CHAPTER 3

Networking Components and Devices

Objectives

This chapter covers the following CompTIA-specified objectives for the "Media and Topologies" and "Protocols and Standards" sections of the Network+ exam:

1.6 Identify the purpose, features, and functions of the following network components:

- **Hubs**
- **Switches**
- **Bridges**
- **Routers**
- **Gateways**
- **CSU/DSU**
- **Network interface cards (NICs)**
- **ISDN adapters**
- **Wireless access points (WAPs)**
- **Modems**
- **Transceivers (media converters)**
- **Firewalls**
- A wide range of devices is used in modern networking. As a Network+ certified technician, you need to have a good understanding of commonly used devices.

2.1 Identify a MAC (Media Access Control) address and its parts.

- MAC addresses are the means by which systems communicate at a base level. As a network administrator, you need to understand the purpose, function, and expression of MAC addresses.

Outline

Study Strategies

- Review the purpose, function, and key characteristics of the various networking components.
- Review the component summary provided in Table 3.3.
- Review the types of routing protocols, link state and distance vector.
- Distinguish between RIP and OSPF as routing protocols.
- Identify the protocols used within TCP/IP, IPX/SPX, and AppleTalk that provide routing functionality.
- Practice identifying the MAC address of a network card using the appropriate utility.
- Remember to review the Notes, Tips, Tables, and Exam Alerts in this chapter. Make sure that you understand the information in the Exam Alerts. If you don't understand the topic referenced in an Exam Alert, refer to the information in the chapter text and then read the Exam Alert again.

Introduction

So far this book has examined topologies, media access methods, networking standards, and cable types and connectors. To complete our examination of networking on a physical level, this chapter looks at the network devices used to create networks.

Objective:

1.6 Identify the purpose, features, and functions of the following network components:

- Hubs
- Switches
- Bridges
- Routers
- Gateways
- CSU/DSU
- Wireless access points (WAPs)
- Modems
- Network interface cards (NICs)
- ISDN adapters
- Transceivers
- Firewalls

Each of these devices fulfills a specific role in a network; however, only the largest and most complex environments use all of them. We'll begin our discussion of networking devices with perhaps the most simple and common network device used today: the hub.

NOTE

Repeaters Traditionally, any discussion of networking components would include repeaters, but today repeaters are a little outdated. Repeaters were once used to increase the usable length of the cable, and they were most commonly associated with coaxial network configurations. Because coaxial networks have now fallen out of favor, and because the functionality of repeaters has been built in to other devices, such as hubs and switches, repeaters are rarely used. For this reason, CompTIA has elected to leave them out of the required knowledge for the Network+ exam.

Hubs

Hubs are simple network devices, and their simplicity is reflected in their low cost. Small hubs with four or five ports (often referred to as *workgroup hubs*) cost less than $50; with the requisite cables, they provide everything needed to create a small network. Hubs with more ports are available for networks that require greater capacity. Figure 3.1 shows an example of a workgroup hub, and Figure 3.2 shows an example of the type of hub you might see on a corporate network.

FIGURE 3.1 A workgroup hub.

FIGURE 3.2 A high-capacity, or high-density, hub.

Computers connect to a hub via a length of twisted-pair cabling. In addition to ports for connecting computers, even an inexpensive hub generally has a port designated as an uplink port that enables the hub to be connected to another hub to create larger networks. The "Working with Hubs and Switches" section later in this chapter presents a detailed discussion of this feature.

NOTE

Token Ring and MSAUs Both hubs and switches are used in Ethernet networks. Token Ring networks, which are few and far between, use special devices called *multistation access units (MSAUs)* to create the network. In some cases, MSAUs are referred to as *Token Ring switches*; but because of the way Token Ring operates, these devices perform a different function from the hubs and switches discussed in this section.

Most hubs are referred to as either active or passive. *Active* regenerate a signal before forwarding it to all the ports on the device and requires a power supply. Small workgroup hubs normally use an external power adapter, but on larger units the power supply is built in. *Passive* hubs, which today are seen only on older networks, do not need power and they don't regenerate the data signal.

Regeneration of the signal aside, the basic function of a hub is to take data from one of the connected devices and forward it to all the other ports on the hub. This method of operation is inefficient because, in most cases, the data is intended for only one of the connected devices. You can see a representation of how a hub works in Figure 3.3.

NOTE

Broadcasting The method of sending data to all systems regardless of the intended recipient is referred to as *broadcasting*. On busy networks, broadcast communications can have a significant impact on overall network performance.

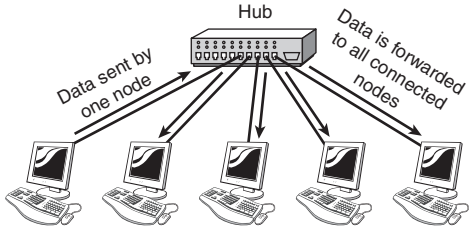

FIGURE 3.3 How a hub works.

Due to the inefficiencies of the hub system and the constantly increasing demand for more bandwidth, hubs are slowly but surely being replaced with switches. As you will see in the next section, switches offer distinct advantages over hubs.

Switches

On the surface, a *switch* looks much like a hub. Despite their similar appearance, switches are far more efficient than hubs and are far more desirable for today's network environments. Figure 3.4 shows an example of a 32-port Ethernet switch. If you refer to Figure 3.2, you'll notice few differences in the appearance of the high-density hub and this switch.

FIGURE 3.4 A 32-port Ethernet switch. (Photo courtesy TRENDware International, www.trendware.com.)

As with a hub, computers connect to a switch via a length of twisted-pair cable. Multiple switches are often interconnected to create larger networks. Despite their similarity in appearance and their identical physical connections to computers, switches offer significant operational advantages over hubs.

As discussed earlier in the chapter, a hub forwards data to all ports, regardless of whether the data is intended for the system connected to the port. This arrangement is inefficient; however, it requires little intelligence on the part of the hub, which is why hubs are inexpensive.

Rather than forwarding data to all the connected ports, a switch forwards data only to the port on which the destination system is connected. It looks at the Media Access Control (MAC) addresses of the devices connected to it to determine the correct port. A *MAC address* is a unique number that is stamped into every NIC. By forwarding data only to the system to which the data is addressed, the switch decreases the amount of traffic on each network link dramatically. In effect, the switch literally channels (or *switches*, if you prefer) data between the ports. Figure 3.5 illustrates how a switch works.

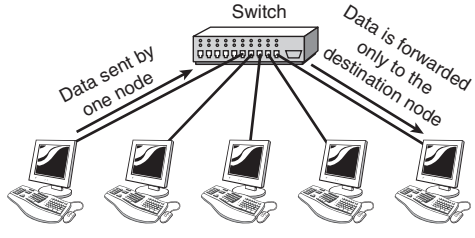

FIGURE 3.5 How a switch works.

You might recall from the discussions of Ethernet networking in Chapter 2, "Cabling Standards, Media, and Connectors," that collisions occur on the network when two devices

attempt to transmit at the same time. Such collisions cause the performance of the network to degrade. By channeling data only to the connections that should receive it, switches reduce the number of collisions that occur on the network. As a result, switches provide significant performance improvements over hubs.

Switches can also further improve performance over the performance of hubs by using a mechanism called *full-duplex*. On a standard network connection, the communication between the system and the switch or hub is said to be *half-duplex*. In a half-duplex connection, data can be either sent or received on the wire but not at the same time. Because switches manage the data flow on the connection, a switch can operate in full-duplex mode—it can send and receive data on the connection at the same time. In a full-duplex connection, the maximum data throughput is double that for a half-duplex connection—for example, 10Mbps becomes 20Mbps, and 100Mbps becomes 200Mbps. As you can imagine, the difference in performance between a 100Mbps network connection and a 200Mbps connection is considerable.

EXAM ALERT

Half-Duplex It's important to remember that a full-duplex connection has a maximum data rate of double the standard speed, and a half-duplex connection *is* the standard speed. The term *half-duplex* can sometimes lead people to believe that the connection speed is half of the standard, which is not the case. To remember this, think of the half-duplex figure as half the full-duplex figure, not half the standard figure.

The secret of full-duplex lies in the switch. As discussed previously in this section, switches can isolate each port and effectively create a single segment for each port on the switch. Because only two devices are on each segment (the system and the switch), and because the switch is calling the shots, there are no collisions. No collisions means no need to detect collisions—thus, a collision-detection system is not needed with switches. The switch drops the conventional carrier-sense multiple-access with collision detection (CSMA/CD) media access method and adopts a far more selfish (and therefore efficient) communication method.

NOTE

Microsegmentation The process that switches perform is referred to as *microsegmentation.*

To use a full-duplex connection, you basically need three things: a switch, the appropriate cable, and a NIC (and driver) that supports full-duplex communication. Given these requirements, and the fact that most modern NICs are full-duplex-ready, you might think everyone would be using full-duplex connections. However, the reality is a little different. In some cases, the NIC is simply not configured to use the driver.

TIP

Troubleshooting Network Connection Speed Most NICs can automatically detect the speed of the network connection they are connected to. However, although the detection process is normally reliable, on some occasions it may not work correctly. If you are troubleshooting a network connection and the autodetect feature is turned on, try setting the speed manually (preferably to a low speed) and then give it another go. If you are using a managed switch you might have to do the same thing at the switch end of the connection.

All Switches Are Not Created Equal

Having learned the advantages of using a switch and looked at the speeds associated with the network connections on the switch, you could assume that one switch is just as good as another. This is not the case. Switches are rated by the number of packets per second (pps) they can handle. When you're buying network switches, it may be necessary to look at the pps figures before making a decision.

Switching Methods

Switches use three methods to deal with data as it arrives:

- **Cut-through**—In a cut-through configuration, the switch begins to forward the packet as soon as it is received. No error checking is performed on the packet, so the packet is moved through quickly. The downside of cut-through is that because the integrity of the packet is not checked, the switch can propagate errors.
- **Store-and-forward**—In a store-and-forward configuration, the switch waits to receive the entire packet before beginning to forward it. It also performs basic error checking.
- **Fragment-free**—Building on the speed advantages of cut-through switching, fragment-free switching works by reading only the part of the packet that enables it to identify fragments of a transmission.

As you might expect, the store-and-forward process takes longer than the cut-through method, but it is more reliable. In addition, the delay caused by store-and-forward switching increases with the packet size. The delay caused by cut-through switching is always the same—only the address portion of the packet is read, and this is always the same size, regardless of the size of the data packet. The difference in delay between the two protocols is high. On average, cut-through switching is 30 times faster than store-and-forward switching.

It might seem that cut-through switching is the obvious choice, but today's switches are fast enough to be able to use store-and-forward switching and still deliver high performance levels. On some managed switches, you can select the switching method you want to use.

NOTE

Latency The time it takes for data to travel between two locations is known as the *latency*. The higher the latency, the bigger the delay in sending the data.

Working with Hubs and Switches

Despite the advantages of switches over hubs, hubs are still widely used in older networks. Whether working with hubs or switches, it is important to be aware of some of their characteristics to troubleshoot a network. For instance, if performance-monitoring tools show network bottlenecks or a congested network, the hubs may need to be replaced with switches for increased performance. This is especially important when working with both hubs and switches in a production environment.

NOTE

Production Environments The term *production* is used to describe a working, or live, computing environment.

Hub and Switch Ports

Hubs and switches have two types of ports: medium dependent interface (MDI) and medium dependent interface crossed (MDI-X). The two types of ports differ in their wiring. As the *X* implies, an MDI-X port's wiring is crossed; this is because the transmit wire from the connected device must be wired to the receive line on the other. Rather than use a crossover cable (which is discussed in the next section, "Cables Connecting Hubs and Switches"), you can use the more simple straight-through cable (also discussed in the next section) to connect systems to the switch or hub.

