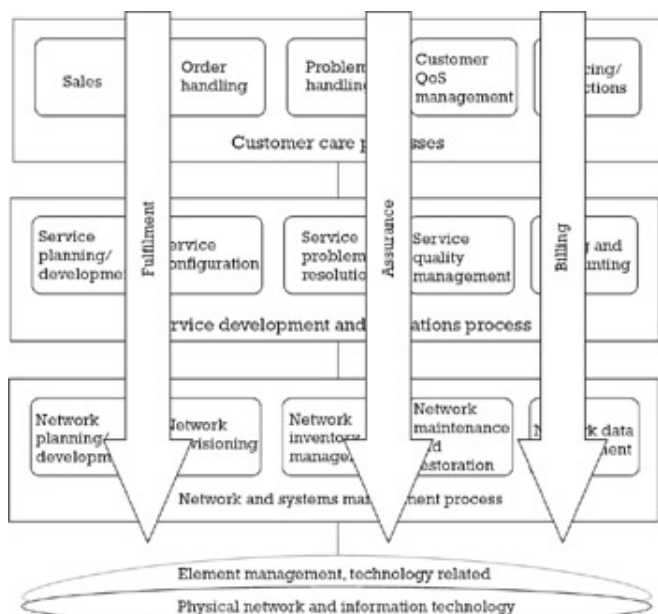


Figure 8.4: FAB in the TOM model.

Each of the arrows in the diagram denotes a set of customer, service, and network activities that support a discernible task, such as clearing a fault, provisioning a new service, or raising a consolidated bill.

If the FAB tasks are to be carried out effectively, a host of practical issues must be attended to. Perhaps the most important are to know exactly what is in the network, know when there is a problem with it (and fix it), be able to manage its configuration, and have some mechanism for correlating alarms. The following discusses each of these in more detail:

- **Inventory**-Invariably, several pieces of management software must be deployed to cover all of the required activities. Each will hold some information about the network or its users-for instance, a billing system needs to know who has used the network, for how long, and for what purpose. Together, the data used by a complete suite of management software covers the network users, services, products, configuration, and physical components so that all of the information needed for an inventory is available. One of the most commonly encountered pitfalls is that the required information is not accessible or, because the same data is usually used in more than one place, it is not clear which is the master source.
- **Isolating network faults**-All but the smallest networks require automated procedures to report problems, raise trouble tickets, and track progress against them. As well as creating a history file for each separate problem report, the systems should record the nature of the problem, its current status, how it was resolved, vendor contacts, and any other information that might be helpful in the future or required for audit purposes.
- **Managing the network configuration**-A basic tenet of management is that you have to know what it is that you are managing! The dilemma that this introduces is that the interrogation of the network to assess what is there can be both time consuming and resource hungry. Hence, a balance has to be struck between gathering information and generating management traffic.
- **Correlation**-One problem deep within a network can manifest itself as a host of broken links and lost services. An effective management solution should be capable of identifying the root cause of a problem from observed or measured effects.

In addition to these capability requirements, several temporal issues should be considered when developing a management solution.

The first is legacy systems-the already installed elements that have to be accommodated. These represent not only a capital investment, but also a solution that is often not totally unsatisfactory, and keeping these systems or modifying them may reduce both costs and risk. Experience shows that just because a system offers color graphics, object-oriented design, and artificial intelligence does not necessarily mean it will do what you want it to. Slightly boring old systems often turn out to be more effective and reliable.

Second is that the cost of a management solution must be in proportion to the benefits it offers, which should normally be a fraction of the network cost. In practice, this fraction tends to average around 5%, a figure that is probably too low to adequately care for today's complex and feature-rich networks.

In [Section 8.2](#), we will see how tools map onto the TOM. This is more than a pure academic exercise, as a growing number of customers ask for management solutions to be presented in terms of how they support the TOM processes. In other words, this is becoming a recognized part of the designer's job.

8.2 Network Management Tools

There are a great number of management systems on the market, and new offerings arrive on a regular basis (at the same time that some disappear—there is considerable churn in this market). A few of the available systems tackle the whole management space (i.e., cover all of the TOM processes), while others cover just a part of it. Some are optimized toward a particular vendor's product and come with proprietary interfaces, while others are built for the open market and use standard interfaces and protocols. The capabilities and quality of management tools varies enormously and there are always bright new stars, as well as established favorites.

Given the size and volatility of the management tools market, it would be impractical and probably quite unfair to review all of them (or even a representative sample). Given this, we now present the main categories of tools and refer to a few of the more widely used examples.

Total management platforms. Over the years, there have been a number of attempts to produce an all-embracing management system. Several major telecommunications operators and equipment suppliers, such as IBM, AT&T, BT, and Nortel, have tried to build the ultimate, all-singing, all-dancing management system. All have tried to construct a crosstechnology platform that would give a concerted, end-to-end view of all technologies used in the network and all the services carried over it. All have been overwhelmed by a combination of complexity and rapid change and have ultimately failed.

IBM's NetView is a good example of the perils of centralizing all management functions. The early NetView product had a strong orientation toward managing transaction monitor programs (i.e., CICS, IMS) and could be used on host systems, as well as on IBM or IBM-compatible network devices. It could be equipped with management links to virtually any type of product or service and was the most pervasive network management system on the market, installed in more than half of all IBM mainframe sites.

NetView was entirely based on IBM's own SNA protocols and relationships and is thus proprietary. As interfaces were added to allow NetView to accommodate other technologies and provide a single view of a communications network, the complexity of the product escalated and, in time, became untenable.

This is probably a somewhat biased view. Despite not fulfilling its original aims, NetView persists as a useful product. The same cannot be said for other attempts. The ServiceView product from BT/Concert and Nortel Network's Preside integration platform suffered in much the same way as NetView before disappearing altogether.

Process-specific tools. Given the difficulties in producing a single tool that could support all of the processes defined in the TOM, most suppliers have tended to focus on one specific area. Combining a number of these process-specific tools, it is possible to create a management solution that suits a specific situation. Using this approach, the quality of this solution depends not only on the quality of the individual tools but also on the way in which they are integrated. In view of this critical dependency on overall design and integration, we look in some detail at how to achieve a viable set of management tools. To start with, we can place the available tools into categories that match the TOM, as follows:

Customer management tools. This category of tool should enable a network operator to deal effectively with customer queries, orders, and other interactions and thereby provide good service (and, hopefully, keep a customer happy). Some of the leading suppliers in this area are Siebel, Vantive, Clarify, and Remedy. The tools from these suppliers are generally referred to as customer relationship managers (CRM). All offer facilities for tracking interactions with a customer and progressing customer requests (e.g., for service activation, orders, problems awaiting resolution) through to completion. Also in this category are billing system suppliers such as Kenan, Portal, and Amdocs.

Service management tools. These are tools that manage at the service level and are probably the least mature of the process-specific set. There are a number of reasonably well-established offerings in this category that allow an operator to turn a customer order into the set of appropriately configured network components. This capability is usually referred to as service design and is supported by tools from suppliers such as Architel and Orchestream (which focus on physical connectivity) and Netscape

iPlanet (which looks after logical configuration by setting router policies). Workflow tools that manage a set of service-related tasks, such as BEA Weblogic, would also fit in this category.

Network management tools. The longest established tools in this set are those that interact directly with network elements via element managers that are associated with each distinct type of equipment. The key attributes of tools in this category are the ability to trap alarms and the retrieval of status and performance information from network devices. The more sophisticated network managers can determine the root cause of faults and tell the operator what the service impact of the fault is likely to be. Leading suppliers at this level are Micromuse and SMARTS, which have developed adapters so that their software works with just about any network technology, and Cisco, which supplies CiscoWorks, which is optimized to manage Cisco router network.

Support tools. In addition to those already mentioned, there are a number of tools that do not have a specific network, service, or customer role, but provide essential support to the tools that do fit into these categories. Inventory, for instance, is not uniquely related to any one category but is required to support all of them (as indicated on the TOM). Ideally, an inventory tool should allow not just network configuration to be recorded but also links between network, service, and customer data to be stored. Leading inventory suppliers include SmallWorld, Cramer, and Metasolv. Other support tools include gateways to other suppliers (where Intergate and Crosskeys are leading suppliers), network planning (where Netplan is a leader), and reporting (where Concord is preeminent).

In practice, many of the leading tool suppliers are building on their core competencies to provide extra functionality from their tools. Often, this means that they will add modules to their base offering that allows more coverage against the TOM. In addition, many suppliers have formed alliances so that their tools are preintegrated with complementary tools from partner suppliers.

Figure 8.5 places some of the better-known suppliers and/or their products on to the TOM framework. This illustrates just how many options there are in satisfying the management requirements for a communications network.

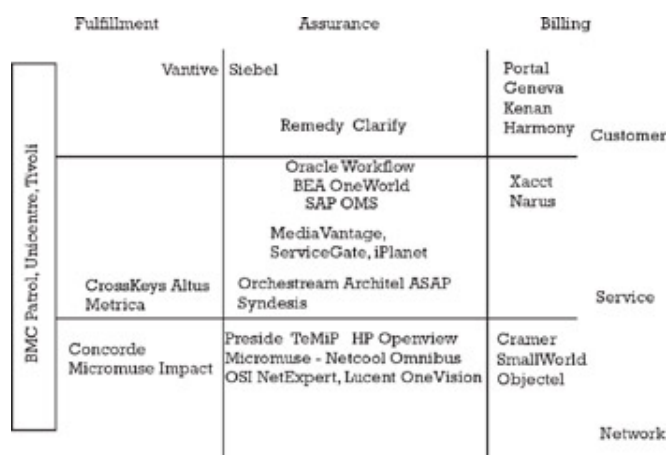


Figure 8.5: Location of major products on the TOM.

The diagram is subjective and intended to illustrate the level of overloading and overlap on the supply side, rather than provide any sort of definitive guide. There are many more systems that could be added to the diagram but the choice is not as wide (or as confusing) as it may seem. In practice, the options are reduced on the basis of preferred suppliers, acceptable price, product quality, level of integration, and so on. It is only when these factors are applied that a reasonable option set comes into view.

Integration frameworks. There are a wide variety of management frameworks on the market that serve to host process-specific tools. HP's Open-View Network Node Manager is probably the most widely used. OpenView is generally deployed by operators who must oversee a number of diverse technologies and supplier-specific systems. To get an overall view of their network, some form of integration across these low-level systems is needed. HP OpenView provides this by collecting information from various sources and presenting the collected information in a single view. Hence, an operator using Cisco routers with Newbridge frame relays can "see" a unified picture of their network,

rather than relying on separate management systems and a swivel chair.

NCR's StarSentry Manager, Compaq's TeMIP, SunConnect's SunNet Manager, and Novell's Network Management System (NMS) are all widely deployed examples of integration framework. These all provide the same sort of framework as OpenView and boast a similar set of additional tools that can be added.

System management tools. Before moving on, it should be noted that the systems that are used to look after communication networks (usually known collectively as OSS) are not perfect. They must be managed, just like any other software. There are a number of well-established system management tools that have been developed to monitor operational software. BMC Patrol, Tivoli, and CA Unicentre are three of the leading market offerings.

8.2.1 Integrating Tools to Create Management Solutions

Most of the commercially available tools support just one of the management processes defined in the TOM. Even the most sophisticated tackle only two or three—no one tool could be considered a complete solution. A comprehensive management solution calls for a significant number of the TOM processes to be supported (rarely fewer than seven or eight) and that these processes need to work with each other. So there is invariably some amount of integration to be done.

In practice, a large part of the budget allocated to management systems installation is invariably allocated to integration. The question that should be asked when allocating this budget is what are the technical considerations in selecting which systems should be adopted in the overall OSS solution? [Table 8.1](#) gives some indication of the pros and cons of each approach.

Table 8.1: Pros and Cons of Different Management System Integration Strategies

Approach	Advantages	Disadvantages
Single source	No integration is required, so faster deployment is likely.	Tied to one supplier. Functionality in some areas is likely to be poor.
Minimize number of systems	Amount of integration is minimized.	Some areas will be less rich in functionality.
Best of breed	Richest functionality throughout.	Considerable integration required, therefore expensive. Some functionality within systems may be paid for but not used.

It should be noted that the amount of functionality available in any system is not just dependent on what that system supports but also on how that system is integrated. For example, a billing system may support customer account hierarchies, but if this data is being managed via a CRM system and the CRM to billing interface does not support account hierarchy information being passed, then the "out of the box" billing system functionality in this area is unusable.

So, however good one management tool is, it is only effective if it cooperates with others. The creation of a management solution calls for intellectual effort not only to ensure that the correct tools are used but also that they work together, communicate in an agreed way, and work on common data. [Section 8.3](#) details the challenge management designers face in creating a viable solution and the way in which that challenge is best met.

8.3 Management Design

In an ideal world, a well-designed network, once installed, would just keep going. In practice, a trouble-free and future-proof network just does not exist—which should be no surprise when you consider that it is comprised of many components from many suppliers that all need to cooperate to meet stringent availability, reliability, and performance targets. On top of this, the uses to which the network is put inevitably change.

Even in stable operations, there are plenty of situations that a management system must be capable of meeting. The following are the main ones that the designer/integrator needs to take into consideration:

- **Chain-reaction failures.** You must know how a failure in one part of the network might affect the total operation. Suppose there is a bug in the database software that keeps track of network addresses. It could block access to critical services and, at the same time, hide this problem from the monitoring system. Correlation and consistency checks must be built in so that you can detect this kind of problem.
- **Traffic congestion.** Any network can suffer from traffic jams. If several network elements fail simultaneously, the combined load of blocked and diverted traffic can overload queues or block switches and bring the entire system to a halt. Often overlooked (and adding to the congestion) are the messages the network generates to report the problems.
- **The unexpected.** A network hit by unexpected events must be able to help itself. It should manage and reroute traffic to avoid trouble spots. The low-level design must ensure that the network reacts properly to duplicate messages or verifies messages from questionable sources. Most systems use time-outs and retransmissions to deal with these problems. Another approach is to display a status flag that warns of impending problems. Good management design has a lot in common with "defensive programming."
- **Packages and bundles.** Management is not just about dealing with problems—it also covers provision of service and the composition and generation of bills. If a product or service being offered to an end user relies on separate network elements, there should be some way of relating reports from those separate elements to the product or service they support.

In dealing with these operational issues, a number of design choices must be made, such as:

- **Centralized or decentralized management.** Central management can ease correlation of information from network elements but it also creates a central point of failure. Decentralized management can be a source of inconsistency but can be a more resilient option. In a large network, it is usual to have elements of both, with their combined drawbacks! This means that issues such as who should be responsible for managing database consistency, standby systems, and database updates must be clearly resolved. Another decision to be made is who should receive status information and error messages? Often, a local group can take care of problems in its own system. There are other times when central management is more effective.
- **Protocol standards.** The choice of network management standards [e.g., SNMP, Common Management Information Protocol (CMIP)] can either improve management or make it harder. If the design is based on recognized standards, it is important to be sure that the entire system follows the same version. Otherwise, it might interpret some messages in strange ways. Standards often make it easier to integrate network management with the network as a whole but, where standards are immature or not widely implemented, proprietary solutions can be a better option. On some occasions, the choice of no standard at all is a valid option (but only if it is indeed a choice, and not a default).
- **Testing.** Testability is one of the fundamental requirements in the development of any technical artifact, yet is often overlooked in network design. A good network management system should include test points and have built-in trace and audit capability.

Whatever approach is taken in the management design, some allowance must be made for change. A reasonable test of this is to check how the following will be accommodated:

- **Growth.** The management system should be able to cope with traffic growth and the addition of

new nodes and networks. It should also be structured so that new technology can be incorporated—an attribute that the layering approach previously outlined helps with.

- Adaptability. The network management system should adapt to system changes and allow the relationships between network elements to be redefined.

This list could be extended, but these are the essential (and inevitable) issues that have to be addressed. Knowing what the likely problems are and having the TOM as a framework for designing a solution, we can now start to look at the options that can be deployed in a management solution.

8.3.1 Systematic Approach to Management Design

The role of the management solution designer is very much that of a consultant who has to work with customers and advise them what is possible from prospective suppliers. Although delivery oriented, it is not a job with clear boundaries, and much of its value derives from this breadth. Done well, it requires a mix of analytical and consultancy skills, knowledge, and persistence from individuals.

There is no formally established basis for management design, as there is for network dimensioning or software creation. Despite this lack of formality, experience indicates that there *are* some reproducible aspects of solution design that can be assembled into a systematic approach. The following are the essential elements of such an approach:

- Some necessary mechanics, primarily requirements management, change and configuration management, design review, and testing;
- Any preferences dictated by supplier agreements, purchasing policy, and so on that covers the use of preferred systems, conformance to (the few) management standards that are available, and reference to previous (preferably similar) designs;
- The definition of interfaces and master data, using process walk-throughs and data consistency checks.

How these elements come together is best illustrated by running through what actually happens to a set of requirements as they are developed into a proposal for a management solution. The sequence of events is typically as follows.

First, the scope of operational requirements has to be clearly defined. Some aspects of management will be more important than others, and some may not be relevant at all (e.g., discounting and rating is not needed if fixed charges are to be levied).

Then, similar design from the archives should be checked. Special attention should be paid to previous requests from the same customer (whose expectations have been set and who will have some basis for comparison) and experience with likely suppliers (whose previous performance gives some input to planning).

Once the requirements and rules of engagement are clear, prospective component systems from the preferred list can then be placed on the TOM. The populated map can then be audited by stepping through each of the key business processes and making sure that the ones that the customer needs are actually supported.

During this step, notes are made of any issues (manual intervention required, unknown interface, missing component, and so on) that are not critical but worth recording. Also, for each of the processes, a note about where data is created, read, updated, or deleted should be made—it is vital that the designer ensures that there is a master source for each of the key data items and that the required records could be retrieved by all systems.

Once a viable management solution design has been documented and reviewed, it can be costed and the delivery can be planned. The delivery plan should show which components, modules, and systems will be delivered and when. Within the design, particular attention should be paid to connectors between components (Do they exist? Are they okay? Do they support the type and volume of interchange you would want?).

Credible management systems designers/integrators should have a process in place that embodies all

of these issues. They should also have reference designs and installations to help customers visualize their specific solution. [Section 8.3.2](#) gives some examples in this area.

8.3.2 Typical Design and Integration Steps

Each network and design has its own characteristics, so any examples given here have to be very much in outline—there are many specifics that influence which systems are chosen and how they are deployed and integrated. Nonetheless, some useful general rules should be followed.

Component systems. This step illustrates how the three main paths through the TOM—assurance, activation, and billing—are catered for and how each of the functions/tasks along that path is carried out. The arrangement of component systems in a typical solution is shown in [Figure 8.6](#).

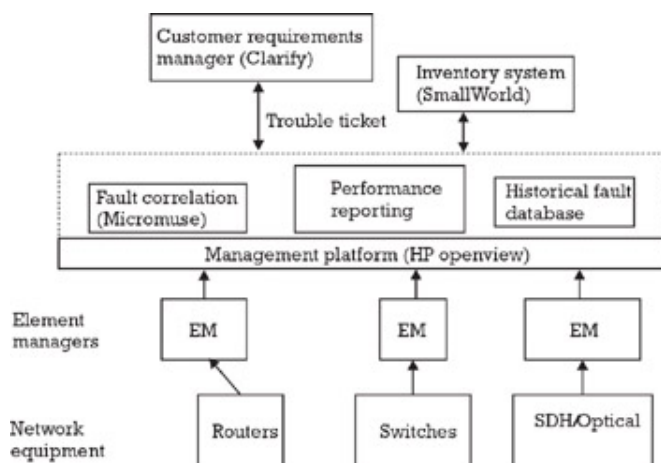


Figure 8.6: Typical systems architecture.

In the diagram, the appropriate network element managers connect the various network elements into a common platform, which supports cross-technology functions such as end-to-end path supervision and root-cause fault detection. Up-to-date information on the network is collected by polling elements managers for their status and receiving unsolicited events from managed devices. This information is used within the integration platform (HP OpenView) as a current view of the network. Hence, the network itself provides basic inventory and configuration information. This information is then exported and overlaid with other data (such as location, and spares) to provide a comprehensive inventory.

Root-cause faulting is provided from the correlation tool (Micromuse), and the alarm map can be viewed through a graphical network browser. The integration platform connects into the customer management system (Clarify) with its fault information and generates a trouble ticket. All of the trouble tickets raised are visible to the network operation center (NOC) via the Clarify system. Hence, all network elements are under control. Following field operations or reconfiguration, updates of network inventory can be initiated from the NOC into the Objectel inventory management system.

