Chapter 2

Storage Devices and Power Supplies

THE FOLLOWING COMPTIA A+ 220-801 EXAM OBJECTIVES ARE COVERED IN THIS CHAPTER:

- ✓ **1.5 Install and configure storage devices and use appropriate media.**
  - Optical drives: CD-ROM, DVD-ROM, Blu-Ray
  - Combo drives and burners: CD-RW, DVD-RW, Dual Layer DVD-RW, BD-R, BD-RE
  - Connection types
    - External: USB, Firewire, eSATA, Ethernet
    - Internal SATA, IDE and SCSI: IDE configuration and setup (Master, Slave, Cable Select), SCSI IDs (0 – 15)
    - Hot swappable drives
  - Hard drives: Magnetic, 5400 rpm, 7200 rpm, 10,000 rpm, 15,000 rpm
  - Solid state/flash drives: Compact flash, SD, Micro-SD, Mini-SD, xD, SSD
  - RAID types: 0, 1, 5, 10
  - Floppy drive
  - Tape drive
  - Media capacity: CD, CD-RW, DVD-RW, DVD, Blu-Ray, Tape, Floppy, DL DVD
- ✓ **1.8 Install an appropriate power supply based on a given scenario.**
  - Connector types and their voltages: SATA, Molex, 4/8-pin 12v, PCIe 6/8-pin, 20-pin, 24-pin, Floppy
  - Specifications: Wattage, Size, Number of connectors, ATX, Micro-ATX
  - Dual voltage options

As a PC technician, you need to know quite a bit about hardware. Given the importance and magnitude of this knowledge, the best way to approach it is in sections. The first chapter introduced the topic via the primary core components, and this chapter follows up where it left off. Specifically, this chapter focuses on storage devices and power supplies.

Identifying Purposes and Characteristics of Storage Devices

What good is a computer without a place to put everything? Storage media hold the data being accessed as well as the files the system needs to operate and data that needs to be saved. The many different types of storage differ in terms of their capacity (how much they can store), access time (how fast the computer can access the information), and the physical type of media used.

Hard Disk Drive Systems

Hard disk drive (HDD) systems (hard disks or *hard drives* for short) are used for permanent storage and quick access. Hard disks typically reside inside the computer, where they are semipermanently mounted with no external access (although there are external and removable hard drives) and can hold more information than other forms of storage. Hard drives use a *magnetic* storage medium and are known as conventional drives to differentiate them from newer solid-state storage media.

The hard disk drive system contains three critical components:

Controller This component controls the drive. The controller chip controls how the drive operates and how the data is encoded onto the platters. It controls how the data sends signals to the various motors in the drive and receives signals from the sensors inside the drive. Most of today's hard disk technologies incorporate the controller and drive into one assembly. The most common and well-known of these are PATA and SATA.

Hard disk This is the physical storage medium. Hard disk drive systems store information on small discs (from under 1 inch to 5 inches in diameter), also called *platters*, stacked together and placed in an enclosure.

Host bus adapter (HBA) This is the translator, converting signals from the controller to signals the computer can understand. Most motherboards today incorporate the host adapter into the motherboard's circuitry, offering headers for drive-cable connection. Legacy host adapters and certain modern adapters house the hard drive controller circuitry.

Figure 2.1 shows a hard disk drive and host adapter. The hard drive controller is integrated into the drive in this case, but it could be resident on the host adapter in other hard drive technologies.

FIGURE 2.1 A hard disk drive system

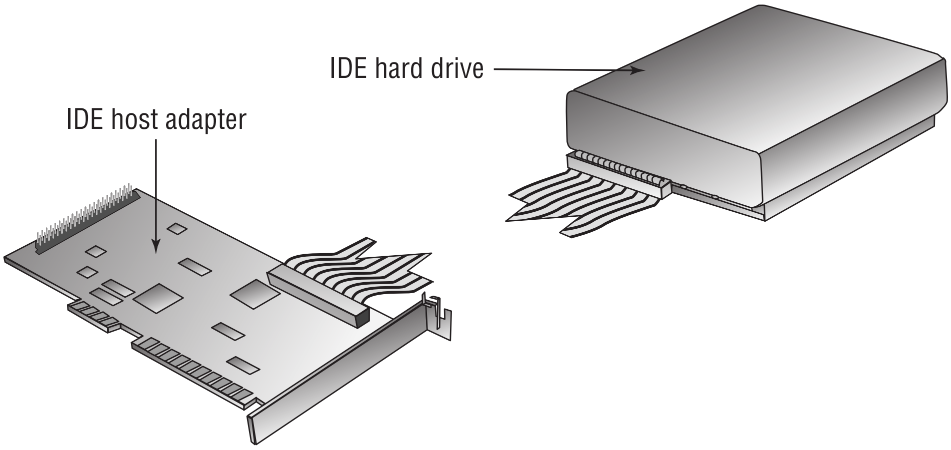

Figure 2.2 shows the 40-pin header (labeled *IDE*) on a motherboard.

FIGURE 2.2 IDE header on a motherboard

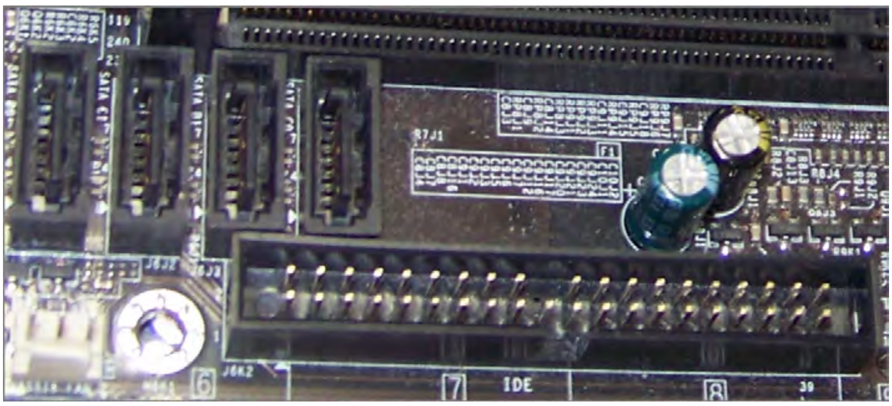

IDE (PATA) drives, discussed later, connect to this interface to access the HBA circuitry in the Southbridge. The four similar, smaller headers at the top left of the photo provide a similar connection for SATA drives, also presented later in this chapter.

Anatomy of a Hard Drive

A hard drive is constructed in a cleanroom to avoid the introduction of contaminants into the hermetically sealed drive casing. Once the casing is sealed, most manufacturers seal one or more of the screws with a warning sticker that removal of or damage to the seal will result in voiding the drive's warranty. Even some of the smallest contaminants can damage the precision components if allowed inside the hard drive's external shell. The following is a list of the terms used to describe these components in the following paragraphs:

- Platters
- Read/write heads
- Tracks
- Sectors
- Cylinders
- Clusters (allocation units)

Inside the sealed case of the hard drive lie one or more platters, where the actual data is stored by the read/write heads. The heads are mounted on a mechanism that moves them in tandem across both surfaces of all platters. Older drives used a stepper motor to position the heads at discrete points along the surface of the platters, which spin at thousands of revolutions per minute on a spindle mounted to a hub. Newer drives use voice coils for a more analog movement, resulting in reduced data loss because the circuitry can sense where the data is located through a servo scheme, even if the data shifts due to changes in physical disc geometry. Figure 2.3 illustrates the key terms presented in this discussion. The four stacked discs shown in the illustration are platters.

FIGURE 2.3 Anatomy of a hard drive

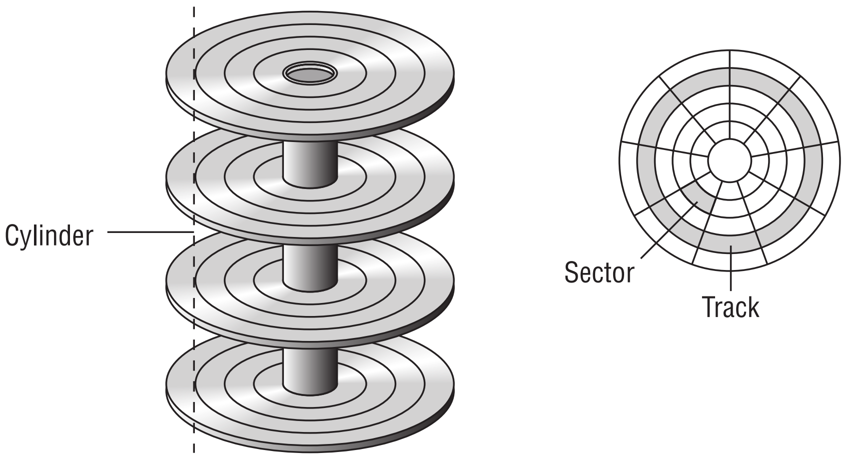

Factory preparation for newer drives or low-level formatting in the field for legacy drives map the inherent flaws of the platters so that the drive controllers know not to place data in these compromised locations. Additionally, this phase in drive preparation creates concentric rings, or *tracks*, which are drawn magnetically around the surface of the platters. Sectors are then delineated within each of the tracks. *Sectors* are the magnetic domains that represent the smallest units of storage on the discs' platters. Magnetic-drive sectors commonly store only 512 bytes (½KB) of data each.

The capacity of a hard drive is a function of the number of sectors it contains. The controller for the hard drive knows exactly how the sectors are laid out within the disk assembly. It takes direction from the BIOS when writing information to and reading information from the drive. The BIOS, however, does not always understand the actual geometry of the drive. For example, the BIOS does not support more than 63 sectors per track. Nevertheless, many hard drives have tracks that contain many more than 63 sectors per track. As a result, a translation must occur from where the BIOS believes it is directing information to be written to where the information is actually written by the controller. When the BIOS detects the geometry of the drive, it is because the controller reports dimensions the BIOS can understand. The same sort of trickery occurs when the BIOS reports to the operating system a linear address space for the operating system to use when requesting that data be written to or read from the drive through the BIOS.

