


MODULE II




INFORMATION TECHNOLOGIES

What challenges do information system technologies pose for business professionals? What basic knowledge should you possess about information technology? The four chapters of this module give you an overview of the hardware, software, and data resource management and telecommunications network technologies used in information systems and their implications for business managers and professionals.

- **Chapter 3: Computer Hardware** reviews history, trends, and developments in microcomputer, midrange, and mainframe computer systems; basic computer system concepts; and the major types of technologies used in peripheral devices for computer input, output, and storage.
- **Chapter 4: Computer Software** reviews the basic features and trends in the major types of application software and system software used to support enterprise and end-user computing.
- **Chapter 5: Data Resource Management** emphasizes management of the data resources of computer-using organizations. This chapter reviews key database management concepts and applications in business information systems.
- **Chapter 6: Telecommunications and Networks** presents an overview of the Internet and other telecommunication networks, business applications, and trends and reviews technical telecommunications alternatives.

CHAPTER 3



COMPUTER HARDWARE

Chapter Highlights

Section I Computer Systems: End User and Enterprise Computing

Introduction

A Brief History of Computer Hardware

Real World Case: AstraZeneca, UnitedHealth, and Others: IT Asset Management—Do You Know What You’ve Got?

Types of Computer Systems

Microcomputer Systems

Midrange Systems

Mainframe Computer Systems

Technical Note: The Computer System Concept

Moore’s Law: Where Do We Go from Here?

Section II Computer Peripherals: Input, Output, and Storage Technologies

Peripherals

Input Technologies

Real World Case: IT in Healthcare: Voice Recognition Tools Make Rounds at Hospitals

Output Technologies

Storage Trade-Offs

Semiconductor Memory

Magnetic Disks

Magnetic Tape

Optical Disks

Radio Frequency Identification

Predictions for the Future

Real World Case: IBM, Wachovia, and PayPal: Grid Computing Makes It Easier and Cheaper

Real World Case: Apple, Microsoft, IBM, and Others: The Touch Screen Comes of Age

Learning Objectives

1. Understand the history and evolution of computer hardware.
2. Identify the major types and uses of microcomputer, midrange, and mainframe computer systems.
3. Outline the major technologies and uses of computer peripherals for input, output, and storage.
4. Identify and give examples of the components and functions of a computer system.
5. Identify the computer systems and peripherals you would acquire or recommend for a business of your choice, and explain the reasons for your selections.

REY
NOLD
S
,
D
O
N
A
L
D
W
A
Y
N
E
2
8
3
3
B
U

SECTION I

Computer Systems: End User and Enterprise Computing

Introduction

All computers are systems of input, processing, output, storage, and control components. In this section, we discuss the history, trends, applications, and some basic concepts of the many types of computer systems in use today. In Section II, we will cover the changing technologies for input, output, and storage that are provided by the peripheral devices that are part of modern computer systems.

Read the Real World Case regarding management of IT assets. We can learn a lot about how different organizations keep track of their IT assets and manage their life-cycle from this case. See Figure 3.1.

A Brief History of Computer Hardware

Today we are witnessing rapid technological changes on a broad scale. However, many centuries elapsed before technology was sufficiently advanced to develop computers. Without computers, many technological achievements of the past would not have been possible. To fully appreciate their contribution, we must understand their history and evolution. Although a thorough discussion of computing history is beyond the scope of this text, a brief consideration of the development of the computer is possible. Let's look quickly into the development of computers.

At the dawn of the human concept of numbers, humans used their fingers and toes to perform basic mathematical activities. Then our ancestors realized that by using some objects to represent digits, they could perform computations beyond the limited scope of their own fingers and toes. Can't you just see in your mind a cave full of cave-men performing some group accounting function using their fingers, toes, sticks, and rocks? It creates a comical, yet accurate, picture.

Shells, chicken bones, or any number of objects could have been used, but the fact that the word *calculate* is derived from *calculus*, the Latin word for "small stone," suggests that pebbles or beads were arranged to form the familiar abacus, arguably the first human-made computing device. By manipulating the beads, it was possible with some skill and practice to make rapid calculations.

Blaise Pascal, a French mathematician, invented what is believed to be the first mechanical adding machine in 1642. The machine partially adopted the principles of the abacus but did away with the use of the hand to move the beads or counters. Instead, Pascal used wheels to move counters. The principle of Pascal's machine is still being used today, such as in the counters of tape recorders and odometers. In 1674, Gottfried Wilhelm von Leibniz improved Pascal's machine so that the machine could divide and multiply as easily as it could add and subtract.

When the age of industrialization spread throughout Europe, machines became fixtures in agricultural and production sites. An invention that made profound changes in the history of industrialization, as well as in the history of computing, was the mechanical loom, invented by a Frenchman named Joseph Jacquard. With the use of cards punched with holes, it was possible for the Jacquard loom to weave fabrics in a variety of patterns. Jacquard's loom was controlled by a program encoded into the punched cards. The operator created the program once and was able to duplicate it many times over with consistency and accuracy.

The idea of using punched cards to store a predetermined pattern to be woven by the loom clicked in the mind of Charles Babbage, an English mathematician who lived in the 19th century. He foresaw a machine that could perform all mathematical calculations, store values in its memory, and perform logical comparisons among values. He called it the *Analytical Engine*. Babbage's analytical engine, however, was never built. It lacked one thing: electronics. Herman Hollerith eventually adapted Jacquard's concept of the punched card to record census data in the late 1880s.

REAL WORLD

CASE

1

AstraZeneca, UnitedHealth, and Others: IT Asset Management—Do You Know What You’ve Got?

Global pharmaceuticals giant AstraZeneca needed some strong medicine of its own to fix a burgeoning IT asset management problem. It was brought about by multiple acquisitions and their nonstandard gear, a high-tech workforce spread across 255 facilities in 147 countries, and a total of more than 67,000 employees using more than 90,000 hardware and software assets ranging from notebooks to SAP and Oracle enterprise applications and databases.

With software vendors becoming more aggressive on audits as sales of new products are generally weak, and with greater internal collaboration requiring a more consistent set of tools to simplify processes and maintenance, the \$31 billion pharmaceuticals company realized a few years ago that Microsoft’s Systems Management Server was simply overmatched for the job of managing the global enterprise’s complex base of IT assets.

So Microsoft recommended the asset management products offered by a French company called PS’Soft, which is a subsidiary of BDNA Corp., a top provider of IT infrastructure inventory and analysis solutions. In the years that AstraZeneca has been steadily getting its IT assets under control, PS’Soft has distinguished itself like few other IT vendors, according to AstraZeneca Global IT Asset lead Bernard Warrington.

“In all my years, our engagement with PS’Soft was one of the first and only times we had an IT vendor show such willingness to work as a true partner and really try to solve our problems with us,” Warrington says. Referring to PS’Soft’s Julian Moreau, Warrington described the uniquely

open collaboration that allowed him and his team to understand the problem, design the solution, and then execute the plan.

“Julian and I worked extremely closely together, and from there our partnership cascaded down to the other members of the team,” Warrington says. “But I need to be very clear about that: in the beginning, the knowledge and expertise were clearly with them—they were teaching and we were learning.”

The problem, Warrington says, is that in the increasingly strategic world of IT asset management, “the toolset itself meets only 30 percent of the overall need: on top of that, you need to build the processes, understand the costs, come up with standards, develop interfaces with other major vendors, and much more—we simply didn’t have all the skills necessary to cover that total lifecycle. But PS’Soft did have those skills, both in-house and through their contacts.”

In addition, says Warrington, PS’Soft and BDNA had the global experiences necessary to help AstraZeneca get its arms around its global sprawl of IT gear, which was essential so the company could (1) begin to gain greater leverage in purchasing negotiations, and (2) be able to fairly, but aggressively, hold its own during audits by software vendors.

“In so many countries where we operate, the tradition has been that budgets are managed locally, making it impossible to see the global aggregate in detail,” Warrington says. “We simply did not have the ability to get a global view. The old tools we used gave us something of a snapshot, but didn’t let us have enough insight to be able to manage the situation. At the same time, the IT vendors are getting very aggressive with audits, and without offering a specific number I can tell you that millions and millions of dollars are at stake—and before our engagement with PS’Soft, no matter how hard we tried with the old toolset, we were just not able to achieve those potential cost savings from vendors.”

Over time, Warrington says, AstraZeneca gained that necessary level of control and knowledge: “Now AstraZeneca is in a position to enter negotiations from a position of strength, confidence, and knowledge.” And that achievement has given the company a new perspective on the realm of IT asset management. Warrington says, “Too many companies just look on IT asset management as nothing more than bean counting, versus looking deeper and understanding the ROI and ROA that can be achieved.

“But we learned first hand that there is a huge opportunity to get control over what you have, to satisfy even the most rigorous audit, and to negotiate better contracts. And that’s a lot more than bean counting,” says Warrington.

IT organizations in diversified companies—particularly those grown through acquisition—wage a seemingly endless battle against unnecessary IT diversity and related costs. Conceived, planned, and executed in 18 months, UnitedHealth Group’s (UHG) Hercules program proves



FIGURE 3.1

Companies are increasingly focusing on managing the myriad of platforms, hardware, and software that make up their IT infrastructures.

that the complexity can be conquered, while protecting or improving IT's service levels. By creating a standard desktop configuration and consistent management processes, Hercules reduced total cost of ownership to \$76 per month per desktop, from more than \$240.

In 2004, with the CEO's support, Alistair Jacques, then SVP of UHG-IT, launched Hercules, focusing it on standardizing and streamlining the processes behind desktop management: procurement, configuration, installation, life cycle, and asset management. In addition to this focus on process, two techniques stand out as key to the program's success. Working with finance, IT developed a chargeback model that imposes a premium on nonstandardized desktop configurations: \$170 per month versus \$45 per month for a standard configuration. This value price encourages business managers to choose the more efficient infrastructure. UHG also reduced costly on-site support by reorganizing it: A central IT team manages high-level support activities, completing 95 percent remotely, while select, on-site end users (often non-IT administrative staff trained by IT) provide basic support to colleagues.

UHG-IT treated desktop management as a business process challenge rather than a technology issue. This approach freed them to use tactics like non-IT staff for desktop support and value pricing. To date, UHG has converted 75,000 out of 90,000 devices to the new standards, delivering \$42 million in annual savings. UHG can now manage nearly four times the number of end users with the same number of IT personnel as in 2004, all while actually improving—not diminishing—service levels. IT now deploys 99.4 percent of releases, updates, and patches in three hours, instead of 65 percent in three weeks.

Indeed, companies that blow off asset management do so at their own peril. At the same time, 99 percent of companies that her organization comes across don't have a proper asset management process in place, according to Elisabeth Vanderveldt, vice president of business development at Montreal-based IT services and consulting firm Conamex International Software Corp.

That's a staggering number, considering the value that life-cycle management can bring to an organization. And it's

indicative of the widespread lack of respect for this important aspect of IT operations.

The ideal time to start considering an asset management program is before the business and its IT infrastructure is even up and running, but the common scenario is that corporations look to asset management after they've encountered a problem running the infrastructure.

Businesses' mentality about asset management is evolving, however. Companies used to consider only reliability, availability, and overall equipment effectiveness in the equation. But now, he said, there is recognition of factors like continuing pressures on cost and green technology. "It really requires a mature organization to understand what's going to be needed to assess and execute a life-cycle management strategy," says Don Barry, associate partner in global business services in the supply chain operations and asset management solutions at IBM.

Why is a life-cycle management program important? For one thing, it puts IT in much better control of its assets, and this can have a number of benefits.

"IT can make really intelligent decisions around what they should get rid of, and they might even find they have more money in the budget and they can start taking a look at newer technology and see if they can bring it in-house. Without that big picture, they just end up spending more and more money than had they been proactive," says Vanderveldt.

Life-cycle management also has value as a risk management tool, and it aids in the disaster recovery process as well, she adds. "It's also beneficial for those moments that are just completely out of your control, like mergers, acquisitions and uncontrolled corporate growth, either organic or inorganic," says Darin Stahl, an analyst at London, Ontario based Info-Tech Research Group. "IT leaders without this tool set are now charged with pulling all this information together on short notice. That could be diminished considerably in terms of turnaround time and effort for IT guys if they have a holistic asset management program in place."

Source: Adapted from Bob Evans, "Global CIO Quick Takes: AstraZeneca Saves Millions with BDNA," *InformationWeek*, February 22, 2010; Rick Swanborg, "Desktop Management: How UnitedHealth Used Standardization to Cut Costs," *CIO.com*, April 28, 2009; and Kathleen Lau, "Asset Management: Do You Know What You've Got?," *CIO Canada*, August 13, 2008.

