Chapter 1

i-Net+ Networking Basics

I-NET+ EXAM OBJECTIVES COVERED IN THIS CHAPTER:

- ✓ **3.7 Create a logic diagram of Internet components from the client to the server. Content may include the following:**
  - Bridge
  - Brouter
  - Router
  - Switch
  - Hub
  - Repeater
  - Network adapter
  - Cable Modem
  - xDSL Modem
  - Modem
  - WAN Link
  - CSU/DSU
  - Firewall
  - Network Address Translation (NAT) server
  - Proxy Server
- ✓ **3.8 Describe various hardware and software connection devices and when to use them. Content could include the following:**
  - Network adapter
  - Bridge
  - Brouter
  - Router

- Switch
- Repeater
- Hub
- Network adapter
- Cable Modem
- xDSL Modem
- Modem
- CSU/DSU
- Firewall
- Network Address Translation (NAT) server
- Proxy Server

✓ **3.10 Understand and be able to describe how common networking topologies are used. Content could include the following:**

- Star
- Bus
- Mesh
- Ring

By most accounts, the Internet is a big network. It contains many of the same components as any corporate network. To that end, before discussing the Internet, it is helpful to understand some of the basic components and concepts of a network. Many of the concepts involved in understanding networks will cross over to understanding the inner workings of the Internet. This chapter will introduce you to some of the more common networking topics you must understand when working with Internet technologies. Some of those topics include definitions of servers and protocols, hardware and software connection devices, and the various bandwidth technologies used to connect Internet sites to one another. This chapter will introduce you to these and other networking components and concepts so that you may have a better understanding of the Internet's underpinnings.

What Is a Network?

In the computer world, the term *network* describes two or more connected computers that can share resources. A *resource* can include data, printers, applications, fax devices, scanners, or any computer device that can be shared. The type used depends on the number of computers (and people) who need access, the geographical and physical layout of the enterprise, and of course, financial resources. Networks can be classified in two different ways:

- Network Topology: The network topology comes in several forms, and is most generally defined by the physical layout of the network.
- Network Type: The geographical location and physical size of the network generally define the type of network.

In this section, we'll discuss each type of network and describe the situation that is most appropriate for its use.

Network Topologies

From the computer itself to the cables that connect them together, the physical layout of a network is known as the network's *topology*. Since the topology is determined by factors such as the building's physical layout and the location of the network devices (also known as *clients*), you can think of a topology as a kind of network map. When choosing a topology for a network, you should choose the one that best facilitates connectivity between the network devices.

FIGURE 1.1 A network with multiple topologies

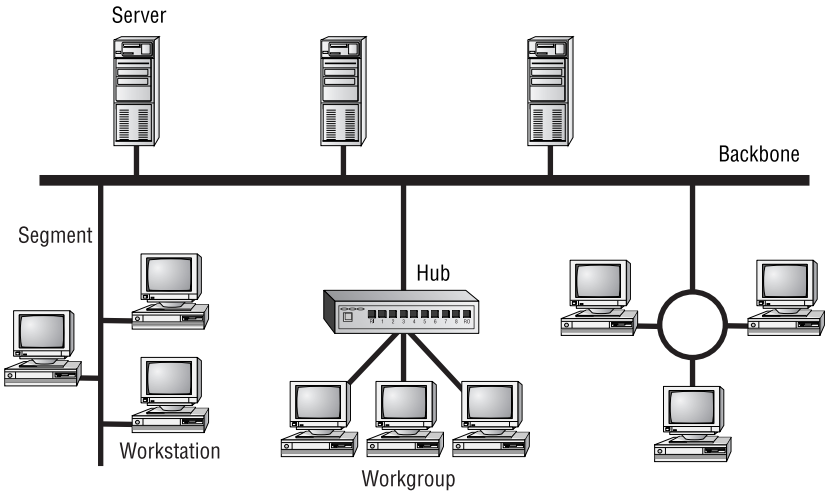

Networks tend to follow one of four basic patterns, although you do have mixtures more commonly in the real world. Figure 1.1 displays a network configured with several different topologies. The topologies that we discuss in this section are:

- Bus
- Ring
- Star
- Mesh
- Mixed Topologies

Bus

Networks generally have a specific area that handles a majority of the network traffic, called the network's *backbone*. The cables that connect the network devices to the backbone are called *segments*. In a *bus topology*, as depicted in Figure 1.1 on the left side of the graphic, the individual clients are directly connected to the backbone with their own cable. If you have relatively few devices in a small area to connect to the network, the bus topology works well; however, the major drawbacks to this configuration are that adding new devices directly into the backbone can take some skill, a cable break affects the entire network, and more than one device could talk on the network at the same time (causing a collision to occur). Bus networks are typically used in Ethernet networks.

Ring

In a *ring topology*, every network device is attached directly to both of its neighboring devices, forming a ring pattern. Ring networks are primarily installed in Token Ring networks, and as we will discuss later in the chapter, have the advantage of preventing two devices from talking on the network at the same time; however, you still have issues with network failure if one of the cable segments fail. Looking back at Figure 1.1, the right side of the graphic displays a ring network.

Star

In a *star topology*, each network device connects directly to a central clustering device, such as a hub or a switch. This method creates a starburst pattern that gives it the name. The advantages you have with a star topology is that if a cable breaks, finding it and replacing it is easy, and that the other network devices are not affected by the broken cable. Star networks are typically seen in Ethernet networks. Figure 1.1 depicts a star network in the middle of the graphic.

Mesh

The least common topology in use is the *mesh topology*. As its name implies, each networking device connects directly to the central server and to every other network device. The main advantage of this topology is that even if one cable breaks, you still have a number of other cables that the device can use to communicate with the network. Unfortunately, the sheer number of connections increase dramatically with the number of network devices—$n(n\text{-}1)/2$ with n being the number of network devices. For example, if you had a network of eight devices, you would need 8(8-1) or 28 connections. Not too

bad, you may think, but if you increased that number to 15, that number becomes 2↑ (15-1) or 98 connections. Trying to find one broken connection becomes the proverbial needle in the haystack. Another drawback to the mesh topology—and the main reason that it isn't commonly used—is that it can get expensive to set up and maintain.

Mixed Topologies

A *mixed topology* is one where you combine two of the four topologies together to gain the advantages of both. While not on the exam, mixed topologies are common enough that you should at least know what they are. The two most common mixed topologies are:

- Star-bus: Combines both the star and the bus topologies, hence its name. A central device, such as a hub or a switch, is used to connect network devices to form the "star" part of the topology. The individual central devices are then connected to the backbone, which in turn forms the "bus" part of the topology. Figure 1.2 shows an example of a star-bus network. As we see later in this chapter, this topology is commonly used in Ethernet networks.

FIGURE 1.2 A network using the star-bus topology

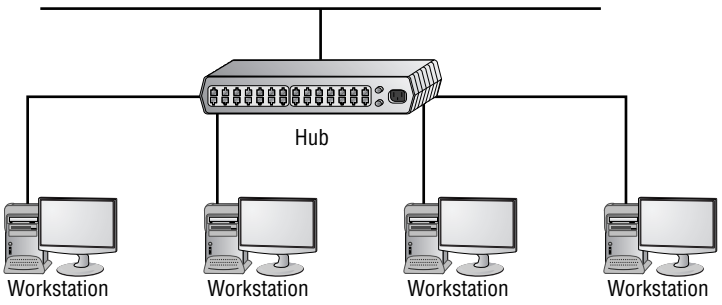

- Star-ring: Offers a central device to cluster individual network components together; however, instead of having a bus topology for a backbone, a ring is used. Star-rings are most commonly used in Token Ring networks, which is discussed later in this chapter.

Physical or Logical?

A network can also be described by two additional methods: physical topology or logical topology. The *physical topology* of a network is the actual, physical cabling of the network devices; however, the *logical topology* describes the manner in which the network really operates. For example, imagine that you had four workstations and one server connected to a hub running an Ethernet network. Physically, you have a star network; however, Ethernet works on a bus technology (more on Ethernet later, so therefore, logically you have a bus network.

The Main Network Types Defined

Now that you know the basic network topologies, you can examine the different types of networks and when they are used. Networks generally fall into one of three main categories:

- Local Area Network (LAN)
- Metropolitan Area Network (MAN)
- Wide Area Network (WAN)

Local Area Network (LAN)

By definition, a *local area network*, or LAN, is limited to a specific area, usually an office, and cannot extend beyond the boundaries of a single building. The first LANs were limited to a range (from a central point to the most distant computer) of 185 meters (about 600 feet) and to no more than 30 computers. Today's technology allows for a larger LAN, but practical administration limitations require dividing it into small, logical areas called *workgroups*. A workgroup is a collection of individual computers that share the same files and databases over the LAN, such as the sales department. Examples of LANs are shown in Figures 1.1 and 1.2.

Theoretically, a LAN can connect a maximum of 1,024 computers at a maximum distance of 900 meters (around 2,700 feet, assuming thinnet cable is used). These figures are based on connecting the segments with special devices to extend the overall range of the network to the backbone, and very light network traffic. If you use a different type of cabling, these maximums can decrease to 30 computers, with the most distant computer connected at a maximum of 100 meters (about 300 feet) from a central point.

Metropolitan Area Network (MAN)

A *metropolitan area network* is a network that generally does not leave the boundaries of a town or a city. The MAN is typically seen in college or university campuses that span one particular area of a town, or used by small companies that have offices in the same city but in different buildings. MANs typically use the local telephone company to connect individual LANs together, but can also use link types such as fiber-optic cabling and various forms of wireless technologies.

Wide Area Network (WAN)

Chances are you are already an experienced WAN user and didn't know it. If you have ever connected to the Internet, you have used the largest WAN on the planet. A *wide area network*, or WAN, is any network that crosses metropolitan, regional, or national boundaries. Most networking professionals define a WAN as any network that uses routers and public network links. The Internet fits both definitions.

WANs differ from LANs in the following ways:

- WANs cover greater distances than LANS or MANs.
- WAN speeds are slower than LAN speeds.
- LANs are limited in size and scope; WANs are not.
- WANs can be connected on demand or can be permanently connected. LANs have permanent connections between stations.
- WANs can use public or private network transports. LANs primarily use private network transports.

The Internet is actually a specific type of WAN. It is a collection of networks that are interconnected, and is therefore technically an *internetwork*. (*Internet* is short for the work *internetwork*.) The Internet will be discussed more fully in Chapter 2.

A WAN can be centralized or distributed. A *centralized WAN* consists of a central computer (at a central site) to which other computers and dumb terminals connect. The Internet, on the other hand, consists of many interconnected central computers in many locations. Thus, it is a *distributed WAN*.

Network Hardware Components

Networks are made up of many components, both hardware and software, and each hardware device on the network performs a different function. In this section, you will learn about some of these devices and their specific functions.

