Ethernet Network Construct

In recent years, networking computers has taken on greater importance as organizations rely on a network for communication applications like electronic mail and for core business operations functions like database applications. This tutorial helps to explain Ethernet and Fast Ethernet, which are two of the most popular technologies used in networking

LANs

Networks are collections of independent computers that communicate with one another over a shared network medium. Local area networks (LANs) are those networks usually confined to a geographic area, such as a single building or a college campus. LANs, however, are not necessarily simple in design, as they may link many hundreds of computers and be used by many thousands of users. The development of various standards for networking protocols and media has made possible the proliferation of LANs in organizations worldwide for business and educational applications.

WANs

Often a network is located in multiple physical locations. Wide area networking is the connecting of multiple LANs that are geographically separate. This is accomplished by connecting the different LANs using services including dedicated leased phone lines, Dial-up phone lines both synchronous and asynchronous, satellite links, and data packet carrier services. Wide area networking can be as simple as providing modems and a remote access server to allow remote employees to dial in; or it can be as complex as linking hundreds of branch offices across the world using special routing protocols and filters to minimize the expense of sending data sent over vast distances.

Internet

With the meteoric rise of demand for connectivity, the Internet has become the communications highway for millions of users. The Internet was initially restricted to military and academic institutions in its infancy but now it is a full-fledged information channel for any and all forms of information and commerce. Internet web site's now provide personal, educational, political and economic resources to every corner of the planet.

Intranet

With the advancements made in browser-based software for the Internet, there is now a phenomenon called an Intranet which corporate or other private organizations have developed. An Intranet is a private network utilizing Internet-type tools, but available only within that organization. For large organizations, an Intranet provides an easy access mode to corporate information for the employees through the same type of tools used to go outside the company.

Ethernet

Ethernet is the most popular physical layer LAN technology in use today. Other Lan types include Token Ring, Fast Ethernet, Fiber Distributed Data Interface (FDDI), Asynchronous Transfer Mode (ATM) and LocalTalk. Ethernet is popular because it strikes a good balance between speed, cost and ease of installation. These strong points, combined with wide acceptance in the computer marketplace and the ability to support virtually all popular network protocols, make Ethernet an ideal networking technology for most computer users today. The Ethernet standard is defined by the Institute for Electrical and Electronic Engineers (IEEE) as IEEE Standard 802.3. This standard defines rules for configuring an Ethernet as well as specifying how elements in an Ethernet network interact with one another. By adhering to the IEEE standard, network equipment and network protocols will interoperate efficiently.

Fast Ethernet

For Ethernet networks that need higher transmission speeds, the Fast Ethernet standard (IEEE 802.3u) has been established. This standard raises the Ethernet speed limit from 10 Megabits per second (Mbps) to 100 Mbps with only minimal changes to the existing cable structure.

There are three types of Fast Ethernet: 100BASE-TX for use with level 5 UTP cable, 100BASE-FX for use with fiber-optic cable, and 100 BASE-T4 which utilizes and extra two wires for use with level 3 UTP cable. The 100BASE-TX standard has become the most popular due to its close compatibility with the 10BASE-T Ethernet standard. For the network manager, the incorporation of Fast Ethernet into an existing configuration presents a host of decisions. Each site in the network must determine the number of users that really need the higher throughput, decide on which segments of the backbone need to be reconfigured specifically for 100BASE-T and then choose the necessary hardware to connect the 100BASE-T segments with existing 10BASE-T segments.

Gigabit Ethernet is a future technology that promises a migration path beyond Fast Ethernet so that the next generation of networks will support even higher data transfer speeds.

Protocols

Network protocols are standards that allow computers to communicate. A protocol defines how computers should identify one another on a network, the form that the data should take in transit, and how this information should be processed once it reaches its final destination. Protocols also define procedures for handling lost or damaged transmissions or "packets." IPX (for Novell NetWare), TCP/IP (for UNIX, Windows NT, Windows 995 and other platforms), DECnet (for networking Digital Equipment Corp. computers), AppleTalk (for Macintosh computers),

and NetBIOS/NetBEUI (for LAN Manager and Windows NT network protocols in use today.

Although each network protocol is different, they all are able to share the same physical cabling. This common method of accessing the physical network allows multiple protocols to peacefully coexist over the network media, and allows the builder of a network to use common hardware for a variety of protocols. This concept is known as "protocol independence," which means that devices that are compatible at the physical and data link layers allow the user to run many different protocols over the same medium.