CASE STUDY QUESTIONS

1. What are the companies mentioned in the case trying to control, or manage, through these projects? What is the problem? And how did they get there?
2. What are the business benefits of implementing strong IT asset management programs? In what ways have the companies discussed in the case benefited? Provide several examples.
3. One of the companies in the case, UnitedHealth Group, tackled the issue by imposing standardization and "charging" those stepping outside standard models. How should they balance the need to standardize with being able to provide business units with the technologies best suited to their specific needs? Justify your answer.

REAL WORLD ACTIVITIES

1. An important metric in this area considered by companies is the Total Cost of Ownership (TCO) of their IT assets. Go online and research what TCO is and how it is related to IT asset management. How are companies using TCO to manage their IT investments? Prepare a presentation to share your research with the rest of your class.
2. What does Don Barry of IBM mean by "life-cycle" in the context of this case? How would this life-cycle management work when it comes to IT assets? Break into small groups with your classmates and create a working definition of life-cycle management and how it works as you understand it from the case.

Census data were translated into a series of holes in a punched card to represent the digits and the letters of the alphabet. The card was then passed through a machine with a series of electrical contacts that were either turned off or on, depending on the existence of holes in the punched cards. These different combinations of off/on situations were recorded by the machine and represented a way of tabulating the result of the census. Hollerith's machine was highly successful; it cut the time it took to tabulate the result of the census by two-thirds, and it made money for the company that manufactured it. In 1911, this company merged with its competitor to form International Business Machines (IBM).

The ENIAC (Electronic Numerical Integrator and Computer) was the first electronic digital computer. It was completed in 1946 at the Moore School of Electrical Engineering of the University of Pennsylvania. With no moving parts, ENIAC was programmable and had the capability to store problem calculations using vacuum tubes (about 18,000).

A computer that uses vacuum tube technology is called a first-generation computer. The ENIAC could add in 0.2 of a millisecond, or about 5,000 computations per second. The principal drawback of ENIAC was its size and processing ability. It occupied more than 1,500 square feet of floor space and could process only one program or problem at a time. As an aside, the power requirements for ENIAC were such that adjacent common area lighting dimmed during the power up and calculation cycles. Figure 3.2 shows the ENIAC complex.


In the 1950s, Remington Rand manufactured the UNIVAC I (Universal Automatic Calculator). It could calculate at the rate of 10,000 additions per second. In 1957, IBM developed the IBM 704, which could perform 100,000 calculations per second.

In the late 1950s, transistors were invented and quickly replaced the thousands of vacuum tubes used in electronic computers. A transistor-based computer could perform 200,000–250,000 calculations per second. The transistorized computer represents the second generation of computer. It was not until the mid-1960s that the third generation of computers came into being. These were characterized by solid-state technology and integrated circuitry coupled with extreme miniaturization.

No history of electronic computing would be complete without acknowledging Jack Kilby. Kilby was a Nobel Prize laureate in physics in 2000 for his invention of the integrated circuit in 1958 while working at Texas Instruments (TI). He is also the inventor of the handheld calculator and thermal printer. Without his work that generated a patent for a “Solid Circuit made of Germanium,” our worlds, and most certainly our computers, would be much different and less productive than we enjoy today.

FIGURE 3.2

ENIAC was the first digital computer. It is easy to see how far we have come in the evolution of computers.



Source: Photo courtesy of United States Army.

In 1971, the fourth generation of computers was characterized by further miniaturization of circuits, increased multiprogramming, and virtual storage memory. In the 1980s, the fifth generation of computers operated at speeds of 3–5 million calculations per second (for small-scale computers) and 10–15 million instructions per second (for large-scale computers).

The age of microcomputers began in 1975 when a company called MITS introduced the ALTAIR 8800. The computer was programmed by flicking switches on the front. It came as a kit and had to be soldered together. It had no software programs, but it was a personal computer available to the consumer for a few thousand dollars when most computer companies were charging tens of thousands of dollars. In 1977 both Commodore and Radio Shack announced that they were going to make personal computers. They did, and trotting along right beside them were Steve Jobs and Steve Wozniak, who invented their computer in a garage while in college. Mass production of the Apple began in 1979, and by the end of 1981, it was the fastest selling of all the personal computers. In August 1982 the IBM PC was born, and many would argue that the world changed forever as a result.

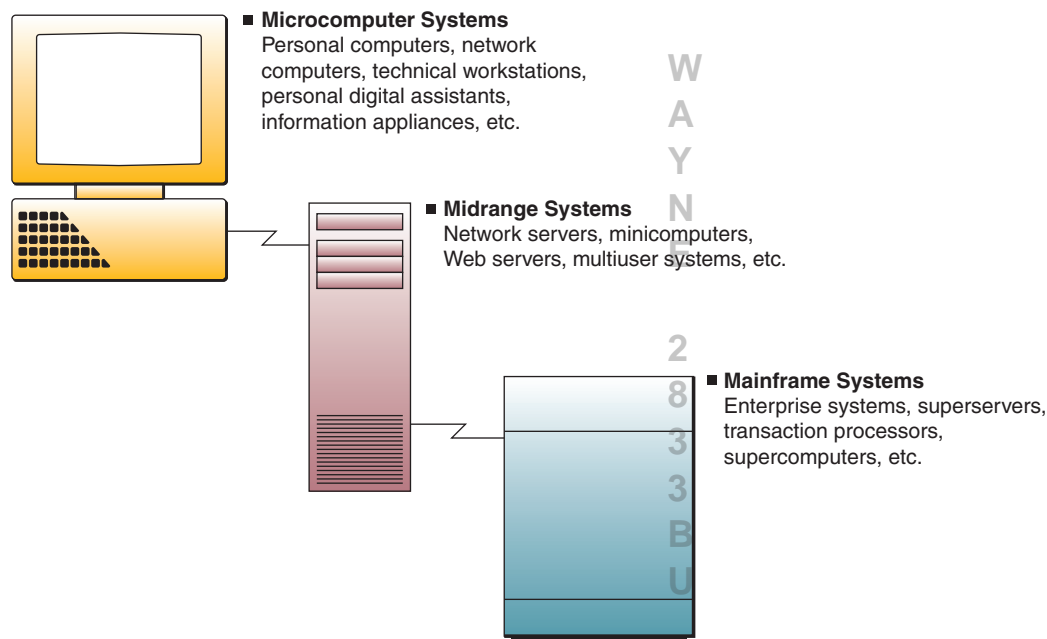
Following the introduction of the personal computer in the early 1980s, we used our knowledge of computer networks gained in the early days of computing and combined it with new and innovative technologies to create massive networks of people, computers, and data on which anyone can find almost anything: the Internet. Today we continue to see amazing advancements in computing technologies.

Okay, it's time to slow down a bit and begin our discussion of today's computer hardware.

Types of Computer Systems

Today's computer systems come in a variety of sizes, shapes, and computing capabilities. Rapid hardware and software developments and changing end-user needs continue to drive the emergence of new models of computers, from the smallest handheld personal digital assistant/cell phone combinations to the largest multiple-CPU mainframes for enterprises. See Figure 3.3.

FIGURE 3.3 Examples of computer system categories.



Source: Courtesy of Hewlett-Packard.

Categories such as *mainframe*, *midrange*, and *microcomputer* systems are still used to help us express the relative processing power and number of end users that can be supported by different types of computers. These are not precise classifications, and they do overlap each other. Thus, other names are commonly given to highlight the major uses of particular types of computers. Examples include personal computers, network servers, network computers, and technical workstations.

In addition, experts continue to predict the merging or disappearance of several computer categories. They feel, for example, that many midrange and mainframe systems have been made obsolete by the power and versatility of networks composed of microcomputers and servers. Other industry experts have predicted that the emergence of network computers and *information appliances* for applications on the Internet and corporate intranets will replace many personal computers, especially in large organizations and in the home computer market. Still others suggest that the concept of *nanocomputers* (computing devices that are smaller than micro) will eventually pervade our entire understanding of personal computing. Only time will tell whether such predictions will equal the expectations of industry forecasters.

Microcomputer Systems

The entire center of gravity in computing has shifted. For millions of consumers and business users, the main function of desktop PCs is as a window to the Internet. Computers are now communications devices, and consumers want them to be as cheap as possible.

Microcomputers are the most important category of computer systems for both businesspeople and consumers. Although usually called a *personal computer*, or PC, a microcomputer is much more than a small computer for use by an individual as a communication device. The computing power of microcomputers now exceeds that of the mainframes of previous computer generations, at a fraction of their cost. Thus, they have become powerful networked *professional workstations* for business professionals.

Consider the computing power on the *Apollo 11* spacecraft. Most certainly, landing men on the moon and returning them safely to earth was an extraordinary feat. The computer that assisted them in everything from navigation to systems monitoring was equally extraordinary. *Apollo 11* had a 2.048 MHz CPU that was built by MIT. Today's standards can be measured in the 4 GHz in many home PCs (MHz is 1 million computing cycles per second and GHz is 1 billion computing cycles per second). Further, the *Apollo 11* computer weighed 70 pounds versus today's powerful laptops weighing in as little as 1 pound. This is progress, for sure.


Microcomputers come in a variety of sizes and shapes for a variety of purposes, as Figure 3.4 illustrates. For example, PCs are available as handheld, notebook, laptop, tablet, portable, desktop, and floor-standing models. Or, based on their use, they include home, personal, professional, workstation, and multiuser systems. Most microcomputers are *desktops* designed to fit on an office desk or laptops for those who want a small, portable PC. Figure 3.5 offers advice on some of the key features you should consider when acquiring a high-end professional workstation, multimedia PC, or beginner's system. This breakdown should give you some idea of the range of features available in today's microcomputers.

Some microcomputers are powerful **workstation computers** (technical workstations) that support applications with heavy mathematical computing and graphics display demands, such as computer-aided design (CAD) in engineering or investment and portfolio analysis in the securities industry. Other microcomputers are used as **network servers**. These are usually more powerful microcomputers that coordinate telecommunications and resource sharing in small local area networks (LANs) and in Internet and intranet Web sites.

FIGURE 3.4 Examples of microcomputer systems:



a. A notebook microcomputer.
Source: Hewlett-Packard.



b. The microcomputer as a professional workstation.
Source: Corbis.



c. The microcomputer as a technical workstation.
Source: Courtesy of Hewlett-Packard.

FIGURE 3.5 Examples of recommended features for the three types of PC users. Note: www.dell.com and www.gateway.com are good sources for the latest PC features available.

Business Pro	Multimedia Heavy or Gamer	Newcomer
To track products, customers, and firm performance, more than just a fast machine is necessary: <ul style="list-style-type: none">• 3–4 GHz dual-core processor• 4–8 GB RAM• 500 GB hard drive• Up to 19-inch flat-panel display• CD-RW/DVD+RW• Network interface card• Color laser printer	Media pros and dedicated gamers will want at least a Mac G4 or a 2–3 GHz Intel dual-core chip, and <ul style="list-style-type: none">• 4–8 GB RAM• 250+ GB hard drive• 19-inch or better flat-panel display• 16× or better DVD+RW• Video cards (as fast and as powerful as budget permits)• Sound cards• Laser printer (color or B&W)	Save some money with a Celeron processor in the 2–3 GHz range while looking for <ul style="list-style-type: none">• 2 GB RAM• 120–160 GB hard drive• 15- to 17-inch flat panel or wide screen• CD-RW/DVD• USB port• Inkjet printer

Corporate PC Criteria

What do you look for in a new PC system? A big, bright screen? Zippy new processor? Capacious hard drive? Acres of RAM? Sorry, none of these is a top concern for corporate PC buyers. Numerous studies have shown that the price of a new computer is only a small part of the total cost of ownership (TCO). Support, maintenance, and other intangibles contribute far more heavily to the sum. Let's take a look at three top criteria.

Solid Performance at a Reasonable Price. Corporate buyers know that their users probably aren't mapping the human genome or plotting trajectories to Saturn. They're doing word processing, order entry, sales contact management, and other essential business tasks. They need a solid, competent machine at a reasonable price, not the latest whizbang.

Many organizations are adopting a laptop, rather than desktop, strategy. Using this approach, the employee uses his or her laptop while in the office and out in the field. With the proliferation of wireless Internet access, this strategy allows employees to take the desktop with them wherever they may be—at their desk, in a conference room, at a meeting off-site, or in a hotel room in another country.

One outcome of this strategy is the development and acquisition of more powerful laptops with larger and higher-quality screens. This demand presents a challenge to laptop manufacturers to provide higher quality while continuing to make the laptop lightweight and portable.

Operating System Ready. A change in the operating system of a computer is the most disruptive upgrade an enterprise has to face. That's why many corporate buyers want their machines to be able to handle current operating systems and anticipate new ones. Although most organizations have adopted Windows XP or Vista, some enterprises still use operating systems of an earlier vintage. Ultimately, they must be able to make the transition to Windows 7 (the newest OS from Microsoft) and even to OS versions expected three to five years from now. Primarily, that demand means deciding what hard disk space and RAM will be sufficient.