Process walkthrough. In this example, we want to show how the order handling function is catered for. The customer order would be recorded by opening a new case in the Clarify system. The complete design would probably also have to indicate exactly how phone, fax, and e-mail orders are captured.

The status of the order is monitored using the Clarify case as the master record. The Architel OMS system is used to turn the order into a network activation schedule. The availability of the necessary network equipment can be checked using the information in the Objectel network inventory. If the order can be fulfilled, it is directly activated through the ASAP system.

Once activated, the inventory is updated. The details of how this is done have to be carefully worked out. Following activation, customer performance reports are derived from the various elements that constitute the circuit. These reports are available in both pdf and XML formats and can be accessed via a browser.

The basic records needed to bill for the service ordered are generated by the service accounting part

of the integration platform and passed to a billing system such as Geneva or Portal for processing.

More could be added (e.g., how end-to-end network circuits are provided, how the call records can be used for marketing) but the level of explanation should be tailored for the specific customer.

Operational issues. By this stage, it should be possible to show (with some conviction) that various components in the management solution will work together to enable the key operating processes. That is not to say that the solution will be all-singing, all-dancing—a good design should be pragmatic and match the available budget and need for automation. Manual intervention and tactical tools are acceptable if they do the job.

The primary system supporting customer care in this example is Clarify e-front office. This is a powerful suite of applications that deals with the management of customer orders, tracking of problems, and maintenance of customer information.

When a new order is received, either over the phone or via the Web, the customer center would enter the order details into the customer database held on Clarify via the input screen. This information is used to monitor the progress of order provision and would be checked against existing customer information as part of the order taking process.

As soon as the order is taken (or received, if input over the Web), the availability of the necessary network equipment and/or resources must be checked. This means that the order has to be broken down in any constituent parts and subservices/network resources identified. Customer center staff does this using established templates.

The information held in the network inventory and service control systems (i.e., those that hold directory and policy management details) must then be checked. Availability of stock and field engineers must be verified.

If all the necessary constituent resources and network components are available, the order can be fulfilled. At this point, an order confirmation would be sent to the customer. If not, the alternatives are to suggest an alternative, put the customer on a waiting list, and so on—the details here would be subject to detailed process design.

Once an order is confirmed, it is scheduled for activation. This can be carried out automatically using a system such as Architel ASAP, which will configure the various network elements required to fulfill the order.

Once activation is complete, a number of master data sources must be modified, including the following:

- The network inventory must be updated with details of this solution, so that any subsequent network changes can be carried out safely.
- Customer records must be updated to show the order and ensure that any modifications or problems are managed.
- If the order is for a billable service, the billing records must be updated.

Once these updates have been made, the customer can be informed that the order has been fulfilled and that service is available. The updated master data sources (network, customer, billing) are made accessible for other systems in the management solution to read. This allows, for instance, the correct association of network equipment and the customer service it supports.

A more detailed design would show how manual and automated processes work together, how data integrity is maintained, and how other processes (assurance, billing) are supported.

Data integrity. It is important to show that the information used in the management system is current, available, accurate, and secure. This means that the following exists:

- A master source of data (i.e., customer, network information);
- A master source that is readily accessed [through an API, a standard like ODBC or Common Object Request Broker Architecture (CORBA) or file exchange];

- A locking mechanism (so that you do not get two people trying to update the same record);
- Adequate measures to assure that the data are not altered without authorization.

By this stage, there should be a reasonable level of confidence that the management system will work as it was intended. In practice, as much time and effort should go into the design of the management solution as goes into the design of the network itself. What has been presented in this section is only part of what must be done. We have covered the logical aspects of managing our network but have said little (apart from which software components will be used) about the physical realization—network operation centers, management overlay networks, and remote monitors are all part of the job. We have, however, strayed from the familiar ground of Ethernet for long enough, so it is time to get back.

Team LiB

[◀ PREVIOUS](#) [NEXT ▶](#)

8.4 Implementing a Management Platform

Perhaps the simplest example of an extended LAN is one that is connected to a remote enterprise network. [Figure 8.7](#) shows this setup equipped with a minimal network management system. In [Figure 8.7](#), we have a workstation connected via a router and core network connection to the target enterprise network sites that are to be managed.

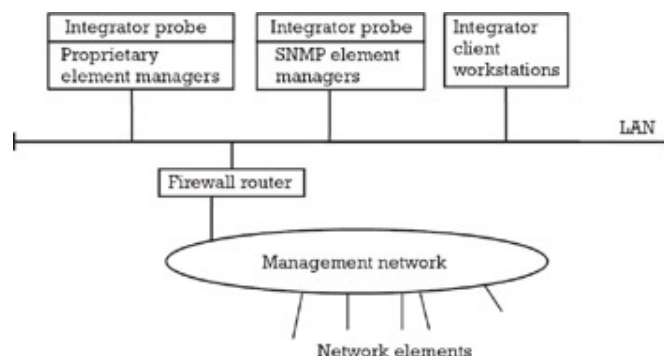


Figure 8.7: Basic extended network.

The element manager will typically be a Unix platform running SNMP management software (e.g., Sun Net Manager or HP Openview). SNMP itself is a straightforward request/response protocol that allows retrieval of information from various agents (a type of monitoring device) on the network through its "get" and "get-next" commands. It also permits network devices to deliver traps (also known as alarms) that can be delivered into the management system, and it has a write capability used in many proprietary network management products for configuration management.

The task of the SNMP management software is to provide a network manager with a graphical view of the network, highlighting nodes with problems. The user will be able to view the graphical user interface to hone in on such nodes to perform further diagnosis (often through using TCP/IP Telnet, an HTTP-enabled browser, or some proprietary graphical user interface (GUI) to log on to the router). The management system is also able to gather statistics and prepare reports to assist in long-term health checking and capacity planning for the network. For instance, the spare memory or processing in a router might be routinely collected and monitored in a tool such as Trend.

The element manager is connected into the enterprise network via the LAN and a router, which also performs packet filtering and perhaps more advanced firewall services (security functions required to protect the management system from misuse). Extra management agents can be accommodated by adding extra workstations, which can be set up to have sessions with the network management software.

How does the management system know the state of various nodes in the network? There are two key methods:

- Alarm or trap processing. Network components tend to produce a rich set of event notification messages (examples would be failure of a token ring attached to a router, or if a router is running short on memory). These messages are typically issued by managed elements using SNMP and logged by the management system. The trick in dealing with such messages can be trying to tell the wood from the trees! Many commercial management software tools have a correlation capability that enables secondary alarms (e.g., a router reporting a lost link when it is the broken link that is the real problem) to be suppressed.

Management systems will require alarm filter capability so that only the most serious events are brought to the attention of the operator. Where possible, information-type alarms should be suppressed at their point of origin to avoid unnecessary traffic—most network components will have the ability to configure just how talkative they are about their status.

- Proactive information gathering through polling. Alarms are not in themselves sufficient. Clearly, if a network component or its communications path to the management system has failed, no

alarms will be seen. This problem is overcome by the management system periodically polling the network component. This can be done using a "ping" (a process whereby a test packet is reflected by the node under test).

Alternatively, we can use SNMP to get one or more variables from the nodes' MIB. The MIB is the repository of information on network devices that contains a description of SNMP-compliant objects on the network and the kind of management information they provide. These objects can be hardware, software, or logical associations, such as a connection or virtual circuit. An object's attributes might include such things as the number of packets sent, routing table entries, and protocol-specific variables for IP routing.

Clearly there is a design trade-off to be made here—if you poll too often, the excessive network bandwidth is consumed. Polling too infrequently means that it takes too long to react to faults. Polling network components every 5 minutes is often a sensible compromise.

Network statistics are also gathered by polling. In this case, we need to make a whole series of SNMP "get" commands to fetch the relevant data. A typical MIB will contain many kilobytes of "useful" data. It is all too easy to consume excessive bandwidth and network management system disk space by collecting too much too often. Collecting statistics every 15 minutes usually provides adequate resolution. Care should be taken to request only the minimum necessary information. This might include the following:

- Router peak processor utilization—to help predict when and where more powerful routers might be needed.
- Router-free memory—to identify where more memory might be needed.
- Network access link utilization—to identify sites needing more capacity, or where we need to investigate how the load can be reduced through tuning.
- Packet drop rate—routers are generally dealing with connectionless IP-type protocols, and they cope with port congestion by simply discarding or dropping any overload. The drop rate is a good warning that undesirable levels of congestion are being reached.
- Errored packet rate—network access links will generally be running protocols based on HDLC framing. These frames include a frame check sequence, which routers use to detect and discard faulty packets. High rates of errored frames will indicate links that need maintenance attention. Note that high rates of line errors may also raise alarms.

8.5 Commercial Management Platforms

While the platform discussed in [Section 8.4](#) may prove adequate for managing a small network, it is unlikely to satisfy the larger network provider user. Some issues that must be catered to when the network grows include:

- Scalability-one workstation, running element manager software, can typically only poll a few hundred network elements. A big network can consist of thousands of elements. An outsourcing supplier with many customers could be faced with handling tens of thousands of elements. And as the number of elements managed grows, so does the bandwidth required to reach them.
- Diversity-not all network components use the same standard management protocols. Some will have proprietary management platforms, providing equivalent information (e.g., a total area network may rely on managed LAN hubs that are managed using Novell protocols, and a long-distance capacity supplier may provide a proprietary information feed to notify the user of the current state of the network). What the network manager needs is an integrated view of the whole network-not a console with a dozen different screens, each showing a different component.
- State of the art-when developing new networking technologies, it is common for the development of the management facilities to lag behind the development of the user data path. This occurs for easily understood economic reasons-the company developing a new data communications product will want to see an early return on its investment and will thus want to focus development efforts toward early product launch. It is possible to sell a new technology that can transport user data in some novel and cost-effective manner, even if the product has rudimentary management ability. You could not sell a product that has excellent management abilities but does not provide effective user data transport! With a collection of state-of-the-art networking components, the management capabilities may well fall short of expectation.

The result of this is that the network designer is unlikely to find a single product that will meet all management needs. An effective platform has to be assembled from (the best features of) a number of products. This is where the design guidelines explained in previous sections come into their own.

A typical large-scale (i.e., national scale) management system is shown in [Figure 8.8](#).

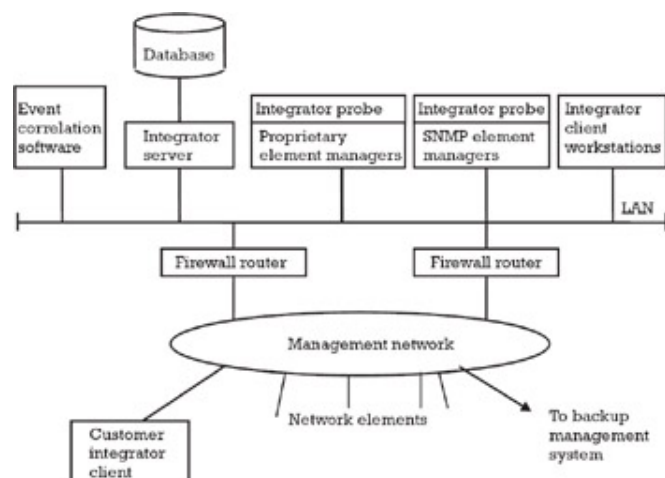


Figure 8.8: A large-scale management solution.

This consists of a number of element managers, each of which will typically support several hundred network elements. The exact number of element managers will depend on a number of factors, including the following:

- The capability of element manager software;
- The power of the hardware platform deployed;
- The frequency of element polling and the volume of data to be collected.

The heart of the platform is the event correlation software, which can be considered to be a "manager of managers." The example of this type of software shown in [Figure 8.9](#) is the Micromuse Netcool/Omnibus product, which verifies and filters network events to give the operator a clear view of what is really going on in the network.

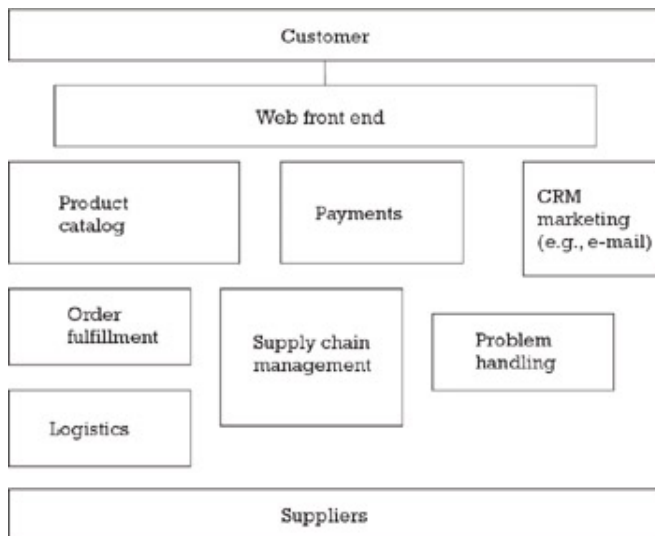


Figure 8.9: The management support needed for effective e-business.

8.6 Distributed Management

In the previous sections, we described highly centralized management systems that rely on collection and analysis of large volumes of data for correct operation. The levels of traffic generated by this process can easily impact the performance of a network, if too much data is collected too often. Some of the recently developed techniques in distributed management systems can help considerably with this problem.

This technology works by placing management probes at the remote sites (generally those with problems or considered likely to be the busiest). These probes can either be software running on routers or LAN hubs or dedicated network analysis hardware. The idea of these units is that they capture large quantities of data about network conditions locally and only send summary reports to the central site on request. The dominant standard in this field is known as remote monitoring (RMON).

RMON probes can also be used to generate smart alarms, through a process known as benchmarking. The reports produced by the probes are analyzed by the network operator to establish what normal peak load conditions are like at that site—in terms of total area network utilization, packet drop rate, router processor utilization, and so on to establish operational norms. The probes are then configured with these norms as thresholds and only raise network alarms when sites exceed thresholds. The alarms can also condition the probe to start much more intensive information gathering, so the network manager has a good chance of discovering what is actually happening at these times.

As network management technology evolves, we can expect to see increasing levels of sophistication in what can be done to monitor the health of a network at the local level, meaning that the center is only contacted when things start to go wrong.

8.7 Managing e-business

In [Section 8.1.2.1](#), we introduced the TOM as a blueprint for developing management solutions. In order to align itself with the most recent network-based applications, the TOM has recently been extended to cover e-business. The eTOM adds the activities needed to operate effective on-line services to the basic set needed in any network.

The area of e-business is epitomized in services such as <http://Amazon.com>. It involves dealing with the customer (business to customer or B2C) and with other suppliers (business to business or B2B) via the Web by supporting on-line transactions for products. It comprises the following features:

- Content management-automation of the publishing of information on the Web;
- Personalization-allowing customers to customize how content is delivered to them (e.g., by putting specific items of interest to them on the first Web page);
- Account management-providing a secure customer account on the system, optionally with validated payment details and so on;
- Ordering-allowing the customer to browse the product catalog (which ideally has information regarding how many of each product are in stock) and manage a "shopping basket" of products, which can subsequently be checked out;
- Order management-processing the order from the customer, including order fulfilment;
- Supply chain management-managing the items in stock and dealing with suppliers when stock levels cross a minimum threshold.
- Invoice production-including support of customer payments;
- Marketing support-through analysis of the customer's preference, market specific products to the customer via advertisements (banners/e-mail);
- Support for customer sales representatives-for helpdesk and so on.

Of course, quite a few of these have already been introduced as they pertain to almost any sort of network. Here we consider the support needed to support the unique operational aspects of e-business. For instance, where an on-line catalog is available from which people can buy products, pay for them on-line, and have them delivered-the aspects of supporting this service are shown in [Figure 8.9](#).

The eTOM provides a comprehensive account of where these newly introduced management requirements lie. For example, items in the product catalog are displayed via the Web to the customer. Those selected by the customer must then be dispatched, and payment collected (usually by credit card). Problems (e.g., returns) and support must be handled, and supply chain management is needed to maintain stock levels.

Although the aspects of providing this service include many of those required to support standard telephony services-ordering, fulfillment of orders, customer invoicing, supply chain management, and so on-it is often better to implement the e-business service as a separate system, for the following reasons:

1. To deploy some of the e-business functions on standard telephony systems would be like trying to use a hammer to crack a nut. For example, using a service provisioning system to support order fulfillment or a billing system to collect payment from a credit card agency is way over the top.
2. The e-business service requires close coupling between the various aspects, which is easier to achieve through using one system. For example, customers will want to know how many of a certain item are in stock and what the delivery time will be.

A host of suppliers have appeared in the marketplace with management solutions for e-business, including Netscape, IBM, Commerce-One, and A Recipe for Success, among many others.

Team LiB

◀ PREVIOUS

NEXT ▶

8.8 Summary

In previous chapters, we have seen how Ethernet technology can be used to build bigger and more capable networks. The more a network offers, the more reliance is placed on it and, in turn, the more important it is to keep it working as it should. In this chapter, we have emphasized the importance of network management systems to the success of a total area network. Although not always the case in reality, the management system merits as much care and attention in its design as the network data path.

Given the physical reach and service diversity that can be achieved with Ethernet, the following issues must be considered in the design of a total area network:

- Scalability of the design—there are finite limits to the number of elements that a single network manager can handle. Large networks will need complex management systems.
- Diversity—the bigger the network, the more likely it is to contain diverse elements. Many useful and important network components have proprietary management feeds, and this diversity must be recognized. A management platform that can integrate various feeds (not just SNMP) into a single big picture is better than multiple consoles, each managing a specific technology.
- Performance impact—with centrally managed systems, it is all too easy to gather too much data too often, impacting the network's overall performance. A distributed approach to management can address this issue.
- Topology—choosing the best way to connect the management system to the network can critically affect the quality of the information gathered and the effectiveness of the management system. In larger networks, there can be a case for multiple views.
- An application view—even if we have designed an optimum management system, this may not be enough. It is all too common for a network to appear perfectly healthy to the management system, but for the customer's end-to-end application to fail. Effective management systems need to increasingly look at the network from the point of view of the end-to-end application. This could be done by systems able to intelligently monitor application traffic or by combining feeds from the host/server management systems with those of the network. In an ideal world, as soon as the management system flags that a customer site has lost its application session with the host/server, the network manager should be proactively investigating to see if it is a network problem!

As stated at the beginning of this chapter, management is not a science. It relies on experience and a systematic approach. This chapter has set some guidelines to help in both of these areas.

Team LiB

◀ PREVIOUS

NEXT ▶

Selected Bibliography

Held,G.,*Network Management-Techniques, Tools, and Systems*,Chichister, England:John Wiley & Sons,1992.

Steinberg,L.,*Troubleshooting with SNMP and Analyzing MIBs*,London:McGraw Hill,2000.

Telecommunications Operation Map,<http://www.tmf.org/clickmap/TOMv2.1/index.htm>.

Team LiB

◀ PREVIOUS

NEXT ▶

Chapter 9: Through the Looking Glass

Overview

We demand rigidly defined areas of doubt and uncertainty.

–Vroomfondel, in *The Hitchhiker's Guide to the Galaxy* by Douglas Adams

The main reason for the popularity of the Ethernet is that it has continued to evolve to meet data application needs over the years. In the process, it has retained *compatibility* (i.e., the essential functional and service elements that allow successive generations of Ethernet to work with another). Given that the best guide to the future is to look at an established track record, it is quite easy to make predictions about Ethernet.

At the most fundamental level, the basic frame format, the techniques for medium access and bandwidth arbitration will all persist. This might appear fairly undramatic until we look at all the variations that are played a basic theme. The extension of the range and speed of Ethernet, along with the addition of EtherLoop, wireless LAN, and others, pushes a once local network technology into many new areas. With optical links and switches on the horizon, there is prospect to exploit Ethernet's compatibility attribute to new levels.

One further prediction would be that all of the technologies and resources that surround and support Ethernet will continue to expand and be refined. An abundance of transmission capacity, a uniformity of communication protocols, and the availability of management and security solutions will make the Ethernet yet more useful and allow it to be a practical total area network solution.

This short chapter takes a look ahead at the prospects for total area networking, post Gigabit Ethernet.

9.1 Lord of the Rings

As we saw in [Chapter 1](#), there is considerable excess capacity in the long-distance market, and this is largely made up of dark or unlit fiber. Advances in transmission technology (notably DWDM) have significantly increased the capacity of existing fiber, and more is promised. So there is a vast amount of untapped capacity available at small incremental cost, and this will continue to drive down the prices of long-distance carriage. In practice, the cost of carriage is becoming almost distance independent because practically all the cost lies in the electronics at either end.