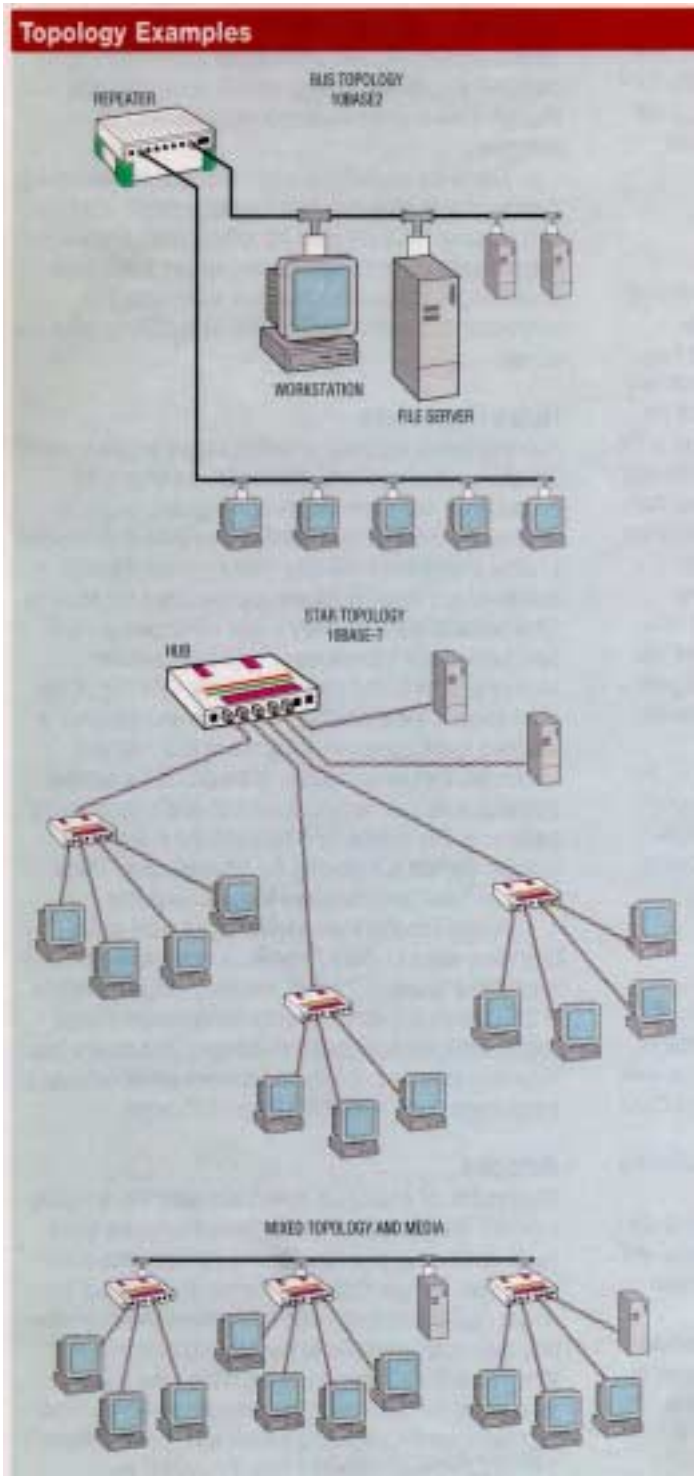

Media

An important part of designing and installing an Ethernet is selecting the appropriate Ethernet medium for the environment at hand. There are four major types of media in use today. Thickwire for 10BASE5 networks, thin coax for 10BASE networks, unshielded twisted pair (UTP) for 10BASE-T networks and fiber optic for 10BASE-FL or Fiber-Optic Inter-repeater Link (FOIRL) networks. This wide variety of media reflects the evolution of Ethernet and also points to the technology's flexibility. Thickwire was one of the first cabling systems used in Ethernet but was difficult to work with and expensive. This evolved to thin coax, which is easier to work with and less expensive.

Today the most popular wiring schemes are 10BASE-T and 100BASE-TX which both use unshielded twisted pair (UTP) cable. This is similar to telephone cable and comes in a variety of grades, with each higher grade offering better performance. Level 5 cable is the highest, most expensive grade, offering support for transmission rates of up to 100 Mbps. Level 4 and level 3 cable are less expensive, but cannot support the same data through put speeds; level 4 cable are less expensive, but cannot support the same data through put speeds; level 4 cable can support speeds of up to 20 Mbps, level 3 up to 16 Mbps. The 100BASE-T4 standard allows for support of 100 Mbps Ethernet over level 3 cable, but at the expense of adding another pair of wires (4 pair instead of the 2 pair used for 10BASE-T); for most users, this is an awkward scheme and therefore 100 BASE-T4 has seen little popularity. Level 2 and level 1 cables are not used in the design of 10BASE-T networks.

For specialized applications, fiber-optic, of 10Base-FL, Ethernet segments are popular. Fiber-optic cable is more expensive, but it is invaluable or situations where electronic emissions and environmental hazards are a concern. Fiber0optic cable is often used in interbuilding applications to insulate networking equipment from electrical damage caused by lightning because it does not conduct electricity. Fiber-optic cable can also be useful in areas where a large amount of electromagnetic interference is present, such as on a factory floor. The Ethernet standard allows for fiber-optic cable segments up to 2 kilometers long, making fiber optic Ethernet perfect for connecting nodes and buildings that are otherwise not reachable with copper media.

Topologies

Ethernet media are used in two general configurations or topologies; "bus" and "star." These two topologies define how "nodes" are connected to one another. A node is an active device connected to the network, such as a computer or a printer. A node can also be a piece of networking equipment such as a hub, switch or a router. A bus topology consists of nodes linked together in series with each node connected to a

long cable or bus. Many nodes can tap into the bus and begin communication with all other nodes on that cable segment. A break anywhere in the cable will usually cause the entire segment to be inoperable until the break is repaired. Examples of bus topology include 10BASE2 and 10BASE5. 10BASE-T Ethernet and Fast Ethernet use a star topology. Generally a computer is located at one end of the segment, and the other end is terminated in a central location with a hub. Because UTP is often run in conjunction with telephone cabling, this central location can be a telephone closet or other area where it is convenient to connect the UTP segment to a backbone. The primary advantage of this type of network is reliability, for if one of these "point-to-point" segments has a break, it will only affect the two nodes on that link. Other computer users on the network continue to operate as if that segment were nonexistent.

Collisions

Ethernet is a shared media, so there are rules for sending packets to avoid conflicts and protect data integrity. Nodes on an Ethernet network send packets when they determine the network is not in use. It is possible that two nodes at different locations could try to send data at the same time. When both PCs are transferring a packet to the network at the same time, a collision will result. Minimizing collisions is a crucial element in the design and operation of networks. Increased collisions are often the result of too many users on the network, which results in a lot of contention for network bandwidth. This can slow the performance of the network from the users point of view. Segmenting the network, where a network is divided into different pieces joined together logically with a bridge or switch, is one way of reducing an overcrowded network.

Ethernet Products

The standards and technology that have just been covered are translated into specific products that network managers use to build Ethernet networks. The following text discusses the key products needed to build a Ethernet Network.

Transceivers

Transceivers are used to connect nodes to the various Ethernet media. Most computers and network interface cards contain a built-in 10BASE-T or 10BASE2 transceiver, allowing them to be connected directly to Ethernet without requiring an external transceiver. Many Ethernet compatible devices provide an AUI connector to allow the user to connect to any media type via an external transceiver. The AUI connector consists of a 15-pin D-shell type connector, female on the computer side, male on the transceiver side. Thickwire (10BASE5) cables also use transceivers to allow connections.

For Fast Ethernet networks, a new interface called the MII (Media Independent Interface) was developed to offer a flexible way to support 1000 Mbps connections. The MII is a popular way to connect 100BASE-FX links to copper-based Fast Ethernet devices.

Network Interface Cards

Network interface cards, commonly referred to as NICs. are used to connect a PC to a network. The NIC provides a physical connection between the networking cable and the computer's internal bus. Different computers have different bus architectures, PCI bus master slots are most commonly found on 486/Pentium PCs and ISA expansion slots are commonly found on 386 and older personal computers. Network interface cards come in three basic varieties; 8 bit, and 32 bit. The larger the number of bits that can be transferred to the NIC, the faster the NIC can transfer data to the network cable.

