PART 3

NETWORK TECHNOLOGIES

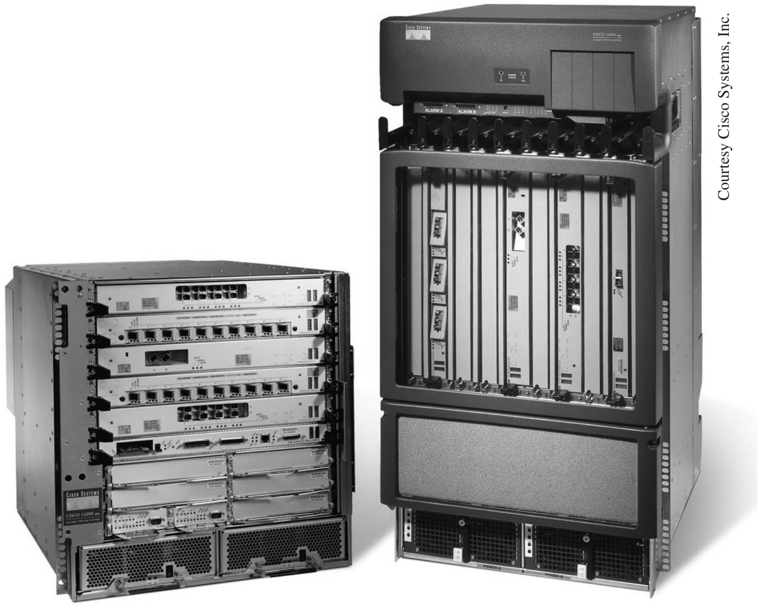

Courtesy Cisco Systems, Inc.

A Cisco switch and router

CHAPTER 6

LOCAL AREA NETWORKS

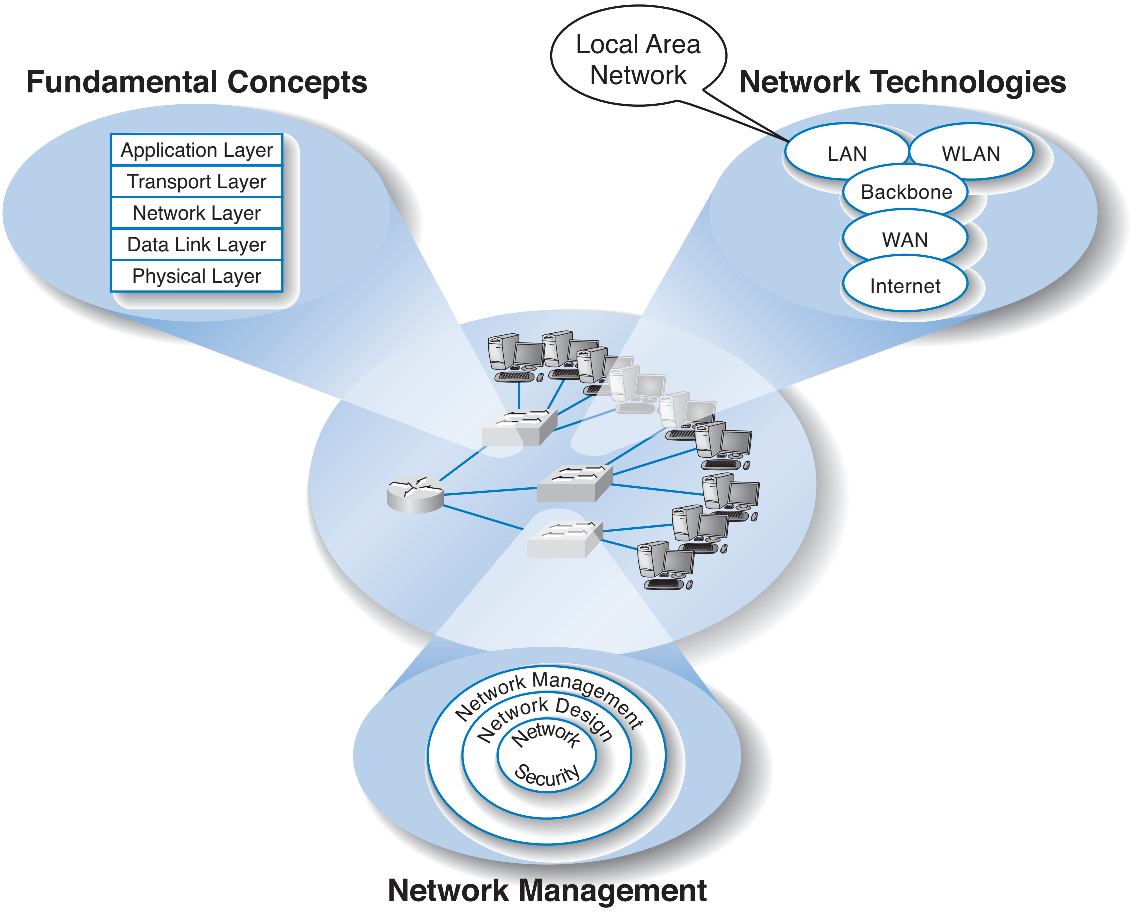

The Three Faces of Networking

THE PRECEDING chapters provided a fundamental understanding of the five basic layers in a typical network. This chapter draws together these concepts to describe a basic LAN. We first summarize the major components of a LAN and then describe the two most commonly used LAN technologies: traditional Ethernet and switched Ethernet. The chapter ends with a discussion of how to design LANs and how to improve LAN performance. In this chapter, we focus only on the basics of LANs; the next chapter describes how LANs and BNs are used together.

OBJECTIVES

- Be aware of the roles of LANs in organizations
- Understand the major components of LANs
- Understand traditional Ethernet LANs
- Understand switched Ethernet LANs
- Understand the best practice recommendations for LAN design
- Be familiar with how to improve LAN performance

CHAPTER OUTLINE

INTRODUCTION

Most large organizations have numerous LANs connected by backbone networks. These LANs also provide access to a variety of servers, mainframe computers, and the Internet. In this chapter, we discuss the fundamental components of a LAN, along with two technologies commonly used in LANs—traditional Ethernet (IEEE 802.3), and switched Ethernet. There used to be many different types of LAN technologies, such as Token Ring, but gradually the world has changed so that Ethernet dominates. Today, very few organizations consider any LAN technology other than Ethernet. Together, traditional Ethernet and its switched and wireless cousins account for almost all LANs installed today.

Why Use a LAN?

There are two basic reasons for developing a LAN: information sharing and resource sharing. *Information sharing* refers to having users access the same data files, exchange information via e-mail, or use the Internet. For example, a single purchase order database might be maintained so all users can access its contents over the LAN. (Many information-sharing applications were described in Chapter 2.) The main benefit of information sharing is improved decision making, which makes it generally more important than resource sharing.

Resource sharing refers to one computer sharing a hardware device (e.g., printer, an Internet connection) or software package with other computers on the network to save costs. For example, suppose we have 30 computers on a LAN, each of which needs access to a word processing package. One option is to purchase 30 copies of the software and install one on each computer. This would use disk space on each computer and require a significant amount of staff time to perform the installation and maintain the software, particularly if the package were updated regularly.

An alternative is to install the software on the network for all to use. This would eliminate the need to keep a copy on every computer and would free up disk space. It would also simplify software maintenance because any software upgrades would be installed once on the network server; staff members would no longer have to upgrade all computers.

In most cases, not all users would need to access the word processing package simultaneously. Therefore, rather than purchasing a license for each computer in the network, you could instead purchase 10 licenses, presuming that only 10 users would simultaneously use the software. Of course, the temptation is to purchase only one copy of the software and permit everyone to use it simultaneously. The cost savings would be significant, but this is illegal. Virtually all software licenses require one copy to be purchased for each simultaneous user. Most companies and all government agencies have policies forbidding the violation of software licenses, and many fire employees who knowingly violate them.

One approach to controlling the number of copies of a particular software package is to use *LAN metering software* that prohibits using more copies of a package than there are installed licenses. Many software packages now come in LAN versions that do this automatically, and a number of third-party packages are also available.

Nonetheless, the *Software Publishers Association (SPA)* in Washington, D.C., estimates that about 40 percent of all the software in the world is used illegally—an annual total of more than $13 billion. North America has the lowest rate of software piracy (28 percent). Although piracy has been on the decline, it still exceeds 75 percent in many parts of the world, with the exception of western Europe (43 percent), Australia (32 percent), New Zealand (35 percent), and Japan (41 percent).

The SPA has recently undertaken an aggressive *software audit* program to check the number of illegal software copies on LANs. Whistleblowers receive rewards from SPA, and the violating organizations and employees are brought to court. SPA will work with companies that voluntarily submit to an audit, and it offers an audit kit that scrutinizes networks in search of software sold by SPA members (see http://www.spa.org).

Dedicated-Server versus Peer-to-Peer LANs

One common way to categorize LANs is by whether they have a dedicated server or whether they operate as a peer-to-peer LAN without a dedicated server. This chapter focuses primarily on dedicated-server LANs because they account for more than 90 percent of all installed LANs, although many of the issues are also common in peer-to-peer networks.

Dedicated Server Networks As the name suggests, a *dedicated-server LAN* has one or more computers that are permanently assigned as network servers. These servers enable users to share files and often are also used to share printers. A dedicated-server LAN can connect with almost any other network, can handle very large files and databases, and uses sophisticated LAN software. Moreover, high-end dedicated-server LANs can be easily interconnected to form enterprisewide networks or, in some cases, can replace a host mainframe computer. Generally speaking, the dedicated servers are powerful microcomputers or minicomputers. Sometimes servers are organized into a large set of servers on one part of the network called a cluster or *server farm.* Server farms can range from tens to hundreds of servers.

In a dedicated-server LAN, the server's usual operating system (e.g., Windows) is replaced by a network operating system (e.g., Linux, Novell Server, Windows Server). Special-purpose network communication software is also installed on each client computer and is the link between the client computer's operating system and the network operating system on the server. This set of communication software provides the data link layer and network layer protocols that allow data transmissions to take place. Three software components must work together and in conjunction with the network hardware to enable communications: the network operating system in the dedicated server, the network communication software on the client, and the application software that runs on the server and client computers.

A LAN can have many different types of dedicated servers, such as mail servers, database servers, and Web servers, as discussed in Chapter 2. Three other common types are file servers, print servers, and remote-access servers (RASs).

File servers allow many users to share the same set of files on a common, shared disk drive. The hard disk volume can be of any size, limited only by the size of the disk storage itself. Files on the shared disk drive can be made freely available to all network users, shared only among authorized users, or restricted to only one user.

Print servers handle print requests on the LAN. By offloading the management of printing from the main LAN file server or database server, print servers help reduce the load on them and increase network efficiency. Print servers have traditionally been separate computers, but many vendors now sell "black boxes" that perform all the functions of a print server at much less than the cost of a stand-alone computer.

