

CHAPTER 8

LOCAL-AREA NETWORKS

INTRODUCTION

A local-area network (LAN) is a communications system designed to transmit and receive digital information between computers. A LAN consists of **nodes** that are interconnected by **links**. Nodes are the hardware connected to the network, such as personal or microcomputers, printers, large capacity hard drives, and so on. Links are the communications media, such as twisted-pair wire, coaxial, or fiber-optic cables that connect the nodes. In most applications, the LAN interconnects a relatively small number of personal computers (PCs), data storage devices, printers, and other peripherals. These nodes and links usually cover a relatively small geographical area, such as an office or a department. Through common usage, the term *local-area network* can also refer to much larger systems, such as the SNAP III system on a ship, which could have literally hundreds of terminals and miles of cables. For our purposes, we will be using a small system in our discussion of LANs.

Any device connected to the network can send and receive data on the network. A majority of data exchanged over a network is text and graphics, which is assembled as structured data that can be manipulated by computers. Unstructured data, such as pictures and facsimile messages, can be stored and retrieved efficiently, but cannot be manipulated easily by the computer.

After completing this chapter, you should be able to:

- **Describe the major components of a LAN.**
- **State the types of cable used in a LAN.**
- **State the function of the network interface card.**
- **Describe the function of the various network servers required by a LAN.**
- **Describe the function of the central mass storage area of a LAN.**
- **Describe the Open System Interconnection (OSI) Reference Model used in the design and implementation of a LAN.**
- **Describe the advantages and disadvantages of the different LAN topologies.**
- **Describe the hardware systems used in LANs.**
- **Describe the function of the software operating system of a LAN.**

LOCAL-AREA NETWORK HARDWARE

The basic hardware components of a LAN are cables, network interface cards, network servers, peripherals, and workstations. These components are covered in the material that follows.

CABLES

Several types of cables can be used in LAN applications. The selection of the type of cable depends on several factors, such as maximum length of a single cable run, security requirements, and the capacity and speed of the system.

Twisted-Pair Cable

The twisted-pair cable is easy to install and costs little on a per-foot basis. In some cases, existing telephone cable may be used. Its disadvantages include limitations in capacity and speed. It is also susceptible to electrical interference unless it is shielded.

Shielded Twisted-Pair Cable

The shielded twisted-pair cable is encased in an RFI shield. The stranded wire used as a conductor is manufactured with greater precision and is capable of greater data transmission rates and longer cable runs.

Coaxial Cable

Coaxial cable networks have gained in popularity because of their use in cable television. The quantities of cable and connectors produced for cable television have greatly reduced the prices of these components for network users. Coaxial cable comes in various thicknesses and is designated by a number: RG-11, RG-58, RG-59, RG-62, and so forth. You can use either baseband or broadband transmission methods with coaxial cable.

Baseband coaxial systems transmit digital signals unchanged over a single channel and have several advantages. They are inexpensive, easy to install, and have low maintenance. They also allow very high

data transmission rates. One disadvantage is that they are limited to transmitting digital signals only.

In contrast, **broadband coaxial systems** require the digital signal to be converted to an analog signal before transmission and then back to digital by modem at the receiving device. Broadband systems support data, voice, and video signals that may be transmitted simultaneously. Disadvantages of broadband systems are their higher installation costs and complex maintenance.

Fiber-Optic Cable

Fiber-optic cable is the best choice if a secure network is needed. Because the cable transmits light, the transmissions are immune to interference caused by electrical or electronic devices. Also, if your network will run through an area of heavy industrial activity or a work place with strong radio frequency interference, fiber-optic cable is the most appropriate choice. Other advantages of the fiber-optic cable are that it lasts longer than other types of cable and can carry many more channels. Its disadvantages include its high price, poor connectivity, and low flexibility.

NETWORK INTERFACE CARD

To attach personal computers to the LAN, you must connect a network interface card (NIC) to each PC and attach the network cable to the NIC. Each individual workstation must have a network interface card. The NIC is nothing more than a circuit board that normally plugs directly into one of the expansion slots inside a PC. Sometimes, the NIC comes as a separate unit. In this case, you would plug it into the back of the PC. Most NICs have their own built-in microprocessor(s) designed to take care of network communications. This relieves the PC's main processor of this responsibility. The type of cable used on the network is determined by the type of LAN to be installed.