Many NIC adapters comply with Plug-n-Play (PnP) specifications. On PnP systems, the NICs are automatically configured without user intervention while on non-PnP systems, configuration is done manually through a setup program and/or manually set DIP switches.

Cards are available to support almost all networking standards, including the latest Fast Ethernet environment. Fast Ethernet NICs are often 10/100 capable, and will automatically set to the appropriate speed. Full duplex networking is another option, where a dedicated connection to a switch allows a NIC to operate at twice the speed.

Hubs/Repeaters

Hubs/repeaters are used to connect together two or more Ethernet segments of any media type. As segments exceed their maximum length, signal quality begins to deteriorate. Hubs provide the signal amplification required to allow a segment to be extended a greater distance. A hub takes any incoming signal and repeats it out all ports. Ethernet hubs are necessary in star topologies. A multi-port, twisted pair hub allows several point-to-point segments to be joined into one network. One end of the point-to-point link is attached to the hub and the other is attached to the computer. If the hub is attached to a backbone, then all computers at the end of the twisted pair segments can communicate with all the hosts on the backbone. The number and type of hubs in any one collision domain is limited by the Ethernet rules. These 'repeater rules" are discussed in more detail later.

A very important fact to note about hubs is that they only allow users to share Ethernet. A network of repeaters is termed a "shared Ethernet". Meaning that all members of the network are contending for transmission of data onto a single network (collision domain). This means that individual members of a shared network will all only get a percentage of the available network bandwidth.

Bridges

The function of a bridge is to connect separate networks together. Bridges can connect different networks types (such as Ethernet and Fast Ethernet) or networks of the same type. Bridges map the Ethernet addresses of the nodes residing on each network segment and then allow only the necessary traffic to pass through the bridge. When a packet is received by the bridge, the bridge determines the destination and source segments. If the segments are the same, the packet is dropped ("filtered"); if the segments are different, then the packet is "forwarded" to the right segment. Additionally, bridges prevent all bad or

misaligned packets from spreading by not forwarding them. Bridges are called "store-and-forward" devices because they look at the whole Ethernet packet before making their filtering or forwarding decisions. Filtering of packets and the regeneration of forwarded packets enables bridging technology to split a network into separate collision domains. This allows for greater distances and more repeaters to be used in the total network design.

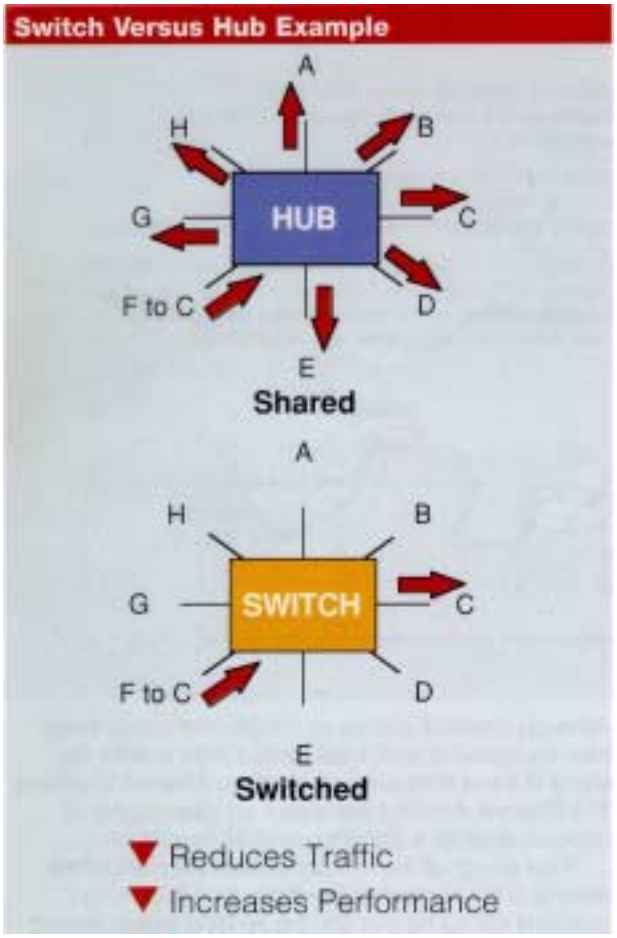

Most bridges are self-learning task bridges, meaning they determine the user Ethernet addresses on the segment by building a table as packets are passed through the network. This address self-learning capability dramatically raises the possibility of creating network loops in networks that have many bridges. As each device learns the network configuration, a loop resents conflicting information on which segment a specific address is located and forces the device to forward all traffic. The Spanning Tree Algorithm is a software standard (found in the IDDD 802.1d specification) for describing how switches and bridges can communicate to avoid network loops.

RMON

(Remote Monitoring MIB) provides a higher level of information than SNMP by itself. When supported on a device, RMON runs continuously and allows the network manager to view statistics, set alarm conditions which can issue "traps" or can be logged in a table and to continuously flag certain events as they occur. RMON will become more commonplace in switches as new chip designs contain the support for ROM directly in their silicon.

Ethernet Switches

Ethernet switches are an expansion of the concepts in Ethernet bridging. If it makes sense to link two networks through a bridge, why not develop a device that can link four, six, 10 or more networks together? That's exactly what a LAN switch does. LAN switches come in two basic architectures, cut-through and store-and-forward. Cut-through switches have, in the past, held a speed advantage because when a packet comes into the switch, it only examines the destination address before forwarding it on to its destination segment. A store-and-forward switch, on the other hand, accepts and analyzes the entire packet before forwarding it to its destination. It takes more time to examine the entire packet, but it allows the switch to catch certain packet errors and keep them from propagating through the network. Today, the speed of store-and-forward switches has caught up with cut-through switches to the point where the difference between the two is minimal. Also, there are a large number of hybrid switches available that mix both cut-through and store-and-forward architectures.

Both cut-through and store-and-forward switches separate a network into collision domains, allowing network design rules to be extended. Each of the segments attached to an Ethernet switch has a full 10 Mbps of bandwidth shared by fewer users which results in better performance (as opposed to hubs that only allow sharing of bandwidth from a single Ethernet).

Newer switches today offer high-speed links, either FDDI, Fast Ethernet or ATM, that can be used to link the switches together or to give added bandwidth to particularly important servers that get a lot of traffic. A network composed of a number of switches linked together via uplinks is termed a "collapsed backbone" network.