The basic hard disk geometry consists of three components: the number of sectors that each track contains, the number of read/write heads in the disk assembly, and the number of cylinders in the assembly. This set of values is known as CHS (for cylinders/heads/sectors). The number of *cylinders* is the number of tracks that can be found on any single surface of any single platter. It is called a cylinder because the collection of all same-number tracks on all writable surfaces of the hard disk assembly looks like a geometric cylinder when connected together vertically. Therefore, cylinder 1, for instance, on an assembly that contains three platters comprises six tracks (one on each side of each platter), each labeled track 1 on its respective surface.

Because the number of cylinders indicates only the number of tracks on any one writable surface in the assembly, the number of writable surfaces must be factored into the equation to produce the total number of tracks in the entire assembly. This is where the number of heads comes in. There is a single head dedicated to each writable surface, two per platter. By multiplying the number of cylinders by the number of heads, you produce the total number of tracks throughout the disk assembly. By multiplying this product by the number of sectors per track, you discover the total number of sectors throughout the disk assembly. Dividing the result by 2 provides the number of kilobytes the hard drive can store. This works because each sector holds 512 bytes, which is equivalent to ½KB. Each time you divide the result by 1024, you obtain a smaller number but the unit of measure increases from kilobytes to megabytes, from megabytes to gigabytes, and so on. The following equation illustrates this computation:

cylinders (tracks/~~surface~~)
X heads (~~surfaces~~/drive)

---

total tracks (~~tracks~~/drive)
X sectors (sectors/~~track~~)

---

total sectors (sectors/drive)

For example, a drive labeled with the maximum allowed CHS geometry of 16383/16/63, respectively, results in only 7.9GB. Using the equation and multiplying the number of cylinders by the number of heads, you arrive at 262,128 total tracks on the drive. Multiplying this number by 63, the result is that there are 16,514,064 total sectors on the drive. Each sector holds ½KB for a total capacity of 8,257,032KB. Dividing by 1024 to convert to MB and again by 1024 to convert to GB, the 7.9GB capacity is revealed. As a result, although drives larger than 8GB still display the 16383/16/63 CHS capacity for devices that must adhere to the CHS geometry, the CHS scheme cannot be used on today's larger drives at the risk of losing the vast majority of their capacity. The solution is to allow the operating system to reference logical blocks of ½KB sectors that can be individually addressed by a 48-bit value, resulting in 128PB of drive capacity, far above the largest drives being manufactured today. A PB is 1024TB; a TB is 1024GB.

File systems laid down on the tracks and their sectors routinely group a configurable number of sectors into equal or larger sets called *clusters* or *allocation units*. This concept exists because operating system designers have to settle on a finite number of addressable units of storage and a fixed number of bits to address them uniquely. Because the units of storage can vary in size, however, the maximum amount of a drive's storage capacity can vary accordingly, but not unless logical drive capacities in excess of 2TB are implemented. Be aware that today's hard drives and volumes created with RAID can certainly exceed 2TB. This is not a hard limit, however. Larger clusters beget larger volumes but result in less efficient usage of space, discussed in the following paragraph.

No two files are allowed to occupy the same sector, so the opportunity exists for a waste of space that defragmenting cannot correct. Clusters exacerbate the problem by having a similar foible: No two files are allowed by the operating system to occupy the same cluster. The larger the cluster size, then, the larger the potential waste. So, although you can increase the cluster size (generally to as large as 64KB, which corresponds to 128 sectors), you should keep in mind that unless you are storing a notable number of very large files, the waste will escalate astoundingly, perhaps negating or reversing your perceived storage-capacity increase. Nevertheless, if you have single files larger than 2TB, increased cluster sizes are for you. A 64KB cluster size results in a maximum volume size in Windows XP Professional's version of NTFS, for example, of 256TB.

HDD Speeds

As the electronics within the HBA and controller get faster, they are capable of requesting data at higher and higher rates. If the platters are spinning at a constant rate, however, the information can only be accessed as fast as a given fixed rate. To make information available to the electronics more quickly, manufacturers increase the speed at which the platters spin from one generation of drives to the next, with multiple speeds coexisting in the marketplace for an unpredictable period until demand dies down for one or more speeds.

The following spin rates have been used in the industry for the platters in conventional magnetic hard disk drives:

- 5400 rpm
- 7200 rpm
- 10,000 rpm
- 12,000 rpm
- 15,000 rpm

While it is true that a higher revolutions per minute (rpm) rating results in the ability to move data more quickly, there are many applications that do not benefit from increased disk-access speeds. As a result, you should choose only faster drives, which are also usually more expensive per byte of capacity, when you have an application for this type of performance, such as for housing the partition where the operating system resides or for very disk-intensive programs. The lower speeds can be ideal in laptops, where heat production and battery usage can be issues with the higher-speed drives.

Solid-State Drives

Conventional hard disk drive platters are still manufactured the same way they have always been. They are metal or glass discs with a magnetic coating on their surface. The read/write heads change the magnetic orientation of each bit location, storing either a binary one or a binary zero. The same head senses the magnetic orientation of each location as the data is read from the disc.

In contrast, solid-state drives (*SSDs*) have no moving parts but use the same solid-state memory technology found in the other forms of flash memory. All solid-state memory is limited to a finite number of write (including erase) operations. Algorithms have been developed to constantly spread the write operations over the entire device. Such "wear leveling" increases the life of the SSD, but lack of longevity remains a disadvantage of this technology.

SSDs read contents more quickly, can consume less power and produce less heat, and are more reliable and less susceptible to damage from physical shock and heat production than their magnetic counterparts. However, the technology to build an SSD is still more expensive per byte, and SSDs are not yet available in capacities high enough to rival the upper limits of conventional hard disk drive technology.

SSDs are separated into two broad categories, volatile DRAM-based and non-volatile flash-based. Flash-based SSDs made with NAND (a transistor-based gate that has the opposite output to an AND gate) memory use considerably less power than HDDs. Those made with DRAM can use every bit as much power as conventional drives, however. The advantage of those made using the standard RAM modules used in desktop motherboards is that the modules can often be upgraded using larger modules, making a larger SSD overall.

When used as a replacement for traditional HDDs, SSDs are most often expected to behave in a similar fashion, mainly by retaining contents across a power cycle. The speed of the HDD is expected to be maintained or exceeded by the SSD as well. The volatility of DRAM-based SSDs can be compensated for by adding a backup power source, such as a battery or capacitor, or by keeping a non-volatile backup of the drive's data that does not detract from the speed of the primary storage location. Flash-based SSDs, while faster during read operations than their HDD counterparts, can be made faster still by including a small amount of DRAM as a cache. DRAM-based SSDs are faster yet.

Floppy Drives

A floppy disk (or floppy diskette) is a magnetic storage medium that uses a diskette made of thin, flexible plastic enclosed in a protective casing. The floppy disk once enabled information to be transported from one computer to another very easily. Today, floppies are a little too small in capacity to be relevant. They were first replaced by writable CDs and DVDs. Today, solid-state storage is the closest analogue to how floppies were originally used. The term *floppy disk* initially referred to the antiquated 8″ medium used with minicomputers and mainframes. The original PC floppy diskette, which used a platter that was 5¼″ in diameter and known as a *minifloppy diskette*, is also obsolete; the *microfloppy diskette* is a diskette that is 3½″ in diameter. A few computers today use microfloppy diskettes, either through internal drives or external USB-attached drives, but most support no floppy at all.

NOTE

Generally speaking, throughout this book we will use the term *floppy drive* to refer to a 3½″ microfloppy diskette drive. Additionally, it is important to understand that the term *floppy* refers to the enclosed disk platter and not to the external enclosure. A mistake made by many uninformed individuals during the transition from 5¼″ floppies to 3½″ floppies was to refer to the newer, smaller diskettes as "hard disks." The microfloppy diskettes are still considered floppy because of the internal platter. The rigidity of the newer enclosure, in contrast to the flexibility of the older diskettes, had nothing to do with the name of the technology.

A *floppy diskette drive*, or FDD (shown in Figure 2.4), is used to read and write information to and from these disks. The advantage of these drives is that they allow portability of data: you can transfer data from one computer to another on a diskette. The downside of a floppy disk drive is its limited storage capacity. Whereas a hard drive can store hundreds or thousands of gigabytes of information, floppy disks allow storage of only one or two megabytes, although formats of much higher capacity exist, using roughly the same physical form factor but largely incompatible technology. Floptical and SuperDisk are examples.

FIGURE 2.4 A floppy diskette drive

Table 2.1 shows five different floppy-diskette formats with their corresponding data capacities supported in PC systems over the years. The following abbreviations are used: DD means double density, HD means high density. and ED means extended density. A sector holds 512 bytes; call it ½KB to make the math easier. Just remember, all your computations result in kilobytes. Unlike other storage capacities, which shift magnitudes by multiplying or dividing by 1024, floppy storage capacities move up one order of magnitude each time you divide by 1000—for example, kilobytes to megabytes, then megabytes to gigabytes. For instance, the 3½″ HD floppy has two sides to its single platter, yielding a total of 160 tracks. With 18 sectors per track, the total number of sectors on such a floppy diskette totals 2880. Because each sector represents ½KB, the resulting disk capacity totals 1440KB. Dividing by an even 1000 changes the magnitude from kilobytes to megabytes and arrives at the 1.44MB capacity commonly associated with these diskettes.

TABLE 2.1 Floppy diskette capacities

Floppy Drive Size	Tracks per Side	Sectors per Track	Capacity
5¼″ DD	40	9	360KB
5¼″ HD	80	15	1.2MB
3½″ DD	80	9	720KB
3½″ HD	80	18	1.44MB
3½″ ED	80	36	2.88MB

Optical Storage Drives

Most computers today have an optical storage drive, such as the latest *Blu-ray Disc* (*BD*), a digital versatile disc—or digital video disc (*DVD*), or the legacy compact disc (*CD*) drive. Each type of optical drive can be expected to also support the technology that came before it. Such optical storage devices began earnestly replacing floppy diskette drives in the late 1990s. Although, like HDDs, these discs have greater data capacity and increased performance over floppies, it is not intended that they replace hard disk drives. HDDs greatly exceed the capacity and performance of optical drives.

