

Chapter 3: Networking Hardware

In this chapter

- Working with network interface cards
- Selecting and installing a NIC
- Using hubs
- Working with PC motherboards
- Understanding processors and PC RAM
- Working with hard drives
- Differentiating server and client hardware

Our Age of Anxiety is, in great part, the result of trying to do today's jobs with yesterday's tools.

–**Marshall McLuhan**

Now that we've discussed the different kinds of networks and looked at network topologies, we should spend some time discussing the hardware involved in networking. This chapter will concentrate on the connectivity devices that define the network topology—the most important being the network interface card. We will also take a look at hubs, routers, and switches.

Another important aspect of building your network is selecting the hardware for your client PCs and your network servers. There are many good primers on computer hardware—for example, the *Absolute Beginner's Guide to PC Upgrades*, published by Que. Also, numerous advanced books, such as *Upgrading and Repairing PCs* (by Scott Mueller, also from Que), are available, so we won't cover PC hardware in depth in this chapter. We will take a look at motherboards, RAM, and hard drives because of the impact these components have on server performance. We will also explore some of the issues related to buying client and server hardware.

Let's start our discussion with the network interface card. We can then look at network connectivity devices and finish up with some information on PC hardware.

Working with Network Interface Cards

As far as networking is concerned, the network interface card (NIC) is one of the most important PC devices. Each computer on the network (including servers and clients) requires one, and it is the NIC that provides the connection between the PC and the network's physical medium (such as copper or fiber-optic cable, both of which are discussed in the next chapter).

An IBM PC (or compatible PC) typically requires that a NIC be added to it (unless you buy your network clients and servers with a NIC as their standard hardware configuration). Even Macintoshes and Apple PowerPCs require NICs. In most cases, newer Apple computers, such as the iMac, have built-in Ethernet NICs. [Figure 3.1](#) shows a PCI NIC from 3Com, which provides the RJ-45 female port for the

CAT 5 twisted pair that connects the computer to a hub or switch (more about these connectivity devices later in the chapter).

Not only does the NIC provide the computer with a connection to the network, but it also handles an important data-conversion function. Data travels in parallel on the PC's bus system; the network medium demands a serial transmission. The transceiver (a transmitter and receiver) on the NIC card is able to move data from parallel to serial, and vice versa. This is not unlike automobiles moving on a multi-lane superhighway that must all merge into one lane of traffic.

FIGURE 3.1

Network interface cards provide the connection between the PC and the network cabling.

Network interface cards also supply the basic addressing system used to get data from one computer to another on the network. This hardware or MAC address is burned into a ROM chip on the NIC. It is referred to as the *MAC address* because the Media Access Control (MAC) layer is actually a sublayer of the OSI model's Data Link layer (the OSI model is discussed in detail in Chapter 5, "Network Protocols: Real and Imagined," in the section "Understanding the OSI Model").

Tip - The hardware address on the NIC actually provides the definitive address for a computer on the network. Different protocol stacks, such as TCP/IP, use a logical addressing system (in the case of TCP/IP, it is IP addresses). In this case, the logical address must be resolved to the NIC's hardware address before the data can actually be received.

Selecting a NIC

NICs come in different types, depending on the network architecture being used on the network (such as Ethernet or token ring, which are both discussed in the next chapter). NICs also differ in the type of motherboard slot they fit into. PCI NICs fit into a PCI slot on a motherboard. ISA NICs fit into an ISA slot. [Figure 3.1](#) shows a PCI card (PCI cards are shorter than ISA cards, as are their respective slots on the motherboard).

When you select a NIC, it needs to be compatible with the architecture of the network. The Ethernet network is, by far, the most common network architecture type and is used for both peer-to-peer networks and large-scale corporate networks. Ethernet connectivity devices, such as hubs, are also much cheaper than the connectivity devices deployed on a token-ring network.

Tip - Token-ring networks are typically found in larger corporations. The token-ring architecture does offer some advantages over Ethernet (particularly on high-traffic networks, where bandwidth is at a premium), which we discuss in the next chapter.

The NIC that you select must also fit in an empty expansion slot on your computer's motherboard. On an IBM or compatible PC, you have a number of choices, such as an ISA, PCI, or EISA slot. The slots available will depend on the motherboard (motherboards are discussed later in the chapter). [Figure 3.2](#) shows a motherboard that provides both ISA and PCI expansion slots.

