NETWORK CHECKLISTS

Introduction

Libraries use more than one type of network. They can have local area networks (LANs), networks connecting computers in a small area. One or multiple LANs are possible. Where libraries have multiple LANs, or where libraries are in one building connected to other buildings of the library's parent organization, libraries may also connect to campus or backbone networks. The LAN or the backbone network is connected to the Internet.

Of the two checklists below, the first specifies what is needed for library LANs and backbone networks, and the second indicates what is needed to connect to the Internet.

Checklist Ingredients

Each checklist identifies key network or service components. These items may be standards to be adopted, required physical hardware, or services that need to be provided by an external organization. For each network component, its nature and role in the network is described. There may be multiple ways to provide the needed component. Where multiple ways exist, the primary current alternatives are described and the best is identified. The measure of 'best' being the ability to meet the requirements of audio and video use, as well as limited cost and complexity. The impacts of these best alternatives on library services and their associated costs are described.

Where such cost information is given, remember that the cost information provided is not specific. Actual costs vary dramatically with the scale of network. Networks with larger numbers of users and higher amounts of bandwidth are likely to more expensive. Costs can also change dramatically as technologies become obsolete. What was once an expensive innovative networking product can become low cost as manufacturing costs go down and widespread use occurs. Also remember that the costs do not represent the total cost of ownership. Labor costs associated with building and maintaining a network can be substantial. Even if a library doesn't maintain its own network, it will need to have access to that expertise from a firm and to pay the resulting costs. Maintenance costs also add to the costs identified. The costs identified give a range of costs. More specific information emerges as library plans are developed.

Each checklist varies slightly. Libraries have more control over their LANs than over Internet services. The LAN is made up of components a library can usually choose. All the cables, equipment, and other parts are selected to provide a certain level of performance. The checklist needs to give a detailed list of these components. On the other hand, the Internet is someone else's network, so the need is reduced to provide an itemized list of hardware. The Internet services checklist focuses on key services needed from someone else's network and the hardware needed to connect the LAN to the Internet.

Introduction to Networking Protocols

Protocol Suite (or stack) is the layered tasks that manage commuication between two computers.

Before proceeding to the checklists, you must understand what a protocol stack or suite is. This understanding helps simplify an otherwise complex discussion of networking and creates a model for how different technologies combine to do the several tasks that must happen before network communication can occur. By understanding how a protocol suite works, you can understand the roles each technology plays in the network. Several of the key items in the checklists are explained in terms of their place within a protocol suite.

A protocol is an agreed way to respond to a situation or to communicate with particular kinds of people, for example how to greet the Queen of England or write a letter to a Congressman. Likewise with networking, a protocol is an agreed way to manage interaction between computers. But computer networking is not so simple that a single protocol can be used. Multiple tasks must be accomplished for one computer to communicate with another. These tasks are not independent, each happening in parallel. Rather, one task depends on another. These different communication tasks appear layered. This set of layered tasks that must occur for two devices to communicate is a protocol stack.

What specifically are these tasks that need to be handled for network communication to occur? Think about three computers connected by cables. For communication to occur, there at least needs to be agreement on how a stream of bits is sent over the physical cable. This agreement includes how the physical cable is designed to connect to the computer, and how an electrical signal is sent over that cable. There also must be agreement on how to interpret this stream of bits, sent as an electrical signal. Further, for more than two computers, an agreed way is needed to identify the computer to which a message is being sent. These computers must agree how to handle many more tasks, such as handling errors to knowing when a communication is completed.

OSI Reference Model is the ISO standard model for protocol stacks within networks.

An international standard model exists for the protocol layers involved in computer networks. This model, an ISO (International Standards Organization) standard, is called the OSI (Open Systems Interconnect) Reference Model. This model specifies seven protocol layers. It is a theoretical model of how different networking tasks are handled that is frequently used in explanations of networking.

To understand the networking technologies included in this checklist, only the bottom three layers are important. The bottom layer is the physical layer, Layer 1. The protocol at this layer has to set a standard for how to use the physical media to send electrical signals between computers, delivering a stream of bits to another device.

The second layer is the data link layer, Layer 2. This layer must be able to receive the stream of bits from Layer 1 and interpret it as meaningful chunks of data. Layer 2 protocols organize data into frames. They must also support an addressing system, a way to specify the computer to which frames are being sent. With these two protocol layers specified, there is agreement on how to use cables to communicate and how to package data and send it to specific computers over a network.

Data link layer addresses offer a way to address communications to other computers on a network, but their addressing systems are flat. There is no way to create a route for data traffic among networks based on these addresses. This requires an addressing scheme with addresses that are hierarchical, providing assurance that a group of addresses are used by one network. The third layer provides a way to cope with more complicated networks. The network layer, Layer 3, specifies how to address, format, and route communications that must go among multiple networks.

Network Layer—Layer 3

Data Link Layer—Layer 2

Physical Layer—Layer 1

Protocol Stack: the first three layers

The OSI protocol model is a theoretical model. It doesn't specify how these different layers are implemented. Many alternative protocols actually implement these protocol layers. The checklists identifies the protocols used in LANs and on the Internet to perform these layers of tasks. It describes how devices on the network interact in terms of the functional requirements these different layers represent.

Don't overestimate the importance of understanding the different layers of the network and how they function together. Much of the confusion generated by networking terminology results from different technologies being used at the same time to perform different layers of the protocol stack. The use of terms such "Ethernet," and "IP" in networking discussions becomes easier to understand when you understand the OSI network model and the concepts represented by these terms in terms of their relationship to that model.

LAN Checklist

The primary definition of a LAN is a network connecting a group of computers in a small area. LANs are interconnected using backbone networks. These may connect multiple LANs within one building or among several buildings occupied by one organization. These several buildings are called a campus and their interconnecting network called a campus network. This is not a reference to a university campus. Any group of buildings occupied by the same organization can be called a campus.

This checklist covers list what is needed both for the library LAN and for campus networks. It includes the following items.

Local Area Network Protocol–The networking standard that is used on the LAN to govern communications.

Cabling—The kind of wiring to be used through the LAN.

Network Interface Cards—The hardware installed into computers to allow them to communicate over the LAN.

Structured Cabling System—The way in which the network cabling is installed in a building.

Network Topology and Connecting Devices—The logical structure of the network and the hardware to which the computers are connected by the cabling and through which communications flows.

Network Operating System and Server—The software and hardware that provides computers with access to shared services like printers and file servers.

Wireless LAN—The means for staff and patrons to work without being tied down by network cables.

Campus or Backbone Networks—Networks to connect library LANs or to connect to campus networks.

The LAN does not need all these items. A wireless LAN is optional. Some libraries may not have a large enough infrastructure to require a backbone network.

The checklist items are interrelated. The choices for one component in this checklist constrains later components. Once a recommendation is made, the alternatives described later are restricted to those compatible with previously recommendations.

#1. Local Area Network Protocol

The first choice to make when building a local area network (LAN) is the basic technology to be used. This is the choice of the technologies or protocols to be used for Layers 1 and 2 of the OSI model.

Alternative Solutions

For more information on Ethernet standard: www.ieee.org

Ethernet. Some say the only choice for these protocols is Ethernet, which is almost universally chosen. Ethernet is an IEEE (Institute of Electrical and Electronic Engineers) standard, IEEE 802.3. It specifies how to perform the functions represented by Layers 1 and 2 of the protocol stack. It specifies how particular types of physical media are used, how data is packaged into frames for delivery on the network, and the system of addresses used at Layer 2. Ethernet frames use a form of address called a MAC (Media Access Control) address. These addresses are uniquely assigned to each Ethernet card, the network interface cards installed into the devices on an Ethernet LAN, by card manufacturers.

Different forms of Ethernet networks have different amounts of bandwidth. As originally developed, Ethernet operated at 10 Mbps. It has since become considerably faster. Variants are available that have 100 Mbps bandwidth, called Fast Ethernet, and 1 Gbps (1 Gbps = 1,000 Mbps), called Gigabit Ethernet. The 10 and 100 Mbps Ethernet are the most common forms of Ethernet. Gigabit Ethernet is relatively new. Ethernet continues to become faster; 10 Gbps Ethernet has been developed. Some 10 Gbps Ethernet devices are already on the market.

Ethernet supports many different forms of physical media. It can use copper wire, fiber-optic cable, or coaxial media. It can also be implemented as a wireless network. Together the bandwidth and the physical media of the Ethernet LAN specify the type of Ethernet employed on a LAN, for example 10Base-T for 10 Mbps Ethernet over twisted pair cable or 100Base FX for 100 Mbps Ethernet with fiber-optic cable. An Ethernet LAN must employ equipment meeting one of the Ethernet standards. It may use more than one form of Ethernet. A 100Base-T version could be used in one part of the LAN and 10Base-T in another.

ATM (Asynchronous Transfer Mode). Some alternatives to Ethernet exist. One alternative was Token Ring, but it is no longer widely employed. A second, more important alternative—important because at one time it was thought to be the successor to Ethernet—is ATM (Asynchronous Transfer Mode), which at one point was faster than Ethernet. Subsequently to the appearance of ATM, Ethernet became much faster.

ATM also has built-in control over Quality of Service (QoS). ATM continues to have an advantage over Ethernet in this area. Ethernet has no way to govern what applications receive certain amounts of bandwidth through the LAN. ATM can do this, so it can guarantee that applications like videoconferencing receive priority over other data traffic. For example, in a library, an ATM LAN can provide a dedicated amount of bandwidth for a videoconference and then share the remaining bandwidth among other applications such as Web browsing. Moreover, unlike Ethernet, ATM can be implemented to guarantee a certain level of QoS across the LAN and the Internet to other distant ATM sites.

Ethernet, on the other hand, can only protect QoS through careful network design. QoS can only be guaranteed by making sure the network has sufficient bandwidth and the right design, so users do not contend for network resources. Every application must have as much bandwidth as it needs.

ATM (asynchronous transfer mode) is the dedicated connection-switching technology that organizes digital data and transmits it over a physical medium using digital signal technology. (Definition compiled from www.whatis.com.)

Recommended Alternative: Ethernet Local Area Network

Despite the attractions of ATM, Ethernet remains the better known and more frequently chosen network. Ethernet equipment is more common, less expensive, and easier to find than ATM network equipment. It is easier to find staff that can understand Ethernet than ATM. ATM has lost its edge as a faster form of LAN. The ATM QoS advantage has not been sufficient to make it widespread as a LAN technology. Most organizations haven chosen to over-provision Ethernet networks instead.

