The World of Computer Networking

In the last 15 years, LANs have gone from being an experimental technology to becoming a key business tool used by companies worldwide.

The World of Computer Networking

In the last 15 years, **LANs** have gone from being an experimental technology to becoming a key business tool used by companies worldwide. A LAN is a high-speed communications system designed to link computers and other data processing devices together within a small geographic area such as a workgroup, department, or a single floor of a multistory building. Several LANs can also be interconnected within a building or campus of buildings to extend connectivity.

Some Background on LANs

LANs have become popular because they allow users to share vital computing resources electronically, including expensive hardware such as printers and CD-ROM drives, application programs, and, most importantly, the information the users need to do their jobs. Prior to the development of LAN technology, individual computers were isolated from each other and limited in their range of applications. By linking these individual computers over LANs, their usefulness and productivity have been increased enormously. But a LAN by its very nature is a local network, confined to a fairly small area such as a building or even a single floor of a building. To realize the full benefit of computer networking, it is critical to link the individual LANs into an enterprise-wide backbone network that connects all of the company's employees and computing resources, no matter how geographically dispersed they may be.

Today's LANs and **LAN internetworks** are powerful, flexible, and easy to use, but they incorporate many sophisticated technologies that must work together flawlessly. For a LAN to really benefit an organization it must be designed to meet the organization's changing communications requirements. Building a LAN is a process of choosing different pieces and matching them together. This primer is designed to help first-time LAN equipment buyers and users understand the fundamentals of how LANs operate, what the different technology choices are for building a LAN, and the ramifications of choosing one option over another. Also discussed is the concept of internetworking or connecting disparate and geographically dispersed LANs together to form an enterprise system, the different technologies and products available to do so, and the benefits and limitations of each.

To aid readers unfamiliar with networking terminology, terms in boldface appear in a Glossary in the back of this primer.

Figure 1 | **A Basic LAN Bus Network**

When Station B sends a packet to another station on the LAN, it passes by all of the stations connected to that LAN. On the bus network illustrated here, the electrical signal representing the packet travels away from the sending station in both directions on the shared cable. All stations will see the packet, but only the station it is addressed to will pay attention to it.

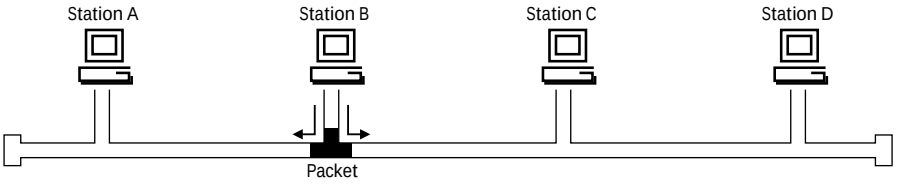

The Basics of Local Area Networking

Today local area networking is a **shared access** technology. This means that all of the devices attached to the LAN share a single communications medium, usually a coaxial, twisted pair, or fiber optic cable. Figure 1 illustrates this concept: Several computers are connected to a single cable that serves as the communications medium for all of them. The physical connection to the network is made by putting a **network interface card (NIC)** inside the computer and connecting it to the network cable. Once the physical connection is in place, it is up to the network software to manage communications between stations on the network.

In a shared media network, when one station wishes to send a message to another station it uses the software in the workstation to put the message in an "envelope." This envelope, called a **packet**, consists of message data surrounded by a **header** and **trailer** that carry special information used by the network software to the destination station. One of the pieces of information placed in the packet header is the address of the destination station.

The NIC then transmits the packet onto the LAN. The packet is transmitted as a stream of data bits represented by changes in electrical signals. As it travels along the shared cable, all of the stations attached to it see the packet. As it goes by the NIC in each station, the NIC checks the destination address in the packet header to determine if the packet is addressed to it. When the packet passes the station it is addressed to, the NIC at that station copies the packet and then takes the data out of the envelope and gives it to the computer.

Figure 1 shows one source station sending a single message packet to one destination station. If the message the source station wants to send is too big to fit into one packet, it will send the message in a series of packets. On a shared access LAN, however, many stations all share the same cable. Since each individual packet is small, it takes very little time to travel to the ends of the cable where the electrical signal dissipates. So after a packet carrying a message between one pair of stations passes along the cable, another station can transmit a packet to whatever station it needs to send a message. In this way, many devices can share the same LAN medium.

Ethernet

The most widely used LAN technology in use today is **Ethernet**. It strikes a good balance between speed, price, ease of installation, and supportability. Approximately 80 percent of all LAN connections installed use Ethernet.

The Ethernet standard is defined by the Institute of Electrical and Electronics Engineers (IEEE) in a specification commonly known as **IEEE 802.3**. The 802.3 specification covers rules for configuring Ethernet LANs, the types of media that can be used, and how the elements of the network should interact. The Ethernet protocol provides the services called for in the Physical and Data Link Layers of the **OSI reference model** (please refer to the "Standards and Protocols" sidebar).

One element of the 802.3 specification states that Ethernet networks run at a data rate of 10 million bits per second (10 Mbps) or 100 million bits per second (100 Mbps) in the case of Fast Ethernet. This means that when a station transmits a packet onto the Ethernet medium it travels along that medium at 10 Mbps.

