COMPUTER NETWORKING PRIMER: TECHNOLOGY TODAY

This chapter explains basic computer networking concepts and technology and also introduces computer networking terminology to help you understand the uses and benefits of Novell products. We believe this primer provides a fairly thorough foundation for understanding any computer networking–related discussion. Understanding computer networking technology illuminates its potential. It also enables one to understand the viability of global networking, which will eventually provide any computer user a digital link to all other computer users and to digital information and services worldwide.

What Is a Computer Network?

On the most basic level, a computer network is a collection of devices that can store and manipulate electronic data, interconnected in such a way that network users can store, retrieve, and share information.

Commonly connected devices include microcomputers, minicomputers, mainframe computers, terminals, printers, fax machines, pagers, and various data storage devices. In the near future, numerous other types of devices will be network connectable, including interactive TVs, videophones, and navigational and environmental control systems. Eventually, devices everywhere will give you two-way access to a vast array of resources on a global computer network.

In today’s business world, a computer network is much more than a collection of interconnected devices. For many businesses, the computer network is the resource that enables them to gather, analyze, organize, and disseminate the information that is essential to their profitability. The rise of intranets and extranets—a recent development in computer networking—is the latest indication of the crucial importance of computer networking to businesses. Intranets and extranets are private business networks that are based on Internet technology. Intranets, extranets, and the Internet will be treated in more detail in a later section. For now, it is enough to understand that businesses are currently implementing intranets at a breakneck pace and for one reason only—an intranet enables a business to collect, manage, and disseminate information more quickly and easily than ever before. Many businesses are implementing intranets simply to remain competitive; businesses that delay are likely to see their competition outdistance them.

What Are the Benefits of Computer Networking?

The most obvious benefit of computer networking is that you can store and retrieve virtually any kind of information on a computer network, including textual information such as letters and contracts, audio information such as voice messages, and visual images such as facsimiles, photographs, medical x-rays, and even video.

In addition to information storage and retrieval, there is a host of other important benefits of networking computers. Having a computer network enables you to combine the skills of different people and the power of different equipment, regardless of the physical locations of the people or the equipment. And computer networking enables people to easily share information, allowing them to work more securely, efficiently, and productively.

Powerful, Flexible Collaboration

A well-designed computer network enables users to collaborate effectively.

For example, a managing editor, associate editors, writers, and artists may need to work together on a publication. With a computer network, they can share the same electronic files, each from his or her own computer, without copying or transferring files. If the applications they are using feature even basic integration with the network operating system, they can perform such tasks as opening, viewing, and printing the same file simultaneously.

Using applications that are designed to take full advantage of network capabilities and services, network users can collaborate with ease and speed. For example, users can engage in real-time teleconferencing, talking face-to-face, while simultaneously viewing and editing the same document, adding and deleting notes and comments, and instantaneously viewing each other’s changes as they are made. And, they can do this without having to worry about accidentally changing or deleting the work of others.

To be able to collaborate electronically from widely separate physical locations has significant advantages. It enables people to avoid the considerable time investments and costs connected with traveling. It enables people to communicate instantaneously, regardless of the distance, and to act before their competitors do. It frees people from having to reconcile the differences in multiple information files. Electronic collaboration enables people to minimize the amount of work required to complete projects—it frees them from redoing work they would do correctly in the first place if they had instantaneous access to up-to-date information and instructions.

Freedom to Choose the Right Tool

If you choose an open networking environment, this adds another dimension to the information-sharing capabilities inherent in computer networking. Open networking products enable users to work on the type of computer best suited to the job they must do, without placing restrictions on their file-sharing capabilities.

The design of any particular computer can make it well suited for some tasks and not as well suited for others. In an open environment, you can combine many kinds of computers to take advantage of the special strengths of each type of machine. For example, Novell network users can use IBM PCs running any version of Windows or DOS, Macintosh computers running a version of the Macintosh operating system, Sun workstations running the UNIX operating system, and many other types of computers, all on the same network. Scientists, secretaries, doctors, lawyers, writers, editors, artists, engineers—everyone can use the type of computer equipment best suited to the type of work he or she does, yet each can still easily share information with everyone else.

Cost-Effective Resource Sharing

A very important reason for having a computer network is that it enables users to share expensive equipment.

Equipment sharing has significant benefits. It enables you to buy equipment with features that you wouldn’t otherwise be able to afford and to ensure that the equipment is used to its full potential. A correctly implemented network can result in both increased productivity and lower equipment costs.

For example, suppose you had a number of unconnected computers. People using these computers would not be able to print unless you purchased a printer for each computer or unless users manually transferred files from computers without printers to those with printers. In choosing between these alternatives, you would be choosing between significant expenses for hardware or significant expenses for labor.

But networking the computers would give you other alternatives. Because all users could share any networked printer, you would not need to buy a printer for every computer. Therefore, rather than buying numerous inexpensive printers, none of which had top-end productivity features and all of which would sit idle most of the time, you could buy a few inexpensive printers and a few printers with top-end productivity features. The more powerful printers might be able to print 20 times more pages per minute than the inexpensive printers. And, the more powerful printers might also be able to print in color and to sort, staple, or bind any number of pages, and to produce large numbers of completed documents.

On a Novell network, all users could share the various printers, accessing whichever printer was most appropriate for the job they were doing. The network software would enable users to print whenever they wanted. The network would print documents in the order they were received, on the printer the user selected. Whenever necessary, users would be able to change the order in which documents were to be printed and where they were to be printed.

By selecting the right mix of printers and allowing each network user appropriate access to them, you could have enough printing power to take care of the needs of all users; you could ensure that expensive equipment was not standing idle; and you could provide users with the latest, most powerful productivity features, freeing them from many tasks they would otherwise have to do manually—all for a significantly lower cost than if you were to buy an inexpensive printer for each of the computers connected to your network.

A network enables you to share any networkable equipment or software and realize the same benefits that you would enjoy from sharing printers. On a network, users can share modems; data storage devices, such as hard disks and CD-ROM drives; data backup devices, such as tape drives; E-mail systems; facsimile machines; and all networkable software. When you compare sharing these resources to purchasing them for each computer, the cost savings can be enormous.

When you implement an intranet, you can share network resources with suppliers, consultants, and other outside partners. Soon, you will be able to allow your employees to rent applications over the Internet. Businesses have just begun to explore the possibilities of intranet resource sharing.

Secure Management of Sensitive Information

There is another advantage to computer networking that may be even more important than instantaneous, coordinated information and resource sharing. The best networks have extremely powerful security features that enable you to exercise flexible control of who will have access to sensitive data, equipment, and other resources.

Effective Worldwide Communications

If you choose a networking company that offers a full suite of products—including robust directory services—and that supports open standards, you will be able to connect heterogeneous computing equipment at distant geographic locations into one cohesive network. As a result, you will be able to disseminate critical information to multiple locations anywhere in the world, almost instantaneously.

Easy, Immediate Information Dissemination

When you implement a business intranet, you can create or update information and easily and immediately make it accessible to all company employees. With a World Wide Web server running on your intranet and with today’s powerful Web publishing tools, you can create or change any information using a favorite, familiar application, and you can have that information automatically and instantaneously published on your Web server. This information will then be available to anyone who has rights to access it, anywhere in the world.

Worldwide, Instantaneous Access to Information

With access to your business's intranet and Web server, your employees will be able to easily and inexpensively access any new or updated information, from anywhere in the world, within a few seconds after it is published. The Internet provides the low-cost backbone for global access to your intranet, and existing Web browsers and other intranet tools make it easy for even the most novice computer user to access the information and intranet resources they need.

Integrated, flexible information sharing; instantaneous information updating and access; lower equipment costs; flexible use of computing power; secure management of sensitive information—these are the benefits of computer networking. And these benefits help us produce the results we are all looking for: increased efficiency, productivity, and profitability.

Hardware Technology

The purpose of computer networks is to enable users to manipulate data so it can be stored, retrieved, and shared. To understand how available technology enables us to do this, we need to define a few terms and understand some basic concepts.

Definition: Data Versus Information

Although we routinely use the terms “data” and “information” interchangeably, they are not technically the same thing.

Data are entities that convey meaning. Computer data is stored as a series of electrical charges arranged in patterns to represent information. In other words, data refers to the form of the information (the electrical patterns). It is not the information itself.

Information means decoded data, in human-readable form. In other words, information is the real-world, useful form of data. For example, the data in an electronic file can be decoded and displayed on a computer screen or printed onto paper as a business letter.

Encoding and Decoding Data

To store meaningful information as data and to retrieve the information, we use an encoding scheme—we agree on a series of electrical patterns that will represent each of the discrete pieces of information we want to store and retrieve. For example, we agree that a particular series of electrical patterns will represent the alphabetic character “A.” There are many encoding schemes in use. One common data-encoding scheme is ASCII code; an appendix at the end of this chapter gives a description of the ASCII data code.

To encode information into data and later decode that data back into information, we use electronic devices (one of which is the computer) that generate electronic signals. Signals are simply the electric or electromagnetic encoding of data. Various components in a computer enable it to generate signals to perform the encoding and decoding task.

How Network Data Is Transferred

At the most general, conceptual level, to have a computer network we need only three components. First, we need a data code. Next, we need computers and auxiliary devices to encode information into data signals and to decode data signals back into information. Finally, we need the means to transfer the signals between the computer devices.

Let's suppose we've agreed on a coding scheme and we have several computers that are all capable of encoding and decoding the information we want to save.

Now, to have a computer network, we need only the means of transferring the generated signals between the computers. To transfer signals between the computers, we need two things: (1) a transmission medium to carry the signals and (2) devices to propagate (send) and receive the signals.

Network Transmission Media

Electrical signals are generated as electromagnetic waves (analog signaling) or as a sequence of voltage pulses (digital signaling). To be propagated (sent) from one location to another, a signal must travel along a physical path. The physical path that is used to carry a signal between a signal transmitter and a signal receiver is called the "transmission medium."

There are two types of transmission media: guided media and unguided media.

Guided media are manufactured so that signals will be confined to a narrow path and will behave predictably. Commonly used guided media include twisted-pair wiring, similar to common telephone wiring; coaxial cable, similar to that used for cable TV; and optical fiber cable.

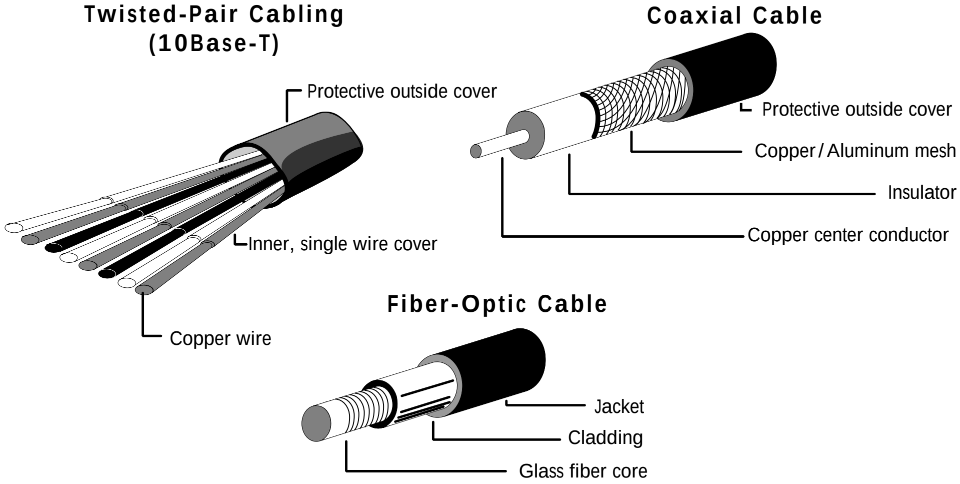

Figure 1

Common guided transmission media

Unguided media are natural parts of the existing environment that can be used as physical paths to carry electrical signals. Earth's atmosphere and outer space are examples of unguided media that are commonly used to carry signals. These media can carry such electromagnetic signals as microwave and infrared light waves.

Regardless of the type of medium, the network signal is transmitted through it as some kind of waveform. When transmitted through wire and cable, the signal is an electrical waveform. When transmitted through fiber-optic cable, the signal is a light wave, somewhere in the spectrum of visible or infrared light. When transmitted through Earth's atmosphere or outer space, the signal can take the form of waves in the radio spectrum, including VHF and microwaves, or it can be light waves, including infrared or visible light (for example, lasers).

When planning a computer network, designers choose a transmission medium, or a combination of media, based on the physical circumstances involved in building the network and the reliability and data-handling performance required of the network. The objective is to keep costs to a minimum yet provide all parts of the network with the required reliability and performance.

For example, if you needed to build a network consisting of two subnetworks located in separate buildings several miles apart, you might use two or more transmission media. If you did not require the same level of performance on both subnetworks, you might use a different type of wire or cable as the transmission medium on each.

To connect the two subnetworks across town and ensure a reliable connection even in rain and fog, you might use a third medium, Earth's atmosphere, and connect the subnetworks through a microwave link. Or, you might use a T1 or T3 connection. T1 and T3 are dedicated lines (basically special telephone lines) that support high-speed communications. They can be leased from private companies that specialize in providing communication services.

Transmitting and Receiving Devices

Once you have a transmission medium, you need devices that can propagate signals across the medium and devices that can receive the signals when they reach the other end of the medium. There are a number of such devices used in computer networking. Such devices are designed to propagate a particular type of signal across a particular type of transmission medium. Transmitting and receiving devices used in computer networks include network adapters, repeaters, wiring concentrators, hubs, switches, and infrared, microwave, and other radio-band transmitters and receivers.

Network Adapters

A network adapter is the hardware installed in computers that enables them to communicate on a network. Network adapters are manufactured in a variety of forms. The most common form is the printed circuit board, which is designed to be installed directly into a standard expansion slot inside a microcomputer. Other network adapters are designed for mobile computing. They are small and lightweight and can be connected to standard connectors on the back of portable (laptop and notebook) computers so that the computer and network adapter can be easily transported from network to network. Network adapters are now being built into many computers, especially portable and laptop computers.

Network adapters are manufactured for connection to virtually any type of guided medium, including twisted-pair wire, coaxial cable, and fiber-optic cable. They are also manufactured for connection to devices that transmit and receive visible light, infrared light, and radio microwaves, to enable wireless networking across the unguided media of Earth’s atmosphere and outer space.

The connection hardware used to make connections between network adapters and different transmission media depends on the type of medium used. For example, twist-on BNC connectors are commonly used for connection to coaxial cable, while snap-in telephone-type jacks are ordinarily used for connection to twisted-pair wiring. Figure 2 shows two different types of network adapters connected to different computers and media, using different types of connectors.

Figure 2

Network adapters are manufactured in a variety of forms, for virtually every kind of communication medium.

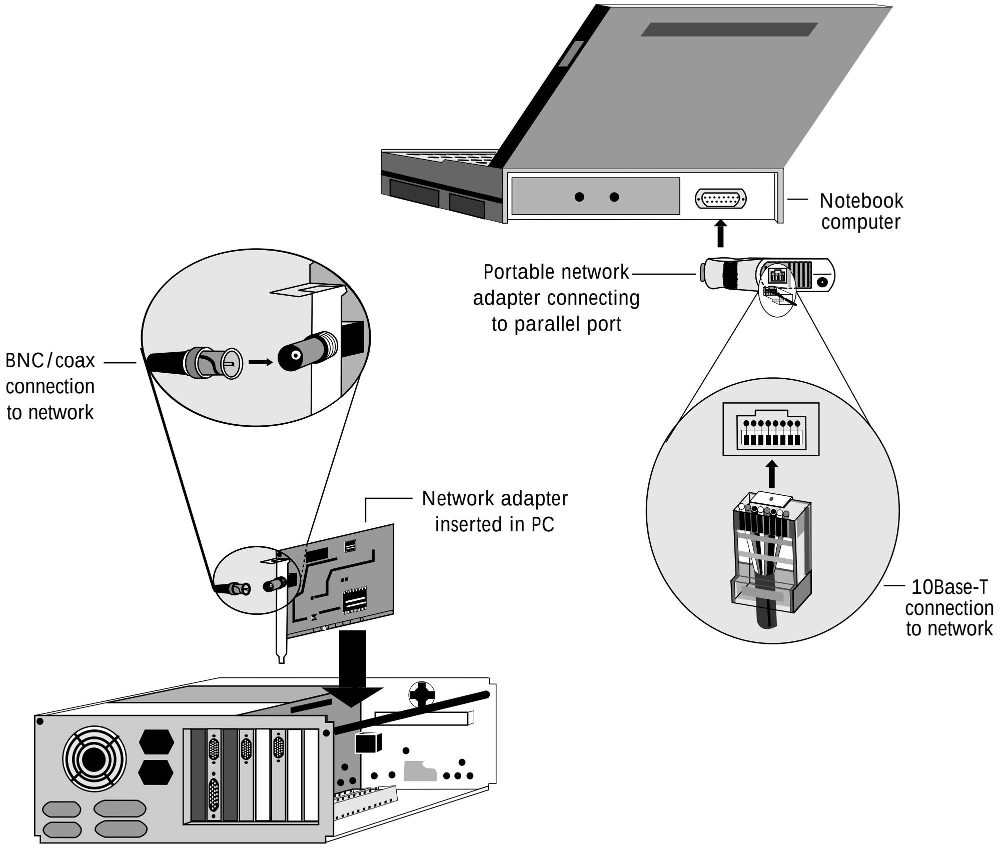
Notebook computer
Portable network adapter connecting to parallel port
BNC/coax connection to network
Network adapter inserted in PC
10Base-T connection to network

Repeaters

Repeaters are used to increase the distance over which a network signal can be propagated.

As a signal travels through a transmission medium, it encounters resistance and gradually becomes weak and distorted. The technical term for this signal weakening is “attenuation.” All signals attenuate, and at some point they become too weak and distorted to be reliably received. Repeaters are used to overcome this problem.

A simple, dedicated repeater is a device that receives the network signal and retransmits it at the original transmission strength. Repeaters are placed between other transmitting and receiving devices on the transmission medium, at a point where the signal will not have attenuated too much to be reliably received.

In today’s networks, dedicated repeaters are seldom used. Repeating capabilities are built into other, more complex networking devices. For example, virtually all modern network adapters, hubs, and switches incorporate repeating capabilities.

Wiring Concentrators, Hubs, and Switches

Wiring concentrators, hubs, and switches provide a common physical connection point for computing devices. (We limit this discussion to devices used for making physical connections. The term “concentrator” can mean something different in a mainframe or minicomputer environment.) Most hubs and all wiring concentrators and switches have built-in signal repeating capability and thus perform signal repair and retransmission (along with other complex functions).

In most cases, hubs, wiring concentrators, and switches are proprietary, standalone hardware. There are a number of companies that manufacture such equipment. Occasionally, hub technology consists of hub cards and software that work together in a standard computer.

Figure 3 shows two common hardware-based connection devices: a token-ring switch and an Ethernet 10Base-T concentrator.

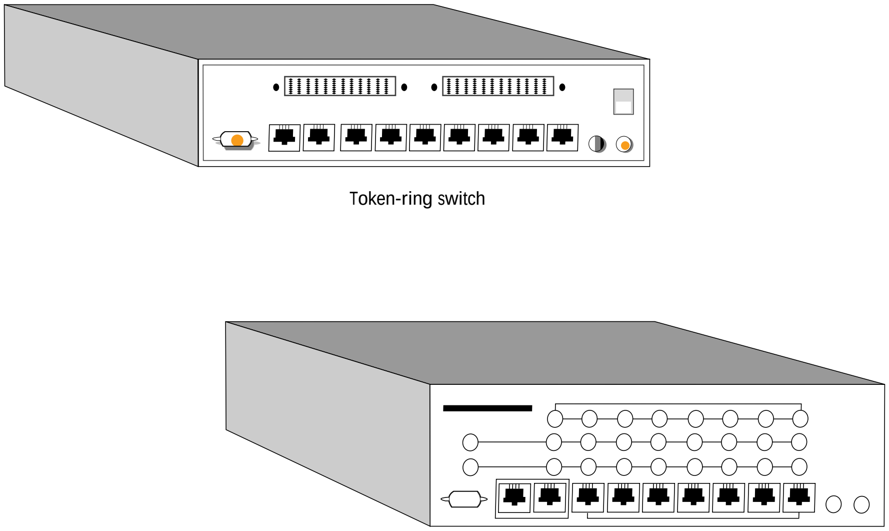

Token-ring switch

10Base-T concentrator

Figure 3

Token-ring switch and Ethernet 10Base-T concentrator

Microwave Transmitters

Microwave transmitters and receivers, especially satellite systems, are commonly used to transmit network signals over great distances. A microwave transmitter uses the atmosphere or outer space as the transmission medium to send the signal to a microwave receiver. The microwave receiver either relays the signal to another microwave transmitter, which sends it to another microwave receiver, or the receiving station translates the signal to some other form, such as digital impulses, and sends it along on some other suitable medium. Figure 4 shows a satellite microwave link.