FIGURE 3.2

Your motherboard will typically provide ISA and PCI expansion slots.

The Industry Standard Architecture (ISA) bus (a *bus* is the internal wiring system on the motherboard) was the standard used on the IBM PC and the IBM XT class of personal computers. This slot type supported 8-bit throughput for expansion devices. The ISA bus speed was increased to 16 bits in 1984. A large number of devices, such as modems, sound cards, and NICs, are available for this bus type. Even some new motherboards provide ISA slots. However, a day will come when the ISA slot is considered obsolete and only the higher-throughput PCI slots will be available on motherboards.

The PCI bus, on the other hand, was introduced by Intel in 1993 and has really become the industry standard. PCI provides a 32-bit bus, and PCI version 2.1 supports 64-bit and faster bus specifications. As with ISA, PCI cards support a large number of peripheral devices.

Tip - There is a slight chance you might run into three other computer bus architectures: EISA, VESA, and MCA. Extended Industry Standard Architecture (EISA) was developed by Compaq and other PC clone manufacturers as a suggested PC slot standard (it never became the industry standard) when IBM rolled out PCs embracing Micro Channel Architecture (which is now nothing more than a footnote in bus architecture literature). EISA provided 32-bit bus throughput as well as ISA compatibility. The Video Electronics Standards Association (VESA) bus provided support for 486 processors. This 32-bit bus system has been abandoned by motherboard manufacturers in favor of the PCI bus. Micro Channel Architecture (MCA) was a new bus system launched by IBM for its PS/2 family of computers. MCA was not compatible with ISA, and the availability of MCA peripherals was limited when the PS/2 family was first introduced. MCA expansion cards were also extremely expensive. The MCA bus was eventually abandoned by IBM in favor of the PCI bus.

Although ISA NICs are generally going to be cheaper than PCI NICs, PCI NICs are obviously going to provide better overall performance because of their potential for faster throughput. You can get some other enhancements with a NIC that improve the transfer of data from the NIC to the computer's processor:

- *Buffering*. Memory chips are actually placed on the NIC. This RAM is used as a buffer. It can hold data that is awaiting services provided by the computer's processor. It can also hold data waiting to be broadcast out onto the LAN medium.
- *Direct Memory Access (DMA)*. Computers that support DMA allow devices such as network interface cards to transfer and accept data directly from the computer's RAM. This means that the processor does not have to become involved in the process of moving data to or from the NIC and RAM.
- *Bus mastering*. NICs can be designed to directly access computer RAM without using the computer's processor as an intermediary. Bus mastering (which is similar to DMA) allows the NIC to control the bus and send and receive data to and from the computer's RAM.

NICs that provide buffering and support DMA or bus mastering are going to be more expensive than

NICs that do not. However, when selecting a NIC for a mission-critical server, you will probably want to buy one with all the performance bells and whistles.

Installing a NIC

You certainly don't need an engineering degree to install a network card in a PC. However, before you tackle the job, you need to be aware of one thing: You need to have a handle on a subject area that an old colleague of mine referred to as "static awareness." Static electricity can actually produce enough voltage to ruin a motherboard or expansion card, such as a NIC. This means that you should use some sort of antistatic wristband that attaches to the PC case or do your work while standing on an antistatic mat. The bottom line is that you should avoid working on the computer in a room that is carpeted.

You will also want to be sure you have a couple of decent screw drivers. These screw drivers should not have magnetized tips. Now, here are the steps to follow:

1. Open up the case on the computer and install the NIC in an open expansion slot.
2. Close up the case and attach the network medium (typically twisted-pair cabling).
3. Boot up the computer. If you purchased a plug-and-play network card and are using an OS that supports plug-and-play technology, the most you will have to do is provide a disk or CD-ROM that contains the driver for the NIC.
4. If you are using an operating system that does not detect new hardware devices, you will have to manually install the NIC. If the card came with installation software, use that software to install the necessary drivers.

Because most new computers and the operating systems that they run embrace plug-and-play technology, installing a NIC in a computer is really just a matter of placing the NIC in an empty motherboard slot and then restarting the computer. Plug-and-play NICs will be recognized by the operating system and the appropriate software driver will be loaded on the system. [Figure 3.3](#) shows the Windows 2000 Professional OS identifying a plug-and-play NIC.