Connectivity. Networked machines are a given in corporate life, and Internet-ready machines are becoming a given. Buyers need machines equipped with reliable wireless capabilities. With fewer cables to worry about, wireless networks, especially when combined with laptop PCs, contribute to the flexibility of the workplace and the simplicity of PC deployment. Many organizations are planning for Internet-based applications and need machines ready to make fast, reliable, and secure connections.

Security-Equipped. Most of the data that is processed by networked workstations in a modern corporate environment can be considered proprietary, if not mission-critical. A major criterion for corporate purchase is the degree to which the device can accept or conform to the myriad of security measures in use in that organization. Can it accept a USB dongle, smartcard reader, biometric access device, and so forth? We will cover this aspect in greater detail in Chapter 13.

Computer Terminals

Computer terminals, essentially any device that allows access to a computer, are undergoing a major conversion to networked computer devices. *Dumb terminals*, which are keyboard/video monitor devices with limited processing capabilities, are being replaced by *intelligent terminals*, which are modified networked PCs or network computers. Also included are **network terminals**, which may be *Windows terminals* that depend on network servers for Windows software, processing power, and storage, or *Internet terminals*, which depend on Internet or intranet Web site servers for their operating systems and application software.

Intelligent terminals take many forms and can perform data entry and some information processing tasks independently. These tasks include the widespread use of **transaction terminals** in banks, retail stores, factories, and other work sites. Examples are automated teller machines (ATMs), factory production recorders, airport check-in kiosks, and retail point-of-sale (POS) terminals. These intelligent terminals use keypads, touch screens, bar code scanners, and other input methods to capture data and interact with end users during a transaction, while relying on servers or other computers in the network for further transaction processing.

Network Computers

Network computers (NCs) are a microcomputer category designed primarily for use with the Internet and corporate intranets by clerical workers, operational employees, and knowledge workers with specialized or limited computing applications. These NCs are low-cost, sealed microcomputers with no or minimal disk storage that are linked to the network. Users of NCs depend primarily on network servers for their operating system and Web browser, application software, and data access and storage.

One of the main attractions of network computers is their lower TCO (total cost of ownership), that is, the total of all costs associated with purchasing, installing, operating, and maintaining a computer. Purchase upgrades, maintenance, and support cost much less than for full-featured PCs. Other benefits to business include the ease of software distribution and licensing, computing platform standardization, reduced end-user support requirements, and improved manageability through centralized management and enterprisewide control of computer network resources.

Information Appliances

PCs aren't the only option: A host of smart gadgets and information appliances—from cellular phones and pagers to handheld PCs and Web-based game machines—promise Internet access and the ability to perform basic computational chores.

Handheld microcomputer devices known as **personal digital assistants** (PDAs) are some of the most popular devices in the **information appliance** category. Web-enabled PDAs use touch screens, pen-based handwriting recognition, or keypads so that mobile workers can send and receive e-mail, access the Web, and exchange information such as appointments, to-do lists, and sales contacts with their desktop PCs or Web servers.

Now a mainstay of PDA technology is the RIM BlackBerry, a small, pager-sized device that can perform all of the common PDA functions, plus act as a fully functional mobile telephone. What sets this device apart from other wireless PDA solutions is that it is always on and connected. A BlackBerry user doesn't need to retrieve e-mail; the e-mail finds the BlackBerry user. Because of this functionality, there is no need to dial in or initiate a connection. The BlackBerry doesn't even have a visible antenna. When a user wishes to send or reply to an e-mail, the small keyboard on the device allows text entry. Just like a mobile telephone, the BlackBerry is designed to remain on and continuously connected to the wireless network, allowing near real-time transfer of e-mail. Furthermore, because the BlackBerry uses the same network as most mobile telephone services, the unit can be used anywhere that a mobile phone can be used.

A relatively new entrant to this field (although gaining favor in leaps and bounds) is the Apple iPhone (Figure 3.6). The iPhone essentially combines three products—a revolutionary mobile phone, a wide-screen iPod music and video player with touch controls, and a breakthrough Internet communications device with desktop-class e-mail, Web browsing, maps, and searching—into one small and lightweight handheld device. The iPhone also introduces an entirely new user interface based on a large, multitouch display and pioneering new software, letting users control everything with just their fingers.

The genesis of the iPhone began with Apple CEO Steve Jobs's direction that Apple engineers investigate touch screens. Apple created the device during a secretive and unprecedented collaboration with AT&T Mobility—called Cingular Wireless at the time of the phone's inception—at a development cost of \$150 million, by one

FIGURE 3.6

The Apple iPhone—a revolutionary player in the information appliance and PDA marketplace.



Source: Lourens Smak/Alamy.

estimate. During development, the iPhone was code-named “Purple 2.” The company rejected an early “design by committee” built with Motorola in favor of engineering a custom operating system, interface, and hardware.

The iPhone went on sale on June 29, 2007. Apple closed its stores at 2:00 p.m. local time to prepare for the 6:00 p.m. iPhone launch, while hundreds of customers lined up at stores nationwide. They sold 270,000 iPhones in the first 30 hours on launch weekend.

In Germany, Deutsche Telekom signed up 70,000 iPhone customers during the 11-week period of November 9, 2007, to January 26, 2008. In the United Kingdom, it has been estimated that 190,000 customers signed with O2 during an eight-week period from the November 9, 2007 launch date to January 9, 2008.

The newest generation of iPhone is the 3G. This version accesses data from the much faster 3G network and provides for the download of literally thousands of applications that allow the iPhone to perform tasks ranging from accessing online banking services to acting as a sophisticated leveling device and everything in between.

The iPhone has truly ushered in an era of software power and sophistication never before seen in a mobile device, completely redefining what people can do on a mobile phone. We can expect to see even more sophisticated mobile PDA-type devices in the future as Moore’s law continues to prevail and the marketplace continues to demand more functionality (see the discussion on Moore’s law at the end of Section I for more details on this concept).

Information appliances may also take the form of video-game consoles and other devices that connect to your home television set. These devices enable people to surf the World Wide Web, send and receive e-mail, and watch television programs or play video games, at the same time. Other information appliances include wireless PDAs and Internet-enabled cellular and PCS phones, as well as wired, telephone-based home appliances that can send and receive e-mail and access the Web.

Midrange Systems

Midrange systems are primarily high-end network servers and other types of servers that can handle the large-scale processing of many business applications. Although not as powerful as mainframe computers, they are less costly to buy, operate, and maintain than mainframe systems and thus meet the computing needs of many organizations. See Figure 3.7.

FIGURE 3.7

Midrange computer systems can handle large-scale processing without the high cost or space considerations of a large-scale mainframe.



Source: China Foto Press/Getty Images.

Burgeoning data warehouses and related applications such as data mining and online analytical processing are forcing IT shops into higher and higher levels of server configurations. Similarly, Internet-based applications, such as Web servers and electronic commerce, are forcing IT managers to push the envelope of processing speed and storage capacity and other [business] applications, fueling the growth of high-end servers.

Midrange systems have become popular as powerful network servers (computers used to coordinate communications and manage resource sharing in network settings) to help manage large Internet Web sites, corporate intranets and extranets, and other networks. Internet functions and other applications are popular high-end server applications, as are integrated enterprisewide manufacturing, distribution, and financial applications. Other applications, like data warehouse management, data mining, and online analytical processing (which we will discuss in Chapters 5 and 10), are contributing to the demand for high-end server systems.

Midrange systems first became popular as **minicomputers** for scientific research, instrumentation systems, engineering analysis, and industrial process monitoring and control. Minicomputers could easily handle such uses because these applications are narrow in scope and do not demand the processing versatility of mainframe systems. Today, midrange systems include servers used in industrial process-control and manufacturing plants and play major roles in computer-aided manufacturing (CAM). They can also take the form of powerful technical workstations for computer-aided design (CAD) and other computation and graphics-intensive applications. Midrange systems are also used as *front-end servers* to assist mainframe computers in telecommunications processing and network management.

And the Oscar Goes to . . . Penguins and 2,000 Blade Servers

An initial implementation of 500 blade servers soon grew to 2,000 to meet the processing capacity requirements for creating the Oscar-winning animated film *Happy Feet*. The 108-minute computer-generated animated feature, which won an Academy Award in 2006, was put together by digital production company The Animal Logic Group.

“We needed huge numbers of processors in a form factor and price level that would work for our business,” says Xavier Desdoigts, director of technical operations.



“We had to render 140,000 frames, and each frame could take many hours to render. The photorealistic look of the movie made our computational requirements soar to new heights.”

For example, Mumble, the main character in the movie, had up to 6 million feathers. “There were six shots in the movie that had more than 400,000 penguins in them,” Desdoigts explained. This added up to more than 17 million CPU hours used throughout the last nine months of *Happy Feet* production. “We were initially concerned about our ability to build and manage a processing capacity of that scale.”

Animal Logic and IBM built a rendering server farm using BladeCenter HS20 blade servers, each with two Intel Xeon servers. Rendering was completed in October 2006, and the film was released the following month in the United States. Management tools to deploy and control the servers while in production included an open-source package for administering computing clusters. For Animal Logic, the biggest sign of success from an IT perspective was that the entire server farm was managed by a single person.

“We have to make sure we choose solutions that aren’t overly complex to set up or manage, so our focus can stay on realizing the creative visions of our clients,” Desdoigts said. *Happy Feet* quickly became one of the Australian film industry’s greatest box-office successes, taking the No. 1 spot in the United States for three consecutive weeks. It made more than \$41 million (U.S.) on its opening weekend and showed on 3,800 cinema screens.

Source: Adapted from Sandra Rossi, “And the Oscar Goes to . . . Jovial Penguins and 2,000 Blade Servers,” *Computerworld Australia*, March 6, 2007.

Mainframe Computer Systems

Several years after dire pronouncements that the mainframe was dead, quite the opposite is true: Mainframe usage is actually on the rise. And it’s not just a short-term blip. One factor that’s been driving mainframe sales is cost reductions [of 35 percent or more]. Price reductions aren’t the only factor fueling mainframe acquisitions. IS organizations are teaching the old dog new tricks by putting mainframes at the center stage of emerging applications such as data mining and warehousing, decision support, and a variety of Internet-based applications, most notably electronic commerce.

Mainframe systems are large, fast, and powerful computer systems. For example, mainframes can process thousands of million instructions per second (MIPS). Mainframes can also have large primary storage capacities. Their main memory capacity can range from hundreds of gigabytes to many terabytes of primary storage. Mainframes have slimmed down drastically in the last few years, dramatically reducing their air-conditioning needs, electrical power consumption, and floor space requirements—and thus their acquisition and operating costs. Most of these improvements are the result of a move from cumbersome water-cooled mainframes to a newer air-cooled technology for mainframe systems. See Figure 3.8.

Thus, mainframe computers continue to handle the information processing needs of major corporations and government agencies with high transaction processing volumes or complex computational problems. For example, major international banks, airlines, oil companies, and other large corporations process millions of sales transactions and customer inquiries each day with the help of large mainframe systems. Mainframes are still used for computation-intensive applications, such as analyzing seismic data from oil field explorations or simulating flight conditions in designing aircraft. Mainframes are also widely used as *superservers* for the large client/server networks and high-volume Internet Web sites of large companies. As previously mentioned, mainframes are becoming a popular business computing platform for data mining and warehousing, as well as electronic commerce applications.

FIGURE 3.8

Mainframe computer systems are the heavy lifters of corporate computing.



Source: © Royalty Free/Corbis.

Supercomputer Systems

Supercomputers have now become “scalable servers” at the top end of the product lines that start with desktop workstations. Market-driven companies, like Silicon Graphics, Hewlett-Packard, and IBM, have a much broader focus than just building the world’s fastest computer; and the software of the desktop computer has a much greater overlap with that of the supercomputer than it used to, because both are built from the same cache-based microprocessors.

The term **supercomputer** describes a category of extremely powerful computer systems specifically designed for scientific, engineering, and business applications requiring extremely high speeds for massive numeric computations. The market for supercomputers includes government research agencies, large universities, and major corporations. They use supercomputers for applications such as global weather forecasting, military defense systems, computational cosmology and astronomy, microprocessor research and design, and large-scale data mining.

Supercomputers use *parallel processing* architectures of interconnected microprocessors (which can execute many instructions at the same time in parallel). They can easily perform arithmetic calculations at speeds of billions of floating-point operations per second (*gigaflops*). Supercomputers that can calculate in *teraflops* (trillions of floating-point operations per second), which use massive parallel processing (MPP) designs of thousands of microprocessors, are now in use. Purchase prices for large supercomputers are in the \$5 million to \$50 million range.

The use of symmetric multiprocessing (SMP) and distributed shared memory (DSM) designs of smaller numbers of interconnected microprocessors has spawned a breed of *minisupercomputers* with prices that start in the hundreds of thousands of dollars. For example, IBM’s RS/6000 SP system starts at \$150,000 for a one-processing-node SMP computer. However, it can be expanded to hundreds of processing nodes, which drives its price into the tens of millions of dollars.