ATM has been implemented to the desktop by a few organizations. In Ohio, ATM has been used for the state's K-12 educational network. ATM was selected primarily because it could be used to guarantee quality of service for educational video. But ATM has primarily been implemented in the backbone networks that join multiple LANs and in the networks of Internet service providers.

Costs

This network component has no specific cost. The choice of Ethernet over other alternatives generally decreases the cost of the LAN. Products for Ethernet LANs are more broadly available than products for alternative forms of LAN. The knowledge of how to work with Ethernet products is also more widely available and therefore less expensive.

#2. Cabling

Half the specification of a form of Ethernet is the bandwidth, "100Base …", the other is the type is the cabling to be used. Ethernet supports a wide variety of forms of cabling. One type must be selected for the LAN.

Alternative Solutions

Three alternatives for network cabling in an Ethernet LAN are available. These are coaxial cable, unshielded twisted pair (UTP), and fiber-optic cable. The choices for LAN cabling differ in the cost of their associated hardware, the bandwidth they can support, and the distance over which they can be used.

Coaxial Cable. Coaxial cable is of more historic than current importance. Some networks still have this form of cabling, but it is increasingly rare. Coaxial cable was the original type of cabling used for Ethernet LANs.

Ethernet coaxial cable is similar to the type of cable used for cable television. Two types of coaxial cable are used in Ethernet LANs, thick and thin. These are often called Thicknet and Thinnet.

The form of Ethernet with thicknet coaxial cable is called 10Base5. The 5 indicates that 10Base5 Ethernet can use up to 500-meter cable segments. Any type of Ethernet cable has a limit on the length of any individual segment of cable, but multiple segments may be connected using a device called a repeater. Specific forms of Ethernet have both a limit on individual segment length and on the number of segments that can be connected. Thicknet can connect segments with repeaters up to a distance of 2,500 meters. Thicknet supported a maximum bandwidth of 10 Mbps. Thicknet lost popularity partly because its thickness made it difficult to work with and install.

Thinnet improved on Thicknet by using a thinner cable that was easier to install. Like Thicknet, Thinnet is limited to 10 Mbps. Ethernet using this form of cable is called 10Base2. The maximum segment length is 185 meters but was rounded up to 200 meters for the basis of the Ethernet designation. Five segments can be connected together to cover a distance of 925 meters.

Unshielded Twisted Pair (UTP). Unshielded twisted pair (UTP) is the most prevalent form of LAN cabling used today. Although unshielded twisted pair is often just called twisted pair, shielded twisted pair does exist. Shielded twisted pair has been used for Ethernet LANs, but it is so rare that twisted pair in networking discussions usually means UTP.

Like coaxial cable, multiple forms of UTP are used in Ethernet LANs. These cables are designated by a series of labels, from Category 1 on up. All UTP cables contain two pairs of copper wires. Each pair of wires is twisted together. Two categories of cable are important, Categories 3 and 5. There are higher categories available, Cat 6 and 7, but these types of cable have not been standardized yet.

Category 3, also called Cat 3, is the minimal grade of UTP implemented in buildings for voice services. Many LANs have this type of cable. It can be used to implement 10Base-T Ethernet using two twisted pairs. With 10Base-T, whether using Cat 3 or a higher grade of cabling, segments may be up to 100 meters in length. It can support up to five contiguous segments, 500 meters. Cat 3 cables are used to implement 10Base-T Ethernet using two twisted pairs. They can also be used to implement a form of 100Base-T Ethernet, 100Base-T4, providing a bandwidth of 100 Mbps using four pairs, but Cat 3 cables are typically only used to implement 10Base-T Ethernet.

The next step up from Cat 3 is Cat 5, or Cat 5 enhanced. The higher category number designates the cable has a higher level of performance. Cat 5 enhanced, or Cat 5e cable, must meet higher performance standards than Cat 5. This type of cable can still be used to implement 10 Mbps Ethernet, but because of its better performance characteristics than Cat 3, it can more easily support 100 Mbps (100Base-T) and can support Gigabit (1000Base-T) Ethernet. Cat 5 cables can be used to implement 100Base-TX Ethernet, a form of 100Base-T that only requires two pairs. This is the most common form of 100Base-T. Like Cat 3 cable, Cat 5 segments can be no longer than 100 meters. 100Base-T is limited to a maximum combined segment length of 210 meters.

Gigabit Ethernet (1000Base-T) can only be implemented with Cat 5 or higher cable. Again, one segment can be 100 meters. Gigabit Ethernet over Cat 5 cable requires the use of a cable with four pairs of wires. Gigabit Ethernet is primarily used for backbone networks not in the LAN.

Fiber-optic Cabling. Coaxial and UTP cables use electrical signals to transmit data; fiber-optic cables use light. Fiber-optic cabling has an advantage in that network segments can be longer than with coaxial cable or UTP. Two types of fiber can be used for Ethernet LANs. The first is single-mode fiber. In this type of fiber, light takes one path down the fiber. Single mode fiber can be used to span long distances. Even with Gigabit Ethernet, segments can be up to 5 km. The second type of fiber is multimode fiber. Light can take many paths down the multimode fiber bouncing off the sides of the fiber. Multimode fiber cannot cover the same distances as single-mode fiber. Used to implement Gigabit Ethernet, its maximum segment length can be 220 meters. Both types of fiber can be used to implement Gigabit Ethernet or slower forms of Ethernet.

Recommended Alternative: Category 5e UTP (Unshielded Twisted Pair) Cabling

Category 5e UTP and fiber-optic cabling are the alternative forms of cabling used currently to implement high-speed LANs. When considering the choice between these two alternatives, you must understand the particular difficulties that relate to the investment in a cable infrastructure. Unlike many network components—for example network interface cards, which can be changed rather quickly—replacing the cabling used in a network is more time-consuming and, because of the increased labor cost, it is also more expensive. Thus the decision about a choice of cabling for a LAN cannot easily be changed. It has to be more of a forward-looking decision than other parts of the network infrastructure.

This need to look to the future in selecting a cable infrastructure seems to favor the use of fiber over UTP. Fiber has significantly greater potential bandwidth than any form of UTP will ever have. In the later discussion of optical connections to the Internet, how fiber can provide a bandwidth of many 10s of Gigabits (Gbps) is discussed. Fiber, though, is also perceived as having several problems. Vacca (1998) lists "four myths about fiber optics": 1.) Fiber optic cable is fragile; 2.) Fiber-optic cable is hard to work with; 3.) Fiber is expensive; and 4.) Fiber has no place to the desktop. Although Vacca sees these as myths, other authors say fiber-optic cable is more expensive, not because the cable is more expensive per foot, but because the related pieces of the infrastructure—for example, network interface cards and terminators are more expensive. They also believe fiber is more difficult to work with.

Where the author of this report works, thinnet cabling is being replaced by multimode fiber. Libraries and other organizations could combine a fiber-optic cable infrastructure with low-cost desktop UTP network cards. Devices, called converters, are available to be used to connect fiber-optic and UTP cables together. The fiber running to the wall jack can be plugged into a fiber-to-UTP converter. UTP can then be plugged into the converter and run to the desktop. This lowers the cost of network interface cards but adds to the cost of the converters and the overall cost of the network.

The disagreement over the costs of fiber indicates that fiber is more of a cutting edge technology than UTP and an industry consensus is that the amount of fiber used in LANS will grow. The decision to choose either UTP or fiber to the desktop needs to be made based on a library's probable bandwidth needs and an analysis of the cost of working with fiber. If the added cost of equipment is an issue, or the library does not have staff or a networking contractor familiar with fiber, then UTP is the better choice. It is also the better choice if the library is risk averse. As the more conventional choice, Category 5e UTP is recommended.

Costs

The cost per foot of Cat 5e UTP is around $0.20 per foot for PVC cable and $0.55 per foot for plenum cable. The cost of multimode fiber is around $0.35 per foot for PVC cable and $0.40 per foot for plenum cable.

#3. Network Interface Cards (NIC)

There must be network interface cards to create a connection between a computer and the LAN. When a decision has been made to have an Ethernet network and to wire it with Category 5e UTP cable, the characteristics of the network interface cards required have been fully specified. Use 10Base-T or 100Base-TX network cards is possible. Most NICs can run at both 100 Mbps and 10 Mbps, thus providing backward compatibility with older types of Ethernet LAN. Many computers already have this type of cards built in. The choice of Cat 5e UTP and 10/100 network interface cards for an Ethernet LAN allows you to run that LAN at 100 Mbps. This helps over-provision the network.

Costs

Individual cards can be obtained for about $20.

#4. Structured Cabling System

A structured cabling system is the alternative to an unstructured system. An unstructured cabling system requires ad hoc deployment of network cable among computers and the devices connecting them. With an unstructured network, when a computer moves you must move the cabling with it.

A structured cabling infrastructure provides a careful way to manage the cabling in a building and to ensure maximum flexibility when moving computers around—without having to change or move the cabling infrastructure. The Electronic Industry Association (EIA) and the Telecommunications Industry Association (TIA) set a standard for structured cabling: EIA/TIA-568-A, the Commercial Building Telecommunications Cabling Standard.

Implement cabling infrastructures by following the EIA/TIA standards. These standards specify how the cabling network of the building is to be physically laid out. A structured wiring system has many components. It has wiring closets, also called telecommunications closets, or intermediate distribution frames. The wiring from these closets to the wall jack near a user is called horizontal cabling or wiring. Multiple wiring closets may be on different floors. The wiring that connects the LANs on different floors is called vertical or backbone cabling. The vertical cabling connects the telecommunications closets to the main distribution frame.

The EIA/TIA standards specify how these series of closets should connect to each other and to the computers or other network devices. For UTP Ethernet, it requires that the horizontal cable, the cable that runs from horizontally in a building, should be no more than 90 meters from face plate (wall jack) to a patch panel in the wiring closet. The patch panel is then connected to network devices by the use of patch cables. Patch cables can be no more than 10 meters. Similarly, the potential length of the cable used between the faceplate and the desktop is limited. By regulating this use of cables, the standard ensures the LAN cabling infrastructure meets the Ethernet requirement of no more than 100-meter segments with UTP cable. It also makes the cable plant easier to manage and to cope with changes in locations of computers or in network equipment.