FIGURE 3.3

Operating systems that embrace plug-and-play technology make it easy to add a plug-and-play NIC to the system.

Not all operating systems embrace plug-and-play technology, however, and these types of operating systems will require that you select an IRQ (or *Interrupt Request*) for the new NIC. Each device in your computer, including the mouse, keyboard, and NIC, is assigned an Interrupt Request line that the device uses to alert the microprocessor (CPU) that it needs data processed. Each device must be assigned a unique IRQ; otherwise, you will have an IRQ conflict. Although PCI motherboards do allow devices to share an IRQ, common practice is to have a unique IRQ for every device on the system.

Finding the available IRQs is not that difficult, and each operating system (including PC operating systems and network operating systems) provides you with a tool to view the used and available IRQs on the system. For example, [Figure 3.4](#) shows the Windows 2000 Device Manager, which displays the IRQs on the system.

Tip - If you are using Windows 2000 Server, getting to the Device Manager is very straightforward. Right-click the My Computer icon on the desktop and then choose Properties. In the Properties dialog box, select the Hardware tab. Then click the Device Manager button.

FIGURE 3.4

An operating system provides a utility that allows you to examine the IRQ configuration of a system.

Table 3.1 shows the standard IRQ settings for a PC. As you can see, a number of the IRQs are reserved for particular system devices.

Table 3.1 IRQ Settings

IRQ	Use
0	System timer
1	Keyboard
2	Cascade to secondary IRQ controller
3	COM port 2 and 4 (serial port)
4	COM port 1 and 3 (serial port)
5	LPT2 (printer port)
6	Floppy disk controller
7	LPT1 (printer port)
8	Real-time clock
9	Free
10	Primary SCSI adapter (or free)
11	Secondary SCSI adapter (or free)
12	PS/2 mouse
13	Floating-point math coprocessor
14	Primary hard disk controller
15	Secondary hard disk controller (or free)

Obviously, in cases where the computer does not have a second COM port or an LPT2 port, these IRQs will be available. Each computer's hardware configuration will vary, so you will have to determine the available IRQs for each computer in which you are installing a NIC.

Tip - Most operating systems do a pretty good job of helping you get a newly installed NIC up and running. Windows 2000 Server and Windows 2000 Professional both embrace Microsoft's Plug and Play scheme for plug-and-play hardware devices. This means that both of these operating systems will, in most cases, identify and install the appropriate drivers for a number of the network interface cards available on the market. And although you can't

call what they do "plug and play," Novell NetWare 5, Sun Microsystems Solaris OS, and the various Linux clones do a pretty good job of helping you set up the appropriate network card in your network server or client machines.

Working with Network Connectivity Devices

Depending on the type of topology your network uses and the type of cabling you use (cabling is discussed in Chapter 4, "Building the Network Infrastructure," in the section "Choosing a Network Connectivity Strategy"), your LAN might require some sort of connectivity device to connect the various network computers, printers, and other devices together. In cases where you need to extend your LAN (say, to the second floor of an office building) or add a large number of new users to the LAN, other connectivity devices might be required. Some of these connectivity devices merely serve to connect devices; others are used to boost the data signal traveling on the network medium, and still others actually participate in determining how data traffic should flow on the network.

Let's start our discussion of network connectivity devices with the hub, which is a device you would use on a small network, or even in a peer-to-peer networking situation, to connect computers. The other devices that we will look at, such as repeaters, switches, and routers, are often lumped under the blanket term *internetworking devices*. An *internetwork* is a network of LANs, meaning that some sort of connectivity technology is used to extend a LAN beyond its typical size or to connect different LANs together into one large network.s

Hubs

Hubs are commonly used LAN connectivity devices. They serve as the central connection points for LANs (hubs are used on LANs that embrace the star topology discussed in Chapter 2, "Different Needs, Different Networks"). A basic hub contains no active electronics and therefore cannot be used to extend a LAN (that is, extend it past the cabling distance specifications discussed in the next chapter). It basically organizes your cables and relays data signals to all the computers on the LAN.

Hubs are most commonly used on networks that use twisted-pair cabling. Ports available on the back of the hub provide the connection points for the devices on the network. Computers and other devices are attached to the hub by individual network cables. Hubs come in many sizes and shapes and supply different numbers of ports.

In cases where the LAN outgrows the size of the hub, a new hub can be attached (the hubs are "daisy chained" together using a short connection cable often referred to as a *rat tail*) to the current hub, thus providing greater port density. [Figure 3.5](#) shows two different types of Ethernet hubs.