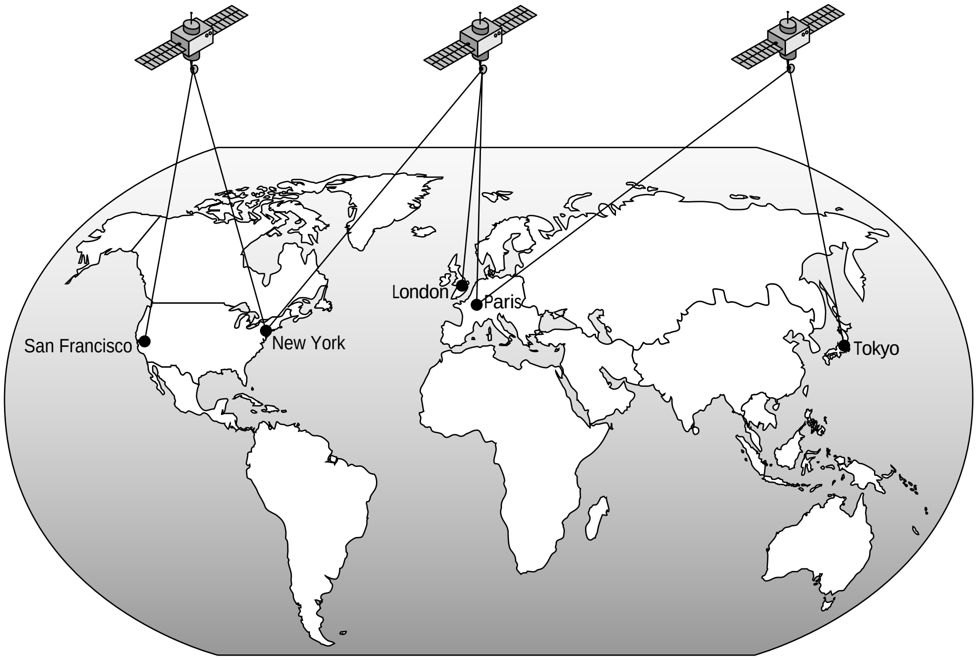

Figure 4

Satellite microwave link

Infrared and Laser Transmitters

Infrared and laser transmitters are similar to microwave systems. They use the atmosphere and outer space as the transmission media and require a line-of-sight transmission path. The major difference is that they transmit light waves rather than radio waves. Infrared and laser transmissions are useful for signaling across short distances where it is impractical to lay cable—for instance, when networks are at sites a few miles apart. Because infrared and laser signals are in the light spectrum, rain, fog, and other environmental factors can cause transmission problems.

Modems

Modems convert digital (computer) signals to analog (audio) signals, and vice versa, by modulating and demodulating a carrier frequency. The most common modems transmit and receive data across ordinary voice-grade telephone lines.

A transmitting modem converts (modulates) the encoded data signal to an audible signal and transmits it. A modem connected at the other end of the line listens to the audible signal and converts it back into a digital signal (demodulates it) for the computer on the receiving end of the communication link. Modems are commonly used for inexpensive, intermittent communications between geographically isolated computers and a main network.

Connecting Network Devices

Now that you understand what a transmission medium is, we will discuss ways in which devices are connected, both physically and electronically, so that they can communicate with each other.

First, we will look at how guided transmission media are commonly connected physically to form the physical topology of a local area network. Then, we will examine three logical topologies, the electronic scheme used to connect network devices.

Common Physical Topologies: Bus, Star, and Star-Wired Ring

Physical topology is applicable to networks in which the devices are connected with some type of guided transmission medium (or media). Physical topology is the physical layout of the guided transmission media. The most common physical topologies are the bus, the star, and the star-wired ring.

Physical Bus

The simplest form of a physical bus topology consists of a trunk (main) cable with only two end points. When the trunk cable is installed, it is run from area to area and device to device—close enough to each device so that all devices can be connected to it with short drop cables and T-connectors. This simple “one wire, two ends” physical bus topology is illustrated in Figure 5.

Figure 5

Physical bus topology

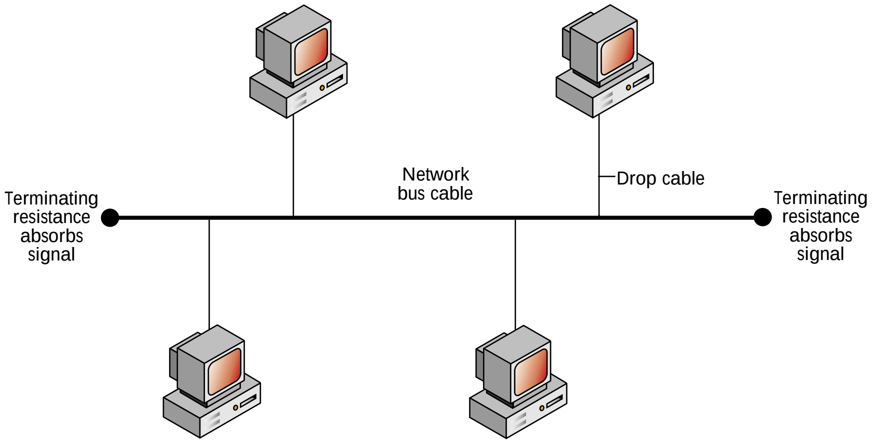
Network bus cable
Drop cable
Terminating resistance absorbs signal
Terminating resistance absorbs signal

A more complex form of the physical bus topology is the distributed bus (also called the tree topology). In the distributed bus, the trunk cable starts at what is called a “root,” or “head end,” and branches at various points along the way. (Thus, unlike the simple bus topology described above, this variation uses a trunk cable with more than two end points.) Where the trunk cable branches, the division is made by means of a simple connector (as opposed to the star physical topology discussed below, where connections are made to a central, somewhat sophisticated connection device). The distributed bus topology is illustrated in Figure 6.

Figure 6

Distributed bus topology

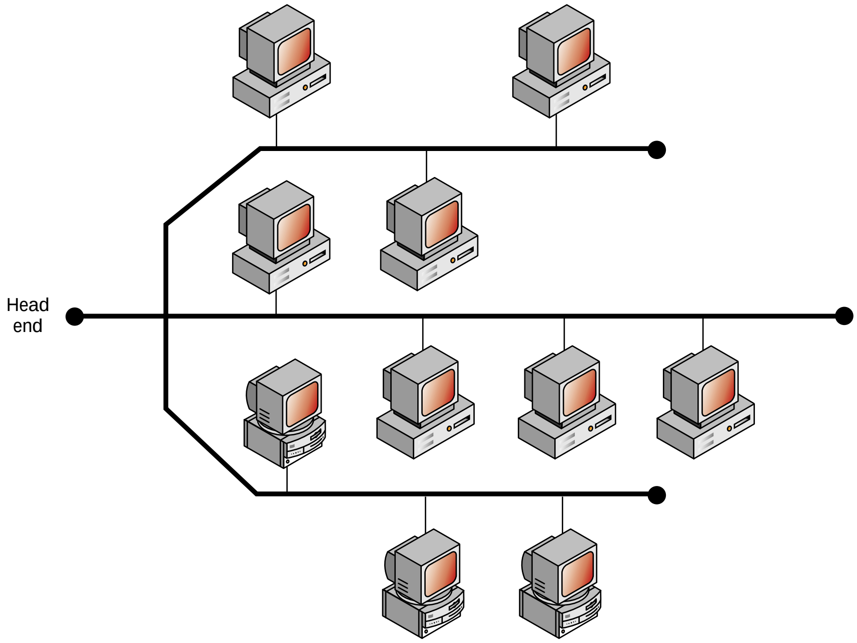

Physical Star

The simplest form of the physical star topology consists of multiple cables—one for each network device—attached to a single, central connection device. For example, 10Base-T Ethernet networks are based on a physical star topology—each network device is attached to a 10Base-T hub by means of twisted-pair cable.

In a real-life implementation of even a simple physical star topology, the actual layout of the transmission media need not form a recognizable star pattern; the only required physical characteristic is that each network device be connected by its own cable to the central connection point.

The simplest form of the physical star topology is illustrated in Figure 7.

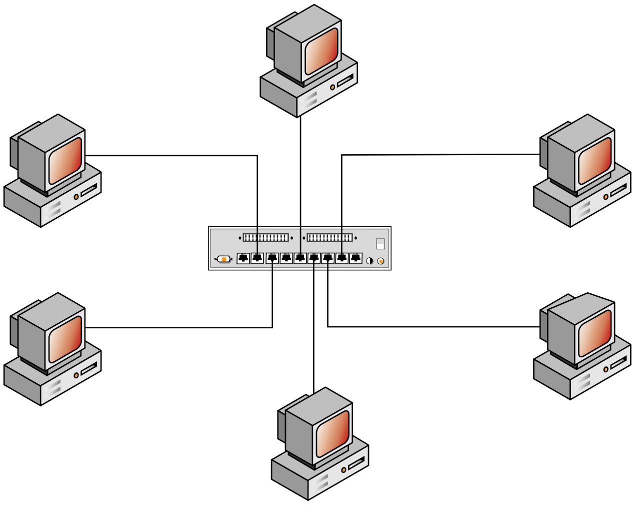

Figure 7

Physical star topology

A more complex form of the physical star topology is the distributed star. In this topology, there are multiple central connection points, which are all connected to form a string of stars. This topology is illustrated in Figure 8.

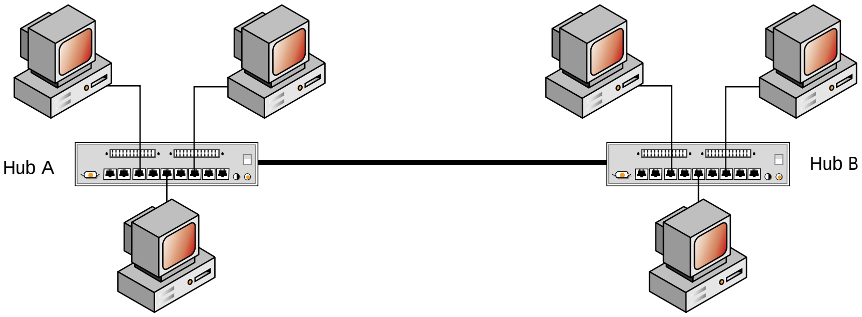

Figure 8

Distributed star topology

Physical Star-Wired Ring

In the star-wired ring physical topology, individual devices are connected to a central hub, as they are in a star or distributed star network. However, within each hub the physical connections form a ring. Where multiple hubs are used, the ring in each hub is opened, leaving two ends. Each open end is connected to an open end in some other hub (each to a different hub) so that the entire network cable forms one physical ring. This physical topology, which is used in IBM's Token-Ring network, is illustrated in Figure 9.

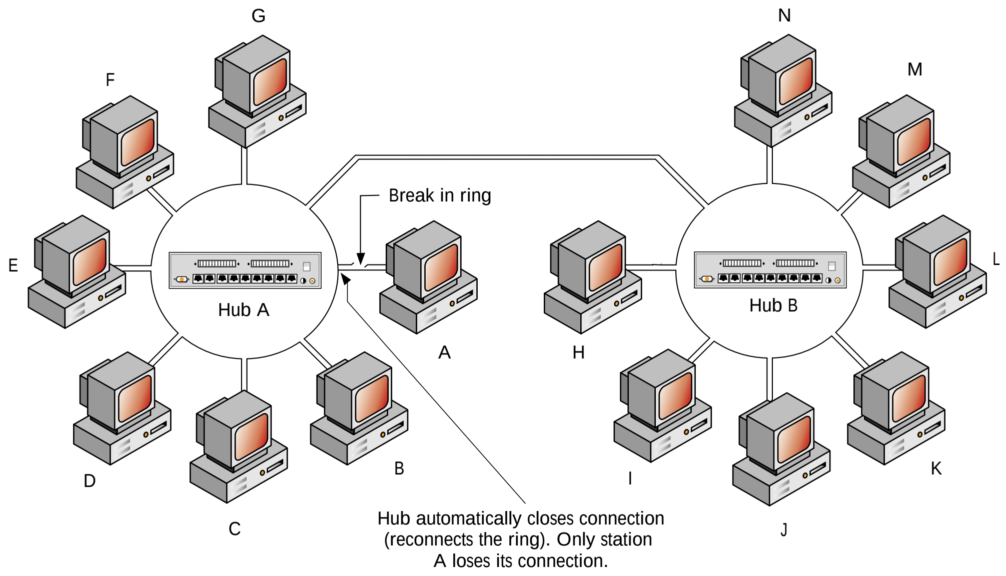
G
N
F
M
Break in ring
E
Hub A
Hub B
L
A
H
D
C
B
I
J
K
Hub automatically closes connection (reconnects the ring). Only station A loses its connection.

Figure 9

Physical star-wired ring topology

In the star-wired ring physical topology, the hubs are "intelligent." If the physical ring is somehow broken, each hub is able to close the physical circuit at any point in its internal ring so that the ring is restored. Refer to details shown in Figure 9, hub A, to see how this works.

Currently, the star topology and its derivatives are most preferred by network designers and installers because using these topologies makes it simple to add network devices anywhere. In most cases, you can simply install one new cable between the central connection point and the desired location of the new network device, without moving or adding to a trunk cable or making the network unavailable for use by other stations.

Common Logical Topologies: Bus, Ring, and Star (Switching)

While the star, distributed star, and star-wired ring are currently the most commonly used physical topologies, there is considerably more variety (and parity) in the use of logical topologies. A logical topology is the electronic scheme used to enable network devices to transmit and receive data across the transmission media without interfering with each other.

There are three basic logical topologies, each of which has distinct advantages in specific situations. As you study the figures representing these topologies, remember that the figures represent a logical (electronic), not a physical, connection scheme.

Logical Bus

In the logical bus topology, transmissions (called frames) are broadcast simultaneously in every direction to every point on the transmission media. Every network station checks each frame to determine whether the frame is intended for it. When the signal reaches any end point on the transmission media, it is absorbed (removed from the media) by appropriate electronics. Removing the signal prevents it from being reflected back along the transmission media and interfering with subsequent transmissions.

On a logical bus network, the transmission media are shared. To prevent transmission interference, only one station may transmit at a time. Thus, there must be a method for determining when each station is allowed to use the media. This method is called the media access control (MAC).

The media access control method most commonly used for a logical bus network is a contention method called “carrier sense multiple access with collision detection (CSMA/CD).” This media access control method is similar to the access scheme used on a telephone party line. When any station wants to send a transmission, it “listens” (carrier sense) to determine if another station is currently transmitting on the media. If another station is transmitting, the station that wants to transmit waits. When the media become free, the waiting station transmits. If two or more stations determine that the media are free and transmit simultaneously, there is a “collision.” All transmitting stations detect the collision, transmit a brief signal to inform all other stations there has been a collision, and all stations then wait a random amount of time before attempting to transmit.

A logical bus network may also use token passing for media access control. In this MAC method, each network station is assigned a logical position in an ordered sequence, with the last number of the sequence pointing back to the first (the logical order that the stations are assigned need not correspond with any physical order). A control frame, called a “token,” is used to control which station can use the media. A station can transmit only when in possession of the token. Furthermore, a station can have the token only a limited time before it must pass the token to the next station. The token starts at the first station in the predefined logical order. While the first station has the token, it transmits, polls stations, and receives responses (gives other stations permission to use the media) until the allotted time expires; or, it passes the token when it no longer needs control of the media, whichever happens first. The first station passes the token to the second station in the logical sequence. This token passing (in sequence) continues nonstop while the network is running—thus, every station gets equitable access to the transmission media.

The logical bus transmission scheme is used in combination with both the physical bus and physical star topology, and the MAC method can vary in different cases. For example, the cable on thin Ethernet networks is laid out as a physical bus and the transmission scheme is a logical bus, but the cable on 10Base-T Ethernet networks and on ARCnet networks is laid out as a physical star, although both use the logical bus transmission scheme. And thin Ethernet (physical bus) and 10Base-T Ethernet (physical star) both use the CSMA/CD MAC method, but ARCnet (physical star) uses token passing as its MAC method.

Figures 10 shows a thin Ethernet network (physical bus, logical bus), and Figure 11 shows a 10Base-T Ethernet network (physical star, logical bus). In both figures, notice that the network signal (shown by the arrows) emanates from the sending station and travels in all directions, to all parts of the transmission media (the determining criterion for a logical bus topology).

Figure 10

Thin Ethernet network (physical bus, logical bus)

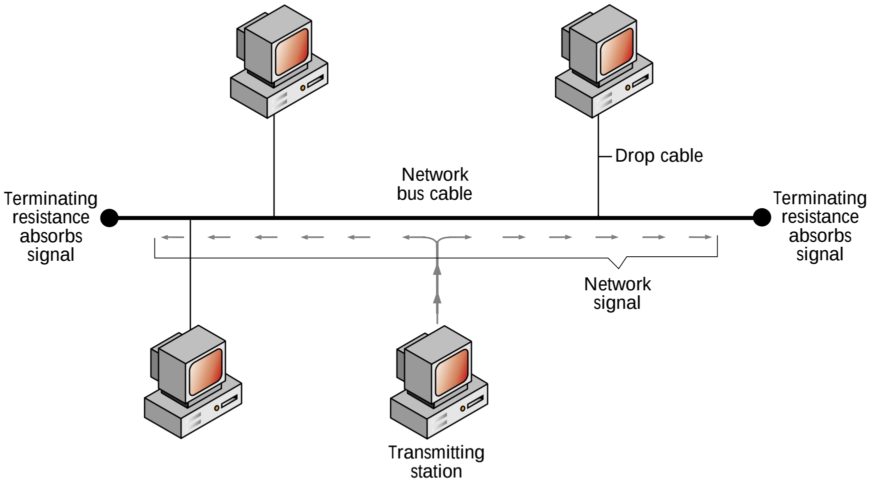
Drop cable
Network bus cable
Terminating resistance absorbs signal
Terminating resistance absorbs signal
Network signal
Transmitting station

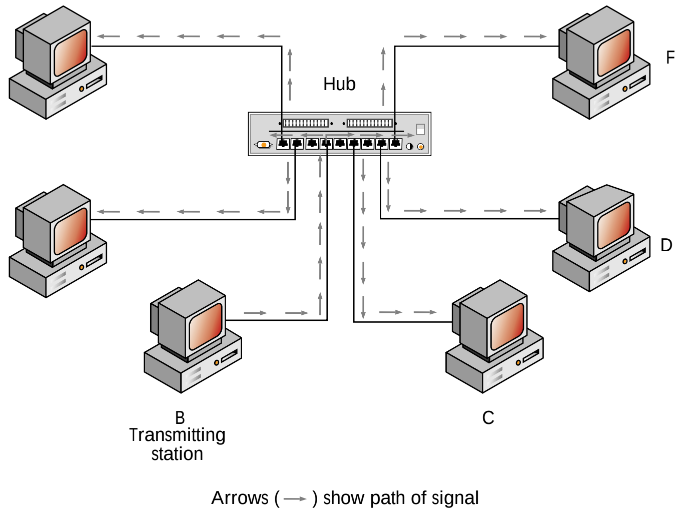

Figure 11

10Base-T Ethernet network (physical star, logical bus)

Logical Ring

In the logical ring topology, frames are transmitted in one direction, around a physical ring, until they have passed every point on the transmission media. (The logical ring must be used in combination with a physical ring topology, such as the star-wired ring explained earlier.) Each station on the physical ring receives the signal from the station before it and repeats the signal for the next station. When a station transmits data, it gives the data the address of some other station on the ring. The data is circulated around the ring through each station's repeater until it reaches the station to which it is addressed and is copied. The receiving station adds an acknowledgment of receipt to the frame. The frame continues on around the ring until it returns to the station from which it was originally transmitted, which reads the acknowledgment and removes the signal from the ring. Figure 12 shows how data would flow on a logical ring network with a star-wired ring physical topology.