Costs

In 1998, *PC Magazine* (Cable Media Choices 1998) estimated the cost of implementing cabling in a structured environment for both Cat 5 UTP and multimode fiber. Each of these included the "Cost of 150 feet of cable plus hardware, labor, and testing (1998 prices)," and did "...not include connectivity at the network hub or network interface card." The two cost estimates are $100 to $150 per drop for Cat 5 UTP and $175 to $260 per drop for multimode fiber. A drop is a wall jack, so the cost is $100 to $260 per jack. These costs also do not include the cost of building and maintaining wiring closets.

These cost estimates need to be taken with a grain of salt as cost estimation is an imperfect science. Costs for cabling library facilities can be dramatically different depending on the type of library building. Installing cable in a new building with a dropped ceiling costs far less than installing a cabling infrastructure in a library in a historic structure. With these differences come differences in labor costs and costs related to locating and building wiring closets. Clear pricing only can be obtained by consulting electrical contractors.

A fiber-optic network infrastructure can decrease the cost of a structured wiring infrastructure. Longer segments are permitted with fiber-optic cables than UTP, so a greater distance between a computer and the wiring closet is possible. This can mean fewer wiring closets in a building, which can lead both to savings in the cost of building and supplying power or cooling services to the wiring closets and in the cost of equipment. Fewer network devices are required.

#5. Network Topology and Connecting Devices

So far, the checklist has identified the need for a 100 Mbps Ethernet LAN and UTP cabling implemented within a structured cabling system. Given such an infrastructure, there are alternative methods to connect the network segments from the wiring closet to the desktop. The shape of the network created by the connecting devices is called a topology. The devices' joining segments must be selected and the right topology implemented.

Alternative Topologies

Topology is a technical term that refers to the shape of a network. The physical topology of the LAN is essentially the geometric shape created by the cable segments of the LAN. The devices connected by these segments determine the logical topology of a LAN. Three logical or physical topologies can be used in a LAN. To introduce the notion of topology, the physical topologies are described first, then the logical topologies are described in the following discussion of the hubs and switches that connect network segments.

Topology is the shape of a network created by the connecting devices.

Bus Topology. This is original type of physical topology used with Ethernet. In this form of topology, a single cable connects a series of computers. Having a bus topology requires less cable than other forms of topology, but it has some serious disadvantages. The network can be brought down by a break in the cable. All the computers connected to the cable share the available bandwidth. If the Ethernet were operating at 10 Mbps, all the computers would have a part of that bandwidth.

Ring Topology. Physical ring topology networks are not used in Ethernet. They are found in token ring networks and in some forms of backbone networks. With a ring topology, a circular network joins all the computers connected to a network. A physical ring topology is more complex to cable because of the need to deploy a circular infrastructure of cables. The computers share the capacity of the ring. If the ring fails, the LAN does not usually go down because rings are usually implemented to combat cable failure with multiple redundancies or the bidirectional capabilities in the ring structure.

Star Topology. In a physical star topology, all computers are connected to a central device either a hub or a switch. A cabling infrastructure creates a physical star topology network. Each computer has a dedicated connection to this device. More cable is required than if the computers were all connected by one wire as in the bus topology. Star topology networks are more reliable than bus topology networks. The network only goes down if the central device goes down. Failures in cables only affect the individual computer connected to that cable.

Alternative Devices for Connecting Network Segments

At the center of the star topology LAN, there needs to be a device through which communications among the computers passes. Network cabling does not connect directly from one computer to another. This is the task of a hub or a switch. Routers are another type of network device with which people may be familiar. They do not play a role as a means of connecting computers together in a LAN. They have a role at the edge of the LAN and are discussed later in the Internet services checklist.

Hubs. Hubs operate at Layer 1, the physical layer, of the network. They receive Ethernet frames at one port and send them out on all the hub's other ports. This means that even with a structured cabling system and a physical star topology, the use of a hub creates a logical bus. Every computer connected to the hub shares the bandwidth available from the hub, since all communications are sent to every computer. Hubs are also sometimes called concentrators or multiport repeaters. Hubs are available with varying numbers of ports. Each computer connects to a port. Hubs may also simultaneously support multiple speeds of Ethernet. The 10/100 Mbps hubs are common. They can automatically sense what the computer connected to them supports.

Switches. A switch does not send Ethernet frames from one port to all the other ports. It only sends out data from one port to the port of the computer requesting the data. With a structured cabling system, it creates a logical star network. It does not divide the bandwidth available among the computers communicating through the switch.

Switches are Layer 2 devices. Because they work at the Data Link Layer, they can take advantage of the MAC addresses of the computers connected to the switch and map computer addresses to ports. The switch sends frames to a particular port based on the address in the frame.

Another distinction between switches and hubs relates to their effect on the limitations on Ethernet segment lengths. In a 10Base-T network, each Ethernet segment with UTP cable must be no more than 100 meters long.

Only four segments may connect; hubs can be used to connect those segments. But no more than four segments in any combination of segments and hubs may connect. Switches are able to regenerate a signal so the limit of four segments can be violated: several segments could be before the switch and then four segments afterward, making a switch valuable where long distances need to be covered.

Like hubs, switches may support multiple speeds. They may combine 10 Mbps, 100 Mbps, and Gigabit Ethernet interfaces. This support for multiple speeds lets you use the higher speed ports to connect to servers or network backbones where traffic from multiple desktops accumulates. Switches may be managed switches, meaning they support network management protocols, SNMP or RMON and may be capable of being remotely managed or monitored by network management software.

Several types of switches are available. In order of their increasing capacity, these switches are desktop, workgroup, and enterprise/backbone switches. Desktop switches have a small number of ports and may be used at a desktop to connect several computers or networked peripherals, such as a printer. A workgroup has a larger number of ports. It is the right size to connect all the computers in a small office or all the computers connected to one wiring closet. Enterprise switches are described in the checklist item "Backbone or Campus Network."

Recommended Alternative: Switches, Star Topology, LAN

Networks may use switches and hubs together, but wherever there are hubs, there is the potential for contention between the computers connected to that hub. The devices connected to a hub share the available bandwidth, so the quality of service applications like audio and video may be affected. A physical and logical star topology network is needed to ensure each computer receives dedicated bandwidth from the LAN.

Costs

Estimating costs for network switches is difficult. The cost varies dramatically from desktop to enterprise class switches. Workgroup switches can cost $1,000. Enterprise switches can cost more than $100,000. Switches may have different costs depending on how they are configured. Additional modules can be added to a switch to add ports or interfaces for specific types of Ethernet.

#6. Network Operating System and Server

If the LAN will be used to share resources such as a printer or files, a network operating system and server are required. The network operating system is capable of accepting and placing requests for access to shared resources. The main network server controls the rights of clients to access those resources. The clients are the desktop computers. In addition to the main network server a network may have several types of servers, file servers, print servers, communications servers, and applications servers, such as a database server. A server may both manage the network and run other services, such as shared file services or a database.

Three network operating systems are in widespread use. They are Novell Netware, Macintosh OS, and Windows NT/2000. Novell Netware has a diminishing share of the network market. The number of Macintosh computers compared with IBM compatibles is minute. For these reasons, network operating system and Windows are almost equivalent terms. Even if a library uses Netware for its network operating system, Windows 2000 is likely used within the library for servers that hold shared database systems or for other purposes, such as Web servers. The familiarity with Windows 2000 that results from these activities creates the strong likelihood it is also used as the main network operating system for the library LAN.

Windows network servers vary dramatically in their capabilities, from a computer not much better equipped then a desktop computer, to large servers with four or more processors. General characteristics of network servers include a large disk drive, a fast network connection, a backup device, and a UPS. The backup device and the UPS protect the computer from data loss or failure due to power problems or disk failure. A faster network connection than other computers on the LAN enables the server to cope with the high volume of communication between computers and the server.

Costs

The server version of Windows 2000 has a cost as does each client license. Microsoft 2000 server with a five-client license starts at $800. Network servers vary dramatically in price depending on the size of the LAN and number of users. As of April 2, 2001, one popular computer manufacturer was offering server models with the low-end model starting at $900 dollars and a high-end model starting at $6,700.

#7. Wireless LAN

An optional element of a LAN is a wireless access infrastructure. A wireless local area network still requires a traditional cable infrastructure. Wireless network access supplements the fixed cable infrastructure, allowing network users to work without being tied down to the vicinity of a network jack. Users can move around a building with their portable computers, working wherever convenient but still be connected to the Internet.

Two types of wireless networks have importance for library LANs. One is an already well-established technology. The standard type of wireless LAN is based on the IEEE 802.11b standard. Wireless LANs use hardware called a network access point that is plugged into a traditional cable network. The network access point or base station radiates a signal at a frequency of 2.4 GHz with a radius of 100 to 200 feet around its antenna. IEEE 802.11b Wireless LANS (WLAN) can provide up to 11 Mbps of bandwidth. Signal strength decreases with distance. Computers more distant from that antenna receive less bandwidth. The users of an access point share the available bandwidth. A base station can support up to three channels each with 11 Mbps. Different users can use each channel to reduce competition for bandwidth. Specific channels could also be used for particular types of users. In a library, one channel could be used by staff and provide access to the staff LAN. Another channel could be used by the public and be given access to the public LAN.

A new type of wireless LAN, 802.11a, is being standardized. This new wireless LAN uses a different frequency. It has significantly more bandwidth than the 802.11b wireless LANS, up to 54 Mbps.

The second current type of local area wireless, Bluetooth, is only just appearing. It also operates at 2.4 GHz. It is intended to cover a smaller area than an IEEE 802.11b LAN. The type of network created by a Bluetooth environment has been called a personal area network. Bluetooth is intended as a means of providing wireless connectivity among a user and multiple local peripherals like printers.

Costs

Wireless access points cost about $300. Portable computer cards cost about $150.

#8. Backbone or Campus Network

Large buildings or a set of buildings on one site creates the need to connect multiple LANs. Several LANs may be in a building, one on each floor, or one LAN in each of several buildings. The networks that connect LANs in these settings are called backbone or campus networks. A library staff likely has limited control over the backbone networks. Librarians may be able to choose the backbone that connects LANs in a library building, but the larger campus network is out of the library staff's control.

Backbone networks carry more traffic than ordinary LANs. The traffic from multiple LANs trying to reach resources on other LANs aggregate on the backbone network. To meet the higher performance requirements of a backbone network, different types of equipment are used to create that backbone network, which may also use a different type of networking than the LAN.