On most modern hubs and switches, a special port called the *uplink port* allows you to connect two hubs and switches to create larger networks. Because the aim of this type of network connection is to make each hub or switch think that it is simply part of a larger network, the connection for the port is not crossed; a straight-through network cable is used to connect the two hubs or switches together. Figure 3.6 shows the uplink port on an Ethernet switch.

In the absence of an uplink port, you can connect two hubs or switches together by using MDI-X ports, but you must use a crossover cable to do so.

NOTE

Hub Ports Instead of having a dedicated uplink port, some switches and hubs have a port that you can change between MDI and MDI-X by pushing a button. If you are using the port to connect a computer, make sure that it is set to MDI-X. If you're connecting to another hub or switch, make sure that it's set to MDI.

FIGURE 3.6 The uplink port on an Ethernet switch.

Cables Connecting Hubs and Switches

Two types of cables are used to connect devices to hubs and switches: crossover cables and straight-through cables. The difference between the two types is that in a crossover cable, two of the wires are crossed; in a straight-through cable, all the wires run straight through.

Specifically, in a crossover cable, Wires 1 and 3 and Wires 2 and 6 are crossed: Wire 1 at one end becomes Wire 3 at the other end, Wire 2 at one end becomes Wire 6 at the other end, and vice versa in both cases. You can see the differences between the two cables in Figures 3.7 and 3.8. Figure 3.7 shows the pinouts for a straight-through cable, and Figure 3.8 shows the pinouts for a crossover cable.

How Many Is Too Many?

Although Ethernet standards state that you can have as many as 1,024 nodes on a network, the practical maximum may be much lower. The number of nodes you can accommodate depends on a number of factors. Using switches instead of hubs makes a *huge* difference, particularly if you are using the full-duplex features of these devices. The amount of traffic generated by clients also has a significant effect, as does the type of traffic. On a more subtle level, you must consider the quality of the networking components and devices you use.

NOTE

Switches—Read the Label Switches are often labeled as being 10/100 switches. This label normally means that the ports on the switch are capable of operating at 10Mbps or 100Mbps. Don't take it for granted, though. Some older switches have 10Mbps ports for connecting systems and 100Mbps ports for uplinking. Because there are no guidelines for labeling devices, some of those older switches are referred to as 10/100 switches. Always check the specifications before buying a switch.

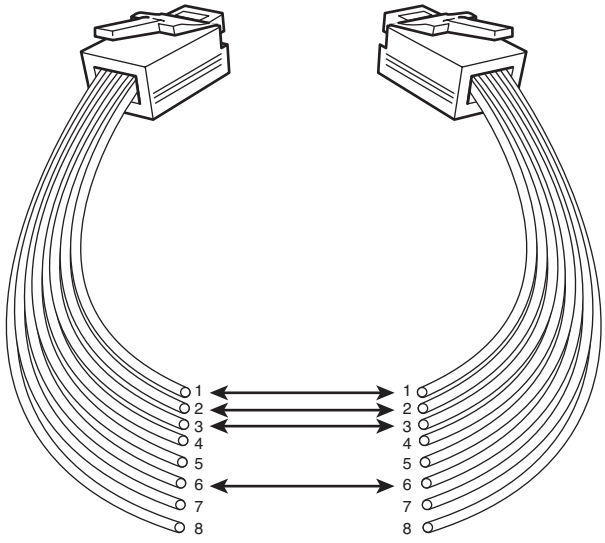

FIGURE 3.7 Pinouts for a straight-through twisted-pair cable.

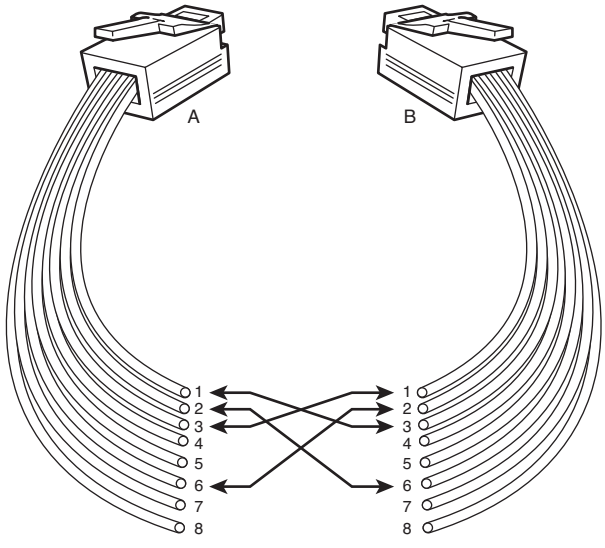

FIGURE 3.8 Pinouts for a crossover twisted-pair cable.

Hubs and switches are sometimes equipped with a network connection for another cable type, such as coaxial. Such switches that accommodate different media types such as fiber-optic cable and UTP, are referred to as *hybrid switches*. Other higher-end devices simply have empty sockets into which you can plug connectivity modules of choice. This approach lets you create very fast networks. For example, three 24-port 10/100 Ethernet switches could be connected to each other by a Gigabit Ethernet fiber-optic connection. This would create a very fast network structure in which switch-to-system communication can occur at 200Mbps (in

full-duplex mode) and switch-to-switch communication can occur at Gigabit Ethernet speeds. The result is a very fast local area network (LAN).

EXAM ALERT

Switch Options Some switches are available with ports that support different media types. These switches are referred to as hybrid switches.

Hub and Switch Indicator Lights

Both hubs and switches use light-emitting diodes (LEDs) to indicate certain connection conditions. At the very least, a link light on the hub will indicate the existence of a live connection. On higher-end devices, additional lights might indicate activity, the speed of the connection, whether the connection is at half- or full-duplex, and sometimes errors or collisions. The LEDs provide an immediate visual indicator about the status of the device, so familiarizing yourself with their function is a worthwhile exercise. A further discussion of hub and switch LEDs is provided in Chapter 13, "Troubleshooting Tools and Utilities."

Rack-Mount, Stackable, and Freestanding Devices

Some hubs and switches, as well as many other networking devices, are designed to be placed in a rack, whereas others are labeled as stackable or freestanding. Rack-mount devices are designed for placement into equipment racks, which are a common sight in computer rooms. The racks are approximately 19 inches wide; devices designed to be rack-mounted are slightly smaller than freestanding devices, so they can fit in the racks. Small metal brackets are screwed to the sides of the devices to allow them to be fitted into the racks.

If you don't have racks, you need to use stackable or freestanding devices. These devices can literally be placed on top of one another. Many network equipment manufacturers realize that not everyone has racks, and so they make their equipment usable in either a rack or a freestanding configuration.

Managed Hubs and Switches

Both hubs and switches come in managed and unmanaged versions. A managed device has an interface through which it can be configured to perform certain special functions. For example, it may allow for port mirroring, which can be useful for network monitoring, or allow ports to be specified to operate at a certain speed. Because of the extra functionality of a managed device, and because of the additional components required to achieve it, managed devices are considerably more expensive than unmanaged devices. When you're specifying switches or hubs, consider the need for manageability carefully. If a switch will be used to connect

servers to the network, a managed device might make the most sense—the extra functionality might come in handy. On parts of the network that accommodate client computers, unmanaged devices generally suffice.

NOTE

Port Density Excluding the small workgroup hubs, hubs and switches normally have 8, 16, 24, or 32 ports each, although variations are available. To help you compare prices between devices, manufacturers often quote a price per port. In some cases, a higher-density device with more ports may cost significantly less per port than a device with fewer ports. Typically, the more ports on a device, the lower the price per port.

At the time of this writing, switches are still quite a bit more expensive than hubs with equivalent capacity, but the gap is narrowing quickly. Some manufacturers have stopped producing hubs and instead are putting all their efforts into developing switches. This would seem to be a sound strategy. In all but the smallest networks or companies with the most restrictive budgets, hubs are rapidly being replaced by switches. In new implementations, hubs are unlikely to be specified and installed.

Bridges

Bridges are networking devices that connect networks. Sometimes it is necessary to divide networks into subnets to reduce the amount of traffic on each larger subnet or for security reasons. Once divided, the bridge connects the two subnets and manages the traffic flow between them. Today, network switches have largely replaced bridges.

A bridge functions by blocking or forwarding data, based on the destination MAC address written into each frame of data. If the bridge believes the destination address is on a network other than that from which the data was received, it can forward the data to the other networks to which it is connected. If the address is not on the other side of the bridge, the data is blocked from passing. Bridges "learn" the MAC addresses of devices on connected networks by "listening" to network traffic and recording the network from which the traffic originates. Figure 3.9 shows a representation of a bridge.

NOTE

Manual Bridge Configuration Some early bridge implementations required you to enter the information for each device on the network manually. Fortunately, bridges are now of the learning variety, and manual configuration is no longer necessary.

The advantages of bridges are simple and significant. By preventing unnecessary traffic from crossing onto other network segments, a bridge can dramatically reduce the amount of

network traffic on a segment. Bridges also make it possible to isolate a busy network from a not-so-busy one, thereby preventing pollution from busy nodes.

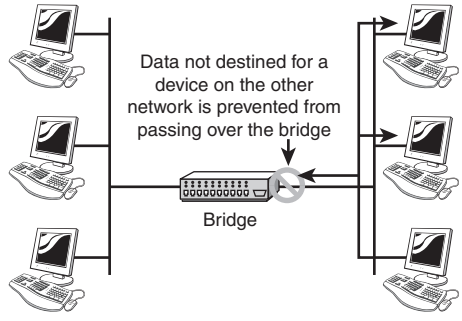

FIGURE 3.9 How a bridge works.

Bridge Implementation Considerations

Although implementing bridges can offer huge improvements in performance, you must factor in a number of considerations. The first is bridge placement. Generally, you should follow the 80/20 rule for bridge placement: 80% of the traffic should not cross the bridge, and 20% of the traffic should be on the other side of the bridge. The rule is easy to understand, but accurately determining the correct location for the bridge to accommodate the rule is another matter.

Another, potentially more serious, consideration is bridging loops, which can be created when more than one bridge is used on a network. Multiple bridges can provide fault tolerance or improve performance. Bridging loops occur when multiple bridges become confused about where devices are on the network.

As an example of bridging loops, imagine that you have a network with two bridges, as depicted in Figure 3.10. During the learning process, the north bridge receives a packet from Interface A (step 1 in Figure 3.11) and determines that it is for a system that is not on Network Z; therefore, the bridge forwards the packet to Network X (step 2 in Figure 3.11). Now, the south bridge sees a packet originating on Network X on Interface C (step 3 in Figure 3.11); because it thinks the destination system is not on Network X, it forwards the packet to Network Z (step 4 in Figure 3.11), where the north bridge picks it up (step 5 in Figure 3.11). The north bridge determines that the destination system is not on Network Z, so it forwards the packet to Network X—and the whole process begins again.

You can work around the looping problem by using the Spanning Tree Algorithm (STA). When STA is used, each interface on a bridge is assigned a value. As the bridge forwards the data, the value is attached to the packet. When another bridge sees the data, if the STA value for the interface is higher than that assigned to its interfaces, the bridge doesn't forward the data, thus eliminating the possibility of a bridging loop. STA eliminates the bridging loop but

still provides the fault tolerance of having more than one bridge in place. If the bridge with the higher STA value (sometimes referred to as the *primary bridge*) fails, the other bridge continues functioning because it becomes the bridge with the higher STA value. All this is achieved by the Spanning Tree Protocol (STP).