Another important element defined by the 802.3 specification is the access method to be used by stations connected to an Ethernet LAN, called **carrier sense multiple access with collision detection (CSMA/CD)**. In this method, each station contends for access to the shared medium. It is possible for two stations to try sending packets at the same time, which results in a **collision** on the LAN. In Ethernet networks, collisions are considered normal events and the CSMA/CD access method is designed to quickly restore the network to normal activity after a collision occurs.

Ethernet Media and Topologies An important part of designing and installing a LAN is selecting the appropriate medium and topology for the environment. Ethernet networks can be configured in either a star or bus topology and installed using any of three different media.

Coaxial cable was the original LAN medium and it is used in what is called a **bus topology** (see Figure 1 for a typical bus topology). In this configuration, the coaxial cable forms a single bus to which all stations are attached. This topology is rarely used in new LAN installations today because it is relatively difficult to accommodate adding new users or moving existing users from one location to another. It is also difficult to troubleshoot problems on a bus LAN unless it is very small.

Figure 2 illustrates a star topology LAN — which is a more robust topology. In a star topology, each station is connected to a central wiring concentrator, or hub, by an individual length of twisted pair cable. The cable is connected to the station's NIC at one end and to a port on the hub at the other. The hubs are placed in wiring closets centrally located in a building.

Ethernet networks can be built using three different types of media: shielded and unshielded twisted pair, coaxial, and fiber optic cables. By far the most common is twisted pair because it is associated with the more popular star topology. It is inexpensive, and very easy to install, troubleshoot, and repair. Twisted pair cable comes both unshielded and shielded. **Unshielded twisted pair (UTP)** cable used for LANs is similar to telephone cable but has somewhat more stringent specifications regarding its susceptibility to outside **electromagnetic interference (EMI)** than common telephone wire. **Shielded twisted pair (STP)**, as its name implies, comes with a shielding around the cable to provide more protection against EMI.

Of the two types of twisted pair cable, UTP is by far the most commonly used. The specification for running Ethernet on UTP is called **10BASE-T**. This stands for 10 Mbps, baseband signaling (the signaling method used by Ethernet networks), over twisted pair cable. Other Ethernet specifications include **10BASE5**, which uses a thick coaxial cable, and **10BASE2**, which uses a thin coaxial cable media. Today, 10BASE5 is seldom installed in new Ethernet networks, and 10BASE2 is used only in very small office networks. An additional standard allows 10BASE-F Ethernet to run on fiber optic cable.

Fast Ethernet An extension of the popular 10BASE-T Ethernet standard, Fast Ethernet transports data at 100 Mbps. With rules defined by the IEEE 802.3u standard, Fast Ethernet leverages the familiar Ethernet technology and retains the CSMA/CD protocol of 10 Mbps Ethernet. Two types of Fast Ethernet are available: 100BASE-TX, which runs over Category 5 UTP; and 100BASE-FX, which operates over multimode fiber optic cabling.

Figure 2 | **Basic Star Topology LAN**

In a star topology all stations are wired to a central wiring concentrator called a hub. Similar to a bus topology, packets sent from one station to another are repeated to all ports on the hub. This allows all stations to see each packet sent on the network, but only the station a packet is addressed to pays attention to it.

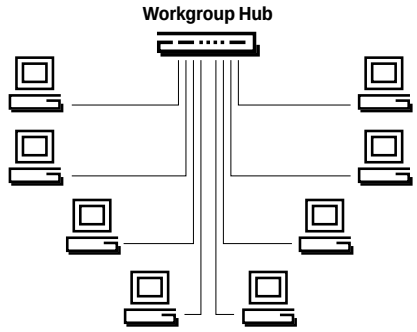

Token Ring

Another major LAN technology in use today is **Token Ring**. Token Ring rules are defined in the **IEEE 802.5** specification. Like Ethernet, the Token Ring protocol provides services at the Physical and Data Link Layers of the OSI model. Token Ring networks can be run at two different data rates, 4 Mbps or 16 Mbps.

The access method used on Token Ring networks is called **token passing**. Token passing is a deterministic access method in which collisions are prevented by assuring that only one station can transmit at any given time. This is accomplished by passing a special packet called a **token** from one station to another around a ring. A station can only send a packet when it gets the free token. When a station gets a free token and transmits a packet, it travels in one direction around the ring, passing all of the other stations along the way. As with Ethernet, the packet is usually addressed to a single station, and when it passes by that station the packet is copied. The packet continues to travel around the ring until it returns to the sending station, which removes it and sends a free token to the next station around the ring.

Token Ring Topology and Media Token Ring networks use what is called a **ring topology**. However, it is actually implemented in what can best be described as a **collapsed ring** that looks like a physical star topology (see Figure 3). In Token Ring LANs, each station is connected to a Token Ring wiring concentrator, called a **multistation access unit (MAU)**, using an individual run of twisted pair cable. Like Ethernet hubs, MAUs are located in wiring closets.

Figure 3 | **Basic Ring Toplogy LAN**

The ring topology used in Token Ring networks is a collapsed ring that looks like a physical star. Each station is connected to a Token Ring wiring connector by a single twisted pair cable with two wire pairs. One pair serves as the "inbound" portion of the ring (also known as the receive pair) and the other pair serves as the "outbound" or transmit pair.