Fiber-optic cabling is more often used for backbone networks because of its higher carrying capacity. Backbone networks also require higher-performance network devices. A special class of switch, an enterprise switch, may be used in this environment. Enterprise switches have much more complex feature sets and a larger number of ports than workgroup switches. They have more high-speed interfaces and are capable of supporting interfaces to backbone network technologies such as ATM.

Ethernet is often used on backbone networks. Fast Ethernet, 100 Mbps, has been used to connect LANs in building backbone networks. With 100 Mbps running to desktops, higher speed versions of Ethernet are starting to be used for backbone networks. Gigabit Ethernet has an advantage as a backbone technology. It is well-known and organizations already have experience with Ethernet. Using fiber, particularly single-mode fiber, the backbone can also span a large distance.

ATM is another backbone choice. It can provide high-speed connections between LANs with up to 622 Mbps bandwidth. As was noted during the discussion of local area network protocols, ATM has the merit of being able to manage the bandwidth in a network, dedicating a part of the bandwidth to the connection between two computers, or in the case of a backbone network between two networks.

LAN Checklist Conclusion: The Multimedia LAN and its Effects on Library Applications

The recommended technologies in this checklist together create a picture of a multimedia LAN. The LAN would be a switched Ethernet LAN having a star topology. The cabling between the user desktop and the network switches is likely a Category 5e UTP but might also be fiber-optic cable. The cabling is implemented to meet the standards for structured cabling systems. Users have 100Base-TX network interface cards. A network server and network operating system provides controlled access to file servers and other shared LAN network resources. This LAN may be supplemented by an IEEE 802.11b standard wireless LAN. For multiple LANs, those LANs are connected to a backbone network, the LANs in one building may be connected to the networks in other buildings through a campus network. The backbone or campus networks may be Ethernet or ATM.

The LAN creates the ability to support library audio and visual applications by over-provisioning the network. Since an Ethernet LAN cannot guarantee specific users or applications a dedicated amount of capacity, it ensures those applications the bandwidth and low latency they require by reducing the possibility of contention for network resources. Category 5e cabling makes possible 100 Mbps at the user's desktop. By connecting the cable segments running to each desktop with a switch, creating a logical and physical star topology, the desktop computer user, whether library staff or patron, has all this bandwidth available to their applications.

Since most library applications require the use of the Internet, the question that must be asked is "How is a similar amount of network capacity obtained for the library's Internet connection?" The answer, discussed in the next section, is that providing the same level of Internet capacity as a LAN is nearly impossible.

Internet Checklist

When computers at one site need to be connected to computers at a geographically remote site, a wide area network (WAN) must be used. It is necessary to use the facilities of a telecommunications company to bridge the distance between these two sites. A private WAN could be created by leasing dedicated telecommunications lines from a telephone company, building a private WAN. More common is the use a public WAN. These carry the traffic of many organizations at the same time. The Internet is the best-known and most widely used public WAN.

What is the Internet? The Internet is a public network, a WAN, based on the IP (Internet Protocol). The IP is a protocol that operates at Layer 3 of the OSI protocol model. It provides a means of routing IP packets among a series of networks that may use different Layer 2 technologies. The Internet Protocol has a system of addresses, IP addresses, which are hierarchical and enable routing. The Internet is a worldwide, interconnected set of networks using the same protocol, IP.

Checklist Contents

This checklist focuses exclusively on listing what libraries need to use the Internet. Internet access is a problem for libraries. Although libraries can build a LAN with the capability of supporting the most advanced new applications, using the Internet for this purpose is more difficult. Libraries do not have the control over the Internet that they have over their LANs. The bandwidth required to support audio and video applications is more difficult to obtain. The digital divide makes Internet access difficult to obtain for many libraries, but even the libraries that do not have a digital divide problem often cannot afford the amount of bandwidth they need.

Because obtaining sufficient Internet bandwidth is a key problem for Internet services, the first item in this checklist is access to the Internet, which describes the many ways a library can connect to the Internet: cable modems, DSL, T-1 lines, and satellites. The costs of these modes of access are presented, as is a discussion of the E-Rate program that can help fund Internet access.

The second item in the checklist is a group of items needed to make the physical connection between a library LAN and the Internet. This includes routers but goes further to discuss the equipment needed to connect to the Internet securely: firewalls, virtual private networks, and proxy servers.

Then, the process of getting a place on the Internet is discussed. Internet access requires an Internet service provider. It also requires the library to obtain a block of IP addresses. And, if the library wants to host Internet services, it needs a domain name. These three elements, an ISP, IP addresses, and domain names, are the third element of the checklist.

This Internet checklist concludes with a discussion of Internet2. This new network is a subject of interest to many librarians who wonder what relevance it has to libraries and what relation it has to the current Internet. This interest is timely, since access to Internet2 is being extended to a broader array of organizations, and because some of the technologies being developed for Internet2 address fundamental problems of the Internet.

#1. Internet Access

At one time, libraries primarily provided patrons with access to electronic information resources housed in the library. In that environment, all that mattered was the LAN and its performance. Now, most—if not all—electronic information resources are located outside the library. These resources are accessed through the Internet. Libraries need connections to the Internet that provide patrons with access to information resources online. The mandate to libraries is clear—get a fast Internet connection, and get it now.

The mandate may be clear, but the solution is often more murky. The connection between the Internet and a business, library, or an person is often called the last mile. Much has been written about the problem of providing connectivity across this last mile. Many politicians, newspapers, and government organizations have talked about the digital divide, the gap between those who have the benefits of the Internet and those who do not. Bertot and McClure (2000) determined that 93.3% of rural libraries have Internet access. They also found that only 35.4% of rural libraries have greater than 56 kbps services up from 22.2% in 1998. Libraries have less connectivity than would be desired either to meet the needs of bursty Web applications or to support audio and video.

The digital divide is the result of two factors. The first is money; the farther away a library is from the place where it connects to an Internet service provider's network, the more access costs or the more limited the connection options are. For many in rural areas, those connection points are too far away and the cost of a high-speed network connection is too great. Access costs can be great even for libraries that benefit from e-rate discounts. The digital divide also results from the shortcomings of the local technical infrastructure. In many areas, local telephone infrastructure may not support high-speed Internet service, and DSL (Digital Subscriber Line) or cable Internet connections may not be available.

Librarians who want to overcome the financial or digital divide need to be aware of alternative means of connecting to the Internet. Librarians whose libraries have good Internet connections should also be interested in these developments. Implementing many of the audio and video applications described in the first section of this report requires more bandwidth than is currently affordable for many libraries. New technologies and the increasing deployment of current technologies offer the promise of eliminating the digital divide and making bandwidth more affordable.

Alternative Solutions

The following is a catalog of the ways a library can connect to the Internet, and though extensive, it's not complete. Researchers are investigating every possible means of connecting people to the Internet.

www.starcom.com

www.mediafusion.com

To appreciate the extent to which companies are trying alternative means of connecting to the Internet, consider Starcom LLC. This company invented a way of sending radio signals to computers by bouncing them off meteor trails. Another company, Media Fusion, is working on a technology for sending Internet traffic over power lines.

Although this catalog doesn't discuss meteor-trail-based Internet service, it covers the other more conventional ways of connecting to the Internet. These Internet access methods are grouped into several categories based on how they move the data through the last mile to a library. The first type uses the cable television system. The second group uses wireless for the last mile. It includes satellite, fixed broadband wireless, and mobile or cellular. The third group include access mechanisms that use leased lines, such as ISDN (Integrated Services Digital Network), DSL, T-1 lines, and fiber-optic networks.

Cable Modems. Accessing the Internet via the network of a cable television company is an alternative for small or under-funded libraries. Cable television networks are not normally built to support two-way communications. They must be upgraded for this purpose. Estimates vary on the number of cable systems that have been upgraded. A recent estimate (Walters 2001) is about 20%.

Cable modems provide about 1 Mbps downstream and can provide around 500Kbps upstream. Cable systems offer asymmetric bandwidth. There is more downstream bandwidth than upstream bandwidth, which is typical of many newer Internet technologies, including most forms of DSL. Actual performance varies because of a significant limitation with the design of cable Internet systems—their bandwidth is shared. Unlike a DSL connection where each house receives a dedicated amount of bandwidth, all the houses in a neighborhood share the bandwidth of the cable network. When many users are on the Internet, performance can drop dramatically.

Costs

See the TimeWarner cable service Web site for information about one cable system, www.rr.com.

The cost of consumer cable modem service is around $39.95 a month, with additional charges for more IP addresses—but a cable router could be used to provide more IP addresses rather than purchasing them from the cable company.

Wireless

For most ways of connecting to the Internet, one of the requirements is some type of physical cable running between the library and the facility of a telecommunications company. This infrastructure is expensive. Infrastructure development and maintenance comprises a great part of the overhead expenses of telecommunications companies. The time taken to install new equipment and cable delays the deployment of services. For telecommunications firms, the idea that Internet access or telephone service could be provided without a cable infrastructure is attractive. The idea of Internet access without a cable infrastructure has fueled the vision of wireless telecommunications.

Building on the knowledge and technologies produced by a long history of radio communications, a number of forms of wireless Internet access have been developed. These are in addition to the wireless LAN technologies described in the LAN part of this checklist. The two basic types of wireless Internet access are fixed and mobile. They are distinguished by the nature of the receiving antenna. Cellular telephone networks are an example of a mobile wireless network. Television broadcast systems are an example of fixed wireless; the recipient has a fixed antenna. Satellite systems are a form of fixed wireless. This analysis of wireless Internet access is divided into three parts: fixed wireless, mobile wireless, and satellite Internet access. Although qualifying as a form of fixed wireless, Internet access via satellite is addressed separately because of the differences in the marketing and development of satellite systems.

Mobile wireless may not seem to be relevant to libraries, who are not going to use mobile wireless networks to connect to the Internet. Mobile wireless is included, though, to highlight wireless network's potential impact on how patrons access Internet services. Describing both fixed and mobile wireless also helps reduce the confusion about what wireless Internet means. Discussions of wireless Internet access in the newspapers and other popular media refer as often to mobile wireless as to fixed wireless, but the same term wireless is used in both cases. An analysis of both forms of wireless help clarify their common and distinguishing features.

Fixed Wireless. Fixed wireless comes in two forms: point-to-point fixed wireless and point-to-multipoint wireless. The first requires two fixed antennas for a direct broadcast between those points. The second type requires one antenna broadcasting to multiple fixed receiving antennas.