Remote-access servers (RASs) enable users to dial into and out of the LAN by telephone. A RAS lets users dial into the LAN and perform all the same functions as though they were physically connected to the LAN itself. RASs are best for applications that move only small amounts of information and do not require high speed beyond the limited capabilities of regular voice-grade telephone lines. (LANs typically provide data transmission rates of between 10 and 100 Mbps whereas telephone lines typically provide between only 28.8 and 128 Kbps.)

Peer-to-Peer Networks *Peer-to-peer networks* do not require a dedicated server. All computers run network software that enables them to function both as clients and as servers. Authorized users can connect to any computer in the LAN that permits access and use its hard drives and printer as though it were physically attached to their own computers. Peer-to-peer networks often are slower than dedicated server networks because if you access a computer that is also being used by its owner, it slows down both the owner and the network.

In general, peer-to-peer LANs have less capability, support a more limited number of computers, provide less sophisticated software, and can prove more difficult to manage than dedicated-server LANs. However, they are cheaper both in hardware and software. Peer-to-peer LANs are most appropriate for sharing resources in small LANs. We should note that peer-to-peer has become popular for application layer software file sharing on the Internet. This is conceptually similar to peer-to-peer LANs, but quite different in practice.

A Day in the Life: LAN Administrator

Most days start the same way. The LAN administrator arrives early in the morning before most people who use the LAN. The first hour is spent checking for problems. All the network hardware and servers in the server room receive routine diagnostics. All the logs for the previous day are examined to find problems. If problems are found (e.g., a crashed hard disk) the next few hours are spent fixing them. Next, the daily backups are done. This usually takes only a few minutes, but sometimes a problem occurs and it takes an hour.

The next step is to see if there are any other activities that need to be performed to maintain the network. This involves checking e-mail for security alerts (e.g., Windows updates, anti-virus updates). If critical updates are needed, they are done immediately. There are usually e-mails from several users that need to be contacted, either problems with the LAN, or requests for new hardware or software to be installed. These new activities are prioritized into the work queue.

And then the real work begins. Work activities include tasks such as planning for the next roll out of software upgrades. This involves investigating the new software offerings, identifying what hardware platforms are required to run them, and determining which users should receive the upgrades. It also means planning for and installing new servers or network hardware such as firewalls.

Of course, some days can be more exciting than others. When a new virus hits, everyone is involved in cleaning up the compromised computers and installing security patches on the other computers. Sometimes virus attacks can be fun when you see that your security settings work and beat the virus.

With thanks to Steve Bushert

LAN COMPONENTS

There are six components in a traditional LAN (Figure 6.1). The first two are the client computer and the server (but see the section above on peer-to-peer networks). Clients and servers have been discussed in Chapter 2, so they will not be discussed further here. The other components are network interface cards (NICs), network cables, hubs, and the network operating system. In recent years, a new form of LAN called switched Ethernet has become popular that uses switches instead of hubs; the role of switches is discussed in a later section.

Network Interface Cards

The *network interface card (NIC)* is used to connect the computer to the network cable and is one part of the physical layer connection among the computers in the network. Most computers come with a NIC built in, but sometimes a separate NIC must be installed. Some laptops have a special port that enables network cards to be installed without physically opening them (i.e., PCMCIA [Personal Computer Memory Card International Association] slot).

Network Cables

Each computer must be physically connected by network cable to the other computers in the network. Just as highways carry all kinds of traffic, the perfect cabling system also

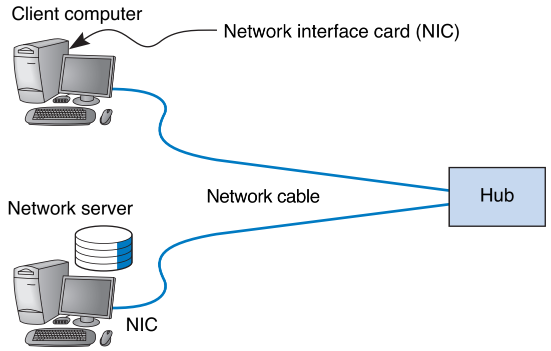

FIGURE 6.1 Local area network components.

should be able to carry all kinds of electronic transmissions within the building. But in practice, it isn't that simple. The selection of a LAN can be influenced greatly by the type of cable that already exists in the building where the LAN is to be installed.

Most LANs are built with *unshielded twisted-pair (UTP)* wires, *shielded twisted-pair (STP), or fiber-optic cable* (although fiber-optic cable is far more commonly used in BNs, which are discussed in the next chapter). Wireless LANs run on infrared or radio frequencies, eliminating the need for cables. (Common cable standards are discussed on the next page. We should add that these cable standards specify the minimum quality cable required; it is possible, for example, to use category 5 UTP wire for a 10Base-T Ethernet.)

Many LANs use a combination of STP and UTP wire. Although initially it appeared that twisted-pair would not be able to meet long-term capacity and distance requirements, today UTP is one of the leading LAN cabling technologies. Its low cost and the availability of shielded wiring make it very useful. STP is only used in special areas that produce electrical interference, such as factories near heavy machinery or hospitals near MRI scanners.

Fiber-optic cable is even thinner than UTP wire and therefore takes far less space when cabled throughout a building. It also is much lighter, weighing less than 10 pounds per 1,000 feet. Because of its high capacity, fiber-optic cabling is perfect for BNs, although it is beginning to be used in LANs.

Network Hubs

Network *hubs* serve two purposes. First, they provide an easy way to connect network cables. A hub can be thought of as a junction box, permitting new computers to be connected to the network as easily as plugging a power cord into an electrical socket (Figure 6.2). Each connection point where a cable can be plugged in is called a *port.* Each port has a unique number.

Simple hubs are commonly available in 4-, 8-, 16-, and 24-port sizes, meaning that they provide anywhere between 4 and 24 ports into which network cables can be plugged. When no cables are plugged in, the signal bypasses the unused port. When a cable is plugged into a port, the signal travels down the cable as though it were directly connected

TECHNICAL FOCUS

6-1 Commonly Used Network Cable Standards

Name	Type	Maximum Data Rate (Mbps)	Often Used By	Cost[1] ($/foot)
Category 1[2]	UTP	1	Modem	.04
Category 2	UTP	4	Token Ring-4[3]	.35
Category 3	UTP	10	10Base-T Ethernet	.06
Category 4	STP	16	Token Ring-16[3]	.60
Category 5	UTP	100	100Base-T Ethernet	.07
Category 5	STP	100	100Base-T Ethernet	.18
Category 5e[4]	UTP	100	1,000Base-T Ethernet	.10
Category 6	UTP	250	1,000Base-T Ethernet	.15
Category 7[5]	STP	600	1,000Base-T Ethernet	.25
X3T9.5	Fiber	100	FDDI[6]	.25

Notes

1. These costs are approximate costs for cable only (no connectors). They often change but will give you a sense of the relative differences in costs among the different options.
2. Category 1 is standard voice-grade twisted-pair wires but it can also be used to support low-speed analog data transmission.
3. Token ring is an old local area network technology seldom used today.
4. Category 5e is an improved version of category 5 that has better insulation and a center plastic pipe inside the cable to keep the individual wires in place and reduce noise from cross-talk, so that it is better suited to 1000Base-T.
5. The standards for category 7 have not been finalized.
6. FDDI (fiber distributed data interface) is a backbone technology discussed in Chapter 8.

to the cables attached to the hub. Some hubs also enable different types of cables to be connected and perform the necessary conversions (e.g., twisted-pair wire to coaxial cable, coaxial cable to fiber-optic cable).

Second, hubs can act as repeaters or amplifiers. Signals can travel only so far in a network cable before they attenuate and can no longer be recognized. (Attenuation was discussed in Chapter 4.) All LAN cables are rated for the maximum distance they can be

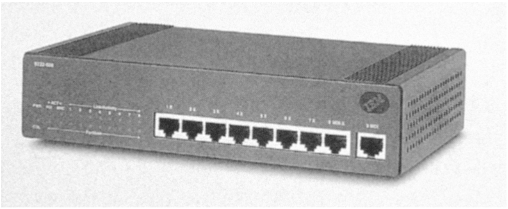

Courtesy Cisco Systems, Inc.

FIGURE 6.2 Network hub.

used (typically 100 meters for twisted-pair wire, and several kilometers for fiber-optic cable).

In the early days of LANs, it was common practice to install network cable wherever it was convenient. Little long-term planning was done. Hubs were placed at random intervals to meet the needs of the few users, and cable was laid where it was convenient. The exact placement of the cables and hubs was often not documented, making future expansion more difficult—you had to find the cable and a hub before you could add a new user.

With today's explosion in LAN use, it is critical to plan for the effective installation and use of LAN *cabling*. The cheapest point at which to install network cable is during the construction of the building; adding cable to an existing building can cost significantly more. Indeed, the costs to install cable (i.e., paying those doing the installation and additional construction) are usually substantially more than the cost of the cable itself, making it expensive to reinstall the cable if the cable plan does not meet the organization's needs.

Most buildings under construction today have a separate LAN *cable plan*, as they have plans for telephone cables and electrical cables. The same is true for older buildings in which new LAN cabling is being installed. Most cable plans are similar in style to electrical and telephone plans. Each floor has a telecommunications wiring closet that contains one or more network hubs. Cables are run from each room on the floor to this wiring closet. It is common to install 20 to 50 percent more cables than you actually need, to make future expansion simple. Any reconfiguration or expansion can be done easily by adding a network hub and connecting the unused cables in the wiring closet. This saves the difficulty and expense of installing new cables.

MANAGEMENT FOCUS

6-1 CABLE PROBLEMS AT THE UNIVERSITY OF GEORGIA

Like many organizations, the Terry College of Business at the University of Georgia is headquartered in a building built before the computer age. When local area network cabling was first installed in the early 1980s, no one foresaw the rapid expansion that was to come. Cables and hubs were installed piecemeal to support the needs of the handful of early users.

The network eventually grew far beyond the number of users it was designed to support. The network cable gradually became a complex, confusing, and inefficient mess. There was no logical pattern for the cables, and there was no network cable plan. Worse still, no one knew where all the cables and hubs were physically located. Before a new user was added, a network technician had to open up a ceiling and crawl around to find a hub. Hopefully, the hub had an unused port to connect the new user, or else the technician would have to find another hub with an empty port.

To complicate matters even more, asbestos was discovered. Now network technicians could not open the ceiling and work on the cable unless asbestos precautions were taken. This meant calling in the university's asbestos team and sealing off nearby offices. Installing a new user to the network (or fixing a network cable problem) now took 2 days and cost $2,000.

The solution was obvious. The university spent $400,000 to install new category 5 twisted-pair cable to every office and to install a new high-speed fiber-optic backbone network between network segments.