NETWORK SERVERS

Your understanding the concept of a **server** is important to understanding how LANs work. A server is a combination of hardware and software that

is used to manage the shared resources of the network. The hardware may be a PC or a computer designed specifically to act as a server. A **network server** is able to control network traffic as well as the sharing of other resources, such as application programs, disk space, data files, and printers. In either case, the computer normally has a hard disk and the software needed to run the network system. There are several different types of servers, and each has a particular function. In newer systems, some separate server functions are combined into a central file server. The servers we will look at are the disk server, the file server, and the print server.

Disk Server

The disk server was the first of the network operating systems. In the early days of PC networks, very few computers were equipped with a hard disk. When the hard disk became affordable, manufacturers were asked to develop a system to allow several users to share a single hard drive. The earliest disk servers were multiplexers that polled each connected computer for requests to write a file on the hard drive or to retrieve a file from the disk. The multiplexer then responded accordingly. A major problem with this process was that it did not allow for any type of security, data organization, or disk management.

As LAN technology evolved, the development of the disk server software in the early 1980s addressed some of these issues. The disk server is a software routine that was installed on each computer in the network. The disk server software allowed each PC to access the shared hard drive as if it were a local drive. In other words, the computer thought the drive was installed in the computer, but in reality, the drive was remotely located on the network.

The disk server also provided for some information sharing. One purpose of a network is to allow multiple users access to the same information. One problem encountered with early disk servers occurred when two or more users updated the same file at the same time. When the file was saved by both users, the updates of one of the users was lost.

A method of preventing this information loss is file locking. File locking means that when one user accesses a file, all the other users are prevented from accessing that file until the first user is finished with it. As you can see, this method severely limits the number of users able to access the information.

Another method used to prevent data loss is record locking. In a data-base environment, many users could access the same data file, but when a record was being modified by one user, the other users were locked out of the record being modified. A data file can be updated by several users without threatening the integrity of the data by using this technique.

Although the disk server was used in most LANs developed before 1985, a major problem still existed in maintaining data integrity. The two methods covered in the previous paragraphs provided for data management, but not for reliable disk management. A disk drive stores information on the next available block on the disk. When the disk server was used, it was not uncommon for two users to try to write data to the same block at the same time. When this happened, the second user would overwrite the data just written by the first user, causing a loss of data. The development of the file server in 1983 solved all of the problems encountered with the disk server.

File Server

Currently, all local-area networks require some type of file server. In most cases, the file server is a dedicated PC or minicomputer. The file server performs the processing of the network control software and the central processing and storage point of the application software and data files of the network. The file server has a hard disk with a very large storage capacity.

The file server manages the hard disk and ensures that multiple requests for the same file do not conflict with each other. In the disk server environment, each PC workstation manages its I/O with the disk through low-level sector calls. In the file server environment, each workstation communicates with the central disk through the use of high-level calls to the file server. A high-level call can be a request to open a particular file or to store a file, while a low-level call maybe to

write this file to sector *xyz* on the disk. The file server converts the high-level calls from the users to low-level disk commands, thus providing effective disk management. The file server maintains the list of privileges and authorizations for each user. This protects the data files from unauthorized access and protects the data. An example of this is that one user may be authorized total access to a data-base file to update the file, while another user may be authorized read-only access to the information. Still a third user may be denied access to the file altogether.

A network file server is a special-purpose unit that can reside in either a dedicated computer, or one of the workstations (a PC) that has a hard disk containing the software of the network. When the network server is used solely for serving the network and is **NOT** used as a workstation, it is referred to as a **dedicated server**. If the server can also be used as a workstation, it is referred to as a **nondedicated server**.

Some networks do **NOT** have a single dedicated file server. Instead, they use a **distributed** approach in which **any** of the nondedicated servers may make available files that reside on their hard disks. Under these circumstances, any workstation on the network can use or copy these files. Moving files back and forth on such a network establishes a temporary relationship, you might say, between the two PCs involved. One PC acts as the server, and the other PC acts as the receiving workstation.