Point-to-point wireless is a custom solution, implemented by arranging for a microwave link between an antenna tower and a customer site. A series of towers may be required to reach the distanced needed. Radio frequencies for use for radio communication are collectively called the spectrum.

A part of the spectrum set aside by the FCC for point-to-point connections can be licensed from the Federal Communications Commission (FCC). The operators of the point-to-point wireless networks usually license special microwave frequencies from the FCC. Unlicensed spectrum—frequencies the FCC lets be used for any purpose—can also be used. Point-to-point systems can provide a lot of dedicated bandwidth, up to 45 Mbps to a single location.

Point-to-multipoint wireless is more common as an Internet service. An ISP can serve many people through this infrastructure. The ISP uses an antenna, called a base station or cell site, to interact with many customer sites. As with point-to-point wireless, the base station is connected to the Internet by a physical link. The area covered by the antenna is called a cell. The cell is divided into a number of sectors. The base station antenna is designed so it provides all the bandwidth available to each sector. The frequencies are reused for each sector. All users in that sector share the bandwidth. The amount of bandwidth available within a point-to-multipoint system depends on the frequencies used and the technical design of the broadcast system. It also depends on the number of users in a sector.

Companies use different approaches to set up point-to-multipoint wireless networks. These approaches are distinguished by the type of frequency used. Two types of frequencies are available. The first are unlicensed spectrum. In an unlicensed part of the radio frequency spectrum, the FCC lets anyone use those frequencies for any application. The second type of wireless uses spectrum licensed from the FCC. Licensed spectrum can only be used by a specific organization in a limited geographic area.

The FCC allocates parts of the spectrum for licensed and unlicensed use. Specific bands are allocated for each of these types of use. Each band is a specific range of frequencies. The frequency matters because it affects how large an area can be served. The higher the frequency, the smaller the area within a cell is. At a high frequency, 25 or more GHz (gigahertz), the range of a signal is about 3 miles. At the other end, with a frequency of 2.4 GHz, the range of an antenna can be 30 miles. A smaller radius for the cell means more antennas are needed to cover any area, making development more expensive and the build out slower, thus reducing the likelihood of rural availability.

The frequency used for a wireless network can also affect the extent to which the service is limited by line of sight. Higher frequencies do not travel through things as well as lower frequencies. For most wireless services, a direct line of sight is needed between the base station and the customer antenna. Frequency can also affect whether the service has problems with weather. Higher frequencies are more likely to be affected by foliage, rain, and fog.

Two bands have been set aside for unlicensed use. Unlicensed bands can be used for any purpose, but the FCC offers no assurance that another application won't generate interference with a wireless system. These are the ISM (Industrial, Scientific, and Medical band), and the UNII (Unlicensed National Information Infrastructure) band. There are two ISM bands, one at 900 MHz and the other at 2.4 GHz. Most products use the 2.4 GHz band. This is the same frequency used for wireless LANs and Bluetooth. It is also used for microwave ovens and cordless phones. The potential for interference does not seem to be sufficient to decrease the willingness of some ISPs to use this frequency. The use of spread spectrum is supposed to limit the potential for this interference. The second unlicensed band is the UNII band, with three bands in the range of 5.1 to 5.8 GHz. UNII isn't used as much as ISM. Both unlicensed bands can cover a distance of up to 30 miles. One vendor, WiLAN, claims to be able to deliver up to 192 Mbps to a cell using unlicensed frequencies.

www.wilan.com

There are four licensed wireless bands: Instructional Television Fixed Service (ITFS), Metropolitan Distribution System (MDS), Multichannel Multipoint Distribution Service (MMDS), and Local Multipoint Distribution Service (LMDS).

The ITFS band had its origins in educational television broadcasting. It was set aside for use by religious and other nonprofit organizations in television broadcasting. A number of schools got ITFS licenses for this purpose. Recently, the FCC began to allow ITFS licensees to use their spectra for two-way Internet services. The ITFS band is at 2.1 GHz.

MDS and MMDS were similarly created for television broadcast. These two parts of the spectra are in the range of 2.1 to 2.7 GHz. They can provide up to 30 Mbps of bandwidth in a 6 MHz part of this spectrum. The cells can be 25 to 35 miles radius in size. Like the ITFS, use of these frequencies for wireless data services is new. One MMDS company, WavePath, is offering a two-way MMDS-based service with 1.5 Mbps.

www.wavepath.com

Several LMDS bands are available at 25 GHz and higher. The FCC estimates they can provide up to 155 Mbps to a cell. At this high frequency, the cells become dramatically smaller. For reliable service, given rain and other environmental problems, the cell radius may have to be 3 miles. LMDS will probably be used both for voice and data services.

Wireless data services have made limited use of LMDS and MMDS. These services are only now being piloted in large cities. The FCC restricts who can use licensed frequencies in any given area. They have been auctioning this right. These auctions have resulted in high prices for LMDS and MMDS licenses. The combination of high prices for spectrum and the cost of building out the infrastructure will probably delay the deployment of wireless using licensed spectrum in rural areas—where wireless is most needed. Unlicensed frequencies are more likely to be used in rural areas.

Mobile Wireless. Cell phones are ubiquitous. These devices receive and send data through the cellular telephone networks. But their bandwidth is limited—right now digital cellular networks can provide 14.4 Kbps. This limited bandwidth is a significant barrier to realizing the vision many technologists have for portable computing. They envision a single device to do all the tasks performed by cell phones, pagers, and PDAs with high-speed data access and the ability to communicate with video that this access would provide.

PDA is a personal digital assistant.

This has provoked a worldwide movement to develop a new generation of cellular technology, one to deliver a much higher bandwidth. This new type of cellular service is called 3G or third-generation cellular. This new service is expected to offer up to 2 Mbps. Creation of a 3G infrastructure will take time: It is not expected to be available until 2003 in the United States. It will reach rural areas after that time. 2.5G networks will appear first offering 100+ Kbps of bandwidth. To build 3G systems, there will need to be a transition from the current three types of cellular systems. These are GSM (Global System for Mobile Communications), CDMA, and TDMA.

The development of 3G cellular systems and the vision of one personal portable device will change how library patrons access library services. In the future, in addition to those patrons who search databases using library computers, there will be patrons wandering around the library searching those databases with their PDAs. This will change how library information systems need to be designed. Information may need to be formatted differently for these portable devices from traditional Web browsers. The future also brings the patrons wandering the library with personal devices who will not be using the library's bandwidth to access its information systems. They will be using the bandwidth of the 3G cellular systems. This may relieve some of the load on the library's Internet connection.

Satellite. There is a long history of using satellites for communication. Satellites were developed for the specific purpose of enabling communications between two ground stations. The use of satellites for the last mile of Internet access is new but builds on this long history. With some satellite Internet access products available now, and many more in development, this form of access will become more important during the next few years.

The three types of satellites are distinguished based on altitude. GEO (geosynchronous orbit) satellites orbit at about 22,300 miles above the earth. MEO (mid-earth orbit) satellites orbit 6,000 to 13,000 miles above the Earth, and LEO (low-earth orbit) satellites orbit 500 to 1,500 miles high. Current Internet satellite services are based on GEO satellites, but most new projects are building systems of LEO satellites.

The height of a satellite matters because the higher the satellite, the longer the signal takes to get to the satellite and back. The higher the satellite, the greater the latency is, too. With greater latency, supporting real-time audio and video interaction is more difficult. GEO satellites have a latency of about 0.25 second for a round trip. This is near the limit of the latency acceptable for voice communication. MEO satellites have a latency of 0.1 second. LEO satellites have negligible latency, around 20 milliseconds. Lower latency is one of the reasons for the development of LEO satellites. LEO makes possible the use of satellites for worldwide Internet and telephone services. Closer satellites can use less power and their antennas can be smaller. The higher the satellite, the higher the power of the satellite broadcast, and the larger the antenna needed.

Until recently, satellite Internet systems were unidirectional. It was possible to receive Internet communications from a satellite, but upstream communication required the use of a phone system. This is just now changing, and these systems are all changing to high-speed, two-way communication. The bandwidth they offer is asymmetric with bandwidth around 500K downstream and 150K upstream. One vendor, Optistreams, has suggested different services will be developed for the small office/home office (SOHO), and that these services will offer greater bandwidth than services for individual users. These SOHO products have yet to appear.

Two prominent LEO systems are in development: Teledesic and Skybridge. Teledesic will orbit at around 435 miles. Skybridge uses the Ku band, 10 to 18GHz. Its satellites will orbit at 913 miles. The Teledesic service is supposed to be launched in 2005. Skybridge is supposed to become operational in 2003. The dishes for the LEO systems will be from 10 to 24 inches in diameter.

Some problems come with LEO satellite-based Internet access. Because they are higher, GEO satellites can cover a larger portion of the earth and fewer satellites are needed. With LEO satellites a smaller area is covered, so more satellites are needed. This raises the cost of a LEO satellite infrastructure. Such satellite systems cost billions of dollars but only last a limited number of years because of the limited life of their batteries. No one knows if the developers of these systems have the right economic equation to succeed. The failure of the worldwide satellite system, Iridium, was because the economic equation was not right. This cost-effectiveness may plague future systems. The attraction of satellite Internet systems is that they can reach anywhere, but if there is widespread development of land-based wireless Internet access there will be a strong competitor to satellite.

Costs

Satellite service vendors:

www.directpc.com

ww.starband.com

www.tachyon.net

www.optistreams.com

Several satellite products available are using GEO satellites. The most well-known are DirectPC, Starband, Tachyon.net, and Optistreams. Starband became available in December 2000. It costs $69 a month with a $640 installation charge and charges of $399 for the dish and the receiver and $199 for the installation. The dish is 3 feet by 2 feet, and together with its related hardware weighs 55 pounds. Starband can be used in the continental United States. It requires unblocked exposure to the southern sky. All these systems have, or will soon offer, two-way Internet access.

Leased Lines From Telecommunications Carriers

Cable, wireless, and satellite are alternative ways of connecting to the Internet. The traditional way is to use the local loop between a building and a telephone company central office. The local loop is the cable that carries a number of pairs of copper wires between these two locations. It is the same infrastructure used for telephone voice services. For larger organizations, this cable may be a coaxial or fiber-optic cable rather than a cable with copper wires. The central office is the location where the telephone company has switches that switch phone calls to connect the participants in the call. The companies providing this service are called local exchange carriers. The telephone infrastructure connects the customer premises to the central office and then connects to an Internet service provider's network point of presence (POP).