Network Adapter

The *network interface card (NIC)*, as its name suggests, is the device in your computer that connects (interfaces) your computer to the network. This device provides the physical, electrical, and electronic connections to the network media. It is responsible for converting the information your computer processes into the special electrical signals for the type of network technology your network uses. Also known as a *network adapter*, a NIC is either an expansion card (currently the most popular implementation) or built into the motherboard of the computer. If the NIC is built-in, it is called an *integrated NIC*, and is fast becoming one of the standard features offered to both the corporate and consumer markets for desktop computers. An example of a NIC is shown in Figure 1.3.

FIGURE 1.3 A sample NIC

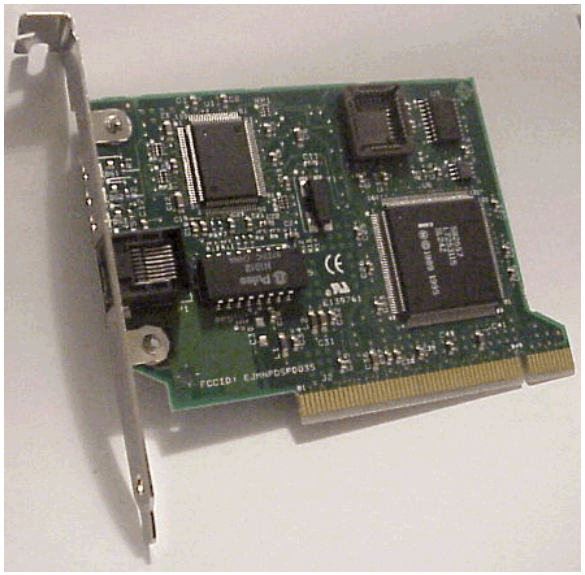

In some cases, a NIC must be added to the computer. It is usually installed into an *expansion slot* on the computer's motherboard. In notebook computers, NIC adapters can be connected to the printer port (known as a *parallel port*), through a built-in PC card slot (currently the most common method), or built-in (the latest trend in laptop computers).

To be used on a network, the NIC must have at least one protocol bound to it within the operating system. *Binding* a protocol means to logically associate a particular protocol with that instance of a NIC within an operating system so that the OS can communicate with the rest of the network using that protocol.

Regardless of which type of NIC you choose, the important thing to remember when buying a NIC for your computer is to buy one that matches the bus type in your computer and the type of network that you have. It sounds rather obvious, but you can't get a Token Ring card to communicate on an Ethernet network, no matter how hard you try. It just won't work because a Token Ring NIC wasn't designed for an Ethernet network. The electrical signals are in a completely different format.

It is never a good idea to mix-and-match NICs from different vendors on the same network. While you can get most NICs from one vendor to work in harmony with another vendor's product, sometimes the two "collide" and can cause problems throughout the entire network. Stick with one (two at the most) vendor's product, and you will enjoy fewer network issues.

Cabling

Although it is possible to use several forms of wireless networking, such as radio and infrared, most current networks communicate via some sort of cable. Although the i-Net+ exam doesn't test you on cabling technologies, it is important that we at least cover some of the common network cabling because, without cabling, most networks have no pathway to transmit data. In this section, we'll look at three types of cables commonly found in LANs:

- Coaxial
- Twisted-Pair
- Fiber Optic

Coaxial Cable

Coaxial cable (or coax) contains a center conductor made of copper, and is surrounded by a plastic jacket, with a braided shield over the jacket (as shown in Figure 1.4). A plastic, such as either PVC or Teflon, covers this metal shield. The Teflon-type covering is frequently referred to as a *plenum-rated* coating. That simply means that the coating does not produce toxic gas when burned and is rated for use in air plenums that carry breathable air. This type of cable is more expensive, but may be mandated by electrical code whenever cable is hidden in walls or ceilings.

FIGURE 1.4 Construction of a coaxial cable

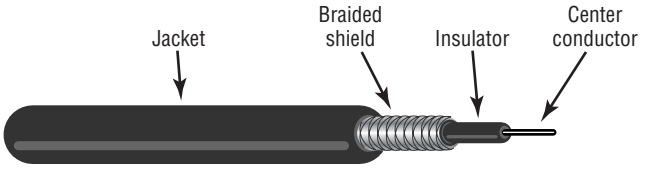

Plenum rating applies to all types of cabling.

Coaxial cable is available in different specifications that are rated according to the *RG* Type system. Different cables have different specifications and, therefore, different RG grading designations (according to the U.S. military specification MIL-C-17). Distance and cost are considerations when selecting coax cable. The thicker the copper, the farther a signal can travel—and you pay higher costs and receive a less-flexible cable.

There are two main categories of coaxial cable, Thick Ethernet (or *thicknet*) and Thin Ethernet (or *thinnet*). The primary difference between the two is the diameter of the cable and the distance they can carry a signal in a single segment. Thinnet coaxial can carry a signal 185 meters in a single segment, and thicknet can carry a signal 500 meters in a single segment. Thicknet cable has approximately the same diameter as a small garden hose and is difficult to bend. Thinnet cable, on the other hand, has approximately the same diameter as a pencil, is much more flexible, and thus easier to install. Of the two, thinnet is much more common in newer installations.

The main consideration with the installation of coaxial cable is the phenomenon of signal bounce. With coaxial cable, the signal travels up and down the entire length of the wire. When the signal reaches the end of the

wire, the electrical change from copper to air prevents the conversation from simply falling out the end. So the signal bounces back down the wire it just traversed. This creates an echo, just as if you were yelling into a canyon. These additional signals on the wire make communication impossible. To prevent this, you must place a *terminator* on each end of the wire to absorb the unwanted echo.

Proper termination requires that one terminator be connected to a ground. Connecting both terminators to a ground can create a *ground loop,* which can produce all kinds of bizarre, ghostlike activity, for example, a network share that appears and disappears.

Coaxial cable primarily uses BNC connectors. BNC has many definitions in the computer world. Some think British Naval Connector, citing its origins. Others would say Bayonet Nut Connector, after its function. Still others would say Bayonet Neill Concelman, after its authors. Suffice to say, it's just easier to call it a BNC connector and know that it's used on 10Base-2 Ethernet connections to RG-58 cable.

Twisted-Pair Cable

Twisted-pair cable consists of multiple, individually insulated wires that are twisted together in pairs. Sometimes a metallic shield is placed around the twisted pairs, hence the name *shielded twisted-pair (STP)*. (You might see this type of cabling in Token Ring installations.) More commonly, you see cable without the metallic shield, called *unshielded twisted-pair (UTP)*. UTP is commonly used in 10BaseT, star-wired networks.

The wires in twisted-pair cable are twisted to minimize electromagnetic interference. When electromagnetic signals are conducted on copper wires that are in close proximity (such as inside a cable), some electromagnetic interference occurs. In cabling parlance, this interference is called *crosstalk*. Twisting two wires together as a pair minimizes such interference and provides some protection against interference from outside sources. This cable type is the most common today, and is popular for several reasons:

- It's cheaper than other types of cabling.
- It's easy to work with.
- It permits transmission rates considered impossible 10 years ago.

UTP cable, the more common type of twisted-pair cable, is rated in the following categories:

Category 1 Two twisted-pair (four wires). Voice grade (not rated for data communications). This is the oldest category of UTP and it is frequently referred to as *POTS*, or *Plain Old Telephone Service*. Before 1983, this was the standard cable used throughout the North American telephone system. POTS cable still exists in parts of the Public Switched Telephone Network (PSTN).

Category 2 Four twisted-pair (eight wires). Suitable for up to 4Mbps. Typically used for telephone wiring and some older token-ring networks.

Category 3 Four twisted-pair (eight wires), with three twists per foot. Acceptable for up to 10Mbps. The popular cable choice for a long time, and is used in 10Base-T Ethernet networks.

Category 4 Four twisted-pair (eight wires) and suitable for 16Mbps. Used for Token-Ring and 10Base-T networks.

Category 5 Four twisted-pair (eight wires) and acceptable for 100Mbps. Used in 100Base-T and 10Base-T Ethernet networks.

Category 6 Four twisted-pair (eight wires) and rated for 155Mbps. Commonly used in fast Ethernet Networks.

Category 7 Four twisted-pair (eight wires) and rated for up to 1000 Mbps (Gigabit). Latest specification.

Frequently, you will hear *Category* shortened to *Cat*. Today, any cable that you install should be a minimum of Cat 5. We say "a minimum" because some cable is now certified to carry a bandwidth signal of 350MHz or beyond. This allows unshielded twisted-pair cables to reach a speed of 1Gbps, which is fast enough to carry broadcast-quality video over a network.

UTP cables use RJ (Registered Jack) connectors rather than BNC connectors. The connector used with UTP cable is called RJ-45, which is similar to the RJ-11 connector used on most telephone cables, except RJ-45 is larger. The RJ-11 has four wires, or two pair, and the network connector RJ-45 has four pair, or eight wires.

Signaling Methods

How much of a cable's available *bandwidth* (overall capacity, such as 10Mbps) is used by each signal depends on whether the signaling method is baseband or broadband. Baseband uses the entire bandwidth of the cable for each signal (using one channel). It is typically used with digital signaling.

In broadband, multiple signals can be transmitted on the same cable simultaneously by means of frequency division multiplexing (FDM). *Multiplexing* is dividing a single medium into multiple channels. With FDM, the cable's bandwidth is divided into separate channels (or frequencies), and multiple signals can traverse the cable on these frequencies simultaneously. FDM is typically used for analog transmissions, such as cable television. Another method, time division multiplexing (TDM), can also be used to further divide each individual FDM frequency into individual time slots. Additionally, TDM can be used on baseband systems.

Fiber-Optic Cable

If your data runs are measured in kilometers, or if you have gigabits of data to move each second, fiber-optic is your cable of choice because copper cannot reach more than 500 meters (around 1,600 feet—that's six football fields to you and me) without electronics regenerating the signal. Additionally, fiber-optic is the only cabling technology that can support the high data transfer speeds that the backbone of the Internet requires. You may also want to opt for fiber-optic cable if an installation requires high security because it does not create a readable magnetic field. The most common use of fiber-optic cable these days is for high-speed telephone lines.

Ethernet running at 10Mbps over fiber-optic cable is normally designated 10BaseF; the 100Mbps version of this implementation is 100BaseFX.

Although fiber-optic cable may sound like the solution to many problems, it has pros and cons just as the other cable types.

The pros are as follows:

- It's completely immune to electromagnetic interference (EMI) or radio frequency interference (RFI).
- It can transmit up to 4 kilometers.

Here are the cons:

- It's difficult to install.
- It requires a bigger investment in installation and materials.