Print Server

The print server is a software routine that allows all the workstations on the LAN to use a single printer. When the laser printer was introduced to the market, the extremely high-quality print and multiple fonts made it desirable for all correspondence. Unfortunately, the cost of a laser printer often exceeded that of an individual workstation and made it impractical for each workstation to have a dedicated printer. The print server solved that problem by accepting requests for print jobs from the network users and directing them to the printer. The print server makes sure one job is completed before a new job is started. Print server routines are included in

almost all network operating systems on the market today.

WORKSTATIONS

Workstations is another name for the PCs used on a network. The PCs can be of the same brand, such as Zenith, or they can be a combination of different brands, such as IBM, Zenith, Compaq, along with other PC compatible computers (clones). Each PC can be configured differently. Some might have their own hard disk drives; others might have expanded memory. Still others might **NOT** even have diskette drives or printer ports of their own. Instead, these less expensive workstations use the storage and printing resources available through the network. Even though a PC may be part of a LAN system, you can use it independently as a stand-alone PC at any time or you can use it as part of the LAN.

THE OPEN SYSTEM INTERCONNECTION (OSI) REFERENCE MODEL

Over the past few years, a number of network standards or protocols (rules to live by) have been developed by the International Standards Organization (ISO) to provide some level of uniformity among computer manufacturers and network vendors. ISO is one of several governing organizations in this field that has developed such protocols. ISO has defined a five level, seven-layer architecture. These seven layers of standards, shown in figure 8-1 define a generalized architecture called the **Reference Model of Open Systems Interconnection**. It is also known as the **OSI reference model** or **OSI model**. The primary purpose of the OSI model is to provide a basis for coordinating the development of standards that relate to the flexible interconnection of incompatible systems using data communications facilities.

The OSI model does **NOT** define anyone vendor's particular network software as such, nor does it define detailed standards for any given software. It simply defines the broad categories of functions that each of the seven layers should perform. The OSI model can include different sets of standards at each layer that are appropriate for given situations. For example, in a very simple data communications system, one that

TRANSPORT LEVEL

The next three layers of the OSI reference model make up the transport level, also known as the *subnet*. The transport level defines the software protocols necessary to exchange data on the network. The three layers of the transport level are the network layer, the transport layer, and the session layer.

Network Layer

The network layer decides which physical pathway the data should take based on network conditions, priorities of service, and other factors. Software on the network interface card must build the data packet, so the network layer can recognize and route the data to the correct destination address. It relieves the upper layers of the need to know anything about the data transmission and switching technologies used to connect the systems. It is responsible for establishing, maintaining, and terminating connections across the intervening communications facility as follows:

- Addresses messages
- Sets up the path between communicating nodes on possibly different networks
- Routes messages among networks
- Is concerned with the sequence delivery of data packets
- Controls congestion if too many packets are on the network
- Translates logical addresses or names into physical addresses
- Has accounting functions to count packets or bits sent by users to produce billing information

Transport Layer

The transport layer makes sure data units are delivered error-free, in sequence, without losses or

duplications. It relieves higher layer protocols from any concern with the transportation of data between them as follows:

- Message segmentation. Accepts data from the session layer, splits it up into smaller units, and passes the units down to the network layer
- Establishes and deletes host-to-host connections across the network
- Multiplexes several message streams onto one channel and keeps track of which message belongs to which connection
- Provides reliable end-to-end delivery with acknowledgment
- Provides end-to-end flow control and window management

Session Layer

The session layer allows users on different machines to establish sessions between one another. It performs the functions that enable two or more applications to communicate across the network, performing security, name recognition, logging, administration, and other similar functions. Unlike the network layer, this layer is dealing with the programs in each machine to establish conversations between them as follows:

- Allows two applications processes on different machines to establish, use, and terminate a connection (or session)
- Performs synchronization between end-user tasks by placing checkpoints in the data stream so that if the network fails, only the data after the last checkpoint has to be retransmitted
- Provides dialogue control (who speaks, when, how long, and so on)