Having all of the bandwidth you will ever need in place is a good start, but there is more to networks than plumbing. Control has to be applied to the raw capacity.

A true "all-optical network" is not something that is just about to happen in the Ethernet space. The reason for this is that there are no technologies established to enable packet switching at the optical layer today. However, purely optical Ethernet switches have already been demonstrated, in the sense that all of the Ethernet interface ports are optical (while the internal switching remains electrical). And there is pressure for progress in this direction.

When a carrier's customers are spread over metropolitan distances (10 km-50 km), then the lowest-cost service-provider network is Ethernet, with all of the device ports being optical. Also, businesses are requesting native Ethernet services to interconnect their facilities into virtual LANs. This trend will be accelerated by the fiber-to-the-consumer move, as the lowest-cost optical service will once again be Ethernet. So the bulk of early fiber-to-the-home (FTTH), and possibly the majority of fiber-to-the-business (FTTB) services, may well be optical Ethernet.

In order to turn this trend into a preferred option, there has to be some way of ensuring that the network has requisite levels of performance, reliability, and availability. The practicalities of total area network as we have it today are slightly challenged in this respect, as it carries with it the use of existing wide area technology, a hangover from the telephone network.

For example, because of its QoS and virtual circuit capabilities, ATM is frequently used, with SONET as the physical layer, to transport packet data. This is not only restrictive in terms of network flexibility (as circuits usually have to be set up in advance of use), it also introduces inefficiency because of the protocol overheads in ATM.

A step forward can be made by dropping the ATM layer and inserting packets directly into the SONET format. Unsurprisingly, this approach is known as POS, which uses the Point-to-Point Protocol (PPP) for data encapsulation at the data link layer. This provides efficient packet delineation and error control between routers at either end of an optical link. Like the other protocols we discussed earlier, PPP is defined by a frame, inside which information is wrapped.

PPP is actually a slightly modified version of the well-known and long-established HDLC. The use of PPP encapsulation over SONET/SDH links is specified in RFC 2615. Although designed for use on point-to-point links, PPP is also suitable for SONET/SDH links, which are provisioned as ring topologies. POS frames are mapped into SONET/SDH frames, and they sit in the payload envelope as octet streams aligned on octet boundaries.

Having dropped one layer from our transmission link, the obvious next step is to look at dropping another. Despite providing a high degree of resilience, SONET is a throwback to telephony networks. If the packets and frames that constitute the information across an Ethernet could be carried direct onto fibers, then complexity is reduced and efficiency enhanced.

While not formally a part of Ethernet, the IEEE 802.17 committee is creating a standard for packet transport over fiber-optic rings. This is known as RPR. Current directions include using Ethernet framing in SONET-style rings. The goal of RPR is simple: to define a high-performance, high-availability optical transport suitable for carrier networks in metropolitan service areas. One of the tenets for RPR is that it is easier to implement a fast and robust link-failure recovery in a ring topology than in a mesh topology; this is because in a ring, the alternate route is always known.

There are, of course, practical issues yet to be resolved with RPR, such as speed of network restoration in the event of failure. That said, the trend is toward fewer layers of protocol, and any

problems that arise when a layer is dropped are usually overcome.

With the abundance of fiber and the prospect of this being suitable for the direct transmission of packets, the total area network based on Ethernet appears yet more likely.

It would be wrong to suggest that Ethernet's major breakthrough will be in the core network. It is well set to play a huge part in the development of the local access network. The average residential user has had to wait about 20 years for ISDN to move from the laboratory to deployment. DSL has been a little faster, but not much. In both cases, heavy reliance is placed on the condition of the local loop, and the offered bandwidth can hardly be considered "broadband"-between a few hundred kilobits per second and one or two megabits per second. By way of contrast, Etherloop and 802.11 have developed quickly, are very competitively priced, and (perhaps most importantly) offer the sort of bandwidth that really does enable new services and applications to be offered-several megabits per second and more.

If there is an immense amount of capacity in a core network, then there will inevitably be demand for access to it. The development of the core and local network go hand in hand-the former as supplier, the latter as consumer. With Ethernet uniquely present in both markets, it seems more of an inevitability than a prediction that it will be the basis for the total area network.

Team LiB

[◀ PREVIOUS](#) [NEXT ▶](#)

9.2 From Here to Eternity

At the same time that life is becoming less complex in the area of data transmission, it is also becoming a little less complicated at the applications end. In [Chapter 7](#), we briefly explained the concept of Web services, a framework for publishing, distributing, and using software components. It seems that there is every prospect of uniformity in the applications area as well.

This is an important issue, as there is a trend to fragment and reform organizations into an internal marketplace of ideas, skills, and projects. There is an erosion of existing "command and control" structures into a flatter and looser federation of relatively independent units. Some corporations will turn into shells of their former selves, mainly concerned with maintaining brand recognition, while a loosely coupled network of suppliers spread around the world does all the real work. In the absence of the uniformity enabled by Web service, people would most likely be spending all their time unpicking and rebuilding bespoke systems and services as organizations change.

So, in both the applications and network areas, less complexity, leading to greater flexibility, is a prime aim. If a technology cannot cope with the flux in requirements brought on by continual organizational change, then it is not likely to last long. Flexibility has doubtless been the reason for the success of Ethernet in the past and is likely to be the key to its continued success. But will it disappear into the background as a useful but mundane support to trendier technology, such as Web services?

The answer is probably no. Throughout the history of communications, connections between people have consistently been more important (and lucrative) than access to information. For instance, the U.S. Postal System of the 1800s was awash with newspapers that needed to be delivered. This was the dominant content, at least in terms of volume, weighing about 20 times as much as the letters carried. But it was the letters that brought in most of the money needed to run the postal system—around 85%. On the Internet, e-mail is still the king, even if its volume is relatively small.

The real demand, and, hence, most of the money, is in peer-to-peer communication.

Team LiB

◀ PREVIOUS

NEXT ▶

9.3 Summary

The simple premise in this chapter is that Ethernet will be best served in expanding its sphere of usefulness by simplification. The specific example we use is the way in which packets are carried over the now abundant fiber links that span the globe. Removing the extra protocol layers, such as ATM and SONET, that have been used to add resilience and assurance leads to a more efficient and, critically, more flexible solution.

Charles Darwin observed that it is not the strongest that survive but those most adaptable to change. If there is one thing that Ethernet has demonstrated in spades over the years, it is adaptability.

Team LiB

◀ PREVIOUS

NEXT ▶

Appendix A: Complementary Technologies

Overview

The highest type of efficiency is that which can utilize existing material to the best advantage.
--Jawaharlal Nehru

Flexible as Ethernet is, there are still a few things that it does not do. A complete communications solution usually involves more than the mechanism for connecting someone with others. In a total area network, there has to be some way of transmitting information over long distances, of controlling an end-to-end dialogue, and of building resilience into the network.

In this appendix, we look at the technologies that would tend to be used alongside, rather than instead of, Ethernet. Of course, this is by no means a comprehensive guide, nor is it any sort of indication that Ethernet plus those explained in this appendix form a complete set. Nonetheless, each of the technologies presented can add something to an Ethernet-based design.

The main aim in this appendix is to characterize each of the chosen technologies, not to assess them. As far as possible, we stick to straightforward facts. The suitability of a technology for a specific deployment relies too much on target cost, reliability, flexibility, ease of integration, and many other factors for any sort of judgment to be made.

A.1 DWDM

The exploitation of DWDM has fueled an explosion in transmission capacity. The amount of information that can be sent over the fiber cables that span the world [such as Fiber Link Across the Globe (FLAG)] has increased so much that there is now a glut of available capacity.

With DWDM, the light within an optic fiber is split up into a number of discrete wavelengths, or "colors." Each color can carry a considerable payload: 2.4 Gbps is typical. Systems with 16 colors are now in common use. An extension to the bandwidth of a basic color from 2.4 Gbps to 10 Gbps is already available. Commercial systems operating at 40 Gbps have been demonstrated, and 80-Gbps equipment is only a few years off. In addition, fibers carrying 64 and more wavelengths are likely to be available in the same timeframe.

The technical aspects of DWDM have been defined by the International Telecommunication Union (ITU). A laser grid for point-to-point (DWDM) systems has been specified based on a pattern of 100-GHz wavelength spacing. This means that there are 45 defined wave lengths in a range from 196.1 THz, which equates to a far infrared laser wavelength of 1528.77 μm , to 191.7 THz (1563.86 μm).

In practice, more can be wrung out of DWDM systems by extending the upper or lower bounds of the available transmission window or by spacing wavelengths more closely, typically at 50 GHz, or even to 25 GHz. In doing this, suppliers can double or triple the number of channels. Each optical channel can currently be routinely used for transmission of light pulses at 10 Gbps, or even higher data rates at 100-GHz spacing, and, with the help of DWDM technology, a pair of fibers can provide data capacity of several hundred gigabits per second.

Advances in DWDM (i.e., more channels, higher rates) depend on developments in fiber terminating and repeater equipment. They do *not* require any upgrade or replacement of the fiber infrastructure that has been put in the ground. Hence, we can upgrade links from one capacity level to the next simply by reconfiguring or upgrading terminal equipment and repeaters.

Of course, the raw capacity carried on the fiber infrastructure must be structured in some way so that it can carry useful traffic and be routed where it needs to go. This is where SDH comes into play. SDH, and SONET, its equivalent in the United States, is a multiplexing transmission carrier system in which lower bit-rate channels are interleaved into a higher level, fixed-length, frame structure and transmitted in a "hierarchy" of successive levels. The levels commonly deployed are STM-1 at 155 Mbps, STM-4 at 622 Mbps, and STM-64 at 2.4 Gbps.

[Table A.1](#) relates the SDH and SONET signal names to the capacity available and traffic that can be carried.

Table A.1: The SDH/SONET Rates

SDH Signal	SONET Signal	Bit Rate	Capacity—Bearer Circuits	Capacity—Voice (Mmins/month)
STM-0	STS-1/OC1	51.48 Mbps	630 ISDN channels	5
STM-1	STS-3/OC3	155 Mbps	63E1 or 3E3	15
STM-4	STS-12/OC12	622 Mbps	252E1 or 4E4	60
STM-16	STS-48/OC48	2.488 Gbps	1008E1 or 16E4	240
STM-64	STS-192/OC192	9.95 Gbps	4032E1 or 64E4	960

A.2SDH and SONET

SDH and SONET are the European and U.S. versions, respectively, of the dominant high-capacity transmission technology in current use. The rates at which they operate were introduced in [Section A.1](#). We will now explain SONET structure and operation.

SONET is the TDM-based standard for high-speed transmission. It is widely used by telecommunications providers to move voice traffic between switches and is increasingly being used as part of both public and private data networks. One of the attractive features of SONET is that it allows access to an individual bit stream within the higher order multiplex. Hence, it is possible to extract a specific signal, and pieces of equipment known as add-drop multiplexers (ADMs) are available to insert/extract specific SONET streams.

The SONET standard is based on a rate of 51.84 Mbps, termed the synchronous transport signal level 1 (STS-1). An STS-1 frame consists of 90 columns and 9 rows of 8-bit bytes. The first three columns are termed "transport overhead" and dedicated to network management information for the section and line parts of a SONET network. The rest of the 87 columns and 9 rows carry the STS-1 synchronous payload envelope (SPE).

The SONET protocol overhead is significantly larger than that found in asynchronous DS1 signaling, which sends a limited amount of information in-band by "robbing" bits from the traffic itself. It is separated into layers that match the segments of a telephone network with the section and line parts contained in the transport overhead and the path part included with the payload.

This layered approach allows different types of equipment to be built to support different functions. The section layer defines the network segment between regenerators (the optical version of an electrical repeater). The section layer's job is to transport overhead traffic for both line and path layers, as well as the actual network traffic. Framing, scrambling, error monitoring, and order wiring (the capability to install or remove a service) are all done in the section layer.

The line is that portion of a network between line terminating equipment where the STS-1 signals are multiplexed to higher rates. It handles synchronization, multiplexing, automatic protective switching, and additional error monitoring. In addition, it provides a data communications channel, a channel for priority installation or removal of service, and room for growth. Automatic protection switching is intended to allow switching traffic from a primary to a backup circuit if the quality of the primary circuit drops below a certain threshold. Line overhead is intended to ensure that the path payload, whether data, video, or voice, is reliably transported.

The path overhead is included as part of the STS-1 payload and is carried from end to end. It includes end-to-end error checking, an identifier for the type of payload being carried, and a status report on maintenance signals. The path overhead also maps services such as DS3s into the SONET format and is accessed by terminating devices such as add/drop multiplexers.

One of the unusual features of SONET is that each overhead layer protects the layer beneath it. So, if the top-level signal is correctly composed and transmitted, the layers under it will also be properly delivered.

SONET uses bit-interleaved parity (BIP), a 1-byte code (consisting of 8 bits) to furnish parity for SONET frames. The section BIP furnishes parity from regenerator to regenerator. If section parity is correctly transmitted, so should the parity bits used by other layers. Parity between terminating devices is done by line BIP, which covers line, path, and traffic segments of a frame. Path BIP sets up parity between line termination equipment and covers path and traffic sections of a frame.

SONET is often deployed in the form of a set of transmission rings. This provides an alternative signal path in the event of failure—if a link is lost, data is transmitted around the "long" part of the ring.

In addition to carrying bulk traffic, nodes on the ring can deliver a low rate bit-stream to a specific location. ADMs can be used to insert and extract from the higher rate hierarchy. Hence, SONET/SDH can be used for both core transmission and point-to-point delivery.

One of the principal benefits of SONET is that it allows for the direct multiplexing of established network carrier rates, such as DS1, E3, and DS3, into the synchronous payload. So, for example, with

direct multiplexing, 28 DS1s or 7 DS2s are combined into the 51.84-Mbps bandwidth.

Team LiB

[← PREVIOUS](#) [NEXT →](#)

A.3 IP

The IP is a connectionless packet protocol that has made it possible for millions of computers to be connected. Since its introduction around 30 years ago, IP has become the protocol that supports just about every application.

The IP (version 4) is described in detail in RFC 791. Like Ethernet, IP is characterized by its packet format. [Table A.2](#) describes the structure of the IP packet (that is, the fields it contains):

Table A.2: Fields in the IPv4Packet

Version	IP version number (currently 4)
Internet header length (IHL)	Packet header length (measured in 32-bit words)
Type of service	Priority and QoS settings
Total length	Length of IP packet (including octets in the header)
Identification	With LAN datagrams of the order of 4.5K bytes (token ring) and 1.5K bytes (Ethernet), fragmentation of the IP packet can occur. This is a unique number assigned to a packet fragment that helps with reassembly.
Flags	Used in the control of fragments
Fragment offset	Specifies offset in the original packet of data being carried in the fragment (measured in 8 octet units)
Time to live (TTL)	Time that the packet can remain in the internetwork before being declared lost (and thus deleted)
Protocol	Specifies the higher layer protocol contained in data field. UDP uses the number 17 and TCP uses number 6.
Checksum	A 16-bit checksum of IP header (which checks only the header)
Source/destination address	The 32-bit IP addresses (often known as a dot address, as it is structured into four bytes and written 138.17.12.1—each byte a number between 0 and 255)

IP packets are routed across networks using established Internet layer adaptive routing protocols such as Routing Information Protocol (RIP) and Open Shortest Path First (OSPF). The routing is actually structured into a number of stages. The taxonomy is that ARP provides the initial discovery needed to attach a user to the Internet so that packets can be sent on their way. RIP/OSPF/ISIS are then used to provide adaptive routing within a routing domain (intradomain or private network), and Border Gateway Protocol (BGP) provides interdomain routing, across autonomous systems, within the global Internet framework. When a packet reaches its final destinations, the local networks on which the destination host resides, the network access layer protocols are able to route the packets directly to physical locations.

The network access layer protocol that performs this function on Ethernet networks is known as the ARP. ARP software constructs a simple "[binding](#)" table in which IP addresses and their corresponding Ethernet addresses are stored. When a packet, with associated IP address, is delivered to this software, ARP will look up its corresponding Ethernet address and route the packet direct to the host. If the IP address does not have an entry in the table, it will broadcast a message to all the devices on the local network and wait for a reply. If one of the hosts on the local network receives a broadcast message containing its IP address, it will respond with its Ethernet address. The ARP software will then cache the IP/Ethernet address pair for the next time that it is needed (as well as periodically updating and flushing the cache).

Reverse Address Resolution Protocol (RARP) is a variation on the ARP protocol. It is used by

computers that do not have disks and, therefore, cannot store any static information about network addresses. Diskless workstations can, however, request network devices for an IP address, providing they have an Ethernet address. Since Ethernet peripheral devices have hardwired addresses, the diskless workstation can use RARP to request an IP address by trading its Ethernet address.

Team LiB

[← PREVIOUS](#) [NEXT →](#)

A.4 Bluetooth

Bluetooth is a de facto open standard for short-range digital radio. It is designed to operate in the unlicensed ISM applications band that is generally available in most parts of the world (see [Table A.3](#)). The specification includes air interface protocols to allow several Bluetooth applications to intercommunicate simultaneously and to overcome external sources of interference such as domestic washing machines and commercial microwave ovens. The short range referred to is defined as up to 10-m in normal operation, although greater range/penetration can be achieved through higher output powers under some circumstances.

Table A.3: Bluetooth Frequency Allocation

Area	Frequency Band (GHz)	Bluetooth Channels
United States, Europe, and most other countries	2.400-2.4835	79
Spain	2.445-2.475	23
France	2.4465-2.4835	23

The aim of Bluetooth's promoters is to enable the intercommunication of just about any piece of apparatus with any other (where this is appropriate, of course!). Consequently, one of the main constraints on the design must be cost. When the Infra Red Interface (IRI), common on mobile phones and PCs today, was conceived, it was appreciated that to persuade equipment manufacturers to implement this interface, the cost of implementation had to be low. The target cost, set at \$5, was achieved, and more than 90% of portable PCs and an increasing number of mobile phones now have an IRI built-in.

A sophisticated radio interface is more complicated (and more flexible) than the IRI and, therefore, more expensive. The price target of \$10 per unit, however, seems to be realistic, especially if all our homes will eventually have half a dozen or so Bluetooth-equipped items operating in them, driving quantities to very high numbers. In addition to cost, size matters. With ever-decreasing form factors and weight, any new addition to a piece of electronic apparatus must be small, light, and consume minimum power from the host system or separate battery. The Bluetooth implementation is feasible in a very small footprint comprising a single chip and associated RF components and should be relatively easy to install in anticipated applications. Its low output power and sophisticated power conservation design ensures minimum power consumption.

Bluetooth has the potential for impacting many areas, including applications that would have been inconceivable a few years ago (e.g., a fridge-freezer telling a microwave oven what ingredients are available, allowing the microwave to suggest menu options!) However, one particular area where Bluetooth will have a significant impact is in the support of other wireless delivery mechanisms such as cellular telephony. While national networks are suited to delivering communication on the move or wireless to any location, purely local interconnection is better handled by a local communication system, the so-called "complementary wireless network" concept.

To deliver telephony-based services from one undefined location to another, and to distribute the services and functions at those locations, requires a hybrid solution, at the core of which is a cellular handset with a built-in Bluetooth transceiver.

Bluetooth is able to simultaneously interconnect up to eight transceivers in a "piconet" over a short range. Each Bluetooth transceiver costs the same (approximately \$10) and there is no requirement for any infrastructure. The simultaneous connectivity limit of eight devices may seem to be a serious constraint; however, several piconets can operate in close proximity, and Bluetooth devices can rapidly move from one piconet to another. In fact, Bluetooth devices need only remain a member of a piconet for the period of time required to complete a communication transaction. So devices can join and leave a local piconet frequently, effectively overcoming the eight-device constraint.

Team LiB

◀ PREVIOUS

NEXT ▶

Team LiB

◀ PREVIOUS

NEXT ▶

A.5 Summary

In this annex, we give an overview of some of the main technologies that tend to complement Ethernet. Those explained here would typically be featured as part of an overall design, either at a lower or higher level on the communications picture than Ethernet. DWDM et al. can be used to transport Ethernet packets and it is usually IP that is carried inside the Ethernet frame.

Team LiB

◀ PREVIOUS

NEXT ▶

Selected Bibliography

Cerf, V., "Networks," *Scientific American*, September 1991, pp. 333-341.

Freeman, R., *Fundamentals of Telecommunications*, John Wiley & Sons, 1999.

Freeman, R., *Telecommunications Transmission Handbook*, John Wiley & Sons, 1998.

Held, G., *Internetworking LANs and WANs*, John Wiley & Sons, 1998.

Minoli, D., *Telecommunications Technology Handbook*, Norwood, MA: Artech House, 1998.