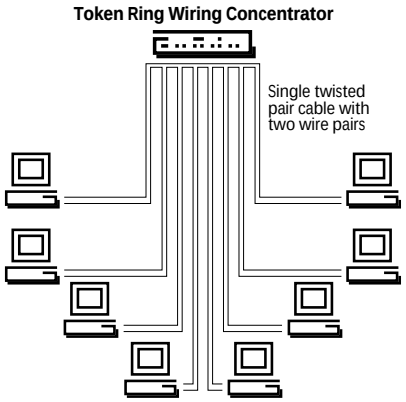

FDDI

Fiber Distributed Data Interface, commonly known as FDDI, provides data transport at 100 Mbps, a much higher data rate than Ethernet or Token Ring. Originally, FDDI networks required fiber optic cable, but today they can be run on UTP as well. Fiber is still preferred in many FDDI networks because it can be used over much greater distances than UTP cable. Like Token Ring, FDDI uses a token passing media access method. It is also usually configured in a collapsed ring, or physical star, topology. FDDI is used primarily as a **backbone**, a segment of network that links several individual workgroup or department LANs together in a single building. It is also used to link several building LANs together in a campus environment.

Structured Wiring Both the Ethernet star topology and the "collapsed ring" topology used in Token Ring LANs are supported by what is called a **structured wiring architecture**. With structured wiring, all of the network stations are physically star wired to **intelligent hubs**. Intelligent hubs are hubs that can be monitored and managed by network operators. This combination of a star topology and intelligent hubs make troubleshooting and fault isolation easier and faster because each endstation is attached to the network on its own individual port, which means it can be monitored easily and, if necessary, can be easily turned off. In addition, structured wiring makes adding users to the network, moving them, or making other physical changes on the network very simple. Since both Ethernet and Token Ring networks can use twisted pair cable and can be configured in a physical star topology, a structured wiring architecture will support either network technology.

Hubs: The Central Connection Point

The **hub** is one of the most important elements of a LAN. It is a central connection point for wiring the network (see Figure 4), and all stations on the LAN are linked to each other through the hub. The term hub is generally associated with 10BASE-T Ethernet networks, while the term multistation access unit (MAU) is used to refer to the Token Ring wiring concentrator. Just as these two LAN technologies use different media access methods, hubs and MAUs perform different media access functions internally, but at one level they perform the same function: They are both network wiring concentrators.

A typical hub has multiple user ports to which computers and peripheral devices such as servers are attached. Each port supports a single 10BASE-T twisted pair connection from a network station. When an Ethernet packet is transmitted to the hub by one station, it is **repeated**, or copied, over onto all of the other ports of the hub. In this way, all of the stations "see" every packet just as they do on a bus network, so even though each station is connected to the hub with its own dedicated twisted pair cable, a hub-based network is still a shared media LAN — picture it as a LAN in a box.

Manageable Hubs Intelligent hubs have been defined as **manageable hubs**, meaning that each of the ports on the hub can be configured, monitored, enabled, or disabled by a network operator from a hub management console. Hub management can also include gathering information on a variety of network parameters, such as the numbers of packets that pass through the hub and each of its ports, what types of packets they are, whether the packets contain errors, and how many collisions have occurred. Each hub vendor has some type of management package it sells with its products. These applications vary in how much information they can gather, what commands can be issued, and how the information is presented to the network operator.

Standalone Hubs Both hubs and MAUs come in three configurations: **standalone hubs**, **stackable hubs**, and **modular hubs**. Some products are combinations of the best configurations. Standalone hubs are — as the term implies — single box-level products with a number of ports. Standalone hubs usually include some method of linking them to other standalone hubs — either by connecting them together with a length of 10BASE5 coaxial cable or cascading them using twisted pair between individual ports on each hub (see Figure 5). Standalone hubs are usually the least expensive type of hub and are often not managed. They are best suited for small, independent workgroups, departments, or offices typically with fewer than 12 users per LAN.

Figure 4 | **Basic LAN with the Hub as the Central Connection Point**

The cornerstone of the network is the intelligent hub, or concentrator, which serves as the control point for systems activity, management, and growth. By integrating any combination of connectivity, internetworking, and management capabilities into intelligent hubs, network managers can create the perfect physical network infrastructure for their environment.

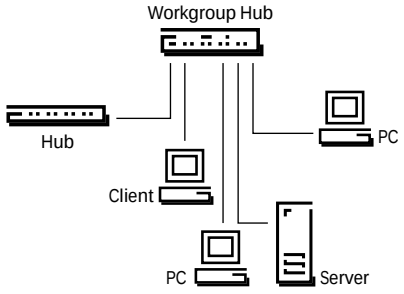

Figure 5 | **Summary of Network Architectures**

Network A illustrates four 10BASE-T hubs connected together by a single cable. This cable could be a coaxial or an optical fiber cable. All of the hubs are part of a single LAN. Network B illustrates two 10BASE-T hubs cascaded. Here the cable connecting the two ports is unshielded twisted pair wire. All of the hubs that are cascaded in this fashion are part of a single LAN.