The ASCI White supercomputer system, shown in Figure 3.9, consists of three IBM RS/6000 SP systems: White, Frost, and Ice. White, the largest of these systems, is a 512-node, 16-way SMP supercomputer with a peak performance of 12.3 teraflops. Frost is a 68-node, 16-way SMP system; and Ice is a 28-node, 16-way SMP system. Supercomputers like these continue to advance the state of the art for the entire computer industry.

FIGURE 3.9

The ASCI White supercomputer system at Lawrence Livermore National Laboratory in Livermore, California.



Source: Image courtesy of Silicon Graphics, Inc.

Supercomputers Aid Satellite Launches



Satellite launches are a noisy affair, especially for the satellite atop the rocket. Vibration and noise, unless compensated, could render it useless before it reaches orbit, so researchers spend a lot of time on complex computer simulations that help them insulate the delicate craft. Now those simulations are about to get much more accurate, thanks to a new supercomputer that recently began work in Japan.

The Fujitsu FX1 computer was inaugurated in 2009 by the Japan Aerospace Explorations Agency (JAXA). It has 3,008 nodes, each of which has a 4-core Sparc64 VII microprocessor. The machine has 94 terabytes of memory and a theoretical peak performance of 120 teraflops. Running standard benchmarks, it achieved a peak performance of 110.6 teraflops, which ranks it not only the most powerful machine in Japan but also the most efficient supercomputer in the world. Its peak performance represents 91.2 percent of its theoretical performance and outranks the previous record holder, a machine at the Leibniz Rechenzentrum in Munich. Ranked below the German computer is another JAXA machine. “Performance is about 15 times higher than the system we had before,” said Kozo Fujii, director of JAXA’s Engineering Digital Innovation Center.

Two rows of computer racks make up the main system, and a third row alongside is a second, less powerful FX1 machine. In an adjoining room sits an NEC SX-9 vector computer for running specialized tasks and the storage that augments the entire system. All together, a petabyte of disk storage space and 10 petabytes of tape storage are connected to the system (a petabyte is a million gigabytes). And between the lot, there are many big, industrial air conditioners to keep the room cool and extract the heat generated by this mass of hardware.

JAXA intends to put it to work on simulations such as the acoustic noise experienced by a satellite at launch, said Fujii. “There is a wide band of frequencies and usually the peak frequencies are located between 60 and 100 Hertz and we can capture at that level of frequencies. But hopefully with the new computer we can capture frequencies of 150 or 200 Hz that are difficult for the current computer.”

Source: Adapted from Martyn Williams, “World’s Most Efficient Supercomputer Gets to Work,” *CIO Magazine*, April 2, 2009.

The Next Wave of Computing

Interconnecting microprocessors to create minisupercomputers is a reality, as discussed in the previous section. The next wave is looking at harnessing the virtually infinite amount of unused computing power that exists in the myriad of desktops and laptops within the boundaries of a modern organization.

Distributed or grid computing in general is a special type of parallel computing that relies on complete computers (with onboard CPU, storage, power supply, network interface, and so forth) connected to a network (private, public, or the Internet) by a conventional network interface. This is in contrast to the traditional notion of a supercomputer, which has many processors connected together in a single machine. The grid could be formed by harnessing the unused CPU power in all of the desktops and laptops in a single division of a company (or in the entire company, for that matter).

The primary advantage of distributed computing is that each node can be purchased as commodity hardware; when combined, it can produce similar computing resources to a multiprocessor supercomputer, but at a significantly lower cost. This is due to the economies of scale of producing desktops and laptops, compared with the lower efficiency of designing and constructing a small number of custom supercomputers.

One feature of distributed grids is that they can be formed from computing resources belonging to multiple individuals or organizations (known as multiple administrative domains). This can facilitate commercial transactions or make it easier to assemble volunteer computing networks.

A disadvantage of this feature is that the computers that are actually performing the calculations might not be entirely trustworthy. The designers of the system must thus introduce measures to prevent malfunctions or malicious participants from producing false, misleading, or erroneous results, and from using the system as a platform for a hacking attempt. This often involves assigning work randomly to different nodes (presumably with different owners) and checking that at least two different nodes report the same answer for a given work unit. Discrepancies would identify malfunctioning and malicious nodes.

Another challenge is that because of the lack of central control over the hardware, there is no way to guarantee that computers will not drop out of the network at random times. Some nodes (like laptops or dial-up Internet customers) may also be available for computation but not for network communications for unpredictable periods. These variations can be accommodated by assigning large work units (thus reducing the need for continuous network connectivity) and reassigning work units when a given node fails to report its results as expected.

Despite these challenges, grid computing is becoming a popular method of getting the most out of the computing resources of an organization.

Technical Note: The Computer System Concept

As a business professional, you do not need detailed technical knowledge of computers. However, you do need to understand some basic concepts about computer systems, which should help you be an informed and productive user of computer system resources.

A computer is more than a high-powered collection of electronic devices performing a variety of information processing chores. A computer is a *system*, an interrelated combination of components that performs the basic system functions of input, processing, output, storage, and control, thus providing end users with a powerful information processing tool. Understanding the computer as a **computer system** is vital to the effective use and management of computers. You should be able to visualize any computer this way, from the smallest microcomputer device to the largest computer networks whose components are interconnected by telecommunications network links throughout a building complex or geographic area.

FIGURE 3.10 The computer system concept. A computer is a system of hardware components and functions.

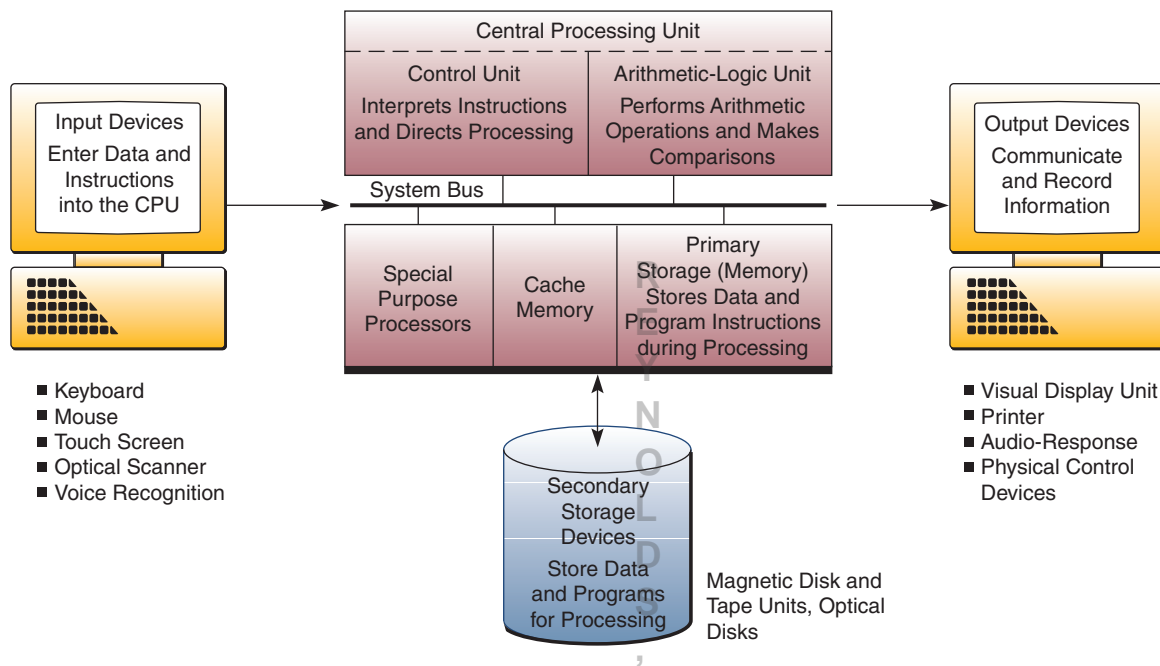



Figure 3.10 illustrates that a computer is a system of hardware devices organized according to the following system functions:

- **Input.** The input devices of a computer system include computer keyboards, touch screens, pens, electronic mice, and optical scanners. They convert data into electronic form for direct entry or through a telecommunications network into a computer system.
- **Processing.** The **central processing unit (CPU)** is the main processing component of a computer system. (In microcomputers, it is the main *microprocessor*. See Figure 3.11.) Conceptually, the circuitry of a CPU can be subdivided into two major subunits: the arithmetic-logic unit and the control unit. The electronic circuits (known as *registers*) of the *arithmetic-logic unit* perform the arithmetic and logic functions required to execute software instructions.
- **Output.** The output devices of a computer system include video display units, printers, and audio response units. They convert electronic information produced by the computer system into human-intelligible form for presentation to end users.
- **Storage.** The storage function of a computer system takes place in the storage circuits of the computer's **primary storage unit**, or *memory*, supported by **secondary storage** devices such as magnetic disk and optical disk drives. These devices store data and software instructions needed for processing. Computer processors may also include storage circuitry called *cache memory* for high-speed, temporary storage of instruction and data elements.
- **Control.** The control unit of a CPU is the control component of a computer system. Its registers and other circuits interpret software instructions and transmit directions that control the activities of the other components of the computer system.

We will explore the various hardware devices associated with each of these system functions in the next section of this chapter.

FIGURE 3.11

Mobile CPU chips, such as the one shown here, can reach speeds up to 3 GHz to bring desktop-like power to a mobile setting.



Source: © Getty Images.

Computer Processing Speeds

How fast are computer systems? Early computer **processing speeds** were measured in **milliseconds** (thousandths of a second) and **microseconds** (millionths of a second). Now computers operate in the **nanosecond** (billionth of a second) range, with **picosecond** (trillionth of a second) speed being attained by some computers. Such speeds seem almost incomprehensible. For example, an average person taking one step each nanosecond would circle the earth about 20 times in one second!

We have already mentioned the *teraflop* speeds of some supercomputers. However, most computers can now process program instructions at **million instructions per second (MIPS)** speeds. Another measure of processing speed is *megahertz* (MHz), or millions of **cycles per second**, and *gigahertz* (GHz), or billions of cycles per second. This rating is commonly called the *clock speed* of a microprocessor because it is used to rate microprocessors by the speed of their timing circuits or internal clock rather than by the number of specific instructions they can process in one second.

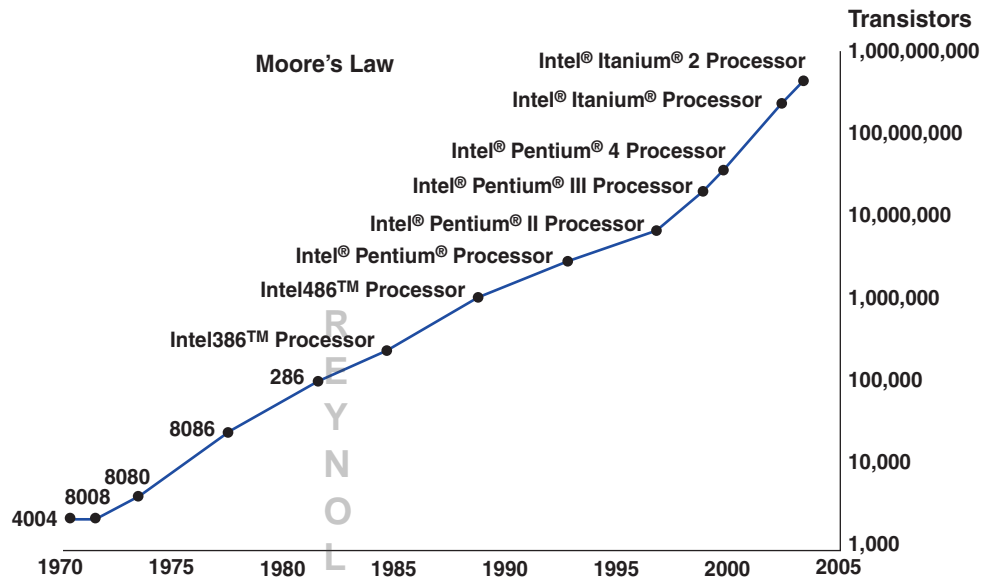
However, such ratings can be misleading indicators of the effective processing speed of microprocessors and their *throughput*, or ability to perform useful computation or data processing assignments during a given period. That's because processing speed depends on a variety of factors, including the size of circuitry paths, or *buses*, that interconnect microprocessor components; the capacity of instruction-processing *registers*; the use of high-speed cache memory; and the use of specialized microprocessors such as a math coprocessor to do arithmetic calculations faster.

Moore's Law: Where Do We Go from Here?