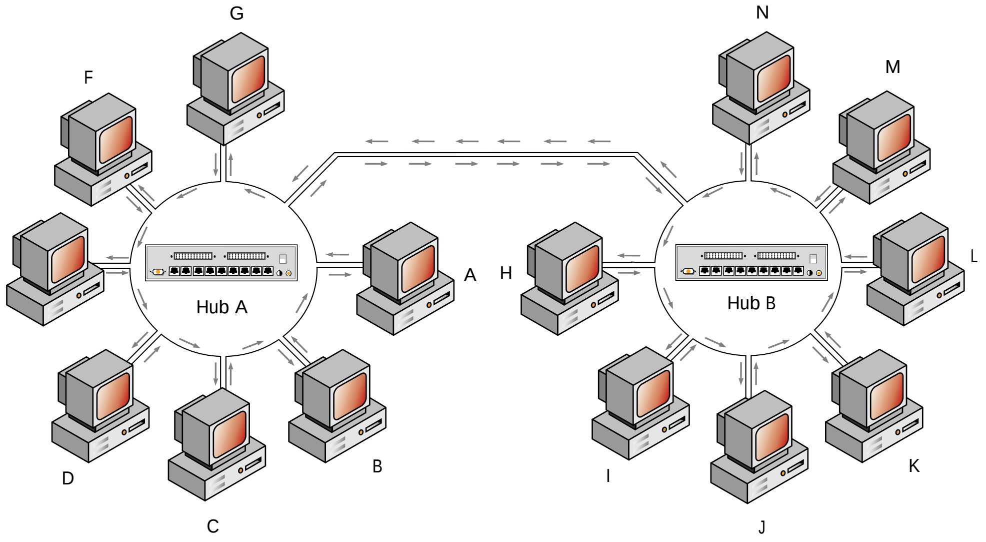

Figure 12

Logical ring topology

Media access control for the logical ring topology is almost always based on a form of token passing, the basics of which are described in the logical bus topology section. (Stations are not necessarily granted media access in the same order in which they receive frames on the physical ring.) IBM's Token-Ring network is a logical ring network based on the star-wired ring physical topology.

Logical Star (Switching)

In the logical star topology, network switches are used to restrict transmissions to a specific part of the transmission medium (transmission path restriction is the identifying characteristic of a logical star).

In its pure form, switching provides a dedicated line for each end station. This means that when one station transmits a signal to another station on the same switch, the switch transmits the signal only on the two paths connecting the sending and receiving station. Figure 13 shows how data would be transmitted from one station to another if two stations were directly connected to the same switch.

Figure 13

Switching

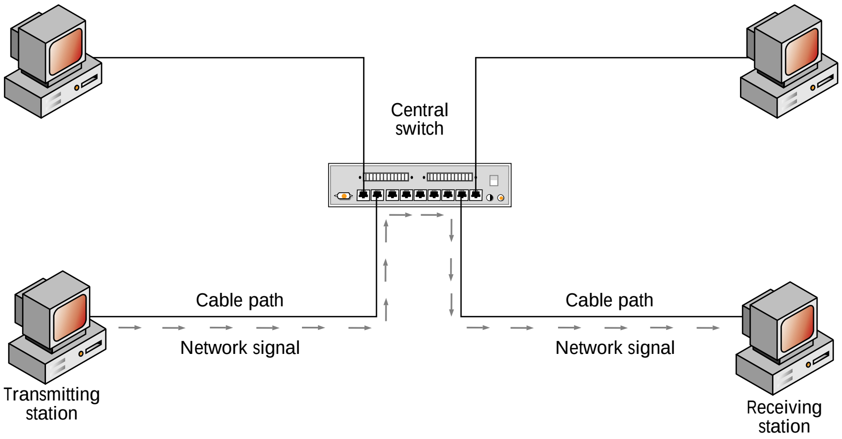

Most switching technology adds switching capability to existing connection standards, incorporating the logical connection schemes (including the media access control methods) of the existing standards.

For example, a 10Base-T Ethernet switch supports the Ethernet CSMA/CD media access control method. Some switches are designed to support and combine multiple network standards. For example, a switch might contain both 10Base-T Ethernet ports and Fiber Distributed Data Interface (FDDI) ports. In this case, the switch would support the logical connection scheme for both standards, including the Ethernet CSMA/CD and the FDDI token-ring MAC methods.

Switches have built-in connection logic and significant amounts of fast memory. This enables them to simultaneously service all connected stations at full access speed. Thus, when you connect a station directly to a switch, you can increase the total throughput of your network—a significant performance advantage.

Switching illustrates well that a logical topology consists of the total of the various aspects of the electronic connection scheme, not just the MAC method. By combining new (switching) capabilities with existing logical connection schemes, engineers create a new logical topology.

Switching can be distributed (multiple switches can be connected using one or more physical topologies). Switches can be used not only to connect individual stations, but also to connect network segments (groups of stations). Thus, in many circumstances, switching can be used to improve the performance of your network.

Connecting a Simple Network

Now that we've seen the hardware pieces that make up a network and discussed the difference between physical and logical topology, let's connect some hardware to form a simple network. Figure 14 shows some of the hardware items we have discussed, connected to form a very basic computer network.

Figure 14

Various networking hardware connected to form a simple network

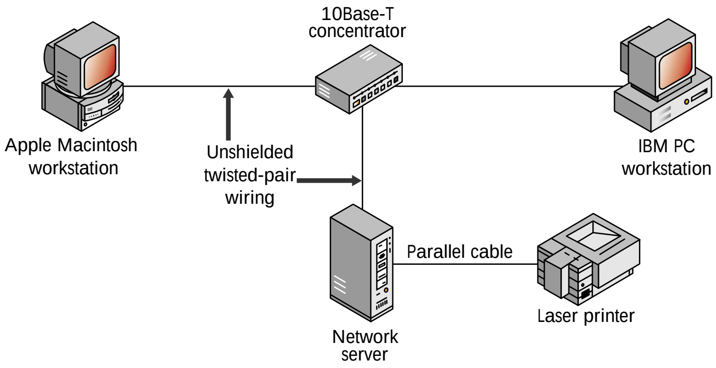

The network in this illustration includes the following components: three computers connected through a 10Base-T concentrator by means of unshielded twisted-pair wiring; three Ethernet 10Base-T network adapters, one installed inside each of the computers; and a laser printer that is connected to one of the computers.

The computer at the bottom center of the illustration is a network server; it controls the network (details will be covered in a following section). The other two computers are workstations. The workstations use the network under the control of the network server. One workstation is an IBM PC and the other is an Apple Macintosh computer.

The 10Base-T concentrator serves as a common connection point for the three computers; it repeats network signals.

The lines between the different components of the network represent the transmission medium, which is twisted-pair wiring. As you may remember from our recent discussion of topologies, this 10Base-T network is connected in a physical star, but it is based on a logical bus that uses a contention scheme as the means for workstations to get access to the transmission medium.

The printer in this network is connected directly to the server by means of a parallel interface cable, which is a standard connection method. The server accepts print jobs from either workstation and sends the jobs through the parallel interface cable to the printer. This is the simplest way to enable both workstations to use the printer. There are other ways to connect printers to a network, including attaching them to a computer set up as a dedicated print server or connecting them to a computer that runs special software enabling it to function as both a workstation and a print server. Many printers are now manufactured with an internal network adapter so that they can be attached directly to the transmission medium at any physical point in the network.

Controlling Data Transmission

Once you have the hardware we've discussed so far, you can start connecting the various pieces into a network.

But simply connecting hardware doesn't make a computer network. Even though the hardware is capable of generating signals and transmitting them across a medium, it must be told when and how to do this. There must be network communication software to tell the hardware when and how to transmit. The software and hardware on all parts of the network must work together to enable the transmission of data from one networked computer to another. We'll explore various networking software a little later. First, let's look at the communication model that is the basis for controlling data transmission on computer networks.

To guarantee reliable transmission of data, there must be an agreed upon method that governs how data is sent and received. For example, how does a sending computer indicate which computer it is sending data to? And, if the data will be passed through intervening devices, how are these devices to understand how to handle the data so that it will get to the intended destination? And, what if the sending and receiving computers use different data formats and data exchange conventions—how will data be translated to allow its exchange? These are only a few of the questions that must be answered before data can be reliably transmitted and received across a computer network.

Understanding the Open Systems Interconnection (OSI) model will allow you to understand how data can be transferred between two networked computers, regardless of whether they are on the same network, or are the same type of computer, or use the same data formats and exchange conventions.

ISO and the OSI Model

The OSI model was developed by the International Standards Organization (ISO) as a guideline for developing standards to enable the interconnection of dissimilar computing devices. It is important to understand that the OSI model is not itself a communication standard. In other words, it is not an agreed method that governs how data is sent and received; it is only a guideline for developing such standards.

The OSI Model: What It Is and Why It’s Important

It would be difficult to overstate the importance of the OSI model. Virtually all networking vendors and users understand how important it is that network computing products adhere to and fully support the networking standards the model has spawned. The reasons are logical.

First, when a vendor’s products adhere to the standards the OSI model has spawned, connecting those products to other vendors’ products is relatively simple. Conversely, the further a vendor departs from those standards, the more difficult it becomes to connect that vendor’s products to those of other vendors. Second, if a vendor were to depart from the communication standards the model has spawned, software development efforts would be very difficult because the vendor would have to build every part of all necessary software, rather than often being able to build on the existing work of other vendors.

The first two problems give rise to a third significant problem for vendors: A vendor’s products become less marketable as they become more difficult to connect with other vendors’ products unless the introduction of the vendor’s products is well ahead of the introduction of other such products into the general marketplace.

Now, keeping in mind the purpose of the OSI model, let’s take a look at its structure.

The Seven Layers of the OSI Model

Because the task of controlling communications across a computer network is too complex to be defined by one standard, the ISO divided the task into seven subtasks. Thus, the OSI model contains seven layers, each named to correspond to one of the seven defined subtasks.

Each layer of the OSI model contains a logically grouped subset of the functions required for controlling network communications. The seven layers of the OSI model and the general purpose of each are shown in Figure 15.

Figure 15

The OSI model

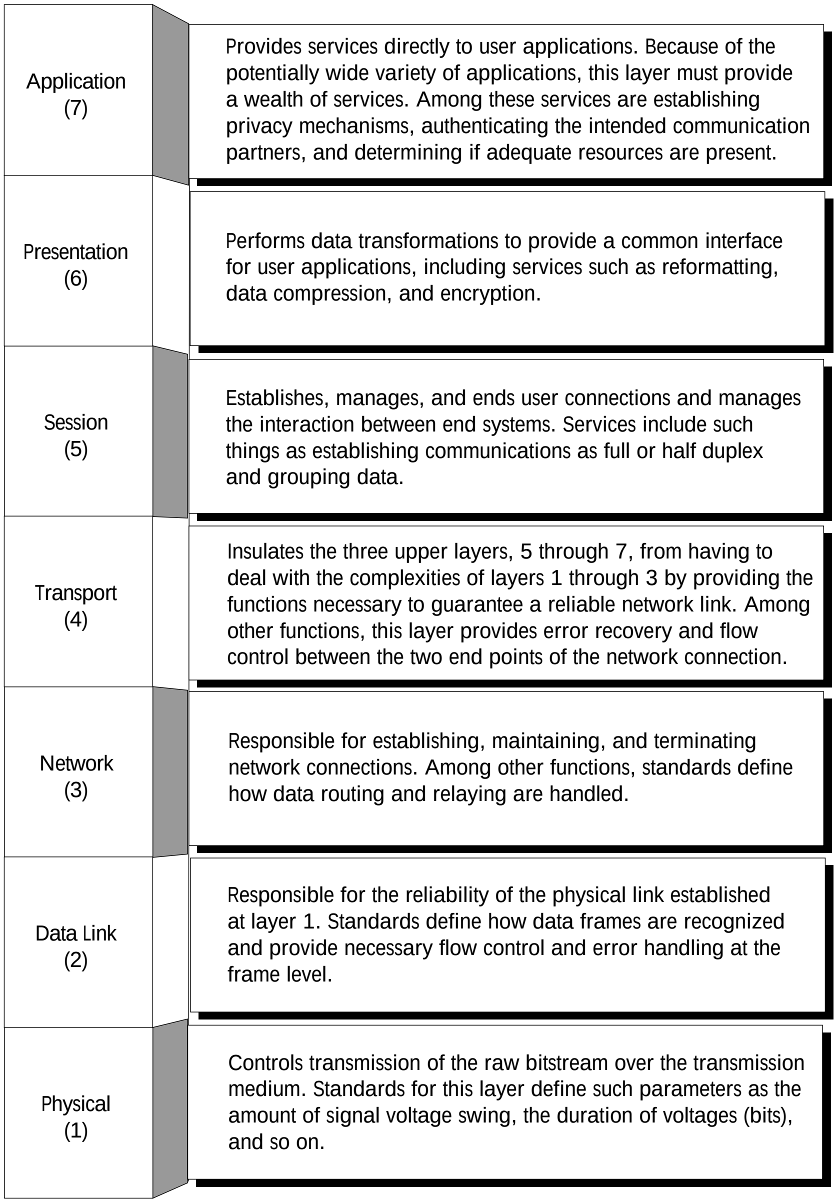

Application (7)

Provides services directly to user applications. Because of the potentially wide variety of applications, this layer must provide a wealth of services. Among these services are establishing privacy mechanisms, authenticating the intended communication partners, and determining if adequate resources are present.

Presentation (6)

Performs data transformations to provide a common interface for user applications, including services such as reformatting, data compression, and encryption.

Session (5)

Establishes, manages, and ends user connections and manages the interaction between end systems. Services include such things as establishing communications as full or half duplex and grouping data.

Transport (4)

Insulates the three upper layers, 5 through 7, from having to deal with the complexities of layers 1 through 3 by providing the functions necessary to guarantee a reliable network link. Among other functions, this layer provides error recovery and flow control between the two end points of the network connection.

Network (3)

Responsible for establishing, maintaining, and terminating network connections. Among other functions, standards define how data routing and relaying are handled.

Data Link (2)

Responsible for the reliability of the physical link established at layer 1. Standards define how data frames are recognized and provide necessary flow control and error handling at the frame level.

Physical (1)

Controls transmission of the raw bitstream over the transmission medium. Standards for this layer define such parameters as the amount of signal voltage swing, the duration of voltages (bits), and so on.

Standards and Protocols

National and international standards organizations have developed standards for each of the seven OSI layers. These standards define methods for controlling the communication functions of one or more layers of the OSI model and, if necessary, for interfacing those functions to the layer above and below.

A standard for any layer of the OSI model specifies the communication services to be provided and a protocol that will be used as a means to provide those services. A protocol is a set of rules network devices must follow (at any OSI layer) to communicate. A protocol consists of the control functions, the control codes, and the procedures necessary for successfully transferring data.

For every layer of the OSI model, there is more than one protocol standard. This is because a number of standards were proposed for each layer and because the various organizations that defined those standards—specifically, the standards committees inside these organizations—decided that more than one of the proposed standards had real merits. Thus, they allowed for the use of different standards to satisfy different networking needs.

Network Communications Through the OSI Model

Using the seven layers of the OSI model, we can explore more fully how data can be transferred between two networked computers. Figure 16 uses the OSI model to illustrate how such communications are accomplished.

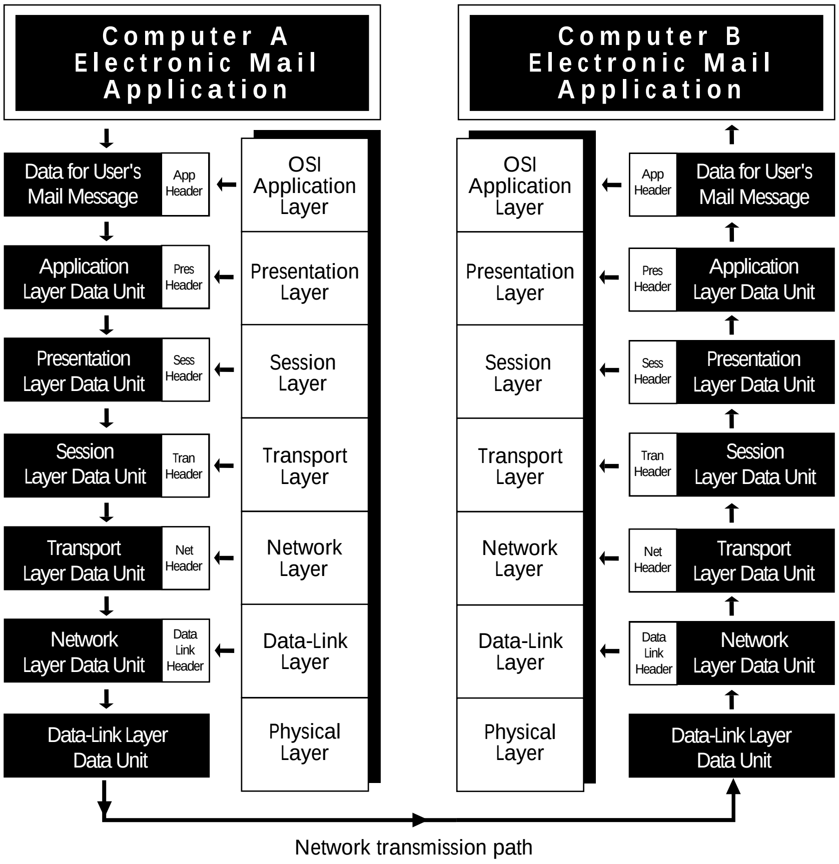
Computer A Electronic Mail Application
Computer B Electronic Mail Application
Data for User's Mail Message
App Header
Application Layer Data Unit
Pres Header
Presentation Layer Data Unit
Sess Header
Session Layer Data Unit
Tran Header
Transport Layer Data Unit
Net Header
Network Layer Data Unit
Data Link Header
Data-Link Layer Data Unit
OSI Application Layer
Presentation Layer
Session Layer
Transport Layer
Network Layer
Data-Link Layer
Physical Layer
Network transmission path

Figure 16

Networked computers communicating through the OSI model

Our figure represents two networked computers, each of which is running various pieces of software (most not shown). Running together, the various pieces of software implement the seven OSI layers. These computers are identical: They are running identical software, and they are using identical protocols at all OSI layers. Above the OSI application layer, each computer is running an E-mail program. The E-mail program enables the users of the two computers to exchange messages. Our figure represents the transmission of one brief message from computer A to computer B.

The transmission starts with the user of computer A pressing a key to send a mail message to the user of computer B. The E-mail application is designed to talk to the OSI application layer—it knows the proper protocol for doing so. The E-mail application transfers the message to the OSI application layer. Using the functions built into its protocol, the application layer accepts the message data and adds an application-layer header to it. The application-layer header contains the information necessary for the application layer in computer B to correctly handle the data when computer B receives it.

After adding its header, the application layer in computer A passes the data to the presentation layer below. The presentation layer treats everything received as data, including the application-layer header, and appends its own header (the technical term for this is "encapsulation"). The presentation-layer header contains the information necessary for the presentation layer in computer B to correctly handle the data. After adding its header, the presentation layer transfers the new data unit to the session layer.

This process is repeated through all layers in computer A until a final header is added at the data-link layer. After the data-link–layer header is added, the data unit is known as a "frame." The data, or frame, is passed from the data-link layer to the physical layer and is transmitted across the transmission medium connecting the two computers.

When the signal reaches computer B, layer one in computer B (the physical layer) copies the data. Now the process is reversed. The physical layer in computer B transfers the data to the data-link layer. The data-link layer removes the header information that was attached by the corresponding layer in computer A, acts upon the information the header contains, and transfers the data unit up to the network layer. This process continues, with the headers being stripped off at each layer and the instructions contained therein carried out, until the original data from computer A (the message) is finally passed from the application layer to the E-mail application in computer B. When the E-mail application receives the message, it displays the message on the screen for the user of computer B to read.

Now look at Figure 16 and imagine what would be possible if the software implementing different layers of the OSI model were able to handle not just one communication protocol at any one layer, but almost any communication protocol used at any layer, by any computer—there would be no limits to the interconnection of dissimilar computing devices. This is the kind of power that will be the basis for a global network—the networking of all kinds of business and personal devices into the Information Superhighway. And this is the kind of power built into NetWare products.

Commonly Used Standards

When you read about NetWare products, you will find references to various standards and communication protocols supported by NetWare networks.

To understand the capabilities of NetWare products, it will help to know the OSI layer at which a particular protocol operates and why the standard is important. As you shall see later, by converting protocols or using multiple protocols at different layers of the OSI model, it is possible to enable different computer systems to share data, even if they use different software applications, operating systems, and data-encoding techniques.

Figure 17 shows some commonly used standards and the OSI layer at which they operate.