The CDs, DVDs, and BDs used for data storage are virtually the same as those used for permanent recorded audio and video. The way data, audio, and video information is written to consumer-recordable versions makes them virtually indistinguishable from such professionally manufactured discs. Any differences arise from the format used to encode the digital information on the disc. Despite the differences among the professional and consumer formats, newer players have no issue with any type of disc used. The encoding schemes used to store data on such discs are incompatible with the schemes used to record audio and video to the same discs.

CD-ROMs, DVD-ROMs, BD-ROMs, and Capacities

The amount of data that can be stored on the three primary formats of optical disc varies greatly, with each generation of disc exceeding the capacity of all previous generations. The following sections detail the science behind the capacities of all three formats.

CD-ROM

The *CD-ROM* (read-only memory) was designed for long-term storage of data. CD-ROMs are read-only, meaning that information written in the factory can't be erased or changed. CD-ROMs became so popular because they made a great software distribution medium. Programs are always getting larger and increasingly require more room to install, version after version. Instead of installing a program of the day using 100 floppy disks, you could use a single CD-ROM, which can hold approximately 650MB in its original, least-capable format. Although CDs capable of storing 700MB eventually became and continue to be the most common, discs with 800MB and 900MB capacities were standardized as well. See Table 2.2 after the following two sections for a list of optical discs and their capacities.

DVD-ROM

For even more storage capacity, many computers feature some form of DVD or BD drive, such as the original DVD-ROM drive. The basic *DVD-ROM* disc is a single-sided disc that has a single layer of encoded information. These discs have a capacity of 4.7GB, many times the highest CD-ROM capacity. Simple multiplication can sometimes be used to arrive at the capacities of other DVD-ROM varieties. For example, by adding another media surface on the side of the disc where the label is often applied, a double-sided disc is created. Such double-sided discs have a capacity of 9.4GB, exactly twice that of a single-sided disc.

Practically speaking the expected 9.4GB capacity from two independent layers isn't realized when those layers are placed on the same side of a DVD, resulting in only 8.5GB of usable space. (BDs do not have this issue; they make use of the full capacity of each layer.) The loss of capacity is due to the space between tracks on both layers being 10 percent wider than normal to facilitate burning one layer without affecting the other. This results in about 90 percent remaining capacity per layer. This technology is known as DVD DL (*DL* for *dual-layer)*, attained by placing two media surfaces on the same side of the disc, one on top of the other, and using a more sophisticated mechanism that burns the inner layer without altering the semitransparent outer layer and vice versa, all from the same side of the disc. Add the DL technology to a double-sided disc and you have a disc capable of holding 17.1GB of information, again twice the capacity of the single-sided version. Figure 2.5 shows an example of an early DVD-ROM drive, which also accepts CD-ROM discs. Modern 5¼″ optical drives are indistinguishable from older ones, aside from obvious markings concerning their capabilities.

FIGURE 2.5 An early DVD-ROM drive

BD-ROM

The next generation of optical storage technology was designed for modern high-definition video sources. The equipment used to read the resulting discs employs a violet laser, in contrast to the red laser used with standard DVD and CD technologies. Taking a bit of creative license with the color of the laser, the Blu-ray Disc Association names itself and the technology Blu-ray Disc (BD), after this "visibly" different characteristic. Blu-ray technology further increases the storage capacity of optical media without changing the form factor. On a 12cm disc, similar to those used for CD-ROMs and DVD-ROMs, BD derives a 25GB storage capacity from the basic disc. When you add a second layer to the same or opposite side of the disc, you attain 50GB of storage. The Blu-ray laser is of a shorter wavelength (405nm) than that of DVD (650nm) and CD (780nm) technologies. As a result, and through the use of refined optics, the laser can be focused on a much smaller area of the disc. This leads to a higher density of information being stored in the same area.

An interesting point to note is that designers of the Blu-ray technology do not have to stop with the common double-layer solution to increasing capacity. Blu-ray discs with more than four layers on a side have been demonstrated, largely owing to the extremely accurate focus attainable with the Blu-ray laser.

In the interest of completeness, it should be mentioned that a high-definition technology directly related to DVD, because it comes from the same forum, and named HD DVD remains only as a footnote to the Blu-ray story. In February 2008, Toshiba, HD DVD's primary champion, gave up the fight, conceding Blu-ray disc as the winner in the high-definition optical-disc race. HD DVD featured red- and blue-laser compatibility and 15GB data storage capacity.

Table 2.2 draws together the most popular optical-disc formats and lists their respective capacities. Some of these formats have already been introduced; others are presented in the upcoming section “Recordable Discs and Burners.” Boldfaced capacities in the table are the commonly accepted values for their respective formats.

TABLE 2.2 Optical discs and their capacities

Disc Format	Capacity
CD SS (includes recordable versions)	650MB, **700MB**, 800MB, 900MB
DVD-R/RW SS, SL	4.71GB (**4.7GB**)
DVD+R/RW SS, SL	4.70GB (**4.7GB**)
DVD-R, DVD+R DS, SL	**9.4GB**
DVD-R SS, DL	8.54GB (**8.5GB**)
DVD+R SS, DL	8.55GB (**8.5GB**)
DVD+R DS, DL	17.1GB
BD-R/RE SS, SL	25GB
BD-R/RE SS, DL	50GB
BD-R/RE DS, DL	100GB

SS: single-sided; DS: double-sided; SL: single-layer; DL: dual-layer

Optical Drive Data Rates

CD-ROM drives are rated in terms of their data transfer speed. The first CD-ROM drives transferred data at the same speed as home audio CD players, 150KBps, referred to as 1X. Soon after, CD drives rated as 2X drives that would transfer data at 300KBps appeared. They increased the spin speed in order to increase the data transfer rate. This system of ratings continued up until the 8X speed was reached. At that point, the CDs were spinning so fast that there was a danger of them flying apart inside the drive. So, although future CD drives used the same rating (as in 16X, 32X, and so on), their rating was expressed in terms of theoretical maximum transfer rate; 52X is widely regarded as the highest multiplier for data CDs. Therefore, the drive isn’t necessarily spinning faster, but through electronics and buffering advances, the transfer rates continued to increase.

The standard DVD-ROM 1X transfer rate is 1.4MBps, already nine times that of the comparably labeled CD-ROM. As a result, to surpass the transfer rate of a 52X CD-ROM drive, a DVD-ROM drive need only be rated 6X. DVD transfer rates of 16X at the upper end of the scale are common.

The 1X transfer rate for Blu-ray is 4.5MBps, roughly 3¼ times that of the comparable DVD multiplier and close to 30 times that of the 1X CD transfer rate. It takes 2X speeds to properly play commercial Blu-ray videos.

Recordable Discs and Burners

Years after the original factory-made CD-ROM discs and the drives that could read them were developed, the industry, strongly persuaded by consumer demand, developed discs that, through the use of associated drives, could be written to once and then used in the same fashion as the original CD-ROM discs. The firmware with which the drives were equipped could vary the power of the laser to achieve the desired result. At standard power, the laser in these drives allowed inserted discs to be read from. Increasing the power of the laser allowed the crystalline media surface to be melted and changed in such a way that light would reflect or refract from the surface in microscopic increments. This characteristic allowed mimicking of the way in which the original CD-ROM discs stored data.

Eventually, discs that could be written to, erased, and rewritten were developed. Drives that contained the firmware to recognize these discs and control the laser varied the laser's power in three levels. The original two levels closely matched those of the writable discs and drives. The third level, somewhere in between, could neutralize the crystalline material without writing new information to the disc. This medium level of power left the disc surface in a state similar to its original, unwritten state. Subsequent high-power laser usage could write new information to the neutralized locations.

The best algorithms for such drives, which are still available today, allow two types of disc erasure. The entire disc can be erased before new data is written (*erased* or *formatted*, in various application interfaces), or the data can be erased on the fly by one laser, just fractions of a second before new data is written to the same location by a second laser. If not properly implemented in a slow, determined fashion, the latter method can result in write errors because the crystalline material does not adequately return to its neutral state before the write operation. The downside to slowing down the process is obvious, and methods exist to allow a small level of encoding error without data loss. This need to move more slowly adds a third speed rating, the rewrite speed, to the read and write speeds a drive is capable of. The following section delves more deeply into this concept. Updates to the drive's firmware can often increase or equalize these speeds.

Recordable CD Formats

CD-recordable (CD-R) and CD-rewritable (*CD-RW*) drives (also known as CD *burners*) are essentially CD-ROM drives that allow users to create (or *burn*) their own CD-ROMs. They look very similar to CD-ROM drives but feature a logo on the front panel that represents the drive's CD-R or CD-RW capability. Figure 2.6 shows the CD-R and CD-RW logos that you are likely to see on such drives.

FIGURE 2.6 CD-R and CD-RW logos

The difference between these two types of drives is that CD-R drives can write to a CD-R disc only once. A CD-RW drive can erase information from a CD-RW disc and rewrite to it multiple times. Also, CD-RW drives are rated according to their write, rewrite, and read times. So instead of a single rating like 64X, in the case of CD-ROM drives, they have a compound rating, such as 52X-32X-52X, which means it writes at 52X, rewrites at 32X, and reads at 52X.

Recordable DVD Formats

A DVD burner is similar to a CD-R or CD-RW drive in how it operates: It can store large amounts of data onto a special, writable DVD. Single-sided, dual-layer (DL) discs can be used to write 8.5GB of information to one single-sided disc. Common names for the variations of DVD burning technologies include DVD+R, DVD+RW, DVD-R, DVD-RW, DVD-RAM, DVD-R DL, and DVD+R DL. The "plus" standards come from the DVD+RW Alliance, while the "dash" counterparts are specifications of the DVD Forum. The number of sectors per disc varies between the "plus" and "dash" variants, so older drives might not support both types. The firmware in today's drives knows to check for all possible variations in encoding and capability. The "plus" variants have a better chance of interoperability, even without the disc being finalized.

A DVD-ROM or recordable drive looks very similar to a CD-ROM drive. The main difference is the presence of one of the various DVD logos on the front of the drive. CD-ROM and recordable CDs are usually able to be read and, if applicable, burned in DVD burners. Figure 2.7 and Figure 2.8 show the most popular data-oriented logos you are likely to see when dealing with DVD drives suited for computers. Figure 2.7 shows the "dash" logos, while Figure 2.8 shows the "plus" logos.