MANAGEMENT FOCUS

6-2 Managing Network Cabling

You must consider a number of items when installing cables or when performing cable maintenance. You should:

- Perform a physical inventory of any existing cabling systems and document those findings in the network cable plan.
- Properly maintain the network cable plan. Always update cable documentation immediately on installing or removing a cable or hub. Insist that any cabling contractor provide "as-built" plans that document where the cabling was actually placed, in case of minor differences from the construction plan.
- Establish a long-term plan for the evolution of the current cabling system to whatever cabling system will be in place in the future.
- Obtain a copy of the local city fire codes and follow them. For example, cables used in airways without conduit need to be plenum-certified (i.e., covered with a fire-retardant jacket).
- Conceal all cables as much as possible to protect them from damage and for security reasons.
- Properly number and mark both ends of all cable installations as you install them. If a contractor installs cabling, always make a complete inspection to ensure that all cables are labeled.

Network Operating Systems

The *network operating system (NOS)* is the software that controls the network. Every NOS provides two sets of software: one that runs on the *network server(s)* and one that runs on the network client(s). The server version of the NOS provides the software that performs the functions associated with the data link, network, and application layers and usually the computer's own operating system. The client version of the NOS provides the software that performs the functions associated with the data link and the network layers and must interact with the application software and the computer's own operating system. Most NOSs provide different versions of their client software that run on different types of computers, so that Windows computers, for example, can function on the same network as Apples. In most cases (e.g., Windows, Linux), the client NOS software is included with the operating system itself.

NOS Server Software

The NOS server software enables the file server, print server, or database server to operate. In addition to handling all the required network functions, it acts as the application software by executing the requests sent to it by the clients (e.g., copying a file from its hard disk and transferring it to the client, printing a file on the printer, executing a database request, and sending the result to the client). NOS server software replaces the normal operating system on the server. By replacing the existing operating system, it provides better performance and faster response time because a NOS is optimized for its limited range of operations. Figure 6.3 summarizes several common NOs.

Microsoft Windows Server

One of the most popular NOS is Windows Server, developed by Microsoft Corporation. It provides good file services and adequate print services, as well as an excellent development environment for application services. Windows Server is very similar to the Windows client operating system, so it is straightforward to learn. It works well with Windows client computers, but requires additional software (and effort) to support Apple and Linux clients.

Linux

Linux is an open source operating system first developed by Linus Torvalds at the University of Helsinki. It is the microcomputer version of UNIX, a popular mainframe operating system. Linux provides excellent file, print, and application services. It is more secure than Windows Server, given its origins as a highly secure mainframe operating system and because it is open source. It has a command driven interface (in contrast to Windows' graphical user interface), so it is harder to learn. It works well with Windows, Apple, and Linux client computers.

Novell Server

Novell was the original and most popular NOS but its influence has declined as Windows Server has improved. It provides excellent file, print, and directory services, but has a limited environment for developing application services. It is arguably more secure than Windows Server, being the target of far fewer viruses and attacks. Novell supports a wide variety of client computers including Windows, Apple, and Linux.

Apple Mac Operating System

The Apple Mac OS is a version of UNIX, integrated with the Apple graphical user interface to make it easy to use. It provides good file and print services, with some ability for application development. It is more secure than Windows Server given its origins as a highly secure mainframe operating system. It works well with Apple client computers, but requires additional software (and effort) to support Windows and Linux clients.

FIGURE 6.3 Several common network operating systems.

NOS Client Software The NOS software running at the client computers provides the data link layer and network layer. To work effectively with the application software, the NOS must also work together with the client's own operating system. Most operating systems today are designed with networking in mind. For example, Windows provides built-in software that will enable it to act as a client computer with a Novell NetWare server or a Windows Server.

One of the most important functions of a NOS is a *directory service.* Directory services provide information about resources on the network that are available to the users, such as shared printers, shared file servers, and application software. A common example of directory services is Microsoft's *Active Directory Service (ADS).*

TECHNICAL FOCUS

6-2 Storage Area Networks and Network-Attached Storage

New ideas and new terms emerge rapidly in data communications and networking. In recent years, a variant on the local area network (LAN) has emerged. A *storage area network (SAN)* is a LAN devoted solely to data storage. When the amount of data to be stored exceeds the practical limits of servers, the SAN plays a critical role. The SAN has a set of high-speed storage devices and servers that are networked together using a very high speed network (often using a technology called *fiber channel* that runs over a series of multi-gigabit point-to-point fiber-optic circuits). Servers are connected into the normal LAN and to the SAN, which is usually reserved for servers. When data are needed, clients send the request to a server on the LAN, which obtains the information from the devices on the SAN and then returns it to the client.

The devices on the SAN may be a large set of *database servers* or a set of network-attached disk arrays. In other cases, the devices may be *network-attached storage (NAS)* devices. A NAS is not a general-purpose computer like a server that runs a server operating system (e.g., Windows, Linux); it has a small processor and a large amount of disk storage and is designed solely to respond to requests for files and data. NAS can also be attached to LANs where they function as a fast database server.

ADS works in much the same manner as TCP/IP's DNS service, and in fact ADS servers, called *domain controllers,* can also act as DNS servers. Network resources are typically organized into a hierarchical tree. Each branch on the tree contains a domain, a group of related resources. For example, at a university, one domain might be the resources available within the business school, and another domain might be the resources in the computer science school, while another might be in the medical school. Domains can contain other domains, and in fact the hierarchical tree of domains within one organization can be linked to trees in other organizations to create a *forest* of shared network resources.

Within each domain, there is a server (the domain controller) that is responsible for resolving address information (much like a DNS server resolves address information on the Internet). The domain controller is also responsible for managing authorization information (e.g., who is permitted to use each resource) and making sure that resources are available only to authorized users. Domain controllers in the same tree (or forest) can share information among themselves, so that a domain controller in one part of the tree (or forest) can be configured to permit access to resources to any user that has been approved by another domain controller in a different part of the tree (or forest).

If you login to a Microsoft server or domain controller that provides ADS, you can see all network resources that you are authorized to use. When a client computer wishes to view available resources or access them, it sends a message using an industry standard directory protocol called *lightweight directory services* (LDAP) to the ADS domain controller. The ADS domain controller resolves the textual name in the LDAP request to a network address and—if the user is authorized to access the resource—provides contact information for the resource.

Network Profiles A *network profile* specifies what resources on each server are available on the network for use by other computers and which devices or people are allowed what access to the network. The network profile is normally configured when the network is established and remains in place until someone makes a change. In a LAN, the server hard disk may have various resources that can or cannot be accessed by a specific network user (e.g., data files, printers). Furthermore, a password may be required to grant network access to the resources.

If a device such as a hard disk on one of the network's computers is not included on the network profile, it cannot be used by another computer on the network. For example, if you have a hard disk (C) on your computer and your computer is connected to this LAN but the hard disk is not included on the network profile assignment list, then no other computer can access that hard disk.

In addition to profiling disks and printers, there must be a *user profile* for each person who uses the LAN, to add some security. Each device and each user is assigned various access codes, and only those users who log in with the correct code can use a specific device. Most LANs keep audit files to track who uses which resource. Security is discussed in Chapter 10.

TRADITIONAL ETHERNET (IEEE 802.3)

Almost all LANs installed today use some form of *Ethernet.* Ethernet was originally developed by DEC, Xerox, and Intel but has since become a standard formalized by the IEEE as *IEEE 802.3.*[1] The IEEE 802.3 version of Ethernet is slightly different from the original version but the differences are minor. Likewise, another version of Ethernet has also been developed that differs slightly from the 802.3 standard. In this section, we describe traditional Ethernet which is sometimes called *shared Ethernet.*

Ethernet is a layer 2 protocol, which means it operates at the data link layer. Every Ethernet LAN needs hardware at layer 1, the physical layer, that matches the requirements of the Ethernet software at layer 2. Ethernet is compatible with a variety of layer 3 protocols but is commonly used with TCP/IP.

Topology

Topology is the basic geometric layout of the network—the way in which the computers on the network are interconnected. It is important to distinguish between a logical topology and a physical topology. A *logical topology* is how the network works conceptually, much like a logical data flow diagram (DFD) or logical entity relation diagram (ERD) in systems analysis and design or database design. A *physical topology* is how the network is physically installed, much like a physical DFD or physical ERD.

Ethernet's logical topology is a *bus topology*. All computers are connected to one half-duplex circuit running the length of the network that is called the bus. The top part of Figure 6.4 shows Ethernet's logical topology. All messages from any computer flow onto the central cable (or bus) and through it to all computers on the LAN. Every computer on the bus receives *all* messages sent on the bus, even those intended for other

[1] The formal specification for Ethernet is provided in the 802.3 standard on the IEEE standards Web site. The URL is http://grouper.ieee.org/groups/802/3.

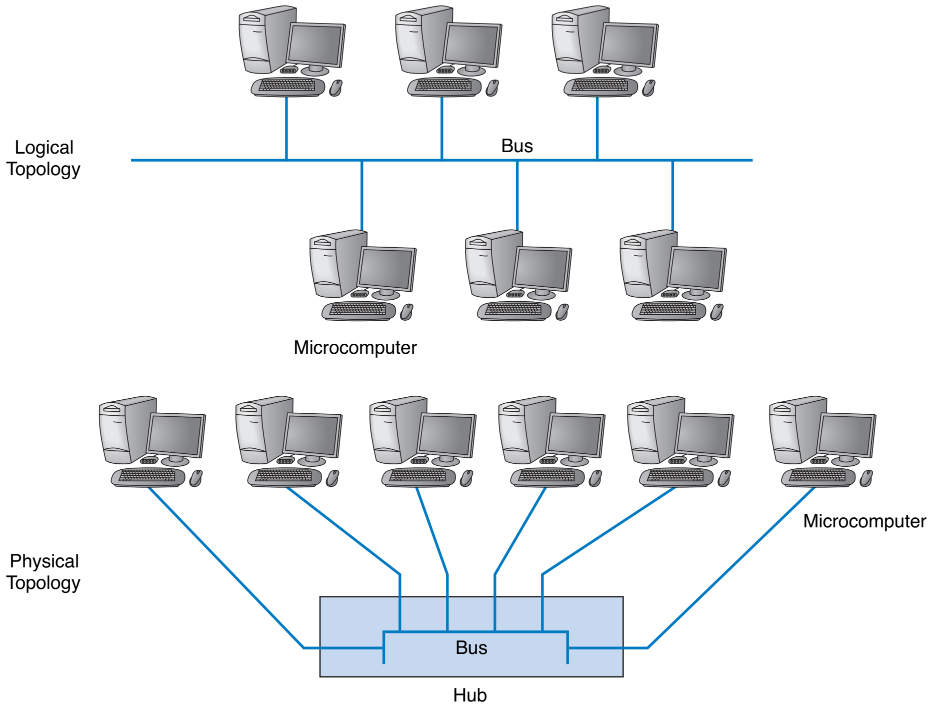

FIGURE 6.4 Ethernet topology.

computers. Before processing incoming messages, the Ethernet software on each computer checks the data link layer address and processes only those messages addressed to that computer.

The bottom part of Figure 6.4 shows the physical topology of an Ethernet LAN when a hub is used. From the outside, an Ethernet LAN *appears* to be a star topology, because all cables connect to the central hub. Nonetheless, it is logically a bus.