Can computers get any faster? Can we afford the computers of the future? Both of these questions can be answered by understanding **Moore's law**. Gordon Moore, co-founder of Intel Corporation, made his famous observation in 1965, just four years after the first integrated circuit was commercialized. The press called it "Moore's law," and the name has stuck. In its form, Moore observed an exponential growth (doubling every 18 to 24 months) in the number of transistors per integrated circuit and predicted that this trend would continue. Through a number of advances in technology, Moore's law, the doubling of transistors every couple of years, has been maintained and still holds true today. Figure 3.12 illustrates Moore's law as it relates to the evolution of computing power.

FIGURE 3.12

Moore's law suggests that computer power will double every 18 to 24 months. So far, it has.



Despite our regular use of exponential growth when predicting the future, particularly the future of technology, humans are often not very good at realizing what exponential growth really looks like. To understand this issue better, let's take a moment to reflect on what Moore's law would mean to us if it applied beyond the number of transistors on a computer chip:

- According to Moore's law, the estimated number of transistors shipped in 2003 was 10^{18} . That's just about 100 times the estimated number of ants in the world.
- In 1978, a commercial flight between New York and Paris cost about \$900 and took about seven hours. If Moore's law could be applied to commercial aviation, that same flight today would cost about a penny and would take less than one second.

Over the years, Moore's law has been interpreted and reinterpreted such that it is commonly defined in a much broader sense than it was originally offered. Nonetheless, its application, and its relative accuracy, is useful in understanding where we have been and in predicting where we are going. For example, one common corollary of Moore's law is that the price of a given level of computing power will be cut in half about every 18 to 24 months. Moore didn't specifically predict this effect, but it has been shown to be rather consistently accurate as well. This trend is also true for the cost of storage (we will explore this further in the next section).

Although Moore's law was initially made in the form of an observation and prediction, the more widely it became accepted, the more it served as a goal for an entire industry. This caused both marketing and engineering departments of semiconductor manufacturers to focus enormous energy on the specified increase in processing power that it was presumed one or more of their competitors would soon actually attain. Expressed as "a doubling every 18 to 24 months," Moore's law suggests the phenomenal progress of technology in recent years. Expressed on a shorter timescale, however, Moore's law equates to an average performance improvement in the industry as a whole of more than 1 percent *per week*. For a manufacturer competing in the processor, storage, or memory markets, a new product that is expected to take three years to develop and is just two or three months late is 10–15 percent slower or larger than the directly competing products, thus rendering it harder to sell.

A sometimes misunderstood point is that exponentially improved hardware does not necessarily imply that the performance of the software is also exponentially improved. The productivity of software developers most assuredly does not increase

exponentially with the improvement in hardware; by most measures, it has increased only slowly and fitfully over the decades. Software tends to get larger and more complicated over time, and Wirth's law (Niklaus Wirth, a Swiss computer scientist) even states humorously that "Software gets slower faster than hardware gets faster."

Recent computer industry studies predict that Moore's law will continue to hold for the next several chip generations (at least another decade). Depending on the doubling time used in the calculations, this progress could mean up to a 100-fold increase in transistor counts on a chip in the next 10 years. This rapid exponential improvement could put 100 GHz personal computers in every home and 20 GHz devices in every pocket. It seems reasonable to expect that sooner or later computers will meet or exceed any conceivable need for computation. Intel, however, suggests that it can sustain development in line with Moore's law for the next 20 years *without* any significant technological breakthroughs. Given the frequency of such breakthroughs in today's marketplace, it is conceivable that Moore's law can be sustained indefinitely. Regardless of what the end of Moore's law may look like, or when it may arrive, we are still moving along at a phenomenal rate of evolution, and the best may be yet to come.

SECTION II

Computer Peripherals: Input, Output, and Storage Technologies

The right peripherals can make all the difference in your computing experience. A top-quality monitor will be easier on your eyes—and may change the way you work. A scanner can edge you closer to that ever-elusive goal: the paperless office. Backup-storage systems can offer bank-vault security against losing your work. CD and DVD drives have become essential for many applications. Thus, the right choice of peripherals can make a big difference.

Read the Real World Case 2 about the use of voice recognition technology in health care settings. We can learn a lot about the future of the human–computer interface and its business applications from this case. See Figure 3.13.

Peripherals

Peripherals is the generic name given to all input, output, and secondary storage devices that are part of a computer system but are not part of the CPU. Peripherals depend on direct connections or telecommunications links to the central processing unit of a computer system. Thus, all peripherals are **online** devices; that is, they are separate from, but can be electronically connected to and controlled by, a CPU. (This is the opposite of **off-line** devices that are separate from and not under the control of the CPU.) The major types of peripherals and media that can be part of a computer system are discussed in this section. See Figure 3.14.

Input Technologies

Input technologies now provide a more **natural user interface** for computer users. You can enter data and commands directly and easily into a computer system through pointing devices like electronic mice and touch pads and with technologies like optical scanning, handwriting recognition, and voice recognition. These developments have made it unnecessary to record data on paper *source documents* (e.g., sales order forms) and then keyboard the data into a computer in an additional data-entry step. Further improvements in voice recognition and other technologies should enable an even more natural user interface in the future.

Pointing Devices

Keyboards are still the most widely used devices for entering data and text into computer systems. However, **pointing devices** are a better alternative for issuing commands, making choices, and responding to prompts displayed on your video screen. They work with your operating system's **graphical user interface** (GUI), which presents you with icons, menus, windows, buttons, and bars for your selection. For example, pointing devices such as an electronic mouse, trackball, and touch pads allow you to choose easily from menu selections and icon displays using point-and-click or point-and-drag methods. See Figure 3.15.

The **electronic mouse** is the most popular pointing device used to move the cursor on the screen, as well as issue commands and make icon and menu selections. By moving the mouse on a desktop or pad, you can move the cursor onto an icon displayed on the screen. Pressing buttons on the mouse initiates various activities represented by the icon selected.

The trackball, pointing stick, and touch pad are other pointing devices most often used in place of the mouse. A **trackball** is a stationary device related to the mouse. You turn a roller ball with only its top exposed outside its case to move the cursor on the screen. A **pointing stick** (also called a *trackpoint*) is a small button-like device, sometimes likened to the eraser head of a pencil. It is usually centered one row above the space bar of a keyboard. The cursor moves in the direction of the pressure you place on the stick. The **touch pad** is a small rectangular touch-sensitive surface usually placed below the keyboard. The cursor moves in the direction your finger moves on the pad.

REAL WORLD

CASE

2

IT in Health Care: Voice Recognition Tools Make Rounds at Hospitals

The infamous doctor's scrawl may finally be on the way out.

Voice technology is the latest tool health care providers are adopting to cut back on time-consuming manual processes, freeing clinicians to spend more time with patients and reduce costs.

At Butler Memorial Hospital, voice-assisted technology has dramatically reduced the amount of time the Butler, Pa., hospital's team of intravenous (IV) nurses spends recording information in patients' charts and on other administrative tasks.

And at the Cleveland Clinic's Fairview Hospital, doctors are using speech recognition to record notes in patients' e-medical records.

Butler recently completed a pilot project where three IV nurses used Vocollect's AccuNurse hands-free, voice-assisted technology along with Boston Software System's workflow automation tools. The nurses were able to cut the time they spent on phone calls and manual processes, including patient record documentation, by at least 75 percent. Now, Butler is rolling out the voice technology for its full IV team of four nurses and seven other clinicians to use for patient care throughout the facility.

The productivity boost from the voice-assisted tools also helps with the hospital's expansion plans, says Dr. Tom McGill. Butler VP of quality and safety.

Butler will soon add about 70 beds—growing from 235 beds now to more than 300—but it won't need to expand the IV nursing team because of the time savings from the voice-assisted technology, McGill says.

In the past, when a patient needed IV care, such as a change in the intravenous medication being administered, an IV nurse would be paged. The nurse would have to call the

patient's nursing station or the doctor requesting the IV to obtain details. The nurse then would prioritize the request with all the existing IV orders. Once IV care was completed, nurses would record what they did in the patient's e-medical record.

With the AccuNurse, which combines the use of speech recognition and synthesis for charting and communication, Butler's IV nurses wear lightweight headsets and small pocket-sized wireless devices that enable them to hear personalized care instructions and other information about patients' IV needs.

IV requests are entered into Butler's computer system, which sends them through the Vocollect system to the appropriate headset. IV nurses listen to details about new orders and use the system to prioritize IV orders.

When they finish caring for a patient, nurses record what they did in the patient's e-medical record using voice commands. "The nurses can document as they're walking to the next patient's room," says McGill. Once they finish with one patient, nurses say "next task" to obtain instructions for the next patient, McGill says.

The system has shown itself to be capable of understanding different accents, he said. Butler is evaluating expanding use of the voice-assisted technology to other clinical areas, including surgery. The technology could be used to help ensure that surgical staff complete patient safety checklists.

McGill wouldn't say how much Butler paid for the system, but he expects the ROI will be realized in 12 to 18 months. "It's very affordable," he notes. Meanwhile, Dr. Fred Jorgenson, a faculty physician at Cleveland Clinic's Fairview Hospital, is using Nuance's Dragon Medical speech recognition technology to speak patient notes into the hospital's Epic EMR (electronic medical records) system.

"I'm not a fast typist," Jorgenson says. "Many doctors over a certain age aren't. If I had to type all the time, I'd be dead." And, at 13 cents to 17 cents per line, dictation transcription services are expensive.

"In primary care, patient notes can be 30 to 40 lines. That adds up," he says. Fairview is saving about \$2,000 to \$3,000 a month that might have otherwise been spent on transcription, Jorgenson said. It cost about \$3,500 to get Dragon up and running.

With transcription services, the turnaround time is 24 to 36 hours before information is available in the EMR. Spoken notes are available immediately.

Jorgenson describes the accuracy of Dragon Medical's speech-to-text documentation as "very good," especially with medical terms and prescriptions. "It rarely gets medical words wrong," he says. "If you see a mistake, it's usually with 'he' or 'she,' and you can correct it when you see it."

Mount Carmel St. Ann's hospital in Columbus, Ohio, has been among the early wave of health care providers using electronic clinical systems bolstered with speech recognition capabilities. About seven years ago, emergency department doctors at Mount Carmel St. Ann's hospital began having access to Dragon's speech recognition software



FIGURE 3.13

Smart use of voice recognition technologies allows hospitals to improve the quality of care while keeping costs under control.

not long after an e-health record system from Allscripts was rolled out there.

When the e-health record was first rolled out—without the voice capabilities—Mount Carmel St. Ann's doctors didn't necessarily see the kind of productivity boost they had been hoping for, in large part because they found themselves spending a lot of time typing notes, says Dr. Loren Leidheiser, chairman and director of emergency medicine at Mount Carmel St. Ann's emergency department. But as more Mount Carmel St. Ann's ER doctors began incorporating the speech recognition capabilities into their workflow—whether speaking notes into a lapel microphone or into a computer in the patient room or hallway—the efficiency picked up tremendously, says Leidheiser.

Also, before using the Dragon software, the ER department spent about \$500,000 annually in traditional dictation transcription costs for the care associated with the hospital's 60,000 to 70,000 patient visits yearly at the time. That was cut down “to zero,” he says. The return on investment on the speech recognition, combined with the use of the e-health record system, was “within a year and a half,” notes Leidheiser.

Leidheiser also makes use of time stuck in traffic to dictate notes that are later incorporated into patient records or turned into e-mails or letters. Using a Sony digital recorder, Leidheiser can dictate a letter or note while in his car, then later plug the recorder into his desktop computer, where his spoken words are converted to text.

Speech recognition technology is also helping U.S. military doctors keep more detailed patients notes while cutting the time they spend typing on their computers. By 2011, the U.S. Department of Defense expects to have implemented its integrated, interoperable electronic medical record system—AHLTA—at more than 500 military medical facilities and hospitals worldwide.

The system will be used for the care of more than 9 million active military personnel, retirees, and their dependents. Military doctors using the AHLTA system also have access to Dragon NaturallySpeaking Medical speech recognition technology from Nuance Communications' Dictaphone health care division, allowing doctors to speak “notes” into patient records, as an alternative to typing and dictation.

Over the last year, the adoption of Dragon has doubled, with about 6,000 U.S. military doctors using the software at

health care facilities of all military branches, including the Air Force, Army, Navy, and Marine Corps.

The use of Dragon Naturally Speaking voice recognition software with the AHLTA e-health record systems is freeing doctors from several hours of typing their various patient notes each week into the AHLTA, he said. Being able to speak notes into an e-health record at the patient bedside—rather than staring at a computer screen typing—also helps improves doctors' bedside manner and allows them to narrate more comprehensive notes, either while the patients are there or right after a visit. That cuts down on mistakes caused by memory lapses and boosts the level of details that are included in a patient record, says Dr. Robert Bell Walker, European Regional Medical Command AHLTA consultant and a family practice physician for the military.