Fiber-optic technology was initially expensive and difficult to work with, but it is now being installed in more places. Some companies with high bandwidth requirements plan to bring fiber-optic speeds to the desktop. At the

time this book is being written, the 10 Gigabit Ethernet Alliance (10GEA) is working on the 10G Ethernet standard for fiber optic cabling (which should be ratified in 2002), and fiber-optic networks will probably take off at an even greater rate when vendors begin shipping products by the end of 2001.

Servers

In the truest sense, a server does exactly what its name implies: It services client requests for access to resources on the network. Servers are typically powerful computers that run the software that controls and maintains the network. This software is known as the *network operating system*, which you will learn about later in this chapter.

Servers are often specialized for a single purpose. This is not to say that a single server can't do many jobs, but more often than not, you'll get better performance if you dedicate a server to a single task. Here are some examples of servers that are dedicated to a single task:

File server Allows for a central storage area that clients can use to share data.

Print server Controls and manages one or more printers for the network.

Proxy server Performs a function on behalf of other computers. Proxy means "on behalf of."

Application server Hosts a network application.

Web server Holds and delivers web pages and other web content and uses the Hypertext Transfer Protocol (HTTP) to deliver them.

Mail server Hosts and delivers e-mail; is the electronic equivalent of a post office.

Fax server Sends and receives faxes (via a special fax board) for the entire network without the need for paper.

Remote access server Hosts modems, or VPN connections, for inbound requests to connect to the network; provides remote users (working at home or on the road) with a connection to the network.

Telephony server Functions as a "smart" answering machine for the network; can also perform call center and call routing functions.

Network Address Translation (NAT) Server Translates a client's network address into an Internet address.

Regardless of the specific role(s) each server plays, they all (should) have the following in common:

- Hardware and/or software for data integrity (such as backup hardware and software)
- The capability to support a large number of clients

Physical resources, such as hard drive space and memory, must be greater in a server than in a workstation because the server needs to provide services to many clients. Also, a server should be located in a physically secure area. Figure 1.5 shows a sample network that includes both workstations and servers. Note that there are more workstations than servers because a few servers can serve network resources to hundreds of users simultaneously.

If the physical access to a server is not controlled, you don't have security. Use this guideline: If anybody can touch it, it isn't secure. The value of the company data far exceeds the investment in computer hardware and software.

FIGURE 1.5 A sample network including servers and workstations

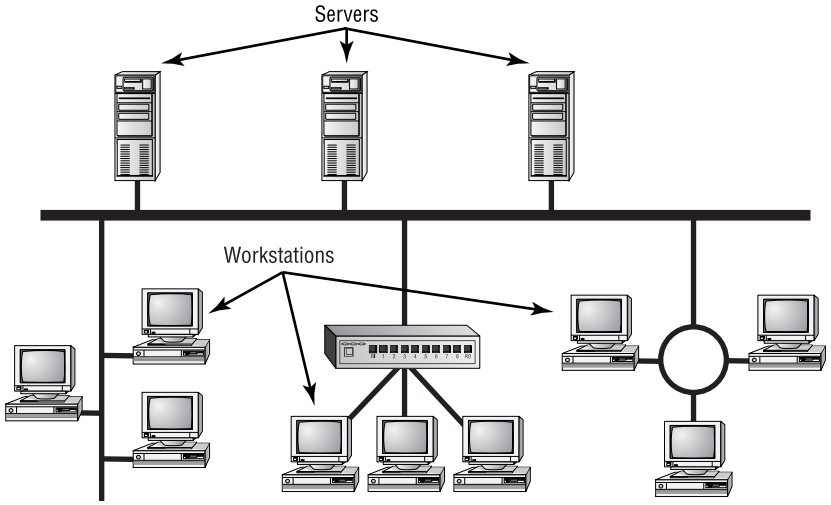

Do you Protect Your Server?

Most institutions list a server backup as a requirement in their disaster recovery plans, or if they don't have a plan at least somewhere in their operational policies; however, if money is tight, they generally forego the expense. This isn't a good idea, because a company's data (and their business) resides on the server(s) that it maintains. In one case, a company's employees failed to check that the backup program had been successful in completing the process. This went on for months until a server crash necessitated the restoration of the data from the backup tapes. The damage was about $5 million dollars and several employee jobs. The moral of the story—always check your backups!

Repeater

While cables connect networks together, it is the electronic signal that contains the information passed between two networking devices. Unfortunately, electronic signals can be interfered with on its way to its destination, and will degrade due to electrical resistance from the cable itself (a process called *attenuation*). Attenuation therefore affects how far away you can place a workstation. A *repeater* is a device that is placed between the workstation and the rest of the network to amplify the signal, thus allowing a workstation to be place further on the network than if it used just one cable. However, you can not use an unlimited amount of repeaters because most networks require a response on the network in a specific amount of time.

Bridge

A *bridge* is a network device, operating at the Data Link layer of the Open Systems Interconnection (OSI) model, that logically separates a single network into two segments but enables the two segments to appear to be one network to connected workstations. The primary use for a bridge is to keep traffic meant for stations on one side on that side of the bridge and not let it pass to the other side. For example, if you have a group of workstations that constantly exchange data on the same network segment as a group of workstations that don't use the network much, the busy group will slow down the performance of the network for the other users. If you put in a bridge to separate the two groups, only traffic destined for a workstation on the other side

of the bridge will pass to the other side. All other traffic stays local. Figure 1.6 shows a network before and after bridging. Notice how the network has been divided into two segments; traffic generated on one side of a bridge will never cross the bridge unless a transmission has a destination address on the opposite side of the bridge.

FIGURE 1.6 A sample network before and after bridging

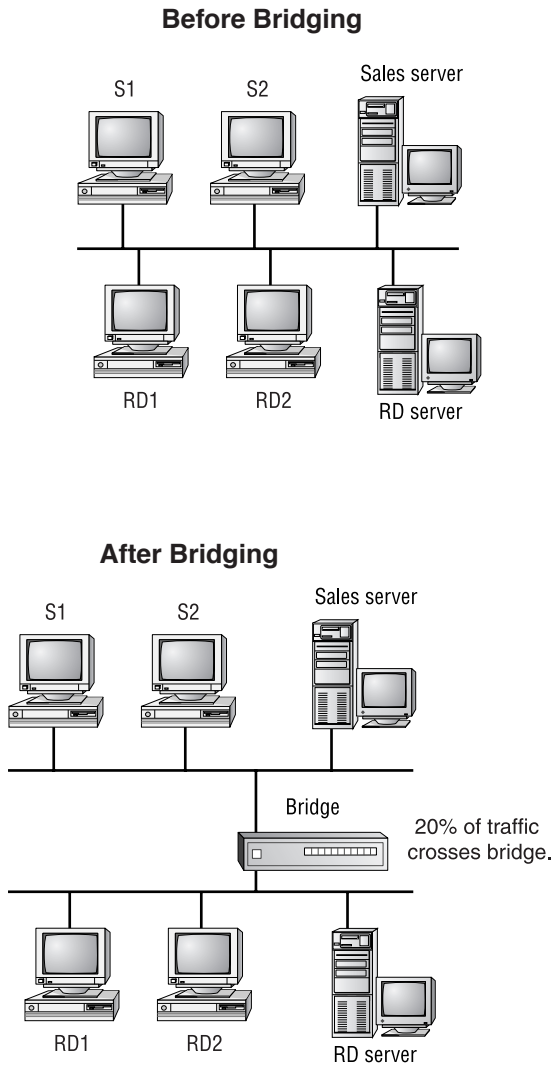

In addition to segmenting a network, bridges also allow you to route non-routable protocols, such as NetBEUI. This can be desirable if your network requires this type of situation, but can re-introduce any traffic problems that you previously experienced.

Hub

After the NIC, a hub is probably the next most common device found on networks today. A *hub* (also called a *concentrator*) serves as a central connection point for several network devices. At its basic level, a hub simply repeats everything it receives on one port to all the other ports on the hub, and doesn't care what stations are connected, thus it provides a communication pathway for all stations connected to it. Figure 1.7 shows an example of a hub.

FIGURE 1.7 A standard hub

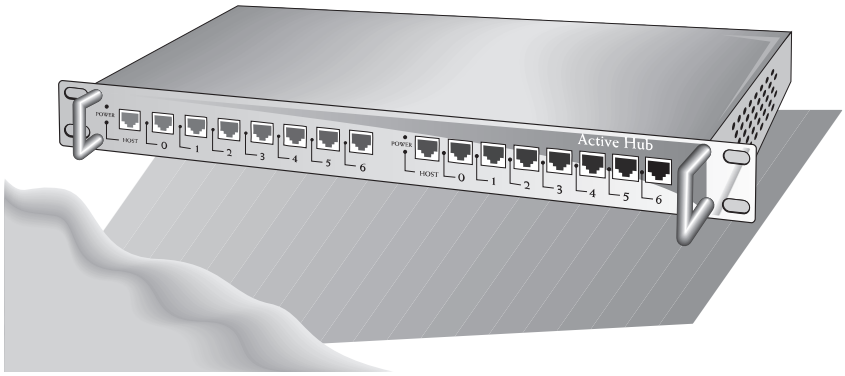

Hubs are found on every twisted-pair Ethernet network, including those found at ISPs. Hubs are used to connect multiple network devices together. ISPs may have several Internet servers connected to a hub, which is in turn connected to the ISP's Internet connection, allowing the servers to communicate with each other as well as with the Internet.

There are many classifications of hubs, but two of the most important are active and passive:

- An *active hub* is electrically powered and actually amplifies and cleans up the signal it receives, thus doubling the effective segment distance limitation for the specific topology (for example, extending an Ethernet segment another 100 meters).

- A *passive hub* typically is unpowered and makes only physical, electrical connections. Normally, the maximum segment distance of a particular topology is shortened because the hub takes some power away from the signal strength to be able to do its job.

Switch

In the past few years, the *switching hub* has received a lot of attention as a replacement for the standard hub. The switching hub is more intelligent than a standard hub in that it can actually understand some of the traffic that passes through it. A switching hub (or *switch* for short) listens to all the stations connected to it and records their network cards' hardware addresses (see Figure 1.8). Then, when one station on a switch wants to send data to a station on the same switch, the data gets sent directly from the sender to the receiver. This is different from the way hubs operate. As mentioned in the previous section, hubs don't care what stations are connected and simply repeat anything they receive on one port out to all the other ports. Because of this difference, there is much less overhead on the transmissions and the full bandwidth of the network can be used between sender and receiver.

Switches have received a lot of attention because of this capability. If a server and several workstations were connected to the same 100Mbps Ethernet switch, each workstation would receive a dedicated 100Mbps channel to the server, and there would never be any collisions.