POP (point of presence) is an ISP's actual point of access to the Internet. (Definition compiled from www.whatis.com.)

These local loops are used to implement three types of Internet service: DSL, ISDN, and the T-Carrier System. DSL is the newest of these three approaches. ISDN and T-Carrier lines, for example T-1 lines, are the traditional type of Internet data service using dedicated lines. T-1 lines are the most commonly used form of Internet connection.

Each of these three types of local loop based data services, T-Carrier, ISDN, and DSL, are described here. Following a description of these services, fiber-optic connections are discussed. These fiber connections provide higher-speed forms of Internet connection than those three services. Fiber-optic Internet connections are often used by large businesses. They are important to libraries because of the way that optical networking technologies are increasing the amount of data that can be carried by the Internet, and because of some fiber-optic projects that may make those fiber-optic connections available to libraries.

T-1 and other T-Carrier Facilities

T-1 lines are familiar to many people with only a little familiarity with telecommunications. They are the most widely used means for businesses and other organization to connect to the Internet. T-1 is one type of T-Carrier facility. T-Carrier is the general term for a family of transmission services. The T-1 carrier refers to the physical interface between a customer and the telephone company central office and the related electronics. A T-1 line carries a DS-1 signal. A DS-1 signal is made up of 24 DS-0 channels. A DS-0 channel is the amount of bandwidth required for digital voice transmission, 64 Kbps. Therefore, a T-1 line has a capacity of 1.544 Mbps. The term T-1 is often used interchangeably for the transmission service, the signal, and the amount of bandwidth. Fractional T-1 connections can be purchased in units of 2 DS-0 channels, 128, 256, 384 Kbps, and so on.

The other well-known form of T-Carrier line is a T-3. A T-3 carrier line can carry a DS-3 signal with 672 DS-0 channels for a total bandwidth of 45 Mbps. Fractional T-3 lines can also be purchased. Unlike T-1 lines that typically use the copper local loop, T-3 lines are usually implemented over fiber-optic or coaxial cables. DS-3 connections are used by large organizations including research universities.

Costs. The costs of Internet services based on a T-Carrier line will always reflect the cost of two services. The first is the cost of the access lines between the customer premises and the ISP's POP. This cost is distance sensitive. The greater the distance between the customer premises and the POP, the

LATA is a local access transport area.

higher the cost is. The area of a state that one local exchange carrier can serve is called a local access transport area (LATA). If the organization and the POP are not in the same LATA, the cost can increase dramatically, making the involvement of three companies necessary to connect the customer to the POP. The local exchange carrier is providing the T-1 line over the local loop. This local exchange carrier needs to send the traffic to the local exchange carrier where the ISP POP is located. This transfer requires the services of an Interexchange Carrier (IXC). The cost of the T-1 reflects the need for the services of all three organizations.

DSL and ISDN are simpler in this regard. There is no need to worry about costs associated with the distance between the library and the POP. All that is necessary is to have the right wiring available between the library and the central office.

The second cost is the actual Internet bandwidth. A physical network connection, such as a T-1 line, could be purchased without any IP bandwidth and could be used for other purposes than carrying Internet traffic. To be able to use the line to carry Internet traffic an Internet services provider is needed. The ISP charges the library for carrying its traffic to the Internet.

Access line charges can vary dramatically based on where an organization is located, from a few hundred to a few thousand dollars. The charges for Internet bandwidth are largely reflective of the average national cost per Mb of Internet bandwidth. These charges are currently around $400 per Mbps per month. They will be higher or lower depending on the length of a contract or the amount of bandwidth purchased. Costs for access lines are only dropping slowly, about 4% a year. Charges for IP bandwidth are dropping more rapidly, around 20% a year.

ISDN

ISDN stands for Integrated Services Digital Network.

ISDN is another means of connecting to the Internet through the telephone network. ISDN can be used to send both data and voice traffic together over a single phone line. The two types of ISDN are BRI (Basic Rate Interface) and PRI (Primary Rate Interface). The former provides up to 128 Kbps, and the latter 1.54 Mbps. A central office must have special equipment to be ISDN capable. PRI ISDN offers 23 channels each of 64 Kbps and one 16 Kbps channel for signaling. The 23 channels can be used partially for data and voice services or all for data services. BRI, in contrast, uses one pair of wires to provide 2 64 Kbps channels and one 16 Kbps signaling channel. With BRI, both channels can be used for data, 128 Kbps, or one for data and the other for voice.

Costs. Unlike a T-1 service, which uses a dedicated connection to the POP, ISDN is a switched service. Charges occur when the service is used to connect to a destination and transmit data traffic. Therefore, it is typically billed based on usage. Some providers offer flat rate plans.

DSL (Digital Subscriber Line)

One of the hottest new approaches to connecting to the Internet is DSL, called Digital Subscriber Line or Digital Subscriber Loop. Although DSL has lagged behind cable modems in adoption, it does not share the fatal flaw of cable television networks—shared bandwidth. All the bandwidth available over the central office to the customer premises site is available to the customer.

Unfortunately, the performance of a DSL connection depends on the distance between the customer site and the central office. As this distance grows the amount of bandwidth available diminishes. It can be available up to 18,000 feet from a central office. Some forms are only available for shorter distances. As with ISDN, the telephone company office needs to have special equipment installed for this service to be available in an area. This equipment is called a DSLAM.

The distance limitations of DSL have led to a new architecture, Digital Loop Carrier (DLC), for local loops that reduces the distance to the customer premises. DLC makes DSL available to more people, and it makes higher-speed bandwidth available. This new architecture puts the equipment that terminates the local loop in the customer's neighborhood.

Here are the main DSL variants:

Integrated Services Digital Subscriber Line (IDSL) is similar to BRI ISDN, but with notable differences. ISDN passes through the phone company's voice network and is switched. IDSL bypasses the voice network. It is an always-on dedicated form of ISDN. Unlike ISDN, IDSL is billed at a flat rate. Like BRI ISDN, IDSL offers up to 128 Kbps of bandwidth. IDSL is the form of DSL that can be available at the greatest distance, and accordingly offers the lowest bandwidth.

Asymmetric DSL (ADSL) is a more common flavor of DSL. It can offer up to 8 Mbps downstream and 1 Mbps upstream. The actual bandwidth depends on the distance from the central office.

DSL Lite is sometimes also called ADSL Lite or Universal ADSL. It is a form of ADSL that has the advantage of not requiring a service call for DSL installation.

High Data Rate Digital Subscriber Line (HDSL) is a symmetric service unlike ADSL and IDSL, which are both asymmetric. As does a T-1 line, HDSL uses two pairs of the local loop. It can provide up to 1.544 Mbps both directions. HDSL is sometimes called repeaterless T-1 and is used to implement T-1/DS-1 services. It has the advantage that it can reach up to 13,000 where T-1 lines require repeaters every 6,000 feet. This reduces the cost of the T-1 service.

Symmetric Digital Subscriber Line (SDSL) is a variant of HDSL. Unlike HDSL, it uses only a single pair of wires. It provides the same amount of bandwidth as HDSL.

Very High Data Rate Digital Subscriber Line (VDSL) is the highest bandwidth form of DSL. It can use fiber or copper to provide up to 52 Mbps downstream at 1,000 feet or 13 Mbps at 4,500 feet with 2 to 52 Mbps upstream. It only works over short distances.

Costs. DSL pricing varies dramatically depending on what flavor is wanted. More bandwidth costs more money. GetSpeed, **www.getspeed.com**, can be used to find out about DSL, cable modem, satellite, and wireless availability. To procure a DSL connection to the author's house could cost from $120 a month for IDSL to $400 a month for SDSL.

Optical Internet Connections and Wave Division Multiplexing

The amount of bandwidth that can be delivered with a copper local loop is limited. Fiber-optic networks can deliver more bandwidth than this. In fact the limit of fiber's data carrying abilities has not yet been reached.

Large businesses have taken advantage of fiber's capacity by connecting via fiber-optic cable to the SONET (Synchronous Optical Network) networks of telecommunications carriers. These carriers themselves have used SONET networks for their long distance and metropolitan area networks. SONET offers a series of levels of bandwidth called the optical carrier (OC) levels. These start at OC-1 with 51.8 Mbps, go to 155.5 Mbps (OC-3), 622.0 Mbps (OC-12), 1.24 Gbps (OC-24), 2.48 Gbps (OC-48), and OC-192 (10 Gbps).

The carrying capacity of SONET networks is the result of an optical communication technology called wave division multiplexing (WDM). The light sent down optical fibers can be divided into multiple wavelengths, for example blue and red. These different wavelengths can each become a different channel. Each color of light sent down the fiber can be a channel carrying a separate stream of communications. Each of these wavelengths is sometimes called a lambda. In the future, SONET will not be the way to gain the benefit of DWDM. Techniques are being developed to carry IP traffic directly over the fiber without the intervention of a SONET layer.

With current technologies and state-of-the-art equipment, each wavelength can provide a bandwidth of 10 Gbps (OC-192) and each fiber can carry 160 wavelengths. This provides an aggregate bandwidth of 1.6 terabits (1.6 million Mbps). This state-of-the-art form of WDM is called DWDM (Dense Wave Division Multiplexing), and it has ongoing development. It is expected to continue to increase the number of wavelengths that can be carried by one fiber.

With the cost of current bandwidth alternatives, networks providing larger amounts of bandwidth at higher costs may seem irrelevant to libraries. After all, if bandwidth costs reflect the general cost of IP bandwidth on the Internet—around $400 per Mbps per month—10 Mbps alone would cost around $4,000 per month. But higher-speed networks do have relevance to libraries. First, optical networks will dramatically increase the carrying capacity of the Internet. As the demand for high-speed Internet access to the home grows, demand for Internet traffic will grow. DWDM will offer the capability of increasing the amount of traffic carried on the Internet backbone. This alternative will provide a way of coping with traffic growth, so the cost of Internet traffic does not grow. Hopefully, the growth in capacity will be sufficient to allow for a steady and significant decrease in bandwidth costs.

A second key way fiber-optic networks impact libraries is through the availability of dark fiber. A dark fiber is a fiber-optic cable to which the electronics needed to transmit light has not been added. Many telecommunications companies have deployed fibers that are not used and so are dark; but these are not available to other organizations. Some projects are underway that would provide libraries and educational institutions with access to dark fibers. A group of educational institutions with access to its own pair of fibers could take advantage of the 1.6+ Tbps that DWDM can provide. Even if you couldn't buy that much bandwidth to the Internet, 1.6 Tbps could still be available for interinstitutional traffic, making possible the transmission of large amounts of data in the form of video.