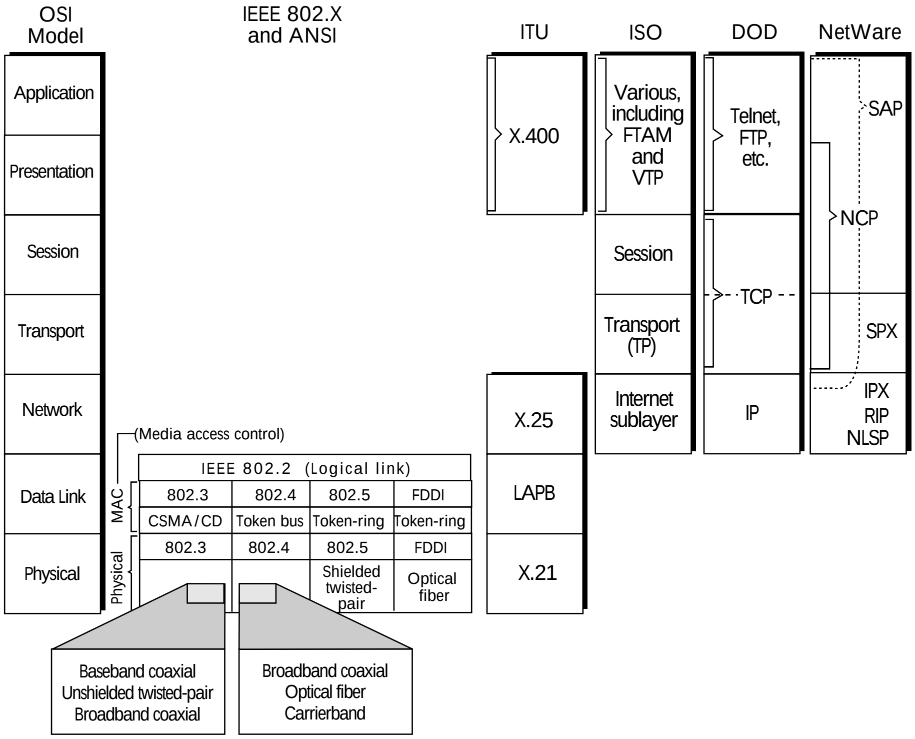

Figure 17

Important standards at various OSI layers

Layer-One Standards: Physical

Standards at the physical layer include protocols for transmitting a bitstream over media such as baseband coaxial cable, unshielded twisted-pair wiring, and optical fiber cable. The most commonly used are those specified in the Institute of Electrical and Electronic Engineers (IEEE) 802.3, 802.4, and 802.5 standards and the American National Standards Institute (ANSI) Fiber Distributed Data Interface (FDDI) standard. Figure 17 shows the transmission media included in each of these standards. An emerging standard for this layer is the Synchronous Optical Network (SONET).

Layer-Two Standards: Data Link (Media Access Control and Logical Link Control)

The most commonly used layer-two protocols are those specified in the IEEE’s 802.2, 802.3, 802.4, and 802.5 standards and the ANSI FDDI standard (all shown in Figure 17). Many microcomputer networking products use one of these standards (or the virtually identical ISO version).

Important recent technologies at layer two include 100Base-T (IEEE 802.2u), 100VG-AnyLAN (802.12), and Asynchronous Transfer Mode (ATM). The ATM standard is not yet fully defined. Also, frame relay is an important layer-two WAN technology. These technologies are treated in greater detail in a later section.

Layer-two standards encompass two sublayers: media access control and logical link control.

Media Access Control

The media access control (MAC) protocol specifies how workstations cooperatively share the transmission medium.

The IEEE 802.3 standard specifies a media access method known as “carrier sense multiple access with collision detection (CSMA/CD).” This media access method is the same as the contention method described in the earlier discussion of topologies, under the heading “Logical Bus.”

The IEEE 802.4, 802.5, and FDDI standards all specify some form of token passing as the media access control method. The basics of the token-passing method were also described earlier, also under the heading “Logical Bus.”

In general, using a form of token passing for the media access control works best when large numbers of computers frequently send small amounts of data—for example, when a number of workstations continually read and write small records to and from a database. Contention schemes work well when computers send large amounts of data intermittently—for example, during desktop publishing or document imaging.

Logical Link Control

The function of the logical link control layer is to ensure the reliability of the physical connection.

The IEEE 802.2 standard is the most commonly used logical link control standard.

The Point-to-Point Protocol (PPP) is an important de facto standard at this OSI level. PPP is used for communications across point-to-point links such as T1 and T3 lines. It is an important protocol for wide area networking, which will be covered later.

Layer-Three Standards: Network

The function of the network layer is to manage communications, most importantly the routing and relaying of data, between nodes.

One important network-layer standard is the Department of Defense (DOD) Internet Protocol (IP), which is part of the Transmission Control Protocol/Internet Protocol (TCP/IP) standard developed by the DOD. This protocol has become extremely important recently because it is the basis for the Internet and for all intranet technology. Also, the Department of Defense will often not purchase networking products that cannot communicate using this protocol.

Because Novell commands a large share of the networking market, its native network-layer protocol, Internetwork Packet Exchange™ (IPX™), is also an important de facto network-layer standard. IPX is a connectionless datagram protocol. A connectionless protocol does not need to establish a connection between two networked computers to transfer information between them. Packet acknowledgment, or connection control, is provided by protocols above IPX, such as Novell's Sequenced Packet Exchange™ (SPX™). SPX will be explained in more detail in a later section. Because IPX is a datagram protocol, each communication packet is treated as an individual entity. IPX does not have to establish a logical or sequential relation between packets. Thus, IPX is very efficient—it addresses and transfers data with minimum control overhead.

IPX uses other NetWare protocols that work at the network layer to accomplish internetwork routing. These protocols, the Routing Information Protocol (RIP), the Service Advertising Protocol (SAP), and the NetWare Link Services Protocol™ (NLSP™), will be explained in more detail in a later section.

The Consultative Committee for International Telegraph and Telephone (CCITT) X.25 standard is another commonly used network-layer standard. It specifies the interface for connecting computers on different networks by means of an intermediate connection made through a packet-switched network (for example, a common carrier network such as CompuServe, Tymnet, or Telnet). The X.25 standard includes the data-link and physical-layer protocols shown below it in Figure 17.

Apple Computer, Inc. has established a set of protocols for its products, referred to collectively as AppleTalk. At the network layer of the OSI model, the Apple protocol is called Datagram Delivery Protocol. Figure 18 shows how the set of AppleTalk protocols fits within the OSI model.

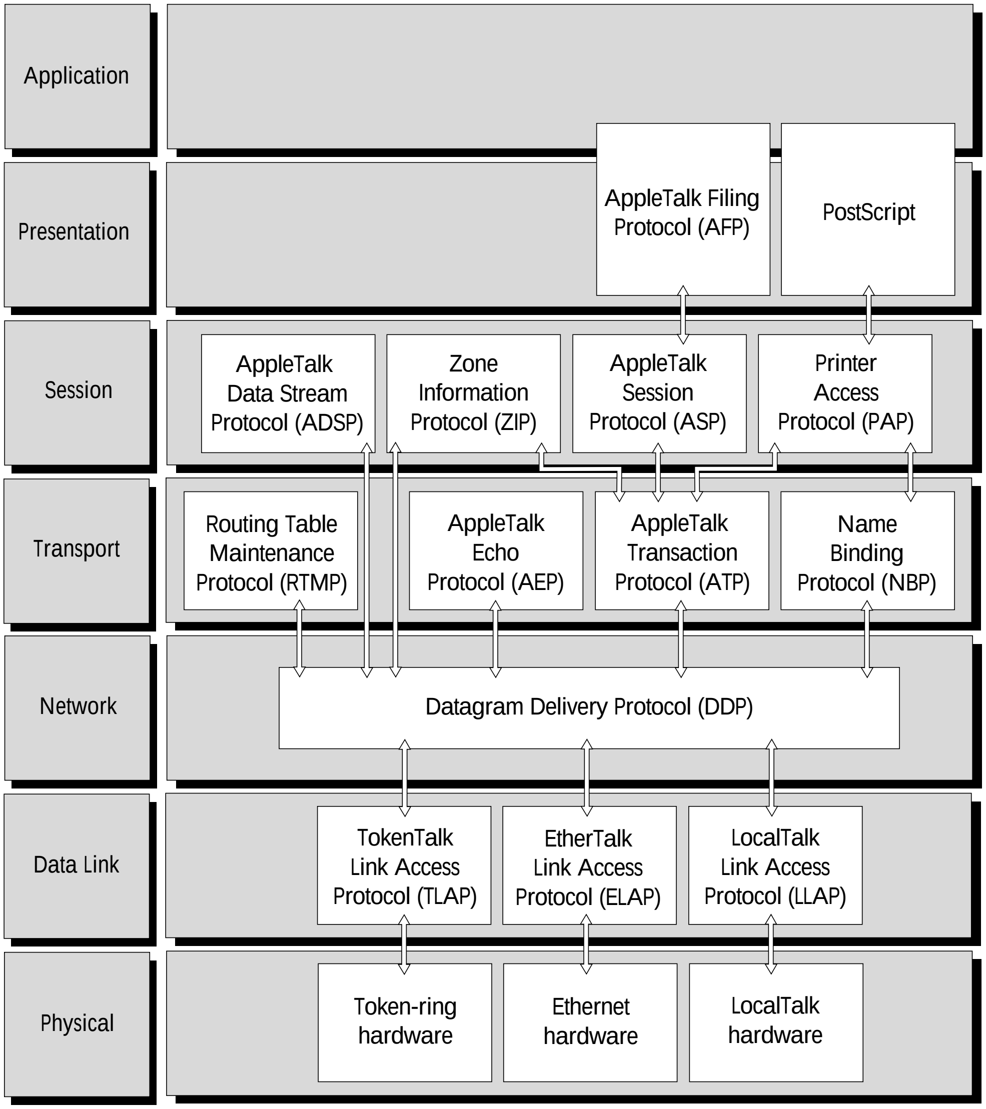
Application
Presentation
Session
Transport
Network
Data Link
Physical
AppleTalk Filing Protocol (AFP)
PostScript
AppleTalk Data Stream Protocol (ADSP)
Zone Information Protocol (ZIP)
AppleTalk Session Protocol (ASP)
Printer Access Protocol (PAP)
Routing Table Maintenance Protocol (RTMP)
AppleTalk Echo Protocol (AEP)
AppleTalk Transaction Protocol (ATP)
Name Binding Protocol (NBP)
Datagram Delivery Protocol (DDP)
TokenTalk Link Access Protocol (TLAP)
EtherTalk Link Access Protocol (ELAP)
LocalTalk Link Access Protocol (LLAP)
Token-ring hardware
Ethernet hardware
LocalTalk hardware

Figure 18

Where AppleTalk protocols fit in the OSI model

Like Novell's native protocols, Apple's de facto standard protocols are important because of Apple's wide acceptance in the microcomputer market.

Layer-Four Standards: Transport

Standards at this OSI layer provide for the reliability of the end-to-end communication link. This layer isolates the upper three layers, which are all concerned with user and application requirements, from knowing the details required to manage the end-to-end connection.

The ISO has issued a transport-layer standard that is simply called the Transport Protocol (TP). Because it is an ISO standard, it is of worldwide importance.

At the transport layer, Novell's native protocol is SPX. SPX provides guaranteed packet delivery and packet sequencing. Although it is basically a transport-layer protocol, it also includes session-layer functions. The NetWare Core Protocol™ (NCP™) and SAP also provide transport-layer functions. SPX, NCP, and SAP will be treated in more detail in a later section.

The AppleTalk protocol set has a number of protocols that operate at the transport layer, including Routing Table Maintenance Protocol, AppleTalk Echo Protocol, AppleTalk Transaction Protocol, and the Name Binding Protocol.

IBM's NetBIOS protocol (not shown in Figure 17) is also an important protocol at this layer and at the session layer above.

The DOD's Transmission Control Protocol, which is part of the TCP/IP standard, is important at the transport layer to the same degree (extremely important) and for the same reasons as the IP standard at layer three. This protocol provides all functions required for this layer (transport) and part of the functions for the session layer above.

Layer-Five Standards: Session

The function of the session layer is to establish, manage, and terminate the connections of individual network users.

The ISO session standard, named simply "session," has the same worldwide importance as the ISO transport standard. The DOD's Transmission Control Protocol performs important functions at this layer.

In a NetWare environment, the NetWare Core Protocol provides most of the necessary session-layer functions. SAP also provides functions at this layer.

Layer-Six and Layer-Seven Standards: Presentation and Application

The presentation layer performs general data transformations useful to a variety of applications, thus providing a useful common interface. Presentation-layer services include data encryption and text compression. The application layer provides user applications with basic (yet complete) services such as file transfer and network management functions.

Two important OSI protocols encompassing both the presentation and application layers are File Transfer, Access, and Management (FTAM) and Virtual Terminal Protocol (VTP). Each of these protocols is exactly what its name implies. FTAM provides user applications with useful file transfer and management functions.
VTP supports applications by converting specific terminal characteristics to a general (virtual) terminal model shared by applications.

X.400 is an important CCITT standard that encompasses both the presentation and application layers. X.400 provides message handling and E-mail services. It is an important standard because it is the basis for a number of pervasive E-mail packages as well as for other widely used messaging products.

An important DOD standard at this level is File Transfer Protocol (FTP), which, again, is named for the service it provides.

The NetWare protocols that provide presentation- and application-layer functions are NCP and SAP. All NetWare protocols will be treated in more detail in a later section.

Further Perspective: Standards and Open Systems

You probably noticed from looking at Figures 17 and 18 that most accepted standards are not neatly packaged to include all (and only) those services specified for any OSI layer. In fact, most common standards encompass parts of multiple OSI layers. This includes most standards adopted by the various government agencies that develop them.

Product vendors' actual implementation of OSI layers is even less neatly divided. Vendors implement accepted standards, which already include mixed services from multiple layers, in different ways.

So why go to all the trouble to agree on a model and then define standards if you are not going to be exact when fitting the standards to the model or in implementing the standards when building a product?

Actually, standards development and implementation have proceeded more or less as expected. The OSI model was never intended to foster a rigid, unbreakable set of rules. It was expected that in implementing the OSI communication model, networking vendors would be free to use whichever standard for each layer they deemed most appropriate. They would also be free to implement each standard in the manner best suited for the purposes of their products.

As noted earlier, however, it is clearly in a vendor's best interest to manufacture products that conform to the intentions behind the OSI model. To do this, a vendor must provide the services required at each OSI model layer in a manner that will enable its system to be simply and easily connected to the systems of other vendors—in other words, vendors must develop open systems. The consequences of not doing so are severe and unavoidable.

Which leads to the next issue—how do you determine if a system is an open system? You can start by getting answers to simple questions such as: (1) Can you establish communications using virtually any accepted communication standard? and (2) How easily can you do this? For example, can you communicate with other networks that are using the TCP/IP protocol, even if your network uses some other protocol at that layer? If you can communicate, what kind of effort is required? And how reliable are such communications?

As you begin asking questions like these, you will find that Novell has the answers you need. NetWare products support every standard we have presented, as well as virtually every other accepted standard. The more you understand NetWare products, the more you will understand that no system is more open than a NetWare system.

The Network Operating System

From our discussion of the OSI model, you have no doubt begun to understand how complex it is to control communications on a computer network. And, you are probably wondering: What is the means of accomplishing this task? The answer is software called the network operating system (NOS).

The network operating system software acts as the command center, enabling all of the network hardware and all other network software to function together as one cohesive, organized system. In other words, the network operating system is the very heart of the network.

Client-Server Network Operating Systems

On a client-server network, the network operating system is installed and runs on a computer called the network server. The server must be a specific type of computer. For example, the most commonly used client-server version of the NetWare network operating system runs on IBM PC and compatible computers.

A client-server operating system is responsible for coordinating the use of all resources and services available from the server on which it is running.

The client part of a client-server network is any other network device or process that makes requests to use server resources and services. For example, network users at workstations request the use of services and resources though client software, which runs in the workstation and talks to the operating system in the server by means of a common protocol.

On a NetWare client-server network, users “log in” to the network server from the workstation. To log in, a user enters a login command and gives his or her user name and password. If the user name and password are valid, the server “authenticates” the user and allows him or her access to all network services and resources to which he or she has been granted rights. As long as the user has proper network rights, the client-server operating system provides the services or resources requested by the distributed applications running in workstations.

The operating system manages various server resources, which include hardware such as hard disks, RAM, printers, and equipment used for remote communications, such as modems. The network file system is also a server resource.

In addition, the network operating system provides many services, including coordinating file access and file sharing (including file and record locking), managing server memory, managing data security, scheduling tasks for processing, coordinating printer access, and managing internetwork communications.

Among the most important functions performed by a client-server operating system are ensuring the reliability of data stored on the server and managing server security.

There are many other functions that can and should be performed by a network operating system. We do not have room to cover them all here. However, many functions might be very important to you, and this means that choosing the right NOS is of paramount importance. NetWare NOSs are robust systems that provide many capabilities not found in less mature systems. NetWare NOSs also provide a level of performance and reliability far above that found in most other network operating systems.

To learn more about client-server operating systems, including the services they can and should provide, read the product sections of this *Buyer's Guide* that cover IntranetWare and NetWare client-server operating systems, including IntranetWare, IntranetWare for Small Business, NetWare 4.11, NetWare 4.11 for OS/2, NetWare 3.12 NOS, and SFT III™ for IntranetWare.

Peer-to-Peer Network Operating Systems

Peer-to-peer network operating systems enable networked computers to function as both a server and a workstation. In a peer-to-peer network, the operating system is installed on every networked computer; this enables any networked computer to provide resources and services to all other networked computers. For example, each networked computer can allow other computers to access its files and use connected printers while it is in use as a workstation.

Peer-to-peer operating systems have both advantages and disadvantages when compared to client-server operating systems. They provide many of the same resources and services as do client-server operating systems, and, under the right circumstances, can provide good performance. They are also easy to install and are usually inexpensive.

However, peer-to-peer networks provide fewer services than client-server operating systems. Also, the services they provide are a great deal less robust than those provided by mature, full-featured client-server operating systems, and the performance of peer-to-peer networks commonly decreases significantly under a heavy load. Furthermore, except in the case of Novell's Personal NetWare™ peer-to-peer network operating system, maintenance is often more difficult: Because there is no method of centralized management, there are often many servers to manage (rather than one centralized server), and many people may have access to and the ability to change the configuration of different server computers.

For more information about the differences between peer-to-peer and client-server networks and the level of services they offer, refer to the Personal NetWare product description in this *Buyer's Guide*.

Desktop Operating Systems

Each workstation on the network must have software that manages its own resources. This software, known as the desktop operating system, enables a workstation to perform such functions as accessing files from its own local hard disks, displaying information on its video display, coordinating local printing, and so on.

There are a number of commonly used desktop operating systems, including Windows 3.x, Windows NT, Windows 95, UNIX, PC-DOS, OS/2, MS-DOS, and various versions of the Macintosh operating system.

Each of the different desktop operating systems has advantages and disadvantages. Unfortunately, for the most part, they are not compatible with each other. Software written for one operating system will not function on another. Furthermore, peripheral hardware (such as modems, facsimile machines, and so on) that is compatible with the hardware required to run one kind of desktop operating system is usually not compatible with hardware required to run other desktop operating systems.

This brings us to another important function of a network operating system—it should be able to interconnect all of the commonly used desktop operating systems to ensure that all network users have access to the computer that they are most familiar with and that is best suited to the job they need to do.

Novell network operating systems enable you to integrate all popular desktop operating systems directly on one network. They allow this because they are able to translate the data from one desktop operating system into data that the other desktop operating systems can read.

Application Software

Once you have your network hardware, a powerful network operating system, and the necessary desktop operating system(s), the final network tool you will need is application software. Application software enables you to do the “real work” you want to do. Commonly used application software includes word processing, accounting, spreadsheet, database, and E-mail programs. You may also need customized applications, or even one-of-a-kind applications, built specifically for your company.

One extremely important issue to consider when selecting commercially built application software is its degree of network and intranetwork integration. To effectively use network and intranet services, application software must be well integrated with the network operating system. The degree of network integration will determine how well the application enables collaboration among network users, whether and how well it provides direct access to all network services, and whether it is as easy as it can be to manage across the network.

To learn more about network-integrated software, read the section in this *Buyer’s Guide* that covers Novell GroupWare, including GroupWise 5.1™, GroupWise WebAccess, GroupWise PhoneAccess, and GroupWise Remote.

Internetworking

As a business grows, it might need to split its network. Or, a business might need to connect two separate networks so that users on each can use resources on either. When a network is split (or when two networks with different addresses are connected), this results in an internetwork. An internetwork has subnetworks (network segments) that have different network addresses. Even a modest-sized business often has several subnetworks operating, each serving a specific portion of the organization.

Why might a business need subnetworks?

The most common reason for segmenting a network is to preserve excellent network performance. On even the fastest and most efficient network, if the network has too many users (devices that need to transmit), the transmission media can become so busy that devices have to wait an unacceptable time to transmit. When this happens, users begin to notice delays when they try to save or open files or perform other operations.

When you segment a network, you give each subnetwork its own network address. This results in two separate transmission media segments, which can be used simultaneously. Each of the two segments will have only half the users of the original network. Thus, you double network performance (on some networks, performance can more than double because on an overloaded network, the overhead required to manage transmission collisions takes a much larger percentage of bandwidth than on a modestly busy network).

Networks are also segmented to enhance data security and to minimize the effect of equipment failure on any part of the network.

Internetworking includes everything from connecting two small workgroup networks, each with perhaps two or three workstations, to connecting thousands of computers—from notebook computers to mainframes—on tens to hundreds of individual segments in a worldwide organization.

Internetworking Devices: Bridges and Routers

Bridges and routers are the devices used to interconnect subnetworks. They can be primarily hardware based or primarily software based.

Software-based routers and bridges can be part of a server's operating system or can at least run in the server with the operating system. Software-based bridges and routers can also be installed on standard computers to create dedicated, standalone devices. For example, IntranetWare MultiProtocol Router software is a family of software-based routing products that can be installed on an IntranetWare, NetWare 4™, or NetWare 3™ server or on a standalone PC.