Routers

Routers work in a manner similar to switches and bridges in that they filter out network traffic. Rather than doing so by packet addresses they filter by specific protocol. Routers were born out of the necessity for dividing networks logically instead of physically. An IP router can divide a network into various subnets so that only traffic destined for particular IP addresses can pass between segments. The price paid for this type of intelligent forwarding and filtering is usually calculated in term of speed of the network. Such filtering takes more time than that exercised in a switch or bridge which only looks at the Ethernet address but in more complex networks efficiency is improved.

Servers

When there is a demand for particular files or device access among network users, a means must be found to allow such resources to be shared. Servers are networked devices that allow their files, devices or other resources to be shared by network users. File servers are computers designed to give users access to files stored on their hard drives. Print servers

are devices that attach a printer to the network and allow all network users access to the printer. Terminal servers allow terminals to attach directly to a network and access any host available.

Network Design Criteria

Ethernets and Fast Ethernets have design rules that must be followed in order to function correctly. The maximum number of nodes, the number of repeaters and maximum segment distances are defined by the electrical and mechanical design properties of each type of Ethernet and Fast Ethernet media

Network Type	Max. Nodes Per Segment	Max. Distance Per Segment
10BASE5	100	500m
10BASE2	30	185m
10BASE-T	2	100m
10BASE-FL	2	2000m

A network using repeaters, for instance, has restrictions having to do with the timing constraints of Ethernet. Although electrical signals on the Ethernet media travel near the speed of light, it still takes a finite time for the signal to travel from one end of a large Ethernet to another. The Ethernet standard assumes it will take roughly 50 microseconds for a signal to reach its destination.

If the design of the network violates the rules for the placing of the number of repeaters, then this timing guideline will not be met and the sending station, having not received an acknowledgment of its sent packet, will resend that packet. This can lead to lost packets, excessive resent packets which can slow network performance and create trouble for applications.

Ethernet is subject to the "5-4-3" rule of repeater placement: the network can only have five segments connected; it can only use four repeaters; and of the five segments, only three can have users attached to them; the other two must be inter-repeater links. Fast Ethernet has modified repeater rules, since the minimum packet size takes less time to transmit than regular Ethernet. The length of the network links and the standard allows a fewer number of repeaters. In Fast Ethernet networks, there are two classes of repeaters. Class I repeaters have a latency of 0.7 microseconds or less and are limited to one repeater per network. Class II repeaters have a latency of 0.46 microseconds or less and are limited to two repeaters per network. The following are the distance (diameter)