Table 2.3 lists the main DVD formats used for storing and accessing data in computer systems as well as their characteristics.

TABLE 2.3 DVD formats and characteristics

Format	Characteristics
DVD-ROM	Purchased with data encoded; not able to be changed
DVD-R, DVD+R	Purchased blank; able to be written to once and then treated like a DVD-ROM
DVD-RW, DVD+RW	Purchased blank; able to be written to and erased multiple times; session usually must be closed for subsequent access to stored data
DVD-RAM	Purchased blank; able to be written to and erased just like a hard or floppy disk; no session to close before subsequent access to stored data

FIGURE 2.7 DVD Forum logos

FIGURE 2.8 DVD+RW Alliance logos

Recordable BD Formats

The Blu-ray Disc Association duplicated use of the R suffix to denote a disc capable of being recorded on once by the consumer. Instead of the familiar RW, however, the association settled on RE, short for re-recordable. As a result, watch for discs labeled BD-R and BD-RE. Dual-layer versions of these discs can be found as well.

The Blu-ray Disc Association decided against creating separate logos for each BD type, resolving instead to use the logo in Figure 2.9 solely. Discs are labeled most often in a sans-serif font with the actual type of disc as well as this generic BD logo.

FIGURE 2.9 The Blu-ray Disc logo

Drive Interfaces and RAID

Storage devices come in many shapes and sizes. In addition to IDE/EIDE and SCSI, two of the older standards, there is now Serial ATA (*SATA*), and you can differentiate between internally and externally attached drives. The following sections look at storage devices from a number of those perspectives.

> **NOTE** Parallel ATA (PATA) is nothing new but rather the name retroactively given to the ATA/IDE standards when SATA became available. PATA uses the classic 40-pin connector for parallel data communications, whereas SATA uses a more modern 7-pin card-edge connector for serial data transfer.

AT Attachment Drives

At one time, integrated drive electronics (*IDE*) drives were the most common type of hard drive found in computers. Though so often thought of in relation to hard drives, IDE was much more than a hard drive interface; it was also a popular interface for many other drive types, including optical drives and tape drives. Today, we call IDE PATA and consider it to be a legacy technology. The industry now favors SATA instead.

IDE/PATA Drives

The design of IDE is simple: Put the controller chip and its related electronics right on the drive, and use a relatively short ribbon cable to connect the drive and controller to an interface on the system. This offers the benefits of decreasing signal loss (thus increasing reliability), eliminating the need for low-level formatting in the field, and making the drive easier to install. The IDE interface can be an expansion board (often referred to as a paddle card because it does little more than transfer pins from the drive to pins on the expansion bus; it has no real intelligence onboard), or it can be built into the motherboard, as was the case on almost all systems over quite a few generations in the past, especially nonserver, desktop systems. Today, similar systems tend to have SATA headers exclusively.

IDE generically refers to any drive that has a built-in controller. Enhanced Small Device Interface (ESDI—an antiquated technology) and, to a certain degree, SCSI drives have drive electronics integrated into them. The IDE we know today was once more properly called advanced technology attachment (ATA); the terms *SATA* and *PATA* were derived from *ATA*. Because ATA encompasses both SATA and PATA, but because IDE is synonymous only with PATA, it is no longer appropriate to equate IDE directly with the term ATA.

There have been many revisions of the IDE standard over the years, and each one is designated with a certain AT Attachment number—ATA-1 through ATA-8, so far. Drives that support ATA-2 and higher are generically referred to as Enhanced IDE (EIDE). SATA specifications appeared at the end of this series but then branched off on their own.

With ATA-3, a technology called ATA Packet Interface (ATAPI) was introduced to help deal with IDE devices other than hard disks. ATAPI enables the BIOS to recognize an IDE CD-ROM drive, for example, or a tape backup or Zip drive. ATA-3 also introduced the Self-Monitoring and Reporting Technology (SMART). SMART allows a hard drive to monitor itself and warn the user during and after boot-up of any impending failure. When heeded, these warnings allow you time to salvage data before it is lost. Generally, backing up the potentially ailing drive, before replacing it and restoring your data, is the best route. Note, however, that drives can still fail, and data loss can still occur, even without a warning from SMART.

Starting with ATA-4, a new technology was introduced called UltraDMA, supporting transfer modes capable of rates of up to 33MBps.

ATA-5 supports UltraDMA/66, with transfer modes having rates of up to 66MBps. To achieve this high rate, the drive must have a special 80-wire ribbon cable (still with only 40 pins, however), and the motherboard or IDE controller card must support ATA-5.

ATA-6 supports UltraDMA/100, with transfer modes capable of up to 100MBps.

NOTE

If an ATA-5 or ATA-6 drive is used with a normal 40-wire cable or is used on a system that doesn't support the higher-speed modes, it reverts to the ATA-4 performance level.

ATA-7 supports UltraDMA/133, with transfer modes of 133MBps and up to 150MBps for serial ATA.

ATA-8 made only minor revisions to ATA-7 and also supports UltraDMA/133 and 150MBps SATA and has the potential to support SATA 300.

IDE Pros and Cons

The primary benefit of IDE is that it's a mature, well-known technology. At one time, almost every motherboard had IDE connectors. A typical motherboard in those days had two IDE connectors, and each connector supported a single channel of up to two drives on the same cable. That means you were limited to four IDE devices per system, unless you add an expansion board containing another IDE interface. In contrast, with SCSI you can have up to seven devices (including drives) per interface, roughly double or quadruple that on some types of SCSI.

Performance also may suffer when certain IDE devices share an interface. It is recommended that you pair like devices on a channel. Otherwise, the slower device can have a negative impact on the faster one. SCSI drives are much more efficient with this type of transfer.

IDE Installation and Configuration

To install an IDE drive, do the following:

1. Set the master/slave jumper on the drive.
2. Install the drive in the drive bay.
3. Connect the power-supply cable.
4. Connect the ribbon cable to the drive and to the motherboard or IDE expansion board. There is a colored (usually red) stripe down one edge of the ribbon cable that is used to correctly orient the cable both where it connects to the drive and where it connects to the motherboard. If there is no marking for pin 1, you'll usually orient the red stripe toward the drive's power connector. Don't rely on that too much, though. Floppy drives are notorious for placing pin 1 away from its power connector.
5. Configure the drive in BIOS Setup if it isn't automatically detected.
6. Partition and format the drive using the operating system.

Each IDE channel can have only one *master* drive on it. If there are two drives on a single cable, one of them must be the *slave* drive. This setting is accomplished via a jumper on the drive. Some drives have a separate setting for Single (that is, master with no slave) and Master (that is, master with a slave); others use the Master setting generically to configure either case. Figure 2.10 shows a typical master/slave jumper scenario, but different drives may have different jumper positions to represent each state.

Another option is to use the Cable Select setting for master/slave selection. Most cables support Cable Select, but it won't work if the cable is not wired properly. Another caveat is that you must never mix Cable Select with Master and Slave settings. If one drive is set manually as Master, the other on the same cable must be set as Slave. If you set one drive to be configured as master or slave automatically with Cable Select, the other must be set for Cable Select as well. The wiring of the cable will result in the drive at the end of the cable being selected as master.

Most BIOS Setup programs today support Plug and Play, so they detect the new drive automatically at startup. If this doesn't work, the drive may not be installed correctly, the jumper settings may be wrong, or BIOS Setup may have the IDE interface set to None or Disable rather than Auto. Enter BIOS Setup and find out. Setting the IDE interface to Auto and then allowing the BIOS to detect the drive is usually all that is required.

FIGURE 2.10 Master/slave jumpers

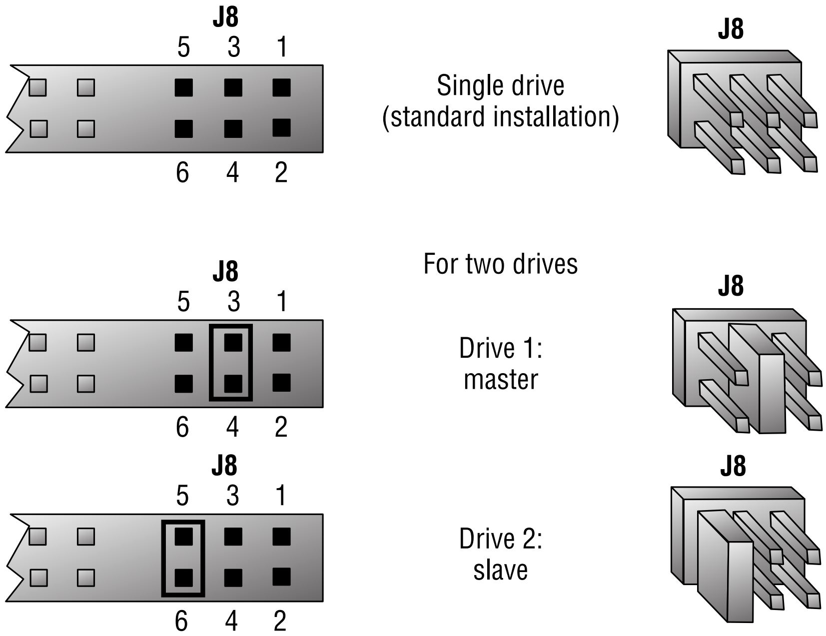

In BIOS Setup for the drive, you might have the option of selecting a DMA or programmed input/output (PIO) setting for the drive. Both are methods for improving drive performance by allowing the drive to write directly to RAM, bypassing the CPU when possible. For modern drives that support UltraDMA, neither of these settings is necessary or desirable.

Now that your drive is installed, you can proceed to partition and format it for the operating system you've chosen. Then, finally, you can install your operating system of choice.

For a modern Windows system, allow the Windows Setup program to partition and format the drive, or use the Disk Management utility in Windows to perform those tasks. To access Disk Management from Control Panel, choose Administrative Tools and then choose Computer Management. You can also right-click Computer (or My Computer in older versions of Windows) and then click Manage.