Muller, N., *Bluetooth Demystified*, McGraw Hill, 2000.

Van Duuren J., P., Kastelein, and F. Schoute, *Fixed and Mobile Telecommunications*, Addison Wesley, 1996.

Whyte, B., *Networked Futures-Trends for Communications Systems Development*, John Wiley & Sons, 1999.

Appendix B: Competing Technologies

Overview

Technology is dominated by two types of people: those who understand what they do not manage, and those who manage what they do not understand.

--Putts Law

There are many ways of achieving a given goal. Nowhere is this more evident than in the world of communications. Every network designer knows that one of the more difficult challenges they face is to select the most appropriate technology. Each one has its merits and its limitations and, whatever the final choice, it is difficult to alter course in midstream.

Given the critical dependency on choosing the right option, it is important to know who the leading contenders are. In this appendix, we give a short introduction to some of the leading alternatives to Ethernet. The technologies explained here are ISDN and xDSL (which provide high-speed local loop access), HiperLAN and GPRS (which are alternatives for mobile data networking), and frame relay and ATM (both leading technologies for wide area data networks). There are, no doubt, others that could be considered and that will become important, so a good number of references are also included here.

As in [Appendix A](#), the aim is to characterize each of the chosen technologies, not to assess them. We stick to straightforward facts as much as possible. The suitability of a technology relies too much on target cost, reliability, flexibility, ease of integration, and many other factors for anything more judgmental.

B.1 ISDN

ISDN is an extension of the PSTN that replaces the analog local loop with a digital one. It was designed around the notion of separate channels operating at 64 Kbps, this number springing from the fact that basic, analog voice transmission requires 8K samples per second, each of which is encoded as 8 bits.

In the United Kingdom and Europe, ISDN is offered in two forms, ISDN2 and ISDN30, where the numeric suffix denotes the number of 64K channels that are provided. ISDN2, also known as basic rate access, gives you two 64-K (B or bearer) channels and a single 16-K signaling (D or delta) channel. ISDN30 is also called primary rate access and provides 30 B channels along with a D channel. In the United States, primary rate access is based around 24 B channels, with one D channel. In both cases, basic rate is intended for home use, and primary rate is meant for businesses.

In practice, there are many applications that require some number of channels between 2 and 30. High-quality videoconferencing, for example, requires around 6 B channels. There are several approaches to getting the right speed to suit a wide variety of services, a technique known as inverse multiplexing. The most common method (called BonDing, for Bandwidth on Demand Interoperability Group) can be used along with standard ISDN channels to support up to 63 combined B channels. Other options include Multilink PPP (designed for Internet traffic over ISDN) and Multirate Service (an Nx64 service, provided as part of the ISDN service).

By way of contrast, there are times when you want to use a 64-K channel to carry data of a lower speed. This is known as rate adaptation and, again, there are standards for this. The two most common are known as V.110 and V.120. Both ensure that slower data is carried safely over the higher speed bearer.

The ISDN equivalent of the telephone socket, at least in the United Kingdom, is called the Network Termination Unit or NT1. This is a box that has copper wires going back to the main telephone network on one side and a socket, like the standard phone socket only a bit wider, on the other. ISDN-compatible equipment plugs directly into the NT1. If not, you need a terminal adapter (TA), which is used to connect ISDN channels to the interfaces you get on most current computing and communications equipment (i.e., RS-232 and V.35).

The precise details of where and how you connect varies from place to place. ISDN standards use a set of reference points (such as the S/T interface between the NT1 and the TA) as a basis for interworking between devices and to define the boundary between the phone network and your private installation.

Up to eight devices can be attached to one ISDN line, and these can be placed anywhere on a "bus" connected to the S/T point, but there are limits on how far this bus stretches (typically about 200m).

Transmission over the local network with ISDN is based on a technique known as echo cancellation. This provides full-duplex operation over the two-wire subscriber loop and works by keeping a record of transmitted signals so any interference (or crosstalk) between sent and received signals is removed. The system is intended for service on twisted-pair cables for operation to around 5.5 km.

With echo cancellation, digital transmission occurs in both directions within the same bandwidth simultaneously. Both transmitter and receiver are connected through a device (known as a hybrid) that allows signals to pass in both directions at the same time. The problem here is that echo is generated by the reflection of the transmitted signal back to the user. Both near-end and far-end echo occurs. The near-end echo arises between the sender's hybrid and the cable; the far-end echo is from the receiver's hybrid device. The magnitude of the echo is such that it cannot be ignored, and the technique used to overcome this problem is echo cancellation. This works by calculating the composite echo signal and subtracting this from the incoming signal, thus restoring the true signal.

B.2DSL

DSL offers high-speed data over the local loop. It defines how a pair of modems—one located at the local telephone exchange and the other at the customer site—can be used to deliver high-speed signals over their established twisted-pair copper connection. There are several varieties of DSL:

- ADSL allocates the available bandwidth in an asymmetric spectrum so that ore data is delivered downstream (toward the user) and then is returned to the exchange in the upstream channel. ADSL is well suited for high-speed Internet/intranet access, video-on-demand, and telecommuter applications. ADSL speeds range from T1 (1.544 Mbps) and E1 (2.048 Mbps) up to 6 Mbps and beyond downstream. Upstream return channel speeds range from 64 to 384 to 640 Kbps. ADSL transmissions operate at distances up to 5 km (between the customer and the local exchange or serving central office switching system) via a single copper twisted pair.
- HDSL is a symmetric technology with speeds of 1.5 or 2.0 Mbps (upstream and downstream). Its main purpose is to replace traditional T1/E1 leased circuits. As per the standards bodies, HDSL is a two-wire implementation with an operating range somewhat more limited than that of ADSL—for distances much over 3 km, telephone companies need to install signal repeaters to deploy the service. Because HDSL is a two-wire implementation, it is deployed primarily for PBX network connections, digital-loop carrier systems, interexchange PoPs, Internet servers, and private data networks.
- SDSL is similar to HDSL in that it delivers 1.5m or 2.0 Mbps (or submultiples), but it does so using a single line, downstream toward the user and upstream. The use of a single line further limits SDSL's operating range; 10,000 feet is the practical limit for SDSL applications. Because of its symmetrical nature, it is well suited for video-conferencing applications or remote LAN access.
- VDSL is asymmetric. Its operating range is limited from 1,000 to 4,500 feet but supports very fast transmission via single twisted-pair copper. Data can travel at rates up to 51.84 Mbps from 330 to 1,000 feet with rates of up to 1.6 Mbps on the upstream return path. VDSL is positioned as the eventual modem of choice for fiber-based full-service networks. The extra bandwidth allows telephone companies to deliver HDTV programming using VDSL technology.
- RADSL, which automatically adjusts to copper quality degradation or can be manually adjusted to run at different speeds up to ADSL rates.
- IDSL or ISDN-based DSL, which "inverse multiplexes" two ISDN 64 Kbps B channels using 2B1Q coding into one 128-Kbps channel.

All of the DSL technologies have been subject to some work within the ANSI, the body that sets U.S. standards for telephone line transmission, and ETSI, its European counterpart. The ITU-T (formerly CCITT), the worldwide telecommunications standards body, has not yet addressed DSL because the systems are intended for local network services and so do not generally cross national boundaries.

In terms of maturity, ADSL is the most advanced of the lot. Implementations of ADSL are based on three types of line modulation scheme: 2B1Q, carrierless amplitude modulation (CAP), and discrete multitone (DMT).

Although fairly widely deployed in some countries, the availability of DSL is limited by the need to have a well-conditioned connection. In many cases, the quality of the local loop is not good enough to support DSL (or, if it is, the transmission speeds are reduced).

B.3GPRS

GPRS is a development of the GSM mobile telephone standard. It uses the same radio channel as a normal voice call, a channel that is 200-kHz wide but can provide additional data services.

The GSM radio channel carries a raw digital radio stream of 271 Kbps. For voice calls, this stream is divided into eight separate data streams, each carrying about 34 Kbps. After protocol and error correction overhead, about 14 Kbps is left for each voice connection—and the same for a data connection. Circuit-switched data today allocates just one of these voice channels for each data user, giving a user only 14 Kbps. GPRS works by combining up to eight channels, and since each of these can deliver up to 14 Kbps of data throughput, the net result is that a group of users can get peak rates over 100 Kbps.

But not all eight voice channels have to be used. In fact, the most economical implementations will be ones that limit the bandwidth offered to each user, leaving some capacity spare for other purposes. The GPRS standard defines a mechanism by which a mobile station can request the amount of bandwidth it desires at the time it establishes a data session.

In order to extend the GSM network to support GPRS, several new components need to be overlaid on the existing ones. These are the following:

- PCU. This is the interface between the packet radio (BTS/BSC), where packets come into and leave the network on a synchronous, connection-oriented link, and the packet network interface (known as the Gb interface), which is asynchronous and connectionless. The PCU is usually located inside the BSC in place of one of its trunking cards.
- GGSN. This holds the routing information for attached GPRS users and provides the external gateway, both to the Internet and to other external packet networks (such as the X.25 network, which is used for many financial applications).
- Serving GPRS support node (SGSN). This is the data equivalent of the mobile switch center (MSC) and visitor location register (VLR). It provides mobility management and authentication, as well as routing the packets.

In operation, a GPRS device appears to work in much the same way as a standard mobile phone—both communicate with a base station and the attached infrastructure, which provide authentication, connection, and service. There are some significant differences, though, and the main one is that GPRS allows users to be continually "connected" to the network. Instead of sending data over a fixed destination using a dial-up connection, GPRS allows packets of data to be inserted into a permanently available stream. The packets from different users in a cell are interleaved, so that the "always-on" transmission capacity is shared, with no permanently preset time slots allocated to an individual. Hence, capacity can be allocated when needed and released when not.

As previously stated, the GSM data rate of about 14.4 Kbps, available through a fixed connection, is replaced in GPRS by access to up to eight time slots with a combined capacity of about 14.4 Kbps for each of the eight time slots. The exact data rate depends on radio conditions. How much of this capacity is available can vary. As well as contention with other users, different versions of GPRS have different characteristics. For instance, class eight GPRS can handle up to five contiguous time slots—four receive and one transmit—to give a downstream data rate in excess of 50 Kbps. Class 12 allows any combination of transmit and receive within the five time slots.

All the packets that are transmitted on the available time slots are sent from the base station (BTS) by the SGSN. An SGSN can support multiple base stations. As previously indicated, the SGSN keeps track of all of the mobiles within its service area. When a mobile device sends packets of data, they go via the SGSN to the GGSN, which converts them for transmission over the desired network, which could be the Internet, one of the X.25 networks, or a private network. Packets received from the Internet (i.e., IP packets) addressed to the mobile device are received by the GGSN, forwarded to the appropriate SGSN, and then transmitted to the mobile user.

To forward packets between each other, the SGSN and GGSN encapsulate them using a specialized protocol called the GPRS tunnel protocol (GTP), which operates over the top of standard TCP/IP

protocols. The details of the SGSN and GGSN are both invisible and irrelevant to the user, who simply experiences a straightforward IP connection—it just happens to be wireless.

There are established protocols and procedures to govern the way in which a mobile device accesses and works with a GPRS network. These are quite complex—the relevant standards (notably ETSI EN 301 344—GPRS service description) give a comprehensive treatment. For now, a couple of details regarding the interaction between the mobile device and network need to be highlighted, as they are directly relevant to mobile data networking.

The first detail is the way in which the mobile device makes itself known to the network. This operation, known as an "attach", establishes a logical link between the device and the SGSN. The "attach" operation also exists in the legacy circuit-switched mobile networks, where it is used to make the mobile terminal's location and authenticity known to the network. It is only when the mobile terminal needs to move traffic over the radio interface that a "connection" is established for the purpose. As well as authenticating the device and establishing its location, the "attach" procedure sets up the address of the mobile device.

One of the options for this is to allocate a static IP address; this is known as static PDP addressing in the GPRS specifications. The static PDP address would usually be held in the Home Location Register [although it is possible to hold it in the subscriber identification module (SIM) card in the mobile device]. Alternatively, a temporary IP address can be allocated. This address, known as the dynamic address, would be linked to the mobile device for a duration determined by several possible events. For instance, the end of a session or the movement of the user away from his/her point of attachment might both trigger the allocation of another dynamic address. There are various options for the dynamic address. It can be allocated by the mobile network being visited or by the user's home network.

The second detail to note at this point is that the GPRS GGSN node will have to connect to some device on another network (e.g., a border router through the Gi interface to connect to the Internet). As well as the physical connection, the protocol used at this network interface needs to be considered part of the overall network design.

B.4 HiperLAN

HiperLAN (High-performance LAN) is actually a family of standards on digital high-speed wireless communication in the 5.15- to 5.3-GHz and the 17.1- to 17.3-GHz spectrum. Four types of HiperLAN have been proposed: HiperLAN types 1 and 2, Hiperaccess and Hiperlink. The HiperLAN standard itself just describes a common air interface and the physical layer for wireless communications equipment, thus ensuring compatible communications systems while leaving the higher level functions open to manufacturers.

HiperLAN type 2 has been developed to mainly have a wired infrastructure providing a short-range wireless access to IP, ATM, and UMTS networks. It operates in the 5.2-GHz frequency band with 100-MHz spectrum. HiperLAN/2 has a very high transmission rate up to 54 Mbps. Connections are TDM and connection-oriented, either bidirectional point-to-point or unidirectional point-to-multipoint connections. There is also a dedicated broadcast channel through which the traffic from an AP reaches all terminals. Each connection can be assigned either a simple relative priority level or a specific QoS in terms of bandwidth, delay, jitter, BER, and so on.

There are three basic layers in the HiperLAN/2: PHY layer, data link control layer (DLC), and the convergence layer (CL).

PHY—The channeling in this layer is implemented by OFDM due to its excellent performance on highly dispersive channels. The basic idea of OFDM is to transmit broadband, high data rate information by dividing the data into several interleaved, parallel bit streams, and let each bit stream modulate a separate subcarrier. The channel spacing is 20 MHz. This is a compromise between high bit rates per channel and a reasonable number of channels: 52 subcarriers are used per channel (48 subcarriers for data and 4 subcarriers tracking the phase for coherent demodulation).

OFDM provides flexibility, considering the realization of different modulation alternatives. Seven different PHY modes are specified in [Table B.1](#)

Table B.1: HiperLAN PHY Models

Mode	Modulation	Code Rate	PHY Bit Rate	Bytes/OFDM
1	BPSK	1/2	6 Mbps	3.0
2	BPSK	3/4	9 Mbps	4.5
3	QPSK	1/2	12 Mbps	6.0
4	QPSK	3/4	18 Mbps	9.0
5	16QAM	9/16	27 Mbps	13.5
6	16QAM	3/4	36 Mbps	18.0
7	64QAM	3/4	54 Mbps	27.0

DLC—This layer includes functions for both medium access and transmission (user plane), as well as terminal/user and connection handling (control plane). It consists of the following sublayers:

- MAC protocol;
- Error control (EC) protocol (or LLC);
- Radio link control (RLC) protocol (also known as RCP), with the associated signaling entities: DLC connection control, radio resource control (RRC), and association control function (ACF) ([Table B.2](#)).

Table B.2: Sublayers of the HiperLAN DLC

LLC	Logical link control	Provide means to cope with unreliable radio link through error detection and retransmission protocol.
MAC	Medium access control	In charge of sharing the capacity of the radio link among different connections.
RCP	Radio Link Control Protocol	Provides following functions:
DCC	DLC connection control	In charge of DLC connection control (e.g., connection setup procedure and connection monitoring).
RRC	Radio resource control	In charge of radio resource handling, channel monitoring, channel selection, and so on.
ACF	Association control function	In charge of association procedure, as well as reassociation procedure.

The MAC protocol is used for access to the medium (the radio link). The air interface is based on time-division duplex (TDD) and dynamic time-division multiple access (TDMA), which allows for simultaneous communication in both downlink and uplink within the same time frame (i.e., the MAC frame). The MAC frame format consists of four elements: BCH, down link (DL), up link (UL), and random access (RA). Except for the broadcast control, the duration of the fields is dynamically adapted to the current traffic situation. The whole DLC is based on scheduling an efficient MAC frame. The MAC frame and transport channels form the interface between the DLC and the PHY.

CL—This layer adapts service requests from higher layers to the service offered by the DLC and converts the higher layer packets (SDUs) into a fixed size used within the DLC. This function makes it possible to implement DLC and PHY that are independent of the fixed network to which the HiperLAN/2 network is connected.

There are currently two types of CLs defined: cell based and packet based. The former is intended for interconnection to ATM networks, the latter is used in a variety of configurations, depending on the fixed-network type.

B.5Frame Relay

Frame relay has its roots in ISDN and shares much of the ISDN philosophy. As a development of the PSTN, ISDN is intrinsically circuit-switched, but packet switching is more relevant for data networking. The first generation of packet switches (which were based on a protocol known as X.25) do not fit the ISDN model of keeping user information and signaling separate and use heavyweight error correction protocols in the comparatively error-free digital environment. As a result, something else was needed in ISDN to support data services effectively, and frame relay was defined to fill the gap.

Frame relay is a simple virtual circuit packet service developed as an ISDN bearer service for use at data rates up to 2 Mbps. It provides both switched virtual calls (SVCs) and PVCs. It follows the ISDN principle of keeping user data and signaling separate from each other.

An ISDN frame relay SVC would be set up in exactly the same way as an ordinary circuit-mode connection using ISDN common-channel signaling protocols. The difference is that in the data transfer (or conversation) phase, the user's information is switched through simple packet switches—known as frame relays—rather than circuit-mode cross-points.

PVCs can be set up, on subscription, by the network operator. Because the links are fixed, user signaling is neither needed nor provided. In frame relay information is transferred in variable length frames. In addition to the user's information, there is a header and trailer. The header contains a 10-bit label agreed between the terminal and the network at call setup time (or at subscription time if a PVC) that uniquely identifies the virtual call.

Terminals can, therefore, support many simultaneous virtual calls to different destinations, or even a mixture of SVCs and PVCs, using the identification facility to identify which virtual call each frame belongs to. Values from 16 to 991 are available to identify the user's SVCs and PVCs (others are reserved for specific purposes, i.e., 0 is used to carry call-control signaling, while 992 to 1007 carry management information).

One of the great merits of the simple data transfer protocol is that it provides a high degree of transparency to the higher layer protocols that are carried. This contrasts with the older X.25 protocol, where the scope for interaction with higher layer protocols often causes problems and can seriously impair performance and throughput.

One issue with frame relay is the absence of flow control that leaves the network open to congestion. Congestion ultimately means throwing frames away. Throwing frames away causes higher layer protocols to retransmit lost frames, which further feeds the congestion leading to the possible collapse of the network. Congestion management is therefore an important issue for network designers and network operators if these serious congestion effects are to be controlled and, preferably, avoided.

Given the flexibility of frame relay, it is important for users and network operators to agree on the nature and quality of the service to be provided. This gives the service provider an estimate of the traffic to be expected, which is essential for properly dimensioning the network, and it gives users defined levels of service that they can select to best match their requirements. The frame relay standards specify a number of parameters that characterize service quality, some relating to the demand the user will place on the network, others specifying the performance targets the network operator is expected to meet.

Until frame relay came along, the choice for wide area LAN interconnection lay between leased lines and X.25. The former tended to be expensive, especially for international interconnection, and not well matched to the bursty nature of LAN traffic. The latter is a complex protocol that tends to interfere destructively with any higher layer protocols being carried, usually degrading throughput seriously, often severely, and occasionally fatally.

Frame relay's high speed and transparency to higher layer protocols make it an almost ideal choice for interconnecting LANs over wide areas. Hence, frame relay provides an ideal service for many Internet users and one that is well matched to the type of traffic being carried in small to medium businesses.

B.6ATM

ATM was once seen by many in the telecommunications industry (author included) as the basis for the total area network. Unlike frame relay, ATM is a technology rather than a service and provides a high-speed, connection-oriented mechanism for carrying a range of services.

The basic operation of ATM is to route short packets of uniform length (called cells) at very high rates. Each cell is 53 octets long (8 bits)—up to 48 octets of data and 5 octets of addressing and control information.

This uniformity of structure allows the switching of cells to be carried out by static hardware rather than software, and it is this feature that underpins the high operational speed of ATM. In action, ATM works by waiting until a cell payload of user information is ready and then adding a cell header before passing the complete cell on to a local switch, where the cell is routed through the network to its destination.

No regard is paid to the content of the cell in this mechanistic part of the process—a uniform switching is presented to all types of traffic, hence, the suitability of ATM to multiservice networks. A variety of control bits are included in the header to secure effective delivery.