To understand internetworking, it is not essential that you understand all the technical differences between a bridge and router. In fact, without some study, this can be a confusing area. For example, if you read about IntranetWare MultiProtocol Routers, you will find that these routers also perform what is called source-route bridging.

However, without a basic understanding of bridging and routing technology (and related terminology), you will find it difficult to understand the capabilities of some products and the reasons such capabilities are useful or important. Please keep in mind throughout the following discussion that bridges and routers have one important thing in common: They both allow the transfer of data packets (frames) between subnetworks with different network addresses.

Bridges

A bridge operates at the data-link layer (layer two) of the OSI model. A bridge acts as an address filter; it relays data between subnetworks (with different addresses) based on information contained at the media access control level.

Simple bridges are used to connect networks that use the same physical-layer protocol and the same MAC and logical link protocols (OSI layers one and two). Simple bridges are not capable of translating between different protocols.

Other types of bridges, such as translational bridges, can connect networks that use different layer-one and MAC-level protocols; they are capable of translating, then relaying, frames.

After a physical connection is made (at OSI layer one), a bridge receives all frames from each of the subnetworks it connects and checks the network address of each received frame. The network address is contained in the MAC header. When a bridge receives a frame from one subnetwork that is addressed to a workstation on another subnetwork, it passes the frame to the intended subnetwork. Figure 19 illustrates, in a general fashion, how a bridge relays frames between subnetworks.

Figure 19

Internetworking through a bridge

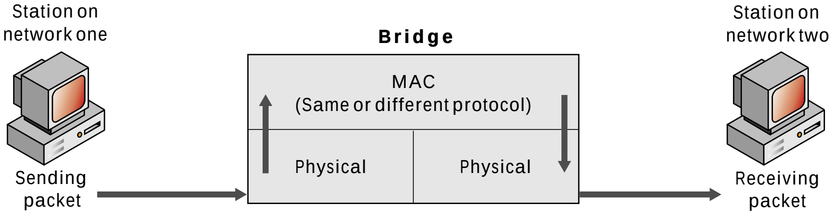

A bridge assumes that all communication protocols used above the data-link layer at which it operates (OSI layers three through seven) are the same on both sides of the communication link. Of course, this must be true, or there must be translation between unlike protocols at layers three through seven for the receiving computer to be able to interpret the transferred data.

Spanning Trees and Source-Route Bridging

There are two terms connected with bridging that will be useful to understand: spanning trees and source-route bridging.

Spanning trees prevent problems resulting from the interconnection of multiple networks by means of parallel transmission paths. In various bridging circumstances, it is possible to have multiple transmission routes between computers on different networks. If multiple transmission routes exist, unless there is an efficient method for specifying only one route, it is possible to have an endless duplication and expansion of routing errors that will saturate the network with useless transmissions, quickly disabling it. Spanning trees are the method used to specify one, and only one, transmission route.

Source-route bridging is a means of determining the path used to transfer data from one workstation to another. Workstations that use source routing participate in route discovery and specify the route to be used for each transmitted packet. Source-route bridges merely carry out the routing instructions placed into each data packet when the packet is assembled by the sending workstation—hence the name “source routing.” In discussions of bridging and routing, do not be confused by the term “source routing.” Though it includes the term “routing,” it is a part of bridging technology. Source-route bridging is important because it is a bridge-routing method used on IBM Token-Ring networks.

You should understand that bridging technologies and routing methods can be combined in various ways. For example, there is an IEEE specification for a source-route transparent bridge, a bridging scheme that merges source-route bridging and transparent bridging in one device.

From this simple discussion of bridging, one thing should be apparent: When choosing internetworking products, it is important to select those that support the various bridging methods—products such as IntranetWare MultiProtocol Router. (For further details, see the IntranetWare MultiProtocol Router 3.1 product section in this *Buyer’s Guide.*)

Routers

Routers function at the network layer of the OSI model (one layer above bridges). To communicate, routers must use the same network-layer protocol. And, of course, the sending and receiving workstations on different networks must either share identical protocols at all OSI layers above layer three, or there must be necessary protocol translation at these layers.

Like some bridges, routers can allow the transfer of data between networks that use different protocols at OSI layers one and two (the physical layer and the data-link layer, which includes sublayers for media access control and logical link control). Routers can receive, reformat, and retransmit data packets assembled by different layer-one and layer-two protocols. Different routers are built to manage different protocol sets. Figure 20 illustrates how a router transfers data packets.

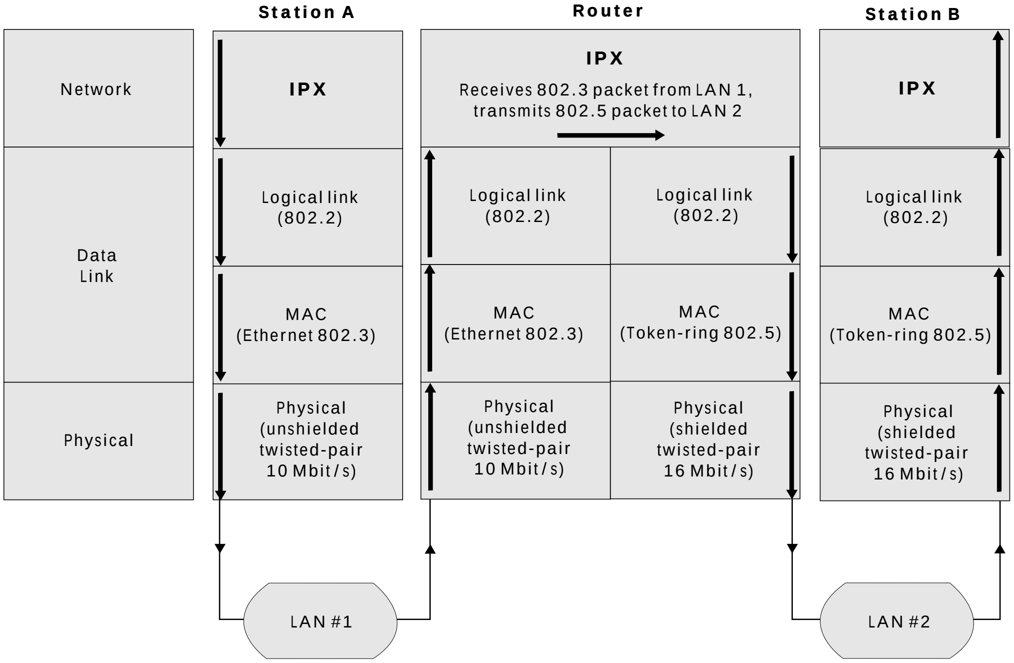
Station A
Router
Station B
Network
Data Link
Physical
IPX
IPX
Receives 802.3 packet from LAN 1, transmits 802.5 packet to LAN 2
IPX
Logical link (802.2)
Logical link (802.2)
Logical link (802.2)
Logical link (802.2)
MAC (Ethernet 802.3)
MAC (Ethernet 802.3)
MAC (Token-ring 802.5)
MAC (Token-ring 802.5)
Physical (unshielded twisted-pair 10 Mbit/s)
Physical (unshielded twisted-pair 10 Mbit/s)
Physical (shielded twisted-pair 16 Mbit/s)
Physical (shielded twisted-pair 16 Mbit/s)
LAN #1
LAN #2

Figure 20

Internetworking through a router

NetWare Internetworking Protocols

Before we conclude with a discussion of host connection, wide area networking technologies, and global networking, you should understand a little more about native NetWare protocols that play a role in NetWare internetworking. Figure 21 shows in greater detail how NetWare protocols fit into the OSI model.

Figure 21

Where NetWare protocols fit in the OSI model

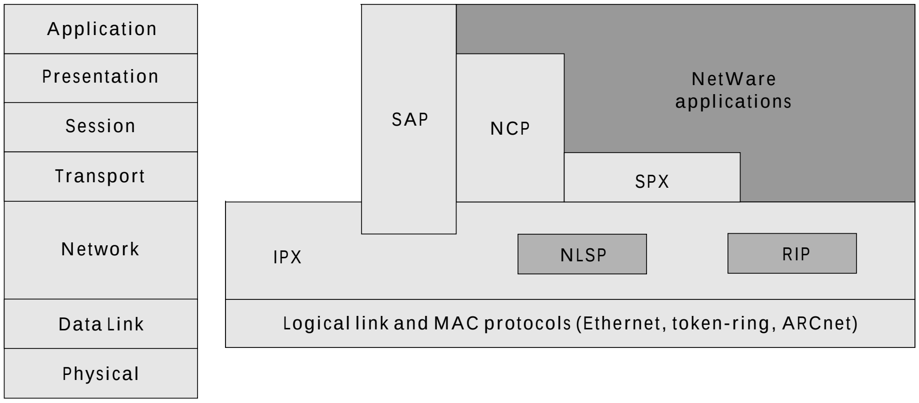

Each of the native NetWare protocols shown in Figure 21 plays a role in NetWare internetworking, either directly or indirectly.

IPX: The Network Layer Protocol

In conjunction with industry-standard media access control protocols, the NetWare IPX protocol provides the NetWare addressing mechanism that delivers communication packets to their destination. IPX works with all important MAC standards. As you can see from Figure 21, IPX operates at the network layer of the OSI model.

In a NetWare environment, internetwork packet routing is accomplished at the network layer. Thus, IPX is the NetWare protocol that addresses and routes packets between internetworked computers.

IPX bases its routing decisions on the address fields in its packet header (provided by the MAC protocol) and on the information it receives from other NetWare protocols. For example, IPX uses information supplied by either RIP or NLSP to forward packets to the destination computer or to the next router. IPX also uses SAP.

RIP and NLSP: The Routing Protocols

NetWare routers use one of two routing protocols, RIP or NLSP, to exchange routing information with neighboring routers.

RIP: A Distance-Vector Routing Protocol

The NetWare RIP is a distance-vector protocol. Distance-vector routing protocols are the traditional method used for router communications.

In an internetwork using distance-vector routing, routers periodically determine if the internetwork configuration has changed. They also periodically broadcast packets to their immediate neighbors; these packets contain all information they currently have about the internetwork's topology.

After receiving any information, distance-vector routers consolidate the information and pass summarized data along to other routers, servers, and end devices, such as printers and workstations. Through this periodic checking and broadcasting, which is performed at regular intervals regardless of whether the internetwork has changed, all routers are kept updated with correct internetwork addresses for all computers and other connected devices, as well as with the best route for transferring data between any two devices.

Because RIP is a distance-vector protocol, NetWare routers that use RIP work in the way described above, performing periodic checking and information exchange and updating their routing tables with any new information.

RIP is one of a number of well-known distance-vector routing protocols. Examples of other such protocols include IP RIP and Cisco IGRP, part of the IP protocol suite, and RTMP, part of the AppleTalk protocol suite.

NLSP: A Link-State Routing Protocol

The NetWare Link Services Protocol is a link-state routing protocol. This type of protocol derives its name from the fact that link-state routers track the status of other routers and links.

Link-state protocols, a relatively recent development, adapt more quickly to network topology changes than do distance-vector protocols. Thus, they are better than distance-vector protocols for managing internetworking on large, complex internetworks.

In an internetwork that uses a link-state routing protocol, each router or server provides information about itself and its immediate neighbors to every reachable router in a routing area. Each router’s map includes all the area’s routers and servers, the links connecting them, and the operational status of each router and link. However, each router builds its own routing map rather than relying on secondhand summaries, as do distance-vector routers. Also, routing transmissions are made only when the internetwork changes, not at predefined intervals. Thus, networks using link-state routing are not burdened by unnecessary routing traffic.

Because NLSP works as explained above, it significantly reduces the communication overhead required for routing. NLSP can significantly improve network performance because it frees resources to be used for transferring data packets rather than routing information. NLSP is particularly efficient for wide area network routing, where available communication bandwidth is ordinarily limited.

Examples of other link-state protocols include the Open Shortest Path First protocol, part of the TCP/IP protocol suite, and the Intermediate System-to-Intermediate System protocol, a router-to-router protocol that is part of the OSI suite.

As a matter of note, various link-state and distance-vector routing protocols can coexist on the same NetWare internetwork and even in the same IntranetWare MultiProtocol Router. Furthermore, individual routers can be configured to accept or to reject individual protocols.

SAP: Service Advertising Protocol

The Service Advertising Protocol is similar in concept to RIP. Just as RIP enables routers to exchange routing information, SAP enables networked devices, such as network servers and routers, to exchange information about available network services.

Servers and routers use SAP to advertise their services and network addresses. SAP enables network devices to constantly correct their information about which network services are available. While servers are running, they use SAP to inform the rest of the network of the services they offer. When a server goes down, it uses SAP to inform the network that its services are no longer available.

Routers gather service information and share it with other routers. Workstations use the information made available through SAP to obtain the network addresses of servers that offer the services they need.

NCP: NetWare Core Protocol

The NetWare Core Protocol is a set of service protocols that a server's operating system follows to accept and respond to service requests.

NCP does not play a direct role in routing. However, it does provide session control and packet-level error checking between NetWare workstations and routers.

SPX: Sequenced Packet Exchange

SPX™ is a transport-layer protocol. Standards at this OSI layer provide for the reliability of the end-to-end communication link. Accordingly, SPX provides guaranteed packet delivery and packet sequencing.

Like NCP, SPX does not play a direct role in routing. SPX is connected with internetworking only in that it guarantees delivery of all routed packets.

Gateways

In contrast to bridges and routers, which function at only one layer of the OSI model, a gateway translates protocols at more than one OSI layer. Therefore, a gateway is used to interconnect computer systems that have different architectures and that therefore use different communication protocols at several OSI layers.

A gateway may connect dissimilar systems on the same network or on different networks (thus, using a gateway does not necessarily involve internetworking). For example, a gateway might translate protocols at several different OSI layers to allow transparent communications between NetWare IPX-based systems and systems based on TCP/IP, System Network Architecture (SNA), or AppleTalk. Figure 22 illustrates how a gateway is used to translate protocols to enable communications between two heterogeneous systems.

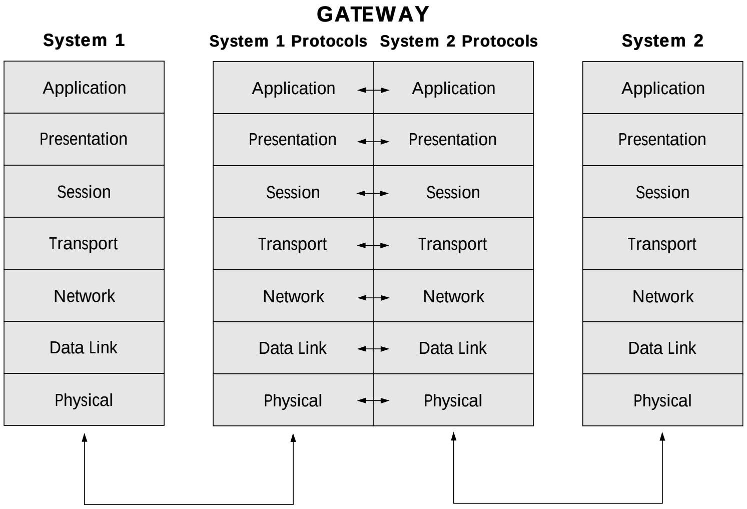

Figure 22

Gateways provide protocol translation between dissimilar systems at more than one OSI layer.

A gateway may consist of hardware, software, or a combination of the two, and it may provide translation at all or at only some of the different OSI layers, depending on the types of systems it connects.

There are a number of NetWare gateways that provide access to computer systems not based on the native NetWare/IPX protocol suite. NetWare for Macintosh is a software-based gateway that connects Macintosh computers to a PC-server–based NetWare network. NetWare for SAA is a gateway that enables NetWare users to transparently access SNA-based IBM hosts.

Real World Networking

In the real world, computer networks can combine a great number of physical and logical topologies. We've already seen a very simple network. Now let's take a look at more complex networks. We'll start with simple internetworks and work our way gradually through more complex situations. For the sake of simplicity and clarity, all of the subnetworks in our internetworks will be based on the NetWare client-server networking model.

One Server, Multiple Networks of the Same Type

The simplest form of an internetwork is two cabling (media) segments of the same media access control type sharing one network server.

For example, one server could contain two Ethernet network adapters, each supporting a different cable segment. There could be several computers connected to each cable segment, in a star physical layout, with each cable segment using contention (CSMA/CD) for the media access control. Each of the cable segments would have a different network address—thus, each would be an independent subnetwork.

Together, the two separate networks would form an internetwork, connected by means of internal routing capabilities built into the server. (Remember, we have already said that in NetWare servers, internetworking is accomplished through routing at the network layer.)

Figure 23 illustrates the one-server internetwork described above.

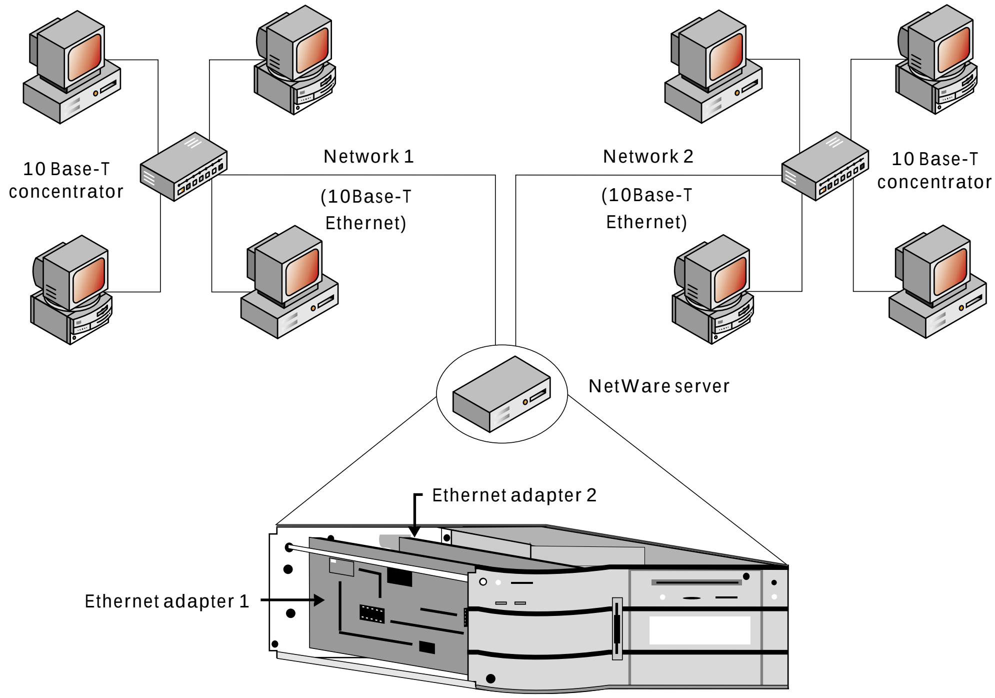

Figure 23

Internetworking two networks using the same type of network adapter (MAC) in one NetWare server, by means of the server's internal routers

In the case of the above network, routing would be accomplished using the NetWare IPX protocol or the NetWare IP protocol, with support from the other NetWare routing protocols, as previously described.

Every NetWare server is capable of using internal routers to accomplish local network routing by means of the NetWare routing protocol set and AppleTalk. All NetWare internal routers operate at layer three of the OSI model and are for use with small workgroup or departmental networks. For larger or more complicated internetworks, or for departments with heavy server-processing requirements, IntranetWare MultiProtocol Routers or dedicated routers from other vendors provide the necessary extra routing power and capabilities.

In a slightly more complex internetwork, a NetWare server could support multiple cable segments using the same physical layouts but different media access controls.

For example, a server could contain one Ethernet network adapter and one token-ring network adapter, with a cable segment attached to each. The Ethernet network might be connected in a physical star and use CSMA/CD for the media access control. The token-ring network might also be connected in a physical star, but it would use token passing for media access control. Like the simpler configuration explained in the previous section, each cable segment would have a different network address. Figure 24 illustrates this more complex one-server internetwork.

Figure 24

Internetworking two networks using different types of network adapters (MAC) in one Netware server, by means of the server's internal routers

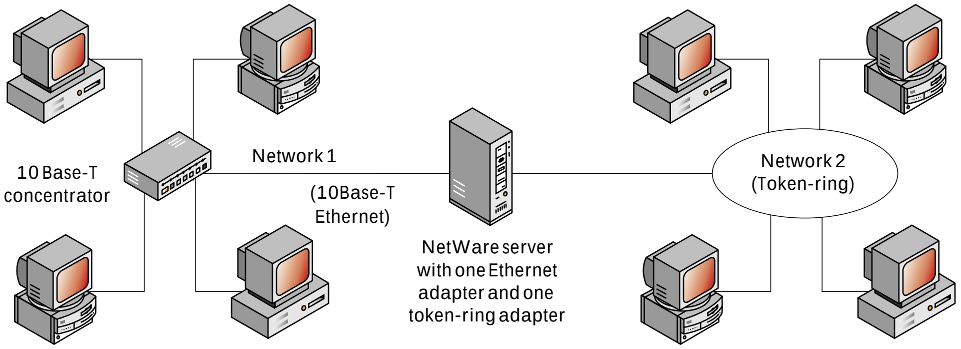

In the case of the internetwork shown above, routing would again be accomplished using the NetWare IPX or NetWare IP protocol, with support from the other NetWare routing protocols.