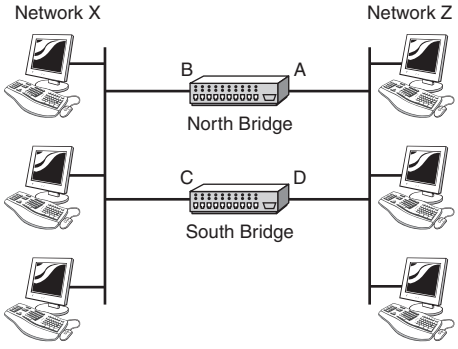

FIGURE 3.10 A network with two bridges.

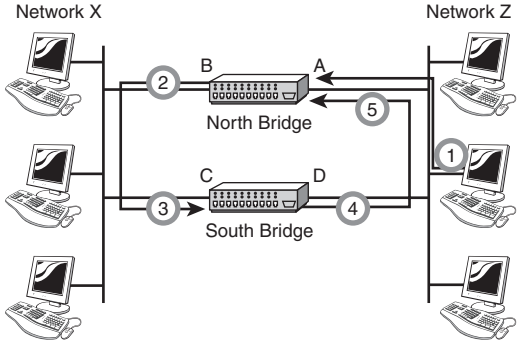

FIGURE 3.11 A bridging loop.

NOTE

STP STP is defined in the IEEE 802.1d standard.

Types of Bridges

Three types of bridges are used in networks. You don't need detailed knowledge of how each bridge works, but you should have an overview:

- **Transparent bridge**—A transparent bridge is invisible to the other devices on the network. Transparent bridges perform only the function of blocking or forwarding data based on the MAC address; the devices on the network are oblivious to these bridges' existence. Transparent bridges are by far the most popular types of bridges.

- **Translational bridge**—A translational bridge can convert from one networking system to another. As you might have guessed, it translates the data it receives. Translational bridges are useful for connecting two different networks, such as Ethernet and Token Ring networks. Depending on the direction of travel, a translational bridge can add or remove information and fields from the frame as needed.
- **Source-route bridge**—Source-route bridges were designed by IBM for use on Token Ring networks. The source-route bridge derives its name from the fact that the entire route of the frame is embedded within the frame. This allows the bridge to make specific decisions about how the frame should be forwarded through the network. The diminishing popularity of Token Ring makes the chances that you'll work with a source-route bridge very slim.

EXAM ALERT

Identify the Bridge On the Network+ exam, you might be asked to identify the purpose of a certain type of bridge.

As switches become ever cheaper, bridges have been overtaken by switches in terms of both functionality and performance. Expect to be working with switches more often than with bridges.

Routers

Routers are an increasingly common sight in any network environment, from a small home office that uses one to connect to an Internet service provider (ISP) to a corporate IT environment where racks of routers manage data communication with disparate remote sites. Routers make internetworking possible, and in view of this, they warrant detailed attention.

Routers are network devices that literally route data around the network. By examining data as it arrives, the router can determine the destination address for the data; then, by using tables of defined routes, the router determines the best way for the data to continue its journey. Unlike bridges and switches, which use the hardware-configured MAC address to determine the destination of the data, routers use the software-configured network address to make decisions. This approach makes routers more functional than bridges or switches, and it also makes them more complex because they have to work harder to determine the information. Figure 3.12 shows basically how a router functions.

The basic requirement for a router is that it must have at least two network interfaces. If they are LAN interfaces, the router can manage and route the information between two LAN segments. More commonly, a router is used to provide connectivity across wide area network (WAN) links. Figure 3.13 shows a router with two LAN ports (marked AUI 0 and AUI 1) and

two WAN ports (marked Serial 0 and Serial 1). This router is capable of routing data between two LAN segments and two WAN segments.

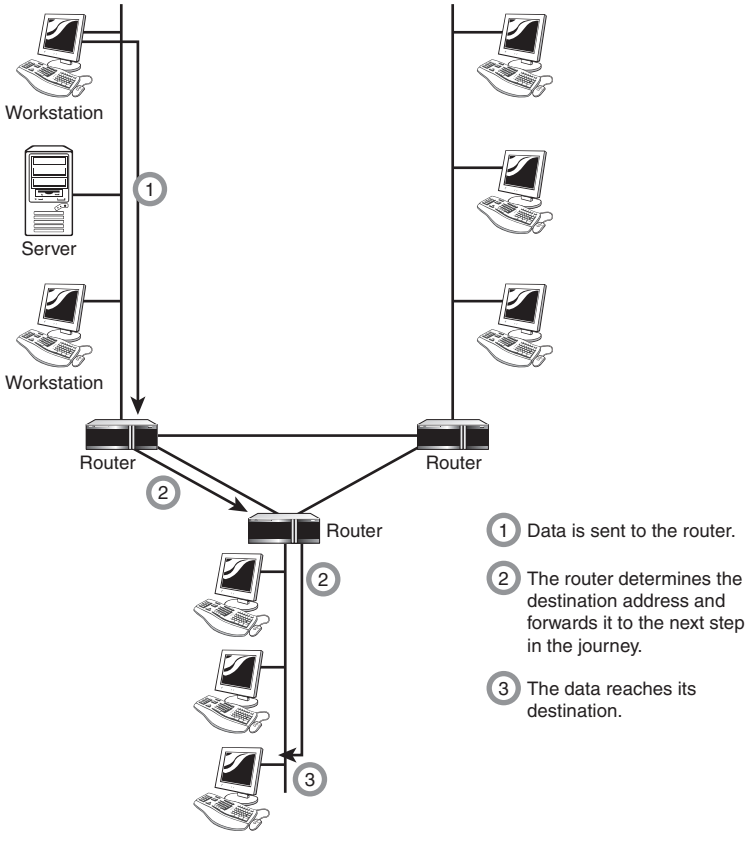

FIGURE 3.12 The basic function of a router.

Routable Protocols and Routing Protocols

Routers rely on two types of network protocols to make the routing magic happen: routable protocols and routing protocols. We'll examine them separately in the next sections.

Routable Protocols

Large internetworks need protocols that allow systems to be identified by the address of the network to which they are attached and by an address that uniquely identifies them on that network. Network protocols that provide both of these features are said to be *routable*. Three routable LAN network protocols are used today:

FIGURE 3.13 A router with two LAN ports and two WAN ports.

- **Transmission Control Protocol/Internet Protocol (TCP/IP)**—TCP/IP was developed in the 1970s by the Department of Defense, which needed a protocol to use on its WAN. TCP/IP's flexibility, durability, and functionality meant that it soon became the WAN protocol of choice and also became the standard for LANs. Today, most networks use TCP/IP in some fashion, even if the main LAN protocol is something other than TCP/IP. TCP/IP is a huge topic, and anyone in networking must understand it. The Network+ exam dedicates an entire exam objective to it. Chapter 6, "Working with TCP/IP," provides complete coverage of TCP/IP.

  TCP/IP is a protocol suite comprised of numerous individual protocols. Within TCP/IP, the two routing protocols used are the Routing Information Protocol (RIP) and Open Shortest Path First (OSPF). RIP is a distance-vector routing protocol, and OSPF is a link-state routing protocol.

- **Internetwork Packet Exchange/Sequenced Packet Exchange (IPX/SPX)**—Created by Novell for use on NetWare networks, IPX/SPX is a routable protocol that was popular for many years. Today, even Novell acknowledges that TCP/IP is the network protocol of choice and so has moved away from IPX/SPX and toward a pure TCP/IP environment. In fact, the last few versions of Novell NetWare have used TCP/IP as the default protocol and allowed IPX/SPX to be enabled if needed.

  Like TCP/IP, IPX/SPX is also a protocol suite. Within IPX/SPX, the NetWare Link State Protocol (NLSP) and the Routing Information Protocol (RIP) manage routing. RIP uses a distance-vector route-discovery method, which calculates routes based on the number of hops. NLSP uses a link-state route discovery method to build routing tables.

- **AppleTalk**—AppleTalk is a full-featured protocol designed to be used with the Macintosh computer systems. AppleTalk has been around since the early 1980s and widely deployed in Apple networks. Like TCP/IP and IPX/SPX, AppleTalk is a protocol suite. Within the suite, the Routing Table Maintenance Protocol (RTMP) provides routing functionality. RTMP is a distance-vector routing protocol similar to the RIP, which is used by IPX/SPX and TCP/IP.

Some routers are capable of routing more than one protocol at a time, a feature known as *multiprotocol routing*. Multiprotocol routing brings with it a number of considerations, not the least of which is the fact that a multiprotocol router may have to work considerably harder than a router working with only a single protocol. This is the case not only because there is more than one protocol but because there may also be multiple routing protocols.

Routing Protocols

Routing protocols are the means by which routers communicate with each other. This communication is necessary so that routers can learn the network topology and changes that occur in it.

NOTE

Static Routing The alternative to using routing protocols is *static routing*, which means that route information must be manually entered by the administrator. There are two main disadvantages of this approach: First, manually entering routes is time-consuming and susceptible to human error. Second, if the topology of the network changes, the routers must be manually reconfigured. Therefore, static routing is generally used only in the smallest of environments. In environments with more than a handful of routers, dynamic routing is the preferred option.

NOTE

Metrics In routing, the term *metric* describes the "cost" of a certain route. The metric can be a combination of factors, including the number of routers between a router's position and the destination, the time it takes to complete the journey, and even a value that can be assigned by an administrator to discourage use of a certain route. Under normal circumstances, routers choose the route with the lowest metric.

The two types of routing protocols are *distance-vector* and *link-state protocols*. Each has a different strategy for dealing with router-to-router communication.

Distance-Vector Protocols

With distance-vector routing protocols, each router communicates all the routes it knows about to all other routers to which it is directly attached (that is, its *neighbors*). Because each router in the network knows only about the routers to which it is attached, it doesn't know

how to complete the entire journey; instead, it only knows how to make the next hop. *Hops* are the means by which distance-vector routing protocols determine the shortest way to reach a given destination. Each router constitutes one hop; so if a router is four hops away from another router, there are three routers, or hops, between itself and the destination. Distance-vector protocols can also use a time value known as a *tick*, which enables the router to make a decision about which path is quickest if given the choice of more than one (a common situation on networks with redundant links).

The frequency with which routers send route updates depends on the routing protocol being used, but it is usually between 10 and 60 seconds. At each update, the entire routing table of the sender is sent to the other connected routers. When the other routers receive the information, they check it against the existing information; if there are any changes, they alter their routing tables accordingly.

This constant update cycle is one of the problems of distance-vector routing protocols because it can lead to large amounts of network traffic. Furthermore, after the initial learning period, the updates should (hopefully) be irrelevant—the chances of the network topology changing every 30 seconds or so are slim, and if you do have such a network, some troubleshooting may be in order.

When a change does occur on the network, it may take some time for all the routers to learn of the change. The process of each router learning about the change and updating its routing tables is known as *convergence*. In a small network, convergence might not take long; but in larger networks, those with, say, more than 20 routers, it might take some time to complete. Rather than cause the routers to wait for the updates, you can configure *triggered updates*, which are sent when a topology change is detected. Using triggered updates can significantly improve the convergence speed of distance-vector–based networks.

You can also use *hold-down timers* to improve convergence. A hold-down timer prevents a router from trying to make too many changes too quickly. When a router receives a change about a route, it makes the change and then applies a hold-down timer to the change. The hold-down timer prevents further changes from being made to that route within the defined time period. Hold-down timers are particularly useful when an unreliable router keeps going on and off the network. If hold-down timers are not applied, updates to the routing tables on routers would continually be changing, and the network might never converge.