The fact that ATM needs to carry a whole range of different types of traffic—voice, video, text—cannot be ignored. Because each of these types of data has different requirements in terms of delay or error tolerance, there are a number of options defined within ATM for putting the raw information into the cells.

Within the descriptive model of ATM, there is a layer known as the adaptation layer that copes with this. The AAL sits between any service-specific function and the basic cell assembly layer.

There are four distinct AALs, each defined to support a different class of service (e.g., connection-oriented with constant bit rate, connectionless with variable bit rate). The features provided within each AAL come at a cost of reduced payload—the more sophisticated the facilities required by a service, the more payload is used to provide it. Over and above these coding options, there are also defined classes of service that allow paths and circuits to be managed effectively.

It was previously mentioned that ATM is connection-oriented. Connections are made by creating suitable entries in lookup tables in every switch en-route. There are two options for doing this. If the entries are made at subscription time, a PVC is created. If, on the other hand, the entries are made at call set-up time, an SVC is created. The latter places greater demands on switch design and has been a less common option.

There is a lot of flexibility built into the way ATM connections are established. The connections established between two sites, for instance, can be further divided into a number of virtual paths. This allows flexible interconnection of user sites—for instance, a connection may support a link between private exchanges, a videoconferencing link, and a frame relay service. ATM cards that enable a PC to connect via a LAN have been available for several years, and ATM network infrastructure developed quickly through the 1990s.

B.7 The Final Few

To complete this appendix, we take a brief look at some of the LAN technologies that have, over the years, provided viable alternatives to Ethernet.

FDDI. FDDI is a LAN technology offering 100-Mbps data transmission rates over two counter-rotating fiber-optic rings. Whereas Ethernet has used CSMA/CD as its MAC methodology, FDDI uses modified token passing. FDDI's MAC methodology is slightly different to that of token ring. Token passing guarantees that a transmitting network device has full bandwidth by ensuring that transmission can only take place once it possesses a 24-bit packet known as a token.

Each computer or group of computers can be assigned different amounts of bandwidth, known as synchronous bandwidth allocation (SBA). Frames are transmitted in a continuous (synchronous) stream and are prioritised according to the predefined SBA. Any spare network capacity can be filled by devices transmitting as and when they can (asynchronously).

FDDI is not standardized by the IEEE; instead, it was ratified by the ANSI X3T9.5 committee in 1984. It is known for its reliability, which comes, in part, from the use of fiber-optic media. Fiber optic is not subject to interference from electromagnetic or radio sources. One fiber ring is known as the primary, the other is a secondary ring that is used if the primary fails. Both rings are attached across a single controller. FDDI can support 500 connections, with the total circumference of the ring nearing 200 km (if repeaters are used, every 2 km). The copper distributed data interface (CDDI) has been developed to support the functionality of FDDI over copper media. Copper wiring cannot support the 2-km segments of fiber, but it can offer the same speeds (100 Mbps) over 100-m segments.

Computers can access the rings either by a single-access controller or a dual-access controller. As the names suggest, a single-access controller attaches to one ring, while dual-access controllers attach to both rings. Single access is cheaper, but the attached computer loses the redundancy benefits afforded by dual attachment. FDDI is interconnectable with Ethernet across FDDI bridges.

FDDI became popular in the 1990s and was positioned to be the industry standard for high-speed networking, but it has now been overtaken by Ethernet.

Token Ring. IBM's name has been synonymous with (and responsible for much of the development of) token ring since the early 1980s. The company made an agreement with Texas Instruments to jointly develop the required chipsets to support a token-passing access methodology.

The IEEE became involved with the standardization of token ring when it ratified the 802.5 standard for token passing. Token ring was originally based on shielded twisted-pair media, but it can operate over UTP and fiber optic. Some manufacturers are hoping to develop token ring so it can use legacy coaxial installations. It is hoped that this will persuade organizations to migrate from older Ethernet to token ring without having to change their wiring infrastructure.

A single computer has to act as a monitoring station for each ring. The monitoring station does not require any extra hardware or software, and all other computers are designed to act as standby monitors if the primary device fails. Monitors are responsible for the destruction of frames not removed by their source computers. Any damaged tokens also have to be renewed by the monitor.

There is a great deal of discussion concerning the advantages of token ring over Ethernet. Ethernet has a greater market share, but much of this is due to the higher cost of token-ring hardware. There have been many studies that contrast the performance characteristics of Ethernet and token ring. CSMA/CD causes a performance tail-off, under high network loading, due to the increased number of collisions and retransmissions. Conversely, token ring does not perform as well under low loading—the monitoring and managing overheads are proportionally greater. Overall though, there is not a great difference in the two architectures.

Token ring is a deterministic architecture that commonly runs at 16 Mbps. It optimizes its bandwidth through the token passing access methodology. Each network device can use the full bandwidth as soon as it possesses a special circulating 24-bit packet (token). Token ring LANs operate as follows:

- When a computer wishes to transmit, it has to receive the circulating token.

- The computer modifies the token into a frame and transmits data.
- The frame circulates to each device on the way to its destination. Each device examines the address field of the frame to determine if the frame is destined for itself. No other station will attempt to transmit; they recognize that the frame is not the token.
- When a frame arrives at its destination, it is copied, processed, and retransmitted slightly modified. The modification is designed to inform the source computer that the data has been received and processed.
- When the source computer receives its modified frame, it removes the frame from the ring and either completes its data transfer or releases the token back onto the ring so that other devices may communicate.

In its heyday, token ring was very popular and, with Ethernet, the leading LAN technology, but its development has not matched that of Ethernet.

Team LiB

[← PREVIOUS](#) [NEXT →](#)

Team LiB

◀ PREVIOUS

NEXT ▶

Selected Bibliography

Goralski, W., *Introduction to ATM Networking*, Addison Wesley, 1995.

Griffiths, J., *ISDN Explained*, Chichester, England: John Wiley & Sons, 1998.

Santamaria, A., and F. Lopez-Hernandez, *Wireless LAN Standards and Applications*, Norwood, MA: Artech House, 2001.

Tomlinson, C., *Telecommunications*, Addison Wesley, 2001.

Team LiB

◀ PREVIOUS

NEXT ▶

Glossary

Language is the dress of thought.

--Samuel Johnson

The flexibility of Ethernet not only makes it a total area networking technology but also a point of convergence for computing and telecommunications. This inevitably invites confusion, as professionals and interested parties alike trade terms that mean different things in different places. The scope for misunderstanding when volumes of specialist words are mixed is great.

Here we list here many of the key abbreviations and concepts in common use. Some of the terms have been explained in the main text, many are not. In either case, the aim is to clarify some of the more complex ideas by providing information about their context, application, and relationship to one another.

A

AAA

Authentication, authorization, and accounting. A remote-access security approach that controls network access by requiring user identification, restricting access to only particular resources, and maintaining records of use for billing and network audits.

Access control method

A methodology of distinguishing between the different LAN technologies. By regulating each workstation's physical access to the transmission medium, it directs traffic around the network and determines the order in which nodes gain access so that each user obtains an efficient service. Access methods include token ring, FDDI and CSMA/CD, a system employed by older versions of Ethernet.

Active-X

Microsoft technology for embedding information objects and application components within one another. For example, an Active-X button can be embedded in an HTML page that is displayed in a browser window.

Address

A common term used in computers, telecommunications, and data communications to designate the destination or origination of data or terminal equipment in the transmission of information. Types of addresses include hardware addresses (e.g., 0321.6B11.3310, for an Ethernet card), logical addresses (e.g., 132.146.29.11, a TCP/IP address for a workstation), or a personal address (mnorris@iee.org, to reach an individual).

Address mask

Also known as a subnet mask. It is used to identify which bits in an IP address correspond to the network address and which bits refer to a local terminal.

Address resolution

The conversion of an Internet address into its corresponding physical address (for instance, a corresponding Ethernet address).

Agent

In systems and network management, a term usually applied to a server specializing in performing management operations on the target system or network node.

A more recent use of the word (sometimes prefixed with the word "intelligent") used to describe a semiautonomous program that roams through a computing network collecting and processing data on behalf of its originator, sending back results as necessary.

Algorithm

A group of defined rules or processes for solving a problem. This might be a mathematical procedure enabling a problem to be solved in a definitive number of steps. Also refers to a precise set of instructions for carrying out some computation (e.g., the algorithm for calculating an employee's take-home pay).

Analog transmission

Transmission of a continuously variable signal as opposed to discretely variable signal. Telephony networks have traditionally been analog.

API

Application programming interface—software designed to make a computer's facilities accessible to an application program. It is the definition of the facilities offered to a programmer. All operating systems have APIs—in a networking environment, it is essential that various machines' APIs are compatible, otherwise programs would be exclusive to the machines on which they reside.

Applet

A small software component of little use on its own but which may be plugged in to a form part of a larger application. Used with World Wide Web applications, Java, and mobile code environments to provide downloadable components.

Application

A collection of software functions (and possibly components) recognized as delivering most of what is need to support a particular activity. Applications can be handcrafted pieces of software but, more often, are commercial products or assemblies of reusable, black box components. Editors, spreadsheets, and text formatters are common examples of applications. Network applications include clients such as those for FTP, electronic mail, and telnet.

Architecture

A high-level conceptual representation showing how systems and components in a domain relate to one another and may be assembled into more complex systems. Any given domain may have a number of different architectures representing different viewpoints. When applied to computer and communication systems, it denotes the logical structure or organization of the system and defines its functions, interfaces, data, and procedures. In practice, architecture is not one thing but a set of views used to control or understand complex systems. A loose definition is that it is a set of components and some rules for assembling them.

Architecture style

A set of components, topological layout, set of interaction mechanisms, environment, and possibly technology (e.g., CORBA).

ARP

Address resolution protocol is a networking protocol that provides a method for dynamically binding a high-level IP address to a low-level physical hardware address. This means, for instance, finding a host's Ethernet address from its Internet address. ARP is defined in RFC 826.

Asynch data transmission

A data transmission in which receiver and transmission and transmitter clocks are not synchronized. Each character (word/data block) is preceded by a start by and terminated by one or more stop bits, which are used at the receiver for synchronization.

Asynchronous

An arrangement where there is no correlation between system time and the data that is exchanged, carried, or transmitted over the system. For instance, an asynchronous protocol sends and receives data whenever it wants—there is no link to a master clock. The penalty for this freedom is that extra information has to be added to announce the start and stop of a communication.

ATM

Asynchronous transfer mode is a standard for high-speed fixed-size packet communications. It provides a basis for multiservice networks—those capable of carrying voice, video, text, and so on.

AUI

Attachment unit interface, autonomous unit interface, or auxiliary unit interface. In networking, the interface to a cable attachment unit that would tap into the thick coax cable. Normally, a 15-pin D-type connector with slide locks. Not used with thin Ethernet or UTP hardware technologies. Defined in IEEE 802.3 as an interface between a media attachment unit (MAU) and a network interface card (NIC). The term AUI can also refer to the rear panel port to which an AUI cable might connect.

Authentication

Authentication is a procedure that establishes the legitimacy of users and defines the parameters of the sessions they establish. As such, authentication can be thought of as a security measure that controls and defines network access. It is always the first part of a session; the range of authentication parameters that can be set depend on the specific authentication system employed.

Authentication header (AH)

A provision of IPSec that adds a "digital signature" to an IP packet. The digital signature is created through a key-controlled hashing of each packet and provides both authentication and integrity.

Automation

Systems that can operate with little or no human intervention. It is easiest to automate simple mechanical processes, hardest to automate those tasks needing common sense, creative ability, judgment, or initiative in unprecedented situations.

B

Backbone

The part of a communications network intended to carry the bulk of the traffic. All systems that have connectivity to the backbone can communicate with each other. This does not stop systems from setting up private arrangements with each other to bypass the backbone (especially a public operator's WAN) for cost, performance, or security.

Bandwidth

The range of signal frequencies that can be carried by a communications channel, measured in megahertz (MHz) for analog systems. The information-carrying capacity of a communications channel measured in kilobits per second (Kbps) for digital systems.

Basic rate interface

An ISDN term that describes the two interfaces, 64-Kbps transmission links and a 16-Kbps signaling channel, referred to as bearer links and the delta channel see also ISDN.

BGP

Border Gateway Protocol. This is the protocol used in TCP/IP networks for routing between different domains.

B channel

The ISDN term used to describe the standard 64-Kbps communications channel.

Binding

The process whereby a procedure call is linked with a procedure address or a client is linked to a server. In traditional programming languages, procedure calls are assigned an address when the program is compiled and linked. This is static binding. With late, or dynamic, binding, the communicating parties are matched at the time the program is executed.

Bits per second

The basic measurement for serial data transmission capacity, abbreviated to bps. Usually has some form of modifier—Kbps is thousands of bits per second, Mbps is millions of bits per second. Typically, a domestic user will have an Internet line running at a few tens of Kbps. Backbone links are usually 2 Mbps and more.

Blocking

A situation when a path or connection is not available because all of those available are busy. Blocking is a phenomenon of circuit-switched networks, where the designer trades off concentration against throughput. Most public-switched networks are designed with sufficient resources to allow users to gain access virtually all the time, without being blocked by other users.

Blocking switch

The switch that enables only a limited number of ports to be connected concurrently, which is less than the number of ports available.

Bridge

A device or technique used to match circuits, thereby minimizing any transmission impairment. Most commonly used to connect two segments of a LAN together.

Browser

A program that allows a person to read hypertext information. The browser gives some means of viewing the contents of nodes and navigating from one node to another. Mosaic, Lynx, and Netscape are browsers for the WWW. They act as clients to the array of remote servers on which Web pages are hosted.

Bug

An error in a program or fault in equipment. Origin of the term is not universally agreed but popular belief is that the first use in a computing context can be attributed to Vice-Admiral Grace Murray Hopper of the U.S. Navy. In the early days of valve-based electronic computing, she found that an error was caused by a genuine bug—a moth fluttering around inside the machine.

Team LiB

◀ PREVIOUS

NEXT ▶

C

Caching

This is a process by which data requested by the operating system of a computer is retrieved from RAM instead of from a hard disk (or some other mass storage media). Caching algorithms will check if the requested data is in its "cache" (or RAM). The significance of this is that RAM access is an order of magnitude faster than today's mass storage devices, so the more accesses to the cache, the faster overall system performance will be.

CCITT

Consultative Committee of the International Telegraph and Telephone. Until the early 1990s, a key standards making body for public network operators. Superseded by International Telecommunications Unit/Telecommunications (ITU/T).

CGI

Common Gateway Interface. A protocol associated with file servers for the WWW. CGI is the logical interface between an HTTP server and an application server. It allows information (e.g., records taken from a database) to be presented to a user in a standard format.

CIR

Committed information rate. In frame relay, this is the guarantee level of throughput.

Circuit

An electrical path between two points generally made up with a number of discrete components.

Circuit switching

The method of communications where a continuous path is first established by switching (making connections) and then using this path for the duration of the transmissions. Circuit switching is used in telephone networks and some newer digital data networks.

Client

Often synonymous with a PC. A client is an entity—for example, a program, process, or person—that is participating in an interaction with another entity and is taking the role of requesting (and receiving) the required service.

Client-Server

The division of an application into (at least) two parts, where one acts as the "[client](#)" (by requesting a service) and the other acts as the "[server](#)" (by providing the service). The rationale behind client-server computing is to exploit the local desktop processing power, leaving the server to govern the centrally held information.

CMIP/CMIS

Common Management Information Protocol/Service. A standard developed by the OSI to allow systems to be remotely managed. Similar in function to SNMP (although more complex to implement and less widely deployed).

COM

Common Object Model (also see [DCOM](#)). The expansion of the Component Object Model (see next entry) to add support for distribution. COM was jointly developed by Microsoft and Digital.

COM

Component Object Model. The nondistributed framework underlying Microsoft's OLE object technology.

Components

Self-contained, recognizable entities that perform well-understood function and can be assembled via known interfaces with other components to build something more complex. Components are often reused and can be replaced with an identical functioning component without affecting overall operation.

Concurrency

The case when two or more systems cooperate on a task in parallel. Concurrent operation can be efficient but is prone to undesirable states (such as deadlock, where all parties are waiting, or livelock, where there is a repeated sequence of activity with no exit defined).

Configuration

A collection of items that bear a particular relation to each other (e.g., the data configuration of a system in which classes of data and their relationships are defined).

Connectionless

This refers to a communication where two or more participants do not have a fixed path between them. Each of the packets that constitute the communication looks after its own routing. This arrangement is subject to the vagaries of network availability but can be a very efficient overall way of using a network.

Connection-oriented

This is the familiar form of communication on the telephone network. A call is initiated by setting up an end-to-end connection between participants, and this connection is kept for the duration of the call. It may not be efficient in terms of network use, but there are some assurances of delivery.

Cookie

A token of agreement between cooperating programs that is used to keep track of a transaction. At a more concrete level, a cookie is a fragment of code that holds some information about your local state—your phone number or home page reference, for instance. You probably have cookies that you do not know about. The Netscape and Explorer browsers both support them, with the cookie being presented to the server to control your dialogue.

CORBA

Common Object Request Broker Architecture. Strictly, the name of a framework specification produced by the Object Management Group describing the main components of a distributed object environment. Informally used to denote any of a number of related specifications produced by the OMG.

CPE

Customer premise equipment. Equipment that is on the customer side of the point of demarcation, as opposed to equipment that is on a carrier side.

CSMA/CD

Carrier sense multiple access with collision detection—a method used in LANs whereby a terminal station wishing to transmit *listens* and transmits only when the shared line is free. If two or more stations transmit at the same time, each backs off for a random time before retransmission. Each station monitors its signal, and if this is different from the one being transmitted, a collision is said to have occurred (collision detection). Each backs off and then tries again later.

D

Daemon

A program that lies dormant, waking up at regular intervals or waiting for some predetermined condition to occur before performing its action. Supposedly an acronym rationalized from Disk And Execution MONitor.

Dark fiber

Optical fiber fully owned and lit by an end user company like Linx, as opposed to optical fiber that is rented from a telco and lit by that telco on the end user company's behalf. The term "dark fiber" comes from the fact that cable is not already lit when you buy it. Generally, companies will rent fiber runs that are lit by a telco and merely present data to the telco's transceiver equipment.

Data

Usually taken to mean the same as information. Data is the raw input that, once interpreted and processed, can be used to provide information. For instance, a spreadsheet will usually contain some data. The fact that it shows your company to be in profit is the information!

Database

A collection of interrelated data stored together with controlled redundancy to support one or more applications. On a network, data files are organized so that users can access a pool of relevant information.

Database server

The machine that controls access to the database using client/server architecture. The server part of the program is responsible for updating records, ensuring that multiple access is available to authorized users, protecting the data, and communicating with other servers holding relevant data.

Datagram

A variety of data packet. A self-contained, independent entity of data carrying enough information to be routed from source to destination without reliance on earlier exchanges between the source and destination.

DBMS

Database management system. A set of software used to set up and maintain a database that will allow users to call up the records they require. In many cases, a DBMS will also offer report and application generating facilities.

DCE

The distributed computing environment. A set of definitions and software components for distributed computing developed by the Open Software Foundation, an industry-led consortia. It is primarily a remote procedure call with a set of integrated services, such as security, time, and directory.

DCOM

Distributed Common Object Model. Microsoft's upgrade to its initial version of COM for a distributed environment.

Deadlock

A condition where two or more processes are waiting for one of the others to do something. In the meantime, nothing happens. A condition (undesirable) that needs to be guarded against, especially in the design of databases.

Directory

A directory provides a means of translating one form of information to another (e.g.,

someone's name into their telephone number). In a distributed system, directory services are a key component. They often perform much the same function as a telephone directory—translating from a symbolic name to a network address. A well-known example is DNS, which translates Internet names (mnorris@iee.org) into their corresponding addresses (142.119.42.17).

Distributed computing

Moving away from having large centralized computers such as minicomputers and mainframes, and bringing processing power to the desktop. Often used as a synonym for distributed processing.

Distributed database

A database that allows users to gain access to records, as though they were held locally, through a database server on each of the machines holding part of the database. Every database server needs to be able to communicate with all the others, as well as be accessible to multiple users.

DLL

Dynamic link library. A set of software utilities that are bound with source code when it is compiled.

DNS

Domain name service. A general-purpose distributed, replicated, data query service used on the Internet for translating host names into Internet addresses (e.g., taking a dot address such as <http://jungle.pdq.com> and returning the corresponding numerical addresses).

DoD

The U.S. Department of Defense. Notable for its sponsorship of the network that became the Internet.