The two one-server networks we have seen each support only two separate subnetworks. As a matter of note, all NetWare servers are capable of supporting as many as four different network adapters (four separate subnetworks), in any combination of same or different types.

Please notice that even though the token-ring network above was described as a physical star, it is drawn as a ring to signify that it is a token-ring network (which uses token passing as the media access control). We will adhere to this convention throughout this *Buyer's Guide* because in virtually all illustrations, it will be more important to make the logical topology clear than to be concerned with the physical topology.

Multiple Servers With Multiple Networks of Different Types

In an even more complex internetwork, there may be multiple servers.

For example, a complex internetwork might consist of two one-server subnetworks connected by a standalone router, such as IntranetWare MultiProtocol Router. Each server might contain multiple network interface adapters.

One server might contain two Ethernet network adapters and one token-ring network adapter, with a cable segment attached to each. One of the Ethernet adapters might support a PC network, and the other Ethernet adapter might connect to both PCs and Macintosh computers. The NetWare for Macintosh product running on the server would support the Macintosh computers.

The other server might contain one Ethernet adapter and one ARCnet adapter, with the Ethernet adapter again supporting both PCs and Macintoshes, and the ARCnet adapter supporting a cable segment with a number of PCs attached.

Each of the two servers would have a unique internal number (server address), and each cable segment in each server would have a unique physical network (cable segment) address.

In this case, there would be five subnetworks on the internetwork, three attached to one server and two attached to the other. The internal server routers would accomplish the routing between any two workstations on subnetworks attached directly to the same server. Both the internal server routers and the intermediate standalone router would be involved in the routing between any two workstations on subnetworks attached to different servers.

Figure 25 illustrates the two-server internetwork described above.

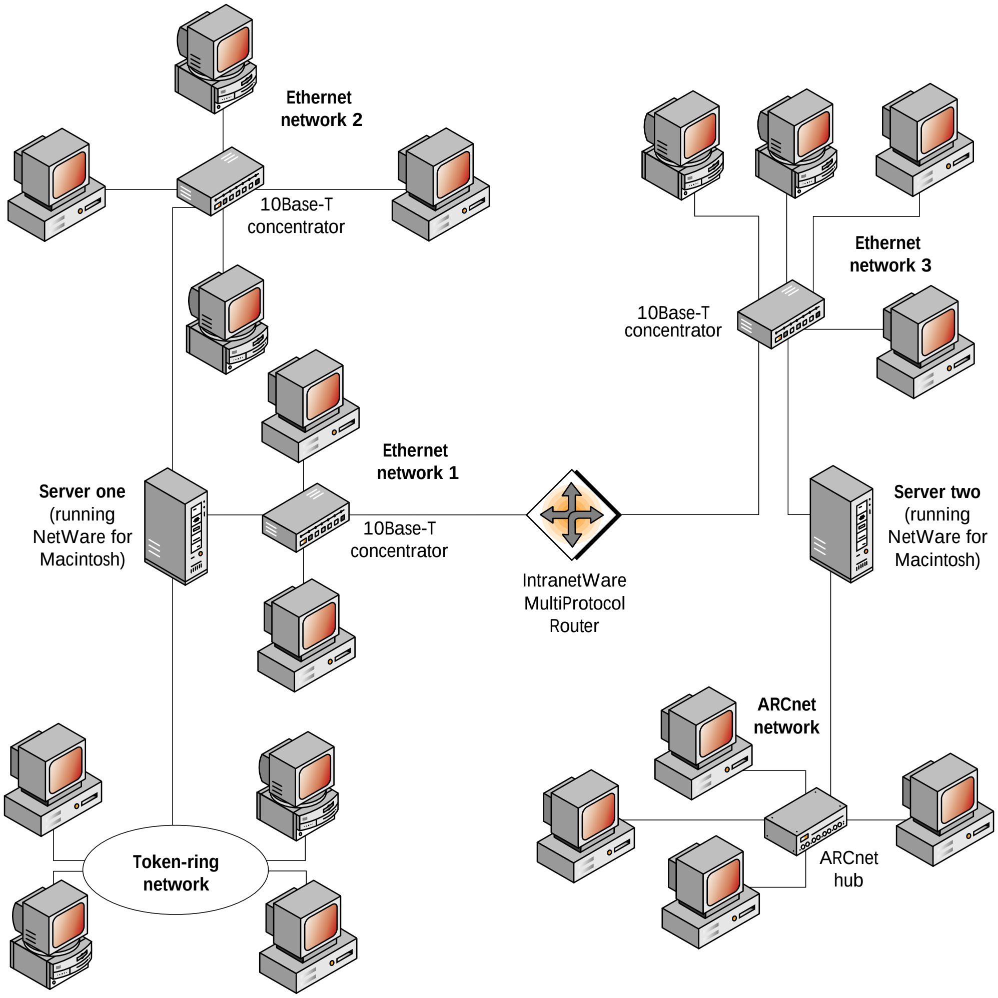

Figure 25

Internetworking multiple networks using different types of network adapters (MAC) in two NetWare servers, by means of internal and standalone routers

Host Access

An already complex multiserver internetwork becomes even more complex with the addition of connections to host computer systems, including mainframe computers such as IBM mainframes, to minicomputers such as IBM's AS/400 or a DEC VAX system, or to other hosts such as UNIX workstations.

Host systems can provide access to additional application software, additional resources such as data storage devices and printers, and additional processing power. For example, you might want to log in to an IBM AS/400 minicomputer to run an application available only on that computer or to use its processing power for one task while you were using the processing power of your own workstation for some other task. Or, you might want to print a large report on a high-speed printer connected to the AS/400.

The illustration in Figure 26 shows a multiserver NetWare network with an IBM mainframe, an IBM AS/400 minicomputer, and several UNIX workstations connected as host computers.

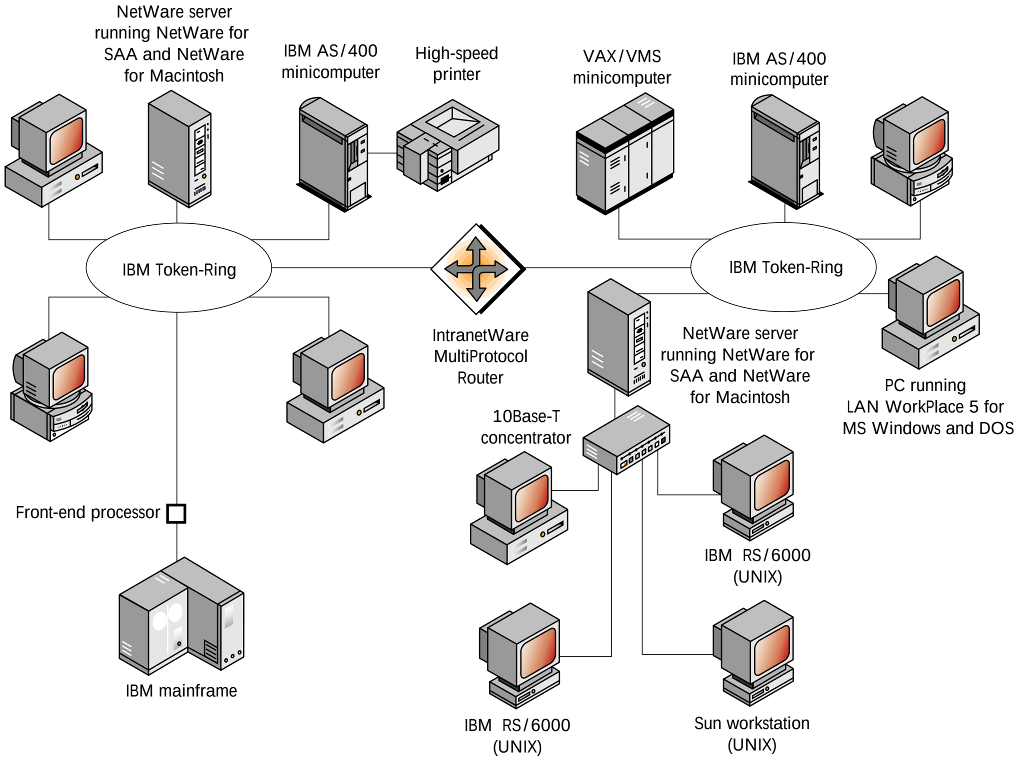

Figure 26

Host systems connected to a complex multiserver NetWare network

To learn more about connecting to host systems from NetWare networks, read the product sections in Chapter 7: Network-to-Network Connectivity; Chapter 8: Host Connectivity; and Chapter 9: UNIX and TCP/IP Connectivity. See the table of contents for listings.

A number of leading networking companies have entered into original equipment manufacturer (OEM) partnerships with Novell. Many provide NetWare connectivity to host-based environments. For a list of Novell's OEM partners, see Appendix C.

Wide Area Networking

So far, we have looked only at internetworks existing at one local site, with computers and other devices directly connected by some type of cabling. This type of network is commonly referred to as a "local area network" (LAN). Now let's add another level of complexity to the internetworking picture: Let's connect complex multiserver internetworks that exist at separate sites, which might be any number of miles apart. Connecting computer networks in these circumstances is called "wide area networking."

The traditional definition of wide area networking has been “connecting two or more networks existing at widely separate geographic sites.” Some traditionalists also prescribe that the separate networks must be connected by means of common carrier telecommunication facilities (private companies that rent resources such as T1 lines and microwave transmission equipment). For the purposes of this primer, we’ll use the first, general definition and let you decide how to apply it in specific internetworking cases. But, to give you some background to help you make such decisions, let’s discuss a few specific internetworking cases and a few terms related to wide area networking.

Of course, like any general term used in connection with rapidly changing technology, not everyone will agree on an exact definition of wide area networking. What is “widely separate”? And, does the connection really have to be through a common carrier? Many major companies now own their own equipment linking networks many miles apart.

Let’s look at some examples. Suppose you connect two networks in two different buildings 100 yards apart by means of asynchronous modems and common telephone lines. Is that wide area networking? Most knowledgeable computer networking people would say no—this would be “one-site” or “campus” networking. What if the networks were two miles apart and separated by a major interstate highway? Or, what if they were 15 miles apart, on opposite sides of a major city? There are many computer networking people who would still not call this wide area networking; they might use a recently coined term—“metropolitan area networking.” Others consider metropolitan area networking a part of wide area networking. Of course, everyone would agree that two networks connected on opposite sides of a continent by means of a satellite microwave link rented from a common carrier is an example of a wide area network.

You can decide for yourself where you think wide area networking begins and ends. Now let’s look at some general possibilities.

Figure 27 shows two separate branch office internetworks connected to a third internetwork at a main corporate office. Each of the three existing internetworks has multiple servers and existing host connections. One of the branch office networks is connected to the corporate network by means of asynchronous modems and regular voice-grade telephone lines. The other branch office network is connected by means of a common carrier–provided intermediate link—in this case an X.25 packet-switching network. Examples of such networks include Tymnet or Telnet.

However, either network could be connected by other means that we have discussed, such as frame relay or a dedicated leased line link, perhaps using PPP.

The following section describes important WAN and LAN technologies in greater detail. To learn more about wide area network connection methods, including wide area communications by means of asynchronous modems, read the Novell Connect Services product section and the product sections in Chapter 7: Network-to-Network Connectivity.

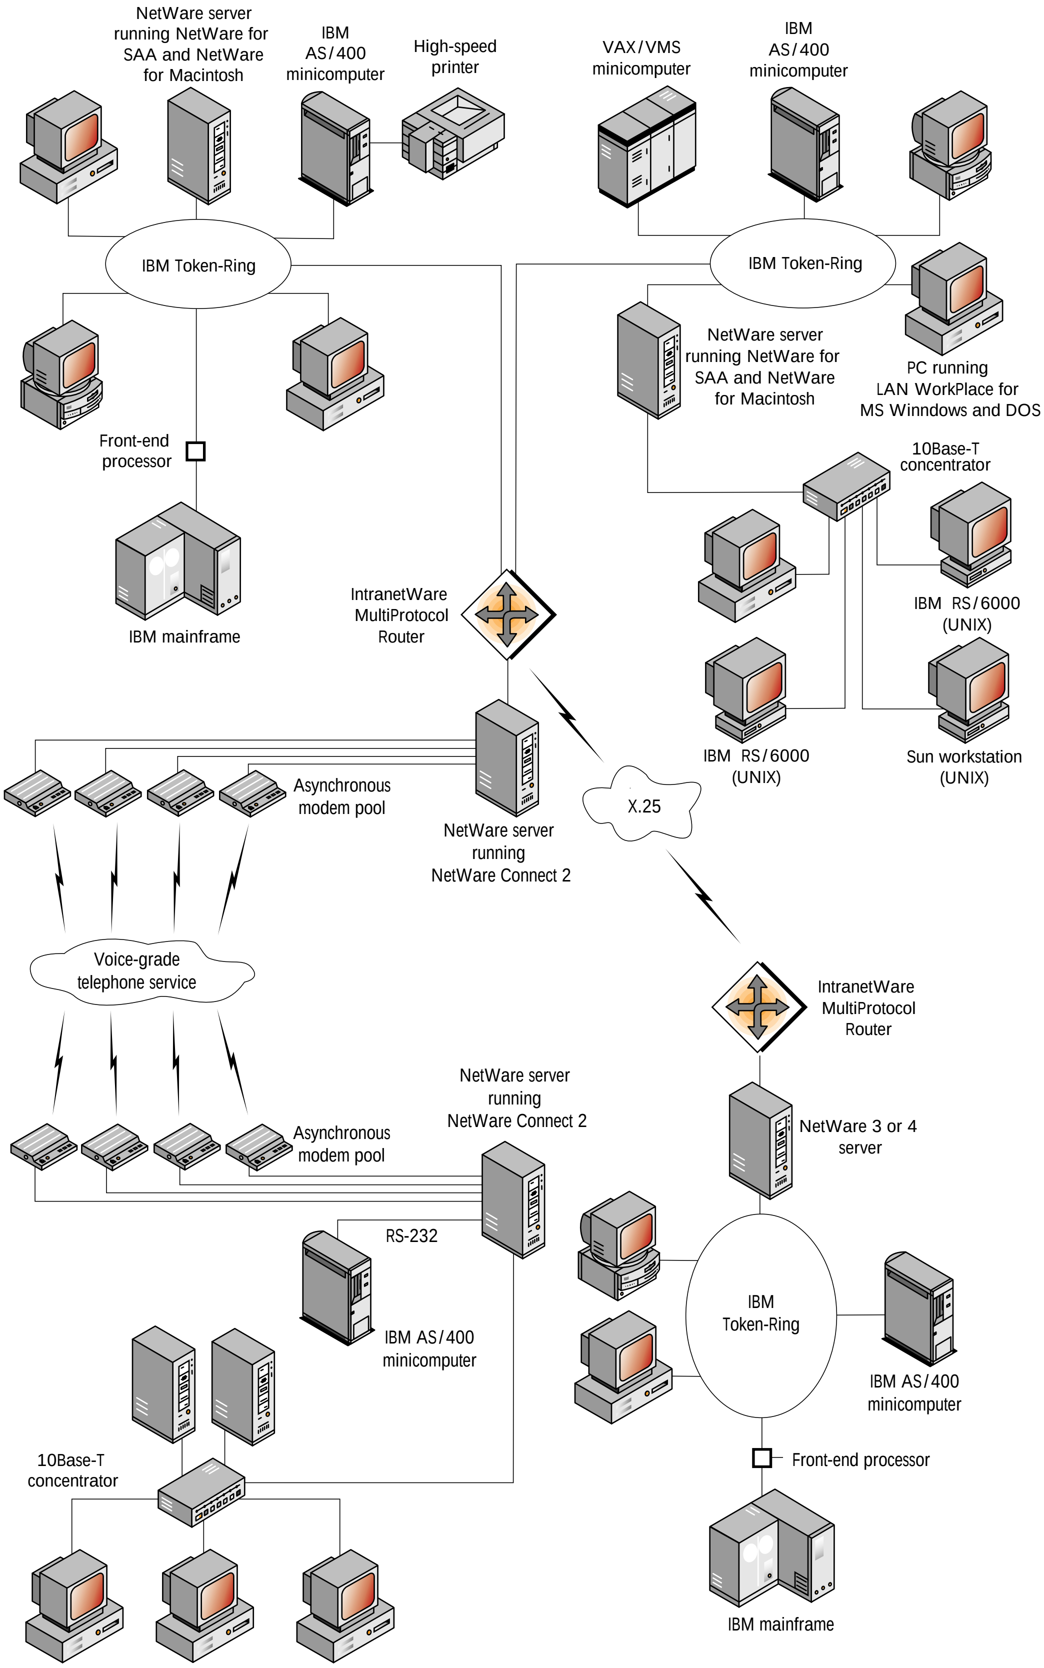

Figure 27

Wide area networking: three networks at widely separated sites connected through asynchronous modems and an X.25 connection

Important WAN and High-Speed Technologies

These days, if you pick up any computer networking magazine, you'll find that among the hot topics are the technologies that make networks faster and the technologies that connect geographically distant networks—technologies such as frame relay, 100VG-AnyLAN, and ATM. We mentioned these and other high-speed and WAN technologies under the "Commonly Used Standards" heading of this primer. Brief explanations will help you more fully understand what these technologies are and why they are important. The following technologies will be treated:

- 100Base-T
- 100VG-AnyLAN
- Fiber Distributed Data Interface (FDDI)
- X.25
- Frame relay
- Asynchronous Transfer Mode (ATM)
- Integrated Services Digital Network (ISDN)
- Synchronous Optical Network (SONET)

100Base-T

100Base-T is a high-speed LAN technology. 100Base-T is officially designated as the IEEE 802.3u standard. It functions at the data-link (OSI level two) layer's media access control sublayer and provides data transfer rates as high as 100 megabits per second (Mbit/s).

Distinguishing Characteristics

Like 10Base-T Ethernet, 100Base-T uses carrier sense multiple access with collision detection as the media access control method. (CSMA/CD was described in the earlier discussion of topologies, under the heading "Logical Bus.") 100Base-T is based on the scalability of CSMA/CD. Scalability means that you can easily enlarge or downsize your network without degrading network performance, reliability, and manageability.

CSMA/CD was known to be scalable before the 100Base-T standard was created. A scaled-down version of Ethernet (1Base-5) uses CSMA/CD, provides data transfer rates of 1 Mbit/s, and enables longer transmission distances between repeaters. If CSMA/CD could be scaled down, then it could be scaled up. Specifying changes such as decreased transmission distances between repeaters produced a reliable data transfer rate of 100 Mbit/s, 10 times faster than traditional 10Base-T Ethernet.

100Base-T supports Category 3 and 5 unshielded twisted-pair (UTP) wiring, Type-1 shielded twisted-pair (STP) wiring, and fiber-optic cable. It uses four wire pairs of Category 3 UTP cable—three for data and one for collision detection. However, 100Base-T uses only two wire pairs of Category 5 UTP cable.

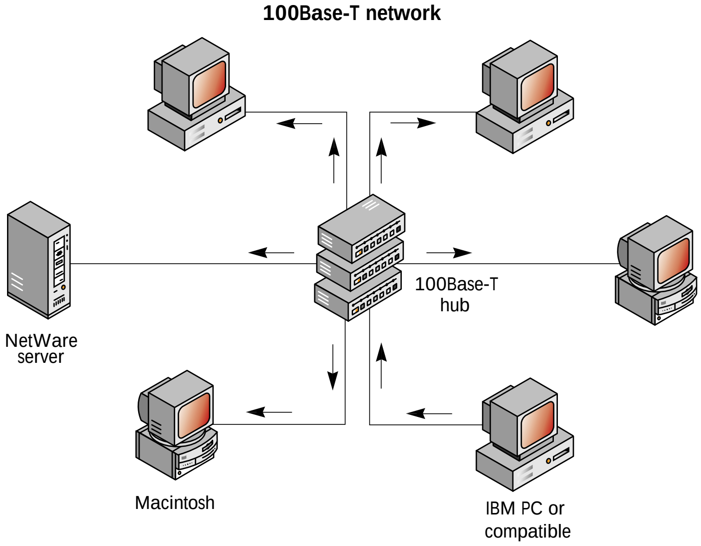

Figure 28

On 100Base-T networks, the physical topology is a star and the logical topology is a bus. A broadcast signal travels to all parts of the cable.

Advantages

100Base-T is widely available. Adapter cards and compatible cable are currently available from various vendors.