PRESENTATION LEVEL/LAYER

The presentation level consists of the presentation layer. The presentation layer formats data to be presented to the application layer. It can be viewed as the translator for the network. This layer provides a common representation for data which can be used between the application processes. The presentation layer relieves the applications from being concerned with data representation, providing syntax independence as follows:

- Encodes data in a standard way (integers, floating point, ASCII, and so on)
- Provides data compression to reduce the number of bits that have to be transmitted
- Provides data encryption for privacy and authentication

APPLICATION LEVEL/LAYER

The final level is the application level, which consists of the application layer. The application layer serves as the window for the application process to access the OSI environment. This layer represents the services that directly support users and application tasks. It contains a variety of commonly needed protocols for the following items:

- Network virtual terminals
- File transfers
- Remote file access
- Electronic mail
- Network management

USING THE OSI MODEL

A communications system that does not use a layered architecture can be designed. A specifically designed communications system is faster, more efficient, requires less software code, and eliminates redundant functions and activities. Why, then, is the

OSI reference model considered the standard in designing networks and writing software? It is considered the standard primarily because the use of a layered architecture, such as the OSI reference model, provides the network with flexibility and migration.

The greatest advantage of your using layer architecture in a network is hardware independence. As advances in technology continue, it is not necessary to scrap a network completely because one component has been superseded. For example, if you have a network and need to upgrade the cable to a type that can handle increased data at a faster rate, the layered architecture of the OSI model will allow you to make this replacement to the physical layer without changing the other layers.

LAN TOPOLOGIES

The physical arrangement of the components of a LAN is called its configuration or topology. The three major types of configurations, or topologies, of a LAN are the **linear bus**, the **star**, and the **ring**. You can also create hybrid topologies by combining features of these configurations. For example, several bus networks can be joined together to form a ring of buses.

Each topology requires the components of a LAN to be connected in a different arrangement. These components are also referred to as nodes. A node is any point on a network where data can be sent (transmitted) or received—a workstation, a server, and so on.

LINEAR BUS NETWORK

The linear bus topology is like a data highway. That is, all components or nodes are connected to the same cable, and the far ends of this cable never meet, as shown in figure 8-2. Linear bus LANs are best suited to applications involving relatively low usage of the bus coupled with the need to pass relatively short messages from one node to another. In many such networks, the workstations check whether a message is coming down the cable before sending their messages. Since all nodes share the bus, all

messages must pass through the other workstations on the way to their destinations. Each node checks the address attached to the message to see if it matches its own address. Bus topologies allow individual nodes to be out of service or to be moved to new locations without disrupting service to the remaining nodes.

Because of the way linear bus cabling is laid out, this type of network is simple. The bus topology is very reliable, because if any node on the bus network fails, the bus itself is **NOT** affected, and the remaining nodes can continue to operate without interruption. Many of the low cost LANs use a bus topology and twisted-pair wire cabling.

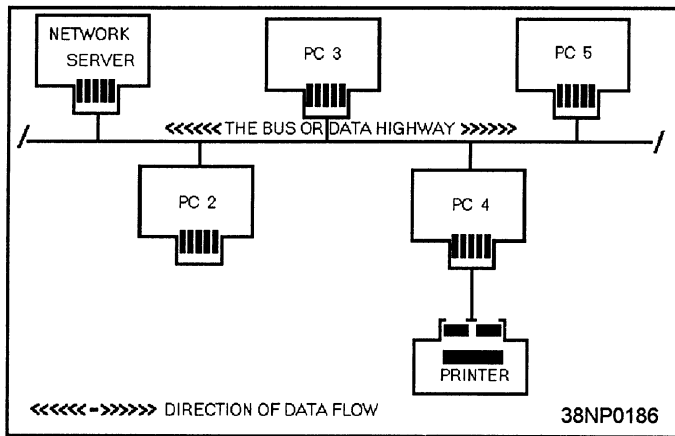


Figure 8-2.—A bus network topology.