Domain

In the broad context, this is a well-understood area of common interest within which common technical terms are understood and common components can be practically applied. When applied to networked communication systems, a domain is part of a naming hierarchy. An Internet domain name consists, for example, of a sequence of names or other words separated by dots.

E

EJB

Enterprise JavaBeans. Components written in the Java programming language intended to be run within a server-based environment (e.g., a WWW server or database). EJBs run within a "container" on the server and appear as objects to the outside world. Clients locate the EJB via the Java Naming and Directory Interface (JNDI).

Electronic mail

Messages automatically passed from one computer user to another, often through computer networks and/or via modems over telephone lines.

Electronic mail address

The coding required to ensure that an electronic mail message reaches its specified destination. There are many formats of mail address, perhaps the best known being the dot address used for Internet mail (e.g., mnorris@iee.org.).

Encapsulation

The transparent enveloping of one protocol within another for the purposes of transport. Encapsulation is, along with tunneling, a favorite method for supporting multiple protocols across linked networks.

Encryption

A means of turning plain text into cipher text, hence protecting content from eavesdroppers. There are many ways of ways of doing this; public and private key encryption are the two main ones.

Enterprise

A term (usually used as a prefix for "[network](#)" or "computing") to denote the resources deployed to suit the operating needs of a particular organization.

F

FDDI

Fiber Distributed Data Interface. An American National Standards Institute (ANSI) LAN standard. It is intended to carry data between computers at speeds up to 100 Mbps via fiber-optic links. It uses a counter rotating token-ring topology and is compatible with the first physical level of the ISO seven-layer model.

Federation

A union of otherwise largely independent systems to support some common purpose. Federated systems share some basic agreements or protocols to enable them to work together but are operated and managed autonomously.

File server

A machine in the LAN dedicated to providing file and data storage to other machines in the network.

Firewall

In general, this refers to the part of a system designed to isolate it from the threat of external interference (both malicious and unintentional).

Firewall machine

This is a dedicated machine that usually sits between a public network and a private one (e.g., between an organization's private network and the Internet). The machine has special security precautions loaded onto it and used to filter access to and from outside network connections and dial-in lines. The general idea is to protect the more loosely administered machines hidden behind the firewall from abuse.

Flow Control

In a packet-switched network, packets compete dynamically for the network's resources—storage, processing, transmission. Flow control is a mechanism for ensuring fairness and controlling congestion.

FTP

File transfer protocol. The Internet standard (as defined in the RFC series) high-level protocol for transferring files from one computer to another. A widely used de facto standard (c.f. the sparingly used *de jure* standard FTP). Anonymous FTP is a common way of allowing limited access to publicly available files via an anonymous login.

G

G.703

The CCITT standard for the physical and logical transmissions over digital circuit. Specifications include the U.S. 1.544 Mbps, as well as the European 2.048 Mbps that use the CCITT recommended physical and electrical data interface.

Gateway

Hardware and software that connect incompatible networks, which enables data to be passed from one network to another. The gateway performs the necessary protocol conversions.

GUI

Graphical user interface. An interface that enables a user to interact with a computer using graphical images and a pointing device rather than a character-based display and keyboard. Such interfaces are also known as "WIMP" (for Windows, Icons, Menus and Pointers) interfaces. The most common pointing device is that electronic rodent—the mouse.

H

Hardware

The physical equipment in a computer system. It is usually contrasted with software.

Heritage system

A euphemism for an old or decrepit system—synonyms include legacy system, cherished system, white elephant and millstone.

Heterogeneous

Of mixed or different type.

Homogeneous

Of the same type.

HTML

Hyper-Text Markup Language. HTML is the language used to describe the formatting in WWW documents. It is an SGML document type definition. It is described in the RFC series of standards.

HTTP

Hyper-Text Transfer Protocol. The basic protocol underlying the WWW. It is a simple, stateless request-response protocol. Its format and use is rather similar to SMTP. HTTP is defined as one of the Internet's RFC series, generated by the IAB.

I

IAB

Internet Activities Board. The influential panel that guides the technical standards adopted over the Internet. Responsible for the widely accepted TCP/IP family of protocols. More recently, the IAB has accepted SNMP as its approved network management protocol.

IDL

Interface Definition Language. A notation that allows programs to be written for distribution. An IDL compiler generates stubs that provide a uniform interface to remote resources. IDL is used in conjunction with remote procedure calls.

IEE

United Kingdom's equivalent of the IEEE.

IEEE

The Institute of Electrical and Electronic Engineers. U.S.-based professional body covering network and computing engineering.

IETF

Internet Engineering Task Force. The IETF is a large, open international community of network designers, operators, vendors, and researchers whose purpose is to co-ordinate the operation, management, and evolution of the Internet and to resolve short-and mid-range protocol and architectural issues. It is a major source of proposals for protocol standards that are submitted to the IAB for final approval.

Interface

The boundary between two things: typically two programs, two pieces of hardware, a computer and its user, or a project manager and the customer. The channel between the two entities is a conduit through which information passes. Information can consist of data or commands. An API defines the commands and data that, when sent through the channel, enables a software application to be controlled.

Internet

A concatenation of many individual TCP/IP sites into one single logical network, all sharing a common addressing scheme and naming convention.

internet

With a lower case "i," this term denotes any set of networks interconnected with routers.

Internet address

The 32-bit host address defined by the IP in RFC 791. The Internet address is usually expressed in dot notation (e.g., 128.121.4.5). The address can be split into a network number (or network address) and a host number unique to each host on the network and sometimes also a subnet address.

The dramatic growth in the number of Internet users over the last few years has led to a shortage of new addresses. This is one of the issues being addressed by the introduction of a new version of IP, IPv6.

Interoperate

The ability of computers from different vendors to work together using a common set of protocols (i.e., Sun computers, IBM PCs, Apple Macintoshes, and so on all work together, allowing each to communicate with and use the resources of the other).

Intranet

A closed internet system running within an organization connected to the Internet and protected from unauthorized access, via a firewall. Increasingly, businesses are using

intranets to provide employees with desktop access to key business systems.

IP

The ubiquitous Internet Protocol, one of the key parts of the Internet. IP is a connectionless (i.e., each packet looks after its own delivery) switching protocol. It provides packet routing, fragmentation, and reassembly to support the Transmission Control Protocol (TCP). IP is defined in RFC 791.

IP address

The Internet Protocol address. This is a 32-bit address that has to be assigned to any computer that connects to the Internet. A typical IP address takes the form 192.61.33.11 and comprises a host component and a network component.

IPv6

The proposed successor to IP. It is a longer address but still compatible with IP. The aims are to extend the available address space, improve security, support user authentication, and cater for delay-sensitive traffic.

ISO

Commonly believed to stand for International Standards Organization. In fact, ISO is not an abbreviation—it is intended to signify commonality (from Greek Iso = same). The ISO is responsible for many data communications standards. A well-known standard produced by ISO is the seven-layer Open Systems Interconnection (OSI) model.

Isochronous

Data transmission in which a transmitter uses a synchronous clock but the receiver does not. The receiver detects messages by start/stop bits as in asynchronous transmission.

ISP

Internet Service Provider. This is most people's first point of contact with the Internet. An ISP usually offers dial-up access via SLIP or PPP connections to a server on the Internet. Most ISPs also offer their customers a range of client software that can be used on the Internet.

J

Java

A programming language and environment for developing mobile code applications. Java is a subset of the C++ language and is widely used to provide mobile code application for use over the Internet.

JavaBean

Components written in the Java programming language, originally intended to be delivered over the Internet and run on a desktop client PC.

JVM

Java Virtual Machine. The ubiquitous engine for running Java code—in essence, a software CPU. The idea is that any computer can equip itself with a JVM, a small program that allows Java applets (which are widely available over the Internet) to be downloaded and used.

Team LiB

◀ PREVIOUS

NEXT ▶

K

Kernel

The level of an operating system that contains the system level commands—the functions hidden from the user. This program is always running on a processor.

Key

The record identifier used in many information retrieval systems (i.e., database keys). The term is also used for the secret codes used to encrypt and decrypt information transmitted over public networks.

Team LiB

◀ PREVIOUS

NEXT ▶

Team LiB

◀ PREVIOUS

NEXT ▶

L

LAN

Local area network. A data communications network used to interconnect data terminal equipment distributed over a limited area.

Life cycle

A defined set of stages through which a development passes over time—from requirement analysis to maintenance. Common examples are the waterfall (for sequential, staged developments) and the spiral (for iterative, incremental developments). Life cycles do not map to reality too closely but do provide some basis for measurement and, hence, control.

Team LiB

◀ PREVIOUS

NEXT ▶

M

MAC

Media access control. This controls access to the shared transmission medium by framing/deframing data units, error checking, and providing access rights. Conceptually, the MAC is part of data link control and is important in the operation of LANs.

Mainframe

A computer (usually a large one, and often manufactured by IBM) that provides a wide range of applications to connected terminals.

Media converter

Device that converts data passing from one media to another, such as from fiber to copper.

Message passing

Communication through the exchange of messages. Although not a rigorously used term, message passing systems usually have the connotation of real-time immediate message exchange.

Message queuing

A message passing technology augmented by a store-and-forward capability.

Messaging

Exchanging messages. Often but not limited to the context of electronic mail.

Method

A way of doing something—a defined approach to achieving the various phases of the life cycle. Methods are usually regarded as functionally similar to tools (e.g., a specific tool will support a particular method).

Methodology

Strictly, the science or study of methods. More frequently used as a more important sounding synonym for method, process, or technique.

MIB

Management information base. The data schema that defines information available in an SNMP-enabled device. MIB (now at version 2) is defined in RFC 1213.

Middleware

Software that mediates between an application program and an underlying set of utilities, such as a database, a network, or a server. It manages the interaction between disparate applications across the heterogeneous platform, masking diversity from the programmer. Object request brokers such as CORBA are an example of middleware, as they manage communication between objects, irrespective of their location.

Mips

Millions of instructions per second—one measure of a computer's processing power is how many instructions per second it can handle.

Mobile code

Programs capable of being run on many different systems (e.g., Java can run on any machine equipped with a Java Virtual Machine). Mobile code is "write once, use anywhere" and gets around the need for porting work to be done every time that the program encounters a different type of computer.

Multiplexing

The sharing of common transmission medium for the simultaneous transmission of a number of independent information signals.

Multiprocessing

Running multiple processes or tasks simultaneously. This is possible when a machine or has more than one processor or processing is shared among a network of uniprocessor machines.

See also [Multitasking](#).

See also [Multithreading](#).

Multiprocessor

A single computer having more than one processor and capable of executing more than one program at once.

Multitasking

Performing (or seeming to perform) more than one task at a time. Multitasking operating systems such as Windows, OS/2, or UNIX give the illusion to a user of running more than one program at once on a machine with a single processor. This is done by "time-slicing", dividing a processor into small chunks that are allocated, in turn, to competing tasks.

Multithreading

Running multiple threads of execution within a single process. This is a lower level of granularity than multiprocessing or multitasking. Threads within a process share access to the process' memory and other resources. Threads may be "time-sliced" on a uniprocessor system or executed in parallel on a multiprocessor system.

N

Network

A general term used to describe the interconnection of computers and their peripheral devices by communications channels. For example, public switched telephone network (PSTN), packet switched data network (PSDN), local area network (LAN), wide area network (WAN).

Network interface

The circuitry that connects a node to the network, usually in the form of a card fitted into one of the expansion slots on the back of the machine. It works with the network software and operating system to transmit and receive messages over the network to other connected devices.

Network operating system

A network operating system (NOS) extends some of the facilities of a local operating system across a LAN. It commonly provides facilities such as access to shared file storage and printers. Examples include Novell @146s NetWare and Microsoft's LAN Manager.

Network topology

The geometry of the network relating to the way the nodes are interconnected.

NFS

Network file system. A method, developed by Sun Microsystems, that allows computers to share files across a network as if they were local.

Nonblocking switch

A switch that facilitates all ports to have simultaneous access through it.

Nonproprietary

Software and hardware that is not bound to one manufacturer's platform. Equipment that is designed to the specification that can accommodate other companies' products. The advantage of nonproprietary equipment is that a user has more freedom of choice and a larger scope. The disadvantage is when it does not work, you may be on your own.

O

Object

An abstract, encapsulated entity that provides a well-defined service via a well-defined interface. An object belongs to a particular class that defines its type and properties. One object can inherit properties from another, and objects can evolve to do specific tasks.

Object-orientation

A philosophy that breaks a problem into a number of cooperating objects. Object-oriented design is becoming increasingly popular in both software engineering and related domains; for example, in the specification of component-based systems.

Object program

The translated versions of a program that have been assembled or compiled. Nothing to do with object-orientation!

ODP

Open Distributed Processing. One of a number of organizations (most of which have the word "open" in their title) that provides standards and/or components that allow computers from different vendors to interwork.

OLE

Object linking and embedding. Microsoft's proprietary object component technology. Often compared to CORBA.

OMG

Object Management Group. An industry consortium responsible for the CORBA specifications.

Open system

A much abused term! The usual meaning of an open system is one built to conform published, standard specifications or interfaces, for example, POSIX. Openness is rather like beauty in that it is often in the eye of the beholder.

Operating system

Software such as VME, MVS, OS/2, Windows, VMS, MS-DOS, or UNIX that manages the computer's hardware and software. Unless it intentionally hands over to another program, an operating system runs programs and controls system resources and peripheral devices.

OSI

Open Systems Interconnection. A model to support the interworking of telecommunications systems. The ISO Reference Model consisting of seven protocol layers. These are the application, presentation, session, transport, network, link, and physical layers.

The concept of the protocols is to provide manufacturers and suppliers of communications equipment with a standard that will provide reliable communications across a broad range of equipment types. Also more broadly applied to a range of related computing and network standards.

OSPF

Open Shortest Path First. A routing protocol for TCP/IP networks.

OSS

Operational support systems. These are all of the behind-the-scenes systems that allow a service to operate reliably and profitably. The usual spread of OSS includes billing, fault and problem management, provisioning, network and service management, customer

handling, and management information.

Overloading

A term used in object-oriented software development to describe the use of one identifier that serves more than one purpose.

Team LiB

[← PREVIOUS](#) [NEXT →](#)

P

Packet

A unit of data sent across a network. The basis for all of the modern data communications networks, a common format for communications between computers.

Packet switching

The mode of operation in a data communications network whereby messages to be transmitted are first transformed into a number of smaller self-contained message units known as packets. Packets are stored at intermediate network nodes (packet-switched exchanges) and are reassembled unto a complete message at the destination. The ITU-T recommendation for packet switching is X.25.

Parallel processing

Performing more than one process in parallel. Usually associated with compute-intensive tasks that can be split up into a large number of small chunks that can be processed independently on an array of relatively inexpensive machines. Many engineering and scientific problems can be solved in this way. It is also frequently used in high-quality computer graphics.

Parameter

A variable whose value may change the operation but not the structure of some activity (e.g., an important parameter in the productivity of a program is the language used). Also commonly used to describe the inputs to and outputs from functions in programming languages. In this context, they may also be known as "arguments."

Peer to peer

Communications between two devices on an equal footing, as opposed to host/terminal or master/slave. In peer-to-peer communications, both machines have and use processing power.

Pipe

A feature of many operating systems, a pipe is a method used by processes to communicate with each other. When a program sends data to a pipe, it is transmitted directly to other processes without ever being written onto a file.

Polling

The process of interrogating terminals in a multipoint network, in turn, in a prearranged sequence by controlling the computer to determine whether the terminals are ready to transmit or receive. If any problems are detected with the normal sequence of operations, the polling sequence is temporarily interrupted while the terminal transmits or receives.

PoP

Point of presence. A site where a collection of telecommunications equipment (i.e., modems, digital leased lines, and multiprotocol routers) exists. The PoP is put in place by an ISP.

An ISP may operate several PoPs distributed throughout its area of operation to increase the chance that its subscribers will be able to reach one with a low-cost telecommunications circuit. The alternative is for subscribers to use virtual PoPs (virtual points of presence) via some third party.

Port

(n.) A device that acts as an input/output connection. Serial communication ports or parallel printer ports are examples. Also, a pleasant after-dinner drink.

(v.) To transport software from one system to another different system and make the necessary changes so that the software runs correctly, taking into account the specific calls and structures used on that system.

Process

The usual term for a program currently being run by an operating system. A process is assigned resources such as memory and processor time by the operating system. The term "task" is sometimes used as a synonym.

See also [Multiprocessing](#).

See also [Multitasking](#).

See also [Multithreading](#).

Processor

That part of a computer capable of executing instructions. More generally, any active agent capable of carrying out a set of instructions (e.g., a transaction processor for modifying a database).

Proprietary

Any item of technology that is designed to work with only one manufacturer's equipment. The opposite of the principle behind OSI.

Protocol

A set of rules and procedures used to formulate standards for information transfer between devices. Protocols can be low level (e.g., the order in which bits and bytes are sent across a wire) or high level (e.g., the way in which two programs transfer a file over the Internet).

Prototype

A scaled-down version of something, built before the complete item is built, in order to assess the feasibility or utility of the full version.

Q

Quality assessment

A systematic and independent examination to determine whether quality activities and related results comply with planned arrangements and whether these arrangements are implemented effectively and are suitable to achieve objectives.

Quality of service

Measure of the perceived quality of a service. Usually based on tangible metrics such as time to fix a fault, average delay, loss percentages, system reliability, and so on. Quality system The organizational structure, responsibilities, procedures, processes, and resources for implementing quality management.

Quality system standards

A document or documents specifying the elements of a quality system. The ISO 9001 standard (which is generally used to control software development) is a widely known and used quality standard.

Queuing

When a frame or packet is to be transmitted on a link, it may have to wait because another frame is being processed in front of it. The frame is placed in a buffer until the transmitter is free. Hence, queuing systems (i.e., packet-switched systems) require buffers (matched to load and capacity) and introduce delay (as opposed to circuit switching systems, which block).

R

RARP

Reverse Address Resolution Protocol, which provides the reverse function of ARP. RARP maps a hardware address to an Internet address.

Remote procedure call

An RPC provides a distributed programming mechanism where an apparently local procedure call in a client causes an implementation of the procedure provided by a server to be invoked.

Repository

A data store holding (or pointing to) software and systems entities that designers and developers could reuse in the process of delivering new "systems solutions." The repository provides services to manage the creation, use, versions, maintenance, translation, and viewing of these entities.

Resolve

To translate an Internet name into its equivalent IP address or other DNS information.

Reuse

The process of creating software systems using existing artifacts rather than starting completely from scratch. Code, components, designs, architectures, operating systems, and patterns are all examples of artifacts that can be reused. Also methods and techniques to enhance the reusability of software.

RFC

Request for comment. A long-established series of Internet "standards" documents widely followed by commercial software developers. As well as defining common Internet protocols, RFCs often provide the implementation detail to supplement the more general guidance of ISO and other formal standards. The main vehicle for the publication of Internet standards, such as SNMP.

RIP

Routing Information Protocol, a standard gateway protocol defined in RFC 1388 that uses message broadcasts to a destination based on hop count.

RMON

Remote monitoring management information database. Developed by the IETF, this extension of SNMP MIB 2 provides a means for tracking, storing, and analyzing remote network management information.

Routers

A router operates at level 3 of the OSI model. Routers are protocol specific and act on routing information carried out by the communications protocol in the network later. A router is able to use the information it has obtained about the network topology and can choose the best route for packets to follow. Routers are independent of the physical level (layer 1) and can be used to link a number of different network types together.

Routing

The selection of a communications path for the transmission of information from source to destination.

S

Server

An object that is participating in an interaction with another object and is taking the role of providing the required service.

Service

A piece of work done; a facility provided. In the context of components, a service is the process and interface through which the predefined function of a component is accessed. Typically, services are used to "wrap" legacy systems and enable imbedded legacy functions to be used as if they were stand-alone components.

Session

The connection of two nodes on a network for the exchange of data—any live link between any two data devices.

SGML

Standard Graphical Markup Language. An international standard encoding scheme for linked textual information. HTML is a subset.

Signaling

The passing of information and instructions from one point to another for the setting up or supervision of a telephone call or message transmission.

SMTP

Simple Mail Transfer Protocol. The Internet standard for the transfer of mail messages from one processor to another. The protocol details the format and control of messages.

SNMP

Simple network management protocol. Consists of three parts: structure of management information (SMI), management information base (MIB), and the protocol itself. The SMI and MIB define and store the set of managed entities; SNMP transports information to and from these entities.

Socket

A mechanism for creating a virtual connection between processes. At the simplest level, an application opens the socket, specifies the required service, binds the socket to its destination, and then sends or receives data.