In addition, it’s easy to upgrade from 10Base-T Ethernet to 100Base-T Ethernet. Both traditional 10Base-T and 100Base-T Ethernet use CSMA/CD, and some network cards now support both 10 Mbit/s and 100 Mbit/s Ethernet. The adapter cards automatically sense whether it is a 10 Mbit/s or 100 Mbit/s environment and adjust their speed accordingly. Because 100Base-T and 10Base-T Ethernet can coexist, network supervisors can upgrade network stations from 10Base-T to 100Base-T one at a time, as needed. Also, most network supervisors are already familiar with CSMA/CD, so there is no need for expensive retraining.

100Base-T can be an inexpensive way to make your network faster. Adapter cards are not significantly more expensive than 10Base-T cards. In addition, Category 3 and Category 5 UTP cable are relatively inexpensive and many organizations already have either Category 3 or 5 cable installed.

Disadvantages

100Base-T will reduce the maximum network size compared to 10Base-T because the standard specifies shorter transmission distances between repeaters.

In addition, the fact that 100Base-T is based on CSMA/CD creates problems. 100Base-T may scale CSMA/CD to its limit, making 100 Mbit/s the maximum data transfer rate for this standard. To increase data transfer rates, 100Base-T specifies shorter distances between signal repeaters, and these distances may be as short as is practical. Also, because CSMA/CD is a shared media contention scheme, collisions will occur, especially under maximum loads. This results in increased overhead, which reduces actual data throughput.

Furthermore, 100Base-T requires four wire pairs of Category 3 cable, and not all companies that have Category 3 cable have four wire pairs available. Thus, companies that are already using some wire pairs for a different purpose, or that installed cable with fewer than four wire pairs or cable that does not meet Category 3 standards, will have to recable to use 100Base-T.

100VG-AnyLAN

100VG-AnyLAN, which is officially designated as the IEEE 802.12 standard, is a high-speed LAN technology that competes with 100Base-T. Like 100Base-T, 100VG-AnyLAN functions at the data-link layer (OSI level two) and provides data transfer rates as high as 100 Mbit/s. However, 100VG-AnyLAN differs from 100Base-T in several important respects.

Distinguishing Characteristics

Instead of using CSMA/CD as the media access control method, 100VG-AnyLAN uses a method called “demand priority.” Demand priority is not a contention scheme; rather, it uses a form of token passing that assigns the token (permission to broadcast data packets) based on a priority scheme that the network supervisor establishes for different types of LAN traffic and on the order in which an intelligent hub receives requests.

For example, when a workstation needs to transmit, it signals the intelligent hub that it needs access to the transmission media. If the intelligent hub receives several requests, it will give access to the workstation that has the highest priority. (100VG-AnyLAN will also function without a prioritization scheme.) If the workstations requesting access have the same priority, the intelligent hub will assign the token to the workstations in the order they request access to the transmission media.

In addition, 100VG-AnyLAN supports both Ethernet and token-ring networks. It also supports Category 3 and 5 UTP, Type-1 STP, and fiber-optic cable. 100VG-AnyLAN uses four wire pairs of Category 3 or Category 5 UTP cable.

Figure 29

On 100VG-AnyLAN networks, both the physical and logical topologies are stars. The signal from one node goes to the intelligent hub and is routed only to the correct destination node.

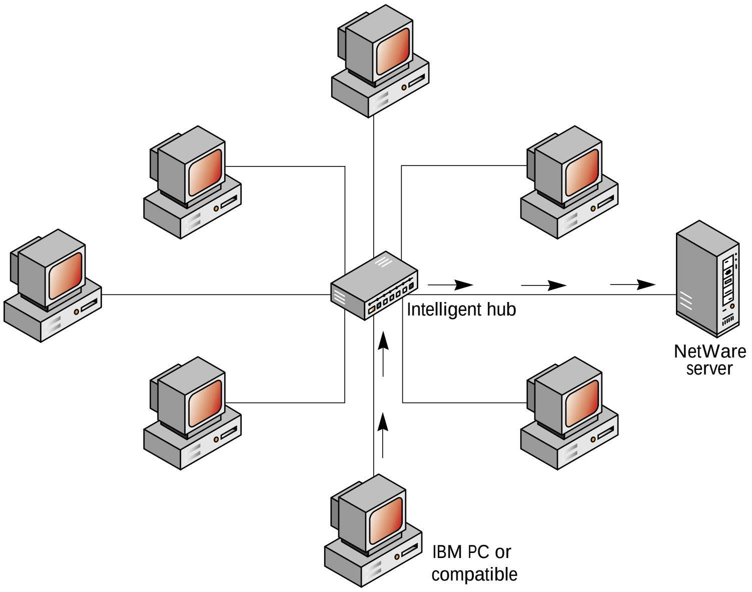

Advantages

Because 100VG-AnyLAN uses demand priority, it has several advantages over 100Base-T. First, 100VG-AnyLAN provides the necessary bandwidth and timing (low latency) needed by applications such as multimedia applications. Network supervisors can, if desired, assign higher priority to the hub ports used to connect workstations or servers that frequently transmit time-sensitive data, such as audio and video. Second, the demand priority media access method avoids the collisions that can occur on 100Base-T networks, ensuring that control overhead will not soar. Third, 100VG-AnyLAN networks are switched (they do not broadcast packets to all workstations), which makes 100VG-AnyLAN networks more secure against eavesdroppers than 100Base-T networks.

Furthermore, unlike 100Base-T, 100VG-AnyLAN supports token-ring networks as well as Ethernet, providing data transfer rates as high as 100 Mbit/s to the former.

100VG-AnyLAN and 100Base-T also share many advantages. The cost of 100VG-AnyLAN is comparable to 100Base-T: Adapter cards that support both 10 and 100 Mbit/s are not priced significantly higher than traditional 10Base-T Ethernet cards. Both standards also support the same types of transmission media. In addition, both are easy to upgrade.

Disadvantages

Unlike 100Base-T's CSMA/CD, which is familiar to many network supervisors, demand priority is new, and network supervisors will require some training to use it effectively. Also, 100VG-AnyLAN has a smaller market share than 100Base-T. Consequently, it is not supported by as many vendors, which means that fewer products are available for 100VG-AnyLAN.

Previously installed cable may be problematic for 100VG-AnyLAN networks, as it is for 100Base-T. 100VG-AnyLAN uses all four wire pairs of Category 3 or 5 UTP cable. Thus, companies that are already using some wire pairs for a different purpose, or that installed cable with less than four wire pairs or cable that does not meet Category 3 standards, will have to recable to use 100VG-AnyLAN.

Fiber Distributed Data Interface

Fiber Distributed Data Interface is also a high-speed LAN technology. It is not generally used for direct connection to desktop computers, but rather as a backbone technology. A backbone connects two or more LAN segments to provide a path for transmitting packets among them. A simple backbone might connect two servers through a high-speed link consisting of network adapter cards and cable.

FDDI is officially designated as ANSI X3T9.5 and operates at the physical and data-link layers (levels one and two) of the OSI model. Like 100Base-T and 100VG-AnyLAN, FDDI provides data transfer rates as high as 100 Mbit/s.

Figure 30

A simple server-based backbone connecting two LAN segments

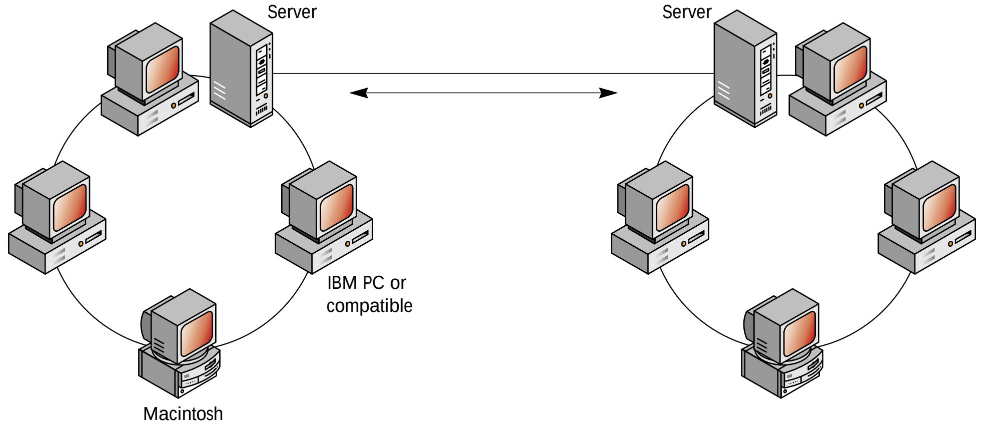

Distinguishing Characteristics

FDDI networks have a dual, counter-rotating ring topology. This topology consists of two logical closed signal paths called “rings.” Signals on the rings travel in opposite directions from each other. Although both rings can carry data, the primary ring usually carries data while the secondary ring serves as a backup.

On FDDI networks, every node acts as a repeater. FDDI supports four kinds of nodes: dual-attached stations (DASs), single-attached stations (SASs), single-attached concentrators (SACs), and dual-attached concentrators (DACs). DASs and DACs attach to both rings; SASs and SACs attach only to the primary ring. Several SASs often attach to the primary ring through a concentrator so that an SAS failure will not bring down the entire network. If the cable is cut or a link between nodes fails, DASs or DACs on either side of the failure route signals around the failed segment using the secondary ring to keep the network functioning.

FDDI uses token passing for the media access control method and is implemented using fiber-optic cable.

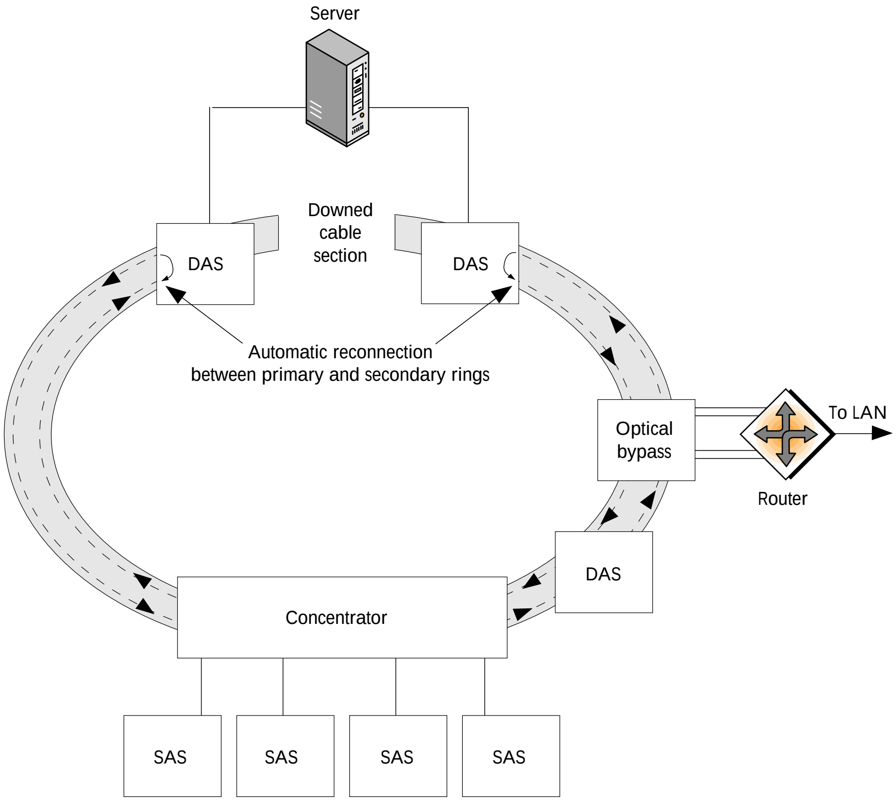

Figure 31

If a cable section on an FDDI network goes down, DASs on either side of the failed section automatically reconnect the primary and secondary rings. Also note that the server has a redundant connection to improve reliability.

Advantages

FDDI is a fast, reliable standard. The dual, counter-rotating ring topology increases the network’s reliability by keeping the network functioning even if a cable is damaged. FDDI also offers network management support, which was designed directly into the standard. Also, the standard includes the Copper Distributed Data Interface (CDDI) specification for building a network using UTP cable (which is less expensive than fiber-optic cable).

Disadvantages

FDDI’s main disadvantage is price. FDDI adapter cards and fiber-optic cable are both relatively expensive compared to other technologies offering the same speed. Fiber-optic cable installation also requires more expert technicians. Even CDDI adapters (for copper wire), which are less expensive than FDDI adapters, are more expensive than either 100Base-T or 100VG-AnyLAN adapters.

X.25

X.25 is a commonly used WAN standard at the network layer (level three) of the OSI model. It is a CCITT (now known as the International Telecommunications Union (ITU)) standard and includes data-link and physical layer protocols (LAP-B and X.21), as shown in Figure 17. X.25 provides data transfer rates of 9.6 kilobits per second (kbit/s) to 256 kbit/s, depending on the connection method.

Distinguishing Characteristics

X.25 specifies the interface for connecting computers on different networks by means of an intermediate connection through a packet-switched network (for example, CompuServe, Tymnet, or Telnet). X.25 was defined when the quality of transmission media was relatively poor. As a result, the standard specifies that each node in the packet-switched network must fully receive each packet and check it for errors before forwarding it.

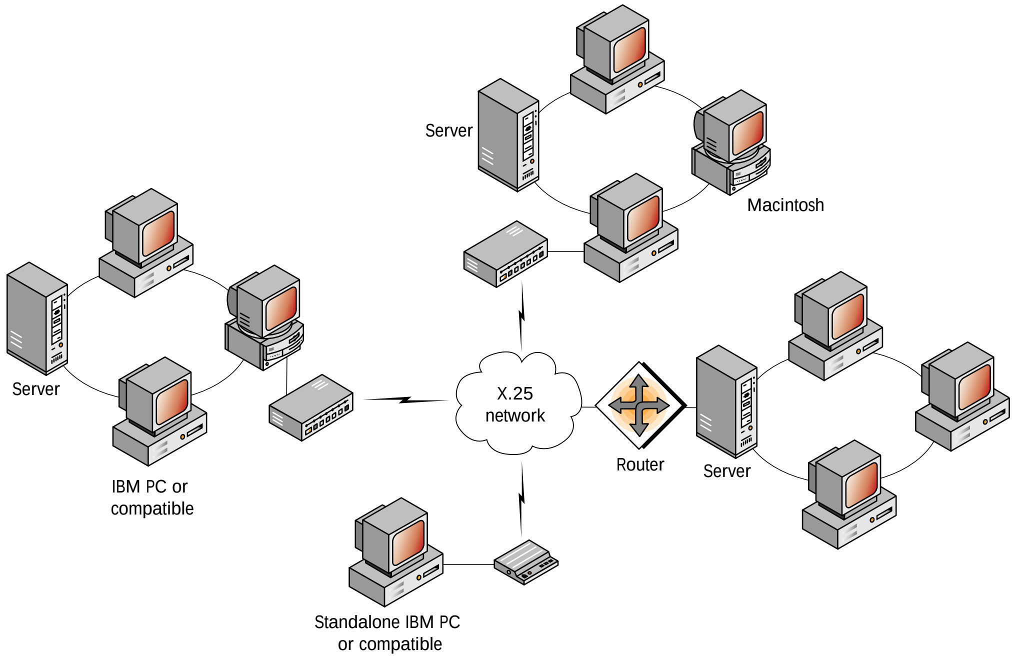

Figure 32

X.25 networks are often provided by telecommunication carriers. CompuServe uses X.25 on its network.

Advantages

X.25 is well understood and reliable. Connections to X.25 networks can be made through the existing telephone system, ISDN, and leased lines. Because access is so simple, it is comparatively inexpensive. X.25 is also available worldwide. In countries with little digital telecommunications infrastructure, X.25 may be the best WAN technology available.

Disadvantages

Although it is widely available, X.25 is slow compared to newer technologies. The process of checking each packet for errors at each node limits data transfer rates. X.25 also uses variable-size packets, which can cause transmission delays at intermediate nodes. In addition, many people connect to X.25 networks through modems, which limit data transfer rates from 9.6 kbit/s to 56 kbit/s. Although X.25 is likely to remain in common use for some time, newer, faster standards are already replacing it.

Frame Relay

Frame relay, like X.25, is a WAN technology. Approved by ANSI and ITU, frame relay works at the data-link layer (level two) of the OSI model, providing data transfer rates from 56 kbit/s to 1.544 Mbit/s.

Distinguishing Characteristics

Frame relay is an interface specification for connecting LANs over public packet-switched networks. This standard can be thought of as a simplified version of X.25 designed to take advantage of digital transmission media.

Frame relay services are typically provided by telecommunications carriers. Customers install a router and lease a line (often a T1 or fractional T1 line) to provide a permanent connection from the customer’s site to the telecommunications carrier’s network. This connection enables frame relay to use permanent virtual circuits (PVCs), which are predefined network paths between two locations.

With frame relay, the router encapsulates (or frames) network layer packets, such as IP and IPX packets, directly into a data-link level protocol and sends them on to the packet-switched network. Like X.25, frame relay uses variable-size frames, but it eliminates the error checking required on X.25 networks. A frame relay switch simply reads the header and forwards the packet, perhaps without even fully receiving a frame before forwarding it. Intelligent end stations must identify missing or corrupted frames and request retransmission.

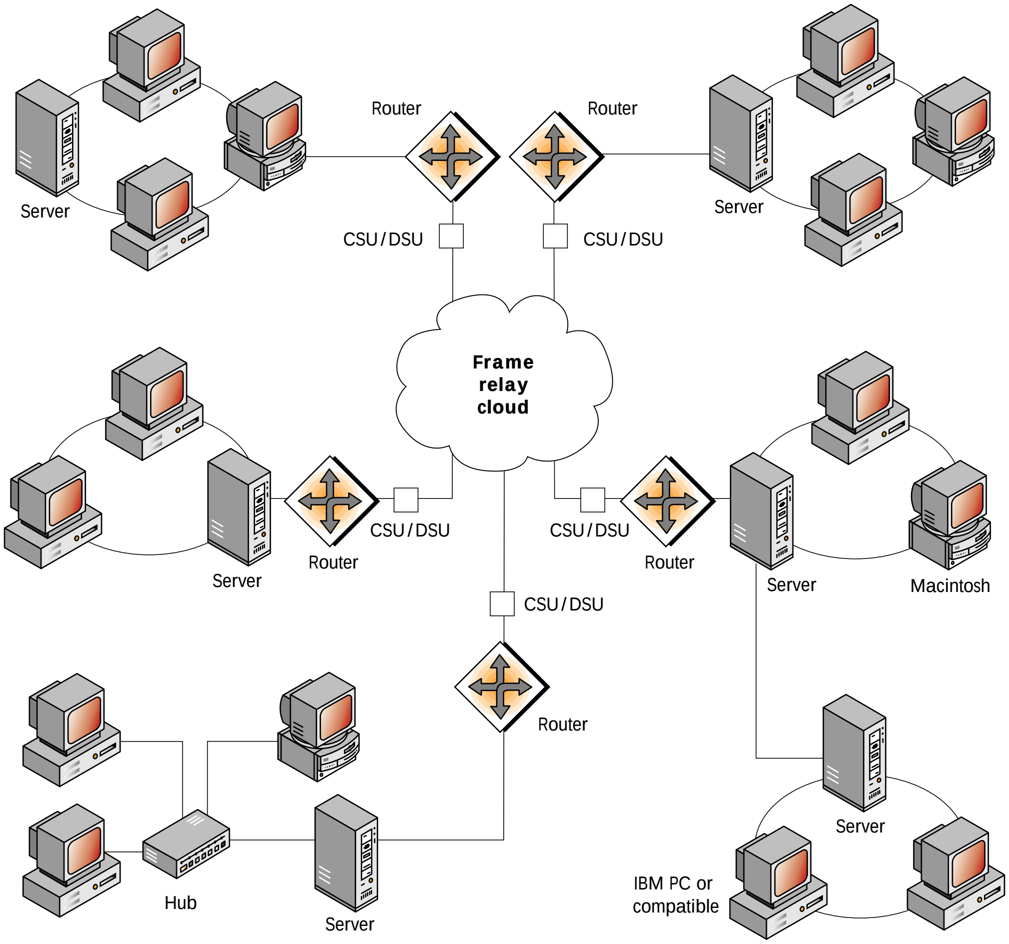

Figure 33

Frame relay is a WAN technology that enables companies to connect LANs through a telecommunications carrier's network. AT&T WorldNet Intranet Connect Service currently uses this technology.

Advantages

Frame relay offers several advantages over X.25. Most importantly, frame relay is faster than X.25. Frame relay uses PVCs over leased lines rather than a modem connection. Unlike modem connections, PVCs transmit and receive data immediately, eliminating the call setup and handshaking that modems must perform. In addition, as mentioned above, frame relay does not require error checking and flow control at the switches, reducing overhead and leaving more bandwidth for data transmission. Also, although not as prevalent as X.25, frame relay is a common standard in many countries. Finally, frame relay is less expensive than other WAN technologies because it provides bandwidth on demand, rather than dedicating bandwidth whether data is being transmitted or not.