In some configurations, distance-vector routing protocols can lead to routing loops. *Routing loops* occur when a router tells another router about a route that it heard about from the same router. For example, consider the router layout in Figure 3.14. If Router C becomes unable to access Router D through Network 1, it removes the route from its table and sends the update to Router B; Router B removes the route. But if Router B receives an update from Router A before it sends an update to Router A, the route is reinstated because according to Router A, it can still access Network 1. Now Router B begins to send anything destined for Network 1 back to Router A, which duly sends it back to Router B, and so on, thus creating a routing loop.

Each time the route is added to the table, the hop count for the route increases—a problem known as the *count to infinity*.

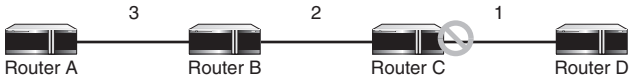

FIGURE 3.14 How routing loops occur.

You can use two strategies to prevent routing loops when using distance-vector routing protocols:

- **Split horizon**—The split horizon algorithm addresses the problem of routing loops by not advertising routes back on the interface from which they are learned. In other words, using Figure 3.14 as an example, Router C would not advertise back to Router B any route that it learned from Router B. Basically, Router C figures that, because it learned about the route from Router B, Router B must be nearer to the destination than it is.
- **Split horizon with poison reverse**—With this strategy, also known simply as poison reverse, routers do advertise routes back on the interfaces from which they were learned, but they do so with a hop count of infinity. The value used for infinity (which seems like an impossible situation) depends on the routing protocol being used. Again using the example from Figure 3.14, Router C would advertise to Router B the routes it learned from Router B, but it would also add the infinite hop count. In other words, Router C would say, "I know about Router A, but I can't reach it myself." This way, Router B would never try to add the route to Router A through Router C, because according to Router C, it can't reach Router A.

The most popular distance-vector routing protocols are both called Routing Information Protocol (RIP). The distance-vector routing protocol for TCP/IP is called RIP, as is the one for IPX/SPX. To set the two apart, the IPX version is often called IPX RIP.

Link-State Protocols

A router that uses a link-state protocol differs from a router that uses a distance-vector protocol because it builds a map of the entire network and then holds that map in memory. On a network that uses a link-state protocol, routers send out link-state advertisements (LSAs) that contain information about what networks they are connected to. The LSAs are sent to every router on the network, thus enabling the routers to build their network maps.

When the network maps on each router are complete, the routers update each other at a given time, just like with a distance-vector protocol, but the updates occur much less frequently with link-state protocols than with distance-vector protocols. The only other circumstance under which updates are sent is if a change in the topology is detected, at which point the routers

use LSAs to detect the change and update their routing tables. This mechanism, combined with the fact that routers hold maps of the entire network, makes convergence on a link-state–based network occur very quickly.

Although it might seem like link-state protocols are an obvious choice over distance-vector protocols, routers on a link-state–based network require more powerful hardware and more RAM than those on a distance-vector–based network. Not only do the routing tables have to be calculated, but they must also be stored. A router that uses distance-vector protocols need only maintain a small database of the routes accessible by the routers to which it is directly connected. A router that uses link-state protocols must maintain a database of the routers in the entire network.

Two of the most popular link-state routing protocols are Open Shortest Path First (OSPF) and NetWare Link State Protocol (A.K.A NetWare Link Services Protocol) (NLSP). The former is used on TCP/IP networks, and the latter is used on networks that use IPX/SPX.

EXAM ALERT

Identify the Protocols Be prepared to identify both the link-state and distance-vector routing protocols used on both TCP/IP and IPX/SPX networks.

NOTE

Multiprotocol Routing In this section and the previous section, we discussed routing protocols as they apply to single protocols. But remember that one router may be routing more than one protocol; it may, for example, use OSPF and NLSP.

Dedicated Hardware Versus Server-Based Routers

A router can be either a dedicated hardware device or a server system that has at least two network interfaces installed in it. All common network operating systems offer the capability to act as routers as part of their functionality.

Dedicated hardware routers offer greater performance levels than server-based solutions, but they have the disadvantage of offering a limited range of features for their cost. However, the attraction of a dedicated hardware device often outweighs this factor.

The following are some of the advantages of dedicated hardware routers:

- Typically faster than server-based routers
- Generally more reliable than server-based routers
- Easier to harden against attacks than server-based routing solutions

The following are some of the disadvantages of dedicated hardware routers:

- More expensive than server-based router solutions; extra functionality may have to be purchased
- Often require specialized skills and knowledge to manage them
- Limited to a small range of possible uses

The capabilities of a router depend on the features it has. A basic router may route only one protocol between two network interfaces of the same type. A more advanced router may act as a gateway between two networks and two protocols. In addition, it may offer firewall services, security and authentication, or remote access functionality such as virtual private networking.

NOTE

Brouters A *brouter* is a device that can route traffic that can be routed and bridge anything that cannot be routed. As bridges have been replaced by the more flexible routers, brouters have also fallen out of favor. In today's networking world, routers rule. Just ask Cisco.

The topic of routing is complex, and the routing information provided in this chapter is the most basic of tutorials. Although we've told you what you need to know for the exam, if you're working with routers on a daily basis, you will want to seek out further sources of information—and there is no shortage of such sources (including the Cisco Press and *Exam Cram* titles you'll find at www.informit.com).

Gateways

The term *gateway* is applied to any device, system, or software application that can perform the function of translating data from one format to another. The key feature of a gateway is that it converts the format of the data, not the data itself.

NOTE

Gateways Versus Default Gateways Don't confuse gateways, which are networking devices, with default gateways, which are discussed in Chapter 6, "Working with TCP/IP." The two perform different roles on a network.

You can use gateway functionality in many ways. For example, a router that can route data from an IPX network to an IP network is, technically, a gateway. The same can be said of a translational bridge that, as described earlier in this chapter, converts from an Ethernet network to a Token Ring network and back again.

Software gateways can be found everywhere. Many companies use an email system such as Microsoft Exchange or Novell GroupWise. These systems transmit mail internally in a certain format. When email needs to be sent across the Internet to users using a different email system, the email must be converted to another format, usually to Simple Mail Transfer Protocol (SMTP). This conversion process is performed by a software gateway.

Another good (and often used) example of a gateway involves the Systems Network Architecture (SNA) gateway, which converts the data format used on a PC to that used on an IBM mainframe or minicomputer. A system that acts as an SNA gateway sits between the client PC and the mainframe and translates requests and replies from both directions. Figure 3.15 shows how this would work in a practical implementation.

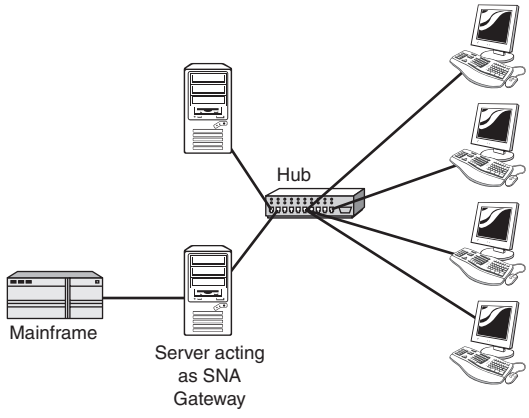

FIGURE 3.15 An SNA gateway.

If it seems from the text in this section that we are being vague about what a gateway is, it's because there is no definite answer. The function of a gateway is very specific, but how the gateway functionality is implemented is not.

No matter what their use, gateways slow the flow of data and can therefore potentially become bottlenecks. The conversion from one data format to another takes time, and so the flow of data through a gateway is always slower than the flow of data without one.

CSUs/DSUs

A Channel Service Unit/Data Service Unit (CSU/DSU) acts as a translator between the LAN data format and the WAN data format. Such a conversion is necessary because the technologies used on WAN links are different from those used on LANs. Some consider a CSU/DSU as a type of digital modem; but unlike a normal modem, which changes the signal from digital to analog, a CSU/DSU changes the signal from one digital format to another. Figure 3.16 shows how a CSU/DSU might fit into a network.

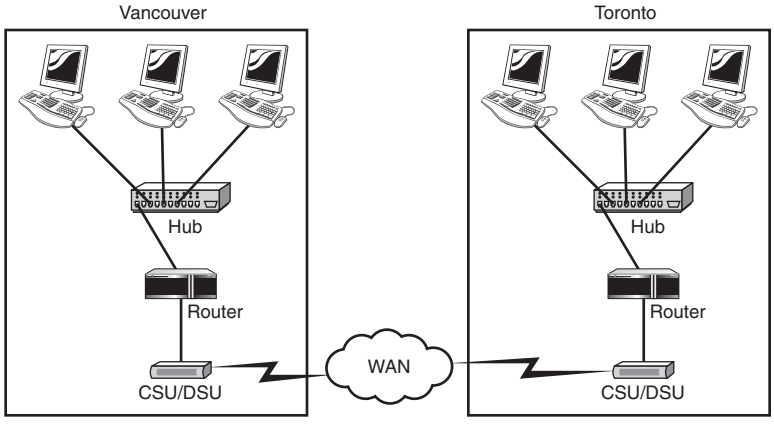

FIGURE 3.16 How a CSU/DSU is used in a network.

A CSU/DSU has physical connections for the LAN equipment, normally via a serial interface, and another connection for a WAN. Traditionally, the CSU/DSU has been in a separate box from other networking equipment; however, the increasing use of WAN links means that some router manufacturers are now including the CSU/DSU functionality in routers or are providing the expansion capability to do so.

Modems

Modem is a contraction of the terms *modulator* and *demodulator*. Modems perform a simple function: They translate digital signals from a computer into analog signals that can travel across conventional phone lines. The modem modulates the signal at the sending end and demodulates at the receiving end.

Modems provide a relatively slow method of communication. In fact, the fastest modem available on the market today has a maximum speed of 56Kbps. Compare that to the speed of a 10Mbps network connection, and you'll find that the modem is approximately 180 times slower. That makes modems okay for browsing web pages or occasionally downloading small files but wholly unsuitable for downloading large files. As a result, many people prefer to use other remote access methods, including ISDN (which is discussed later in this chapter, in the section "ISDN Terminal Adapters") and cable/DSL access.

Modems are available as internal devices that plug into expansion slots in a system; external devices that plug into serial or USB ports; PCMCIA cards designed for use in laptops; and specialized devices designed for use in systems such as handheld computers. In addition, many laptops now come with integrated modems. For large-scale modem implementations, such as at an ISP, rack-mounted modems are also available. Figure 3.17 shows an internal modem and a PCMCIA modem.

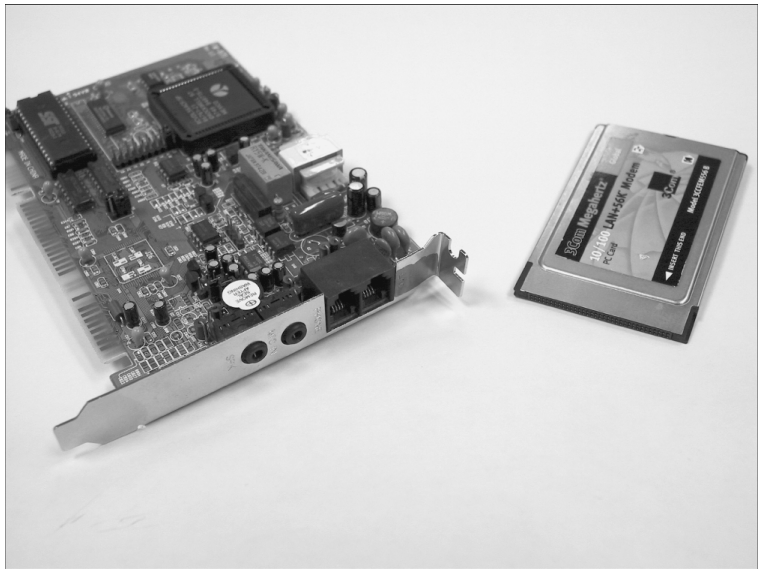

FIGURE 3.17 An internal modem (left) and a PCMCIA modem (right).