Most Ethernet LANs span sufficient distance to require several hubs. In this case, the hubs are connected via cable in the same manner as any other connection in the network (Figure 6.5).

Media Access Control

When several computers share the same communication circuit, it is important to control their access to the media. If two computers on the same circuit transmit at the same time, their transmissions will become garbled. These collisions must be prevented, or if they do occur, there must be a way to recover from them. This is called media access control.

Ethernet uses a contention-based media access control technique called *Carrier Sense Multiple Access with Collision Detection (CSMA/CD).* CSMA/CD, like all contention-based techniques, is very simple in concept: wait until the circuit is free and then transmit. Computers wait until no other devices are transmitting, then transmit their data. As an analogy, suppose you are talking with a small group of friends (four or five

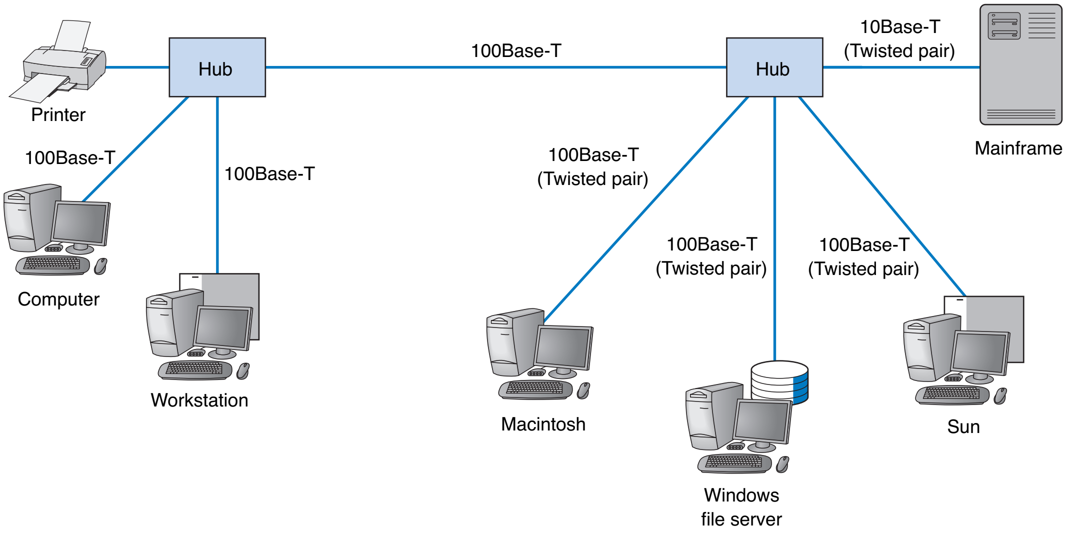

FIGURE 6.5 An example of an Ethernet local area network with two hubs.

people). As the discussion progresses, each person tries to grab the floor when the previous speaker finishes. Usually, the other members of the group yield to the first person who jumps right after the previous speaker.

Ethernet's CSMA/CD protocol can be termed "ordered chaos." As long as no other computer attempts to transmit at the same time, everything is fine. However, it is possible that two computers located some distance from one another can both listen to the circuit, find it empty, and begin simultaneously. This simultaneous transmission is called a *collision.* The two messages collide and destroy each other.

The solution to this is to listen while transmitting, better known as *collision detection (CD).* If the NIC detects any signal other than its own, it presumes that a collision has occurred and sends a jamming signal. All computers stop transmitting and wait for the circuit to become free before trying to retransmit. The problem is that the computers that caused the collision could attempt to retransmit at the same time. To prevent this, each computer waits a random amount of time after the colliding message disappears before attempting to retransmit. Chances are both computers will choose a different random amount of time and one will begin to transmit before the other, thus preventing a second collision. However, if another collision occurs, the computers wait a random amount of time before trying again. This does not eliminate collisions completely, but it reduces them to manageable proportions.

Types of Ethernet

Figure 6.6 summarizes the many different types of Ethernet in use today. *10Base-T* runs on very cheap twisted-pair cable up to 100 meters. It was the 10Base-T standard that revolutionized Ethernet and made it the most popular type of LAN in the world. The ex-

Name	Maximum Data Rate	Cables
10Base-T	10 Mbps	UTP cat 3, UTP cat 5
100Base-T	100 Mbps	UTP cat 5
1000Base-T	1 Gbps	UTP cat 5, UTP cat 5e, UTP cat 6
1000Base-F	1 Gbps	fiber
10 GbE	10 Gbps	UTP cat 5e, UTP cat 6, UTP cat 7, fiber
40 GbE	40 Gbps	fiber

FIGURE 6.6 Types of Ethernet. UTP = unshielded twisted-pair.

tremely low cost of 10Base-T made it very inexpensive compared to its foremost competitor, Token Ring. *100Base-T* is the most common form of Ethernet today.

Three other types of Ethernet have been introduced: *1000Base-T* and 1000Base-F (which run at 1 Gbps and are sometimes called *1 GbE*), *10 GbE* (which runs at 10 Gbps), and *40 GbE* (which runs at 40 Gbps). They can use Ethernet's traditional half-duplex approach, but most are configured to use full duplex. Each is also designed to run over fiber-optic cables, but some may also use traditional twisted-pair wire cables (e.g., Cat 5, Cat 5e). For example, two common versions of 1000Base-F are *1000Base-LX* and *1000Base-SX,* which both use fiber-optic cable, running up to 440 meters and 260 meters,

MANAGEMENT FOCUS

6-3 Hospital Leaps to 10GbE

The good news was that the LAN at the North Bronx Healthcare Network (NBHN) was predictable; unfortunately that was the bad news, too. With zero network downtime in five years, the old network was "a phenomenally stable environment," says Dan Morreale, CIO. But doctors and nurses using the system also could count on phenomenal delays over its 10 Mbps hubs.

A standard prescription for such a network problem might call for a gigabit Ethernet upgrade. Instead, NBHN skipped a step and upgraded its network to 10 GbE. Morreale says he feared that even 1GbE might be outpaced by the hospital's ballooning LAN capacity needs. In recent years, the hospital added digitized medical-imaging technology, which allows X-rays, MRIs, and other images to be viewed and stored on computers instead of film and videotape. Also, doctors and clinicians commonly dictated notes into their desktop PCs instead of onto dictation minicassettes. That prompted the IT staff to set up servers and storage for the bulky voice note files. Videoconferencing among NBHN staff in separate buildings also was taking off.

In addition to updating its LAN and backbone segments, gigabit Ethernet to the desktop also will be in place to support new medical-imaging systems. "The bandwidth involved with that is not insignificant," Morreale says. For a doctor to view a graphic file, such as an X-ray or cardiology image, involves a 200M-byte file download.

SOURCE: "Bronx Hospital Leaps to 10G," *Network World,* August 8, 2003.

respectively; *1000Base-T,* which runs on four pairs of category 5 twisted-pair cable, but only up to 100 meters[2]; and *1000Base-CX,* which runs up to 24 meters on one category 5 cable. Similar versions of 10 GbE and 40 GbE that use different media are also available.

Some organizations use *10/100 Ethernet,* which is a hybrid that uses either 10Base-T or 100Base-T. 10/100 Ethernet NICs have the ability to run at either 10Base-T or 100Base-T, depending on how they are configured. 10/100 autosense hubs (and switches, as we will discuss shortly) detect the signal transmitted by the client's NIC and will use 10 Mbps or 100 Mbps, depending on which the client uses. 10/100 is useful in the short term as organizations move from 10Base-T to 100base-T or if they are uncertain where they want to use which standard.

SWITCHED ETHERNET

Switched Ethernet is identical to traditional Ethernet, except that a switch replaces the hub (Figure 6.7). In traditional shared Ethernet, all devices share the same multipoint circuit and must take turns using it. When a message is sent from one computer to another, it enters the hub, and the hub retransmits it to *all* the computers attached to the hub (Figure 6.7). Each computer looks at the Ethernet address on incoming packets, and if the address on the packet does not match its address, it discards the packet. This process ensures that no two computers transmit at the same time, because they are always listening and do not transmit when they are receiving a message, even if the message is not addressed to them.

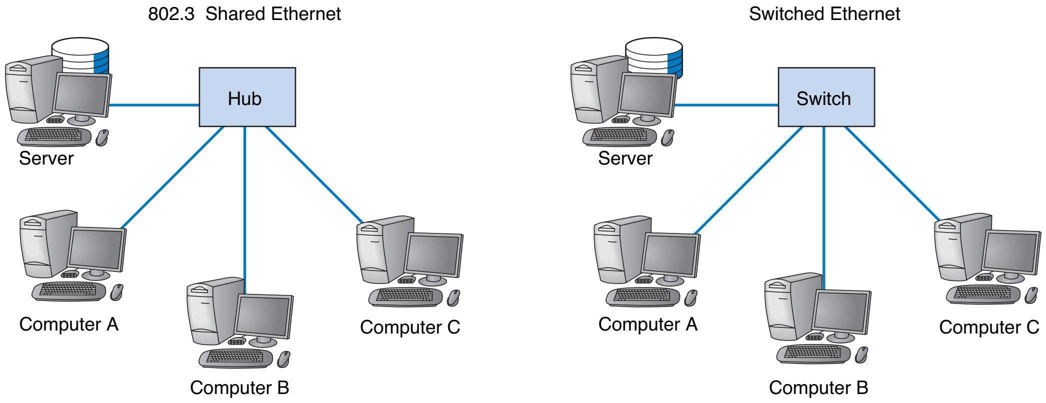

FIGURE 6.7 802.3 Ethernet versus switched Ethernet.

[2] It would be reasonable to think that 1000Base-T would require 10 category 5 cables because 10 × 100Mps = 1000 Mbps. However, it is possible to push 100-Mbps cables to faster speeds over shorter distances. Therefore, the category 5 flavor of 1000Base-T uses only 4 pairs of category 5 (i.e., 8 wires) running at 125 Mbps, but over shorter distances than would be normal for 100Base-T. A special form of category 5 cable (called category 5e) has been developed to meet the special needs of 1000Base-T. This same approach is used to run 10 GbE over category 5.

If the hub did not send the message to all computers, a computer could begin transmitting at the same time as another computer and never be aware of it.

Topology

With switched Ethernet, the hub is replaced by a *switch* (Figure 6.6). This type of switch is often called a *workgroup switch* because it is designed to support a small set of computers (often 16 to 24) in one LAN. From the outside, the switch looks almost identical to a hub, but inside, it is very different. A switch is an intelligent device with a small computer built-in that is designed to manage a set of separate point-to-point circuits. That means that each circuit connected to a switch is *not* shared with any other devices; only the switch and the attached computer use it. The physical topology looks essentially the same as Ethernet's physical topology: a star. On the inside, the logical topology is a set of separate point-to-point circuits, also a star.