SATA Drives

Serial ATA began as an enhancement to the original ATA specifications, also known as IDE and, today, PATA. Technology is proving that the orderly progression of data in a single-file path is superior to placing multiple bits of data in parallel and trying to synchronize their transmission to the point that each bit arrives simultaneously. In other words, if you can build faster transceivers, serial transmissions are simpler to adapt to the faster rates than are parallel transmissions.

The first version of SATA, known as SATA 1.5Gbps, and also known by the less-preferred terms SATA I and SATA 150, used an 8b/10b encoding scheme that requires 2 non-data overhead bits for every 8 data bits. The result is a loss of 20 percent of the rated bandwidth, but the silver lining is that the math becomes quite easy. Normally you have to divide by 8 to convert bits to bytes. With 8b/10b encoding, you divide by 10. Therefore, the 150MBps throughput that this version of SATA was nicknamed for is easily derived as ⅒ of the 1.5Gbps transfer rate. The original SATA specification also provided for hot swapping at the discretion of the motherboard and drive manufacturers.

Similar math works for SATA 3Gbps, also recklessly tagged as SATA II and SATA 300, and SATA 6Gbps, which is not approved for being called SATA III or SATA 600, but the damage is already done. Note that each subsequent version doubles the throughput of the version before it. Figure 2.11 shows a SATA connector on a data cable followed by the headers on a motherboard that will receive it.

FIGURE 2.11 SATA cable and headers

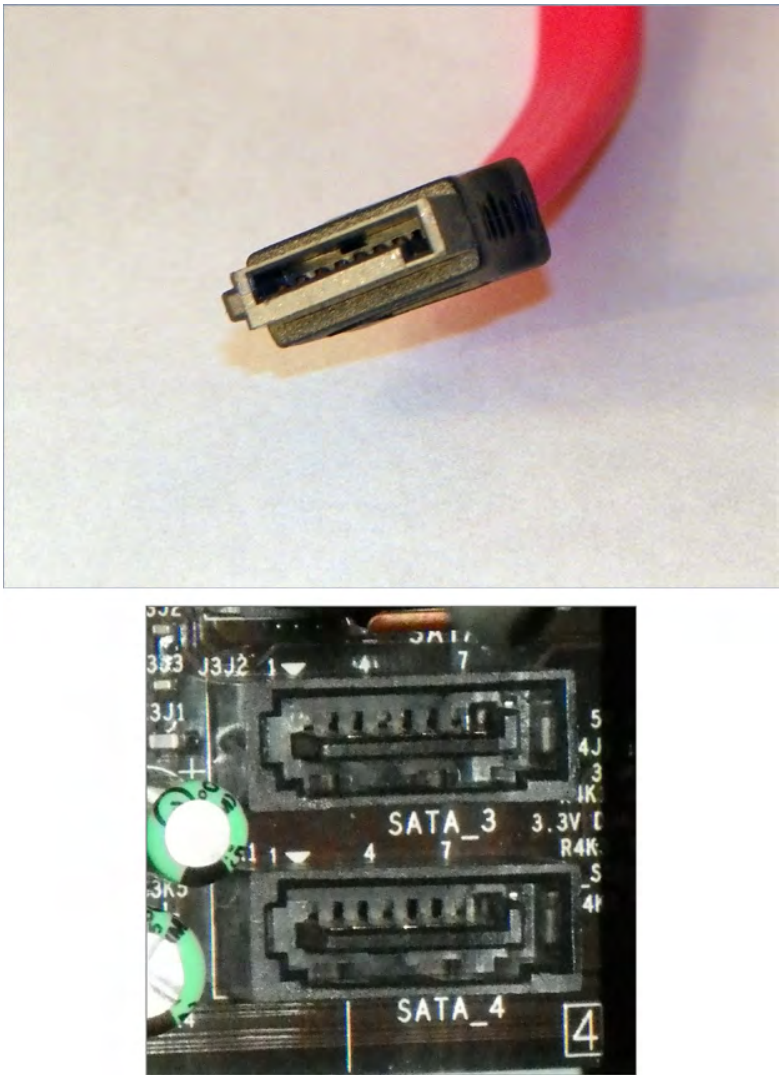

The card-edge style connectors for data and power are arranged in such a manner on the back of SATA drives that no cables are required, although desktop and server systems almost certainly employ them. The same interface, however, can be used in laptops without the adapters needed to protect the delicate pins of the parallel interfaces found on the preceding generation of small form factor drives. The lack of the adapter also leads to less space reserved in the bays for drives of the same size, giving designers and consumers the choice between smaller systems or more complex circuitry that can move into the newly available space.

SCSI Drives

Small Computer System Interface (SCSI) devices can be either internal or external to the computer. Eight-bit SCSI-1 and SCSI-2 internal devices use a SCSI A cable, a 50-pin ribbon cable similar to that of an IDE drive. Sixteen-bit SCSI uses a SCSI P cable, with 68 wires and a 68-pin D-subminiature connector. There is also an 80-pin internal SCA connector ideal for use in hot-swapping environments.

Like IDE and floppy-drive cables, 50-pin SCSI ribbon cables have a colored stripe (usually blue or red, but it depends on the color of the rest of the cable) down one side to indicate the orientation of pin 1. The 68-wire ribbon requires no indicator because the connector is keyed by its shape. The cable is often a multicolored ribbon braided for noise reduction. External SCSI connectors depend on the type of standard in use. SCSI-1 uses a 50-pin Centronics connector similar to the 36-pin version used on the older printers with parallel interfaces. SCSI-2 uses a 25-, 50-, or 68-pin connector. SCSI-3 uses a 68- or 80-pin connector.

To configure SCSI, you must assign a unique device number (often called a SCSI address, SCSI ID, or SCSI device ID) to each device on the SCSI bus. These numbers are configured through jumpers, DIP switches, and up/down pushbuttons with the selected ID displayed through a hole on a wheel, among other ways. When the SCSI controller needs to send data to the device, it activates the wire dedicated to signaling that address.

A device called a *terminator* (technically a terminating resistor pack) must be installed at both ends of the bus to keep the signals "on the bus." The device then responds with a signal that contains the number of the device that sent the information and the data itself. The terminator can be built into the device and activated/deactivated with a jumper, or it can be a separate block or connector hooked onto the device when termination is required.

Termination can be either active or passive. A passive terminator works with resistors driven by the small amount of electricity that travels through the SCSI bus. Active termination uses voltage regulators inside the terminator. Active termination is much better, and you should use it whenever you have fast, wide, or Ultra SCSI devices on the chain and/or more than two SCSI devices on the chain. It may not be obvious from looking at a terminator whether it's active or passive.

SCSI Device Installation and Configuration

Installing SCSI devices is more complex than installing an IDE drive. The main issues with installing SCSI devices are cabling, termination, and addressing.

We'll discuss termination and cabling together because they're closely tied. There are two types of cabling:

- Internal cabling uses a 50-, 68-, or 80-wire ribbon cable with several keyed connectors. These connectors are attached to the devices in the computer (the order is unimportant), with one connector connecting to the adapter.
- External cabling uses thick, shielded cables that run from adapter to device to device in a fashion known as *daisy-chaining*. Each device has two ports on it (most of the time). When hooking up external SCSI devices, you run a cable from the adapter to the first device. Then you run a cable from the first device to the second device, from the second to the third, and so on.

Because there are two types of cabling devices, you have three ways to connect them. The methods differ by where the devices are located and whether the adapter has the terminator installed. The guide to remember here is that *both ends* of the bus must be terminated. Let's look briefly at the three connection methods:

Internal devices only When you have only internal SCSI devices, you connect the cable to the adapter and to every SCSI device in the computer. You then install the terminating resistors on the adapter and terminate the last drive in the chain. All other devices are unterminated. This is demonstrated in Figure 2.12.

FIGURE 2.12 Cabling internal SCSI devices only

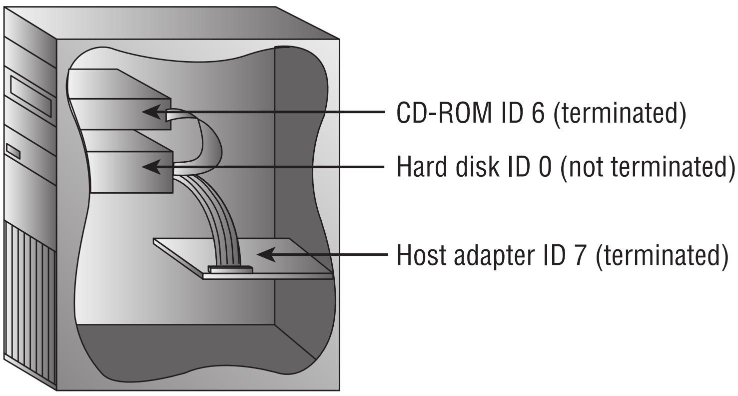

NOTE Some devices and adapters don't use terminating resistor packs; instead, you use a jumper or DIP switch to activate or deactivate SCSI termination on such devices. Check the documentation to find out what type your device uses.

External devices only In the next situation, you have external devices only, as shown in Figure 2.13. By external devices, we mean that each has its own power supply. You connect the devices in the same manner in which you connected internal devices, but in this method you use several very short (less than 0.5 meters) *stub* cables to run between the devices in a daisy chain (rather than one long cable with several connectors). The effect is the same. The adapter and the last device in the chain (which has only one stub cable attached to it) must be terminated.

FIGURE 2.13 Cabling external SCSI devices only

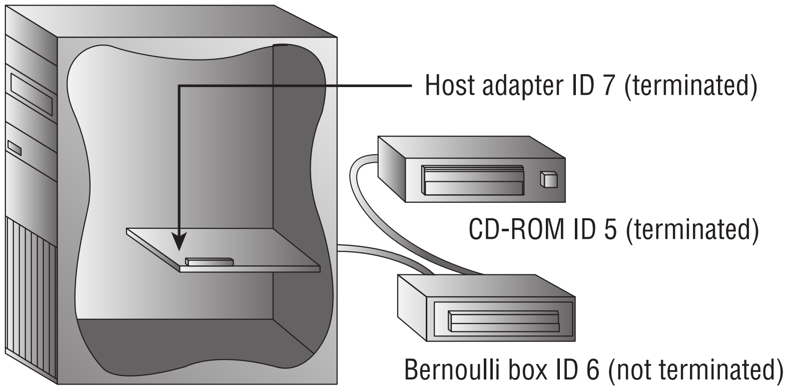

Both internal and external devices Finally, there's the hybrid situation in which you have both internal and external devices (Figure 2.14). Most adapters have connectors for both internal and external SCSI devices—if yours doesn't have both, you'll need to see if anybody makes one that will work with your devices. For adapters that do have both types of connectors, you connect your internal devices to the ribbon cable and attach the cable to the adapter. Then you daisy-chain your external devices off the external port. You terminate the last device on each chain, leaving the adapter unterminated.