A disadvantage of the bus topology is that generally there must be a minimum distance between workstations to avoid signal interference. Another disadvantage is that the nodes must compete with each other for the use of the bus. Simultaneous transmissions by more than one node are **NOT** permitted. This problem, however, can be solved by using one of several types of systems designed to control access to the bus. They are collision detection, collision avoidance, and token passing, which we will cover shortly. Also, there is no easy way for the network administrator to run diagnostics on the entire network. The bus network can be easily compromised by an unauthorized user, since all messages are sent along a common data bus. For this reason, it is difficult to maintain network security.

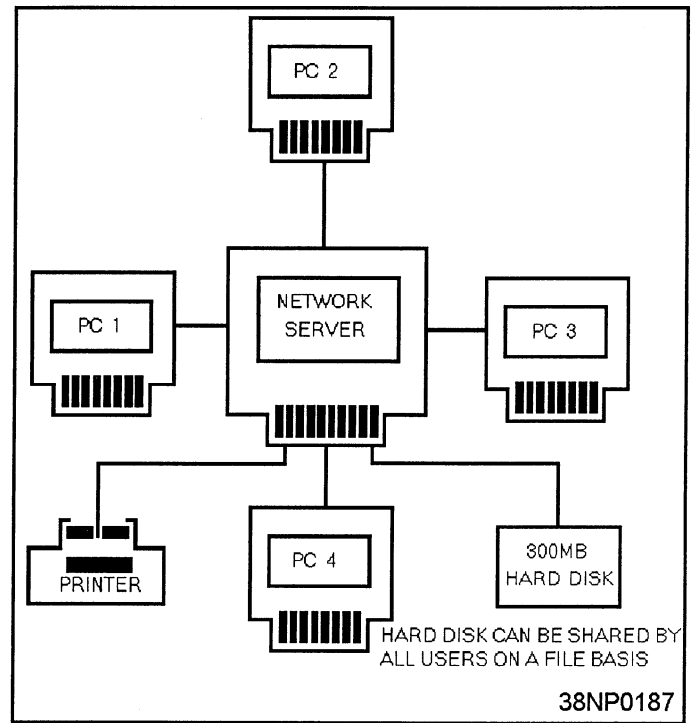


Figure 8-3.—A star network topology.

STAR NETWORK

In a star network, each component is connected directly to the central computer or network server, as shown in figure 8-3. Only one cable is required from the central computer to each PC's network interface card to tie that workstation to the LAN. The star is one of the earliest types of network topologies. It uses the same approach to sending and receiving messages as our phone system. Just as a telephone call from one person to another is handled by a central switching station, all messages must go through the central computer or network server that controls the flow of data. New workstations can be easily added to the network without interrupting other nodes. This is one of the advantages of the star topology.

Another advantage of star topology is that the network administrator can give selected nodes a higher priority status than others. The central computer looks for signals from these higher priority workstations before recognizing other nodes. The star topology also permits centralized diagnostics (troubleshooting) of all functions. It can do this because all messages must first go through the central

computer. This can prove invaluable in making sure that network security has not been breached.

The main disadvantage of the star topology is its reliance on the central computer for performing almost all the functions of the network. When the central computer fails, all nodes also stop functioning, resulting in failure of the entire network.

DISTRIBUTED STAR

The distributed star, or tree, topology is shown in figure 8-4. It provides many of the advantages of both bus and star topologies. It connects workstations to a central point called a hub. This hub can support several workstations or hubs which, in turn, can support other workstations. Distributed star topologies can be easily adapted to the physical arrangement of the facility site. If the site has a high concentration of workstations in a given area, the system can be configured to more closely resemble a star topology. If the workstations are widely dispersed, the system can use inexpensive hubs with long runs of shared cable between hubs, similar to the bus topology.

RING NETWORK

In a ring network, all the components or nodes are connected to the main cable, and the cable forms a ring, as shown in figure 8-5. This topology allows a node to send a message to another node on the ring.