When a switch receives a packet from a computer, it looks at the address on the packet and retransmits the packet only on the circuit connected to that computer, not to all circuits as a hub would. For example, in Figure 6.7, if computer A sends a packet to the switch destined for computer C, the switch retransmits it only on the circuit connected to computer C.

So how does a switch know which circuit is connected to what computer? The switch uses a *forwarding table* that is very similar to the routing tables discussed in Chapter 5. The table lists the Ethernet address of the computer connected to each port on the switch. When the switch receives a packet, it compares the destination address on the packet to the addresses in its forwarding table to find the port number on which it needs to transmit the packet. Because the switch uses the Ethernet address to decide which port to use and because Ethernet is a data link layer or layer-2 protocol, this type of switch is called a *layer-2 switch.* In Chapter 8, we describe other types of switches.

When switches are first turned on, their forwarding tables are empty; they do not know what Ethernet address is attached to what port. Switches *learn* addresses to build the forwarding table. When a switch receives a packet, it reads the packet's data link layer source address and compares this address to its forwarding table. If the address is not in the forwarding table, the switch adds it, along with the port on which the message was received.

If a switch receives a packet with a destination address that is not in the forwarding table, the switch must still send the packet to the correct destination. In this case, it must retransmit the packet to all ports, except the one on which the packet was received. In this case, the attached computers, being Ethernet and assuming they are attached to a hub, will simply ignore all messages not addressed to them. The one computer for whom the message is addressed will recognize its address and will process the message, which includes sending an ACK or a NAK back to the sender. When the switch receives the ACK or NAK, it will add this computer's address and the port number on which the ACK or NAK was received to its forwarding table and then send the ACK or NAK on its way.

So, for the first few minutes until the forwarding table is complete, the switch acts like a hub. But as its forwarding table becomes more complete, it begins to act more and more like a switch. In a busy network, it takes only a few minutes for the switch to learn most addresses and match them to port numbers.

There are three modes in which a switch can operate. The first is *cut through switching.* With cut through switching, the switch begins to transmit the incoming packet on the proper outgoing circuit as soon as it has read the destination address in the packet. In other words, the switch begins transmitting before it has received the entire packet. The advantage of this is low *latency* (the time it takes a device from receiving a packet to transmitting it) and results in a very fast network. The disadvantage is that the switch begins transmitting before it has read and processed the frame check sequence at the end of the packet; the packet may contain an error, but the switch will not notice until after almost all of the packet has been transmitted. Cut through switching can only be used when the incoming data circuit has the same data rate as the outgoing circuit.

With the second switching mode, called *store and forward switching,* the switch does not begin transmitting the outgoing packet until it has received the entire incoming packet and has checked to make sure it contains no errors. Only after the switch is sure there are no errors does the switch begin transmitting the packet on the outgoing circuit. If errors are found, the switch simply discards the packet. This mode prevents invalid packets from consuming network capacity, but provides higher latency and thus results in a slower network (unless many packets contain errors). Store and forward switching can be used regardless of whether the incoming data circuit has the same data rate as the outgoing circuit because the entire packet must be stored in the switch before it is forwarded on its way.

The final mode, called *fragment-free switching,* lies between the extremes of cut through and store and forward switching. With fragment-free switching, the first 64 bytes of the packet are read and stored. The switch examines the first 64 bytes (which contain all the header information for the packet) and if all the header data appears correct, the switch presumes that the rest of the packet is error free and begins transmitting. Fragment-free switching is a compromise between cut through and store and forward switching because it has higher latency and better error control than cut through switching, but lower latency and worse error control than store and forward switching. Most switches today use cut through or fragment-free switching.

Media Access Control

Each of the circuits connected to the switch is a separate point-to-point circuit connecting the switch to one computer (or another network device, such as another switch). The switch and the attached computer (or other network device) must share this circuit. Media access control is done in the same manner as traditional Ethernet: each computer (or device) listens before it transmits, and if no one is transmitting, it transmits.

Unlike a hub, in which all attached cables form one shared circuit so that the hub can process only one packet at a time (forcing all attached computers to wait until the one packet is transmitted and it is someone else's turn), a switch is built so that it can simultaneously send or receive packets on *all* the attached circuits. In Figure 6.7, computer A could be sending a packet to the server at the same time as computer B sends one to computer C.

It is possible that two computers may attempt to transmit a packet to the same computer at the same time. For example, both A and B send a packet to C. In this case, the switch chooses which packet to transmit first (usually, the first packet it receives is sent first) and temporarily stores all other packets for that circuit in its internal memory. When the packet is finished and the circuit is again free, the switch then retransmits (or forwards) the temporarily stored packets.

Performance Benefits

In planning a network, it is generally accepted that hub-based 10Base-T LANs can run effectively only to about 50 percent of their capacity. Once the total amount of traffic exceeds 50 percent, so many collisions occur that response time becomes unacceptable. This would mean, for example, that a standard hub-based LAN using 10Base-T is really only capable of providing a total network capacity of only 5 Mbps. This capacity is shared by all computers on the LAN. So if we had 10 computers on one 10Base-T hub, each computer could realistically use about 500 Kbps on average.

As speeds increase, packets take less time to transmit on the circuit and the probability of collisions decreases. Tests have shown that 100Base-T can run close to 90 percent of capacity with few problems.

Switched Ethernet dramatically improves network performance because each computer has its own dedicated point-to-point circuit, rather than the one common shared multipoint circuit in traditional hub-based Ethernet. Because there are only two devices on each point-to-point circuit (e.g., the switch and a computer), the probability of a collision is lower. We do not yet have extensive experience with Ethernet switches, but some experts believe we can effectively use up to about 95 percent of the switched Ethernet capacity before performance becomes a problem. So each 10Base-T switched circuit effectively has a maximum capacity of about 9.5 Mbps. Therefore, if we have 10 computers on one 10base-T switch, this would mean that on average, each computer could realistically use about 9.5 Mbps, giving a total network capacity of about 95 Mbps.

In most LANs, the majority of network traffic is to and from the server, or to and from the connection from the LAN to the BN (the gateway in TCP/IP terminology used in Chapter 5, or more commonly, a device called a router, as discussed in Chapter 8). In most LANs, this circuit is the network bottleneck. Each computer is transmitting at 10 Mbps, but if the circuit to the server is also 10 Mbps, there is often a traffic jam. The solution to this is to use a 10/100 switch, which provides 10-Mbps circuits to the client computers but a 100-Mbps circuit to the server or BN. Although traffic jams will still occur, the higher speed on the bottleneck circuit will mean they will clear up much more quickly.

THE BEST PRACTICE LAN DESIGN

The past few years have seen major changes in LAN technologies (e.g., gigabit Ethernet, switched Ethernet). As technologies have changed, so too has our understanding of the best practice design for LANs.[3]

[3] We thank our friends at Cisco Systems Inc., the market leader in LAN and backbone networking, for helping us think about this.

Effective Data Rates

The *effective data rate* of the hardware layers is the maximum practical speed in bits that the hardware layers can be expected to provide. The effective data rate depends on four basic factors. The first factor is the nominal data rate provided by the physical layer; that is, the data rate specified by the hardware (e.g., 10Base-T provides a nominal rate of 10 Mbps). The second is the error rate because this determines how many retransmissions must occur. The third is the efficiency of the data link layer protocols used. As discussed in the previous chapters, efficiency is the percentage of a transmission that contains user data and is dependent on the number of overhead bytes in the transmission. The final factor is the efficiency of the media access control protocol; that is, how well the media access control protocol can use the nominal data rate.

Data Link Protocol Efficiency Shared Ethernet and switched Ethernet share the same data rates, the same types of cables that can be assumed to have the same error rates, and the same data link protocol with the same efficiency. The efficiency of the Ethernet data link protocols (excluding higher-level protocols such as TCP/IP) is fairly good. For every 1,500-byte packet transmitted, there are 33 bytes of overhead on the packet itself. Thus assuming we have no errors requiring a retransmission, we have an efficiency of about 98 percent if we send 1,500-byte packets (1467/1500 = 97.8%). If we use jumbo packets (9,000 bytes), then the efficiency is about 99.6 percent. Conversely, if we transmit mostly small packets (e.g., 150-byte Web requests), then data link protocol efficiency is only about 82 percent (150/183). (Remember that these calculations do not include the overhead imposed by higher-level packets such as TCP/IP.)

Average efficiency depends on typical pattern of packet sizes and thus differs from LAN to LAN, depending on the number of users and what applications they use. To estimate an average efficiency, we must make some assumptions about the nature of traffic in a "typical" LAN, thus any estimate we derive could differ from the actual efficiency of a specific LAN if the pattern of traffic in the LAN is different from our assumptions. Generally speaking, the pattern of traffic in most LANs for Web or e-mail applications is a small HTTP or SMTP request sent from the client to a server, followed by a long series of large packets from the server to the client providing a Web page or e-mail message. Thus, most traffic is large packets. If we assume that each short packet is followed by 20 large packets (e.g., each Web request produces a set of files totaling 30–50 K in response), then our average efficiency is about 97 percent. Thus we will use 97 percent as a reasonable estimate of Ethernet's data link layer protocol efficiency for typical LAN traffic. It is also important to note that this assumes that virtually no errors occur, which is a reasonable assumption for most LAN environments today.

Media Access Control Protocol Efficiency Shared Ethernet and switched Ethernet differ in the media access control protocol. It is generally accepted that Ethernet's CSMA/CD media access control protocol works very well in low-traffic networks. As traffic increases and network utilization increases, collisions become more common. Several mathematical models, simulations, and real experiments with shared and switched Ethernet running at different data rates using different assumptions about the number of computers on the network and the types of traffic they generate (e.g., large packets versus

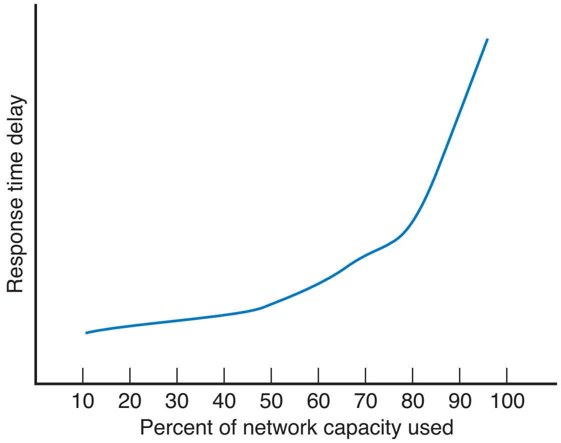

FIGURE 6.8 Performance of Ethernet LANs.

short packets) have been done. Ethernet performance varies based on the assumptions one uses, but a general pattern does emerge.

As shown in Figure 6.8, the response time delays experienced by users are low when there is little traffic (lower delays are better). Response time delays increase slowly as traffic increases to about 50 percent of the nominal data rate. Once the 50 percent capacity mark is reached, response time delays increase much more quickly as traffic increases, until about 80 percent of capacity is reached. Past 80 percent, delays increase exponentially as traffic increases.