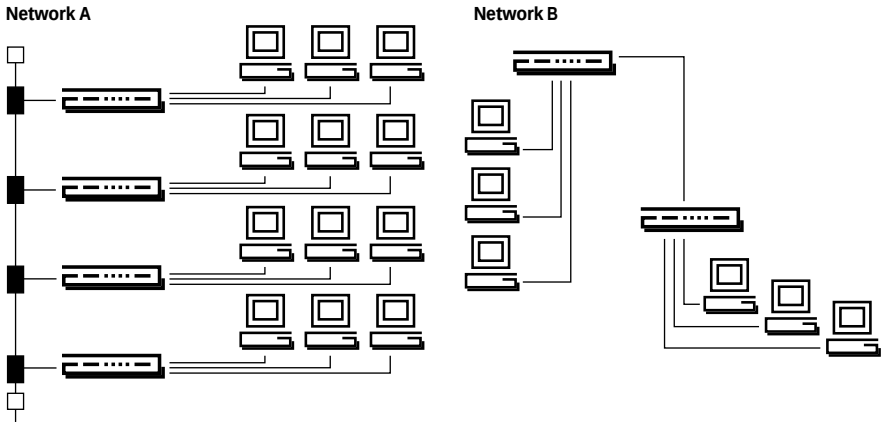

Stackable Hubs A third type of hub is the stackable hub. Stackable hubs look and act like standalone hubs except that several of them can be stacked or connected together, usually by short lengths of cable. When they are linked together, they act like a modular hub in that they can be managed as a single unit. One manageable hub, used within a stack, can typically provide the management for all other hubs in the stack. These hubs are ideal when an organization wants to start with a minimal investment but knows that its LAN will grow. By utilizing stackable hubs, an organization doesn't need to invest in a large chassis, which may only have one or two cards in it for a considerable length of time until more are needed.

Linking Hubs Each hub usually represents a single LAN. In most organizations it is desirable to interconnect all of the LANs, which means linking hubs in some way. One way to link hubs is to use an **interrepeater link** or cascaded segment. This type of connection simply repeats all of the packets from one hub to the other hub it is linked to, so that in effect the two hubs are part of the same LAN.

Modular Hubs Modular hubs are popular in networks because they are easily expanded and always have a management option. A modular hub starts with a chassis, or card cage, with multiple card slots, each of which accepts a communications card, or module. Each module acts like a standalone hub; when the communications modules are placed in the card slots in the chassis, they connect to a communications backplane that links them together so that a station connected to a port on one module can easily communicate with a station on another module. Figure 6 illustrates a modular hub.

Modular hubs typically range in size from four to 14 slots, so the network can be easily expanded. Typically, several slots in a modular hub will be filled with 10BASE-T Ethernet modules. For instance, with 10 modules, each supporting 12 users, a single hub could support 120 users over 10BASE-T. The modules are linked by the high-speed backplane, which can also be used to connect the communications modules to a management module that manages all of the cards in the chassis. In addition to using one management module for a large number of ports, all of the modules share a common power supply. Another advantage of some modular hubs is that Ethernet, Token Ring, and even FDDI communications modules can be placed in the same chassis, using the same common power supplies.

Figure 6 | **Modular Hubs**

Modular hubs provide a central point where multiple concentrators located in different wiring closets can be united into a LAN or WAN. The modular hub can be equipped with a wide variety of connectivity and network management modules designed to provide a customized solution for the creation of enterprise-wide LANs and WANs.

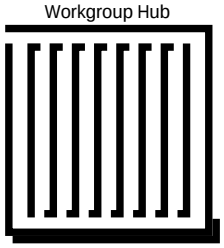

Internetworking

The term **internetworking** refers to linking individual LANs together to form a single internetwork. This internetwork is sometimes called an enterprise network because it interconnects all of the computer networks throughout the entire enterprise. Workgroup LANs on different floors of a building or in separate buildings on a business campus can be linked together so that all of the computing systems at that site are interconnected. Geographically distant company sites can also be tied together in the enterprise-wide internetwork.

An individual LAN is subject to limits on such things as how far it can extend, how many stations can be connected to it, how fast data can be transmitted between stations, and how much traffic it can support. If a company wants to go beyond those limits — link more stations than that LAN can support, for example — it must install another LAN and connect the two together in an internetwork.

There are two main reasons for implementing multiple LANs and internetworking them. One is to extend the geographic coverage of the network beyond what a single LAN can support — to multiple floors in a building, to nearby buildings, and to remote sites. The other key reason for creating internetworks is to share traffic loads between more than one LAN. A single LAN can only support so much traffic. If the load increases beyond its carrying capacity, users will suffer reduced throughput and much of the productivity achieved by installing the LAN in the first place will be lost. One way to handle heavy network traffic is to divide it between multiple internetworked LANs.

There are three major types of devices used for internetworking: **bridges**, **routers**, and **switches**. Today the most commonly used internetworking devices are high-speed routers, especially in wide area internetworks linking geographically remote sites. But routers are also heavily used in building and campus internetworks. Bridges have also been popular, even though they offer less functionality than routers, because they are less expensive to purchase, implement, and maintain.

LAN switches are a new class of internetworking device, and many people believe that switched internetworks will become the most common design for linking building and campus LANs in the future. Today's LAN switches and switching hubs are the first steps on a migration path to something called **asynchronous transfer mode (ATM)** switching, an emerging technology that will be widely implemented in both LANs and wide area networks in the coming years.