Modems are controlled through a series of commands known as the Hayes AT command set. Hayes was a company that, for many years, led the field in the development of modems and modem technology. The AT commands allow you to control a modem as well as configure and diagnose it. Table 3.1 lists some of the most commonly used AT commands.

EXAM ALERT

Know the AT Command On the Network+ exam, you might be asked to identify the correct AT command to be used in a given situation.

TABLE 3.1 Commonly Used AT Modem Commands

Command	Result
ATA	Answers an incoming call
ATH	Hangs up the current connection
ATZ	Resets the modem
ATI3	Displays modem identification information

Modem Connection Speeds

The actual speed you obtain on a modem connection depends on a variety of factors, including the quality of the line you are using and the speed of the modem. For example, you might

find (as we often do) that even with a 56Kbps modem, the most you can get on a certain connection is 49Kbps. If you try the same connection again on a different phone line, you might get a higher or lower rate. Quality of the connection aside, two factors govern the maximum speed attainable by your modem: the speed of the Universal Asynchronous Receiver/Transmitter (UART) chip in your system (which controls the serial ports) and the speed of the modem itself.

In older systems, the UART chips were capable of only slow speeds, making them unable to keep up with fast modems. Today, most systems have UART chips capable of speeds well in excess of those offered by modems. Now the modem, not the UART chip, is the bottleneck. Table 3.2 lists the types of commonly used UART chips and their associated speeds.

TABLE 3.2 UART Chips and Their Associated Speeds

UART Chip	Speed (bps)
8250	9,600
16450	115,200
16550	115,200
16650	430,800
16750	921,600
16950	921,600

EXAM ALERT

Know the UART Speed On the Network+ exam, you might be asked to identify the maximum speed of a given UART chip.

Modem speeds can be expressed in either baud rate or bits per second (bps). The *baud rate* refers to the number of times a signal changes in each second, and the *bps rate* is the number of bits of data that can be sent or received in a second. Although the figures are identical in some modems, in others the bps rate is higher than the baud rate. The baud rate is actually not as important, and the higher the bps figure, the better. Most modern modems offer bps rates far greater than the baud rate.

To make it easier to compare modems, standards have been created that define the data throughput of the modem and what features it provides. These are sometimes referred to as the *V standards*, and you can use them when buying a modem to determine the modem's capabilities.

Network Interface Cards (NICs)

NICs—sometimes called network cards—are the mechanisms by which computers connect to a network. NICs come in all shapes and sizes, and they come in prices to suit all budgets. Consider the following when buying a NIC:

NOTE

NIC Terminology Many different terms are used to refer to NICs, such as *network card*, *network adapter*, and *LAN adapter*. All refer to the same thing.

- **Network compatibility**—Perhaps this is a little obvious, but sometimes people order the wrong type of NIC for the network. Given the prevalence of Ethernet networks, you are likely to have to specify network compatibility only when buying a NIC for another networking system.
- **Bus compatibility**—Newly purchased NICs will almost certainly use the Peripheral Component Interconnect (PCI) bus, although if you are replacing a card in an older system, you might have to specify an Industry Standard Architecture (ISA) bus card instead. If the card you are buying is PCI, check to see what kind of PCI interface is being used. Many high-end server systems now come with 64-bit PCI slots; if you have them, it is definitely worth taking advantage of the extra performance they offer. Such 64-bit PCI slots can be easily identified because they are the same color and width as 32-bit PCI slots but are longer. 64-bit slots are referred to as PCI-X and are backward compatible with 32-bit PCI. Figure 3.18 shows 32-bit PCI slots on a system board.
- **Port compatibility**—Generally a NIC has only one port, for twisted-pair cabling. If you want some other connectivity, you need to be sure to specify your card accordingly; for example, you might need a fiber-optic or coaxial cable port.

NOTE

Combo Cards Sometimes a NIC has a twisted-pair socket, a coaxial connector, and an attachment unit interface (AUI) port. These cards are referred to as *combo* cards. Today, the dominance of twisted-pair cabling means that most NICs have only a twisted-pair connection.

- **Hardware compatibility**—Before installing a network card into a system, you must verify compatibility between the network card and the operating system on the PC in which you are installing the NIC. If you are using good-quality network cards from a recognized manufacturer, such verification should be little more than a formality.

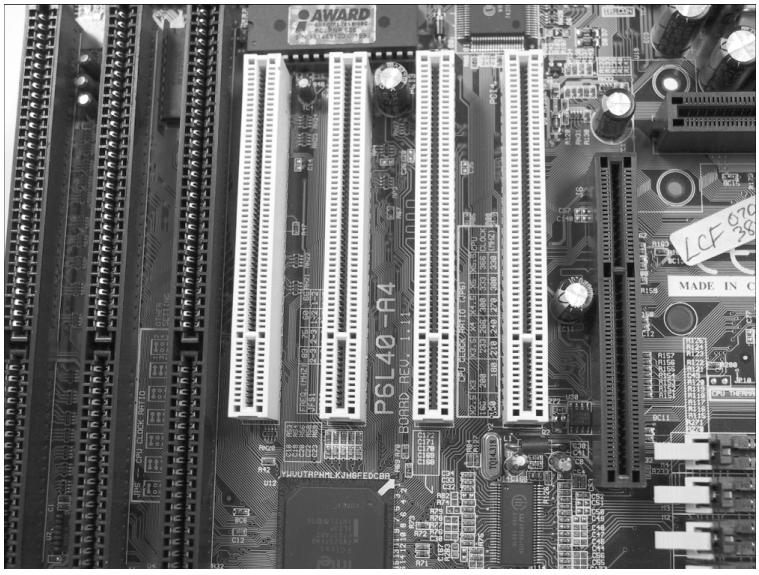

FIGURE 3.18 32-bit PCI slots on a system board. (Photo copyright © Intel Corporation.)

Types of Network Interfaces

Network interfaces come as add-in expansion cards or as PCMCIA cards used in laptop systems. In some cases, rather than have an add-in NIC, the network interface is embedded into the motherboard. Figure 3.19 shows an example of an add-in NIC, Figure 3.20 shows a PCMCIA network card, and Figure 3.21 shows a built-in network interface in a laptop system.

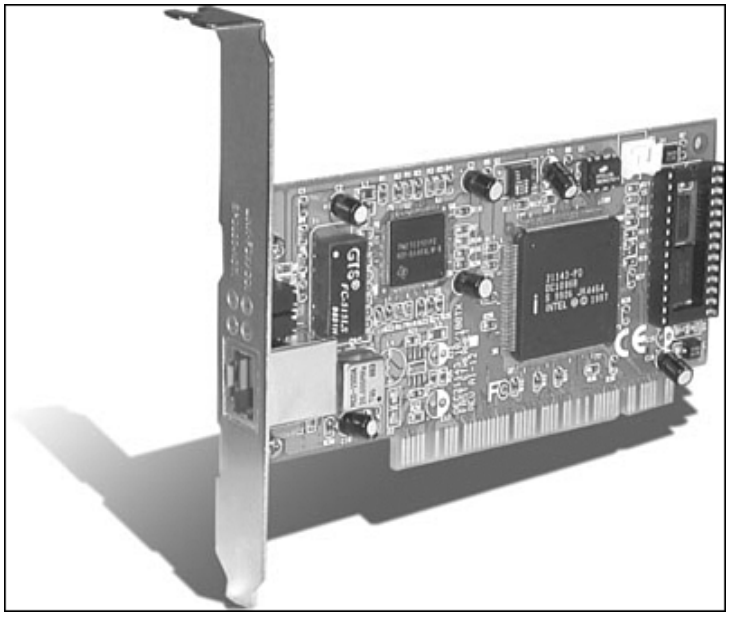

FIGURE 3.19 An expansion NIC.

FIGURE 3.20 A PCMCIA NIC.

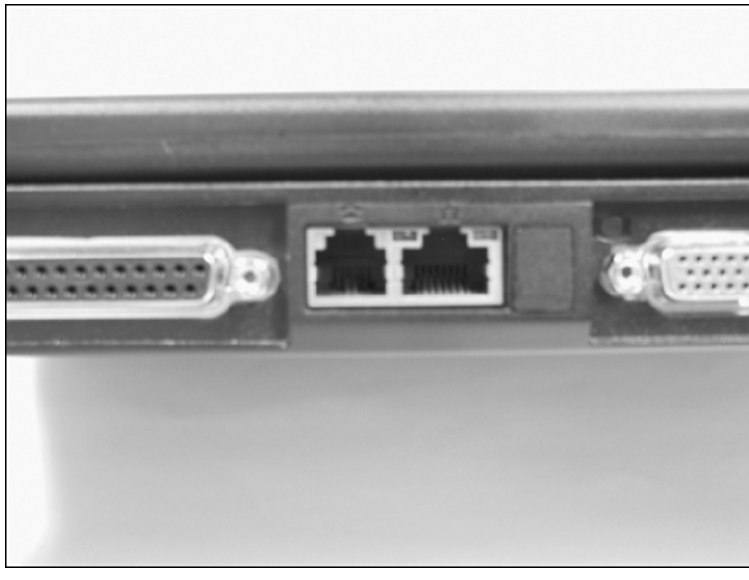

FIGURE 3.21 A built-in network interface on a laptop system.

The False Economy of NICs

The difference between an inexpensive network card and an expensive one is less than you might think; but even so, people are tempted to go for the low-cost option. In many cases, this turns out to be a false economy. Not only do higher-end cards tend to be easier to install, they are generally easier to

(continues)

(continued)

troubleshoot as well. An hour trying to troubleshoot a misbehaving inexpensive network card can negate any cost savings from the purchase. This is particularly relevant on server systems, where a problem network card will not only cause you frustration but also will most likely cause the users of the server problems. In fact, if you are working on server systems, it's worth investigating fault-tolerant network card configurations, such as adapter teaming.

A network interface typically has at least two LEDs that indicate certain conditions:

- **Link light**—This LED indicates whether a network connection exists between the card and the network. An unlit link light is an indicator that something is awry with the network cable or connection.
- **Activity light**—This LED indicates network activity. Under normal conditions, the light should flicker sporadically and often. Constant flickering may indicate a very busy network or a problem somewhere on the network that is worth investigating.
- **Speed light**—This LED indicates that the interface is connected at a certain speed. This feature is normally found on Ethernet NICs that operate at 10Mbps/100Mbps—and then only on certain cards.

Some network cards combine the functions of certain lights by using dual-color LEDs. PCMCIA cards sometimes have no lights, or the lights are incorporated into the media adapter that comes with the card. You can see an example in Figure 3.22.

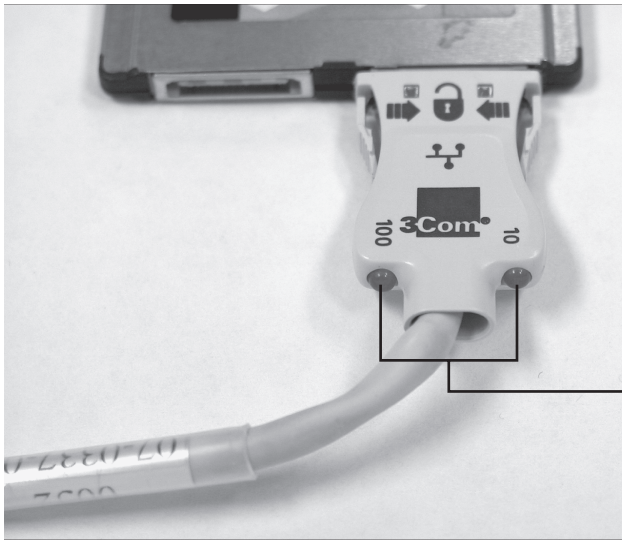

FIGURE 3.22 Indicator lights on a media adapter for a PCMCIA NIC.