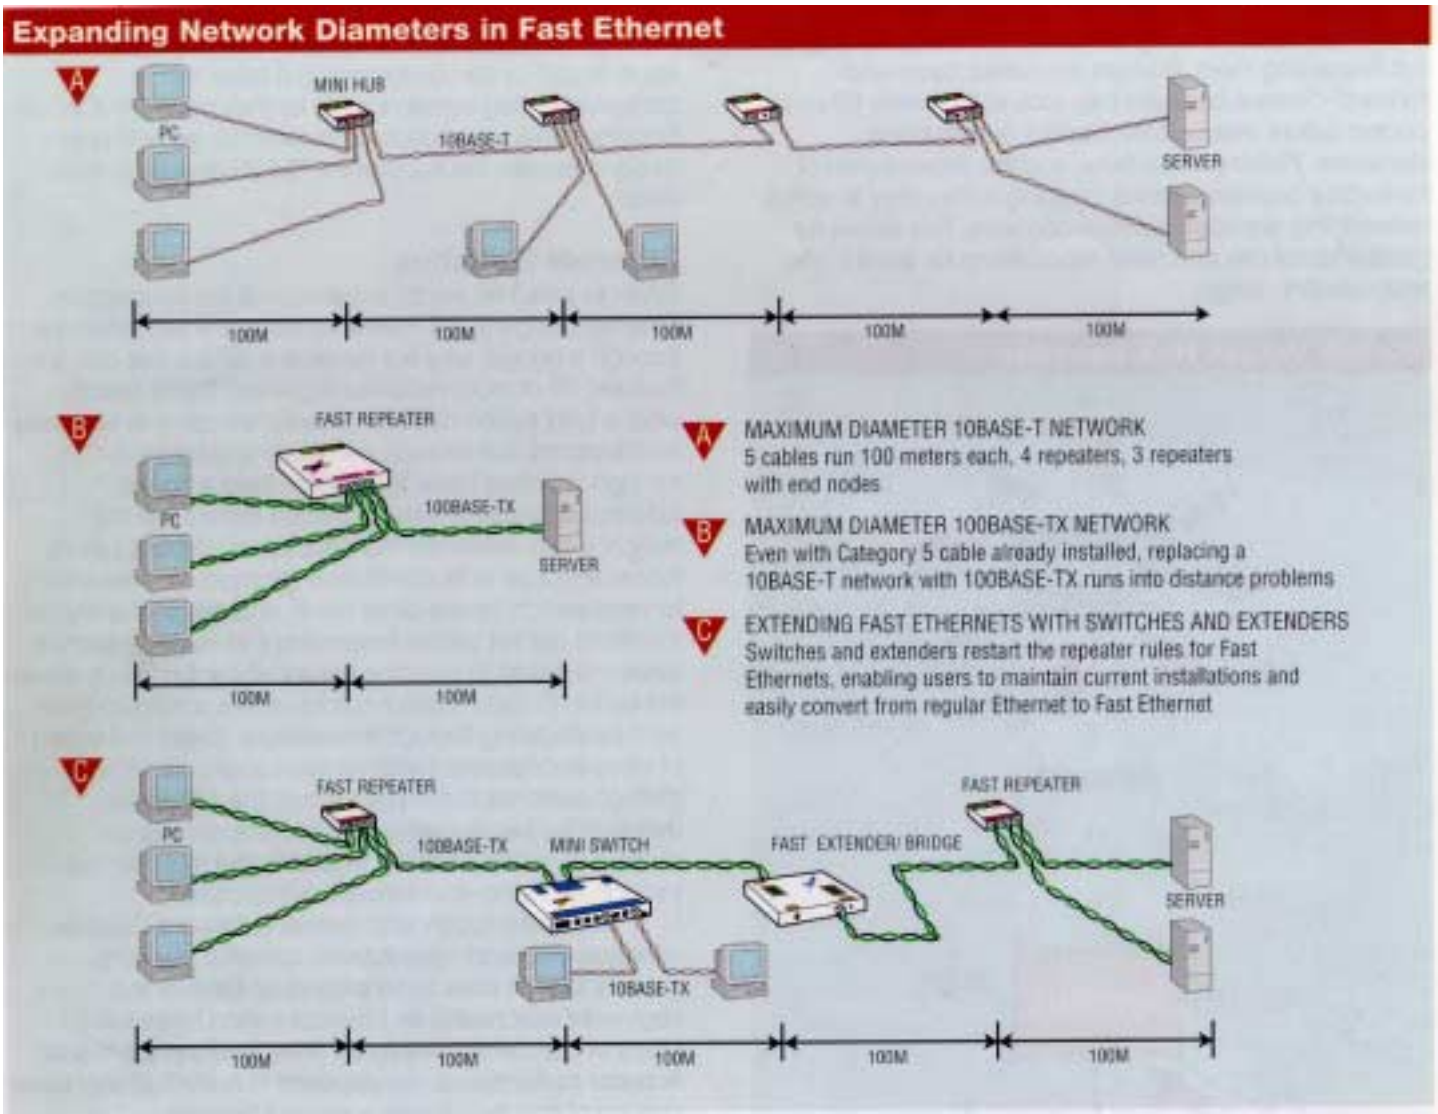

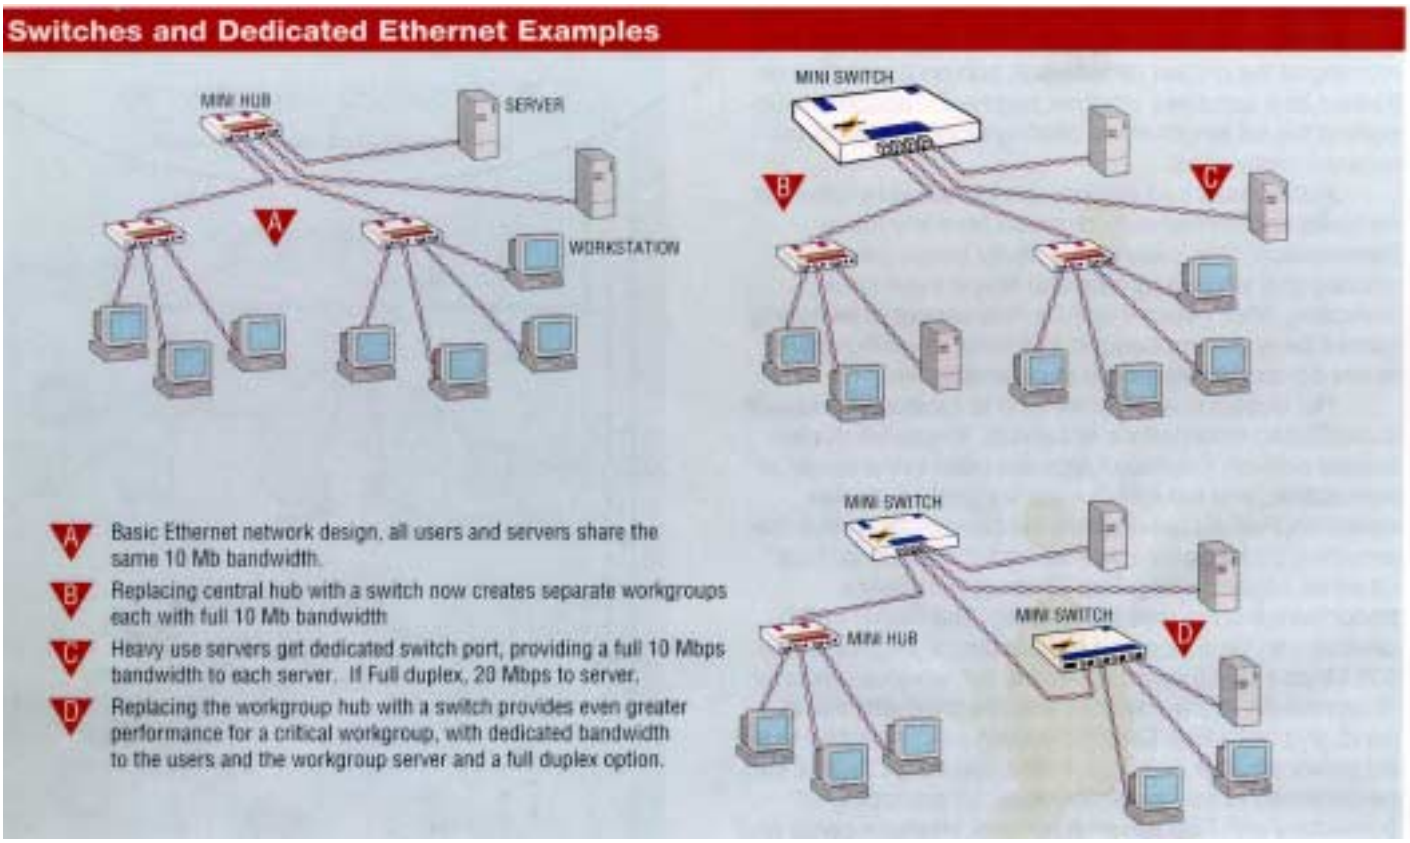

Characteristics for these types of Fast Ethernet Repeater combinations:

Fast Ethernet	Copper	Fiber
No repeaters	100m	412m*
One Class I repeater	200m	272m
One Class II repeater	200m	272m
Two Class II repeaters	205m	228m

*Full Duplex Mode 2 km

When conditions require more distance or an increase in the number of nodes/repeaters, a bridge, router, or switch can be used to connect multiple networks together. These devices essentially "join" two separate networks, allowing the network design criteria to be restarted. With switches, network designers can build large networks that function well. Each network connected via one of these devices is referred to as a separated collision domain in the overall network. The reduction in costs of bridges and switches has reduced the impact of repeater rules on network design.

When Ethernets Become Too Slow

As more users are added to a shared network or as applications requiring more data are added, performance deteriorates. This is so because all users on a shared network are competitors for the Ethernet bus. On a moderately loaded 10 Mbps Ethernet network being shared by 30050 users, that network will usually only be able to sustain throughput in the neighborhood of 2.5 Mbps after accounting for packet overhead, interpacket gaps and collisions. Increasing amounts of users (and therefore packet transmissions) create increasing potential for collisions. Collisions occur when two or more nodes attempt to send information at the same time – when they realize that a collision has occurred, each node backs off for a random time before attempting another transmission. With shared Ethernet, the likelihood of collision increases as more nodes are added to the shared collision domain of the shared Ethernet.