FIGURE 2.14 Cabling internal and external SCSI devices together

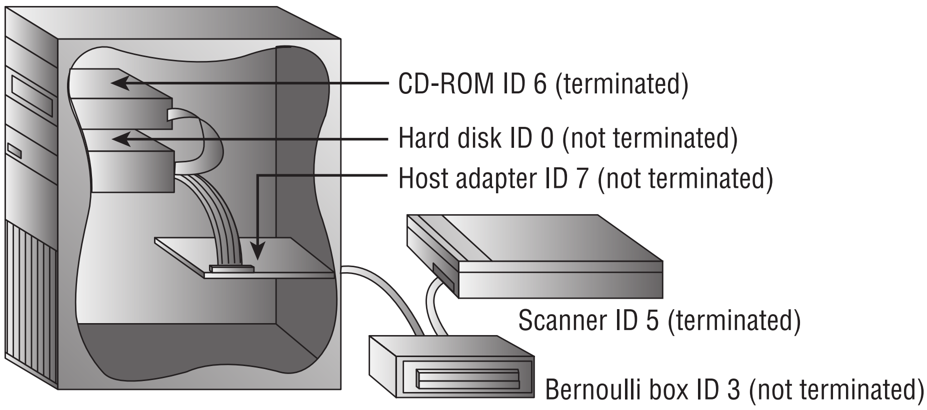

NOTE

Even though the third technique described is the technically correct way to install termination for the hybrid situation (in which you have both internal and external devices), some adapter cards still need to have terminators installed. Both ends of a SCSI chain must be terminated.

Each device must also have a unique SCSI ID number. This number can be assigned by the jumper (with internal devices) or with a rotary switch (on external devices). You start by assigning your adapter an address, if necessary. This can be any number from 0 to 7 on an 8-bit bus, 0 to 15 on a 16-bit bus, and 0 to 31 on a 32-bit bus, as long as no other device is using that ID. An ID of 7 is always recommended for the host adapter. This ID has the highest priority for arbitration and therefore the host adapter will always take priority.

Here are some recommendations that are commonly accepted by the PC community. Remember that these are guidelines, not rules:

- Generally speaking, give slower devices higher priority so they can access the bus whenever they need it. Higher numbers are higher priority.
- Set the bootable (or first) hard disk to ID 0.
- Set the CD-ROM to ID 3.

After setting the IDs and the devices are cabled and terminated, you have to get the PC to recognize the SCSI adapter and its devices. The SCSI adapter manages all SCSI device resource allocation, so generally all that is required is to make sure the operating system is able to see the SCSI adapter. This involves installing a driver for the adapter in Windows, for example.

However, if you want to boot from a SCSI drive, the system must be able to read from that drive in order to load the operating system; you must enable the SCSI adapter's own BIOS extension so that the PC can read from it at startup without a driver. Check the documentation for the adapter; sometimes the BIOS Setup program for the SCSI adapter is activated via a function key at startup.

If the SCSI adapter has no BIOS, some versions of Windows will create the `NTBOOTDD.SYS` file in the root of the system partition. Pressing F6 during Windows setup and installing the driver for the device causes this to happen. Once the drive is installed and talking to the computer, you can high-level format the media and install the operating system.

> **TIP** If there are problems, double-check the termination and ID numbers. Termination will most likely be the problem, but you might need to make sure no two devices are set to the same ID.

RAID

RAID stands for *Redundant Array of Independent Disks*. It's a way of combining the storage power of more than one hard disk for a special purpose, such as increased performance or fault tolerance. RAID was once more commonly done with SCSI drives, but it can be done with other drives. RAID can be implemented in software or in hardware, but hardware RAID is more efficient and offers higher performance at an increased cost.

There are several types of RAID. The following are the most commonly used RAID levels:

RAID 0 Also known as *disk striping*, where a striped set of equal space from at least two drives creates a larger volume. This is in contrast to unequal space on multiple disks being used to create a simple *volume set*, which is not RAID 0. *RAID 0* is not RAID in every sense because it doesn't provide the fault tolerance implied by the *redundant* component of the name. Data is written across multiple drives, so one drive can be reading or writing while the next drive's read-write head is moving. This makes for faster data access. However, if any one of the drives fails, all content is lost. Some form of redundancy or fault tolerance should be used in concert with RAID 0.

RAID 1 Also known as *disk mirroring*. *RAID 1* is a method of producing fault tolerance by writing all data simultaneously to two separate drives. If one drive fails, the other contains all the data and will become the primary drive. However, disk mirroring doesn't help access speed, and the cost is double that of a single drive. If a separate host adapter is used for the second drive, the term *duplexing* is attributed to RAID 1. Only two drives can be used in a RAID 1 array.

RAID 5 Combines the benefits of both RAID 0 and RAID 1, creating a redundant striped volume set. Unlike RAID 1, however, *RAID 5* does not employ mirroring for redundancy. Each stripe places data on *n*–1 disks, and parity computed from the data is placed on the remaining disk. The parity is interleaved across all the drives in the array so that neighboring stripes have parity on different disks. If one drive fails, the parity information for the stripes that lost data can be used with the remaining data from the working drives to derive what was on the failed drive and rebuild the set once the drive is replaced.

The same process is used to continue to serve client requests until the drive can be replaced. The loss of an additional drive, however, results in a catastrophic loss of all data in the array. Note that while live requests are served before the array is rebuilt, nothing needs to be computed for stripes that lost their parity. Recomputing parity for these stripes is required only when rebuilding the array. A minimum of three drives is required for RAID 5. The equivalent of one drive is lost for redundancy. The more drives in the array, the less of a percentage this single disk represents.

Although there are other implementations of RAID, such as RAID 3 and RAID 4, the three detailed here are by far the most prolific. RAID 6 is somewhat popular as well because it is essentially RAID 5 with the ability to lose two disks and still function. RAID 6 uses the equivalent of two parity disks as it stripes its blocks across all disks in a fashion similar to the way RAID 5 does. A minimum of four disks is required to make a RAID 6 array.

There are also nested or hybrid implementations, such as *RAID 10* (also known as RAID 1+0), which adds fault tolerance to RAID 0 through the RAID 1 mirroring of each disk in the RAID 0 striped set. Its inverse, known as RAID 0+1, mirrors a complete striped set to another striped set just like it. Both of these implementations require a minimum of four drives and, because of the RAID 1 component, burn half of your purchased storage space for mirroring.

Removable Storage and Media

Many additional types of storage are available for PCs today. Among the other types of storage are tape backup devices, solid-state memory, and advanced optical drives. There are also external hard drives and optical drives as well as new storage media, such as USB keys that can store many gigabytes (more all the time) on a single small plastic device that can be carried on a lanyard around your neck or on a keychain.

Removable storage once meant something vastly different from what it means today. Sequential tape backup is one of the only remnants of the old forms of removable storage that can be seen in the market today. The more modern solution is random-access, solid-state removable storage. The following sections present details of tape backup and the newer removable storage solutions.

Tape Backup Devices

An older form of removable storage is the tape backup. Tape backup devices can be installed internally or externally and use either a digital or analog magnetic tape medium instead of disks for storage. They hold much more data than any other medium but are also much slower. They are primarily used for batch archival storage, not interactive storage.

With hard disks, it's not a matter of "if they fail"; it's "when they fail." So you must back up the information onto some other storage medium. Tape backup devices were once the most common choice in larger enterprises and networks because they were able to hold the most data and were the most reliable over the long term. Today, however, tape backup systems are seeing competition from writable and rewritable optical discs, which continue to advance in technology and size. Nevertheless, when an enterprise needs to back up large amounts of data on a regular basis, some form of tape media is the most popular choice. Table 2.4 lists the best-known tape formats in order of market release dates, oldest first, and capacities they are known for. Note that capacities are not associated with the format names but instead with models of tape within each format family.

TABLE 2.4 Sequential tape formats

Format Name	Representative Capacity
Quarter-inch Cartridge (QIC)	200KB to 525MB
Digital Linear Tape (DLT)	Up to 160GB
Eight Millimeter (Exabyte)	Up to 800GB
Digital Audio Tape (DAT)/Digital Data Storage (DDS)	Up to 300GB
Linear Tape-Open (LTO)	Up to 1.5TB (12.8TB planned)

Flash Memory

Once only for primary memory usage, the same components that sit on your motherboard as RAM can be found in various physical sizes and quantities in today's solid-state storage solutions. These include older removable and nonremovable flash memory mechanisms, Secure Digital (SD) cards and other memory cards, and USB flash drives. Each of these technologies has the potential to reliably store a staggering amount of information in a minute form factor. Manufacturers are using innovative packaging for some of these products to provide convenient transport options (such as keychain attachments) to users. Additionally, recall the SSD alternatives to magnetic hard drives mentioned earlier in this chapter.

For many years, modules and PC Card devices known as *flash memory* have offered low- to mid-capacity storage for devices. The name comes from the concept of easily being able to use electricity to instantly alter the contents of the memory. The original flash memory is still used in devices that require a nonvolatile means of storing critical data and code often used in booting the device, such as routers and switches.

For example, Cisco Systems uses flash memory in various forms to store its Internetwork Operating System (IOS), which is accessed from flash during boot-up and, in certain cases, throughout operation uptime and therefore during an administrator's configuration sessions. Lesser models store the IOS in compressed form on the flash and then decompress the IOS into RAM, where it is used during configuration and operation. In this case, the flash is not accessed again after the boot-up process is complete, unless its contents are being changed, as in an IOS upgrade. Certain devices use externally removable PC Card technology as flash for similar purposes.