FIGURE 1.8 A switch builds a table of all addresses of all connected stations

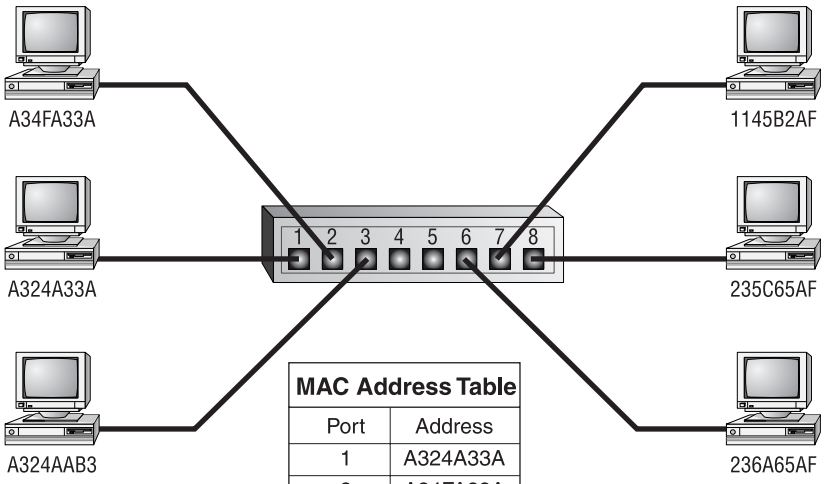

MAC Address Table	
Port	Address
1	A324A33A
2	A34FA33A
3	A324AAB3
4	—
5	—
6	236A65AF
7	1145B2AF
8	235C65AF

Router

Routers play a major part in the Internet. As a matter of fact, the structure of the Internet is made up of two major items: routers and phone connections. (Phone connections are discussed later in this chapter.) A *router* is a network device that connects multiple, often dissimilar, network segments into an internetwork. The router, once connected, can make intelligent decisions about how best to get network data to its destination based on network performance data that it gathers from the network itself. Because the router is somewhat intelligent, it is much more complex and thus more expensive than other types of network connectivity devices.

Router Ports

A router is not much to look at. Most routers have metal cases and are roughly 19 inches wide, approximately 14 inches deep, and anywhere from 1.5 inches high to 2 feet high with the more complex models. A typical router has multiple *ports*, or connection points, so that it can connect to all kinds of different network segments and route traffic between them. But at the bare minimum, most routers have at least three ports, and each has a different use.

Each port connects to a different device. For example, the most common port found on a router (there may be many of these ports) is a high-speed serial port (usually labeled something like WAN 0 or Serial 0). This port usually connects to either a modem bank or a WAN connection device like a Channel Service Unit /Data Service Unit (CSU/DSU), which is used to connect a router to a T1 phone line, discussed later in this chapter.

The second type of port is the port that connects the router to the LAN. It is usually an Ethernet port that you would connect to a hub so that the router could communicate with the rest of the LAN. It is usually labeled something like LAN 0 or Eth0 (for an Ethernet router).

The third type of port that some routers have is what is called an *out-of-band management port*. This port is a serial port (that most often uses an RJ-45 connector) that you connect to a terminal or PC running terminal software so you can configure the router. Some routers forego this port in favor of in-band management, meaning that you run the management software on a PC connected to the network and configure the router over the network. Some routers have one or the other, but many high-end routers have both to allow you the most flexibility in configuration.

Figure 1.9 shows an example of a router and some of the most common items found on routers today. Note that the router shown in Figure 1.9 has two serial ports, a LAN port, and an out-of-band management port.

FIGURE 1.9 A sample router

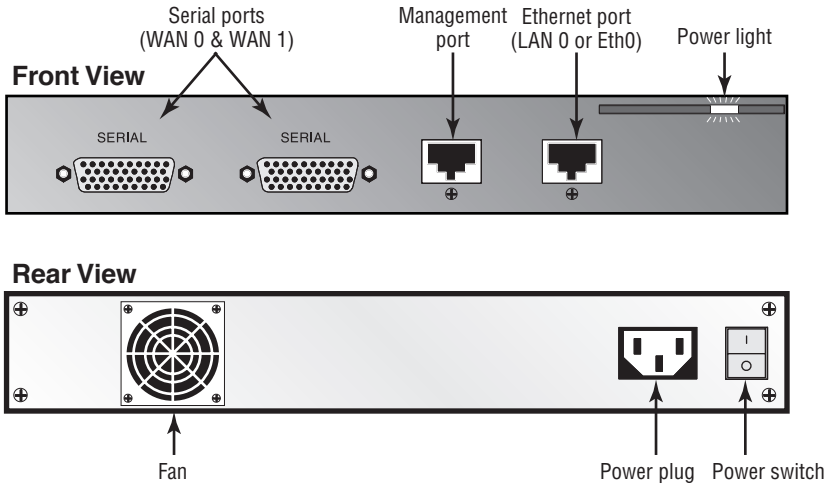

Brouter

A *brouter* is a cross between a bridge and a router. As discussed previously, a bridge segments a network to keep local traffic on the right side of the bridge while a router can connect multiple networks together and make intelligent decisions on where to forward traffic. A brouter combines the best of both worlds, but is not used as often as a switch.

Modems

The device most commonly used to connect computers over a public medium is a *modem* (a contraction of *mo*dulator/*dem*odulator). A modem changes digital signals, which are in the form of ones and zeros, from the computer into analog signals that can be transmitted over phone lines and other analog media. On the receiving end, the modem changes the analog signals back to digital signals. The pattern of these analog signals encodes the data for transmission to the receiving computer. The receiving modem then takes the analog signals and turns them back into ones and zeros. Using this method, which is slower than completely digital transmissions, data can travel over longer distances with fewer errors.

A modem can be either internal or external. The key difference between the two is the amount of configuration required. You must configure internal modems with an IRQ and an I/O address as well as a virtual COM port address to ensure that they function properly. External modems simply connect to a serial port and don't require nearly as much configuration. When deciding which type to purchase, consider how many IRQs and I/O addresses are free on the machine where you are installing the modem. This is a key factor if you are already running a significant number of peripherals on the computer.

The three most common types of modems that are discussed in this chapter are:

- Analog
- Cable
- xDSL

Cable and xDSL modems are not really modems, but rather are adapters. The different between a modem and an adapter is that a modem converts digital signals into analog, and vice versa. Cable and xDSL modems convert digital data from the computer into a digital format that is understood by that technology, and vice versa. Even though the i-Net+ exam lists them as modems, be aware of the difference.

Analog Modems

Analog modems are the most common form of modem in use today and are rated to transmit at 56 Kbps; however, the best transmission rate that you can get with a standard analog modem is actually 52 Kbps due to FCC regulations. Analog modems use the standard public phone line from your home to send and receive data, and are typically a connect-on-demand device (although you could maintain a permanent connection, if you wanted).

Cable Modems

Cable television has become rather popular in the United States because of the number of channels a subscriber can receive and the relatively few reception problems compared to the old "rabbit-ears" method. Cable systems have now evolved to include Internet services directly from your cable box.

A cable modem is designed to convert the digital signals from your computer and translate them into an acceptable format that the cable provider's system can understand.

Cable modems are easy to install and use. You can either lease one from the cable company for a small monthly charge, or you can save some money and buy one at your local computer store. External cable modems have one connection to your cable TV converter (or wall outlet) and one connection to your PC, while internal versions require only the cable connection. Typically, transmission rates can get up to 10Mbps, which is much higher than an analog modem and comparable to an xDSL modem.

Do not expect to receive a 1.5 Mbps data rate very frequently. The data rate is actually dependent upon how many clients in your area have subscribed to the same service and whether or not they are using the service at the same time. Cable Internet services are similar to a hub in that they are *shared* services.

Cable modems are becoming increasingly popular in metropolitan areas for two reasons:

- They are easy to install and require very little waiting time for the cable company to turn on the service. This is because metropolitan areas already have cable lines installed for television, and the phone company basically just has to "flip a switch."
- They are inexpensive. For about the cost of a second (or third) phone line per month, cable offers a high-speed connection to the Internet without interfering with your phone line. (So who needs call waiting?)

xDSL Modems

xDSL modems are based on *digital subscriber line* technology, which is discussed later in the chapter, and is cable's main competition. The *x* in xDSL can stand for several different versions of DSL technology, but xDSL modems all work the same way. All xDSL configurations require a modem, called an *endpoint*, and a NIC. Often, the modem and NIC are on a single expansion card that helps cut down on the number of card's or peripherals that attach to the computer. The modem is then hooked up to the phone line.

CSU/DSU

A *channel service unit / data service unit (CSU/DSU)* is a LAN-to-WAN network device that converts the digital signals from your LAN into the format required by the WAN communications link, and vice versa. You can think of a CSU/DSU as a type of translator between two different technologies. The CSU/DSU typically has one port that is connected to the LAN, and one port that connects to your WAN modem or adapter.

Firewall

Networks that are connected to the Internet are subject to possible attacks from outside malicious entities located elsewhere on the Internet. To protect a network against attacks, a device called a firewall is employed. *Firewalls* reside between a company's LAN and the Internet and monitor all traffic going into and out of the network. Any suspicious or unwanted activity is monitored and, if necessary, quelled. Firewalls are usually combinations of hardware and software with multiple NICs (one for the Internet side, another for the LAN side, and possibly a third for a DMZ, discussed in a moment). Some firewalls are stand-alone hardware devices; others consist of special software that turns the server into a firewall. Both types can be generalized as firewalls. The major difference between the two is that the latter may run a commercially-available NOS, like NT, NetWare, or Unix, whereas the former is running its own highly specialized operating system.

Most firewalls in use today implement a concept called a *demilitarized zone (DMZ)* or *screened subnet*, which is a network segment that is neither public nor local, but halfway between. People outside your network primarily access your public Web servers, FTP servers, and mail-relay servers. Because hackers tend to go after these servers first, place them in the DMZ. A standard DMZ setup has three network cards in the firewall computer. The first goes to the Internet. The second goes to the network segment where the aforementioned servers are located, the DMZ. The third connects to your intranet.

Never put your intranet server into a DMZ. By doing so, you're allowing a hacker access to your corporate information and thereby defeating the purpose of a DMZ.

When hackers break into the DMZ, they can see only public information. If they break into a server, they are breaking into a server that holds only public information. Thus, the entire corporate network is not compromised. In addition, no e-mail messages are vulnerable; only the relay server can be accessed. All actual messages are stored and viewed on e-mail servers inside the network. As you can see in Figure 1.10, the e-mail router, the FTP server, and the Web server are all in the DMZ, and all critical servers are inside the firewall.