FIGURE 3.5

Hubs provide the central connection points for LANs.

Hubs come in all sizes and shapes and are available in a wide range of prices. Typically, the more ports on the hub, the more expensive the hub. Hubs that support faster varieties of Ethernet, such as Fast Ethernet (which we discuss in the next chapter in the section "Understanding Ethernet"), will also cost more.

Repeaters

As you will find out in the next chapter, the different types of network cabling all have a maximum distance that they can move a data signal. In cases where a LAN must be extended beyond the maximum run for a particular cabling type, repeaters are used. A *repeater* takes the signal that it receives from computers and other devices on the LAN and regenerates the signal so that the signal maintains its integrity along a longer media run than is normally possible.

Repeaters don't have any capabilities for directing network traffic or deciding the particular route certain data should take; they are simple devices that just sit on the network, boosting the data signals they receive. The problem with repeaters is that they amplify the entire signal that they receive, including any line noise. So, in worst-case scenarios, they pass on data traffic that is barely discernable from the background noise on the line.

Bridges

A *bridge* is an internetworking device used to help conserve the bandwidth available on the network. When LANs really start to grow, network data traffic can begin to overwhelm the available bandwidth on network media.

One strategy for conserving network bandwidth is to chop the network up into smaller segments. These segments are connected to a bridge. Bridges are smarter than hubs and repeaters and actually use some software to help get the job done. A bridge is able to read the MAC address (also known as the *hardware address*—remember it's burned onto the NIC in each computer on the network) of each data packet circulating on the network segments connected to the bridge. By learning which MAC addresses live on each of the network segments, the bridge can help keep data traffic that is local to a particular segment from spreading to the other network segments that are serviced by the bridge.

Switches

A *switch* is another internetworking device used to manage the bandwidth on a large network. A switch, which is often referred to as a "bridge on steroids," controls the flow of data by using the MAC address that is placed on each data packet (which coincides with the MAC address of a particular computer's network card). Switches divide networks into what are called *Virtual LANs* or *VLANs*. The great thing about a VLAN, which is a logical grouping of computers on the network into a sort of communication group, is that the computers don't have to be in close proximity or even on the same floor. This allows you to group computers that serve similar types of users into a VLAN. For example, even if your engineers are spread all over your company's office building, their computers can still be made part of the same VLAN, which would share bandwidth.

Switches use a combination of software and hardware to switch packets between computers and other devices on the network. Switches have their own operating system. [Figure 3.6](#) shows the status of a VLAN (VLAN1) on a Cisco 2900 switch. Understanding what is being shown in this figure requires an understanding of the switch's OS. Basically, this particular screen shows the hardware/MAC address of the switch and the IP address of the switch. Other statistics relate to the number of packets sent and received by the switch (which are all at 0 because the switch has just been placed on the network).

Because switches can offer a high density of connection ports, they can even replace the hub on a network. This means that each computer on the network can be connected to its own port on the switch. When PCs are directly connected to a switch, it can supply each PC with a dedicated amount of bandwidth. For example, users on a 10Mbps Ethernet network can realize bandwidth of 10Mbps. The computers don't compete for the bandwidth the way computers do on a network that is connected via a

hub.

FIGURE 3.6

An operating system is used to configure and monitor a switch.

Switch hardware can also take advantage of full-duplex access to the network media, which allows for the sending and receiving of data simultaneously on the network. This provides network access on an Ethernet network that would essentially be collision free (Ethernet networks experience data collisions pretty much as a rule; Ethernet is discussed in the next chapter). A computer on a Fast Ethernet network, which runs at 100Mbps, would actually realize a net total of 200Mbps throughput because sending and receiving can take place simultaneously on the full-duplex media.

Switches (because of the aforementioned reasons) are becoming very popular on large networks. They have all but replaced bridges as the internetworking devices for conserving network bandwidth and expanding LANs into larger corporate internetworks.

Routers

Routers are even smarter than bridges and switches (routers operate at the Network layer—a higher level in the OSI conceptual model than bridges and switches, which operate at the Data Link layer; we discuss all the layers of OSI in Chapter 5). A *router* uses a combination of hardware and software to actually "route" data from its source to its destination. (By *software*, I mean an operating system.) Routers actually have a very sophisticated OS that allows you to configure their various connection ports. You can set up a router to route data packets from a number of different network protocol stacks, including TCP/IP, IPX/SPX, and AppleTalk (protocols are discussed in Chapter 5).