Bridges and Routers

Bridges and routers are both special kinds of devices used for internetworking LANs — that is, linking different LANs or LAN segments together. Many organizations have LANs located at sites that are geographically distant from each other. Routers were originally designed to allow users to connect these remote LANs across a wide area network, but bridges can also be used for this purpose. By placing routers or bridges on LANs at two distant sites and connecting them with a telecommunications link, a user on one of the LANs can access resources on the other LAN as if those resources were local.

Bridges and routers link adjacent LANs. Local bridges and routers were first used to extend the area a network could cover by allowing users to connect two adjacent LANs to maintain performance by reducing the number of users per segment. Both Ethernet and Token Ring specify limits on maximum distances between workstations and hubs, hubs and hubs, and a maximum number of stations that can be connected to a single LAN. To provide network connectivity for more people, or extend it to cover a larger area, it is sometimes necessary to link two different LANs or LAN segments. Bridges and routers can both provide this function.

Today, however, these internetworking devices are also increasingly used to **segment** LANs to maintain performance by reducing the number of users per segment. When users on a single LAN begin to experience slower response times, the culprit is often congestion: too much traffic on the LAN. One method users are employing to deal with this is to break large LANs with many users into smaller LANs, each with fewer users. Adding new network users may require the organization to create new LANs to accommodate them. Implementing new applications on an existing LAN can create so much incremental traffic that the organization may need to break the LAN into smaller LANs segments to maintain acceptable performance levels.

In all of these cases, it is still critical that users on one LAN be able to reach resources on other LANs within the organization. But the LANs must be connected in such a way that packets are **filtered**, so that only those packets that need to pass from one LAN to another are forwarded across the link. This keeps the packets sent between two stations on any one LAN from crossing over onto the other LANs and thereby congesting them. A general rule of thumb suggests that 80 percent of the packets transmitted on a typical workgroup or department LAN are destined for stations on that LAN. Both bridges and routers can be used to segment LANs.

Bridges Bridges are the simpler, and often less expensive, type of device. Bridges filter packets between LANs by making a simple forward/don't forward decision on each packet they receive from any of the networks they are connected to. Filtering is done based on the destination address of the packet. If a packet's destination is a station on the same segment where it originated, it is not forwarded. If it is destined for a station on another LAN, it is connected to a different bridge port and forwarded to that port. Many bridges today filter and forward packets with very little delay, making them good for large traffic volumes.

Routers Routers are more complex internetworking devices and are also typically more expensive than bridges. They use Network Layer Protocol Information within each packet to route it from one LAN to another. This means that a router must be able to recognize all of the different Network Layer Protocols that may be used on the networks it is linking together. This is where the term multiprotocol router comes from — a device that can route using many different protocols. Routers communicate with each other and share information that allows them to determine the best route through a complex internetwork that links many LANs.

Switches

Switches are another type of device used to link several separate LANs and provide packet filtering between them. A LAN switch is a device with multiple ports, each of which can support a single endstation or an entire Ethernet or Token Ring LAN. With a different LAN connected to each of the switch's ports, it can switch packets between LANs as needed. In effect, it acts like a very fast multiport bridge — packets are filtered by the switch based on the destination address.

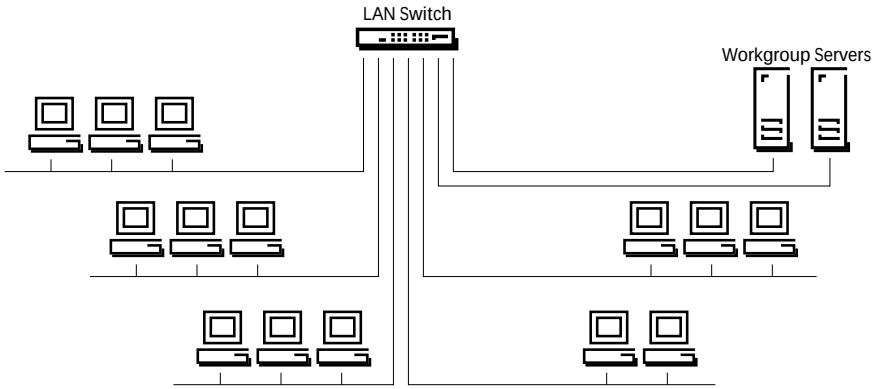

Figure 7 | **Switches**

Using LAN switches allows a network designer to create several small network segments. These smaller segments mean that fewer stations are competing for bandwidth, thereby diminishing network congestion.

Switches are used to increase performance on an organization's network by segmenting large networks into many smaller, less congested LANs, while still providing necessary interconnectivity between them. Switches increase network performance by providing each port with dedicated bandwidth, without requiring users to change any existing equipment, such as NICs, hubs, wiring, or any routers or bridges that are currently in place. Switches can also support numerous transmissions simultaneously.

Deploying technology called **dedicated LANs** is another advantage of using switches. Each port on an Ethernet switch supports a dedicated 10 Mbps Ethernet LAN. Usually, these LANs comprise multiple stations linked to a 10BASE-T hub (see Figure 7), but it is also possible to connect a single high-performance station, such as a server, to a switch port. In this case, that one station has an uncontested 10 Mbps Ethernet LAN all to itself. Packets forwarded over it from other ports on the switch will never produce any collisions because there are no other stations on the LAN at that port.