The voice capability “saves a lot of time and adds to the thoroughness of notes from a medical and legal aspect,” says Dr. Craig Rohan, a U.S. Air Force pediatrician at Peterson Air Force base in Colorado. The ability to speak notes directly into a patient's electronic chart is particularly helpful in complicated cases, where a patient's medical history is complex, he says. Text pops up on the computer screen immediately after words are spoken into the system, so doctors can check the accuracy, make changes, or add other details.

Also, because spoken words are immediately turned into text, the medical record has “a better flow” to document patient visits. Previously, “the notes that had been created by [entering] structured text into the AHLTA system looks more like a ransom note,” says Walker, with information seemingly randomly pasted together.

Doctors can speak into a microphone on their lapels to capture notes in tablet PCs during patient visits, or speak into headsets attached to desktop or wall-mounted computers. The storage requirement of voice notes is “small,” especially when compared with other records, such as medical images, says Walker. By adding spoken notes to medical records, e-mails, and letters, “it's easier to tell the story,” remarks Leidheiser.

Source: Adapted from Marianne Kolbasuk McGee, “Voice Recognition Tools Make Rounds at Hospitals,” *InformationWeek*, September 17, 2009; Marianne Kolbasuk McGee, “Doctors Use Speech Recognition Tools to Enhance Patient E-Health Records,” *InformationWeek*, May 19, 2008; and Matt Hamblen, “Doctors' Notes Get Clearer with Speech Recognition Software,” *Computerworld*, May 16, 2008.

CASE STUDY QUESTIONS

1. What are some of the benefits afforded to organizations implementing voice recognition technologies in these settings? How can you quantify these benefits to assess the value of the investment? Provide several examples from the case.
2. There is no margin for error when working in a health care setting. How would you go about implementing these technologies in this high-risk environment? What precautions or approaches would you take to minimize risks? Develop some recommendations.
3. In what other areas of medicine would you expect technology to make inroads next? Where do you think it would be most beneficial, and how would it change the way doctors and nurses work today? Provide several examples.

REAL WORLD ACTIVITIES

1. The case talks about electronic medical or health records systems. These are slowly becoming standard in many hospitals and clinics, both private and public. Go online and search for reports of these implementations. What are the main benefits derived from their adoption? What have been the major roadblocks preventing their acceptance? Prepare a report to share your findings.
2. The case above was presented from the perspective of practitioners and hospital administrators. How comfortable would you feel, as a patient, knowing that your health care providers are using these technologies? Would you have any concerns? Break into small groups with your classmates to discuss this issue.

FIGURE 3.14

Some advice about peripherals for a business PC.

Peripherals Checklist
<ul style="list-style-type: none">• Monitors. Bigger is better for computer screens. Consider a high-definition 19-inch or 21-inch flat screen CRT monitor, or LCD flat-panel display. That gives you much more room to display spreadsheets, Web pages, lines of text, open windows, and so on. An increasingly popular setup uses two monitors that allow multiple applications to be used simultaneously.
<ul style="list-style-type: none">• Printers. Your choice is between laser printers and color inkjet printers. Lasers are better suited for high-volume business use. Moderately priced color inkjets provide high-quality images and are well suited for reproducing photographs; per-page costs are higher than for laser printers.
<ul style="list-style-type: none">• Scanners. You'll have to decide between a compact, sheet-fed scanner and a flatbed model. Sheet-fed scanners will save desktop space, while bulkier flatbed models provide higher speed and resolution.
<ul style="list-style-type: none">• Hard Disk Drives. Bigger is better; as with closet space, you can always use the extra capacity. So go for 80 gigabytes at the minimum to 160 gigabytes and more.
<ul style="list-style-type: none">• CD and DVD Drives. CD and DVD drives are a necessity for software installation and multimedia applications. Common today is a built-in CD-RW/DVD drive that both reads and writes CDs and plays DVDs.
<ul style="list-style-type: none">• Backup Systems. Essential. Don't compute without them. Removable mag disk drives and even CD-RW and DVD-RW drives are convenient and versatile for backing up your hard drive's contents.

Trackballs, pointing sticks, and touch pads are easier to use than a mouse for portable computer users and are thus built into most notebook computer keyboards.

Touch screens are devices that allow you to use a computer by touching the surface of its video display screen. Some touch screens emit a grid of infrared beams, sound waves, or a slight electric current that is broken when the screen is touched. The computer senses the point in the grid where the break occurs and responds with an appropriate action. For example, you can indicate your selection on a menu display just by touching the screen next to that menu item.

Pen-Based Computing

Handwriting-recognition systems convert script into text quickly and are friendly to shaky hands as well as those of block-printing draftsmen. The pen is more powerful than the keyboard in many vertical markets, as evidenced by the popularity of pen-based devices in the utilities, service, and medical trades.

FIGURE 3.15 Many choices exist for pointing devices including the trackball, mouse, pointing stick, and touch screen.



Source: (left to right) Courtesy of Logitech, Microsoft®, International Business Machines Corporation, and © Don Wright/AP Images.

FIGURE 3.16

Many PDAs accept pen-based input.



Source: © Comstock/PunchStock.

Pen-based computing technologies are still being used in many handheld computers and personal digital assistants. Despite the popularity of touch-screen technologies, many still prefer the use of a stylus rather than their fingertip. *Tablet PCs* and PDAs contain fast processors and software that recognizes and digitizes handwriting, handprinting, and hand drawing. They have a pressure-sensitive layer, similar to that of a touch screen, under their slate-like liquid crystal display (LCD) screen. Instead of writing on a paper form fastened to a clipboard or using a keyboard device, you can use a pen to make selections, send e-mail, and enter handwritten data directly into a computer. See Figure 3.16.

Various pen-like devices are available. One example is the *digitizer pen* and *graphics tablet*. You can use the digitizer pen as a pointing device or to draw or write on the pressure-sensitive surface of the graphics tablet. Your handwriting or drawing is digitized by the computer, accepted as input, displayed on its video screen, and entered into your application.

Speech Recognition Systems

Speech recognition is gaining popularity in the corporate world among nontypists, people with disabilities, and business travelers, and is most frequently used for dictation, screen navigation, and Web browsing.

Speech recognition may be the future of data entry and certainly promises to be the easiest method for word processing, application navigation, and conversational computing because speech is the easiest, most natural means of human communication. Speech input has now become technologically and economically feasible for a variety of applications. Early speech recognition products used *discrete speech recognition*, for which you had to pause between each spoken word. New *continuous speech recognition* software recognizes continuous, conversationally paced speech. See Figure 3.17.

Speech recognition systems digitize, analyze, and classify your speech and its sound patterns. The software compares your speech patterns to a database of sound patterns in its vocabulary and passes recognized words to your application software. Typically, speech recognition systems require training the computer to recognize your voice and its unique sound patterns to achieve a high degree of accuracy.

FIGURE 3.17

Using speech recognition technology for word processing.



Source: © Tim Pannell/Corbis.

Training such systems involves repeating a variety of words and phrases in a training session, as well as using the system extensively.

Continuous speech recognition software products like Dragon Naturally Speaking and ViaVoice by IBM have up to 300,000-word vocabularies. Training to 95 percent accuracy may take several hours. Longer use, faster processors, and more memory make 99 percent accuracy possible. In addition, Microsoft Office Suite 2007 has built-in speech recognition for dictation and voice commands of a variety of software processes.

Speech recognition devices in work situations allow operators to perform data entry without using their hands to key in data or instructions and to provide faster and more accurate input. For example, manufacturers use speech recognition systems for the inspection, inventory, and quality control of a variety of products; airlines and parcel delivery companies use them for voice-directed sorting of baggage and parcels. Speech recognition can also help you operate your computer's operating systems and software packages through voice input of data and commands. For example, such software can be voice-enabled so you can send e-mail and surf the World Wide Web.

Speaker-independent voice recognition systems, which allow a computer to understand a few words from a voice it has never heard before, are being built into products and used in a growing number of applications. Examples include *voice-messaging computers*, which use speech recognition and voice response software to guide an end user verbally through the steps of a task in many kinds of activities. Typically, they enable computers to respond to verbal and Touch-Tone input over the telephone. Examples of applications include computerized telephone call switching, telemarketing surveys, bank pay-by-phone bill-paying services, stock quotation services, university registration systems, and customer credit and account balance inquiries.

One of the newest examples of this technology is Ford SYNC. SYNC is a factory-installed, in-car communications and entertainment system jointly developed by Ford Motor Company and Microsoft. The system was offered on 12 different Ford, Lincoln, and Mercury vehicles in North America for the 2008 model year and is available on most 2009 Ford offerings.

Ford SYNC allows a driver to bring almost any mobile phone or digital media player into a vehicle and operate it using voice commands, the vehicle's steering wheel,

or manual radio controls. The system can even receive text messages and read them aloud using a digitized female voice named “Samantha.” SYNC can interpret a hundred or so shorthand messages, such as LOL for “laughing out loud,” and it will read swear words; it won’t, however, decipher obscene acronyms. Speech recognition is now common in your car, home, and workplace.

Optical Scanning

Few people understand how much scanners can improve a computer system and make your work easier. Their function is to get documents into your computer with a minimum of time and hassle, transforming just about anything on paper—a letter, a logo, or a photograph—into the digital format that your PC can read. Scanners can be a big help in getting loads of paper off your desk and into your PC.

Optical scanning devices read text or graphics and convert them into digital input for your computer. Thus, optical scanning enables the direct entry of data from source documents into a computer system. For example, you can use a compact desktop scanner to scan pages of text and graphics into your computer for desktop publishing and Web publishing applications. You can scan documents of all kinds into your system and organize them into folders as part of a *document management* library system for easy reference or retrieval. See Figure 3.18.

There are many types of optical scanners, but all use photoelectric devices to scan the characters being read. Reflected light patterns of the data are converted into electronic impulses that are then accepted as input to the computer system. Compact desktop scanners have become very popular due to their low cost and ease of use with personal computer systems. However, larger, more expensive *flatbed scanners* are faster and provide higher-resolution color scanning.

Another optical scanning technology is called **optical character recognition** (OCR). The OCR scanners can read the characters and codes on merchandise tags, product labels, credit card receipts, utility bills, insurance premiums, airline tickets,

FIGURE 3.18 A modern document management system can serve as an optical scanner, copier, fax, and printer.



Source: Courtesy of Xerox.

FIGURE 3.19

Using an optical scanning wand to read bar coding of inventory data.



Source: © Jeff Smith/Getty Images.

and other documents. In addition, OCR scanners are used to automatically sort mail, score tests, and process a wide variety of forms in business and government.

Devices such as handheld optical scanning **wands** are frequently used to read *bar codes*, codes that use bars to represent characters. One common example is the Universal Product Code (UPC) bar coding that you see on just about every product sold. For example, the automated checkout scanners found in supermarkets read UPC bar coding. Supermarket scanners emit laser beams that are reflected off a code. The reflected image is converted to electronic impulses that are sent to the in-store computer, where they are matched with pricing information. Pricing information is returned to the terminal, visually displayed, and printed on a receipt for the customer. See Figure 3.19.

Forget the ATM: Deposit Checks Without Leaving Home

First, we didn't need to visit the bank teller anymore. Then we were able to stick our checks right into the ATM without an envelope. Now we won't have to leave the house to make deposits.

Sacramento, California-based Schools Financial Credit Union is one of the latest banks to allow customers to scan checks at home and deposit them over the Internet. Golden One Credit Union, also from California, had introduced scanner-based check deposits in July 2009. "Banking's not the way it was five or 10 years ago," said Nathan Schmidt, a vice president at Schools Financial. "With any type of technology, it becomes more convenient to self-service."

Even with the widespread use of direct deposit and online banking, people still write and receive millions of paper checks each year. And for the most part, when we have to deposit a paper check, we still need to go to an ATM to do it.

Businesses have been making deposits over the Internet far longer, ever since the passage in 2004 of the federal Check 21 Act, which made a digital image of a check legally acceptable for payment. Businesses quickly saw the benefits of the new law. Sending checks as digital images eliminated courier costs and paperwork.

The extension of the service to consumers has come much more slowly. Cary Whaley, a director at Washington, D.C.-based Independent Community Bankers of America, says financial institutions have been wary about potential fraud.

"For many banks, it remains a business application," Whaley says. "The next step is the consumer side, but a lot of community banks are a little wary. When you're getting

into thousands of consumers, the challenge for banks and credit unions is not only monitoring risk, but monitoring for changes in transactions and transaction amounts.”

But some bankers say consumers are increasingly demanding the same convenience given to their business counterparts, and it’s simply a matter of time before remote deposits become much more widespread.

When Schools Financial Credit Union decided to take the plunge, it included safeguards to prevent abuse. Customers must use their existing secure online banking log-in, and they can’t transmit items more than twice a day.

Users have a time limit to scan and deposit the check online, and checks must meet specific requirements before they are deposited. Post-dated, damaged, or lightly printed checks, for instance, will not scan properly and cannot be deposited.