SONET

Synchronous Optical Transmission Protocol (the United Kingdom's equivalent is SDH). SONET is intended to be able to add and drop lower bit rate signals from the higher bit rate signal without needing demultiplexing. The standard defines a set of transmission rates, signals, and interfaces for fiber-optic transmission.

SQL

Structured Query Language. A widely used means of accessing the information held in a database. SQL enables a user to build reports the data held.

Stateful

When applied to a server, this term implies that it maintains knowledge about and context for its clients between requests.

Stateless

When applied to a server, this term implies that it maintains no knowledge about its clients between requests—each request is handled in isolation from all preceding and following requests.

Switch

A piece of equipment with multiple ports providing dynamic connections. Alternatively, providing a fabric of scalable (full) bandwidth with high-speed routing to each of the connected ports.

Switch port

A hardware unit connecting a node to a network.

Synchronization

The actions of maintaining the correct timing sequences for the operation of a system.

Synchronous transmission

Transmission between terminals where data is normally transmitted in blocks of binary digit streams and transmitter and receiver clocks are maintained in synchronism.

Syntax

The set of rules for combining the elements of a language (e.g., words) into permitted constructions (e.g., phrases and sentences). The set of rules does not define meaning (this is covered by semantics), nor does it depend on the use made of the final construction.

System

A collection of independently useful objects that happen to have been developed at the same time.

A collection of elements that works together, forming a coherent whole (e.g., a computer system consisting of processors, printers, disks, and so on).

System assembly

Another name for system integration.

System design

The process of establishing the overall (logical and physical) architecture of a system and then showing how real components are used to realize it.

System integration

The process of bringing together all of the components that form a system with the aim of showing that the assembly of parts operates as expected. This usually includes the construction of components to carry out missing or forgotten functions and glue to interconnect all of the components.

T

TA

Terminal adapter. A piece of equipment used with an ISDN connection to allow existing terminals to hook up. The equivalent of a modem.

TCP

Transmission control protocol. The most common transport layer protocol used on Ethernet and the Internet. TCP is built on top of IP, and the two are nearly always seen in combination as TCP/IP (which implies TCP running on top of IP). The TCP element adds reliable communication, flow-control, multiplexing, and connection-oriented communication to the basic IP transport. It is defined in RFC 793.

TCP/IP

Transmission Control Protocol/Internet Protocol. The set of data communication standards adopted, initially on the Internet, for interconnection of dissimilar networks and computing systems.

Telnet

A TCP/IP-based application that allows connection to a remote computer.

Throughput

A way of measuring the speed at which a system, computer, or link can accept, handle, and output information.

Topology

A description of the shape of a network, for example, star, bus, and ring. It can also be a template or pattern for the possible logical connections onto a network.

TP

Transaction processing. Concerned with controlling the rate of inquiries to a database. Specialist software—known as a TP monitor—allows a potential bottleneck to be managed.

Trading

Matching requests for services to appropriate suppliers of those services, based on some constraints.

Transaction

A single, atomic unit of processing. Usually a single, small "parcel" of work that should either succeed entirely or fail entirely.

Transaction processing

Originally a term that mainly applied to technology concerned with controlling the rate of inquiries to a database. Specialist software—known as a TP monitor—allowed potential bottlenecks to be managed.

Transparency

Distribution transparencies provide the ability for some of the distributed aspects of a system to be hidden from users. For example, location transparency may allow a user to access remote resources in exactly the same way as local ones.

Tunneling

Usually refers to a situation where a public network is used to connect two private domains so that the privacy of the overall link is maintained. To all intents and purposes, the same as encapsulation.

U

UML

Unified Modeling Language. The object-oriented notation adopted by the OMG and devised by Booch, Rumbaugh, and Jacobson. UML unifies several of the main (and formerly disparate flavors) of object-oriented notation.

URL

Uniform resource locator. Essentially, this is the form of address for reaching pages on the WWW. A typical URL takes the form <http://www.interesting.com/>.

V

Value chain

A process description of a business' key activities showing where business functions add value to the products or services the business provides.

Vendor independent

Hardware or software that will work with hardware and software manufactured by different vendors-the opposite of proprietary.

Virtual circuit

A logical connection across a network (e.g., the transmission path through an X.25 packet-switched data network established by exchange of setup messages between source and destination terminals).

Virtual device

A module allocated to an application by an operating system or network operating system, instead of a real or physical one. The user can then use a computer facility (keyboard, memory, disk, or port) as though it was really present. In fact, only the operating system has access to the real device.

Virtual LAN

A group of devices on one or more LANs that communicate as if they were connected to the same wire even though they are physically located on different LAN segments. Because virtual LANs are configured through software rather than hardware, they are extremely flexible.

Virtual machine

A software program that provides an implementation of an abstract processing environment. It supplies an execution engine for other programs compiled into byte code that it interprets.

Virus

A program, passed over networks, that has potentially destructive effects once it arrives. Packages such as VirusGuard are in common use to prevent infection from hostile visitors.

See also [Worm](#).

VoIP

Voice over Internet Protocol. A term applied to a set of facilities for managing the delivery of voice information using the IP. Voice information is sent in digital form in discrete packets over the Internet instead of in analog form over the public switched telephone network (PSTN). A major advantage of VoIP is that it avoids the tolls charged by ordinary telephone service. Large companies often use it to carry their voice traffic over their corporate data network.

VPN

Virtual private network. A combination of public and private resources that has been combined to give users a network that looks like a coherent resource, suited to their particular needs. For all intents and purposes, a VPN is an enterprise network.

W

W3

Common abbreviation for World Wide Web (WWW), the Internet-based distributed information retrieval system that uses hypertext to link multimedia documents. This makes the relationship of information that is common between documents easily accessible and completely independent of physical location.

WWW is a client-server system. The client software takes the form of a "[browser](#)" that allows the user to easily navigate the information on-line. Well-known browsers are Netscape and Mosaic. A huge amount of information can be found on WWW servers.

WAP

Wireless application protocol. A specification for a set of communication protocols to standardize the way that wireless devices, such as cellular telephones and radio transceivers, can be used for Internet access, including e-mail, the WWW, and newsgroups.

WDM

Wavelength division multiplexing. An optical transmission technique in which a number of separate wavelengths (each carrying its own information), are combined for transmission over a single optical fiber. At the receiving end, the wavelengths are separated and directed to separate receivers.

WAN

Wide area network. Any network that covers a wide geographic area (i.e., a region or country). The PSTN is wide area, as is a data network that extends a LAN outside a building or beyond a campus, over leased lines or a bearer network to link to other LANs at remote sites. A data WAN is typically created by using bridges or routers to connect geographically separated LANs.

Window

A flow control mechanism the size of which determines the number of packets that can be sent before an acknowledgement of receipt is needed and before more can be transmitted.

Windows

A way of displaying information on a screen so that users can do the equivalent of looking at several pieces of paper at once. Each window can be manipulated for closer examination or amendment. This technique allows the user to look at two files at once or even to run more than one program simultaneously. Also the generic name (though not a registered trademark) for Microsoft's family of operating systems-Windows 97®, Windows NT®.

WINS

Windows Internet Name Service. A Microsoft product that manages the mapping between resource names (in the form of easy-to-remember nicknames) and IP addresses. The DNS service used on the Internet cannot map between IP addresses and local resource names dynamically. However, through dynamic database updates, WINS lets users access network resources via more user-friendly names instead of IP addresses.

Wireless communications

Technologies that provide mobile communications for home or office and "in-building wireless" for extended mobility around the work area, campus, or business complex. It is also used to mean "cellular" for in-or out-of-building mobility services.

Wireless modem

A modem that uses radio transmission technology to transmit data between remote

locations. A wireless modem is often used by mobile clients in a location where access to a landline connection is not feasible.

Wireless technology

A communications system in which electro-magnetic waves carry the signal. Examples of wireless equipment include cellular telephones, pagers, the cordless mouse, and wireless transceivers for connecting to the Internet.

Wireline

In the United States, FCC regulations restrict providers of cellular telephone service in any given market to one wireline carrier (the local telephone company) and one nonwireline carrier (any company other than the local telephone company).

WLAN

Wireless LAN. Generic term for the type of network built to the IEEE 802.11b or HiperLAN specifications. Also known as WiFi or 4G mobile.

WLL

Wireless local loop. A mechanism for delivering normal telephony and other services to an end customer.

Worm

A computer program that replicates itself-a form of virus. The Internet worm was probably the most famous-it successfully, and accidentally, duplicated itself across the entire system.

X

X.121

An International Telecommunication Union Telecommunication Standardization Sector (ITU-T) standard that specifies the addressing conventions for any data terminal equipment (DTE) connected to an X.25 network. International numbering plan for packet-switched public data networks.

X.25

A widely used standard protocol suite for packet switching that is also approved by the International Standards Organization. For many years, X.25 was the dominant public data communications packet standard. Many X.25 products are available on the market, and networks have been installed all over the world.

X.400

A store-and-forward message handling system (MHS) standard that allows for the electronic exchange of text, as well as other electronic data such as graphics and fax. It enables suppliers to terwork between different electronic mail systems. X.400 has several protocols, defined to allow the reliable transfer of information between user agents and message transfer agents.

X.500

A directory services standard that permits applications such as electronic mail to access information, which can either be central or distributed.

XDR

eXternal Data Representation. A standard for machine-independent data structures developed by Sun Microsystems. Similar to ASN.1.

xDSL

Generic term for digital subscriber line. Covers asynchronous DSL, ADSL, and so on.

XML

eXtensible Markup Language. Developed by the Worldwide Web Consortium for describing and exchanging information between applications and environments that are based on different technologies and have different semantic representations. It uses the Internet HTTP for data exchange, but has a higher level metadata representation.

X/Open

An industry standards consortium that develops detailed system specifications, drawing on available standards. It has produced standards for a number of distributed computing technologies. X/Open also licenses the UNIX trademark and, thereby, brings focus to its various flavors (e.g., HP-UX, AIX from IBM, Solaris from Sun, and so on).

Team LiB

◀ PREVIOUS

NEXT ▶

Y

Yahoo

Yet another hierarchically organized offering. One of the many search utilities that can be used to trawl and crawl the information held on the WWW.

Team LiB

◀ PREVIOUS

NEXT ▶

Team LiB

◀ PREVIOUS

NEXT ▶

Z

Zip

A compression program, from PKWare, to reduce files to be sent over a network to a more reasonable size. This was originally popularized on MS-DOS but has now spread to other operating systems.

Zone of authority

Term used in the Internet domain name system to refer to the group of names for which a given name server is an authority. Two name servers must supply each zone that has no common point of failure.

Team LiB

◀ PREVIOUS

NEXT ▶

Index

Numbers

- 3G, [150-51, 154](#)
 - active mobile phones, [150](#)
 - IEEE 802.11 marriage, [157](#)
 - mobile phone carriers, [157](#)
 - as shared transmission media, [150](#)
- 10Base-2, [29, 30](#)
- 10 Base-5, [29-30](#)
 - advantages/disadvantages, [29](#)
 - attributes, [30](#)
 - defined, [29](#)
- 10Base-FL, [31](#)
- 10Base-T, [29, 30-31](#)
 - characteristics, [31](#)
 - comparison, [49-50](#)
 - defined, [30](#)
 - levels, [31](#)
- 10-Gigabit Ethernet, [58, 69-79](#)
 - chip interface, [75-76](#)
 - data rate, [77-78](#)
 - defined, [69](#)
 - LAN PHY, [74-75](#)
 - MAC, [76-78](#)
 - MII, [78-79](#)
 - pacing mechanism, [78](#)
 - parallel implementation, [71-73](#)
 - physical layer, [74-75](#)
 - physical layer architecture, [70-76](#)
 - PMD, [73](#)
 - protocol stack, [70](#)
 - serial implementation, [70-71](#)
 - WAN PHY, [74-75](#)
 - See also [Gigabit Ethernet](#)
- 100Base-T, [44, 45-47](#)
 - autonegotiation, [46](#) defined, [44](#)
 - MAC, [46-47](#)
 - MII, [46](#)
 - physical layer, [45-46](#)
 - See also [Fast Ethernet](#)
- 100VG-AnyLAN, [44-45, 47-49](#)
 - comparison, [49-50](#)
 - defined, [44](#)
 - demand priority, [48-49](#)
 - frame support, [47](#)
 - hub layers, [49](#)
 - pros/cons, [48](#)
 - See also [Fast Ethernet](#)
- 1000Base-T, [60](#)
- 1000Base-X, [60](#)

Index

A

Account management, [192](#)

Add-drop multiplexers (ADMS), [205](#)

Address resolution protocol (ARP), [28](#)

 cache, [29](#)

 defined, [208](#)

 reverse, [209](#)

 software, [208](#)

ALOHA

 lessons, [27](#)

 slotted, [26](#)

Application programming interfaces (APIs), [136](#)

Application user interface (AUI), [34](#)

Asymmetrical DSL (ADSL), [119,216](#)

ATM, [106-7,147,148,199](#)

 adaptation layer, [224](#)

 basic operation, [224](#)

 cell structure, [106](#)

 connection orientation, [225](#)

 defined, [106,224](#)

 flexibility, [225](#)

 for multiservice networks, [224](#)

 predictions, [107](#)

 protocol overheads, [199](#)

Attachment unit interface (AUI), [41](#)

Index

B

Backward compatibility, [22](#)

Basic service area (BSA), [93](#)

Basic service set (BSS), [92-93](#)

Bluetooth, [156, 209-10](#)

 defined, [209](#)

 potential, [210](#)

 transceivers, [210](#)

BonDing, [214](#)

Bridges

 defined, [34-35](#)

 self-learning, [35](#)

 as "store-and-forward" devices, [35](#)

 See also [Network components](#)

Broadcast address, [28](#)

Brouters, [35-36](#)

Bus topology, [31-32](#)

Index

C

- Carrierless amplitude modulation (CAP), [217](#)
- Centralized management, [179-80](#)
- CiscoWorks, [175](#)
- Collision domain, [34](#)
- Collision sense multiple access/collision detect. See [CSMA/CD](#)
- Commercial management platforms, 189-91
 - diversity, [189](#)
 - scalability, [189](#)
 - solution illustration, [190](#)
 - state of the art, [189-90](#)
- Common information models (CIM), [136](#)
- Common Management Information Protocol (CMIP), [180](#)
- Communication media, [29-31](#)
 - 10Base-2, [30](#)
 - 10Base-5, [29-30](#)
 - 10Base-FL, [31](#)
 - 10Base-T, [30-31](#)
- Competing technologies, [213-27](#)
 - ATM, [224-25](#)
 - DSL, [215-17](#)
 - FDDI, [225-26](#)
 - frame relay, [222-24](#)
 - GPRS, [217-20](#)
 - HiperLAN, [220-22](#)
 - ISDN, [214-15](#)
 - token ring, [226-27](#)
- Complementary technologies, [203-11](#)
 - Bluetooth, [209-10](#)
 - DWDM, [204-5](#)
 - IP, [207-9](#)
 - SONET/SDH, [205-7](#)
 - summary, [210-11](#)
- Content management, [192](#)
- Copper distributed data interface (CDDI), [226](#)
- CSMA/CD, [44](#)
 - defined, [28](#)
 - Gigabit Ethernet and, [61](#)
- CT2, [155](#)
- Customer management tools, [174-75](#)

Index

D

Decentralized management, [179-80](#)

DECT, [155](#)

Dense wave division multiplexing (DWDM), [204-5](#)
 advances in, [204](#)
 defined, [9](#)
 metropolitan (MDWDM), [113](#)
 technical aspects, [204](#)

DiffServ, [112](#)

Discrete multitone (DMT), [217](#)

Distributed coordinating function (DCF), [93,95](#)
 contention service, [95](#)
 with handshaking, [96](#)
 PCF with, [96](#)

Distributed management, [191-92](#)

DSL, [119-20,215-17](#)
 asymmetrical (ADSL), [119,216](#)
 defined, [215](#)
 high-speed (HDSL), [119,216](#)
 ISDN-based (IDSL), [120,216](#)
 rate-adaptive (RADSL), [120,216](#)
 symmetric (SDSL), [119-20,216](#)
 technologies, [217](#)
 very high-speed (VDSL), [120](#)

DS system, [89-90](#)

Index

E

e-business

- features, [192-93](#)
- management, [192-94](#)
- management support needed for, [193](#)
- separate system implementation, [194](#)

Echo cancellation, [215](#)

EtherCLECs, [149](#)

EtherLoop, [120-25](#)

- in action, [124-25](#)
- advantage, [121](#)
- architecture, [124](#)
- coding, [123](#)
- defined, [120](#)
- Ethernet compliance, [120](#)
- flexible symmetry, [121](#)
- rate adaptation, [121](#)
- routers, [123](#)
- structures and operation, [122-23](#)
- summary, [122](#)

Ethernet

- ALOHA and, [26-27](#)
- alternative, [107-10](#)
- backward compatibility, [22](#)
- bus collisions, [27](#)
- communication media, [29-31](#)
- concepts, [24-36](#)
- CSMA/CD, [28](#)
- design rules, [50-51](#)
- destination address, [25](#)
- Fast, [44-50](#)
- first, clock rate, [22](#)
- generations, compatibility across, [114](#)
- history, [21-24](#)
- IEEE 802.3, [22, 29](#)
- in the last mile, [117-25](#)
- "napkin diagram", [22](#)
- network components, [33-36](#)
- network protocol, [26-29](#)
- packet format, [24-26](#)
- packet-switching technology, [16](#)
- popularity, [55, 197](#)
- repeater placement rule, [50](#)
- shared, [33](#)
- simplicity, [23](#)
- source address, [25](#)
- topology, [31-33](#)
- in the wide area, [104-17](#)
- wireless, [83-101](#) See also [Gigabit Ethernet](#)

Ethernet protocol stack, [40-44](#)

- illustrated, [41](#)
- LLC, [41-42](#)
- MAC, [42-44](#)
- PLS, [41](#)

PMA, [40](#)

Extended service set (ESS), [93-94](#)

eXtensible Markup Language (XML), [162](#)

Team LiB

[◀ PREVIOUS](#) [NEXT ▶](#)

Index

F

Fast Ethernet, [44-50](#)

100Base-T, [44, 45-47](#)

100VG-AnyLAN, [44, 47-49](#)

comparison, [49-50](#)

design rules, [50-51](#)

distance limits, [51](#)

layers, [47](#)

repeater placement rule, [51](#)

transfer rate, [57](#) See also [Ethernet](#)

FDDI, [225-26](#)

CDDI and, [226](#)

defined, [225](#)

synchronous bandwidth allocation (SBA), [225](#) See also [Competing technologies](#)

Fiber Channel, [58, 79-81](#)

advantages/disadvantages, [79](#)

application support, [79](#)

configuration support, [133](#)

defined, [79](#)

FC-0, [80](#)

FC-1, [80](#)

long wavelength lasers, [59](#)

Open Fiber Control (OFC) system, [80](#)

physical layer, [81](#)

running disparity (RD), [81](#)

as SAN transmission technology, [134](#)

structure, [79-80](#)

transfer rates, [79](#) See also [Gigabit Ethernet](#)

Fiber-optic cable, [9](#)

Fiber-to-the-business (FTTB), [198](#)

Fiber-to-the-home (FTTH), [198](#)

Fibre Channel Arbitrated Loop (FCAL), [133](#)

Fibre Channel Industry Association (FCIA), [139](#)

Fibre Channel over IP (FCIP), [140](#)

File transfer protocol (FTP), [15](#)

Forward error correction (FEC), [69](#)

Frame bursting, [63-64](#)

defined, [64](#)

illustrated, [64](#)

Frame check sequence (FCS), [26, 63](#)

Frame relay, [222-24](#)

defined, [222](#)

flexibility, [223](#)

PVCs, [223](#)

SVCs, [222-23](#)

Freedom, [152](#)

Frequency hopping (FH), [86, 89](#)

Fulfillment, assurance, billing (FAB), [171-72](#)

Index

G

General Packet Radio Service (GPRS), [217-20](#)

 components, [217-18](#)

 defined, [217](#)

 device operation, [218](#)

 GGSN, [218, 219](#)

 networks, [150](#)

 SGSN, [218, 219](#)

 tunnel protocol, [219](#)

 versions, [218](#)

Gigabit Ethernet, [57-81](#)

 10-Gigabit, [58, 69-79](#)

 business case, [113-17](#)

 cost advantages, [116](#)

 CSMA/CD and, [61](#)

 customer benefits, [114-15](#)

 enhanced service creation opportunities, [116-17](#)

 frame bursting, [63-64](#)

 growing pains, [112-13](#)

 MAC, [60-65](#)

 MAN built with, [108](#)

 MAN operators in Europe, [109](#)

 from market point of view, [113-14](#)

 modifications for, [58](#)

 operator advantages, [115-16](#)

 PCS, [66](#)

 physical layer, [59-60](#)

 PMA, [66](#)

 PMD, [66](#)

 protocol stack, [65-66](#)

 QoS, [110-12](#)

 signal encoding, [66-69](#)

 slot size, [62, 63](#) See also [Fiber Channel](#)

Gigabit Ethernet Alliance, [108](#)

Index

H

High-performance LAN (HiperLAN), [152,220-22](#)

CL layer, [222](#)

defined, [220](#)

DLC layer, [221-22](#)

PHY layer, [220](#)

type 2, [220](#)

High-speed DSL (HDSL), [119,216](#)

Home phonenumber networking (HPNA), [31](#)

Hubs, [33-34](#)

defined, [33](#)

switching, [51](#) See also [Network components](#)

Index

I-K

IEEE 802 committee, [52-53](#)

IEEE 802.1p/Q, [111](#)

IEEE 802.11, [85,152](#)

3G marriage, [157](#)

802.11a, [86,88](#)

802.11b, [86-87,149-50](#)

coordination functions, [94-96](#)

DS system, [89-90](#)

frequency band, [156](#)

frequency hopping (FH) scheme, [86,89](#)

MAC, [90-98](#)

MAC functions, [96-98](#)

MAC structure, [92-94](#)

open issues, [100](#)

physical layer, [85-90](#)

protocol layers, [87](#)

security, [100](#)

spread-spectrum techniques, [88-89](#)

IEEE 802.3, [22,29,53](#)

802.3ae, [73](#)

802.3z, [58-59](#)

committee, [54](#)

specifications, [54](#)

IEEE .802.17, [199](#)

Information

access, [2](#)

economy, [2-3](#)

sharing problem, [3](#)

volume growth, [4](#)

Integrated services digital network (ISDN), [14,118-19,214-15](#)

defined, [214](#)

echo cancellation, [215](#)

forms, [118,214](#)

inverse multiplexing, [214](#)

motivation behind, [118](#)

Network Termination Unit (NT1), [215](#)

rate adaptation, [214](#)

Integration frameworks, [177](#)

International Standards Organization.