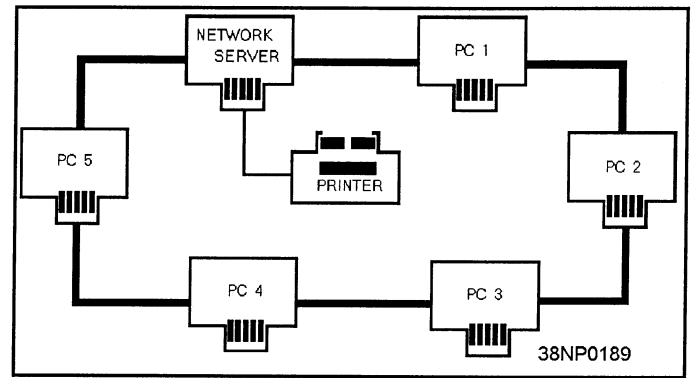


Figure 8-5.—A ring network topology.

However, the message must be transmitted through each node until it reaches its destination. Messages proceed from node to node in one direction only. Should a node fail on the network, data can no longer be passed around the ring unless the failed node is either physically or electronically bypassed. Using bypass software, the network can withstand the failure of a workstation, by bypassing it, and continue to maintain the integrity of the network. One of the major issues in a ring topology is the need for making sure all workstations have equal access to the network.

One of the major disadvantages of ring topologies is the extreme difficulty of adding new workstations while the network is in operation. Normally the entire network has to be brought down while a new node is added and cabling reattached. However, this particular problem can be overcome by the installation




Figure 8-4.—A distributed star (tree) network topology.

of additional connectors when the network is initially set up. These connectors enable you to add or remove

nodes while the network remains intact and in operation.

ACCESS METHODS


Another decision the designer makes is that of which access method to use. Access methods are the arrangements used to make sure each workstation has fair and equal access to the network. The access method used is governed primarily by the topology and the protocol of the network. The principal access methods are contention and token passing.

Contention

The contention method features Carrier Sense Multiple Access (CSMA) and Carrier Sense Multiple Access with Collision Detection (CSMA/CD). The CSMA/CD method is shown in figure 8-6. Access for both is on a first-come, first-served basis. The CSMA access scheme is very similar to that of a citizens band (CB) radio. Stations with data to send listen to the channel and wait until it is clear to transmit. With CSMA/CD, when two or more workstations transmit simultaneously, their messages will collide. As soon as a workstation detects a collision, it ceases transmission, monitors the network until it hears no other traffic, and then retransmits. Most contention networks assign a unique retry algorithm to vary the wait-and-retry period. This algorithm reduces the likelihood that after a collision, two workstations will transmit retries simultaneously.

Token Passing

Token passing is an orderly access method and is shown in figure 8-7. Each workstation passes on the opportunity to transmit to its closest neighbor until a station is found with a message to send. This permission to transmit is called a token. When a workstation with data to send is handed a token, part of the token is changed, indicating it is carrying a message, and then data is transmitted with the token. The token is then passed around the network, and every station checks whether the message is intended for it. The receiving station copies the message from



38NP0190

Figure 8-6.—A bus network using the CSMA/CD access method.

the token, but then passes the unchanged token along the network. When the transmitting station receives the same token, it knows the message has been passed around the network. The transmitting station erases the message and puts the empty token back into circulation on the network. The amount of information that may be transmitted during possession of the token is limited so that all workstations can share the cable equally.

PROTOCOLS

Network protocols are an important component because they define how networks establish

communications between element, exchange information, and terminate communications. Protocols have two major operational functions. They establish the circuit for transmission (handshaking) and for the transmission itself. Transmission is conducted subject to the line discipline. The line discipline is the sequence of operations that actually transmits and receives the data, handles the error-control procedures, handles the sequencing of message blocks, and provides for validation for information received correctly.

Two representative protocols, which control line discipline, are the Binary Synchronous communications Protocol (Bisync) and the Synchronous Data Link Control (SDLC).

Bisync

Bisync is a half-duplex protocol that transmits strings of characters at lower speeds over dial-up circuits. The information movement is in one direction at a time, with each data transfer being answered by an acknowledgement.

SDLC

SDLC is a control procedure that sends multiple blocks of data and returns a single acknowledgement for many blocks, thereby increasing the amount of time spent transmitting data. The bits that are put before and after the message at the transmitting end are removed at the receiving end, so only the message is presented to the user.