“So many people prefer to do self-service. They choose to go online—maybe they’re parents with small kids, or they might not want to go to an ATM at 3 a.m.,” says Golden One’s chief executive officer, Teresa Halleck.

“People are already online,” she says. “They’re comfortable with electronic delivery and they’re looking for more.”

Source: Adapted from Darrell Smith, “Forget the ATM—Some Banks Allow Check Deposits via Scanner, iPhone,” *The Sacramento Bee*, October 26, 2009.

Other Input Technologies

Magnetic stripe technology is a familiar form of data entry that helps computers read credit cards. The coating of the magnetic stripe on the back of such cards can hold about 200 bytes of information. Customer account numbers can be recorded on the magnetic stripe so that it can be read by bank ATMs, credit card authorization terminals, and many other types of magnetic stripe readers.

Smart cards that embed a microprocessor chip and several kilobytes of memory into debit, credit, and other cards are popular in Europe and becoming available in the United States. One example is in the Netherlands, where millions of smart debit cards have been issued by Dutch banks. Smart debit cards enable you to store a cash balance on the card and electronically transfer some of it to others to pay for small items and services. The balance on the card can be replenished in ATMs or other terminals. The smart debit cards used in the Netherlands feature a microprocessor and either 8 or 16 kilobytes of memory, plus the usual magnetic stripe. The smart cards are widely used to make payments in parking meters, vending machines, newsstands, pay telephones, and retail stores.

Digital cameras represent another fast-growing set of input technologies. Digital still cameras and digital video cameras (digital camcorders) enable you to shoot, store, and download still photos or full-motion video with audio into your PC. Then you can use image-editing software to edit and enhance the digitized images and include them in newsletters, reports, multimedia presentations, and Web pages. Today’s typical mobile phone includes digital camera capabilities as well.

The computer systems of the banking industry can magnetically read checks and deposit slips using **magnetic ink character recognition (MICR)** technology. Computers can thus sort and post checks to the proper checking accounts. Such processing is possible because the identification numbers of the bank and the customer’s account are preprinted on the bottom of the checks with an iron oxide-based ink. The first bank receiving a check after it has been written must encode the amount of the check in magnetic ink on the check’s lower righthand corner. The MICR system uses 14 characters (the 10 decimal digits and 4 special symbols) of a standardized design. *Reader-sorter* equipment reads a check by first magnetizing the magnetic ink characters and then sensing the signal induced by each character as it passes a reading head. In this way, data are electronically captured by the bank’s computer systems.

Output Technologies

Computers provide information in a variety of forms. Video displays and printed documents have been, and still are, the most common forms of output from computer systems. Yet other natural and attractive output technologies, such as **voice response**

systems and multimedia output, are increasingly found along with video displays in business applications.

For example, you have probably experienced the voice and audio output generated by speech and audio microprocessors in a variety of consumer products. Voice messaging software enables PCs and servers in voice mail and messaging systems to interact with you through voice responses. Of course, multimedia output is common on the Web sites of the Internet and corporate intranets.

Video Output

Video displays are the most common type of computer output. Many desktop computers still rely on **video monitors** that use a *cathode ray tube* (CRT) technology similar to the picture tubes used in home television sets. Usually, the clarity of the video display depends on the type of video monitor you use and the graphics circuit board installed in your computer. These can provide a variety of graphics modes of increasing capability. A high-resolution, flicker-free monitor is especially important if you spend a lot of time viewing multimedia on CDs, or on the Web, or the complex graphical displays of many software packages.

The biggest use of **liquid crystal displays** (LCDs) has been to provide a visual display capability for portable microcomputers and PDAs. However, the use of “flat panel” LCD video monitors for desktop PC systems has become common as their cost becomes more affordable. See Figure 3.20. These LCD displays need significantly less electric current and provide a thin, flat display. Advances in technology such as *active matrix* and *dual scan* capabilities have improved the color and clarity of LCD displays. In addition, high-clarity flat panel televisions and monitors using *plasma* display technologies are becoming popular for large-screen (42- to 80-inch) viewing.

Printed Output

Printing information on paper is still the most common form of output after video displays. Thus, most personal computer systems rely on an inkjet or laser printer to produce permanent (hard-copy) output in high-quality printed form. Printed output is still a common form of business communications and is frequently required for legal documentation. Computers can produce printed reports and correspondence, documents such as sales invoices, payroll checks, bank statements, and printed versions of graphic displays. See Figure 3.21.

FIGURE 3.20

The flat-panel LCD video monitor is becoming the de facto standard for a desktop PC system.



Source: Courtesy of Hewlett-Packard.

FIGURE 3.21

Modern laser printers produce high-quality color output with high speed.



Source: Courtesy of Xerox.

Inkjet printers, which spray ink onto a page, have become the most popular, low-cost printers for microcomputer systems. They are quiet, produce several pages per minute of high-quality output, and can print both black-and-white and high-quality color graphics. **Laser printers** use an electrostatic process similar to a photocopying machine to produce many pages per minute of high-quality black-and-white output. More expensive color laser printers and multifunction inkjet and laser models that print, fax, scan, and copy are other popular choices for business offices.


Storage Trade-Offs

Data and information must be stored until needed using a variety of storage methods. For example, many people and organizations still rely on paper documents stored in filing cabinets as a major form of storage media. However, you and other computer users are more likely to depend on the memory circuits and secondary storage devices of computer systems to meet your storage requirements. Progress in very-large-scale integration (VLSI), which packs millions of memory circuit elements on tiny semiconductor memory chips, is responsible for continuing increases in the main-memory capacity of computers. Secondary storage capacities are also escalating into the billions and trillions of characters, due to advances in magnetic and optical media.

There are many types of storage media and devices. Figure 3.22 illustrates the speed, capacity, and cost relationships of several alternative primary and secondary storage media. Note the cost/speed/capacity trade-offs as you move from semiconductor memories to magnetic disks to optical disks and to magnetic tape. High-speed storage media cost

FIGURE 3.22

Storage media cost, speed, and capacity trade-offs. Note how cost increases with faster access speeds but decreases with the increased capacity of storage media.



Computer Storage Fundamentals

FIGURE 3.23 Examples of the ASCII computer code that computers use to represent numbers and the letters of the alphabet.

more per byte and provide lower capacities. Large-capacity storage media cost less per byte but are slower. These trade-offs are why we have different kinds of storage media.

However, all storage media, especially memory chips and magnetic disks, continue to increase in speed and capacity and decrease in cost. Developments like automated high-speed cartridge assemblies have given faster access times to magnetic tape, and the speed of optical disk drives continues to increase.

Note in Figure 3.22 that semiconductor memories are used mainly for primary storage, although they are sometimes used as high-speed secondary storage devices. Magnetic disk and tape and optical disk devices, in contrast, are used as secondary storage devices to enlarge the storage capacity of computer systems. Also, because most primary storage circuits use RAM (random-access memory) chips, which lose their contents when electrical power is interrupted, secondary storage devices provide a more permanent type of storage media.

Data are processed and stored in a computer system through the presence or absence of electronic or magnetic signals in the computer's circuitry or in the media it uses. This character is called "two-state" or **binary representation** of data because the computer and the media can exhibit only two possible states or conditions, similar to a common light switch: "on" or "off." For example, transistors and other semiconductor circuits are in either a conducting or a nonconducting state. Media such as magnetic disks and tapes indicate these two states by having magnetized spots whose magnetic fields have one of two different directions, or polarities. This binary characteristic of computer circuitry and media is what makes the binary number system the basis for representing data in computers. Thus, for electronic circuits, the conducting ("on") state represents the number 1, whereas the nonconducting ("off") state represents the number 0. For magnetic media, the magnetic field of a magnetized spot in one direction represents a 1, while magnetism in the other direction represents a 0.

The smallest element of data is called a **bit**, short for *binary digit*, which can have a value of either 0 or 1. The capacity of memory chips is usually expressed in terms of bits. A **byte** is a basic grouping of bits that the computer operates as a single unit. Typically, it consists of eight bits and represents one character of data in most computer coding schemes. Thus, the capacity of a computer's memory and secondary storage devices is usually expressed in terms of bytes. Computer codes such as ASCII (American Standard Code for Information Interchange) use various arrangements of bits to form bytes that represent the numbers 0 through 9, the letters of the alphabet, and many other characters. See Figure 3.23.

Character	ASCII Code	Character	ASCII Code	Character	ASCII Code
0	00110000	A	01000001	N	01001110
1	00110001	B	01000010	O	01001111
2	00110010	C	01000011	P	01010000
3	00110011	D	01000100	Q	01010001
4	00110100	E	01000101	R	01010010
5	00110101	F	01000110	S	01010011
6	00110110	G	01000111	T	01010100
7	00110111	H	01001000	U	01010101
8	00111000	I	01001001	V	01010110
9	00111001	J	01001010	W	01010111
		K	01001011	X	01011000
		L	01001100	Y	01011001
		M	01001101	Z	01011010

FIGURE 3.24
Computers use the binary system to store and compute numbers.

2^7	2^6	2^5	2^4	2^3	2^2	2^1	2^0
128	64	32	16	8	4	2	1
0 or 1	0 or 1	0 or 1	0 or 1	0 or 1	0 or 1	0 or 1	0 or 1

To represent any decimal number using the binary system, each place is simply assigned a value of either 0 or 1. To convert binary to decimal, simply add up the value of each place.

Example:

2^7	2^6	2^5	2^4	2^3	2^2	2^1	2^0
1	0	0	1	1	0	0	1
128	0	0	16	8	0	0	1

128 + 0 + 0 + 16 + 8 + 0 + 0 + 1 = 153

10011001 = 153

Since childhood, we have learned to do our computations using the numbers 0 through 9, the digits of the decimal number system. Although it is fine for us to use 10 digits for our computations, computers do not have this luxury. Every computer processor is made of millions of tiny switches that can be turned off or on. Because these switches have only two states, it makes sense for a computer to perform its computations with a number system that has only two digits: the **binary number system**. These digits (0 and 1) correspond to the off/on positions of the switches in the computer processor. With only these two digits, a computer can perform all the arithmetic that we can with 10 digits. Figure 3.24 illustrates the basic concepts of the binary system.

The binary system is built on an understanding of exponentiation (raising a number to a power). In contrast to the more familiar decimal system, in which each place represents the number 10 raised to a power (ones, tens, hundreds, thousands, and so on), each place in the binary system represents the number 2 raised to successive powers (2^0 , 2^1 , 2^2 , and so on). As shown in Figure 3.24, the binary system can be used to express any integer number by using only 0 and 1.

Storage capacities are frequently measured in **kilobytes** (KB), **megabytes** (MB), **gigabytes** (GB), or **terabytes** (TB). Although *kilo* means 1,000 in the metric system, the computer industry uses *K* to represent 1,024 (or 2^{10}) storage positions. For example, a capacity of 10 megabytes is really 10,485,760 storage positions, rather than 10 million positions. However, such differences are frequently disregarded to simplify descriptions of storage capacity. Thus, a megabyte is roughly 1 million bytes of storage, a gigabyte is roughly 1 billion bytes, and a terabyte represents about 1 trillion bytes, while a **petabyte** is more than 1 quadrillion bytes.


To put these storage capacities in perspective, consider the following: A terabyte is equivalent to about 20 million typed pages, and it has been estimated that the total size of all the books, photographs, video and sound recordings, and maps in the U.S. Library of Congress approximates 3 petabytes (3,000 terabytes).

Direct and Sequential Access

Primary storage media such as semiconductor memory chips are called **direct access** memory or random-access memory (RAM). Magnetic disk devices are frequently called direct access storage devices (DASDs). In contrast, media such as magnetic tape cartridges are known as **sequential access** devices.

The terms *direct access* and *random access* describe the same concept. They mean that an element of data or instructions (such as a byte or word) can be directly stored and retrieved by selecting and using any of the locations on the storage media. They also mean that each storage position (1) has a unique address and (2) can be individually accessed in about the same length of time without having to search through other

FIGURE 3.25 Sequential versus direct access storage. Magnetic tape is a typical sequential access medium. Magnetic disks are typical direct access storage devices.



storage positions. For example, each memory cell on a microelectronic semiconductor RAM chip can be individually sensed or changed in the same length of time. Also, any data record stored on a magnetic or optical disk can be accessed directly in about the same period. See Figure 3.25.

Sequential access storage media such as magnetic tape do not have unique storage addresses that can be directly addressed. Instead, data must be stored and retrieved using a sequential or serial process. Data are recorded one after another in a predetermined sequence (e.g., numeric order) on a storage medium. Locating an individual item of data requires searching the recorded data on the tape until the desired item is located.

Semiconductor Memory

Memory is the coalman to the CPU's locomotive: For maximum PC performance, it must keep the processor constantly stoked with instructions. Faster CPUs call for larger and faster memories, both in the cache where data and instructions are stored temporarily and in the main memory.