Installing Network Cards

At some point in your networking career, it is likely that you will have to install a NIC into a system. For that reason, an understanding of the procedures and considerations related to NIC installations is useful. Here are some of the main things to consider:

EXAM ALERT

Avoid ESD When installing any component in a system, you need to observe proper and correct procedures to guard against electrostatic discharge (ESD). ESD can cause components to fail immediately or degrade so that it fails at some point in the future. Proper ESD precautions include wearing an antistatic wrist strap and properly grounding yourself.

- **Drivers**—Almost every NIC is supplied with a driver disk, but the likelihood of the drivers on the disk being the latest drivers is slim. Always make sure that you have the latest drivers by visiting the website of the NIC manufacturer. The drivers play an important role in the correct functioning of the NIC, so spend a few extra minutes to make sure that the drivers are installed and configured correctly.
- **NIC configuration utilities**—In days gone by, NICs were configured with small groups of pins known as *jumpers*, or with small plastic blocks of switches known as *dip switches*. Figure 3.23 shows an example of jumpers. Unless you are working with very old equipment, you are unlikely to encounter dip switches.

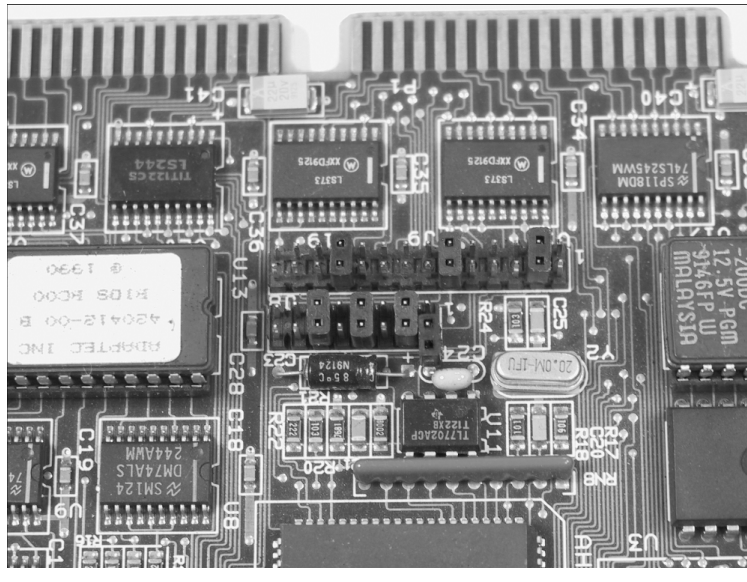

FIGURE 3.23 A block of jumpers.

Although these methods were efficient and easy to use, they have now largely been abandoned in favor of software configuration utilities, which allow you to configure the settings for the card (if any) and to test whether the card is working properly. Other utilities can be used through the operating system to obtain statistical information, help, and a range of other features.

- **System resources**—To function correctly, NICs must have certain system resources allocated to them: the interrupt request (IRQ) and memory addresses. In some cases, you might need to assign the values for these manually. In most cases, you can rely on plug-and-play, which assigns resources for devices automatically.
- **Physical slot availability**—Most modern PCs have at least three or four usable expansion slots. Not only that, but the increasing trend toward component integration on the motherboard means that devices such as serial and parallel ports and sound cards are now built in to the system board and therefore don't use up valuable slots. If you're working on older systems or systems that have a lot of add-in hardware, you might be short of slots. Check to make sure that a slot is available before you begin.
- **Built-in network interfaces**—A built-in network interface is a double-edged sword. The upsides are that it doesn't occupy an expansion slot, and hardware compatibility with the rest of the system is almost guaranteed. The downside is that a built-in component is not upgradeable. For this reason, you might find yourself installing an add-in NIC and at the same time disabling the on-board network interface. Disabling the on-board interface is normally a straightforward process, achieved by going into the BIOS setup screen or, on some systems, a system configuration utility. In either case, consult the documentation that came with the system or look for information on the manufacturer's website.

As time goes on, NIC and operating system manufacturers are making it increasingly easy to install NICs in systems of all sorts and sizes. By understanding the requirements of the card and the correct installation procedure, you should be able to install cards simply and efficiently.

ISDN Terminal Adapters

When the speed provided by a modem just isn't enough, you must seek alternatives. One of the speedier options available is an ISDN link. ISDN is a digital communication method that can be used over a conventional phone line, although certain criteria must be met for an ISDN line to be available (such as the availability of the service and the proximity of your location to the telco's site). (The information in this section is intended to cover only ISDN terminal adapters, not ISDN as a system. Detailed coverage of ISDN is provided in Chapter 7, "WAN and Internet Access Technologies," which covers WAN topics.)

To use ISDN, you need a device called an *ISDN terminal adapter*. ISDN terminal adapters are available as add-in expansion cards installed into computers, external devices that connect to the serial interfaces of PC systems, or modules in a router. You can think of an ISDN terminal adapter as a kind of digital modem. (Remember that a modem converts a signal from digital to analog and vice versa. An ISDN terminal adapter translates the signal between two digital formats.) Figure 3.24 shows an external ISDN terminal adapter, and Figure 3.25 shows an example of an internal ISDN adapter. Notice that an ISDN terminal adapter is similar in appearance to a standard NIC.

FIGURE 3.24 An external ISDN adapter.

FIGURE 3.25 An internal ISDN adapter.

Installing an external ISDN adapter is simple because, like an external modem, an external ISDN adapter plugs in to the serial port of the system and thus uses its resources. You need drivers for an ISDN terminal adapter, so be sure to visit the manufacturer's website and download the latest available drivers. An internal ISDN terminal adapter requires a little more effort: You must make sure that you have physical and logical system resources to accommodate it.

Wireless Access Point (WAP)

Wireless access points, referred to as either *WAPs* or *wireless APs*, are a transmitter and receiver (*transceiver*) device used for wireless LAN (WLAN) radio signals. A WAP is typically a separate network device with a built-in antenna, transmitter, and adapter. WAPs use the wireless infrastructure network mode to provide a connection point between WLANs and a wired Ethernet LAN. WAPs also typically have several ports allowing a way to expand the network to support additional clients.

Depending on the size of the network, one or more WAPs may be required. Additional WAPs are used to allow access to more wireless clients and to expand the range of the wireless network. Each WAP is limited by a transmissions range, the distance a client can be from a WAP and still get a useable signal. The actual distance depends on the wireless standard being used and the obstructions and environmental conditions between the client and the WAP. Figure 3.26 shows an example of a WAP in a network configuration.

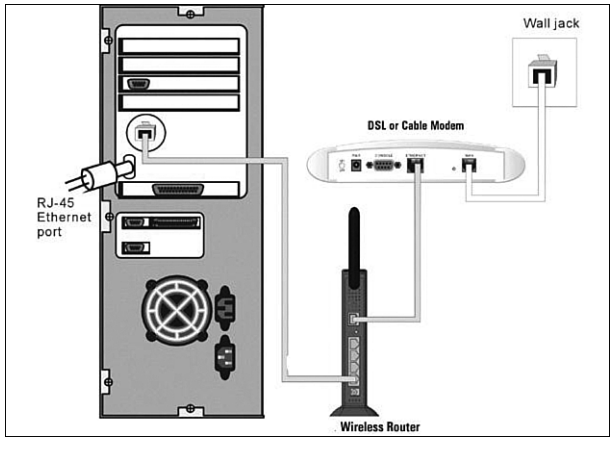

FIGURE 3.26 WAPs connect WLANs and a wired Ethernet LAN.

NOTE

Wireless Access Points—A WAP can operate as a bridge connecting a standard wired network to wireless devices or as a router passing data transmissions from one access point to another.

TIP

WAP Range If you are using a wireless device that loses its connection, it may be that you are too far away from the WAP.

As mentioned, a WAP is used in an infrastructure wireless network design. Used in the infrastructure mode, the WAP receives transmissions from wireless devices within a specific range and transmits those signals to the network beyond. This network may be a private Ethernet network or the Internet. The transmission range a WAP can support and number of wireless devices that can connect to it depends on the wireless standard being used and the signal interference between the two devices. In infrastructure wireless networking, there may be multiple access points to cover a large area or only a single access point for a small area such as a single home or small building.

Figure 3.27 shows an example of an infrastructure wireless network using a WAP.

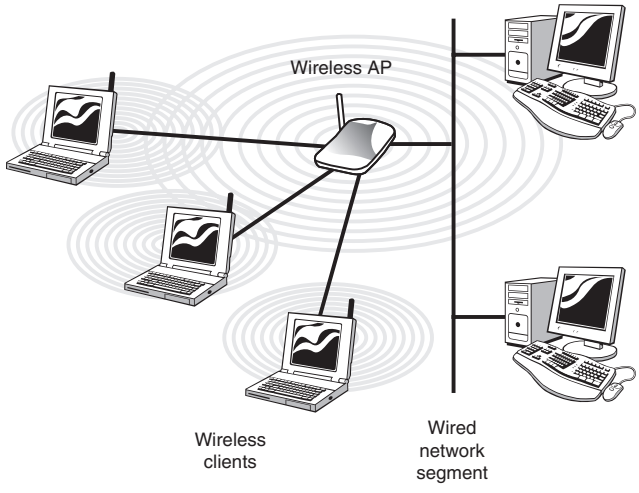

FIGURE 3.27 An infrastructure wireless network uses a WAP.

NOTE

A WAP for All Seasons Because wireless networks are sometimes deployed in environments other than inside a warm, dry building, some manufacturers offer rugged versions of WAPs. These devices are sealed against the elements, making them suitable for placement in locations where nonrugged devices would not survive. If you are implementing a wireless network, consider whether using these rugged devices may be warranted.

Transceivers

The term *transceiver* does not necessarily describe a separate network device but rather an integrated technology embedded in devices such as network cards. In a network environment, a transceiver gets its name from being both a transmitter and a receiver of signals, such as analog or digital. Technically, on a LAN the transceiver is responsible to place signals onto the network media and also detecting incoming signals traveling through the same cable. Given the description of the function of a transceiver, it makes sense that that technology would be found with network cards.

Although transceivers are found in network cards, they can be external devices as well. As far as networking is concerned, transceivers can ship as a module or chip type. *Chip transceivers* are small and are inserted into a system board or wired directly on a circuit board. *Module transceivers* are external to the network and are installed and function similarly to other computer peripherals, or they may function as standalone devices.

There are many types of transceivers: RF transceivers, fiber-optic transceivers, Ethernet transceivers, wireless (WAP) transceivers, and more. Though each of these media types is different, the function of the transceiver remains the same. Each type of the transceiver used has different characteristics such as the number of ports available to connect to the network and whether full-duplex communication is supported.

Listed with transceivers in the CompTIA objectives are media converters. Media converters are a technology that allows administrators to interconnect different media types—for example, twisted pair, fiber, and thin or thick coax—within an existing network. Using a media converter, it is possible to connect newer 100Mbps, Gigabit Ethernet, or ATM equipment to existing networks such as 10Base-T or 100Base-T. They can also be used in pairs to insert a fiber segment into copper networks to increase cabling distances and enhance immunity to electromagnetic interference (EMI).