As was noted earlier, LAN switching is a relatively new technology. Today's switching devices switch relatively large, variable-length LAN packets between different local area networks. ATM is another type of switching technology that switches small, fixed-length cells containing data. ATM networks can be run at much higher data rates than today's LANs. Eventually, they will be used to carry voice, video, and multimedia traffic, as well as computer-generated data over both short and long distances. ATM will be one of the dominant enterprise networking technologies of the future, and many companies are beginning to develop strategies to incorporate ATM in their existing LANs and LAN internetworks.

Networking Today

LAN technology is evolving. In the early 1980s LANs were strictly local area networks, linking small groups of computers in company departments. As workgroup LANs proliferated over the past 10 years, users began connecting them to form internetworks, first with bridges and later with routers. Today's networks typically comprise a combination of workgroup and campus hubs, bridges, and routers. Switches are also beginning to become more prevalent.

The next few years will see networks evolve to include more sophisticated LAN switches and switching hubs. They will be designed using several different types of components, both old and new. Ethernet and Token Ring LANs will be built with stackable workgroup hubs, which, in turn, will be interconnected by larger modular hubs that may incorporate LAN switching functionality. Large networks will include another layer of consolidation with **network center** hubs linking workgroup hubs and switches. Routers will continue to be used as gateways to the wide area network linking other buildings and remote sites.

For networks to deliver the performance today's users require, their many components must work together to deliver seamless connectivity between all of the users and computing systems throughout the enterprise. Flexibility to grow, power to support applications, and seamless connectivity are what users expect in the products they choose to build LANs and enterprise networks.

About Bay Networks

Born from the merger of SynOptics Communications and Wellfleet Communications on October 20, 1994, Bay Networks, Inc., is one of the world's largest networking companies with revenues exceeding $2 billion and earnings of over $250 million.

Headquartered in Santa Clara, California, Bay Networks manufactures and markets a comprehensive line of networking equipment used to build both small and large-scale corporate networks for companies around the world.

Through both direct and indirect channels, the company sells a complete line of intelligent hubs, high-speed switches, multiprotocol routers, and sophisticated network management systems to virtually every Fortune 100 company.

The foundation of Bay Networks networking solutions is its system of intelligent hardware and software products. Designed to meet current and future networking needs, these solutions provide the flexibility to create a network today that can easily grow into a vast, multienterprise network in the future.

Bay Networks product portfolio includes modular, multiprotocol intelligent hubs for both network center and wiring closet applications, highly scalable, high-performance multiprotocol routers for corporate and branch office connectivity, multiservice WAN switches, fixed configuration "stackable" workgroup hubs for Token Ring, FDDI, and Ethernet environments, standalone ATM and 10/100 Mbps Ethernet switches, and a comprehensive network management system that allows for sophisticated control and monitoring of these devices.

Bay Networks markets these products to large and small end-user organizations through a combination of original equipment manufacturers (OEMs), distributors, value-added resellers, and a direct sales force. Typical target users include worldwide retailers, food service companies, financial institutions, technology manufacturers, telecommunications companies, hospitals and universities, and government organizations.

A representative list of Bay Networks customers include: AT&T, Australia Department of Social Services, Bank of International Settlements, Bear Stearns, Boeing Aircraft, British Petroleum, Chase Manhattan Bank, Ford Motor Company, General Motors, McDonald's, MCI, Northwestern Mutual Life, Sprint, 3M, and Wal-Mart.

A major force in the internetworking industry with an installed base of more than 31 million desktop connections, Bay Networks employs over 5,400 people around the world.

The company pioneered the networking industry in the mid-1980s by innovating the ability to run Ethernet networks over common phone wire, as well as being one of the first companies to bring to market high-speed multiprotocol routing.

Additionally, Bay Networks has a number of strategic development and technology partnerships with a variety of industry-leading companies, including IBM, Microsoft, Novell, Intel, Hewlett-Packard, and Sun Microsystems.

Publicly held and traded on the New York Stock Exchange, Bay Networks is led by chairman of the board Paul Severino.

Glossary of Terms

asynchronous transfer mode (ATM) — A type of switching technology in which the switches are small, fixed-length cells containing data.

backbone — A segment of network that links several individual workgroup or department LANs together in a single building. It is also used to link several building LANs together in a campus environment.

bridges — Devices that filter packets between LANs by making a simple forward/don't forward decision on each packet they receive from any of the networks they are connected to.

bus topology — The original coaxial cable-based LAN topology in which the medium forms a single bus to which all stations are attached. The bus topology is rarely used in LAN installations today because it is relatively difficult to add new users or more existing users from one location to another. It is also difficult to troubleshoot a bus-based LAN unless it is very small.

carrier sense multiple access with collision detection (CSMA/CD) — An element defined by the IEEE 802.3 specification. It is an access method that is used by stations connected to an Ethernet LAN. In this method, each station contends for access to the shared medium.

collision — This occurs when two stations try to send packets at the same time. In Ethernet networks, collisions are considered normal events and the CSMA/CD access method is designed to quickly restore the network to normal activity after a collision occurs.

dedicated LAN — Switch configurations in which a port supports a "dedicated" 10 Mbps Ethernet LAN connected to a single high-performance station such as a server, providing an uncontested 10 Mbps Ethernet link all to itself.

EMI — Electromagnetic interference.