FIGURE 1.10 A firewall with a DMZ

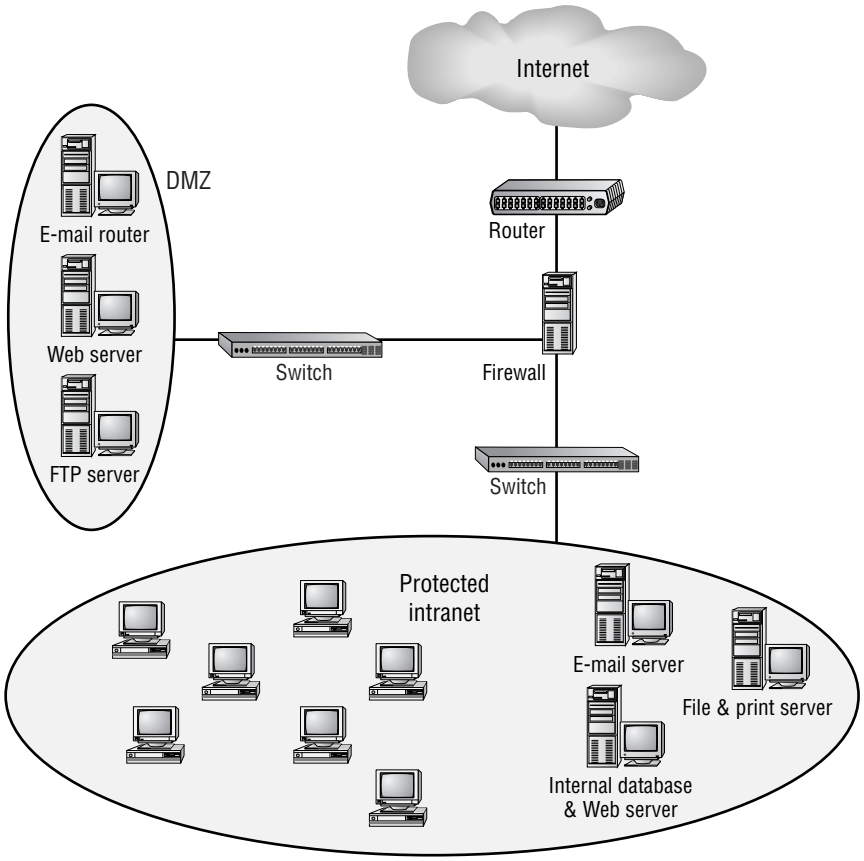

Network Software Components

In addition to all the hardware components, networks use some software components to tie together the functions of the different hardware components. Each software component has a different function on the network. In this section, you will learn about some of the software often found on a network. The most important network software components that you'll learn about include:

- Network operating system (NOS)
- Protocols

Each software component runs on a computer and provides the network with some service.

Network Operating System

Every network today uses some form of software to manage the resources of the network. This software runs on the servers and is called a *network operating system* (or NOS, for short). NOSes are, first and foremost, computer operating systems, which means they manage and control the functions of the computer they are running on. NOSes are more complex than computer operating systems because they manage and control the functions of the network as well. A NOS gives a network its "soul" because each NOS works a bit differently. Different NOSes will need to be administered differently.

The three most popular network operating systems that you will need to know about are:

- Microsoft Windows NT/Windows 2000
- Novell NetWare
- Unix

In the following sections, you will learn background information on each NOS, its current version, its applicability to the Internet/its strength as a NOS for an Internet server, and its system requirements.

Microsoft Windows NT/2000

There has been a buzz in the computer industry as of late about Windows 2000, produced by Microsoft Corporation. Everyone's asking, "Should I be installing it or Windows NT?" With the same graphical interface as other versions of

Windows and simple administration possible from the server console, it is a force to be reckoned with. Microsoft has put its significant marketing muscle behind it, and Windows NT has become a viable alternative in the network operating system market, previously dominated by Novell NetWare and the various flavors of Unix. Currently, Windows 2000 is just beginning to catch up with Windows NT deployments.

Most companies already have a significant investment in Windows NT, and have not migrated to Windows 2000 because of the expense of migrating their current environments. In the business world, upgrades usually take a few years to occur. For more information on Windows NT and 2000 than is discussed in this chapter, visit Microsoft's Web site at `www.microsoft.com`.

Microsoft's Windows NT Server has become the predominant general-purpose server for the industry. Its versatility and familiar graphical user interface (it's the same as Windows 95/98 in NT 4 and 2000) belie its complexity. Using TCP/IP and other protocols, Windows NT can communicate and be integrated with NetWare and Unix servers. Additionally, it is the preferred NOS for the intranet and Internet services of small companies because it's easy to set up and manage for Internet services. Again, this ease comes from the familiarity people have with the client OS, Windows 95/98. Also, Internet services can be installed during NOS setup with a few mouse-clicks and a minimal amount of configuration. The only downside to Windows NT is that it's sometimes unstable and it has much larger hardware requirements than the other NOSes discussed in this chapter (as listed in Table 1.1).

TABLE 1.1 Windows NT Server 4 Hardware Requirements

Hardware	Minimum	Recommended
Processor	Intel 80486 or higher (I386 architecture) or a supported RISC processor (MIPS R4x00, Alpha AXP, or PowerPC)	Pentium 90MHz or higher (the faster the better)
Display	VGA	SVGA
Hard disk space	120MB free	300MB free

TABLE 1.1 Windows NT Server 4 Hardware Requirements *(continued)*

Memory	16MB	32MB or greater
Network card	At least one that matches the topology of your network	At least one that matches the topology of your network
CD-ROM	None	8x or greater
Mouse	Required	Required

Windows 2000 is the latest version of Microsoft's server suite, and is designed to eventually replace both Windows 95/98 (known as *Windows 9x* in the industry) and Windows NT. While Windows NT used a similar architecture as Windows 9x, Windows 2000 only looks similar. Windows 2000 is the biggest release of Windows to date and has the most features, including a new directory service, called *Active Directory Service* and based on X.500 technology, and Plug-and-Play support. The minimum hardware support requirements for Windows 2000 are listed in Table 1.2.

TABLE 1.2 Windows 2000 Server Hardware Requirements

Hardware	Minimum	Recommended
Processor	133MHz or higher Pentium-Compatible CPU	(The faster the better)
Display	VGA	SVGA
Hard disk space	2GB with a minimum of 1GB free space (Additional free hard drive space is required if you are installing over a network.)	40GB or more, depending on the Server's function
Memory	128MB	256MB

TABLE 1.2 Windows 2000 Server Hardware Requirements *(continued)*

Hardware	Minimum	Recommended
Network card	At least one that matches the topology of your network	At least one that matches the topology of your network
CD-ROM	None	8x or greater
Mouse	Required	Required

Note that these are the minimum requirements for Windows 2000 Server, and not Windows 2000 Advanced Server.

Novell Netware

NetWare, made by Novell, Inc., was the first NOS developed specifically for use with PC networks. It was introduced in the late '80s and quickly became the software people chose to run their networks. NetWare is one of the more powerful network operating systems on the market today. It is almost infinitely scalable and has support for multiple client platforms. Although most companies larger than a few hundred stations are running NetWare, this NOS enjoys success in many different types of networks.

Currently, NetWare is at version 5.1 and includes workstation management support, Internet connectivity, Web proxy, and native TCP/IP protocol support, as well as continued support for its award-winning X.500-based directory service, Novell Directory Services (NDS). At the time this book is being written, Novell 6 is scheduled for release at the last quarter of 2001.

As an Internet and intranet NOS, NetWare sees use in large networks for secure intranets. In our tests, with similarly configured servers, NetWare had the best Web server performance over NT and Unix (using the included Netscape Enterprise Server for NetWare). Plus, its web page security is integrated with Novell's directory service (NDS). Hardware requirements are listed in Table 1.3.

For more information on NetWare, check out Novell, Inc.'s web site at www.novell.com.

TABLE 1.3 NetWare 5.1 Hardware Requirements and Recommendations

Hardware	Minimum	Recommended
Processor	Pentium II	None listed
Display	VGA	SVGA
Hard disk space	1.3GB	40GB or more
Memory	128MB for Standard, 384MB to Include WebSphere	256MB for Standard, 768MB to include WebSphere
Network card	At least one	As many as required
CD-ROM	Required	8x or greater
Mouse	Not required	Recommended if using graphical interface. PS/2 style is the best choice

Unix

Of the network operating systems other than Windows NT and NetWare, the various forms of Unix are probably the most popular. It is also among the oldest of the network operating systems. Bell Labs developed Unix, in part, in 1969—*in part* because there are now so many iterations, commonly called *flavors*, of Unix that it is almost a completely different operating system.

Although the basic architecture of all flavors is the same (32-bit kernel, command-line based, capable of having a graphical interface, as in X Windows), the subtle details of each make one flavor better than another in a particular situation.

Unix flavors incorporate a kernel, which constitutes the core of the operating system. The kernel can access hardware and communicate with various types of user interfaces. The two most popular user interfaces are the command-line interface (called a *shell*) and the graphical interface (X Windows). The Unix kernel is similar to the core operating system components of Windows NT and NetWare. In Unix, the kernel is typically simple and, therefore, powerful. Additionally, the kernel can be recompiled to include support for more devices. As a matter of fact, some flavors include the source code so that you can create your own flavor of Unix.

As an Internet platform, Unix has many advantages, mainly because the Internet was first and foremost a Unix-based network. Many services available for the Internet (like Usenet news) work best on the Unix platform because these technologies were first developed on Unix. Additionally, Unix is powerful enough to scale to service hundreds of thousands of web requests per second. Many of the most popular web sites run on Unix.

Each flavor of Unix has widely varied hardware requirements. Some flavors can run on any processor/hardware combination. Others can only run on certain combinations. As an example, hardware requirements for the common PC-based Unix flavor Red Hat Linux 7.1 are covered in Table 1.4. If you need to install any flavor of Unix onto a computer, check the software's packaging or documentation for its respective hardware requirements.

Unix hardware requirements vary from vendor to vendor. As such, they are not currently tested for in the exam.

TABLE 1.4 Red Hat 7.1 Linux Hardware Requirements

Hardware	Minimum	Recommended
Processor	Intel 80486 or higher (I386 architecture), 680x0, or a supported RISC processor (MIPS, AP1000+, Alpha AXP, SPARC, or PowerPC)	Pentium 90MHz or higher (the faster the better)
Display	VGA	SVGA

TABLE 1.4 Red Hat 7.1 Linux Hardware Requirements *(continued)*

Hard disk space	450MB for Workstation, 1620MB for Server	1.5GB free
Memory	16MB	32MB or greater
Network card	None	At least one that matches the topology of your network

Protocols

All network entities must communicate to gain the benefits of being networked. To be able to communicate, each device on the network must understand the same basic rules of that communication. For example, each node must understand a common "language" and the types of "words" to use. Not to imply that computers speak English, but they do need a set of rules to communicate. These rules are called *protocols*. Multiple protocols operating together are called a *protocol suite*. Finally, a software implementation of a protocol is called a *protocol stack*.