Firewalls

Today, firewalls are an essential part of a network's design. A *firewall* is a networking device, either hardware or software based, that controls access to your organization's network. This controlled access is designed to protect data and resources from outside threat. To do this, firewalls are typically placed at entry/exit points of a network. For example, a firewall might be placed between an internal network and the Internet. After the firewall is in place, it can control access in and out of that point.

Although firewalls typically protect internal networks from public networks, they are also used to control access between specific network segments within a network. For example, you might place a firewall between the Accounts Department and the Sales Department.

As mentioned, firewalls can be implemented through software or through a dedicated hardware device. Organizations implement software firewalls through network operating systems (NOS) such as Linux/Unix, Windows servers, and Mac OS servers. The firewall is configured on the server to allow or permit certain types of network traffic. In small offices and for regular home use, a firewall is commonly installed on the local system and configured to control traffic. Many third-party firewalls are available.

Hardware firewalls are used in networks of all sizes today. Hardware firewalls are often dedicated network devices and can be implemented with very little configuration and protect all system behind it from outside sources. Hardware firewalls are readily available and often combined with other devices today. For example, many broadband routers and wireless access points have firewall functionality built in. In such a case, the router or WAP may have a number of ports available to plug systems into.

NOTE

More on Firewalls A complete discussion of firewalls is provided in Chapter 11, "Securing the Network."

Network Devices Summary

The information in this chapter is important for the Network+ exam. To summarize our coverage of network devices, we have placed some of the key points about each device in Table 3.3. You should learn this information well.

REVIEW BREAK

TABLE 3.3 Network Devices Summary

Device	Function/Purpose	Key Points
Hub	Connects devices on a Ethernet twisted-pair network.	A hub does not perform any tasks besides signal regeneration.
Switch	Connects devices on a twisted-pair network.	A switch forwards data to its destination by using the MAC address embedded in each packet.
Bridge	Connects LANs to reduce overall network traffic.	A bridge allows or prevents data from passing through it by reading the MAC address.
Router	Connects networks together.	A router uses the software-configured network address to make forwarding decisions.
Gateway	Translates from one data format to another.	Gateways can be hardware or software based. Any device that translates data formats is called a gateway.

(continues)

TABLE 3.3 *Continued*

Device	Function/Purpose	Key Points
CSU/DSU	Translates digital signals used on a LAN to those used on a WAN.	CSU/DSU functionality is sometimes incorporated into other devices, such as a router with a WAN connection.
Network card	Enables systems to connect to the network.	Network interfaces can be add-in expansion cards, PCMCIA cards, or built-in interfaces.
ISDN terminal adapter	Connects devices to ISDN lines.	ISDN is a digital WAN technology often used in place of slower modem links. ISDN terminal adapters are required to reformat the data format for transmission on ISDN links.
WAP	Provides network capabilities to wireless network devices.	A WAP is often used to connect to a wired network, thereby acting as a link between wired and wireless portions of the network.
Modem	Provides serial communication capabilities across phone lines.	Modems modulate the digital signal into analog at the sending end and perform the reverse function at the receiving end.
Transceiver	A device that can be both a transmitter and a receiver of signals.	A transceiver is a device that functions as a transmitter and a receiver of signals such as analog or digital.
Firewall	Provides controlled data access between networks.	Firewalls can be hardware or software based and are an essential part of a network's security strategy.

Identifying MAC Addresses

Objective:

2.1 Identify a MAC (Media Access Control) address and its parts.

This chapter many times refers to MAC addresses and how certain devices use them. However, it has not yet discussed why MAC addresses exist, how they are assigned, and what they consist of. Let's do that now.

TIP

A MAC Address Is the Physical Address A MAC address is sometimes referred to as a *physical address* because it is physically embedded in the interface. Sometimes it is also referred to as a *network address*, which is incorrect. A *network address* is the logical protocol address assigned to the network to which the interface is connected.

A MAC address is a 6-byte hexadecimal address that allows a NIC to be uniquely identified on the network. The MAC address forms the basis of network communication, regardless of the protocol used to achieve network connection. Because the MAC address is so fundamental to network communication, mechanisms are in place to ensure that there is no possibility of duplicate addresses being used.

To combat the possibility of duplicate MAC addresses being assigned, the Institute of Electrical and Electronics Engineers (IEEE) took over the assignment of MAC addresses. But rather than be burdened with assigning individual addresses, the IEEE instead decided to assign each manufacturer an ID and then let the manufacturer further allocate IDs. The result is that in a MAC address, the first three bytes define the manufacturer, and the last three are assigned by the manufacturer.

For example, consider the MAC address of the computer on which this book is being written: 00:D0:59:09:07:51. The first three bytes (00:D0:59) identify the manufacturer of the card; because only this manufacturer can use this address, it is known as the O*rganizational Unique Identifier (OUI)*. The last three bytes (09:07:51) are then referred to as the *Universal LAN MAC address*: They make this interface unique. You can find a complete listing of organizational MAC address assignments at `http://standards.ieee.org/regauth/oui/oui.txt`.

EXAM ALERT

MAC Address Tip Because MAC addresses are expressed in hexadecimal, only the numbers 0 through 9 and the letters A through F can be used in them. If you get a Network+ exam question about identifying a MAC address and some of the answers contain letters and numbers other than 0 through 9 and the letters A through F, you can discount those answers immediately.

You can discover the MAC address of the NIC in various ways, depending on what system or platform you are working on. Table 3.4 defines various platforms and the method you can use to view the MAC address of an interface.

TABLE 3.4 Methods of Viewing the MAC Addresses of NICs

Platform	Method
Windows 95/98/Me	Run the `winipcfg` utility
Windows 2000/2003/XP	Run `ipconfig /all` from a command prompt
Linux/some Unix	Run the `ifconfig -a` command
Novell NetWare	Run the `config` command
Cisco router	Run the `sh int <interface name>` command

Figure 3.28 shows the `ipconfig /all` command run on a Windows 2000 system. The MAC address is defined on the Physical Address line of the output.

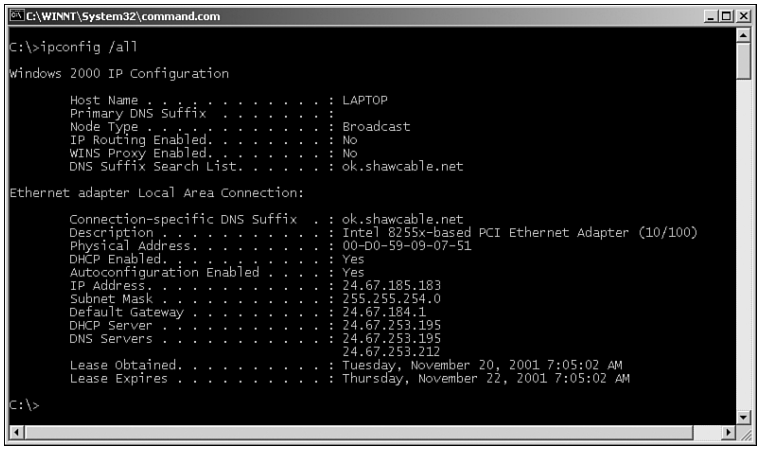

FIGURE 3.28 The output from the ipconfig /all command on a Windows 2000 system.

Challenge

You are concerned about security on your new wireless connection. As a result you decide to implant MAC level security on the wireless AP. You want to allow both of your client systems, one Linux and one Windows XP, to access the Internet through the wireless AP. Identify the procedures and command for obtaining the MAC address on both a Linux and Windows XP system.

Chapter Summary

Many devices are used to create networks. Every network except the simplest, single-segment coaxial networks uses one or more of these devices. Knowledge of the purpose of the devices discussed in this chapter is vital for the Network+ exam, as well as for the real world.

Hubs and switches provide a mechanism to connect devices to a network created with twisted-pair cabling. Switches offer a speed advantage over hubs because they can use full-duplex communications. They also create dedicated paths between devices, reducing the number of collisions that occur. Both hubs and switches are available in managed and unmanaged varieties.

Bridges allow network traffic to be confined to certain network segments, thereby reducing the amount of network traffic. On Ethernet networks, an additional benefit is reduced collisions.

Routers are devices that connect networks and thereby create internetworks. Because routers use software-configured network addresses instead of hardware-defined MAC addresses, they can provide more functionality than bridges. Routers either can be dedicated hardware devices or can be implemented through software on server systems.

A gateway is a device that translates from one data format to another; it can be a hardware device or a software application. A CSU/DSU is an example of a gateway: CSUs/DSUs translate from the data format used on LANs to that used on WANs. A modem, which translates a signal from digital to analog so that it can be transmitted across a conventional phone line, is another example of a gateway.

WAPs are a relative newcomer to the networking equipment field. Wireless network clients use WAPs to connect to the network. WAPs also generally have a connection point that lets them connect to a wired network infrastructure.

NICs are the point of connectivity between devices and the network. NICs can be add-in expansion cards, PCMCIA devices for laptops, or devices built in to the system board. When you install NICs, you must observe ESD best practices and also pay attention to hardware compatibility and bus compatibility issues.

In addition to NICs used to connect to a LAN, ISDN terminal adapters are sometimes used for remote connectivity.

When you're using clustering, a special system-area NIC is applied to network interfaces used to communicate clustering information between servers.

On a network, each NIC is identified by a unique MAC address. MAC addresses are assigned by the manufacturers that produce the devices, although the high-level assignment of addresses is managed and carried out by the IEEE.

If you get a chance to use all the hardware devices discussed in this chapter, count yourself lucky. Almost every environment will use some of them, but few use them all.

Key Terms

- 80/20 rule
- AT modem commands
- bridge
- CSU/DSU
- cut-through
- distance-vector protocols
- fragment-free
- gateway
- hub
- IPX/SPX
- ISDN adapter
- link-state protocol
- MAC address
- MDI
- MDI-X
- modem
- multiprotocol routing
- NIC
- rack mounting
- RIP
- router
- source-route bridge
- split horizon
- split horizon with poison reverse
- STA
- store-and-forward
- STP
- switch
- TCP/IP
- translational bridge
- transparent bridge
- UART chips
- uplink port
- WAP

Apply Your Knowledge

Exercises

3.1 Determining MAC Addresses for Network Cards

This chapter identifies the characteristics and functions of network devices. In an ideal world, this project would require hands-on experience with these devices, but this is not an ideal world, and access to this equipment is not always easy. Therefore, we will include two exercises that you might be required to perform if such devices are used on your network.

This project assumes that you are using Windows 2000 Server or Professional. You will need an installed NIC with a working driver.

Estimated time: 20 minutes

1. Open a command window by selecting Start, Run. In the command box, type **`cmd`** and then click OK.
2. At the command prompt, type **`ipconfig /all`**. The MAC address of your NIC is displayed in the Physical Address line.
3. Open a web browser and go to the following website: http://standards.ieee.org/regauth/ oui/oui.txt.
4. Using the Find functionality in your Web browser, locate the entry that corresponds with the address of your NIC. Is the manufacturer of your NIC the company you expected it to be? Some NIC manufacturers re-brand cards manufactured by another company. For that reason, the MAC address may correspond to a manufacturer that is different from the brand name of the card.

3.2 Using the `tracert` Utility to View the Path to an Internet Destination

One of the tools network administrators have at their disposal is the `tracert` utility. `tracert` allows you to see the hops a network packet takes to get to its destination. At each point along the way, the packet gives information about the route it is taking, along with details of the routers it crosses. More information on the `tracert` utility is provided in Chapter 14, "Troubleshooting Network Connectivity."

NOTE

Firewalls If you are using a system protected by a firewall system, this exercise might not work because firewalls are commonly configured to block `tracert` traffic.