One of the steps to alleviating problems is to segment the traffic with a bridge or switch. A switch can replace a hub and improve network performance. For example, an eight-port switch can support eight Ethernets, each running at a full 10 Mbps. Another option is to dedicate one or more of these switched ports to a high traffic device such as a file server

Multimedia and video applications demand as much as 1.5 Mbps of continuous bandwidth – as we have seen above, a single such user would be hardpressed to get this amount of bandwidth alone as their share of an average 10 Mbps network. If you add in the fact that video will look disjointed or "clunky" if the data rate is not sustained, then the pressure will be on the network manager to provide greater throughput to support this application.

When added to the network, Ethernet switches provide a number of enhancements over shared networks. The foremost enhancement is the ability to divide networks into smaller and faster segments. Ethernet switches examine each packet, determine where that packet is destined and then forward that packet to only those ports to which the packet needs to go. Modern switches are able to do all these tasks at "wirespeed", that is without adding delay.

Aside from deciding when to forward the packet or

when to filter the packet, Ethernet switches also completely regenerate the Ethernet packet. This regeneration and retiming of the packet allows each port on a switch to be treated as a complete Ethernet segment, capable of supporting the full length of the cabling along with all of the repeater restrictions.

Additionally, bad packets are identified by Ethernet switches and immediately dropped from any future transmission. This "cleansing" activity keeps problems isolated to a single segment and keeps them from disrupting other network activity. This aspect of switching cannot be underemphasized in a network environment where hardware failures are to be anticipated.

Full duplex is another method to increase bandwidth to dedicated workstations or servers. To use full duplex, special network interface cards are used in the server or workstation, and the switch must support full duplex operation. Full duplex doubles the bandwidth on that link, providing 20 Mbps for Ethernet and 200 Mbps for Fast Ethernet. Implementing Fast Ethernet to increase performance is the next logical step. The higher traffic devices can be connected to switches or each other via 100 Mbps Fast Ethernet, providing tremendous amounts of bandwidth. Many switches are designed with this in mind, and have Fast Ethernet uplinks for connection to a file server or other switches. Eventually, Fast Ethernet can be deployed to the user' desktops, by equipping all computers with Fast Ethernet network interface cards and using Fast Ethernet switches and repeaters. With an understanding of the underlying technologies and products in Ethernet networks, we can now progress to a discussion of some of the most popular real world applications.