Routers are used to segment LANs that have become very large and congested with data traffic. Routers are also used to connect remote LANs together using different WAN technologies.

[Figure 3.7](#) shows a Cisco 2516 router. This router has a built-in hub and three different WAN connection points. Notice the BRI port marked in the figure. *BRI* stands for ISDN *Basic Rate Interface*, which allows this router to connect to a remote network using an ISDN connection (ISDN and other WAN technologies are discussed in Chapter 10, "Expanding a LAN with WAN Technology").

FIGURE 3.7

Routers are used to segment networks into logical subsets.

Routers divide large networks into logical segments called *subnets*. This division of the network is based on the addressing scheme the network uses, such as IP addresses. Data traffic related to a particular subnet is kept local. The router only forwards data that is meant for other subnets on the extended network. This routing of network data helps conserve network bandwidth.

Routers decide how to forward data packets to their destinations based on a *routing table*. Routers use protocols built into their operating system to identify neighboring routers and their network addresses (such as IP addresses). This allows routers to build a routing table. [Figure 3.8](#) shows the command-line interface used on a Cisco router. This figure also shows the IP routing table for a small network that consists of two connected Cisco routers. Each of the subnets shown at the bottom of the table (notice the list of IP addresses) represents a different router interface. 10.2.0.0 and 10.3.0.0 are on the router that supplied this screen. The subnets 10.1.0.0 and 132.1.0.0 were discovered by the router (using the RIP protocol) on a connected router.

Tip - No matter what operating system you use, computers use pretty much the same strategy to identify themselves on the network. This strategy involves *broadcast messages*, meaning that a message will go out to every other computer on the network, proclaiming a computer's identity or requesting information related to the identity of another computer on the network. Broadcast messages create what is called *broadcast traffic*, and broadcast traffic can suck up valuable network bandwidth that otherwise could be used to move LAN data. Routers keep broadcast traffic from spreading beyond a particular segment, thus conserving a lot of potential network bandwidth.

FIGURE 3.8

Routers build and use a routing table to determine where data packets should be forwarded.

Some Thoughts on PC Hardware

Now that we've discussed NICs and some of the connectivity devices you will run into when you network computers, we will finish out the chapter by looking at some other critical pieces of PC hardware—specifically, motherboards, processors, hard drives, and RAM. We need to discuss these hardware components because they inherently affect a PC's overall performance. Understanding how these hardware components fit into the overall PC hardware puzzle will help you when it is time for you to select hardware for network clients or network servers. Since this book serves as a primer to networking and network hardware, you should consult a heavy-duty reference such as *Upgrading and Repairing PCs* by Scott Mueller, published by Que.

First, let's take a look at some issues related to motherboards. Then we can take a look at processors, hard drives, and memory.

Motherboards

The *motherboard* is the main system board for a PC, and it provides the data highway (or *bus system*) that moves data between components on the motherboard and the PC's processor. The speed of data moving along the bus is measured in megahertz (MHz). Although 66MHz was the standard for PCs running a Pentium (or Pentium clone) processor, motherboards can now operate as fast as 100 to 133MHz (and faster bus systems from a number of motherboard manufacturers are just beyond the horizon). A motherboard is specific to the type of processor run on the computer. An IBM/PC clone motherboard using an Intel processor will be designed differently from an Apple iMac motherboard, which uses a Motorola processor.

Tip - Motherboards have different designations, such as ATX and LPX. These designations refer to the footprint of the motherboard as well as the overall design of the motherboard and will therefore determine certain hardware choices.

The motherboard also provides all the slots for the computer's processor and memory. Areas to connect hard drives (and drives such as CD-ROM drives) with ribbon cables are also present on the motherboard, as are expansion slots (PCI and ISA) for other devices, such as sound cards and network

interface cards.

Motherboards are also characterized by the type of connection provided for the processor. A Socket 7 motherboard uses a processor that is typically a small square attached to a socket on the motherboard with a bunch of little pins. Intel, AMD, and Cyrix make processors for Socket 7 motherboards.