The following sections explain a bit more about today's most popular forms of flash memory, memory cards, and USB flash drives.

SD and Other Memory Cards

Today's smaller devices require some form of removable solid-state memory that can be used for temporary and permanent storage of digital information. Gone are the days of using microfloppies in your digital camera. Even the most popular video-camera media, such as mini-DVD and HDD, are giving way to solid-state multi-gigabyte models. These more modern electronics, as well as most contemporary digital still cameras, already use some form of removable memory card to store still images permanently or until they can be copied off or printed out. Of these, the Secure Digital (*SD*) format has emerged as the preeminent leader of the pack, which includes the older *MultiMediaCard (MMC)* format on which SD is based. Both of these cards measure 32mm by 24mm, and slots that receive them are often marked for both. The SD card is slightly thicker than the MMC and has a write-protect notch (and often a switch to open and close the notch), unlike MMC. Figure 2.15 is a photo of an older SD card with size reference.

Even smaller devices, such as mobile phones, have an SD solution for them. One of these products, known as *miniSD*, is slightly thinner than SD and measures 21.5mm by 20mm. The other, *microSD*, is thinner yet and only 15mm by 11mm. Both of these reduced formats have adapters allowing them to be used in standard SD slots.

Table 2.5 lists additional memory card formats, the slots for some of which can be seen in the images that follow the table.

FIGURE 2.15 A typical SD card

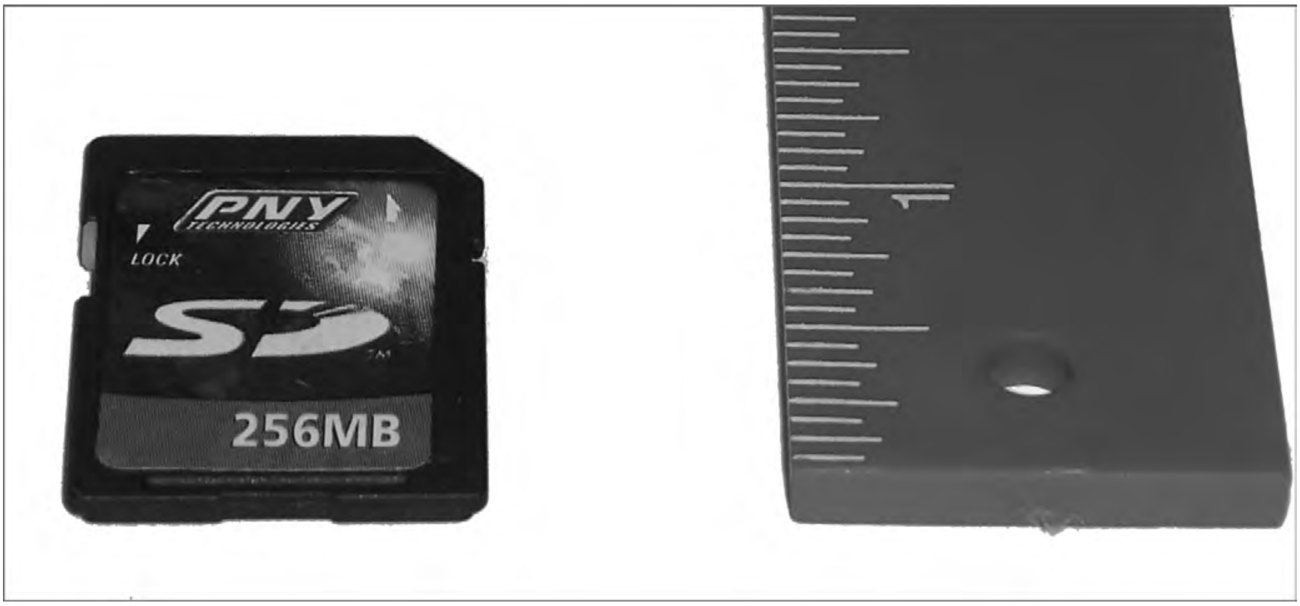

TABLE 2.5 Additional memory card formats

Format	Dimensions	Details	Year Introduced
Subscriber Identity Module (SIM)	25mm by 15mm	Used to store a subscriber's key on a telephone	1991
CompactFlash (CF)	36mm by 43mm	Type I and Type II variants; Type II used by IBM for Microdrive	1994
SmartMedia (SM)	45mm by 37mm	From Toshiba; intended to replace floppies	1995
Memory Stick (MS)	50mm by 21.5mm	From Sony; standard, Pro, Duo, and Micro formats available	1998
xD-Picture Card	20mm by 25mm	Used primarily in digital cameras	2002

Figure 2.16 shows the memory-card slots of an HP PhotoSmart printer, which is capable of reading these devices and printing from them directly or creating a drive letter for access to the contents over its USB connection to the computer. Clockwise from upper left, these slots accommodate CF/Microdrive, SmartMedia, Memory Stick (bottom right), and MMC/SD. The industry provides almost any adapter or converter to allow the various formats to work together.

FIGURE 2.16 Card slots in a printer

Many other devices exist for allowing access to memory cards. For example, 3½″ form factor devices can be purchased—some of which have multiformat floppy drives embedded—and installed in a standard front-access drive bay. One such device is shown in Figure 2.17. External card readers, such as the USB-attached one shown in Figure 2.18 (front first, then back), are widely available in many different configurations.

FIGURE 2.17 An internal card reader with floppy drive

Many of today's laptops have built-in memory-card slots, such as the ones shown in Figure 2.19.

USB Flash Drives

USB flash drives are incredibly versatile and convenient devices that allow you to store large quantities of information in a very small form factor. Many such devices are merely extensions of the host's USB connector, extending out from the interface but adding little to its width, making them easy to transport, whether in a pocket or laptop bag. Figure 2.20 illustrates an example of one of these components and its relative size.

FIGURE 2.18 A USB card reader

FIGURE 2.19 Memory-card slots in a laptop

FIGURE 2.20 A USB flash drive

USB flash drives capitalize on the versatility of the USB interface, taking advantage of the Plug and Play feature and the physical connector strength. Upon insertion, these devices announce themselves to Windows Explorer as removable drives and show up in the Explorer window with a drive letter. This software interface allows for drag-and-drop copying and most of the other Explorer functions performed on standard drives. Note that you might have to use the Disk Management (discussed in Chapter 12, "Operating System Basics") utility to manually assign a drive letter to a USB flash drive if it fails to acquire one itself. This can happen in certain cases, such as when the previous letter the drive was assigned has been taken by another device in the USB flash drive's temporary absence.

USB flash drives emerged as the de facto replacement for other, now legacy, removable storage devices such as floppies, edging out Zip and Jaz offerings from Iomega as well as other proprietary solutions for the honor.

Externally Attached Drives

Before USB, an external drive used a SCSI or proprietary adapter and interface/cable combination or the standard RS-232 serial or parallel printer port often built in to the computer. Since USB, there has been a sense that there was no other way to attach drives externally. The fact is, there were and are other ways, but USB remains ubiquitous in today's systems and continues to increase in speed with each new revision. Nevertheless, FireWire and eSATA present higher performance options to those with the interest.

USB-Attached External Disk Drives

USB devices that comply with the USB mass storage device class (USB MSC or UMS) specification are recognized as drives by the operating system upon connection, and if the external drive is to be used as a backup location, you simply install any additional software you want to use. Windows Vista and Windows 7 have built-in backup utilities that are forms of drive imaging software and that work well with external drives.

Many external optical and hard disk drives today are manufactured into their own chassis and have detachable connectivity for USB (and/or FireWire through the Serial Bus Protocol 2 [SBP-2]). If the power requirement for the unit is high enough, there might also be a separate power connection for the device. Otherwise, the USB or FireWire interface on the host provides all the power for the drive. Figure 2.21 is a photo of a small external hard disk drive with a USB interface and no separate power attachment.

FIGURE 2.21 A self-contained external hard disk drive

More flexibly, USB-attached external disk drives can use the same drives that you might install in a drive bay in your chassis; they simply employ a specialty chassis that houses only the drive and the supporting circuitry that converts the drive interface to USB. Almost always, the 3.5″ drive enclosure has a DC power input and a Type-B USB interface, as shown in Figure 2.22. This external chassis has its cover removed, and you can see the internal protective casing, inside of which the hard drive is mounted.

FIGURE 2.22 External drive enclosure

eSATA-Attached External Disk Drives

Having sung the praises of USB as the savior of the external drive market, let's dispense with the illusion. An external drive-attachment technology based on SATA, called *eSATA* for "external" SATA, promises to offer external attachment with no compromises. Where the very nature of USB can hinder the achievement of maximum SATA drive performance, eSATA, by its nature, pledges to represent SATA faithfully, because it *is* SATA. Many enhancements over the SATA physical interface and signal levels, however, were required with eSATA to accommodate the harsher external environment. A different interface, without the recognizable L-shaped key, had to be specified to avoid accidental or intentional insertion of inadequately shielded internal cables.

Did we say "no compromises"? There is one, and it might be a big one, depending on the application. eSATA doesn't provide power the way USB and FireWire do. External power has to be supplied to the drive outside of the 2m or shorter data cable. The eSATA specification is likely to ride the coattails of internal SATA to appreciable success. Figure 2.23 shows an eSATA interface on an external drive chassis from two angles, followed by the associated cable.