Ethernet — The most widely used LAN technology, accounting for approximately 80 percent of all network connections. Standard Ethernet runs at 10 million bits per second (10 Mbps) and balances speed, price, ease of installation, and availability. The rules of Ethernet are defined by the IEEE 802.3 specification. The most popular form of Ethernet is 10BASE-T.

Fast Ethernet — An extension of 10 Mbps Ethernet, Fast Ethernet runs at 100 million bits per second (Mbps). The rules of Fast Ethernet are defined by the IEEE 802.3u specification. Because they use the same protocol, data can be moved between Ethernet and Fast Ethernet without protocol translation.

Fiber Distributed Data Interface (FDDI) — LAN technology that runs at 100 Mbps, a much higher data rate than Ethernet or Token Ring. Originally, FDDI networks required fiber optic cable, but today they can also be run on UTP.

Standards and Protocols

LANs are complex systems that implement many different services in order to provide communication between all of the types of devices that can be connected to them. A communications model called the Open Systems Interconnect (OSI) reference model was developed by the International Standards Organization (ISO) to define all of the services a LAN should provide (see Figure 8). This model defines seven layers, each of which provides a subset of all of the LAN services. This layered approach allows small groups of related services to be implemented in a modular fashion that makes designing network software much more flexible. A network software module that implements services at the Network and Transport Layers of the model can be paired up with different Physical and Data Link Layer modules depending on the requirements of the user's application.

But the OSI model doesn't say how these services should actually be implemented in LAN equipment. The "how to" part has been defined in a number of different **protocols** that have been developed by international standards bodies, individual LAN equipment vendors, and ad hoc groups of interested parties. These protocols typically define how to implement a group of services in one or two layers of the OSI model. For example, Ethernet and Token Ring are both protocols that define different ways to provide the services called for in the Physical and Data Link Layers of the OSI model. They have both been approved by the Institute of Electrical and Electronics Engineers (IEEE), an international communications standards body.

Because they are approved and published by the IEEE, both the Ethernet and Token Ring protocols are said to be industry standards. Any company can acquire the specifications and design Ethernet or Token Ring NICs. Users can purchase an Ethernet NIC, for example, from any vendor and be assured that it will operate in a network with Ethernet NICs from other vendors. This degree of **interoperability** is highly desirable. However, there are many more

filtering — Occurs when a data packet is examined on the network to determine its destination. By looking at a packet's address, network hardware decides whether it should be retained in the local LAN or copied to another LAN. Filtering, which provides some control over internetwork traffic and security, is usually performed by bridges, switches, and routers.

header — A message at the beginning of a data packet that carries special information used by the network to identify the destination station. It is similar to a trailer, which comes at the end of a packet.

IEEE 802.3 — An Ethernet specification commonly defined by the Institute of Electrical and Electronics Engineers (IEEE). The 802.3 specification covers rules for configuring Ethernet LANs, the types of media that can be used, and how the elements of the network should interact.

IEEE 802.5 — A Token Ring specification commonly defined by the Institute of Electrical and Electronics Engineers (IEEE). The 802.5 specification covers rules for configuring Token Ring LANs, the types of media that can be used, and how the elements of the network should interact.

intelligent hubs —Wiring concentrators that can be monitored and managed by network operators.

interoperability — The ability of software and hardware on multiple machines from multiple vendors to communicate.

interrepeater link — One method of linking hubs. This type of connection simply repeats all of the packets from one hub to the other hub it is linked to, so that in effect the two hubs are part of the same LAN.

LAN internetwork — Connecting disparate and geographically dispersed LANs together to form an enterprise system.

local area network (LAN) — A high-speed communications system designed to link computers and other data processing devices together within a small geographic area such as a workgroup, department, or a single floor of a multistory building.

manageable hubs —Another definition for intelligent hubs. Each of the ports on the managed hub can be configured, monitored, and enabled or disabled by a network operator from a hub management console.

modular hubs — A hub that starts with a chassis, or card cage, with multiple card slots, each of which can accept a communications card, or module. Each module acts like a standalone hub; when the communications modules are placed in the card slots in the chassis, they connect to a high-speed communications backplane that links them together so that a station connected to a port on one module can easily communicate with a station on another module.

protocols for providing services at the higher layers of the OSI model and very few of them have been approved by an international standards bodies. In fact, most upper layer LAN protocols are incorporated into proprietary network operating systems, such as Novell's NetWare, IBM's LAN Server, and Microsoft's LAN Manager. A user has to buy only that vendor's products in order to be assured that they will interoperate on a LAN.

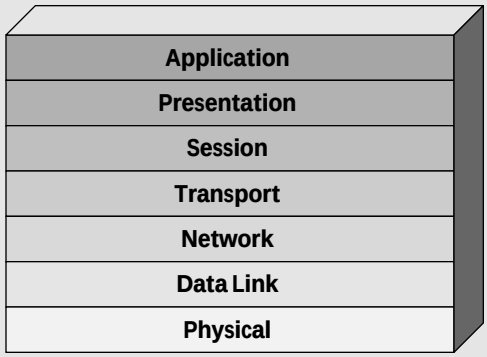

Figure 8 | **ISO Reference Model**

The International Standards Organization (ISO), the primary standard-setting body in the data communications industry, developed a framework for LAN standards called the Open Systems Interconnect reference model. This reference model represents a standard approach to communicate information throughout a network so that a variety of independently developed computer and communications devices can interoperate.

multistation access unit (MAU) — A Token Ring wiring concentrator that connects each station in a Token Ring LAN.

network center — A single, secure, fire-safe location where a company consolidates its network resources.

network interface card (NIC) — The physical connection from the computer to the network is made by putting a NIC inside the computer and connecting it to the shared cable.