There is really only one protocol suite used on the Internet, the Transmission Control Protocol/Internet Protocol (TCP/IP) suite. It was developed at approximately the same time the Internet was developed. When it was being designed, its designers wanted a protocol that could reconfigure itself around possible breaks in the communication channel. Today, TCP/IP is almost ubiquitous because almost every operating system includes a TCP/IP protocol stack so that the operating system can communicate with the Internet. That feature, along with its relatively decent performance, makes TCP/IP a very popular protocol. We'll discuss TCP/IP in more in detail in Chapter 3.

Local Area Network Link Types

Local area networks (LANs) have many ways of delivering data from point A to point B. These "link types" include specifications that dictate how the stations will transmit their data, how the data will travel on the network, and how much data can be transmitted. The majority of networks installed today

Other Protocols

In addition to TCP/IP, there are other protocols available for use on LANs. The protocol suite Internetwork Packet eXchange/Sequenced Packet eXchange (IPX/SPX), developed by Novell for use with NetWare, is probably the second most popular protocol. It is used with both NetWare and Windows NT and is a popular choice because of its ease of configuration. Some other protocol suites you may encounter are the NetBIOS Enhanced User Interface (NetBEUI), DEC Networking (DECNet), and Systems Network Architecture (SNA) protocols, but these see much more limited use in LANs today when compared to TCP/IP and IPX/SPX.

(including the ones at ISPs) use these link types. There are two popular LAN link types you will see on almost every network:

- Ethernet
- Token Ring

Most servers and workstations connect using one of these link types.

Ethernet

Ethernet, the *most* popular network specification, was originally the brainchild of Xerox Corporation. Introduced in 1976, it quickly became the network of choice for small LANs. The Unix market was the first to embrace this easy-to-install network.

Ethernet uses the CSMA/CD (Carrier Sense Multiple Access with Collision Detection) media access method, which means that only one workstation can send data across the network at a time. It functions much like the old party line telephone systems used in rural areas. If you wanted to use the telephone, you picked up the line and listened to see if anyone was already using it. If you heard someone on the line, you didn't try to dial or speak; you simply hung up and waited a while before you picked up the phone to listen again.

If you picked up the phone and heard a dial tone, you knew the line was free. You and your phone system operated by *carrier sense*. You sensed the dial tone or carrier, and if it was present, you used the phone. *Multiple access* means that more than one party shared the line. *Collision detection* means that if two people picked up the phone at the same time and dialed, they would "collide" and both would need to hang up the phone and try again. The first one back on the free line gains control and is able to make a call.

In the case of Ethernet, workstations send signals (frames) across the network. When a collision takes place, the workstations transmitting the frames stop transmitting and wait for a random period of time before retransmitting. Using the rules of this model, the workstations must *contend* for the opportunity to transmit across the network. For this reason, Ethernet is referred to as a *contention-based* system.

Current implementations of Ethernet allow for connection speeds of either 10 or 100Mbps. There are, however, standards being developed for Gigabit Ethernet (one thousand megabits per second).

Token Ring

Token Ring was developed by IBM as a robust, highly reliable network. It is more complex than Ethernet because it has self-healing properties. Token Ring is an IEEE 802.5 standard whose topology is physically a star but logically a ring. Workstations connect to the bus by means of individual cables that connect to a multistation access unit (MSAU) or controlled-access unit (CAU). MSAUs and CAUs are similar to Ethernet hubs in that they exist at the center of the star, but they are for Token Ring networks. The difference between an MSAU and a CAU is that an MSAU is a passive device that has no power plug and no intelligence, whereas a CAU has intelligence and a power plug. A CAU can perform physical network management operations.

The original Token Ring cards were 4Mbps. These were later replaced by 16Mbps cards. The 16Mbps cards are manufactured to work at 4Mbps (for compatibility), but the 4Mbps cards only run at 4Mbps. The 4Mbps version will allow only one token on the ring at a time. The 16Mbps version will allow a card to retransmit a new free token immediately after the last bit of a frame. The term for this is *early token release*.

When configuring a Token Ring network, you must remember that all Token Ring cards must be set to either 4Mbps or 16Mbps. You cannot mix the speeds on the same segment.

In a Token Ring, although the cards attach like a star to the MSAU or CAU, they function logically in a ring. A *free token* (a small frame with a special format) is passed around the ring in one consistent direction. A node receives the token from its *nearest active upstream neighbor (NAUN)* and passes it to its *nearest active downstream neighbor (NADN)*. If a station receives a free token, it knows that it can attach data and send it on down the

ring. This is called *media access*. Each station is given an equal chance to have the token and take control to be able to pass data.

Each station in the ring receives the data from the busy token and repeats the data, exactly as it received it, on to the next active downstream neighbor on the ring. The addressed station (the station the data is intended for) keeps the data and passes it on up to its upper-layer protocols. It then switches 2 bits of the frame before it retransmits the information back on to the ring to indicate that it received the data. The data is sent repeatedly until it reaches the source workstation, and then the process begins again.

Each station in the ring basically acts as a repeater. The data is received and retransmitted by each node on the network until it has gone full circle. This is something like the party game called Rumor or Telephone, in which one person whispers something into one player's ear, who in turn whispers it into someone else's ear, and so on until it has gone full circle. The only difference is that, in the party game, when the person who initiated the message receives it back, it has usually undergone substantial permutations. When the originating node on the network receives the message, it is normally intact except that 2 bits have been flipped to show that the message made it to its intended destination.

Token Ring computers act as repeaters, in contrast to computers in an Ethernet network, where they are passive and therefore not relied on to pass data. This is why Token Ring networks can experience periods of latency when a computer fails and Ethernet networks will not. Also, the token-passing access method will not have collisions because only one token is on the cable at one time; Ethernet networks with CSMA/CD do have collisions.

Internet Bandwidth Link Types

An *Internet bandwidth technology* (or *link)* is the communications pathway between the various LANs that make up the Internet. These links are typically specific types of analog or digital telephone lines that carry data for a corporate WAN and for the Internet. They are leased from the telephone companies that serve the cities at the ends of the link. Hence, these WAN links are often called *leased lines*.

In addition to connecting networks together, the same WAN link technologies are also used to connect entire networks to the Internet and to provide

the Internet with its structure by connecting multiple ISPs together. Wide area network links are commonly grouped into two main types:

- Point-to-point
- Public switched networks

Point-to-Point WAN Connections

Point-to-point WAN connections are WAN links that exist directly between two locations. Point-to-point connections are typically used for WAN connections between a central office and a branch office or from these locations to an ISP for Internet connectivity. These connections come in a variety of connection speeds. The main advantage of point-to-point connections to the Internet is that there is only one "hop" between the two locations, thus much less latency in each transmission, which means more data can be transmitted. The main downside is that these connections are often more expensive than their switched counterparts.

There are seven main point-to-point WAN connections in use today:

- DDS/56Kbps
- T1/E1
- T3/E3
- Asynchronous Transfer Mode (ATM)
- Integrated Services Digital Network (ISDN)
- Digital Subscriber Line (DSL)
- Synchronous Optical Network (SONET)

Each connection type differs primarily in the data throughput rates offered and in the cost. In this section, you will learn about the most popular point-to-point WAN (and Internet) connection types.

DDS/56Kbps

The *Dataphone Digital Service (DDS)* line from AT&T is a dedicated, point-to-point connection with throughput anywhere from 2400bps to 56Kbps. The 56Kbps digital connection is the most common, and this type of line has since obtained the moniker *56K line*. This type of line is used most often for small office connections to the central office. Some small companies may use this for their connection to their ISP for an Internet connection.

If a phone company other than AT&T provides this service, the line is known as a Digital Data Service line. The abbreviation is still DDS, however.

T1/E1

A *T1* is a 1.544Mbps digital connection that is typically carried over two pair of UTP wires. This 1.544Mbps connection is divided into 24 discrete, 64Kbps channels (called DS0 channels). Each channel can carry either voice or data. In the POTS world, T1 lines are used to bundle analog phone conversations over great distances, using much less wiring than would be needed if each pair carried only one call. This splitting into channels allows a company to combine voice and data over one T1 connection. You can also order a fractional T1 channel that uses fewer than the 24 channels of a full T1. An E1 is the same style channel, but it is a European standard and is made up of 32 64Kbps channels for a total throughput of 2.048Mbps.

A T1 connection is used very often to connect a medium-size company (50 to 250 workstations) to the Internet. It is usually cost prohibitive to have a T1/E1 connection for any company smaller than that, and it doesn't have the bandwidth that larger companies would require for high-speed WAN connections. Smaller ISPs that mainly provide residential dial-up connections may only have a T1 connection.

T3/E3

A *T3* line and a T1 connection work similarly, but a T3 line carries a whopping 44.736Mbps. This is equivalent to 28 T1 channels (or a total of 672 DS0 channels). E3 is a similar technology for Europe that uses 480 channels for a total bandwidth of 34.368Mbps. Currently these services require fiber-optic cable or microwave technology. Many local ISPs have T3 connections to the major ISPs, such as SprintNet, AT&T, and MCI. Also, very large, multinational companies use T3 connections to send voice and data between their major regional offices.

Asynchronous Transfer Modem (ATM)

Of the link types we have discussed so far, *Asynchronous Transfer Mode (ATM)* is one link type that is used on both LANs and WANs. ATM uses cell-switching technology, which means that it works by dividing all data to be transmitted into special 53-byte packets called *cells* and sending them over a switched, permanent virtual circuit. Because all the packets are the same length, and because they are very small, ATM is a highly efficient, and

very fast, set of WAN standards. It can support transmissions of voice and video in addition to data at speeds of from 1.5 to 2488Mbps. Additionally, ATM supports the ability to reserve bandwidth to ensure Quality of Service (QoS) so that voice and data transmissions won't interfere with each other. Several Internet backbone ISPs use ATM to move massive amounts of Internet data quickly.

ISDN

ISDN is a digital, point-to-point network capable of maximum transmission speeds of about 1.4Mbps, although speeds of 128Kbps are more common. Because it is capable of much higher data rates, at a fairly low cost, ISDN is becoming a viable remote Internet connection method, especially for those who work out of their homes and require high-speed Internet access but can't afford a T1 or higher. ISDN uses the same UTP wiring as your residential or business telephone wiring (also known as Plain Old Telephone Service, or POTS), but it can transmit data at much higher speeds. That's where the similarity ends, though. What makes ISDN different from a regular POTS line is how it uses the copper wiring. Instead of carrying an analog (voice) signal, it carries digital signals. This is the source of several differences.