In other words, Ethernet LANs work very well and their users experience few response time delays as long as the total amount of traffic in the LAN remains under 50 percent of the nominal data capacity. As traffic increases to between 50 percent and 80 percent of capacity, users experience noticeable delays but can still use the network. Once capacity hits 80 percent, the delays make the network effectively unusable.

This means, for example, that a shared hub-based LAN using 10Base-T is really only capable of providing a total network capacity of just under 5 Mbps (97% efficiency × 50% capacity × 10 Mbps = 4.85 Mbps). This capacity is shared by all computers on the LAN. So in order to estimate the effective data rate of shared Ethernet, we must make some assumptions about the number of computers that will be active—that is, *simultaneously* be sending and receiving data over the network. The key word here is *simultaneously;* a typical shared Ethernet LAN today has about 20 computers, and except for computer labs, most computers are not simultaneously sending and receiving data. Even when users are actively using the computer, they are seldom constantly sending and receiving data; most users pause to read the Web pages or e-mail messages they retrieve.

In a low-traffic network, we might expect only one or two of the attached users to simultaneously attempt to send or receive data over the network. With two users, the total capacity is divided among both users. So if we had two active computers in a low-traffic 10Base-T shared Ethernet environment, this would mean that on average, each computer could realistically use about 2.5 Mbps. In a moderate-traffic LAN, we might have five active users, meaning each computer could realistically use about 1 Mbps on average. In a high-traffic environment with 10 active computers on one 10Base-T hub, this would mean that, on average, each computer could realistically use about 500 Kbps on average (Figure 6.9).

Tests have shown that shared 100Base-T can run close to 80 percent of capacity with very few delays. On a high-traffic LAN with 10 active computers using shared 100Base-T, this would mean that each computer could realistically use about 7.5 Mbps on average (97% efficiency × 80% capacity × 100 Mbps = 7.8 Mbps) (Figure 6.9).

Switched Ethernet dramatically improves network performance because each computer has its own dedicated circuit rather than the one common shared multipoint circuit in shared Ethernet. Because there are only two devices on each half-duplex point-to-point circuit (e.g., the switch and a computer), the probability of a collision is lower. Most experts believe we can effectively use up to about 95 percent of the switched Ethernet capacity before performance becomes a problem. In 10Base-T switched LAN, each computer circuit would have an effective capacity of about 9 Mbps (97% efficiency × 95% capacity × 10 Mbps = 9.2 Mbps). In a 100Base-T switched LAN, each computer would have about 92 Mbps (95% efficiency × 95% capacity × 100 Mbps = 92 Mbps). Because each computer has its own circuit connecting it to the switch, it is unaffected by the amount of traffic generated by the other computers on the switch—assuming, of course, that not all computers are trying to send a message to the same computer or device attached to the switch, which is sometimes the case.

Gigabit Ethernet is most often implemented in full-duplex switched environments, which means it provides 1 Gbps in both directions simultaneously. It provides a data rate of about 900 Mbps, but one could argue that since this is full-duplex and available in both directions simultaneously, a better relative number might be 1.8 Gbps per computer. Ten GbE is similar, so it provides about 18 Gbps per computer.

Figure 6.9 provides a summary of the effective data rates. These rates provide a general guide because, as we noted above, one must make certain assumptions about the typical frame sizes, error rates, reasonable response time expectations of users, number of active users, and so on. It is also important to note that these numbers do not include the effects of higher-layer packets (e.g., TCP/IP) in the calculations—they focus only on the hardware layers.

Costs

When new technologies are first introduced, they are expensive. As time passes, their prices drop as new technologies appear that outperform them. Today, shared 10Base-T Ethernet equipment is very cheap and shared 100Base-T is relatively inexpensive because both are quite old in design. Switched Ethernet, both 10Base-T and 100Base-T, are also relatively inexpensive. 1 GbE and 10 GbE are both quite expensive.

Recommendations

Given these trade-offs in costs and effective data rates, there are several best practice recommendations (Figure 6.10). For most networks, shared 100Base-T provides the best trade-off between cost and performance. As the cost of technology continues to drop, pure 10Base-T devices are starting to disappear. The difference in manufacturing cost between

	Effective Data Rate per User		
Technology	Low Traffic	Moderate Traffic	High Traffic
Shared 10Base-T	2.5 Mbps	1 Mbps	500 Kbps
Shared 100Base-T	37.5 Mbps	15 Mbps	7.5 Mbps
Switched 10Base-T	9 Mbps	9 Mbps	9 Mbps
Switched 100Base-T	92 Mbps	92 Mbps	92 Mbps
Full Duplex 1 GbE	1.8 Gbps	1.8 Gbps	1.8 Gbps
Full Duplex 10 GbE	18 Gbps	18 Gbps	18 Gbps

Assumptions:

1. Most packets are 1,500 bytes or larger
2. No transmission errors occur
3. Low traffic means 2 active users, moderate traffic means 5 active users, high traffic means 10 active users

FIGURE 6.9 Effective data rate estimates for Ethernet.

10Base-T and 100Base-T devices is small, so some vendors are discontinuing 10Base-T-only devices and selling 10/100 autosensing devices that run at 10 Mbps or 100 Mbps at almost the same cost as 10Base-T devices.

Most network managers install category 5 or 5e cables (rated to 100 Mbps) even though category 3 cables are sufficient for 10Base-T because the additional cost for cat 5/5e is very small and this provides room for upgrades to 100Base-T or 1000Base-T.

For very small networks, such as home networks connecting only a handful of computers, traditional shared 10Base-T over cat 5/5e cable should prove sufficient because of their low traffic demands (although, as we noted above, this technology is dying out). For networks with very high traffic needs, switched 100Base-T or 1 GbE over fiber is recommended, although as the price of gigabit Ethernet drops, it will become the recommended best practice.

Most networks	Shared 100Base-T Ethernet over Category 5e cables
Very small networks (e.g., home networks)	Shared 10Base-T Ethernet over Category 5 or Category 5e cables
Networks with high demands (e.g., multimedia networks)	Switched 100Base-T Ethernet over Category 5e cables or full duplex 1 GbE over fiber

FIGURE 6.10 Best practice LAN recommendations.

IMPROVING LAN PERFORMANCE

When LANs had only a few users, performance was usually very good. Today, however, when most computers in an organization are on LANs, performance can be a problem. Performance is usually expressed in terms of throughput (the total amount of user data transmitted in a given time period). In this section, we discuss how to improve throughput. We focus on dedicated-server networks because they are the most commonly used type of LANs, but many of these concepts also apply to peer-to-peer networks.

To improve performance, you must locate the *bottleneck,* the part of the network that is restricting the data flow. Generally speaking, the bottleneck will lie in one of two places. The first is the network server. In this case, the client computers have no difficulty sending requests to the network server, but the server lacks sufficient capacity to process all the requests it receives in a timely manner. The second location is the network circuit, often the circuit connecting the LAN to the corporate BN. In this case, the server can easily process all the client requests it receives, but the circuit lacks enough capacity to transmit all the requests to the server. It is also possible that the bottleneck could lie in the client computers themselves (e.g., they are receiving data too fast for them to process it), but this is extremely unlikely—unless, of course, you are still using old computers!

The first step in improving performance, therefore, is to identify whether the bottleneck lies in the circuit or the server. To do so, you simply watch the utilization of the server during periods of poor performance. If the server utilization is high (e.g., 60 to 100

TECHNICAL FOCUS

6-3 ERROR CONTROL IN ETHERNET

Ethernet provides a strong error control method using stop and wait ARQ with a CRC-32 error detection field (see Chapter 4). However, the normal way of installing Ethernet doesn't use stop and wait ARQ.

In the early days of Ethernet, LAN environments were not very reliable, so error control was important. However, today's LAN environments are very reliable; errors seldom occur. Stop and wait ARQ uses considerable network capacity because every time a packet is transmitted, the sender must stop and wait for the receiver to send an acknowledgment. By eliminating the need to stop and wait and the need to send acknowledgments, Ethernet can significantly improve network performance—almost doubling the number of messages that can be transmitted in the same time period. Ethernet does still add the CRC and does still check it for errors, but any packet with an error is simply discarded.

If Ethernet doesn't provide error control, then higher layers in the network model must. In general, TCP is configured to provide error control by using continuous ARQ (see Chapter 5) to ensure that all packets that have been sent are actually received at the final destination. If a packet with an error is discarded by Ethernet, TCP will recognize that a packet has been lost and ask the sender to retransmit. This moves responsibility for error control to the edges of the network (i.e., the sender and receiver) rather than making every computer along the way responsible for ensuring reliable message delivery.

percent), then the bottleneck is the server; it cannot process all the requests it receives in a timely manner. If the server utilization is low during periods of poor performance (e.g., 10 to 40 percent), then the problem lies with the network circuit; the circuit cannot transmit requests to the server as quickly as necessary. Things become more difficult if utilization is in the midrange (e.g., 40 to 60 percent). This suggests that the bottleneck may shift between the server and the circuit depending on the type of request, and it suggests that both should be upgraded to provide the best performance.

Now we will focus attention on ways to improve the server and the circuit to remove bottlenecks. These actions address only the supply side of the equation—that is, increasing the capacity of the LAN as a whole. The other way to reduce performance problems is to attack the demand side: reduce the amount of network use by the clients, which we also discuss. Figure 6.11 provides a performance checklist.

Improving Server Performance

Improving server performance can be approached from two directions simultaneously: software and hardware.

Software The NOS is the primary software-based approach to improving network performance. Some NOSs are faster than others, so replacing the NOS with a faster one will improve performance.

Performance Checklist

Increase Server Performance

- Software
- Fine-tune the network operating system settings
- Hardware
- Add more servers and spread the network applications across the servers to balance the load
- Upgrade to a faster computer
- Increase the server's memory
- Increase the number and speed of the server's hard disk(s)
- Upgrade to a faster network interface card

Increase Client Capacity

- Upgrade to a faster circuit
- Segment the network

Reduce Network Demand

- Move files from the server to the client computers
- Increase the use of disk caching on client computers
- Change user behavior

FIGURE 6.11 Improving local area network performance.

Each NOS provides a number of software settings to fine-tune network performance. Depending on the number, size, and type of messages and requests in your LAN, different settings can have a significant effect on performance. The specific settings differ by NOS but often include things such as the amount of memory used for disk caches, the number of simultaneously open files, and the amount of buffer space.

Hardware One obvious solution if your network server is overloaded is to buy a second server (or more). Each server is then dedicated to supporting one set of application software (e.g., one handles e-mail, another handles the financial database, and another stores customer records). The bottleneck can be broken by carefully identifying the demands each major application software package places on the server and allocating them to different servers.

Sometimes, however, most of the demand on the server is produced by one application that cannot be split across several servers. In this case, the server itself must be upgraded. The first place to start is with the server's CPU. Faster CPUs mean better performance. If you are still using an old computer as a LAN server, this may be the answer; you probably need to upgrade to the latest and greatest. Clock speed also matters: the faster, the better. Most computers today also come with CPU-cache (a very fast memory module directly connected to the CPU). Increasing the cache will increase CPU performance.