Open Systems Interconnect reference model (OSI) — A communications model developed by the International Standards Organization (ISO) to define all of the services a LAN should provide. This model defines seven layers, each of which provides a subset of all of the LAN services. This layered approach allows small groups of related services to be implemented in a modular fashion that makes designing network software much more flexible.

packet — In a shared media network, when one stations wishes to send a message to another station, it uses the network software to put the message in an "envelope." This envelope is called a packet.

protocols — Developed by international standards bodies, individual LAN equipment vendors, and ad hoc groups of interested parties to define how to implement a group of services in one or two layers of the OSI model.

repeaters — Devices that amplify and regenerate signals so they can travel farther on a cable. The term "repeater" is often used to describe hubs.

ring topology — A network whose nodes are connected in a continuous loop.

routers — These are more complex internetworking devices that are also typically more expensive than bridges. They use Network Layer Protocol Information within each packet to route it from one LAN to another.

segmentation — The act of improving network performance by dividing a single large network into multiple smaller, less congested LANs while maintaining connectivity between them. Switches offer an effective segmentation tool by providing each port with dedicated bandwith without requiring users to change any existing equipment such as NICs, hubs, wiring, or any routers or bridges that are currently in place. Switches can also support numerous transmissions simultaneously.

shared access — Shared media technology means that all of the devices attached to the LAN share a single communications medium, usually a coaxial, twisted pair, or fiber optic cable.

shielded twisted pair (STP) — Cable that has shielding around it to provide more protection against electromagnetic interference (EMI).

Network Operating Systems

Ethernet and Token Ring technologies are just one part of a complete LAN. They provide the services specified in the Physical and Data Link Layers of the OSI model, but several other services must be added on top of the connectivity of Ethernet or Token Ring. Network operating systems (NOSs) are most often used to provide the additional communications services.

A NOS defines client and server systems. Clients are individual user workstations attached to the network where application programs are run and data is generated. Servers are shared network resources that provide hard disk space for users to store files, printer services, and a number of other network services. The network operating system provides a set of protocols in software that run on both servers and client systems and allow them to communicate with each other, share files, printers, and other network resources.

stackable hubs — Hubs that look and act like standalone hubs except that several of them can be "stacked" or connected together, usually by short lengths of cable. When they are linked together they can be managed as a single unit.

standalone hubs — Single box-level hubs with a number of ports. Standalone hubs usually include some method of linking them to other standalone hubs — either by connecting them together with a length of 10BASE5 coaxial cable or cascading them using twisted pair between individual ports on each hub.

structured wiring architecture — A wiring architecture that physically star-wires all network stations to intelligent hubs.

switches — A device that links several separate LANs and provides packet filtering between them. A LAN switch is a device with multiple ports, each of which can support an entire Ethernet or Token Ring LAN.

token — a signal used in a Token Ring network that coordinates the transmission of data among the nodes. The token travels around the network, and a node can transmit data only when it has a token.

token passing — The access method used on Token Ring networks.

Token Ring — A major LAN technology in use today. Token Ring rules are defined in the IEEE 802.5 specification. Like Ethernet, the Token Ring protocol provides services at the Physical and Data Link Layers of the OSI model. Token Ring networks can be run at two different data rates, 4 Mbps or 16 Mbps.

trailer — A message a the end of a data packet that carries special information used by the network to identify the destination station. It is similar to a header, which comes at the beginning of a packet.

10BASE-T — The specification for running Ethernet on UTP. This stands for 10 Mbps, baseband signaling (the signaling method used by Ethernet networks), over twisted pair cable.

10BASE5 — An Ethernet specification that uses a thick coaxial cable. 10BASE5 is seldom installed in new Ethernet networks today.

10BASE2 — An Ethernet specification that uses a thin coaxial cable medium. 10BASE2 is only used in very small office networks.

unshielded twisted pair (UTP) — UTP cable is similar to telephone cable but has somewhat more stringent specifications regarding its susceptibility to outside EMI than common telephone wire. UTP is used much more often than STP.

For more sales and product information, please call **1-800-8-BAYNET**.

United States

Bay Networks, Inc.
4401 Great America Parkway
Santa Clara, CA 95054
1-800-8-BAYNET

Bay Networks, Inc.
8 Federal Street
Billerica, MA 01821-5501
1-800-8-BAYNET

Europe, Middle East, and Africa

Bay Networks EMEA, S.A.
Les Cyclades – Immeuble Naxos
25 Allée Pierre Ziller
06560 Valbonne, France
+33-4-92-96-69-96 Fax
+33-4-92-96-69-66 Phone

Pacific Rim, Canada, and Latin America

Australia +61-2-9927-8888
Brazil +55-11-247-1244
Canada 416-733-8348
China +8610-238-5177
Hong Kong +852-2-539-1388
India +91-11-301-0404
Japan +81-3-5402-7001
Mexico +52-5-202-7599
Singapore +65-323-3522

World Wide Web: **http://www.baynetworks.com**

10/96 WP510-2620WC-C

Printed on recycled paper.