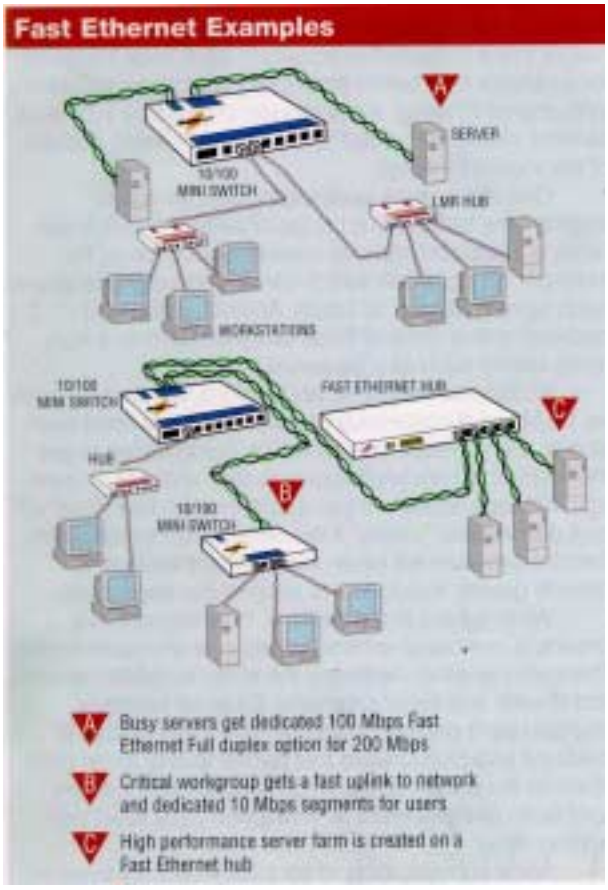

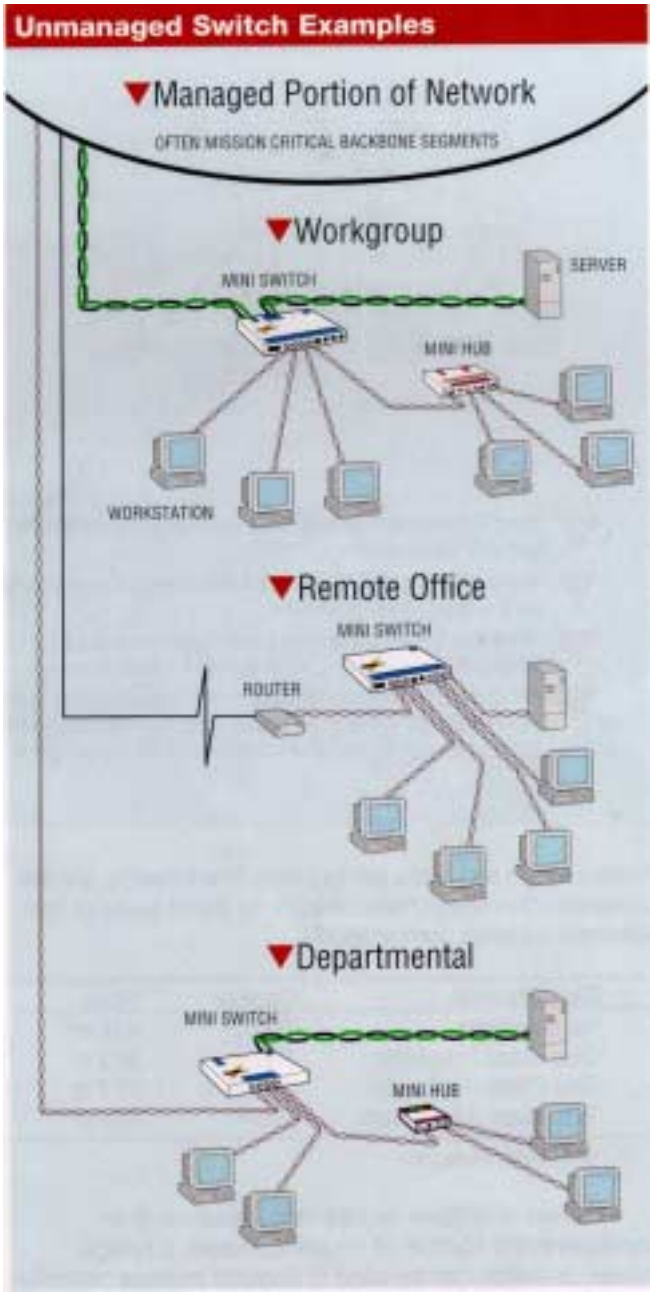

the remote office and the remote user the economy and flexibility of "pay as you go" telephone services. ISDN is a special telephone service that offers three channels, two 64 Kbps "B" channels for user data and a "D" channel for setting up the connection. With ISDN, the B channels can be combined for double the bandwidth or separated and used for different applications or users.

With asynchronous remote access, regular telephone lines are combined with modems and remote access servers to allow users and networks to dial anywhere in the world and have data access. Remote access servers provide connection points for both dial-in and dial-out applications on the network to which they are attached these hybrid devices are capable of routing and filtering protocols and offer other services such as modem pooling and terminal/printer services. For the remote PC user, there is the flexibility of connecting from any available telephone jack, including those in a hotel or on an aircraft.

Remote Access Servers

While Ethernet is local to a geographic area, like a building , remote users, such as traveling sales people, are requesting access to network-based resources. Remote LAN accesses or remote access is quickly becoming a popular way to provide this connectivity. Remote access solutions use telephone services to link a remote user or office with an office network. For demanding applications, where speed and full-time access is crucial, a leased-line solution should be considered. This involves purchasing a router and a special leased line service, which essentially sets up a dedicated telephone line with a set amount of bandwidth – ranging from 56 Kbps to many megabits per second. This solution is limited to the two connected offices and can be very expensive.

Dial-up remote access solutions such as ISDN or asynchronous dial up introduce more flexibility into a remote access solution. Dial-up remote access offers both

Remote Access Applications

Remote access technology is optimized for a number of remote applications. Remote node and remote control applications are when a remote user on a PC or workstation dials into a network and is able to function as if he or she were directly attached to the network. A remote access server provides dial-in services and support for PPP to allow full functionality of the remote user as a network peer (remote node) or to allow the remote user to take over a local node (remote control).

LAN-to-LAN is when an entire remote network is supported over a dial-up connection. Remote access servers on each end act as routers to automatically generate a connection when remote resources are requested. The dial-up connection is maintained according to parameters established by the network manager for timeouts, allowed protocols and for connection duration. Internet access applications involve the use of a remote access server as a router to "firewall" the local network from security problems

present on the Internet. Filters are configured by the network manager to ensure that only authorized traffic is allowed to pass between the local network and the Internet. These applications are actually a hybrid form of LAN-to-LAN connections.

Modem sharing is the ability of the remote access server to provide access for network users to a bank of modems for both dial-in and dial-out applications. Software running on networked hosts allows them to connect to modems attached to a remote access server, providing cost-effective communications from the central site and preserving the investment in modem and communications hardware.

The key to controlling costs is the ability of the remote access server to route the desired protocols and to implement policy-based decisions on how the dial-up connections between sites are managed. In a LAN-to-LAN application, IP and IPX protocol traffic on the network is monitored by a server and when a connection to resources on a remote network is required, the server automatically dials up and connects to that network. Once the network connection is established, the server will monitor the link according to criteria defined by the network manager and manage the link to those specifications. These parameters include: the

amount of time the link is to remain connected if no data is being passed; whether the link is to remain connected if only certain types of traffic are present (i.e. disconnect if only the keep alive or broadcast messages are being transmitted); whether or not to allow a particular protocol or packet type to travel the link between the two networks. Additional convenience features are automatic redialing in case of a busy answering modem or an unplanned disconnect, and time-of-day limits for dial-in/dial-out operations.

Printer Servers

Printer servers allow printers to be shared by other nodes on the network. Supporting with parallel or serial interfaces (sometimes both), a printer server accepts print jobs from any node on the network using the supported protocols and manages the printing of those jobs on the appropriate printer.

The earliest printer servers were external devices, which supported printing via parallel or serial ports on the device. Typically, only one or sometimes two protocols were supported. The latest generation of printer servers supports to transmit data to and receive data from host computers across local area networks, without requiring each terminal to have its own direct connection. And while the terminal server'' existence is still usually justifies by convenience and cost considerations, its inherent intelligence provides many more advantages. Among these is enhanced remote monitoring and control. Terminal servers that support protocols like SNMP make networks easier to manage.

Devices that are attached to a network through a terminal server can be shared between terminals and hosts at both the local site and throughout the network. A single terminal may be connected to several hosts at the same time (in multiple concurrent sessions), and can switch between them. Terminal servers can also be used to link devices that have only serial outputs over a network. A network connection between serial ports on different servers is opened, allowing data to move between to two devices.

multiple protocols, has multiple parallel and serial connection options and, in some cases, are small enough to fit directly on the parallel port of the printer itself. Some printers have printer servers that are internal to the printers themselves, this type of design has an integral communication benefit between the printer and the printer server, but lacks flexibility if a printer has physical problems.

Printer servers as a rule do not contain a large amount of memory. Rather than store each print job in memory, they simply store the information about the host and the protocol involved in a queue. When the desired printer becomes available, then they allow the host to transmit the data to the appropriate printer port on the server. The printer server can then simply queue and print each job in the order in which print requests are received, regardless of protocol used or the size of the job.