Optical networks could have significant impact for distance education programs. One project in Cincinnati, Ohio, is proposing to give the higher education community access to a dark fiber from a bundle of fibers being laid around that city. This is probably happening elsewhere.

Comparing the alternatives

Which Internet connection alternative is right for a library? No one single answer is possible. For some libraries, whatever connection they can afford is best. Many areas have few, if any, affordable choices. In other areas, with multiple choices for how to connect to the Internet, a number of criteria could be used to judge the options.

When you have a choice, the first criteria to apply is whether the technology provides enough bandwidth. Several ways could be used to assess the amount of bandwidth a library needs for its Internet connection. The first is to reflect on the multimedia applications described earlier in this report. To support interactive videoconferencing, at least 384 Kbps would be needed to support one H.323 videoconference. Although supporting more bandwidth intensive forms of video would be nice, H.323 is of more immediate impact. Additional bandwidth would be needed to support Web browsing at the same time.

Another way to assess how much bandwidth is needed is to look at other libraries. For public libraries, examine the statistics gathered by John Carlo Bertot and Charles McClure (2000). They have conducted two studies for the National Commission on Libraries and Information Science. These studies have analyzed how much bandwidth public libraries use. The results of these studies are as follows: 36.2% of public libraries has T-1 service compared with 21.9% in 1998. Overall, 45.9% of the libraries has bandwidth of 128 Kbps or greater. No equivalent data is available for other types of libraries. In higher education, universities and colleges tend to have no less than a T-1 line, with larger campuses having several T-1 lines, and the largest a DS-3. Talk to peer institutions to determine what other libraries or their parent organizations have in terms of connectivity.

www.nclis.gov

When thinking about how much Internet bandwidth is enough, plan for increasing consumption of bandwidth. What is sufficient this year will rarely be sufficient the next. The continued increase in bandwidth consumption is shown by such phenomenon as the doubling of Internet traffic every eight to nine months. It is also demonstrated by phenomena such as the need of OARnet, Ohio's higher education ISP, to double its aggregate network capacity every year. The McClure and Bertot survey has also identified a year-to-year increase in the bandwidth purchased by libraries. To cope with a growing need for bandwidth, libraries must over-buy. Enough room is needed for growth in bandwidth use.

These hints on how to think about bandwidth do not provide an exact number, but they underline the need for bandwidth at least in the several hundreds of Kbps. They also suggest the desirability of the most common minimum amount of dedicated Internet bandwidth, 1.54 Mbps, used by nonlibrary organizations with dedicated Internet connections. Most of the alternative means of connecting to the Internet can provide this much bandwidth downstream. Not all can do so upstream. This limitation is another next reason for preferring one service to another. If other things are equal, most importantly cost, a symmetric technology is better. Symmetric bandwidth is needed for applications such as audio and video. It is also needed for libraries that would serve up data to the Internet. Although browsers in the library mostly download information, users at home and in their office are doing the reverse if they are searching the library catalog or digital library collection.

For more information on e-rate funding, go to: www.ed.gov/technology/eratemenu.html.

E-Rate Funding for Internet Access

The E-rate discount program has become essential to many libraries. It has given libraries the means of reducing the cost of telecommunications services and their internal network infrastructure. The E-rate program was created because of the Telecommunications Act of 1996 and the May 1997 FCC adoption of the Universal Service Order. The program's goal is to reduce the digital divide by ensuring schools and libraries have access to the Internet at a reasonable cost.

Through funds collected by the FCC from telecommunications companies, discounts are given to schools and libraries for purchase of telecommunications services—Internet access services, basic and long-distance telephone services—and for internal connections—internal LAN or WAN hardware. Only libraries whose budgets are not part of the budget of a larger organization, such as a school or a university, can directly apply for E-rate funds. This includes public libraries and other independent research libraries. School libraries can also benefit but do so as a result of discounts given to their schools. Libraries who can participate may do so directly or as part of a consortium.

The amount of discount available to a library depends on two factors: the library's location (urban or rural) and a measure of the economic disadvantage of the area. Economic disadvantage is measured by the percentage of students eligible for the national student lunch program. By combining these two factors, a discount ranging from 20% to 90% may be available. The following chart gives the numbers from the Universal Service Administrative Company E-rate Program Description:

Economic Disadvantage % of students eligible for school lunch program	**Discount Level**	
< 1	20	25
1-19	40	50
20-34	50	60
35-49	60	70
50-74	80	80
75-100	90	90

Limited funding is available for these discounts. Not every library that applies for discounts receives all the discounts it requests. Discounts are allocated in a series of waves based on a system of priorities. Priority is given to telecommunications services and Internet access. Priority is also given to organizations with a high level of economic disadvantage. The funding for internal connections to less needy organizations has been more limited.

An application process lets you seek e-rate funds. Part of the process is creating a technology plan that describes your technical infrastructure and goals for that infrastructure. A sample technology plan developed by the State Library of Ohio is included as Appendix A of this report. The creation of such plans would be worthwhile even for libraries that don't plan to seek e-rate funding. In combination with the checklist provided by this report, the planning process helps assess and evaluate the status of the library's technical infrastructure and the investments that will be needed during the next several years. The Public Library Association has created a guide to the

www.pla.org

development of technology plans (Mayo and Nelson 1999). It is entitled Wired for the Future: Developing Your Library Technology Plan.

#2. Equipment to Securely Connect to the Internet

Router

Libraries need routers between their LANs or backbone networks and the Internet. These network devices operate at Layer 3 of the protocol stack. They are able to make decisions about the route IP packets should take over the Internet. Routers open Layer 2 Ethernet frames to obtain an IP packet and then send it on to the next router. They also perform the reverse process with packets they receive from other routers that are destined for the LAN, putting the IP packet into an Ethernet frame.

Routing decisions are made based on an examination of the IP address in the IP packet. The address for the destination of the packet is looked up in a routing table. The table indicates which route to take. Routing tables are built up by the exchange of routing protocol messages between routers.

Three main routing protocols are used:

Open Shortest Path First (OSPF)

Border Gateway Protocol (BGP)

Routing Information Protocol (RIP).

Edge routers sit at the edge of a LAN, between the LAN and the Internet. Edge routers need to be equipped with the right type and number of interfaces to connect to the LAN, such as 100Base-T, and to a particular type of Internet access service, such as ISDN. They may have built-in modules for interconnection to popular Internet access lines, such as a T-1 line. But, additional hardware may be needed for that access outside the router. For example, a CSU/DSU (Channel Service Unit/Data Service Unit) needed to connect to a T-1 line may be built into a router or be purchased separately.

Routers can be placed between a LAN and a network backbone. They create a means to interconnect two kinds of networks with different Layer 2 protocols. A new form of a router, called a Layer 3 switch, which performs routing functions in hardware, can also be used for this task. Layer 3 switches are needed for high-speed backbone networks. Backbone switches may actually combine Layer 2 and Layer 3 switching, using either approach depending on the networks with which they are interconnected, and where the traffic is going.

Firewall

Essential library services provided by the network are not only threatened by hardware failures or poor network design, but they are also threatened by poor security. An inadequately protected network creates the risk that hackers could disrupt library services by attacking the library network or the computer systems it supports. Libraries can choose to ignore this and hope hackers aren't interested in their network, or to implement the kind of infrastructure that protects the library LAN. A long history of hacker activities directed against libraries shows the foolishness of the former strategy. An explicit security strategy is needed. Firewalls are needed as a basic building block of network security.

Firewalls protect the network from security threats coming from the Internet. They use three techniques to add protection. First, the firewall can hide information about the library network. It can implement IP masquerading, hiding the IP addresses of library machines. This may be done through network address translation (NAT). Second, the firewall can filter traffic only letting specific kinds of traffic pass from the Internet to the local network or vice versa. It can filter based on the origin or destination of network packets. It may also restrict traffic to specific computers and ports. Traffic that does not meet one of these criteria doesn't pass through the firewall.

Firewall features can be implemented within a router, in a special firewall appliance, or via the use of a host (computer) dedicated to this purpose. The firewall may be implemented so everything is inside the firewall or so some servers are located outside the firewall. When some computers are outside the firewall, the network is said to have a DMZ (Demilitarized Zone). By placing a Web server outside the firewall, services can be made available without giving access to the protected network.

Firewalls have caused some problems for libraries, such as when the need to communicate with a certain port of a library OPAC is blocked by the firewall. The configuration of the firewall needs to be changed to let this traffic through. Although firewalls can cause problems, these can be overcome given a good working relationship between the firewall administrators and the library staff. Having no firewall is not a better alternative than resolving these issues. Firewalls, given some cooperation with network managers, can coexist with library services.

Virtual Private Networks (VPN)

If a library has a secure network with a firewall, staff or other users working outside the network, likely from home, can't access to a LAN or other resources within the secure network. In the past, such connections were created by the use of private wide area network where the remote site was the only user of a dedicated network connection. Without other users of that dedicated line, security problems were limited. The Internet can be used to create the same type of remote access to the LAN, but without using expensive dedicated wide area network services. A protected, virtual private network (VPN) connection can be built over the Internet.

A VPN uses encryption to protect data traveling over the Internet. The encrypted transmission of information between two machines is said to form an encrypted tunnel over the Internet. Specific tunneling protocols are used for this purpose. Implementation of a VPN requires devices or software at both ends to encrypt and unencrypt data, which can be done through a dedicated host computer, router, or a specialized VPN gateway appliance.

Proxy Servers

Proxy servers are usually needed as part of a firewall infrastructure. They hide information about the network by channeling communications through the proxy. They can also perform other useful functions for a library.

University libraries have a problem with the combination of library vendor's use of IP addresses for authentication, and a growing number of off-site users who are accessing the Internet through commercial ISPs. Libraries like to use IP authentication as a means of securing access to their own and remote vendor Web sites. This is a simpler method of security than the use of passwords. When most university or college students used a university dial-up service to access the Internet, knowing what IP address their communications would come from was possible, and the system could to pass those IP addresses on to remote vendors. This no longer works so well now that many of those users are using commercial ISPs rather than the university or college dial-up services. Users do not have the right IP addresses to be given access to licensed resources.