Slot 1 and Slot 2 motherboards provide a larger "slot" for the processor, and the slot processor is typically a rectangle (it looks like a very small brick). [Figure 3.9](#) shows an Abit slot-1 IBM/compatible motherboard (Abit is the manufacturer of the motherboard).

FIGURE 3.9

Motherboards come in a number of different configurations and sizes.

A number of the Slot 1 motherboards available on the market also supply space for dual processors. This makes the Slot 1 motherboard appropriate for high-end client workstations or network servers.

Although how fast a computer can work with data is really ultimately decided by the computer's processor, deciding on a particular motherboard should probably be determined, in part, by the features the motherboard's chipset provides. The chipset determines the ceiling for the bus speed. Chipsets provide DMA support (discussed earlier in this chapter in the section "Selecting a NIC") and other features, such as an Accelerated Graphics Port (AGP provides faster video response). These features will also boost the overall performance of the computer.

Processors

The processor is really the brain of the computer. It processes binary information input by the user or received from other devices on the computer, such as the network interface card. Processor speed is measured in megahertz (MHz). The Intel processor used on the original IBM PC ran at 8MHz.

Processors are now available that run in excess of 1,000MHz (that is, 1GHz or *gigahertz*). [Figure 3.10](#) shows an Intel Celeron 400MHz processor (on the left) for use on a Socket 7 motherboard and a Slot 1 Intel 300MHz Pentium II processor.

FIGURE 3.10

Processors are built to fit into a particular place on a motherboard.

Processors are manufactured by a number of different companies, including Intel, Motorola, Cyrix, and AMD. The selection of the processor type and speed, as far as networking is concerned, will depend on whether you are configuring a client machine or a network server. Faster (and even multiple) processors are a must on a server that needs to process a large number of calls for data from users on the network.

Basically, you should choose the processor you want to use for a computer and allow that choice to determine the motherboard you use for the computer. This helps narrow the choice of motherboards down to those that support the processor.

Memory

There seems to be as many memory types for personal computers as there are types of apples. Computer memory or *RAM (Random Access Memory)* is the working storage area. It is used by the processor and other devices to temporarily store information, and it's also accessed by software as we work on our computers.

You can't just install any type of RAM on your motherboard; it is actually dictated by the chipset the manufacturer places on the motherboard. This means you need to know the type of RAM that is compatible with a particular motherboard.

RAM comes in two basic formats: SIMMS and DIMMS. A *SIMM (Single Inline Memory Module)* is a epoxy-coated silicon wafer that contains a number of memory chips. The SIMM fits into a slot on the motherboard. SIMMs must be placed on the motherboard in identical pairs (there are typically four SIMM slots on a motherboard). This means that to have 64MB of memory on the motherboard, two 32MB SIMMs would be installed.

A *DIMM (Dual Inline Memory Module)* contains memory chips like a SIMM but actually provides a greater density of memory chips and therefore more RAM. DIMMs do not have to be installed in pairs, and there are typically three DIMM slots on a motherboard. This means that a computer with 128MB of memory would only require the installation of one 128MB DIMM. All the newer motherboards use DIMMs (although you might find some motherboards that have slots for both DIMMs and SIMMs).

RAM speed has been historically measured in nanoseconds (ns). The lower the nanosecond rating for the RAM, the faster the RAM. For example, 10ns RAM would be faster than 30ns RAM. The speed of newer RAM memory is now measured in MHz. Currently, 100 and 133MHz RAM is available. These two RAM speeds are roughly equivalent to 12ns and 8ns, respectively.

As mentioned earlier, a number of different RAM types are available, and the type used is dictated by the motherboard used on the computer. Although memory has changed dramatically since the advent of the PC, the RAM types listed here are all types you might find on a motherboard designed for a Pentium processor (clone processor motherboard):

- *Fast Page Mode (FPM) memory.* This type of RAM is mounted in SIMM modules of 2, 4, 8, 16, or 32MB and is the traditional RAM type. FPM RAM is typically found in 60ns and 70ns versions. You cannot mix different speeds on the same motherboard.
- *Extended Data Output (EDO) RAM.* This RAM type is an improvement of FPM RAM that provides for the faster reading of data. EDO RAM is usually sold in 60 and 50ns versions. The 50ns version is available at a higher cost. EDO RAM is mounted in SIMM modules.
- *Synchronous Dynamic RAM (SDRAM).* This RAM type is a newer memory type and is used in most new PCs. This memory type is typically called *PC100* or *PC133 RAM*, depending on its speed. SDRAM comes in DIMMs and has an access time of only 12 to 8ns (PC100 and PC133 RAM, respectively) This type of RAM is called *synchronous* because it's actually able to synchronize itself with the speed of the motherboard.
- *Double Data Rate-Synchronous DRAM (DDR-SDRAM).* The newest RAM type, which is still only supported by a few motherboard manufacturers, is DDR-SDRAM. This type of RAM is actually able to transfer data twice during the motherboard's clock cycle (which is measured in MHz). This means that this type of RAM can be twice as fast as other RAM types. DDR-SDRAM comes in DIMMs.