Terminal Servers

The original role of terminal servers was to enable terminals

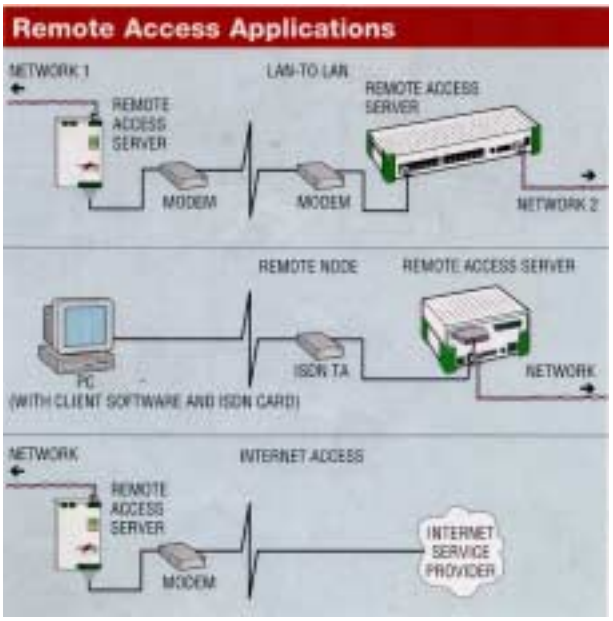

With the advent of multiprotocol terminal servers, the problem of a user needing two terminals to reach hosts that used different communications protocols was alleviated. As long as the terminal server supports the protocol used by the host, the terminal attached to that server can access that host as if it were using the terminals own native protocol. Economically, it also makes sense to have a single connection to the network instead of individual interface cards and transceivers for each terminal.

Digital systems using the LAT protocol and Unix systems using TCP/IP have no natural means to communicate with each other, in spite of how common it is to have VAX and Sun workstations in the same facility. Given its natural translation ability, a multi-protocol terminal server can

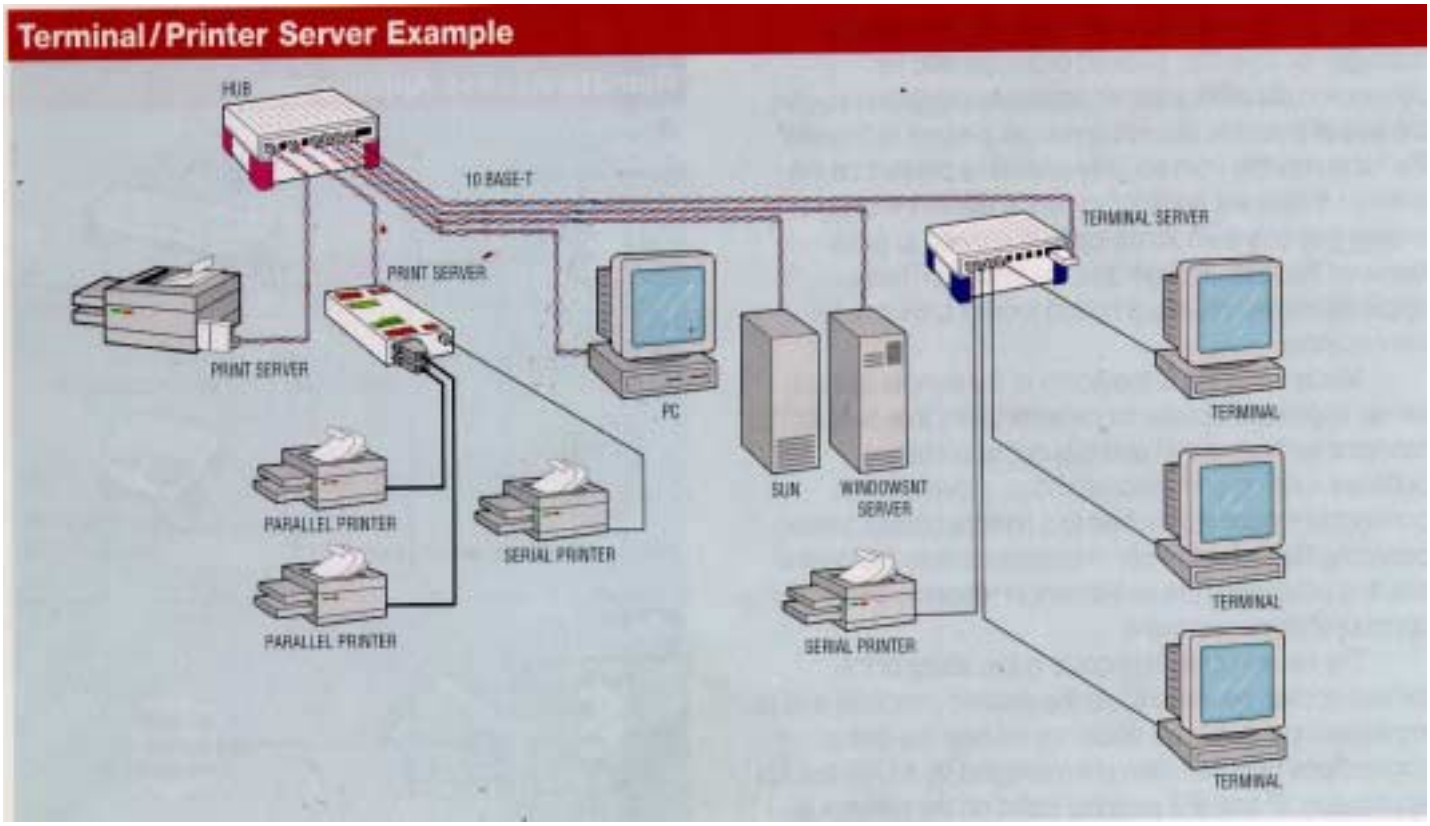

perform conversions between the protocols it knows, like LAT and TCP/IP, at least for those that are set up to work with terminals. While terminal server bandwidth isn't adequate for large file transfers, it can easily handle host-to-host inquiry/response applications, electronic mailbox checking, etc. And it is far more economical than the alternatives of acquiring expensive host software and special-purpose converters. Terminal and printer servers give their users great flexibility in configuring and managing their networks.

Whether it is moving printers and other peripherals from one network to another, expanding the dimensions of interoperability, or preparing for growth, terminal servers can fulfill your needs. You can do it all without major rewiring. The demand for dial up remote access applications is causing terminal and server functionality to evolve. The requirement for support of PPP and SLIP connections has created the need for a "communication" server which does not offer the routing capabilities of a true remote access server, but still offers sophisticated dial up modem support.

Now What?

We hope this introduction to local area networks has been helpful and most informative. Unfortunately we cannot explain everything there is to know about planning, installing, administering and troubleshooting A LAN in a few pages, or even a hundred, pages. The Internet, many books and magazines exist that explain all aspects of computer networks, From LANs to WANs, from network applications, to running cable. Check your local bookstore, software retailer or newsstand.