FIGURE 2.23 External drive and cable with eSATA interface

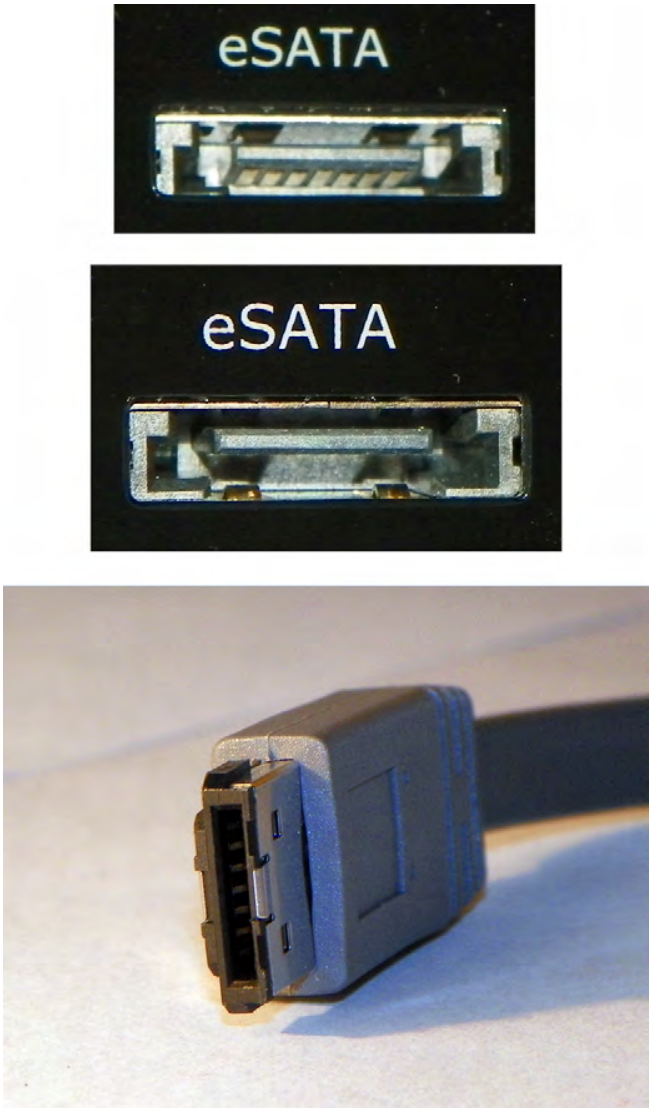

Network Attached Storage

Chapter 6, "Networking Fundamentals," discusses network attached storage (NAS) in more detail, but it bears mentioning in this section that you can use the same Ethernet connection you use to put your computer on the network to place a specialty file-serving appliance on the same Ethernet network. This appliance is essentially a stand-alone computer with one or more internal drives and self-contained intelligence for attaching to the network and supporting common file-sharing protocols without the need for an operating system license. If your computer uses server message block (SMB), for instance, to request resources from a file server, the NAS appliance can field that request through its mutual support for SMB.

Hot-Swappable Devices

Many of the removable storage devices mentioned are *hot swappable*. This means that you can insert and remove the device with the system powered on. Most USB-attached devices without a file system fall into this category. Non-hot-swappable devices, in contrast, either cannot have the system's power applied when they are inserted or removed or they have some sort of additional conditions for their insertion or removal. One subset is occasionally referred to as cold swappable, the other as warm swappable. The system power must be off before you can insert or remove cold-swappable devices. An example of a cold-swappable device is anything connected to the PS/2-style mini-DIN connector, such as a keyboard or mouse. Insertion with the power on generally results in lack of recognition for the device and might damage the motherboard. AT keyboards and the full-sized DIN connector have the same restriction.

Warm-swappable devices include USB flash drives and external drives that have a file system. Windows and other operating systems tend to leave files open while accessing them and write cached changes to them at a later time, based on the algorithm in use by the software. Removing such a device without using the Safely Remove Hardware utility can result in data loss. However, after stopping the device with the utility, you can remove it without powering down the system, hence the *warm* component of the category's name. These are officially hot-swappable devices.

RAID systems benefit from devices and bays with a single connector that contains both power and data connections instead of two separate connectors. This is known as Single Connector Attachment (SCA). SCA interfaces have ground leads that are longer than the power leads so that they make contact first and lose contact last. SATA power connectors are designed in a similar fashion for the same purpose. This arrangement ensures that no power leads make contact without their singular ground leads, which would often result in damage to the drive. Drives based on SCA are hot swappable. RAID systems that have to be taken offline before drives are changed out but the system power can remain on are examples of warm swapping.

Installing, Removing, and Configuring Storage Devices

The removal and installation of storage devices, such as hard drives, floppy drives, CD/DVD drives, and tape drives, is pretty straightforward. There really isn't any deviation in the process of installing or exchanging the hardware. Fortunately, with today's operating systems, little to no configuration is required for such devices. The Plug and Play BIOS and operating system work together to recognize the devices. However, you still have to partition and format out-of-the-box hard drives before they will allow the installation of the operating system. Nevertheless, today's operating systems allow for a pain-free partition/format/setup experience by handling the entire process if you let them.

Removing Storage Devices

Removing any component is frequently easier than installing the same part. Consider the fact that most people could destroy a house, perhaps not safely enough to ensure their well-being, but they don't have to know the intricacies of construction to start smashing away. On the other hand, very few people are capable of building a house. Similarly, many could figure out how to remove a storage device, as long as they can get into the case to begin with, but only a few could start from scratch and successfully install one without tutelage.

In Exercise 2.1, you'll remove an internal storage device.

NOTE This section details the removal of internal storage devices, and the section "Installing Storage Devices" details their installation. Be aware that external storage devices exist, but today's external storage devices are eSATA-, USB-, and FireWire-attached devices, making them completely Plug and Play. Only the software preparation of external hard drives is a consideration, but the same procedure for the software preparation of internal devices works for external devices as well.

EXERCISE 2.1

Removing an Internal Storage Device

1. With the power source removed from the system, ground yourself and the computer to the same source of ground.
2. Remove the cover from the system, exposing the internal components.
3. Unplug all connections from the storage device you wish to remove. These include data and power connections as well as any others, such as audio connections to the sound card or motherboard. The beveled Molex power connectors fit very tightly, so don't worry about how hard removing them seems to be. There is no clip to release. Do, however, be sure to grip the connector, not the wires.
4. Gather the appropriate antistatic packaging to plan ahead for all static-sensitive components that will be reused in the future, including any adapter cards that the storage device plugs into.
5. Remove any obstructions that might hinder device removal, such as component cables attached to adapter cards or adapter cards themselves, storing them to be reused in antistatic packaging.
6. Remove related adapter cards from the motherboard, storing them to be reused in antistatic packaging.
7. Remove the machine screws holding the storage device to the chassis. These could be on the side of the device or on the bottom.
8. Some devices, especially hard drives because they have no front access from the case, pull out of the chassis toward the rear of the case, while others, such as CD/DVD and floppy drives, generally pull out from the front. A gentle nudge from the rear of the device starts it on its way out the front. Go ahead and remove the device from the case. If you discover other components that obstruct the storage device's removal, repeat step 5.

Installing Storage Devices

An obvious difference among storage devices is their *form factor*. This is the term used to describe the physical dimensions of a storage device. Form factors commonly have the following characteristics:

- 3½″ wide vs. 5¼″ wide
- Half height vs. full height vs. 1″ high and more
- Any of the laptop specialty form factors

You will need to determine whether you have an open bay in the chassis to accommodate the form factor of the storage device you want to install. Adapters exist that allow a device of small size to fit into a larger bay. For obvious reasons, the converse is not also true.

In Exercise 2.2, you'll install an internal storage device.

EXERCISE 2.2

Installing an Internal Storage Device

1. With the power source removed from the system, ground yourself and the computer to the same source of ground.
2. Remove the cover from the system, exposing the internal components.
3. Locate an available bay for your component, paying attention to your device's need for front access. If you do not see one, look around; some cases provide fastening points near the power supply or other open areas of the case. If you still do not see one, investigate the possibility of sacrificing a rarely or never used device to make room.
4. Remove any obstructions that might hinder device installation, such as component cables attached to adapter cards or adapter cards themselves, storing them to be reused in antistatic packaging.
5. Find the proper screws for the storage device and set any jumpers on the drive while it is in-hand. Then insert the device into the bay. Keep in mind that some insert from the rear of the bay and some from the front.
6. Line up the screw holes in the device with the holes in the bay. Note that many devices rarely insert as far as they can before lining up with the chassis's holes. So don't be surprised when pushing the device all the way into the bay results in misalignment. Other devices that require front access stop themselves flush with the front of the case, and still others require you to secure them while holding them flush.
7. Use at least two screws on one side of the device. This keeps the device from sliding in the bay as well as from rotating, which happens when you use only one screw or one screw on each side. If the opposite side is accessible, go ahead and put at least one screw in the other side. Most devices allow for as many as four screws per side, but eight screws are not necessary in the vast majority of situations.

EXERCISE 2.2 *(continued)*

8. Connect the data cable from the device to the adapter card or motherboard header. ATA devices, such as those that are designated as IDE drives (compatible hard drives and CD/DVD drives, for example) use a 40-pin connector. Floppy drives and some tape backup drives that connect through the floppy subsystem use a 34-pin connector. They look the same except for the three rows of two pins that differentiate them. Note that if you use the master/slave and not the Cable Select feature of IDE drives on the same chain, it does not matter which device connects to which connector on the cable. However, with floppy drives, the A: drive must always be attached to the connector after the twist in the cable.
9. Attach a power connector from the power supply to the device, bearing in mind that there are two connector styles that are not very close in appearance. You should have no trouble telling them apart. Be sure to fully insert the connector. Watch out for the smaller connector. (See the sidebar "Do You Smell Something?")

Real World Scenario

Do You Smell Something?

In 1990, author Toby Skandier started a PC sales and repair business. Those were the days when you could build a computer from scratch for relatively little cost and sell it with a great markup and still come in way under the prices of the name-brand systems.

One customer was especially price conscious. In those days, a floppy drive was not the afterthought that it is today, both in use and price. You needed a floppy drive and could actually save a bit of money if you were buying quite a few units, just by opting for a cheaper model. This customer was buying 45 computers. So, one of the corners that was cut to keep the invoice amount down was floppy-drive quality. They went with a brand that Toby had never heard of but that his distributor listed as the cheapest. How bad could it be? How much of a difference could there be between brands and models? Toby found out. The customer didn't.

The cheaper drive worked just like any other, from the perspective of the user, but the difference showed while they were building the systems. The manufacturer scrimped in the production of the power connector. Where most manufacturers create a casing to receive the power supply's connector with little chance of inserting the connector upside down, this manufacturer allowed the four pins of the connector to protrude in a nondescript manner without any keying or guidance for the power supply's connector.