Bottom line, the more RAM you have on a system, the better the system's performance. Adding RAM to any computer will increase system throughput. The addition of RAM, in many cases, will even provide more of a performance jump than switching to a faster processor.

Tip - Another RAM type that you may see in the near future is Rambus DRAM (RDRAM). RDRAM was developed by Rambus, Inc. RDRAM is a very fast type of DRAM that can run as fast as 600MHz. RDRAM is currently used in some graphics accelerator cards. However, Intel has a licensing agreement with Rambus to use RDRAM technology on future Intel motherboards.

Hard Drives

The first hard drive available for the original IBM PC boasted a capacity of 10MB (10 million bytes). Now it is common for hard drives or fixed disks to have capacities in excess of 30GB (that's 30 billion bytes).

Hard drives come in two flavors: IDE drives and SCSI drives. An *Integrated Drive Electronics (IDE)* drive is a hard drive (or other device such as a CD-ROM) where the controller for the drive is built into the drive itself. An IDE drive is connected to the motherboard using a ribbon cable. Each motherboard IDE connection supports up to two IDE drives. Motherboards typically have two IDE connections, meaning a maximum of four IDE drives (this includes hard drives and CD-ROM drives) could be installed on the computer.

Tip - The newest drive standard available is Enhanced IDE (EIDE). It provides faster data rates and can support larger storage devices than the IDE standard.

Small Computer System Interface (SCSI) hard drives are attached to a SCSI controller card (placed in one of the motherboard's expansion slots or built right onto the motherboard, in the case of servers). SCSI controllers allow for the attachment of more drives (up to seven); therefore, SCSI drives are pretty much the standard for server computers. [Figure 3.11](#) shows an IDE drive (on the left) and a SCSI drive; note that the attachment pins are different (due to the fact that different ribbon cables are used to attach these drives to the motherboard or SCSI controller card, respectively).

FIGURE 3.11

Hard drives can be IDE or SCSI drives.

In terms of server hardware, SCSI hard drives are preferred over IDE drives because SCSI drives boast a 12% performance boost over their IDE counterparts. Also, only two IDE drives can be connected to the motherboard via the same ribbon cable (limiting a computer to a maximum of four IDE drives connected to the motherboard). On the other hand, one SCSI controller will support up to seven drives, which makes it easier to deploy multiple-drive configurations, such as a RAID array (RAID is discussed in the section titled "Working with RAID," in Chapter 14, "Protecting Network Data").

Tip - Network storage capacity is no longer limited to the hard drives you have on your network servers. The recent introduction of network attached storage (NAS) provides a way to add storage capacity to a network without adding an actual server computer. An NAS device is really just a box of hard drives that is directly connected to the network. NAS devices, such as the Quantum SnapServer, run their own proprietary OS and can be

configured for file sharing using remote management software that can be run from a network client or server. NAS devices don't need a monitor or keyboard because you remotely configure them over the network. NAS devices provide an easy way to add hard drive storage capacity to a large network or even a peer-to-peer network without the expense of a network operating system and dedicated server hardware.

Network Server Considerations

Network servers require large storage capacity, fast processors, and ample memory. Server tower cases also are larger than client computers and provide much more space for additional hard drives and other peripherals.

Server computers often have motherboards that provide space for more than one processor. Most network operating systems are built to take advantage of multiple processors. In fact, some network operating systems (such as Windows 2000 Advanced Server) support up to 32 processors.

Servers also typically contain many more hard drives than you would find on a client machine and usually contain SCSI hard drives rather than IDE drives (which you would typically find on a client computer). This is not only to provide adequate storage space but also to build in some redundancy for storing files. We will talk about how drive redundancy can be used to help keep important data safe in Chapter 14. The SCSI controllers are also typically part of the server's motherboard rather than an add-on card.