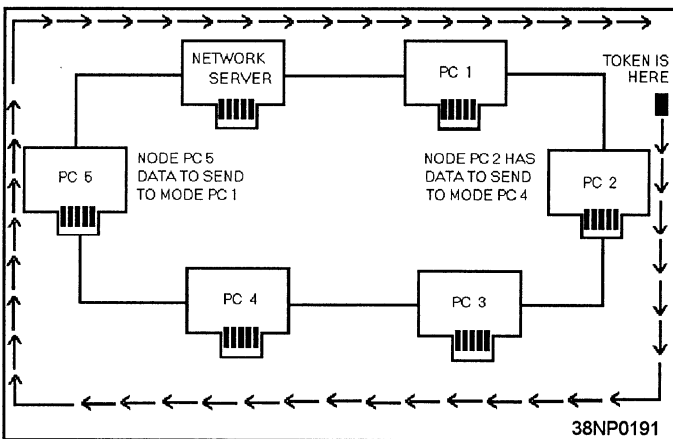


Figure 8-7.—A ring network using the token passing access method.

The hardware chosen for the network plays a part in the choice of network protocol. Most users and many of the vendors who build the clone type of equipment would like to see universal interfaces, while others think that the availability of different specifications will lead to a proprietary set of equipment, even though they favor the overall OSI specifications.

LAN SYSTEMS

When you decide to install a LAN system, the type of topology used in the initial wiring of the system will have a major effect on the type of system that can be used. There are many LAN systems available, each with advantages and disadvantages. In the following paragraphs, we briefly examine some of the available LAN systems.

The Institute of Electrical and Electronics Engineers (IEEE) has developed a set of standards for local-area networks. These standards encourage the use of common approaches for LAN protocols and interfaces. The IEEE LAN standards were developed by a committee of engineers and classified as the 802 standards. The 802 standards are broken down even further to define the protocols and topology used in a LAN. Some of the standards we are concerned with are the following:

- IEEE 802.3—Carrier sense multiple access/collision detection (CSMA/CD)
- IEEE 802.4—Token Bus
- IEEE 802.5—Token ring

ETHERNET

The EtherNet local-area network was developed by Xerox, the Intel Corporation, and the Digital Equipment Corporation. It became the model for the development of the IEEE 802.3 standard. The original standard defined a maximum throughput for EtherNet of 10 Mbit/s, but it has been revised to support throughput of much higher rates. When operating over coaxial cable, EtherNet has a 20-Mb per second throughput speed. For high-demand

environments, such as engineering or graphics, EtherNet is often the choice. It is a bus topology and uses CSMA/CD protocol. It is available in the following three versions: standard EtherNet, ThinNet, and twisted-pair EtherNet.

Standard EtherNet and ThinNet both use coaxial cable. Standard EtherNet is somewhat more expensive and more difficult to install than ThinNet, but it allows networking over greater distances with more users. Twisted-pair EtherNet uses a distributed star topology with wiring concentrator hubs, not the bus topology characterizing standard EtherNet and ThinNet. Connecting more than 100 users on a standard EtherNet trunk or on a series of twisted-pair concentrators is not uncommon, while ThinNet LANs usually support less than 50 users.

All versions of EtherNet create a LAN with high interconnectivity options. A number of products are available for connecting EtherNet LANs to minicomputers and mainframe computers and for bridging to other LANs; examples are STARLAN, ARCnet, and IBM Token Ring Network.

STARLAN

STARLAN uses a star topology with a CSMA/CD protocol. Its throughput speed is 1Mb per second over twisted-pair cable. If buildings are already wired with twisted-pair cable meeting AT&T premise cabling specifications, STARLAN can be installed easily. It is considered to be a low cost-per-user network and its star topology makes it a flexible network.

ARCnet

ARCnet is a distributed star topology that uses a token passing protocol and either twisted-pair or coaxial cabling. Its throughput speed is 2.5Mb per second. Although ARCnet does not conform to an IEEE standard, it closely resembles the 802.4 standard for a token bus system. It can easily handle up to 75 users. If user demand is low, it can handle additional users. It is considered an extremely reliable network and it is easy to install, expand, and modify.