In this exercise, you use the `tracert` utility to view the path to an Internet destination. This project assumes that you are using Windows 2000/2003 or Windows XP with a working Internet connection.

Estimated time: 5 minutes

1. Open a command window by selecting Start, Run. In the command box, type **`cmd.exe`** and then click OK.
2. At the command prompt, type **`tracert www.novell.com`**. The route to the Novell web server is displayed.
3. How many hops are you from the destination?

In addition to performing a `tracert` on a remote location, you can use the command on your internal network. If you have a small network set up with a wireless AP or a hub/switch, try doing this exercise using these network devices.

Exam Questions

1. You have configured a 100Mbps network connection between your computer and the switch as half-duplex. What will be the maximum speed of the connection?

   - ❍ **a.** 50Mbps
   - ❍ **b.** 100Mbps
   - ❍ **c.** 100MBps
   - ❍ **d.** 200Mbps

2. You want to create a larger network by connecting two switches together. One of the switches has a port that can be switched from MDI to MDI-X as needed. The other switch doesn't have such a port or a dedicated uplink port. Which type of cable should you use, and how should you configure the switchable port to create the larger network?

   - ❍ **a.** Use a straight-through cable and set the port to MIDI
   - ❍ **b.** Use a crossover cable and set the port to MDI-X
   - ❍ **c.** Use a straight-through cable and set the port to MDI-X
   - ❍ **d.** Use a crossover cable and set the port to MDI

3. Of the following, which represents a valid MAC address?

   - ❍ **a.** 00:D0:59:09:07:51
   - ❍ **b.** 000:D00:599:099:071:512
   - ❍ **c.** 00:D0:59:09:07:51:C4:56
   - ❍ **d.** 00:H0:59:09:07:51

4. A bridge makes forwarding decisions based on what information?

   - ❍ **a.** IP address
   - ❍ **b.** MAC address
   - ❍ **c.** Binary address
   - ❍ **d.** IRQ address

5. What information does a switch use to determine the port to which data should be sent?

   - ❍ **a.** The IP address of the connected device
   - ❍ **b.** The priority of the connected device
   - ❍ **c.** The MAC address of the connected device
   - ❍ **d.** The Ethernet address of the connected device

6. Which of the following is a link-state routing protocol used on TCP/IP networks?

   ❍ **a.** RIP

   ❍ **b.** ARP

   ❍ **c.** OSPF

   ❍ **d.** NLSP

7. On a Windows 2000 system, what command would you use to view the MAC address?

   ❍ **a.** `ifconfig -a`

   ❍ **b.** `ipconfig /all`

   ❍ **c.** `ipconfig`

   ❍ **d.** `config /all`

8. What is the purpose of the uplink port on a hub or switch?

   ❍ **a.** It allows for satellite connections.

   ❍ **b.** It allows hubs or switches to be connected together.

   ❍ **c.** It allows computers to connect to the device.

   ❍ **d.** It provides a spare port, which can be used if another port fails.

9. By what method does a router determine the destination address for a packet?

   ❍ **a.** It looks at the MAC address of the sender.

   ❍ **b.** It looks for the MAC address of the destination.

   ❍ **c.** It looks for the software-configured network address for the destination.

   ❍ **d.** It looks at the FCS field of the packet.

10. Which of the following statements best describes split horizon?

    ❍ **a.** Routes are advertised back on the interface from which they were learned, with a metric of 16.

    ❍ **b.** Routes are advertised back on the interface from which they were learned, with a metric of 0.

    ❍ **c.** Routes are not advertised back on the interface from which they were learned.

    ❍ **d.** Routes are advertised back on the interface from which they were learned, with a metric of 16, and on all other interfaces they are advertised back on the interface from which they were learned, with a metric of 0.

11. In a network that uses distance-vector routing protocols, what information is included in the update sent out by each router?

- ❍ **a.** Details of the routers to which it is directly managed by
- ❍ **b.** A map of the entire network, with hop counts valued from its current position
- ❍ **c.** Details of all the routers it knows about
- ❍ **d.** Details of its own configuration

12. What is the difference between an active hub and a passive hub?

- ❍ **a.** An active hub has management capabilities.
- ❍ **b.** An active hub forwards the data only to the ports that need it.
- ❍ **c.** An active hub channels bandwidth to a given connection if the connection becomes too slow.
- ❍ **d.** An active hub regenerates the signal before forwarding it.

13. What condition can arise if routers advertise a route back to the router from which it was learned?

- ❍ **a.** Count to infinity
- ❍ **b.** Road to nowhere
- ❍ **c.** Loop de loop
- ❍ **d.** Count to 16

14. What term is used by routers to describe each step necessary to reach a destination?

- ❍ **a.** Hop
- ❍ **b.** Jump
- ❍ **c.** Skip
- ❍ **d.** Leap

15. What is the maximum speed of a 16550 UART chip?

- ❍ **a.** 64,000bps
- ❍ **b.** 115,200bps
- ❍ **c.** 430,800bps
- ❍ **d.** 921,600bps

16. What is the name of the bridging method used to segregate Ethernet networks?

- ❍ **a.** Source-route
- ❍ **b.** Invisible
- ❍ **c.** Cut-through
- ❍ **d.** Transparent

17. Which of the following is a distance-vector routing protocol used on TCP/IP networks?

- ❍ **a.** ARP
- ❍ **b.** NLSP
- ❍ **c.** OSPF
- ❍ **d.** RIP

18. A CSU/DSU is used in which of the following network configurations?

- ❍ **a.** When converting from a Token Ring network to an Ethernet network
- ❍ **b.** When converting a digital signal to an analog signal
- ❍ **c.** When converting from the digital signals used on a LAN to the digital signals used on a WAN
- ❍ **d.** When converting from the digital signal format used on a LAN to the analog signal format used on a WAN

19. A router makes its forwarding decisions based on which of the following information?

- ❍ **a.** IP address
- ❍ **b.** ARP address
- ❍ **c.** Binary address
- ❍ **d.** Frame address

20. You are tasked with upgrading a new NIC in the company file and print server. Which of the following should you determine before buying a replacement card? (Choose the three best answers.)

- ❍ **a.** Bus compatibility
- ❍ **b.** Network compatibility
- ❍ **c.** Hardware compatibility
- ❍ **d.** Cooling requirements

Answers to Exam Questions

1. **b.** A half-duplex connection operates at the normal speed of the link. Thus, a 100Mbps network connection in a half-duplex configuration would operate at a maximum of 100Mbps. All the other answers are invalid. For more information, see the section "Working with Hubs and Switches" in this chapter.

2. **b.** Because one of the switches does not have MDI capability, the switchable port should be set to MDI-X. Then, a crossover cable should be used to cancel out the crossing between the two devices. None of the other options would result in a successful connection. For more information, see the section "Working with Hubs and Switches" in this chapter.

3. **a.** A MAC address comprises 6 bytes presented in a hexadecimal format. The letters A through F and numbers 0 through 9 are the only valid characters. Therefore, all the other answers provided are incorrect. For more information, see the section "Identifying MAC Addresses" in this chapter.

4. **b.** Bridges make forwarding decisions based on the destination MAC address embedded in each packet. Routers use software addresses, such as IP addresses, to make forwarding decisions. Answers c and d are not valid options. For more information, see the section "Bridges" in this chapter.

5. **c.** A switch uses the MAC address of the connected device to determine the port to which data is forwarded. Routers use software addresses, such as IP addresses, to make forwarding decisions. Answer b is not valid. Although there are many addressing schemes used on networks, *Ethernet address* is not a valid term. Therefore, Answer d is incorrect. For more information, see the section "Switches" in this chapter.

6. **c.** OSPF is a link-state routing protocol used on TCP/IP networks. RIP is a distance-vector routing protocol used on both TCP/IP and IPX/SPX networks; ARP is a component of the TCP/IP protocol suite. NLSP is a link-state routing protocol used on IPX/SPX networks. For more information, see the section "Routers" in this chapter.

7. **b.** The `ipconfig /all` command shows a range of network-related information, including the MAC addresses of any installed NICs. Answers a and d are valid, and using the `ipconfig` command without the `/all` switch shows limited information. For more information, see the section "Identifying MAC Addresses" in this chapter.

8. **b.** The uplink port can be used to connect hubs and switches together, using a standard twisted-pair cable. All the other answers are invalid. For more information, see the section "Working with Hubs and Switches" in this chapter.

9. **c.** Routers use the software-configured network address to make routing decisions. Bridges use MAC addresses to make decisions. Answer d is not valid. The FCS (that is, frame checksum) field is used for error detection. For more information, see the section "Routers" in this chapter.

10. **c.** Split horizon is a routing algorithm that dictates that routes are not advertised back on the interface from which they were learned. Answer a describes the operation of the split horizon with poison reverse algorithm. None of the other answers are valid. For more information, see the section "Routers" in this chapter.

11. **c.** In a network that uses distance-vector routing protocols, routers advertise details of the routers they know about. These updates are sent to all the neighbor routers. Answer a describes the actions on a link-state-based network. Answers b and d are invalid. For more information, see the section "Routers" in this chapter.

12. **d.** An active hub regenerates the data signal before forwarding it to all connected devices. Active hubs come in both managed and unmanaged varieties. Answer b describes the action of a switch. Answer c is invalid. For more information, see the section "Hubs" in this chapter.

13. **a.** A count to infinity occurs when two routers provide information on the same destination and so create a routing loop. All the other answers are invalid. For more information, see the section "Routers" in this chapter.

14. **a.** Each step in the path between a router and its destination is called a hop. The other terms are not used in networking. For more information, see the section "Routers" in this chapter.

15. **b.** A 16550 UART chip is capable of speeds up to 115,200bps. None of the other answers represent the speed for the 16550 UART chip. For more information, see the section "Modems" in this chapter.

16. **d.** The bridging method used on Ethernet networks is called *transparent* because the other network devices are unaware of the existence of the bridge. Source-route bridges are used on Token Ring networks, invisible is not a type of bridge, and cut-through is a switching method, not a type of bridge. For more information, see the section "Bridges" in this chapter.

17. **d.** RIP is a distance-vector routing protocol used on TCP/IP networks. ARP is a component of the TCP/IP protocol suite. NLSP is a link-state routing protocol used on IPX networks, and OSPF is a link-state routing protocol used on TCP/IP networks. For more information, see the section "Routers" in this chapter.

18. **c.** CSUs/DSUs are used to convert the digital signals used on a LAN to the digital signals used on a WAN. The process described in Answer a would be performed by a gateway, and the process described in Answer b would be performed by a modem. Answer d is not valid because WANs commonly use digital signals. For more information, see the section "CSU/DSU" in this chapter.

19. **a.** Routers make routing decisions based on the software-configured network address, which is protocol dependent. There is no such thing as an ARP address. Answers c and d are invalid. For more information, see the section "Routers" in this chapter.

20. **a, b, c.** You should verify bus compatibility, network compatibility, and hardware compatibility before you buy a new NIC. You do not typically need to concern yourself with cooling requirements of a component. For more information, see the section "Network Interface Cards (NICs)" in this chapter.

Suggested Readings and Resources

1. Olexa, Ron. *Implementing 802.11, 802.16, and 802.20 Wireless Networks: Planning, Troubleshooting, and Operations* (Communications Engineering). Newnes Publishing, 2004.
2. Computer networking products and information, www.alliedtelesyn.com.
3. Computer networking device information, www.3com.com.
4. "Computer Networking Tutorials and Advice," compnetworking.about.com.
5. "TechEncyclopedia," www.techencyclopedia.com.
6. "Networking Technology Information from Cisco," www.cisco.com/public/products_tech.shtml.