Although frame relay is fairly complex to implement, value-added resellers and some telephone companies will assist customers in determining their needs and will help install the technology.

Disadvantages

Although frame relay is faster than X.25, its speed is limited because it uses variable-size frames, which can cause delays at switches along the frame's path. As a result, frame relay cannot support applications that require low latency, such as real-time video.

In addition, frame relay is more complex to implement than X.25. Customers must negotiate a service agreement with the phone company, lease a line, and have it installed. They must also purchase and install a frame relay–compatible router.

Asynchronous Transfer Mode

Asynchronous Transfer Mode is both a LAN and a WAN technology. It is generally implemented as a backbone technology. The exact relationship of the ATM layers to the OSI model is currently undefined, although ATM LAN emulation works at the data-link layer (level two).

ATM is extremely scalable; data transfer rates range from 25 Mbit/s to 2.4 gigabits per second (Gbit/s). This wide range of data transfer rates reflects the various ways in which ATM can be used. The 25 Mbit/s rate is a new offering meant for desktop environments. In LAN backbones, ATM provides data transfer rates of 100 Mbit/s and 155 Mbit/s. At the high end, WAN implementations using ATM and SONET together have achieved data transfer rates of 2.4 Gbit/s. (For more information about SONET, see the "Synchronous Optical Network" heading later in this primer.)

Distinguishing Characteristics

ATM is a cell relay technology, meaning that it uses standard-sized packets called cells. The size of an ATM cell is 53 bytes.

In a LAN implementation, ATM functions at the data-link layer's media access control sublayer. It further divides the MAC sublayer into three layers: LAN Emulation, ATM Adaptation Layer (AAL), and ATM. LAN Emulation enables you to integrate ATM with Ethernet and token-ring networks without modifying existing Ethernet or token-ring protocols.

On a mixed network, LAN Emulation hardware sits between the Ethernet or token-ring segment and the ATM part of the network. It uses the three layers mentioned above to convert packets moving toward the ATM segment into cells and to assemble cells moving toward the Ethernet or token-ring segment into packets. AAL and ATM put data into standard-sized cells. In most network computing situations, ATM Adaptation Layer 5 breaks packets into 48-byte blocks that are then passed to the ATM layer, where the five-byte header is attached to form a complete 53-byte cell.

Advantages

Many people believe that ATM will become the industry-standard transmission technology for LANs and WANs. The scalability, discussed above, seems to be limitless. Data transfer rates have climbed into the gigabit range and are still growing.

One reason that ATM is so fast is its use of cells. Because cells are a standard size, ATM networks handle data in a predictable, efficient manner at the switches. Standard-sized cells and high-bandwidth media like fiber-optic cable also enable ATM to support real-time voice, video, and data traffic.

ATM also offers flexibility in its transmission media. As many as 22 ATM specifications exist for media like unshielded twisted-pair, shielded twisted-pair, and fiber-optic cable. (ATM is generally implemented with fiber-optic cable.)

Although it is seen as a technology of the future, ATM can currently be integrated with Ethernet and token-ring networks, through use of LAN Emulation.

Disadvantages

ATM standards are still developing. Without industry standards, interoperability between equipment from different vendors is not guaranteed. Furthermore, ATM is more expensive than the other high-speed LAN technologies. The expense is preventing many companies from taking ATM to the desktop.

Integrated Services Digital Network

Integrated Services Digital Network is a set of protocols defined by CCITT to integrate data, voice, and video signals into digital telephone lines. It functions at the physical, data-link, network, and transport layers (levels one through four) of the OSI model. ISDN offers data transfer rates between 128 kbit/s and either 1.544 Mbit/s or 2.048 Mbit/s, depending on the country where it is implemented.

Distinguishing Characteristics

ISDN makes end-to-end digital connections over telephone lines. Although many telephone networks are almost completely digital, the local loop that connects a home or office to the telephone company’s network is not: Most local loops send analog rather than digital signals. ISDN replaces local analog signaling with digital signaling, enabling end-to-end digital communications.

ISDN offers Basic Rate Interface (BRI) for individuals or small branch offices and Primary Rate Interface (PRI) for larger companies.

BRI uses two bearer, or B, channels (providing 64 kbit/s each) to transmit and receive data and one delta, or D, channel for call setup and management.

PRI is the same thing as a T1 line. A T1 line in the United States consists of 23 B channels and one D channel, providing a total data transfer rate of 1.544 Mbit/s. A T1 line in Europe consists of 30 B channels and one D channel, providing a total data transfer rate of 2.048 Mbit/s. A fractional T1 uses only some of the B channels in a T1 line (and thus offers some fraction of the total T1 data transfer rate).

ISDN requires special equipment at the customer's site, including a digital phone line and a network termination unit (NT-1). An NT-1 converts the bandwidth coming over the line into the B and D channels and helps the phone company with diagnostic testing. The NT-1 also provides a connection for terminal equipment, such as ISDN telephones and computers that have an ISDN interface. In addition, the NT-1 provides terminal adapter (TA) equipment to connect equipment that is not compatible with ISDN. TA equipment provides an intermediary connection point: Such equipment has an ISDN interface, for connection to the NT-1, and a non-ISDN interface, for connection to non-ISDN equipment.

Advantages

ISDN increases speed and broadens data transmission capabilities, especially for those currently using analog modems to remotely connect to an office or to access the Internet. It offers faster call setup and faster data transfer rates. The transfer rates are acceptable for transmitting voice, data, limited video, fax, and images. ISDN can also be used for limited LAN-to-LAN communications.

With ISDN, you can transmit voice and data traffic simultaneously: An ISDN user can simultaneously talk on the phone and download a data file to his or her computer, over the same telephone line. For example, one BRI ISDN configuration enables users to use the two B channels (128 kbit/s) for data and part of the D channel for a phone conversation.

Disadvantages

Although widely available in Australia, Japan, and Western Europe, ISDN is available in only 50 percent of the United States. Presently, telephone companies are working to make it available throughout the United States.

Acceptance of ISDN in the United States has been slow for several reasons. First, to understand ISDN well enough to even order services requires considerable effort. Furthermore, configuration can be difficult. In addition, ISDN lacks the standards that ensure interoperability. As a result, customers must be careful to purchase equipment that is compatible with the local phone company's equipment. Another problem is that not all phone companies offer the same services, so customers must ensure that the services they need are available in their area. Finally, to take full advantage of ISDN, customers must communicate with others who also have ISDN.

Synchronous Optical Network

Synchronous Optical Network, also known in some countries as Synchronous Digital Hierarchy, is a WAN technology that functions at the physical layer (level one) of the OSI model. Telecommunications companies are implementing SONET on some of their networks: A typical business would not implement this standard on its network. SONET has been accepted by ANSI and recommended by ITU. It specifies a number of data transfer rates from 51.8 Mbit/s to 2.48 Gbit/s.

Distinguishing Characteristics

SONET defines a fiber-optic standard for high-speed digital traffic. This standard provides the flexibility to transport many digital signals with different capacities.

Data communications sometimes prove difficult because digital signaling rates can vary. For example, in the United States, a T1 line provides 1.544 Mbit/s; in Europe, a T1 line (sometimes called an E1 line) provides 2.048 Mbit/s. SONET resolves such problems by defining how switches and multiplexers coordinate communications over lines with different speeds, including defining data transfer rates and frame format.

SONET defines a number of Optical Carrier (OC) levels. Each level defines an optical signal and a corresponding electrical signal called Synchronous Transport Signal (STS). The base level is OC-1/STS-1 or 51.84 Mbit/s. Each level's rate is a multiple of 51.84 Mbit/s. The table below shows the OC levels and the corresponding data transfer rates that SONET defines.

OC Level	Data Rate
OC-1	51.8 Mbit/s
OC-3	155.5 Mbit/s
OC-9	466.5 Mbit/s
OC-12	622.0 Mbit/s
OC-18	933.1 Mbit/s
OC-24	1.24 Gbit/s
OC-36	1.86 Gbit/s
OC-48	2.48 Gbit/s

SONET also provides easy access for low-speed signals, such as DS-0 (64 kbit/s) and DS-1 (1.544 Mbit/s) by assigning them to sub-STS-1 signals called Virtual Tributaries.

Advantages

The SONET standard defines data transfer rates and a frame format that all vendors and telephone companies throughout the world can use, creating the potential for global networking. SONET also includes management capabilities for telephone company equipment. Cell relay technologies such as Switched Multimegabit Data Services and ATM operate above SONET, making SONET the expected foundation for future broadband service.

Disadvantages

Some telephone companies are currently using SONET in their networks, but they are not yet offering it to the public on a tariffed basis. Unless your company is a large corporation in a metropolitan area, you probably cannot get dedicated SONET service. Also, some countries do not yet have a digital, fiber-optic telecommunications infrastructure, which means they cannot take advantage of SONET.

Global Networking

To conclude this networking primer, we will go beyond LANs and WANs and briefly discuss the exciting concept of global networking. There are four terms that currently encompass the present and future of global networking: the Internet, the intranet, the extranet, and the Information Superhighway. The Internet and the intranet are the main components of global networking today, and the extranet is rapidly becoming a common component. The Information Superhighway is a concept of what global networking could one day become—the global networking ideal.

The global networking ideal is the simple, powerful idea of people around the world connecting to a network on which they can share ideas, exchange information, and access endless electronic resources. Novell will play a major role in making the global networking ideal a reality. Although the ideal is still some years away, a basic form of global networking, based on the Internet and the intranet, exists today.

The Internet

The Internet is the world's largest computer network. Thousands of local networks belonging to organizations such as government agencies, universities, hospitals, private businesses, and other organizations in countries around the world are attached to the Internet. Millions of users send and receive E-mail, download and upload files, do research, and conduct business on the Internet daily.

The Internet is a global network, but in many ways, it does not currently meet the global networking ideal. From a business standpoint, for example, the Internet has several disadvantages. First, the Internet uses packet-switching, so you can never be sure what route a packet will take or how long it will take to arrive. Second, because no one owns the Internet, no one is responsible for ensuring that the network as a whole is functioning properly (or has the authority to require that it is). Third, while improvements have been made, security on the Internet is still a problem. Fourth, the Internet does not offer the fastest data transfer rates available.

Intranets

An intranet is a privately owned, secure, business network that is based on Internet technology.

The term "intranet" sprang up virtually overnight when companies discovered that they could use publicly available Internet technologies to make useful information immediately available to all employees, no matter where the employees were located; that they could still secure the information from unwanted access; and that, along with these other advantages, they could also make the information available at the lowest possible cost.

On a typical intranet, there is a World Wide Web server, on which information is published in an electronic format called Hypertext Markup Language (HTML). Workstations have some type of client software, most often a Web browser, through which they can access any information published (in HTML format) on any Web server. Users of client stations can be given different rights so that they can access only selected information on selected Web servers.

The main reason for a company to implement an intranet is that an intranet enables a business to collect, manage, and disseminate information more quickly and easily than ever before, even much more quickly and inexpensively than with other current means of electronic communications, including E-mail and other types of cross-platform publishing. In fact, intranet publishing is the ultimate in cross-platform publishing because it is based on the Internet technologies that were developed specifically for the purpose of allowing information sharing among dissimilar computing systems.

While even a small company with only one office and a small network can benefit from an intranet, the value of an intranet increases with the number of employees, the size of the network, and the number of geographically separate sites. The reason is that as a company grows, if the company continues to use conventional means of information dissemination, such as printed memoranda and newsletters, the cost of disseminating information to all employees increases exponentially. And other methods of sharing information, such as E-mail and file sharing, also fall short of the cost savings and immediacy that can be obtained through intranet publishing.

On an intranet, any employee with a properly configured workstation and a Web browser can read documents as soon as the files are completed and copied to any Web server, regardless of where the employee is located. If a company were to instead disseminate documents as files in a public directory or by E-mail, the documents would have to be provided in multiple formats to accommodate the various computing platforms and applications used within the company. There would need to be people dedicated to the task of preparing the differently formatted documents and disseminating them to different locations where they could be accessed. In even a small company, this type of effort takes significantly more time and costs far more than does publishing the same information once, in HTML format, on a single Web server. In a large company, the time and cost differences can be enormous.

Intranet publishing has other advantages. One important advantage is that the network can update your intranet documents automatically, in real time. For example, if you published a document that contained the stock price for your company or news about the market in which your company competes, you could create a Web server script that would automatically update the document every 15 minutes with the most current stock price and market news. With immediate access to up-to-date information, employees can respond more quickly to changes in the marketplace (with the result of increased profits). Also, after the script is created, the network continues to update the information—the work isn't forgotten or ignored because employees are too busy—and there is no further cost.

In addition, you can get immediate feedback about the documents published on your intranet. For example, with paper-based documents or publicly available files stored on a server, you cannot determine whether or not people are reading the documents. If you published the documents on an intranet server, however, the network could track how many people read the documents and which documents were used the most.

Businesses are continually finding more ways to use intranets to decrease costs, especially since the specification for World Wide Web documents has been extended to include graphics, audio clips, and movies. For example, many companies have installed applications that allow employees to access company databases directly from a Web browser, thus avoiding the cost of specialized database access programs. Recent products such as Novell’s GroupWise WebAccess even allow employees to read their E-mail messages and schedules directly from a Web browser.

Another factor that makes any intranet valuable is that after it is built, it can be connected to the Internet with very little extra effort. Remote users, such as traveling employees, suppliers, and customers, can then access your intranet documents over the Internet. You can control access to your intranet documents, allowing the general public to view some documents and allowing only authorized users to view others. Furthermore, you can allow employees to connect to the Internet and access a vast pool of information that covers nearly every topic imaginable.

Of course, intranets need not be connected to the Internet: An intranet may be only local, or, if it is a WAN intranet, the various locations might be connected by means other than the Internet. However, many intranets are now connected to the Internet, and in the future many more will be. The most important reason is that the Internet is a ready-made, low-cost WAN backbone. And, as mentioned above, if your intranet is connected to the Internet, all users can easily and speedily access the wealth of information available there.

Extranets

An extranet is two or more intranets connected in such a way that they enable collaboration among the businesses that own the separate intranets.

On an extranet, each connected company usually makes some selected part of its intranet accessible to the employees of one or more of the other companies. For example, several companies might create an extranet to consolidate data gathering and share data, or to jointly develop and share training programs and other material, or to coordinate project management for a common work project. On an extranet, each company uses the security inherent in its own intranet to the keep employees of other companies from accessing information they do not need to see.

The collaborative business application is a powerful extranet tool. Such applications, possibly developed jointly by participating companies, enable the employees of the different companies to work together very effectively without leaving their offices (which might be located in different places all over the world).

For example, a consumer company might work with a supply company to connect their intranets and create a supply ordering system, to allow all employees of the consumer company to order whatever supplies they needed, whenever they needed, directly from the supply company. A consumer company employee might order by using his or her Web browser to look through one or more electronic catalogs that the supply company published on the extranet. The employee might check a box next to each of the items he or she needed. Different employees might be given different rights to different catalogs so that they could see and order only from selected parts of a catalog. Also, different employees might be allowed to see different items in each part. Underlying parts of the collaborative business application could sort all ordered items by company division, group, and employee and fill out one daily purchase requisition containing all items ordered by all employees. Each purchase requisition could be immediately delivered over the extranet. For the supply company, the application could automatically generate a shipping ticket that contained the items to be shipped, broken down by division, group, and the person each item was to be delivered to.

For the consumer company, the end result might be to eliminate the need to stock any supplies and to considerably reduce purchasing costs. The consumer company employees might be able to get any supplies they needed in less time than ever before. And the supply company might be able to sell more supplies and deliver them faster than before, with less staff than before.

Because almost all intranets and extranets will eventually be connected to the Internet, intranet technology should be designed to deal as effectively as possible with the security problems and other problems inherent to the Internet. Thus, Novell is constantly working on new technologies such as IntranetWare Border Services, which you can read about in the Early Access Release section in this *Buyer’s Guide*.

The Information Superhighway

The terms “Internet” and “Information Superhighway” are sometimes used synonymously, but they are not the same thing. The term Information Superhighway describes the global networking ideal, which the Internet is not.

The ideal global network, or Information Superhighway, will include a vastly improved Internet and many other networks, services, and technologies. The Information Superhighway will be pervasive: Low-cost access will be available to virtually everyone worldwide. The Information Superhighway will provide homes, businesses, and other organizations with a myriad of services, such as on-demand video, E-mail, electronic commerce, shopping, research, video conferencing, and voting services. In sum, the Information Superhighway will provide literally every digitally deliverable service to everyone on the globe.

Presently, the Information Superhighway is only a concept, but governments, businesses, and public institutions worldwide are taking steps to make the Information Superhighway a reality.

Novell contributes to the growth of the Information Superhighway with all of its offerings, from core infrastructure and services such as IntranetWare and NetWare 4.11 operating systems; to advanced services such as Novell Directory Services, Novell Connect Services, NetWare Telephony Services™, and Novell Web Server 3.1; to Internet and intranet access products such as LAN WorkPlace® and LAN WorkGroup™ family of products; to advanced groupware and intranet applications such as GroupWise 5.1 and GroupWise WebAccess. You can read about all of these products and more in following sections of this *Buyer’s Guide*.

We hope this primer has been helpful to you. We welcome your comments and suggestions. Happy networking!

Primer Appendix

The table below illustrates the ASCII code set. The second table, ASCII Code-to-Character Conversion, gives the code-to-character mapping.

ASCII Code Set

Bit Settings			7	0	0	0	0	1	1	1	1
			6	0	0	1	1	0	0	1	1
			5	0	1	0	1	0	1	0	1
4	3	2	1	Resulting Codes							
0	0	0	0	0	16	32	48	64	80	96	112
0	0	0	1	1	17	33	49	65	81	97	113
0	0	1	0	2	18	34	50	66	82	98	114
0	0	1	1	3	19	35	51	67	83	99	115
0	1	0	0	4	20	36	52	68	84	100	116
0	1	0	1	5	21	37	53	69	85	101	117
0	1	1	0	6	22	38	54	70	86	102	118
0	1	1	1	7	23	39	55	71	87	103	119
1	0	0	0	8	24	40	56	72	88	104	120
1	0	0	1	9	25	41	57	73	89	105	121
1	0	1	0	10	26	42	58	74	90	106	122
1	0	1	1	11	27	43	59	75	91	107	123
1	1	0	0	12	28	44	60	76	92	108	124
1	1	0	1	13	29	45	61	77	93	109	125
1	1	1	0	14	30	46	62	78	94	110	126
1	1	1	1	15	31	47	63	79	95	111	127

In the ASCII coding scheme, information (a numeral, symbol, or alphabetic character) is represented by the value of a data unit called a “byte.” Each byte can represent one character. There are eight bits in one byte. Bits, short for binary digits, are the data units actually stored in the computer as either a one or a zero. Computers read stored bytes and interpret them as the codes that represent character-based information.

Sample Byte (Bit Settings)
8 7 6 5 4 3 2 1
1 1 0 0 0 0 0 1 = 65 = letter “A”
0 0 1 1 1 0 0 0 = 56 = number “8”

ASCII Code-to-Character Conversion

In the ASCII coding scheme, information (a number or alphabetic character) is represented by the value of a data unit called a “byte.” The following table shows possible byte values and the characters that each value is used to represent..

CODE	CHAR	CODE	CHAR	CODE	CHAR	CODE	CHAR	CODE	CHAR
0	NUL	26	SUB	52	4	78	N	104	h
1	SOH	27	ESC	53	5	79	O	105	i
2	STX	28	FS	54	6	80	P	106	j
3	ETX	29	GS	55	7	81	Q	107	k
4	EOT	30	RS	56	8	82	R	108	l
5	ENQ	31	US	57	9	83	S	109	m
6	ACK	32	SP	58	:	84	T	110	n
7	BEL	33	!	59	;	85	U	111	o
8	BS	34	“	60	<	86	V	112	p
9	HT	35	#	61	=	87	W	113	q
10	LF	36	$	62	>	88	X	114	r
11	VT*	37	%	63	?	89	Y	115	s
12	FF	38	&	64	@	90	Z	116	t
13	CR	39	‘	65	A	91	[	117	u
14	SO	40	(	66	B	92	\	118	v
15	SI	41	)	67	C	93	]	119	w
16	DLE	42	*	68	D	94	^	120	x
17	DC1	43	+	69	E	95	_	121	y
18	DC2	44	,	70	F	96	‘	122	z
19	DC3	45	-	71	G	97	a	123	{
20	DC4	46	.	72	H	98	b	124	|
21	NAK	47	/	73	I	99	c	125	}
22	SYN	48	0	74	J	100	d	126	~
23	ETB	49	1	75	K	101	e	127	DEL
24	CAN	50	2	76	L	102	f		
25	EM	51	3	77	M	103	g		