IBM Token Ring Network

The IBM Token Ring Network uses a star ring topology, and is defined by the IEEE 802.5 specification. It has a throughput speed of 4 Mbits per second and 16 Mbit per second. This allows for flexible expansion of very large networks. Because of its speed and token passing protocol, it is a good choice to meet high-volume requirements. It is a sophisticated LAN technology developed by IBM to be used with an IBM cabling system and it is currently the fastest growing installed network base. The star ring topology also makes use of redundant circuits and loopbacks to handle breaks in the ring and results in high-fault tolerance on the network.

NETWORK OPERATING SYSTEMS

Network operating system software is necessary to control the overall operations of the network. Careful consideration must be given to the various packages on the market to ensure the operating software is fully compatible with your system, topology, and needs.

NETWORK OPERATING SYSTEM BASICS

The two basic components of the network operating system are the network operating system server and the workstation. The network server is usually a dedicated computer that runs the operating system software and processes all requests for services. The workstation computer runs the application software needed by the workstation user and establishes communications with the network server.

The network server operating system consists of the following five subsystems: the **control kernel**, the **network interfaces**, the **file systems**, the **system extensions**, and the **system services**.

Control Kernel

The control kernel is the main subsystem of the network operating software. The control kernel coordinates the various processes of the other subsystems. Some of the functions performed by the control kernel are as follows:

- Optimizing access to services by users

- Maintaining status information of many of the processes
- Error reporting
- Service initialization and service termination of workstations

Network Interfaces

The network interfaces provide the low-level subnet protocols and basic translation for bridging hardware drivers with the network operating system. In sophisticated systems, the network interfaces can also provide for bridging a new network into an operating network without having to rebuild the operating system.

File Systems

The file system controls the way the data is organized, stored, and retrieved from the storage systems available to the network. The files may be stored on hard drives, RAM disks, or optical storage devices, such as CD-ROM or write once, read many (WORM) drives.

File systems are generally designed to provide universal applicability. This means that the file system can be compatible with any application program's expectation of file input/output protocol. When adaptable interfaces are used, the file system can appear to emulate a number of different file systems.

System Extensions

The system extensions define the openness of the network operating system and are used by third party developers to produce add-on products. The extensions are usually high-level protocol handlers that perform operations, such as file access protocol translations required by different operating systems. The extensions available also include network management, system tools, and data-base services.

System Services

Network system services contain all services that are not easily defined by any of the other areas of the network. Examples of network services are security, system reliability features, error conditions, and access violations.

NETWORK OPERATING SYSTEMS SOFTWARE

The most important job of a network operating system (NOS) is to provide file service for the attached computers. This allows information retrieval and usage and the storage of data in a shared environment. A NOS manages the other resources shared by the network and provides the following functions:

- Directory structure for shared hard disk storage devices
- File service for sharing and using data
- Interface to the network for application software/programs
- The means by which the network manager manages the network and its users
- Network security and data protection
- Communications with other networks

The types of network operating systems include full-featured, low-cost, and zero-slot operating systems.

Full-Featured Network Operating Systems

Most full-featured network operating systems allow for high performance, flexibility, and excellent security measures. They require a LAN administrator. They require network interface cards. Also, they can be quite costly. Examples of a full-featured NOS are EtherNet, Novell's NetWare, 3Com's 3+Share, IBM Token Ring Network, and Banyan's Vines.

Low-Cost Network Operating Systems

Most low-cost network operating systems differ from full-featured systems only in the maximum number of users accommodated on the network and the number of security levels incorporated into the operating system. In general, they are much lower in cost and are easier to install and use. Examples of low-cost systems are STARLAN, ARCnet, 10Net, and LANtastic.

Zero-Slot Network Operating Systems

Zero-slot network operating systems are appropriate only for networks with very few users and light usage. They are an inexpensive and simple alternative to the NOSs that require expensive network interface cards. Rather than requiring each workstation computer to have a NIC, the computer's RS-232 serial communications port and twisted-pair cables are used. Because of this, they are very slow and offer limited file transfer abilities. They may not provide disk sharing. An example of a zero-slot system is LANLink.