A second bottleneck is the amount of memory in the server. Increasing the amount of memory increases the probability that disk caching will work, thus increasing performance.

A third bottleneck is the number and speed of the hard disks in the server. The primary function of the LAN server is to process requests for information on its disks. Slow hard disks give slow network performance. The obvious solution is to buy the fastest disk drive possible. Even more important, however, is the number of hard disks. Each computer hard disk has only one read/write head, meaning that all requests must go through this one device. By using several smaller disks rather than one larger disk (e.g., five 20-gigabyte disks rather than one 100-gigabyte disk), you now have more read/write heads, each of which can be used simultaneously, dramatically improving throughput. A special type of disk drive called *RAID (redundant array of inexpensive disks)* builds on this concept and is typically used in applications requiring very fast processing of large volumes of data, such as multimedia. Of course, RAID is more expensive than traditional disk drives, but costs have been shrinking. RAID can also provide fault tolerance, which is discussed in Chapter 11.

A fourth bottleneck is the NIC itself. Simply put, some NICs are faster than others. Some NICs provide built-in CPUs to perform some of the network functions usually handled by the server (much like front-end processors in mainframe networks). Others provide memory and cache to improve the access time to and from the network.

Several vendors sell special-purpose network servers that are optimized to provide extremely fast performance. Many of these provide RAID and use *symmetric multiprocessing (SMP)* that enables one server to use up to 16 CPUs. Each of these CPUs may be an Intel chip such as Pentium, or may be based on reduced instruction set computing (RISC). Such servers provide excellent performance but cost more than a standard microcomputer (often $5,000 to $15,000).

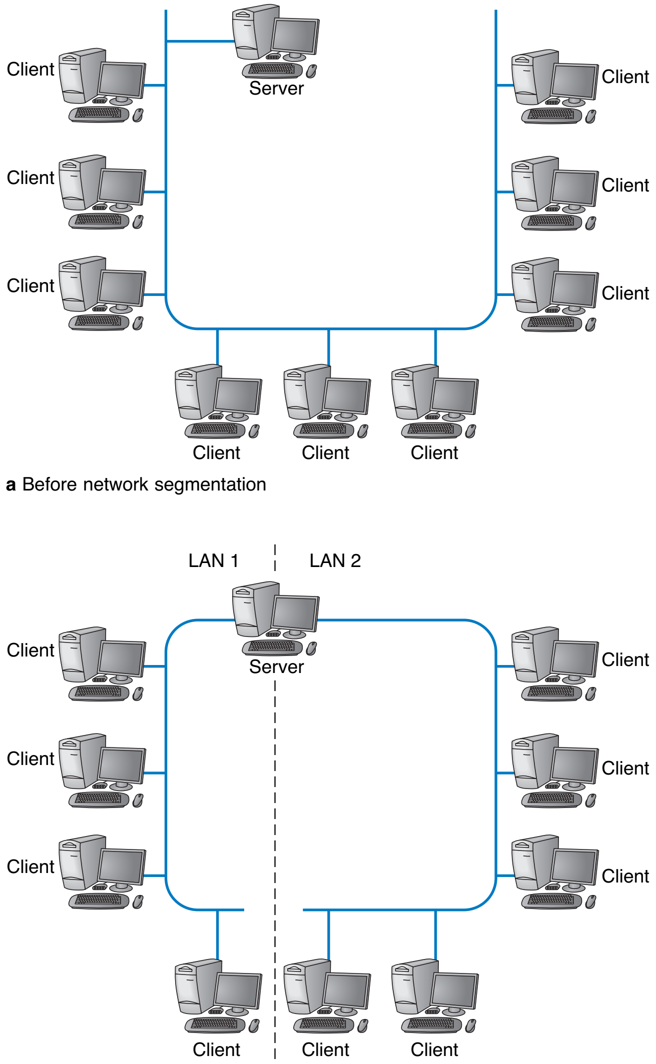

a Before network segmentation

b After segmentation

FIGURE 6.12 Network segmentation example.

Improving Circuit Capacity

Improving the capacity of the circuit means increasing the volume of simultaneous messages the circuit can transmit from network clients to the server(s). One obvious approach is simply to buy a bigger circuit. For example, if you are now using a traditional hub-based 10Base-T LAN, upgrading to 100Base-T or switched 10Base-T will improve capacity.

The other approach is to segment the network. If there is more traffic on a LAN than the network circuit and media access protocol can handle, the solution is to divide the

LAN into several smaller segments. Breaking a network into smaller parts is called *network segmentation.* By carefully identifying how much each computer contributes to the demand on the server and carefully spreading those computers to different network segments, one can often break a network bottleneck.

Figure 6.12 presents an example in which each network segment is connected into the same server. Most servers can support as many as 16 separate networks or network segments simply by adding one NIC into the server for each network. As the number of NICs in the server increases, however, the server spends more of its processing capacity monitoring and managing the NICs and has less capacity left to process client requests. Most experts recommend no more than three or four NICs per server. There are two ways to create more network segments: one is to use more servers, each dedicated to one or more segments, and the other is to use a BN to connect different segments. BNs are discussed in the next chapter.

Reducing Network Demand

Upgrading the server hardware and software, choosing a different LAN protocol, or segmenting the LAN are all strategies to increase network capacity. Performance also can be improved by attempting to reduce the demand on the network.

One way to reduce network demand is to move files to client computers. Heavily used software packages that continually access and load modules from the network can place unusually heavy demands on the network. Although user data and messages are often only a few kilobytes in size, today's software packages can be many megabytes in size. Placing even one or two such applications on client computers can greatly improve network performance (although this can create other problems, such as increasing the difficulty in upgrading to new versions of the software).

Another way is to increase the use of disk-caching software on the client machines to reduce the client's need to access disk files stored on the server. For example, most Web browsers store Web pages in their cache so that they can access previously used pages from their hard disks without accessing the network.

Because the demand on most LANs is uneven, network performance can be improved by attempting to move user demands from peak times to off-peak times. For example, early morning and after lunch are often busy times when people check their e-mail. Telling network users about the peak times and encouraging them to change their habits may help; however, in practice, it is often difficult to get users to change. Nonetheless, finding one application that places a large demand on the network and moving it can have a significant impact (e.g., printing several thousand customer records after midnight).

IMPLICATIONS FOR MANAGEMENT

As LANs have standardized on Ethernet, local area networking technology has become a commodity in most organizations. As with most commodities, the cost of LAN equipment (i.e., network interface cards, cabling, hubs, and switches) has dropped significantly. Some vendors are producing high-quality equipment while some new entrants into the market are producing equipment that meets standards but creates opportunities for problems because it lacks the features of more established brands. It becomes difficult for LAN managers to explain to business managers why its important to purchase higher-quality, more expensive equipment when low-cost "standardized" equipment is available.

As costs for LAN equipment drop, LANs are becoming more common in homes, student apartments, and small offices. What once was groundbreaking new technology in the early 1990s has now become a standard consumer product. As LANs become commonplace in homes, apartments, and offices, new software applications will be developed to take advantage of these new capabilities.

Decreasing costs for LAN equipment also means that network enabled microprocessor controlled devices that have not normally been thought of as computer technology is becoming less expensive. Therefore, we have seen devices such as copiers turned into network printers and scanners. This trend will increase as electrical appliances such as refrigerators and ovens become network devices. Don't laugh; networked vending machines are already in use.

SUMMARY

Why use a LAN? The two basic reasons for developing a LAN are information sharing and resource sharing. *Information sharing* refers to business needs that require users to access the same data files, exchange information via e-mail, or search the Internet for information, as discussed in Chapter 2. *Resource sharing* refers to one computer sharing a hardware device (e.g., a printer) or software package with other computers on the network. The main benefit of resource sharing is cost savings whereas the main benefit of information sharing is improved decision making.

Dedicated-Server versus Peer-to-Peer Networks A dedicated-server LAN has one computer that acts as the network server. It can connect with almost any other network, handle very large databases, and use sophisticated LAN software. Moreover, high-end dedicated-server LANs can be interconnected easily to form enterprisewide networks or, in some cases, replace the host mainframe central computer. Common types of dedicated servers include Web servers, application servers, file servers, database servers, print servers, and remote access servers. All computers on a peer-to-peer LAN run special network software that enables them to function both as a client and as a server.

LAN Components The NIC enables the computer to be physically connected to the network cable and provides the physical layer connection among the computers in the network. Most LANs use UTP wires, STP wires, coaxial cable, and/or fiber-optic cable. Network hubs provide an easy way to connect network cables and act as repeaters or amplifiers. Most new buildings built today have a separate LAN cable plan, just as they have plans for telephone cables and for electrical cables. The NOS is the software that performs the functions associated with the data link and the network layers and interacts with the application software and the computer's own operating system. Every NOS provides two sets of software: one that runs on the network server(s) and one that runs on the network client(s). A network profile specifies what resources on each server are available for network use by other computers and which devices or people are allowed what access to the network.

Ethernet (IEEE 802.3) Ethernet, the most commonly used LAN protocol in the world, uses a logical bus topology that has a shared multipoint circuit used by all attached computers and devices although the physical appearance of the network is a star. It uses a contention-based media access technique called CSMA/CD. There are many different types of Ethernet that use different network cabling (e.g., 10Base-T, 100Base-T, 1000Base-T, 10 GbE).

Switched Ethernet With switched Ethernet, a switch replaces the hub, but otherwise, all other components are identical. The switch provides a series of separate point-to-point circuits to the attached devices, so that no device needs to wait for another device before it transmits. When a packet arrives at the switch, the switch reads the Ethernet address and then forwards the packet to the one destination computer. Switched Ethernet has considerably better performance than traditional Ethernet because computers do not have to share circuits with other computers.

Best Practice LAN Design The best practice LAN design depends on cost and the effective data rate of the LAN hardware layers, which in turn depends on the nominal data rate provided by the physical layer, the error rate, the efficiency of the data link layer protocol, and the efficiency of the media access control protocol. Given the trade-offs in costs and effective data rates, the best LAN design for most networks is shared 100Base-T with category 5/5e cables. For very small networks, such as home networks connecting only a handful of computers, traditional shared 10Base-T over cat 5/5e cable may prove sufficient because of their low traffic demands. For networks with very high traffic, switched 100Base-T is recommended although as the price of gigabit Ethernet drops, it will become the recommended best practice.

Improving LAN Performance Every LAN has a bottleneck, a narrow point in the network that limits the number of messages that can be processed. Generally speaking, the bottleneck will lie in either the network server or the network circuit. Server performance can be improved with a faster NOS that provides better disk caching, by buying more servers and spreading applications among them, or by upgrading the server's CPU, memory, NIC, and the speed and number of its hard disks. Circuit capacity can be improved by using faster technologies (100Base-T rather than 10Base-T) and by segmenting the network into several separate LANs. Overall LAN performance also can be improved by reducing the demand for the LAN by moving files off the LAN, using disk caching on the client computers, and by shifting users' routines.