One solution to this problem is to implement a proxy server where off-site users can access licensed or otherwise protected resources. The proxy server must be designed to authenticate off-site users. Users may perhaps need to use their campus password. Once authenticated, they can use the proxy and appear to have the same IP addresses as the proxy machine. If licensed or locally protected resources are accessible to this IP address, the proxy enables the users to gain access.

Proxy servers can also cache Web resources. For frequently accessed network resources, the proxy stores that page, and, in response to the next request, retrieves the page from its cache, rather than going out on the Internet to collect it, which conserves bandwidth.

Costs

Routers may be provided, and possibly maintained, by the library's ISP (Internet service provider). If not, they can be purchased for few thousand dollars up to many tens of thousands for high-performance routers. Dedicated proxy servers, VPN gateways, and firewalls also fall into the same price range. The cost depends on the set of features available and on the performance needed. The more expensive machines are for larger and more heavily used networks.

#3. A Place on the Internet: IP Addresses, Domain Names, and Internet Service Provider (ISP)

IP Addresses

Libraries need IP addresses when they connect to the Internet. The addresses provide the means of routing traffic to the library site. They are needed even when a library does not intend to host any information on its network. IP addresses are obtained from an Internet service provider. Internet access comes with a fixed number of IP addresses. In the case of a cable service, this may be as little as one address. Other providers may offer larger numbers of addresses. ISPs offer a number of IP addresses and then charge for additional addresses.

At one time, IP addresses were given out in class blocks. These could be Class A, B, or C blocks. Most organizations would obtain a Class C block of addresses with 254 addresses. Larger organizations would obtain a Class B block with more than 65,000 IP addresses. Because of the growth in the Internet, this approach to awarding addresses has ended. Now organizations seeking IP addresses from Internet service providers can get smaller blocks of addresses.

The limited supply of IP addresses has led to the invention of two approaches to managing IP addresses. The first is DHCP (Dynamic Host Configuration Protocol). This conserves IP addresses by only allocating IP addresses to computers when a computer is plugged into the network and turned on. In the past, IP addresses were statically assigned to each computer. The computers used up the available addresses whether they were actually on and connected to the network. The use of DHCP lets you conserve addresses.

Network address translation (NAT) does not conserve existing addresses; instead, it expands the number of IP addresses available to an organization. To the outside world that interacts with a network router, the organization would seem to be using the IP addresses assigned to it by the ISP. Inside the network, the organization uses whatever group of IP addresses it wants. These are mapped to one or more of the official addresses by a router or firewall. The internal addresses can be as numerous as the organization needs.

Domain Names

Use the Internet would be difficult, if, like computers, using IP addresses were necessary to distinguish computers. Every Web site would have to be known by its

IP address, such as 193.44.187.3. To provide a way to name network resources that users can remember, the Internet developed domain names. These can be translated by computers to IP addresses via the use of a DNS.

DNS stands for domain name service.

Domain names are registered with an official domain name registrar. There are several top level domains used by libraries, .org, .net, and .edu. Several new top level domains, .info and .museum, for which domain names are not yet being awarded may also be of interest to libraries. .org and .net addresses are registered by ISPs and other organizations that have been officially accredited as registrars by the Internet Corporation for Assigned Names and Numbers. When a domain name, such as ala.org, has been registered, additional subdomain names can be created, such as www.ala.org.

Internet Service Provider

An ISP is needed to provide the library with Internet connectivity whether wireless, DSL, T-1, or any other form of connection. The ISP plays several roles. Some roles relate to the hardware needed for Internet access and others to the services needed.

ISPs have to provide a customer with the hardware needed for Internet connections, either directly or by leasing it from a telecommunications carrier. In the case of T-1, DSL, and ISDN services, a local exchange carrier provides the local loop. With a wireless or satellite system, the ISP needs to provide the antenna or satellite dish. ISPs may also provide the router needed to connect to the Internet.

In addition to providing or contracting for the needed hardware, the ISP is a service provider. They supply IP addresses and may be able to register domain names. ISPs also supply the critical service of carrying IP traffic to the Internet. Unless the ISP is a national backbone carrier, the ISP contract with a larger ISP to carry its IP traffic. Thus, the ISP has to purchase sufficient bandwidth to serve all its customers. The ability to carry IP traffic is its key service.

A second key service of the ISP is to ensure the network stays operating—and to respond quickly when it does not. The ISP needs to maintain staff that can be on-call or monitoring the network. They respond to problems with the ISP network and with services to customer sites. ISPs may also offer other auxiliary services not related to carrying network traffic. They may host Web sites. Some also offer security services, implementing firewalls or VPNs.

In some areas, or with some types of services, such as cable modem Internet connections, libraries often have no choice of Internet service providers. Where you have a choice of providers of dedicated line services, the guide to evaluating ISPs is included as Appendix B of this report, Choosing an Internet Service Provider. It is a whitepaper written by an Internet Service Provider, Intermedia Communications. The whitepaper provides a detailed discussion of what such Internet service providers can and should do. Although the whitepaper is primarily focuses on T-1 services, much of what it says about the ISP networks and services would apply to the networks of DSL ISDN providers. DSL and ISDN services are not subject to the same problems of the location of ISP POPs, as are T-1 services, and may not be subject to the same kind of telecommunications circuit costs. But they are subject to the quality of the ISP network in the same way as a T-1 line user, and their users share the same concern about network support services.

www.intermedia.com

The Meaning of Internet2

The popular press has been talking about Internet2, particularly in the higher education press. These discussions probably leave librarians wondering: Is Internet2 relevant to libraries and why?

Internet 2 is important to libraries. It affects the library community in two ways. Internet2 was initially created as a research network but is turning into a general educational network. Internet2 was operated as a research test-bed by UCAID (University Consortium for Advanced Internet Development). The network was intended to be a test-bed for new networking technologies. By permitting this kind of research, Internet2 was to function in a similar fashion to the original Internet, now sometimes called Internet1, which early on was used to test networking technologies. This kind of research cannot be done on Internet1 any more because it has become too important as a production network. Now UCAID is talking about changing the goals of Internet2 and extending access to a broader community. They intend to operate two parallel networks, one for research and the other for educational applications.

www.ucaid.net

UCAID is repositioning Internet2 as a national educational network. It will open to a broader range of organizations than its original large research university members, and to a broader range of applications. It can be used to carry traffic among educational institutions at a lower cost than the original Internet. It will become a lower cost means of Internet connectivity than Internet1. Unfortunately, it will only be available for use to connect to other Internet2 sites, so it can't be used to carry the majority of traffic that goes to commercial Internet sites. Internet2 will enable libraries to deploy newer, more sophisticated, and more bandwidth-intensive applications that share information with other educational institutions.

No one can say yet whether public libraries will be able to use the Internet2 this way, but K-12 and higher education institutions will benefit, as will their libraries. K-12 schools and higher education institutions who want to take advantage of Internet2 will not become direct members of the Internet2. They will become a kind of auxiliary member. Full membership will continue to cost a lot of money.

Organizations that want to join will connect through an organization called a GigaPOP. These regional organizations provide aggregate connections to Internet2 for a group of universities. The GigaPOPs will sponsor other organizations that want to connect to Internet2. They will gain associate membership in this way. Internet2 participation will require organizations to have two network connections: one to Internet2 and one to Internet1.

The second impact of Internet2 on libraries will come from the networking research and development taking place on Internet2. Internet2 has an important role in testing new network technologies. It has played an important part as a test-bed for H.323 videoconferencing on IP networks. It is also an important arena for communication about new security technologies such as digital certificates.

A third, and much more important potential area of research on Internet2 has been QoS on IP networks. When ATM has been described in this report, it has been made clear that Ethernet and IP networks do not contain any mechanisms to protect QoS. The only way to ensure applications receive the network resources they need has been to over-provision networks—having high performance network equipment and more bandwidth than needed. In this way, any contention for network resources is nearly eliminated. When contention occurs for network resources, this approach does not solve the problem of whose applications receives priority. IP networks can't determine who receives what class of service. Organizational priorities don't make this determination possible.

The ability to control which network applications have priority is desirable. In a university, faculty doing research or collaborating with other faculty could receive a higher priority than students looking at recreational Web sites or downloading music. In a corporation, preferred treatment can exist for other reasons. Other applications using bandwidth must slow down while bandwidth was reserved for these priority users.

Over-provisioning networks is also problematic because it is an expensive solution. Internet bandwidth is expensive and few organizations can afford to buy more bandwidth than they need. Over-provisioning a LAN may be possible since there is complete control over implementation, but the Internet doesn't provide this kind of control. Increased bandwidth costs more money and will continue to cost more money year after year.

MPLS stands for multiprotocol label switching.

A different solution than over-provisioning is needed, one that also assures certain applications or certain users receive priority in contentions for network resources. Internet2 leads the way in solving this problem for IP networks. Current Internet2 research focuses on a new technology, multiprotocol label switching (MPLS), as a potential solution.

MPLS changes how IP packets are sent across the network. A router at the edge of a network labels each packet. Next, all subsequent routers would examine these labels to determine how to treat the packets. Because opening the packets isn't necessary as they are routed over the network, they can be routed quickly. Any latency introduced by routing is decreased. Each router replaces the label with information about the next destination. Routers can remember how they have treated packets and can treat subsequent packets the same way. This further speeds processing and sets a virtual path through the network.

Different applications or users can receive a different quality of service because packets can be assigned to classes, and those classes are reflected in the label. Classes of traffic may receive the same treatment throughout the network. The classes can be determined based on network requirements of the traffic or on the source or destination of the packets. Association of packets to classes can be based on policies set up by network administrators.

The result of further development of MPLS and its widespread use could be the ability to control QoS on IP networks. You might be able to assign packets from specific computers a certain path through the network based on their IP address. Thus a faculty member's computer could be assigned to a certain class of traffic and receive a certain quality of service because the faculty member because of policies set by network managers.

Although development of this technology appears to be succeeding, some practical difficulties need to be solved for the potential of MPLS to be realized. Multiple organizations would need to agree on policies for this approach to work. This might be possible on Internet2, but finding consensus on how to implement policies will be more difficult on other networks. Traffic crosses multiple networks on the way to its destinations. Consistent treatment of traffic would require the managers of all those networks to agree to treat classes of traffic the same way. For MPLS to be implemented, it would have to be adopted by router manufacturers and those MPLS routers would have to be widely installed. Given these not insuperable problems, MPLS will more likely have a greater and more immediate impact on networks run by one organization, such as a state K-12 network or Internet2, before it has widespread impact on general Internet traffic.