See [ISO model](#)

Internet, [14-16](#)

functioning of, [15](#)

functions, [15](#)

IP, [16](#)

PSTN vs., [14](#)

Internet architecture, [37](#)

application layer, [39](#)

illustrated, [37](#)

Internet layer, [40](#)

network access layer, [40](#)

received packets, [40](#)

transport layer, [39](#)

Internet Protocol (IP), [207-9](#)

defined, [207](#)

packets, [207-8](#)

packet-switching technology, [16](#)

version 4, [207](#)

Inverse multiplexing, [214](#)

ISDN-based DSL (IDSL), [120,216](#)

ISO FCAPS, [167-68](#)

ISO model, [36-39](#)

application layer, [38](#)

data link layer, [38](#)

defined, [36-37](#)

illustrated, [37](#)

network layer, [38](#)

physical layer, [38](#)

presentation layer, [38](#)

session layer, [38](#)

transport layer, [38](#)

Team LiB

◀ PREVIOUS

NEXT ▶

Index

L

LAN PHY, [74-75](#)

Local area networks (LANs), [1](#)

 linking, [8](#)

 opportunities, [18-19](#)

 speed, [6](#)

 virtual (VLAN), [123](#)

 wireless, [84-101](#) See also [Wide area networks \(WANs\)](#)

Local loop access, [12](#)

Logical link control (LLC), [41-42](#)

 classes of service, [41-42](#)

 format, [42](#)

 two-byte address, [42](#)

Long-distance transmission costs, [11](#)

Index

M

MAC, [42-44](#)

100Base-T, [46-47](#)

address, [28](#)

address structure, [43](#)

functions, [42-43](#)

MAC (10-Gigabit Ethernet), [76-78](#)

data rate, [77-78](#)

frame format, [77](#)

full-duplex only, [76-77](#)

pacing mechanism, [78](#)

MAC (Gigabit Ethernet), [60-65](#)

defined, [60](#)

frame bursting, [63-64](#)

full-duplex transmission, [61](#)

half-duplex transmission, [61](#)

modes, [60](#)

MAC (wireless Ethernet), [90-98](#)

association, [97](#)

basic service set (BSS), [92-93](#)

extended service set (ESS), [93-94](#)

fragmentation, [97-98](#)

power management, [96-97](#)

security, [97](#) See also [IEEE 802.11](#); [wireless LANs](#)

Main distribution frame (MDF), [13](#)

Management, [165-95](#)

account, [192](#)

centralized, [179-80](#)

commercial platforms, [189-91](#)

content, [192](#)

decentralized, [179-80](#)

distributed, [191-92](#)

e-business, [192-94](#)

geographically dispersed networks, [166](#)

models, [166-73](#)

order, [192](#)

platform implementation, [186-89](#)

SAN, [131, 134-35, 136-38](#)

supply chain, [192](#)

tools, [173-78](#)

VPN, [147](#)

Management design, [178-86](#)

component systems, [182-83](#)

data integrity, [185-86](#)

operational issues, [179, 184-85](#)

process walkthrough, [183-84](#)

protocol standards, [180](#)

steps, [182-86](#)

systematic approach, [181-82](#)

testability and, [180](#)

Marketplace, changing, [143-63](#)

MB810, [68](#)

Media independent interface (MII)

10-Gigabit Ethernet, [78-79](#)

100Base-T, [46](#)

Metcalf's Law, [23](#)

Metro area networks (MANs), [24](#)

with Gigabit Ethernet, [108](#)

Gigabit Ethernet, operators in Europe, [109](#)

Metropolitan dense wave division

multiplexing (MDWDM), [113](#)

Microsoft Passport, [160](#)

MPLS, [111-12](#)

Team LiB

[◀ PREVIOUS](#) [NEXT ▶](#)

Index

N

Net allocation vector (NAV), [96](#)

NetView, [174](#)

Network attached storage (NAS), [129-31](#)

 concept, [130](#)

 configuration, [130](#)

 limitations, [131](#)

 servers, [131](#)

Network components, [33-36](#)

 bridges, [34-35](#)

 hub, [33-34](#)

NICs, [36](#)

 repeaters, [34](#)

 routers, [35-36](#)

 switches, [35](#)

 transceivers, [34](#)

Network interface cards (NICs), [36](#)

Network management tools, [173-78](#)

 customer management, [174-75](#)

 integrating, for management solution, [177-78](#)

 integration frameworks, [177](#)

NetView, [174](#)

 process-specific, [174](#)

 service management, [175](#)

 support, [175-76](#)

 system management, [177](#)

 total management platforms, [173-74](#) See also [Management](#)

Network protocol, [26-29](#)

Networks

 access link utilization, [189](#)

 basic extended, [186](#)

 cooperative communication over, [158](#)

 infrastructure investments, [7](#)

 investment in, [6](#)

 operational support systems, [153](#)

 providers, [5-6](#)

 statistics, [188-89](#)

 technology, [6](#)

 with wide area and local elements, [7](#) See also [specific networks](#)

Network storage

 growth, [128-29](#)

 options, [129-31](#) See also [Storage area networks \(SANs\)](#)

Index

O

Optical cable

 capacity, [10-11](#)

 construction cost, [11](#)

Optical network providers, [8-9](#)

Optical switching, [10](#)

Order management, [192](#)

Organization, this book, [16-17](#)

Orthogonal frequency-division multiplexing (OFDM), [88](#)

Index

P

Packet over SONET (POS), [115](#)

Packets

- drop rate, [189](#)
- errored rate, [189](#)
- format, [24-26](#)
- illustrated, [24](#)
- IP, [207-8](#)
- preamble, [25](#)

PAM-5 encoding, [68](#)

Personal Handyphone System (PHS), [155](#)

Personalization, [192](#)

Physical address, [28](#)

Physical layer

- 10-Gigabit Ethernet, [74-75](#)
- 100Base-T, [45-46](#)
- Fiber Channel, [81](#)
- IEEE 802.11, [85-90](#)
- ISO model, [38](#)

Physical layer (Gigabit Ethernet), [59-60](#)

- 1000Base-T, [60](#)
- 1000Base-X, [60](#)
- defined, [59](#)

Physical medium attachment (PMA), [40-41,66](#)

Physical signaling sublayer (PLS), [41](#)

Point coordination function (PCF), [93](#)

- DCF with, [96](#)

point coordinator, [96](#)

Points of presence (PoPs), [9](#)

Point-to-Point Protocol (PPP), [199,214](#)

Process-specific tools, [174](#)

Protocol stack

- 10-Gigabit Ethernet, [70](#)
- Ethernet, [40-44](#)
- Gigabit Ethernet, [65-66](#)

Public data network (PDH), [14](#)

Public switched telephone network (PSTN), [13-14](#)

- access layer, [13](#)
- core transmission layer, [13-14](#)
- Internet vs., [14](#)
- switch layer, [14](#)

Team LiB

◀ PREVIOUS

NEXT ▶

Index

Q

Quadrature amplitude modulation (QAM), [123](#)

Quadrature phase shift keying (QPSK), [123](#)

Quality of service (QoS), [110-12](#)

Team LiB

◀ PREVIOUS

NEXT ▶

Index

R

Rate adaptation, [214](#)

Rate-adaptive DSL (RADSL), [120,216](#)

Remote monitoring (RMON), [191](#)

Repeaters

 "5-4-3" placement rule, [50](#)

 defined, [34](#) See also [Network components](#)

Reverse ARP (RARP), [209](#)

Router-free memory, [188](#)

Routers

 defined, [35-36](#)

 EtherLoop-based, [123](#)

 peak processor utilization, [188](#) See also [Network components](#)

RPR, [111,200](#)

Running disparity (RD), [81](#)

Index

S

- Scrambling, [67-68](#)
- SCSI interface, [129-30](#)
 - defined, [130](#)
 - I/O operation, [139](#)
 - throughput rate support, [130](#)
- Security
 - IEEE 802.11, [100](#)
- SANs, [135](#)
 - wireless LANs, [91](#)
- Service level agreements (SLAs), [110](#)
- Service management tools, [175](#)
- Shared Ethernet, [33](#)
- Short Message Service (SMS), [150](#)
- Signal encoding, [66-69](#)
 - 8B/10B, [67](#)
 - 16B/18B, [69](#)
 - forward error correction (FEC), [69](#)
 - MB810, [68](#)
 - PAM-5, [68](#)
 - scrambling, [67-68](#) See also [Gigabit Ethernet](#)
- Simple management protocol (SNMP), [136](#)
 - defined, [187](#)
 - management software, [187](#)
- Simple Object Access Protocol (SOAP), [162](#)
- Slotted ALOHA, [26](#)
- SONET/SDH, [13, 205-7](#)
 - benefits, [105, 207](#)
 - bit-interleaved parity (BIP), [207](#)
 - data rates, [14](#)
 - defined, [104-5, 205](#)
 - deployment, [207](#)
 - limitations, [106](#)
 - overhead, [206](#)
 - PPP and, [199](#)
 - standard, [206](#)
- Spread-spectrum techniques, [88-89](#)
- Star topology, [32](#)
- Statistics, [188-89](#)
- Storage area networks (SANs), [127-41](#)
 - advantages, [132](#)
 - availability, [135](#)
 - business case, [136](#)
 - concept, [129](#)
 - elements, [132-33](#)
 - environment, [131](#)
 - evaluating, [134-36](#)
 - evolution, [138-40](#)
 - extending, [133-34](#)
 - Fiber Channel technology, [134](#)

file administration and, [135](#)
illustrated, [132](#)
interconnection technology, [132](#)
management, [131](#), [134-35](#), [136-38](#)
as proven solution, [134](#)
questions, [134-35](#)
scalability, [134](#)
scope, [137](#)
security, [135](#)
in storage strategy, [140-41](#)
summary, [141](#)

Storage strategy, [140-41](#)

Sub-Network Access Protocol (SNAP), [42](#)

Supply chain management, [192](#)

Support tools, [175-76](#)

Switches, [35](#)

Symmetric DSL (SDSL), [119-20](#), [216](#)

Synchronous digital hierarchy (SDH) See [SONET/SDH](#)

Synchronous optical network (SONET). See [SONET/SDH](#)

System management tools, [177](#)

Team LiB

◀ PREVIOUS

NEXT ▶

Index

T

Telecommunications operations map (TOM), [169,170-73](#)

defined, [169](#)

FAB, [171-72](#)

framework, [170](#)

illustrated, [170](#)

location of major products on, [176](#)

Telemanagement Forum (TMF), [168-73](#)

defined, [168-69](#)

management hierarchy, [168](#)

TOM, [169,170-73](#)

Testability, [180](#)

Thick Ethernet. See [10 Base-5](#)

Time division multiplexing (TDM), [5](#)

Token ring, [226-27](#)

as deterministic architecture, [227](#)

standardization, [226](#) See also [Competing technologies](#)

Topologies, [31-33](#)

bus, [31-32](#)

star, [32](#)

tree, [32](#)

Total area networks, [7,12-13,103-25](#)

defined, [103](#)

design issues, [194-95](#)

economics and, [12](#)

last mile, [117-25](#)

wide area, [104-17](#)

Transatlantic transmission capacity, [10](#)

Transceivers, [34](#)

Tree topology, [32](#)

Team LiB

◀ PREVIOUS

NEXT ▶

Index

U

Universal Description, Discovery, and Integration (UDDI), [162](#)

Universal mobile telecommunication service (UMTS), [149](#)

Unshielded twisted pair (UTP), [22](#)

Team LiB

◀ PREVIOUS

NEXT ▶

Index

V

Very high-speed DSL (VDSL), [120,216](#)

Virtual LAN (VLAN), [123](#)

Voice over IP (VoIP), [112](#)

VPNs

connectivity, [145-46](#)

cost, [146](#)

defined, [145](#)

flexibility, [146](#)

management, [147](#)

QoS, [146-47](#)

scalability, [146](#)

Index

W

Wavelength division multiplexing (WDM), [72](#)

Web services, [157-63](#)

 application, [161](#)

 benefits, [158](#)

 categories, [160-61](#)

 configuration information, [161](#)

 control of, [159](#)

 definitions, [158-59](#)

 federation, [159-61](#)

 interface operations, [161-62](#)

 Microsoft-proposed, [160](#)

 Passport, [160](#)

 personalization information, [160](#)

 software components, [161-63](#)

 summary, [163](#)

 XML, [162](#)

Web Services Description Language (WSDL), [162](#)

Wide area networks (WANs), [1](#)

 bottleneck, bypassing, [8](#)

 cost drop, [12](#)

 examples, [6](#) See also [Local area networks \(LANs\)](#)

Wireless access protocol (WAP) service, [144](#)

Wireless Ethernet, [83-101](#)

Wireless Ethernet Compatibility

 Alliance (WECA), [87](#)

Wireless LANs, [84-101](#)

 ad hoc network, [99-100](#)

 broadcast capability, [91](#)

 collocation, [92](#)

 configurations, [98-100](#)

 connectivity, [91](#)

 defined, [84](#)

 delay, [91](#)

 deployment optimization, [91](#)

 deployment options, [99](#)

 fairness, [91](#)

 features required for, [90-92](#)

 MAC, [90-98](#)

 manageability, [92](#)

 nomadic access, [98-99](#)

 open issues, [100](#)

 as part of larger network, [94](#)

 physical layer, [85-90](#)

 power consumption, [91](#)

 priority marking, [91](#)

 roaming, [91, 155](#)

 robustness, [91](#)

 security, [91](#)

 single/multiple cells, [98](#)

 summary, [100-101](#)

 throughput, [90](#)

 transparency, [90](#)

use of, [150](#)

Team LiB

◀ PREVIOUS

NEXT ▶

Team LiB

◀ PREVIOUS

NEXT ▶

Index

X-Z

XAUI interface, [75-76](#)

XGMII interface, [75](#)

Team LiB

◀ PREVIOUS

NEXT ▶

List of Figures

Chapter 1: The Quiet Revolution

[Figure 1.1:](#) (a) The volume growth of information. (b) The amount of information that resides in software code.

[Figure 1.2:](#) A typical network with wide area and local elements.

[Figure 1.3:](#) Avoiding the WAN bottleneck.

[Figure 1.4:](#) Transatlantic transmission capacity.

[Figure 1.5:](#) The falling cost of long-distance transmission.

Chapter 2: Ethernet-The Story So Far

[Figure 2.1:](#) The original Ethernet "napkin diagram."

[Figure 2.2:](#) The basic Ethernet packet.

[Figure 2.3:](#) Collisions on an Ethernet bus.

[Figure 2.4:](#) A typical Ethernet system.

[Figure 2.5:](#) (a) The ISO layered models of communications, and (b) the Internet layered models of communications.

[Figure 2.6:](#) The Ethernet protocol stack.

[Figure 2.7:](#) The LLC and SNAP formats.

[Figure 2.8:](#) The structure of the MAC address.

[Figure 2.9:](#) The Fast Ethernet layers.

[Figure 2.10:](#) The IEEE view of the Ethernet layers.

Chapter 3: Gigabit Ethernets

[Figure 3.1:](#) Collisions on a link.

[Figure 3.2:](#) Slot sizes.

[Figure 3.3:](#) Carrier extension.

[Figure 3.4:](#) Frame bursting.

[Figure 3.5:](#) Gigabit Ethernet protocol stack.

[Figure 3.6:](#) 10-Gigabit Ethernet protocol stack.

[Figure 3.7:](#) Serial physical layer implementation.

[Figure 3.8:](#) Parallel physical layer implementation.

[Figure 3.9:](#) Where the XAUI fits.

Chapter 4: Wireless Ethernet

[Figure 4.1:](#) The IEEE 802.11 protocol layers.

[Figure 4.2:](#) Options for connectivity.

[Figure 4.3:](#) A wireless LAN as part of a larger network.

[Figure 4.4:](#) Two basic deployment options.

Chapter 5: Total Area Networks

[Figure 5.1:](#) A MAN built with Gigabit Ethernet.

[Figure 5.2:](#) The compatibility across the Ethernet generations.

Chapter 6: Storage Area Networks

[Figure 6.1:](#) The configuration for NAS.

[Figure 6.2:](#) A SAN.

Chapter 7: A Changing Marketplace

[Figure 7.1:](#) The range of operational support systems that need to be provided around a network.

[Figure 7.2:](#) Roaming across a wireless LAN network.

Chapter 8: Managing Total Area Ethernetworks

[Figure 8.1:](#) The TMF management hierarchy for communication networks.

[Figure 8.2:](#) Key concerns at each level of the pyramid.

[Figure 8.3:](#) The TOM.

[Figure 8.4:](#) FAB in the TOM model.

[Figure 8.5:](#) Location of major products on the TOM.

[Figure 8.6:](#) Typical systems architecture.

[Figure 8.7:](#) Basic extended network.

[Figure 8.8:](#) A large-scale management solution.

[Figure 8.9:](#) The management support needed for effective e-business.

List of Tables

Chapter 1: The Quiet Revolution

[Table 1.1:](#) Data Rates for SONET/SDH

[Table 1.2:](#) A Guide to This Book

Chapter 2: Ethernet-The Story So Far

[Table 2.1:](#) Planning Rules for Different Types of Ethernet

[Table 2.2:](#) Distance Limits for Fast Ethernet

[Table 2.3:](#) Main Specifications in the IEEE 802.3 Family

Chapter 3: Gigabit Ethernets

[Table 3.1:](#) Maximum Reach for Various Media Types

[Table 3.2:](#) Target Distances Specified by the IEEE 802.3ae Task Force

Chapter 5: Total Area Networks

[Table 5.1:](#) Capacity and Traffic Offered by Current Transmission Technology

[Table 5.2:](#) Gigabit Ethernet MAN Operators in Europe

[Table 5.3:](#) Cost Advantages of Gigabit Ethernet

Chapter 7: A Changing Marketplace

[Table 7.1:](#) Wireless Technology Ranges When Deployed in Urban Areas, Providing Indoor and Outdoor Coverage

[Table 7.2:](#) The Range of Web Services Proposed by Microsoft

Chapter 8: Managing Total Area Ethernetworks

[Table 8.1:](#) Pros and Cons of Different Management System Integration Strategies

Appendix A: Complementary Technologies

[Table A.1:](#) The SDH/SONET Rates

[Table A.2:](#) Fields in the IPv4Packet

[Table A.3:](#) Bluetooth Frequency Allocation

Appendix B: Competing Technologies

[Table B.1:](#) HiperLAN PHY Models

[Table B.2:](#) Sublayers of the HiperLAN DLC