KEY TERMS

Active Directory Service (ADS)
bottleneck
bus topology
cable plan
cabling
Carrier Sense Multiple Access with Collision Detection (CSMA/CD)
collision
collision avoidance (CA)
collision detection (CD)
cut through switching
database server
dedicated server
domain controller
Ethernet
fiber channel
fiber-optic cable
file server
forwarding table
fragment-free switching
hub
IEEE 802.3
information sharing
LAN management software
LAN metering software
latency
layer-2 switch
lightweight directory services (LDAP)
logical carrier sense method
logical topology
network-attached storage (NAS)
network interface card (NIC)
network operating system (NOS)
network profile
network segmentation
network server
peer-to-peer network
physical topology
print server
port
redundant array of inexpensive disks (RAID)
remote-access server (RAS)
resource sharing
server farm
shared Ethernet
shielded twisted-pair (STP)
software audit
software piracy
Software Publishers Association (SPA)
storage area network (SAN)
store and forward switching
switch
switched Ethernet
symmetric multi-processing (SMP)
topology
transceiver
twisted-pair wiring
unshielded twisted-pair (UTP) wiring
user profile
workgroup switch
1 GbE
10 GbE
40 GbE
10Base-T
100Base-T
1000Base-T
10/100 Ethernet

QUESTIONS

1. Define *local area network.*
2. What are the distinguishing features of a LAN?
3. What are two reasons for developing LANs?
4. What is the function of LAN metering software?
5. Discuss the legal issue of using single-computer license software on networks.
6. Discuss why it is important for organizations to enforce policies restricting use of employee-owned hardware and software and unauthorized copies of software.
7. In some LANs, most of the computers talk with the server, but others use no server. What are these two approaches called?
8. Describe at least three types of servers.
9. What is a NIC? What is a hub?
10. What media do LANs normally use?
11. What type of cables are commonly used in LANs?
12. Compare and contrast category 5 UTP, category 5e UTP, and category 5 STP.
13. What is a cable plan and why would you want one?
14. What does a NOS do? What are the major software parts of a NOS?
15. What is the most important characteristic of a NOS?
16. What is a network profile?
17. What is Ethernet? How does it work?
18. How does a logical topology differ from a physical topology?
19. Briefly describe how CSMA/CD works.
20. Why should CSMA/CD networks be built so that no more than 50 percent of their capacity is dedicated to actual network traffic?
21. Explain the terms 100Base-T, 100Base-F, 1000Base-T, 10 GbE, and 10/100 Ethernet.
22. How does switched Ethernet differ from traditional Ethernet?
23. How do layer-2 Ethernet switches know where to send the packets they receive? Describe how switches gather and use this knowledge.
24. What are the primary advantages and disadvantages of switched Ethernet?
25. What is an effective data rate and how do you calculate it?
26. Under what circumstances does shared Ethernet provide its best performance? At what point does shared Ethernet performance begin to rapidly decline?
27. Compare Ethernet to other data link protocols from previous chapters in terms of efficiency.
28. Why is the effective data rate per user so different between shared Ethernet and switched Ethernet?
29. Why doesn't the data rate available to each user of gigabit Ethernet change as traffic increases?
30. What is a bottleneck and how can you locate one?
31. Describe four ways to improve network performance on the server.
32. Describe four ways to improve network performance on the circuit.
33. Why does network segmentation improve LAN performance?
34. It is said that hooking some computers together with a cable does not make a network. Why?
35. Compare and contrast cut through, store and forward, and fragment-free switching.
36. Is 1 GbE Ethernet really "Ethernet?" Explain.
37. Under what circumstances is switched Ethernet preferred to shared Ethernet? Under what circumstances is shared Ethernet preferred to switched Ethernet?
38. As the cost of 100Base-T Ethernet continues to drop, many people predict that 10Base-T will fade away. What do you think? Why?

EXERCISES

6-1. Survey the LANs used in your organization. Are they Ethernet, switched Ethernet, or some other standard? Why?

6-2. Document one LAN (or LAN segment) in detail. What devices are attached, what cabling is used, and what is the topology? What does the cable plan look like?

6-3. You have been hired by a small company to install a simple LAN for their 18 Windows computers. Develop a simple LAN and determine the total cost; that is, select the cables, hubs/switches, and NICs and price them.

MINI-CASES

I. Designing a New Ethernet

One important issue in designing Ethernet lies in making sure that if a computer transmits a packet, any other computer that attempts to transmit at the same time will be able to hear the incoming packet before it stops transmitting, or else a collision might go unnoticed. For example, assume that we are on earth and send an Ethernet packet over a very long piece of category 5 wire to the moon. If a computer on the moon starts transmitting at the same time as we do on earth and finishes transmitting before our packet arrives at the moon, there will be a collision, but neither computer will detect it; the packets will be garbled, but no one will know why. So, in designing Ethernet, we must make sure that the length of cable in the LAN is shorter than the length of the shortest possible message that can be sent. Otherwise, a collision could go undetected.

a. Let's assume that the smallest possible message is 64 bytes (including the 33-byte overhead). If we use 10Base-T, how long (in meters) is a 64-byte message? While electricity in the cable travels a bit slower than the speed of light, once you include delays in the electrical equipment in transmitting and receiving the signal, the effective speed is only about 40 million meters per second. (*Hint:* First calculate the number of seconds it would take to transmit the message then calculate the number of meters the signal would travel in that time, and you have the total length of the message.)

b. If we use 10 GbE, how long (in meters) is a 64-byte message?

c. The answer in part b is the maximum distance any single cable could run from a switch to one computer in a switched Ethernet LAN. How would you overcome the problem implied by this?

II. Pat's Petunias

You have been called in as a network consultant by your cousin Pat who operates a successful mail-order flower business. She is moving to a new office and wants to install a network for her telephone operators, who take phone calls and enter orders into the system. The number of operators working varies depending on the time of day and day of the week. On slow shifts, there are usually only 10 operators, whereas at peak times, there are 50. She has bids from different companies to install (1) a shared Ethernet 10Base-T network, (2) a switched Ethernet 10Base-T network, or (3) a switched Ethernet 100Base-T network. She wants you to give her some sense of the relative performance of the three alternatives so the can compare that with their different costs. What would you recommend?

III. Eureka!

Eureka! is a telephone and Internet-based concierge service that specializes in obtaining things that are hard to find (e.g., Super Bowl tickets, first-edition books from the 1500s, Fabergé eggs). It currently employs staff members who work 24 hours per day (over three shifts), with usually 5 to 7 staff members working at any given time. Staff members answer the phone and respond to requests entered on the Eureka! Web site. Much of their work is spent on the phone and on computers searching on the Internet. They have just leased a new office and are about to wire it. They have bids from different companies to install (a) a shared Ethernet 100Base-T network, (b) a switched Ethernet 10Base-T network, (c) a switched Ethernet 100Base-T network, or (d) a switched 100Base-F network. What would you recommend? Why?

IV. Tom's Home Automation

Your cousin Tom runs a small construction company that builds custom houses. He has just started a new specialty service that he is offering to other builders on a subcontracting basis: home automation. He provides a complete service of installing cable in all the rooms in which the homeowner wants data access and installs the necessary

(continued)

networking devices to provide a LAN that will connect all the computers in the house to the Internet. Most homeowners choose to install a DSL or cable modem Internet connection that provides a 1–2 Mbps from the house to the Internet (see Chapter 10). Tom has come to you for advice. What type of cabling (e.g., cat 3, cat 5, cat 5e, cat 6, fiber optic) and what type of networking hardware (e.g., hub or switch) would you recommend? Why?

V. Sally's Shoes

Sally Smith runs a shoe store in the mall that is about 30 feet by 50 feet in size, including a small office and a storage area in the rear. The store has one inventory computer in the storage area and one computer in the office. She is replacing the two cash registers with computers that will act as cash registers but will also be able to communicate with the inventory computer. Sally wants to network the computers with a LAN. What sort of a LAN design would you recommend in terms of cabling and hubs or switches? Draw a picture. Should Sally use peer-to-peer networking or use a dedicated server?

CASE STUDY

NEXT-DAY AIR SERVICE

See the Web site.

HANDS-ON ACTIVITY

Windows Peer-to-Peer Networking

In this chapter, we've discussed two types of LANs: peer-to-peer LANs and dedicated server LANs. This activity will show you how to set up a peer-to-peer LAN for your house or apartment. We first describe file sharing and then discuss printer sharing.

Windows File Sharing

Windows file sharing enables you to select folders on your computer that you can permit others users on your LAN to read and write. There are three steps to create a shared folder.

Step 1. Give your computer an Application Layer Name within a Workgroup

1. Go to Settings → Control Panel → System
2. Click on the Computer Name Tab
3. Click Change
4. Type in a New Computer Name and Workgroup Name. All computers must have the same workgroup name to share files. Each computer within a workgroup must have a unique name.

Step 2. Enable File Sharing

1. Go to Settings → Control Panel → Windows Firewall
2. Click on the Exceptions tab
3. Make sure the box in front of File and Printer Sharing is checked
4. Go to Settings → Control Panel → Network Connections
5. Right click on the LAN connection and click Properties
6. Ensure that the box in front of File and Printer Sharing for Microsoft Networks is checked.

Step 3. Create the Shared Folder

1. Open Windows Explorer
2. Create a new folder
3. Right click the folder name and choose Properties

4. Click on the Sharing tab
5. Avoid the Network Wizard and make sure the boxes in front of Share this Folder and Allow Network Users to change are checked

Once you have created a shared folder, other computers in your workgroup can access it. Move to another computer on your LAN and repeat steps 1 and 2 (and step 3 if you like). Now you can use the shared folder:

1. Double click on My Network Places.
2. Double click on a shared folder
3. Create a file (e.g., using Word) and save it in your shared directory
4. Move the file(s) across computers in your workgroup

If you do this on your home network, anyone with access to your network can access the files in your shared folder. It is much safer to turn off file sharing unless you intentionally want to use it (see Step 2 and make sure the boxes are not checked if you want to prevent file sharing).

Windows Printer Sharing

In the same way you can share folders with other computers in your workgroup you can share printers. To share a printer, do the following on the computer that has the printer connected to it:

1. Go to Settings → Control Panel → Printers and Faxes
2. Right click on a printer and select Properties
3. Click on the Sharing tab
4. Click on Share This Printer

Once you have done this, you can move to other computers on your LAN and install the network on them:

1. Go to Settings → Control Panel → Printers and Faxes
2. Click on Add a Printer
3. In the Welcome to Add a Printer Wizard, click Next
4. Click the Radio Button in front of A Network Printer and click Next
5. Click the Radio Button in front of Browse for a Printer and click Next
6. Select the Network Printer and click Next
7. You can make this printer your default printer or not, and click Next