Tip - Servers often will contain drive bays that contain "hot-swappable" drives. This means a drive can be added or swapped while the server is running. These drive bays allow you to access the drive by sliding the drive out of the case; it's not unlike opening a drawer. This means that you can deal with drive issues without powering down the network server or opening the server case.

The hardware configuration for a server will, to a certain degree, be dictated by the network operating system running on that server. Each network operating system provides a baseline configuration for a server that will run that NOS.

For example, the recommended baseline server configuration for Microsoft Windows 2000 Server is as follows:

- Processor: Intel Pentium 166MHz
- RAM: 128MB (256 recommended)
- Hard drive: 2GB
- CD-ROM drive: 12X
- Monitor: Super VGA capable of 800x600 resolution

Novell NetWare 5.1 also has similar hardware requirements as a baseline for a NetWare 5.1 server. Keep in mind that the baseline configuration is often the hardware that is needed just to run the NOS properly and doesn't necessarily take into consideration how many clients you have on your network or the type of resources that will be accessed on a particular server.

Network operating systems are also generally less forgiving when you attempt to run them on computers with odd configurations or atypical hardware. NOS software companies often provide a hardware compatibility list that allows you to view what types of server hardware have been tested with the particular operating system.

Before you actually finalize your server configurations, you need to do some research on the NOS you will be running. Novell, Microsoft, Sun Microsystems, and other NOS providers typically provide white papers and other material that allow you to look at case studies and performance data related to a particular NOS and various hardware configurations.

Tip - Microsoft, Sun, and Novell all provide the hardware requirements to run their network operating systems. Check out these sites:

- www.microsoft.com
 - www.sun.com
 - www.novell.com
-

Network Client Considerations

Network clients don't typically need the processing power and storage capacity required by a server computer. Network clients do need, however, to be able to properly run the client operating system they have been configured with. Be sure that a client machine has at least the minimum hardware configuration to run the client operating system.

Because there are tons of different computer configurations out there (for standalone PCs, with most being quite satisfactory as a network client) and operating systems such as Windows 9x and Windows Me have been created with that fact in mind. They run on many different hardware configurations. Higher-end client software, such as Windows 2000, will be supported on fewer configurations. As is the case with network operating systems, some client operating systems will provide hardware compatibility lists you can review. These lists allow you to choose a compatible configuration for your client computers.

Tip - To buy or to build? That is the question. Although the relatively low price of computer components may seem like a fairly seductive reason to build computers for a network, prebuilt PCs are configured with components that the computer manufacturer knows (or *should* know) work well together. Also, the prospect of building 100 computers for a network is an extremely daunting task, even if you've thrown a few computers together yourself. Adding RAM or hard drives to a PC is one thing; building a computer is something else entirely.

Chapter Summary

- Network interface cards (NICs) supply the connection between the computer and the network's physical medium.
- The NIC supplies the MAC or hardware address that is used to identify a computer on the network.
- Operating systems that embrace plug-and-play technology automatically install and configure plug-and-play NICs.
- In cases where the operating system does not configure a NIC that you've added to the system, you will have to identify an available IRQ for the NIC to use.
- A motherboard is the main circuit board for a PC, and it provides the connection point for the processor and peripheral cards. The motherboard also supplies the data bus used to move data from various devices to the computer's memory and processor, and vice versa.
- Computer memory speed and capacity have changed dramatically as the PC has evolved. The Dual Inline Memory Module (DIMM) is pretty much the rule for outfitting and upgrading PCs with RAM.
- Hard drives come in two basic flavors: IDE (or EIDE) and SCSI. SCSI drives are typically used on network servers because a SCSI controller can support up to seven drives.
- A hub is used as the central connection point for network devices that use copper wire, such as unshielded twisted pair (UTP).
- Bridges can be used on large networks to segment the network and preserve valuable network bandwidth.
- Switches can be used to segment large networks in logical VLANs. Dedicated bandwidth can also be provided to client and server computers connected to a switch. Switches can also support full-duplex data transfers, which allow a computer to send and receive data on the network simultaneously.
- Routers are used to divide networks into logical subnets, which keeps local traffic on each of the subnets. Routers also have the ability to determine the route that data should take as it is moved on the network.
- Network servers require more processing power, RAM, and storage space than network clients.

© Copyright Pearson Education. All rights reserved.