A computer connects to an ISDN line via an ISDN terminal adapter (often incorrectly referred to as an ISDN modem). An ISDN terminal adapter is not a modem because it does not convert a digital signal to an analog signal; ISDN signals are digital.

A typical ISDN line has two types of channels. The first type of channel is called a *Bearer,* or *B,* channel, which can carry 64Kbps of data. A typical ISDN line has two B channels. One channel can be used for a voice call while the other is being used for data transmissions, and this occurs on one pair of copper wire. The second type of channel is used for call setup and link management and is known as the *Signal,* or *D,* channel (also referred to as the Delta channel). This third channel has only 16Kbps of bandwidth.

In many cases, to maximize throughput, the two Bearer channels are combined into one data connection for a total bandwidth of 128Kbps. This is known as *bonding* or *inverse multiplexing*. This still leaves the Delta channel free for signaling purposes. In rare cases, you may see user data such as e-mail on the D line. This was introduced as an additional feature of ISDN, but it hasn't caught on.

ISDN has three main advantages:

- Fast connection.
- Higher bandwidth than POTS. Bonding yields 128Kbps bandwidth.
- No conversion from digital to analog.

ISDN does have a few disadvantages:

- It's more expensive than POTS.
- Specialized equipment is required at the phone company and at the remote computer.
- Not all ISDN equipment can connect to each other.

xDSL

Digital Subscriber Line (DSL) is a hot topic for home Internet access because it is relatively cheap (less than $100/month in most areas), fast (greater than 128Kbps), and available in most major cities in the United States. xDSL is a general category of copper access technologies that is becoming popular because it uses regular, POTS phone wires to transmit digital signals and is extremely inexpensive compared with the other digital communications methods. xDSL implementations cost hundreds instead of the thousands of dollars that you would pay for a dedicated, digital point-to-point link (such as a T1). They include Digital Subscriber Line (DSL), High Data Rate Digital Subscriber Line (HDSL), Single Line Digital Subscriber Line (SDSL), Very High Data Rate Digital Subscriber Line (VDSL), and Asymmetric Digital Subscriber Line (ADSL), which is currently the most popular. It is beyond the scope of this book to cover all the DSL types. Ask your local telephone company which method they provide.

ADSL is winning the race because it focuses on providing reasonably fast upstream transmission speeds (up to 640Kbps) and very fast downstream transmission speeds (up to 9Mbps). This makes downloading graphics, audio, video, or data files from any remote computer very fast. The majority of web traffic, for example, is downstream. The best part is that ADSL works on a single phone line without losing the ability to use it for voice calls. This is accomplished with what is called a *splitter*, which enables the use of multiple frequencies on the POTS line. xDSL modems were discussed earlier in this chapter.

Sonet

Some of the fastest WAN connections are those employed in the Synchronous Optical Network (SONET). *SONET* is a high-speed, fiber-optic system that provides a standard method for transmitting digital signals over a fiber-optic network. Multiple transmission types (such as 64Kbs channels, T1/E1 channels) can be multiplexed together to provide SONET speeds.

SONET is able to achieve maximum transmission speeds of up to 2.488 gigabits per second. It does so by using a fixed frame size of 810 bytes. This fixed frame size makes transmissions very efficient, and thus they can carry more data.

SONET speeds are rated as channels. They are designated with an OC (Optical Carrier) number. The OC lines are designated OC-1 through OC-768. OC-1 channels communicate at 51.84Mbps, OC-3 channels communicate at 155.52Mbps, and OC-768 channels communicate at 40Gbps.

Public Switched Network WAN Connections

The other type of WAN link most commonly in use is the public switched network WAN connection. These connections use the telephone company's analog switched network to carry digital transmissions. Your network traffic is combined with other network traffic from other companies. Essentially, you are sharing the bandwidth with all other companies. The upside to this type of WAN connection is that it is cheaper than point-to-point connections, but because you share the bandwidth with other traffic, it isn't necessarily as efficient.

Let's take a brief look at some of the public switched network connections that companies use to connect to the Internet, including:

- Public Switched Telephone Network (PSTN)
- X.25
- Frame Relay

Public Switched Telephone Network (PSTN)

Almost everyone outside the phone companies themselves refers to *PSTN (Public Switched Telephone Network)* as POTS (Plain Old Telephone Service). This is the wiring system that runs from most people's houses to the rest of the world. It is the most popular method for connecting to the Internet because of its low cost, ease of installation, and simplicity. The majority of the houses in the U.S. that have Internet connections connect to their ISP via PSTN and a modem.

The phone company runs a UTP (unshielded twisted-pair) cable (called the *local loop*) from your location (called the *demarcation* point or *demarc*, for short) to a phone company building called the *Central Office*. All the pairs from all the local loop cables that are distributed throughout a small regional area come together at a central point, similar to a patch panel in a UTP-based LAN.

This centralized point has a piece of equipment called a *switch* attached. The switch functions almost exactly like the switches we mentioned earlier, in that a communications session, once initiated when the phone number of the receiver is dialed, exists until the conversation is closed. The switch can then close the connection. On one side of the switch is the neighborhood wiring. On the other side are lines that may connect to another switch or to a local set of wiring. The number of lines on the other side of the switch depends on the usage of that particular exchange. Figure 1.11 shows a PSTN system that utilizes these components.

FIGURE 1.11 A local PSTN (POTS) network

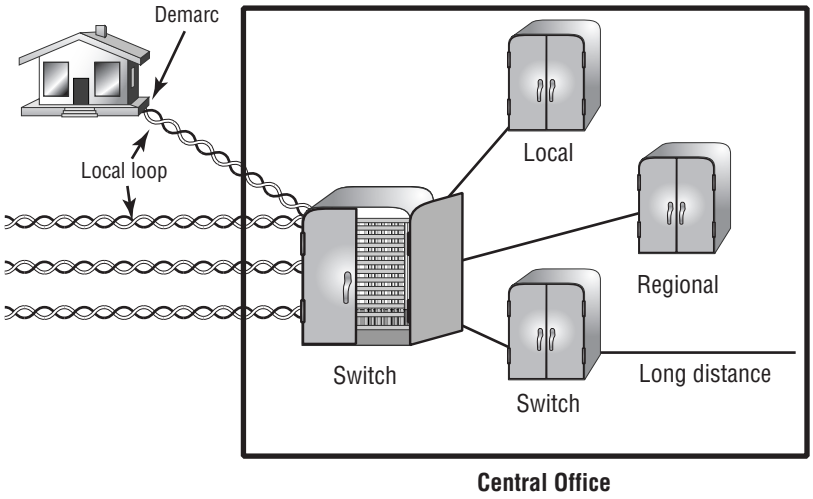

Use caution when working with bare phone wires because they may carry a current. In POTS, the phone company uses a battery to supply power to the line, which is sometimes referred to as *self-powered*. It isn't truly self-powered; the power comes from the phone system.

POTS has many advantages, including:

- It is inexpensive to set up. Almost every home in the United States has or can have a telephone connection.
- There are no LAN cabling costs.
- Connections are available in many countries throughout the world.

POTS is the most popular remote access connection method for the Internet because only two disadvantages are associated with it: limited bandwidth and thus a limited maximum data transfer rate. At most, 64Kbps data transmissions are possible, though rarely achieved by anyone connecting from home to the Internet.

X.25

X.25 was developed by the International Telecommunications Union (ITU) in 1974 as a standard interface for WAN packet switching. It does *not* specify anything about the actual data transmission, however. It only makes specifications about the access to the WAN and just assumes that a route from sender to receiver exists. The original X.25 specification supported transmission speeds of up to 64Kbps, but the 1992 revision supports transmission speeds of up to 2Mbps. It is currently one of the most widely used WAN interfaces.

Frame Relay

Similar to X.25, *Frame Relay* is a WAN technology in which packets are transmitted by switching. Packet switching involves breaking messages into chunks at the sending router. Each packet can be sent over any number of routes on its way to its destination. The packets are then reassembled in the correct order at the receiver. Because the exact path is unknown, a cloud is used when creating a diagram to illustrate how data travels throughout the service. Figure 1.12 shows a Frame Relay WAN connecting smaller LANs.

FIGURE 1.12 A typical frame relay configuration

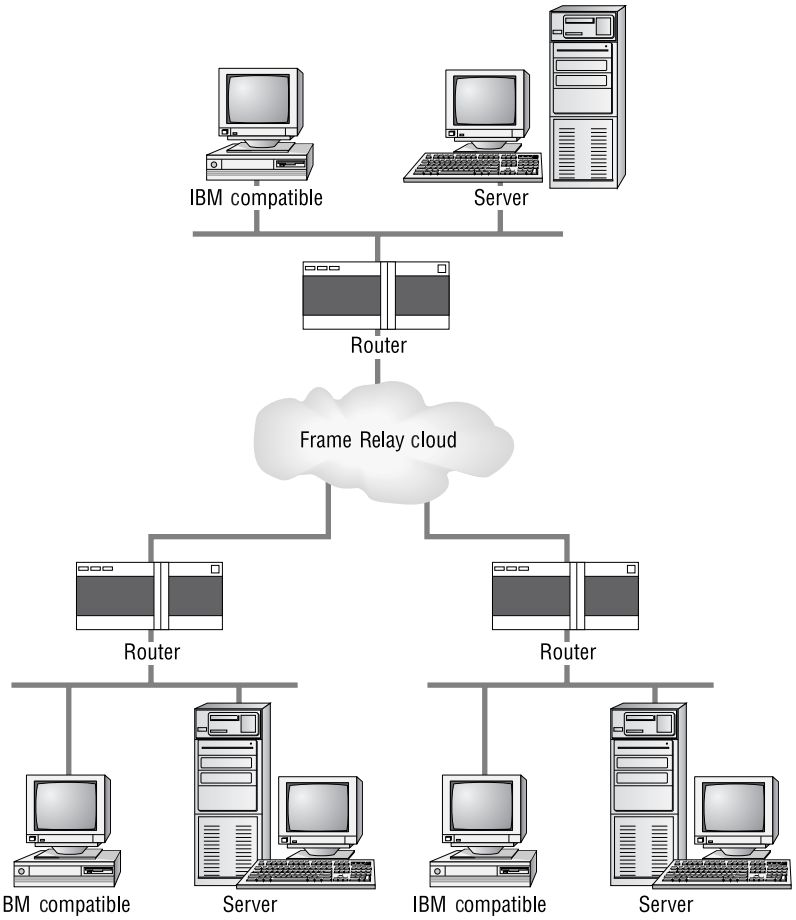

Frame Relay uses permanent virtual circuits (PVCs). PVCs allow virtual data communications circuits between sender and receiver over a packet-switched network. This ensures that all data that enters a Frame Relay cloud at one side comes out at the other over a similar connection.

The beauty of using a shared network is that sometimes you can get much better throughput than you are paying for. When signing up for one of these connections, you specify and pay for a Committed Information Rate (CIR), or in other words, a minimum bandwidth. If the total traffic on the shared network is light, you may get much faster throughput without paying for it.

Frame Relay begins at the CIR speed and can reach as much as 1.544Mbps, the equivalent of a T1 line, which was discussed earlier.

However, the major downside to Frame Relay is that you share traffic with all other people within the Frame Relay cloud. If you aren't paying for a CIR, your performance can vary widely. Despite this disadvantage, Frame Relay is a popular Internet connection method because of its low cost.

Table 1.5 shows all these point-to-point connections and their respective performance, availability, and cost.

TABLE 1.5 Point-to-point WAN and Internet connection types

Connection	Max Throughput	U.S. Availability	Relative Cost
56K/DDS	56Kbps	Widely available	Low
T1/E1	1.544Mbps/ 2.048Mbps	Widely available	Medium
T3/E3	44.736Mbps/ 34.368Mbps	Widely available	High
ATM	2488Mbps	Moderately available	Very high
ISDN	Around 2Mbps	Moderately available	Low
DSL	Greater than 128Kbps	Available in larger cities, becoming more available in rural areas	Low
Frame Relay	1.544Mbps or slower	Widely available	Low
OC-1	51.84Mbps	Moderately available	Very, very high
OC-3	155.52Mbps	Moderately available	Very, very high
OC-48	2.488Gbps	Slim availability	Don't ask

There are many software applications available that help you draw your network, as well as some software that attempts to map the network for you.

Summary

In this chapter, you learned about some of some of the LAN and WAN networking technologies that apply to the business of the Internet and how networks are physically described by its topology. Because most networks are connected to the Internet these days, the concepts contained in this chapter will be valuable to you as an Internet professional. You learned the definitions of a LAN, MAN and WAN as well as the differences between them. You also learned about some of the hardware components that exist on the network, including workstations, servers, NICs, network cables, repeaters, hubs, switches, bridges, routers, brouters, firewalls, and modems. In addition to learning about the hardware components, you learned about some of the software components that work on the network to provide Internet (and other) services, including network operating systems (NOSes) and protocols.

This chapter included a discussion about the link types that carry data from point A to point B on a network. LAN link types include Ethernet and Token Ring. Ethernet is the most common LAN link type. WAN link types include DDS/56Kbps, T1/E1, T3/E3, ATM, ISDN, DSL, and Frame Relay. The WAN link types can be used for connecting to the Internet and vary in speed and link cost.

Exam Essentials

Understand network topologies. Identify the four basic network topologies and be able to identify them graphically.

Know the different network types. Understand the difference between a LAN, MAN, and WAN.

Know how to identify a network device and its function. Understand the difference between a network interface card, server, repeater, bridge, hub, switch, router, and brouter.

Know the different server types and their function. Understand that the name of the server explains the servers' function.

Understand the different WAN connections. Identify each WAN connection type and the speeds that they can achieve.

Key Terms

Before you take the exam, be certain you are familiar with the following terms:

56K line	demarcation	modem	router
Asynchronous Transfer Mode (ATM)	demilitarized zone	Multiple access	self-powered
active hub	Digital Subscriber Line (DSL)	Multiplexing	shell
Bearer	early token release	nearest active downstream neighbor (NADN)	shielded twisted-pair (STP)
Binding	endpoint	nearest active upstream neighbor (NAUN)	Signal
bonding	Firewalls	network	SONET
bridge	Frame Relay	network interface card	splitter
brouter	free token	network operating system	star topology

bus topology	ground loop	out-of-band management port	star-bus topology
carrier sense	hub	passive hub	star-ring topology
Cat	Internet bandwidth technology	physical topology	switch
Category	inverse multiplexing	Plain Old Telephone Service	switch
Central Office	ISDN	Point-to-point WAN connections	switching hub
channel service unit / data service unit	leased lines	ports	T1
clients	local area network (LAN)	POTS	T3
Collision detection	local loop	protocol stack	terminator
concentrator	logical topology	protocol suite	topology
contention-based	media access	protocols	unshielded twisted-pair (UTP)
crosstalk	mesh topology	PSTN (Public Switched Telephone Network)	wide area network (WAN)
Dataphone Digital Service (DDS)	metropolitan area network	resource	X.25
demarc	mixed topology	ring topology	

Review Questions

1. Which network hardware device connects dissimilar network topologies into an internetwork?

   A. Hub

   B. Bridge

   C. Switch

   D. Router

2. Which Internet bandwidth technology is the primary technology used on the Internet backbone?

   A. Token Ring

   B. Ethernet

   C. X.25

   D. ATM

3. Which of the following is a standard interface for Frame Relay?

   A. X.25

   B. ISDN

   C. T1

   D. xDSL

4. Which network hardware device is required for the computer to be able to connect it to a network?

   A. Bridge

   B. NIC

   C. Router

   D. Firewall

5. Which network hardware device protects a LAN against malicious attacks from the Internet?

A. Bridge

B. Switch

C. NIC

D. Firewall

6. Which of the following is the fastest possible Internet communications technology?

A. Ethernet

B. ATM

C. T1

D. T3

7. A T3 connection has a maximum bandwidth of _____ Mbps?

A. 1.544

B. 2.048

C. 34.368

D. 44.736

8. An E1 connection has a maximum bandwidth of _____ Mbps?

A. 1.544

B. 2.048

C. 34.368

D. 44.736

9. Of the following, which Internet connection type for home users is taking off and offers fairly high speed (>128Kbps) for a fairly reasonable price?

   A. DSL

   B. ISDN

   C. Frame Relay

   D. ATM

10. You are running a Token Ring network with five clients and one server on the same floor of an office building. What topology are you configured for?

    A. Bus

    B. Star

    C. Ring

    D. Mesh

11. The _____ server provides address translation services to network clients accessing the Internet from a LAN.

    A. Windows 2000 Server

    B. Firewall

    C. Network Address Translation Server (NAT)

    D. Remote Access Server (RAS)

12. Which network hardware device is used to segment a single network into multiple segments?

    A. Hub

    B. Firewall

    C. NIC

    D. Bridge

13. Which component of the network is responsible for providing network services to the rest of the network?

    A. Server

    B. Bridge

    C. Workstation

    D. NIC

14. A WAN link is depicted in a logic diagram by a _____.

    A. Straight line

    B. Dashed line

    C. Zig-zagged line

    D. Straight line with an arrow

15. Which network hardware device will increase your web browsing performance?

    A. Firewall

    B. Cache

    C. Bridge

    D. Router

16. Which NOS is the oldest NOS currently in use?

    A. Unix

    B. NetWare

    C. Windows NT

    D. OS/2

17. A _____ is used in firewalls to provide a safe area for public data that is not part of the public or private networks.

A. Firewall

B. Internet-in-a-box

C. DMZ

D. Router

18. Based on speed and cost, which Internet bandwidth link type would be the best choice for a small ISP serving 100 dial-up users?

A. 56K/DDS

B. T1

C. T3

D. ATM

19. Which NOS was developed, in part, by Bell Labs and currently has several hundred different "flavors?"

A. OS/2

B. Windows NT

C. NetWare

D. Unix

20. Which device is also known as a concentrator?

A. Router

B. Switch

C. Hub

D. Brouter

Answers to Review Questions

1. D. Hubs, bridges, and switches connect only the same network topologies. Routers are the only devices that connect different topologies (such as Ethernet to Token Ring).

2. D. Although Token Ring and Ethernet are found in ISPs, ATM is the primary WAN technology used on the Internet backbone. X.25 is only a WAN access technology.

3. A. ISDN, T1, and xDSL are all Internet bandwidth technologies; X.25 is the interface for Frame Relay.

4. B. The other devices (bridge, router, and firewall) are all different network connectivity devices, but you absolutely must have a NIC installed in a computer to be able to connect the computer to a network.

5. D. Bridges, switches, and routers are all simply network connectivity devices. Some routers can perform packet filtering, but firewalls are designed specifically to protect a network against malicious activity from the Internet.

6. B. ATM has maximum speeds of 2488Mbps. Ethernet has a maximum transmission speed of 100Mbps. T1 lines are 1.544Mbps, and T3s are 44.736Mbps.

7. D. A T1 connection has a maximum transmission speed of 1.544Mbps. The 2.048 is E1 speed, 34.368 is E3 speed, and 44.736 is T3 speed.

8. B. An E1 connection communicates at 2.048Mbps. T1 connections are 1.544Mbps, E3 connections are 34.368Mbps, and T3 connections are 44.736

9. A. Frame Relay and ATM normally aren't for home users (unless you happen to be Bill Gates J). ISDN is more expensive than DSL and offers more bandwidth.

10. C. Token Ring typically uses a ring or a star-ring topology. With only six networking devices, you would not need a star-ring.

11. C. A remote access server is used for remote clients needing access to your corporate network, but does not provide translation services. While Windows 2000 Server does have an NAT component, it is a network operating system.

12. D. A bridge is the only device of those listed that is used to segment a network. Hubs and NICs are only connection devices and don't divide a network. Firewalls perform security checks on network traffic but don't do any segmenting.

13. A. A server provides network services to the rest of the network. Bridges, workstations, and NICs do not. Bridges segment a network, workstations request the resources a server provides, and NICs allow a workstation to get access to a network.

14. C. The straight line with an arrow may sound correct, but a zig-zag shows a WAN link that is not under company control. WAN links are leased from vendors, and therefore beyond the direct control of the company.

15. B. Of the devices listed, a cache is the only device that can increase a network's web browsing performance. All the others can actually introduce delay into Internet communications.

16. A. Although NetWare, NT, and OS/2 have been in use for some time, Unix is, in fact, the oldest.

17. C. The demilitarized zone (DMZ) is the network segment connected to a firewall where public data is placed so that it is available to both public and private networks. A, B, and D are incorrect because they are all examples of other Internet hardware and software technologies.

18. B. Because the maximum connection speed of today's modems is 56Kbps and the ISP is serving a maximum of 100 users, the maximum throughput needed is 100 × 56, or 5600Kbps (5.6Mbps). A 56K/DDS link would be too slow and a T3 or ATM connection would be way too fast (and probably way too expensive for a small ISP). A T1 (at 1.544Mbps) would be slower than the throughput number figured above, but it is extremely unlikely that all 100 users would be on at the same time. Plus, you can buy multiple T1s for the cost of a single T3.

19. D. The only one of these listed that was developed in any part by Bell Labs is Unix. The others were all developed by other companies, such as IBM (OS/2), Novell (NetWare), and Microsoft (NT).

20. C. A hub serves as a central connection point for several network devices, and so it's also known as a concentrator. Switches, which are also used as central connection points, were developed afterwards and are also known as switching hubs.