The primary storage (main memory) of your computer consists of microelectronic **semiconductor memory** chips. It provides you with the working storage your computer needs to process your applications. Plug-in memory circuit boards containing 256 megabytes or more of memory chips can be added to your PC to increase its memory capacity. Specialized memory can help improve your computer's performance. Examples include external cache memory of 512 kilobytes to help your microprocessor work faster or a video graphics accelerator card with 64 megabytes or more of RAM for faster and clearer video performance. Removable credit-card-size and smaller "flash memory" RAM devices like a jump drive or a memory stick can also provide hundreds of megabytes of erasable direct access storage for PCs, PDAs, or digital cameras.

Some of the major attractions of semiconductor memory are its small size, great speed, and shock and temperature resistance. One major disadvantage of most semiconductor memory is its **volatility**. Uninterrupted electric power must be supplied, or the contents of memory will be lost. Therefore, either emergency transfer to other devices or standby electrical power (through battery packs or emergency generators) is required if data are to be saved. Another alternative is to permanently "burn in" the contents of semiconductor devices so that they cannot be erased by a loss of power.

Thus, there are two basic types of semiconductor memory: **random-access memory (RAM)** and **read-only memory (ROM)**.

- **RAM, random-access memory.** These memory chips are the most widely used primary storage medium. Each memory position can be both sensed (read) and changed (written), so it is also called read/write memory. This is a volatile memory.

FIGURE 3.26
A USB flash memory drive.



Source: Courtesy of Lexar Media.

- **ROM, read-only memory.** Nonvolatile random-access memory chips are used for permanent storage; ROM can be read but not erased or overwritten. Frequently used control instructions in the control unit and programs in primary storage (such as parts of the operating system) can be permanently burned into the storage cells during manufacture, sometimes called *firmware*. Variations include PROM (programmable read-only memory) and EPROM (erasable programmable read-only memory), which can be permanently or temporarily programmed after manufacture.

One of the newest and most innovative forms of storage that uses semiconductor memory is the *flash drive* (sometimes referred to as a *JumpDrive*). Figure 3.26 shows a common flash memory drive.

Flash memory uses a small chip containing thousands of transistors that can be programmed to store data for virtually unlimited periods without power. The small drives can be easily transported in your pocket and are highly durable. Storage capacities currently range as high as 20 gigabytes, but newer flash technologies are making even higher storage capacities a reality. The advent of credit-card-like memory cards and ever-smaller storage technologies puts more data into the user's pocket every day.

Nanochip Inc.: New Memory Process May Overcome Traditional Barriers

A new kind of flash memory technology with potentially greater capacity and durability, lower power requirements, and the same design as flash memory is primed to challenge today's solid-state disk products. Nanochip Inc., based in Fremont, California, said it has made breakthroughs in its array-based memory research that will enable it to deliver working prototypes to potential manufacturing partners by 2009.

Current thinking is that flash memory could hit its limit at around 32–45 nanometers. That describes the smallest possible width of a metal line on the circuit or the amount of space between that line and the next line. The capacity of an integrated circuit is restricted by the ability to print to a smaller and smaller two-dimensional plane, otherwise known as the lithography. That's exactly where Nanochip's technology shines.

"Every two years, you need to buy this new machine that allows you to print something that's smaller and finer," says Stefan Lai of Nanochip. Array-based memory uses a grid of microscopic probes to read and write to a storage material. The storage area isn't defined by the lithography but by the movement of the probes. "If Nanochip can move the probes one-tenth the distance, for example, they can get 100 times the density with no change in the lithography," says Lai. "You don't have to buy all these new machines." IBM has been working on a similar technology for years.

Lai believes that the new memory could herald breakthroughs in mobile devices and biotechnology. "You now need your whole life history stored in your mobile device," he says. "If you want something to store your genome in, it may take a lot of memory, and you'll want to carry it with you." The big question that remains for Nanochip is whether the company can create working prototypes with the cost advantages that array-based technology is supposed to offer over conventional forms of memory.

The challenge for adoption of any new type of memory is that flash itself isn't standing still. "In 2010, it's going to be \$1 per gigabyte . . . so hopefully the cost per gigabyte [of probe-based arrays] is going to be low."

Source: Adapted from Dian Schaffhauser, "A Storage Technology That Breaks Moore's Law," *Computerworld*, March 19, 2008.

Magnetic Disks

Multigigabyte magnetic disk drives aren't extravagant, considering that full-motion video files, sound tracks, and photo-quality images can consume colossal amounts of disk space in a blink.

Magnetic disks are the most common form of secondary storage for your computer system. That's because they provide fast access and high storage capacities at a reasonable cost. Magnetic disk drives contain metal disks that are coated on both sides with an iron oxide recording material. Several disks are mounted together on a vertical shaft, which typically rotates the disks at speeds of 3,600 to 7,600 revolutions per minute (rpm). Electromagnetic read/write heads are positioned by access arms between the slightly separated disks to read and write data on concentric, circular tracks. Data are recorded on tracks in the form of tiny magnetized spots to form the binary digits of common computer codes. Thousands of bytes can be recorded on each track, and there are several hundred data tracks on each disk surface, thus providing you with billions of storage positions for your software and data. See Figure 3.27.

Types of Magnetic Disks

There are several types of magnetic disk arrangements, including removable disk cartridges as well as fixed disk units. Removable disk devices are popular because they are transportable and can be used to store backup copies of your data off-line for convenience and security.

- **Floppy disks**, or magnetic diskettes, consist of polyester film disks covered with an iron oxide compound. A single disk is mounted and rotates freely inside a protective flexible or hard plastic jacket, which has access openings to accommodate the read/write head of a disk drive unit. The 3½-inch floppy disk, with capacities of 1.44 megabytes, was the most widely used version, with a Super-disk technology offering 120 megabytes of storage. Zip drives use a floppy-like technology to provide up to 750 MB of portable disk storage. Today's computers have all but eliminated inclusion of a drive to read floppy disks, but they can be found if necessary.

FIGURE 3.27 Magnetic disk media: a hard magnetic disk drive and a 3½-inch floppy disk.



Source: © Stockbyte/PunchStock.



Source: © Royalty Free/Corbis.

- **Hard disk drives** combine magnetic disks, access arms, and read/write heads into a sealed module. This combination allows higher speeds, greater data recording densities, and closer tolerances within a sealed, more stable environment. Fixed or removable disk cartridge versions are available. Capacities of hard drives range from several hundred megabytes to hundreds of gigabytes of storage.

RAID Storage

RAID computer storage equipment—big, refrigerator-size boxes full of dozens of inter-linked magnetic disk drives that can store the equivalent of 100 million tax returns—hardly gets the blood rushing. But it should. Just as speedy and reliable networking opened the floodgates to cyberspace and e-commerce, ever-more-turbocharged data storage is a key building block of the Internet.

Disk arrays of interconnected microcomputer hard disk drives have replaced large-capacity mainframe disk drives to provide virtually unlimited online storage. Known as **RAID (redundant arrays of independent disks)**, they combine from 6 to more than 100 small hard disk drives and their control microprocessors into a single unit. These RAID units provide large capacities (as high as 1–2 terabytes or more) with high access speeds because data are accessed in parallel over multiple paths from many disks. Also, RAID units provide a *fault-tolerant* capacity, in that their redundant design offers multiple copies of data on several disks. If one disk fails, data can be recovered from backup copies automatically stored on other disks. Storage area networks (SANs) are high-speed *fiber channel* local area networks that can interconnect many RAID units and thus share their combined capacity through network servers with many users.

There are a variety of classifications of RAID, and newer implementations include not only hardware versions, but also software methods. The technical aspects of RAID are beyond the scope of this text and probably beyond the needs of the modern business technologist as well. It is sufficient to note that the storage mechanisms in the modern organization are probably using some type of RAID technology. If you are interested in drilling deeper into this technology and how it works, a wide variety of Internet resources are available.

Magnetic Tape

Tape storage is moving beyond backup. Although disk subsystems provide the fastest response time for mission-critical data, the sheer amount of data that users need to access these days as part of huge enterprise applications, such as data warehouses, requires affordable [magnetic tape] storage.

Magnetic tape is still being used as a secondary storage medium in business applications. The read/write heads of magnetic tape drives record data in the form of magnetized spots on the iron oxide coating of the plastic tape. Magnetic tape devices include tape reels and cartridges in mainframes and midrange systems and small cassettes or cartridges for PCs. Magnetic tape cartridges have replaced tape reels in many applications and can hold more than 200 megabytes.

One growing business application of magnetic tape involves the use of high-speed 36-track magnetic tape cartridges in robotic automated drive assemblies that can directly access hundreds of cartridges. These devices provide lower-cost storage to supplement magnetic disks to meet massive data warehouse and other online business storage requirements. Other major applications for magnetic tape include long-term *archival* storage and backup storage for PCs and other systems.

Optical Disks

Optical disk technology has become a necessity. Most software companies now distribute their elephantine programs on CD-ROMs. Many corporations are now rolling their own CDs to distribute product and corporate information that once filled bookshelves.

FIGURE 3.28
Comparing the capabilities
of optical disk drives.

Optical Disk Drive Capabilities	
• CD-ROM	A CD-ROM drive provides a low-cost way to read data files and load software onto your computer, as well as play music CDs.
• CD-RW	A CD-RW drive allows you to easily create your own custom data CDs for data backup or data transfer purposes. It will also allow you to store and share video files, large data files, digital photos, and other large files with other people that have access to a CD-ROM drive. This drive will also do anything your CD-ROM drive will do. It reads all your existing CD-ROMs, audio CDs, and CDs that you have created with your CD burner.
• CD-RW/DVD	A CD-RW/DVD combination drive brings all the advantages of CD-RW, CD-ROM, and DVD-ROM to a single drive. With a CD-RW/DVD combo drive, you can read DVD-ROM disks, read CD-ROM disks, and create your own custom CDs.
• DVD-ROM	A DVD-ROM drive allows you to enjoy the crystal-clear color, picture, and sound clarity of DVD video on your PC. It will also prepare you for future software and large data files that will be released on DVD-ROM. A DVD-ROM drive can also read CD-ROM disks, effectively providing users with full optical read capability in one device.
• DVD+RW/+R with CD-RW	A DVD+RW/+R with CD-RW drive is a great all-in-one drive, allowing you to burn DVD+RW or DVD+R disks, burn CDs, and read DVDs and CDs. It enables you to create DVDs to back up and archive up to 4.7GB of data files (that's up to 7 times the capacity of a standard 650MB CD) and store up to to 2 hours of MPEG2 digital video.

Source: Adapted from "Learn More—Optical Drives," www.dell.com.

Optical disks, a fast-growing type of storage media, use several major alternative technologies. See Figure 3.28. One version is called **CD-ROM** (compact disk–read-only memory). CD-ROM technology uses 12-centimeter (4.7-inch) compact disks (CDs) similar to those used in stereo music systems. Each disk can store more than 600 megabytes. That's the equivalent of more than 400 1.44-megabyte floppy disks or more than 300,000 double-spaced pages of text. A laser records data by burning permanent microscopic pits in a spiral track on a master disk from which compact disks can be mass produced. Then CD-ROM disk drives use a laser device to read the binary codes formed by those pits.

CD-R (compact disk–recordable) is another popular optical disk technology. CD-R drives or *CD burners* are commonly used to record data permanently on CDs. The major limitation of CD-ROM and CD-R disks is that recorded data cannot be erased. However, **CD-RW** (CD-rewritable) drives record and erase data by using a laser to heat a microscopic point on the disk's surface. In CD-RW versions using magneto-optical technology, a magnetic coil changes the spot's reflective properties from one direction to another, thus recording a binary 1 or 0. A laser device can then read the binary codes on the disk by sensing the direction of reflected light.

DVD technologies have dramatically increased optical disk capacities and capabilities. DVD (digital video disk or digital versatile disk) optical disks can hold from 3.0 to 8.5 gigabytes of multimedia data on each side. The large capacities and high-quality images and sound of DVD technology are expected to replace CD technologies for data storage and promise to accelerate the use of DVD drives for multimedia products that can be used in both computers and home entertainment systems. Thus, **DVD-ROM** disks are increasingly replacing magnetic tape videocassettes for movies and other multimedia products, while **DVD+RW** disks are being used for backup and archival storage of large data and multimedia files. See Figure 3.29.

FIGURE 3.29

Optical disk storage includes CD and DVD technologies.



Source: Photodisc/Getty Images.

Business Applications

One of the major uses of optical disks in mainframe and midrange systems is in **image processing**, where long-term archival storage of historical files of document images must be maintained. Financial institutions, among others, are using optical scanners to capture digitized document images and store them on optical disks as an alternative to microfilm media.

One of the major business uses of CD-ROM disks for personal computers is to provide a publishing medium for fast access to reference materials in a convenient, compact form. This material includes catalogs, directories, manuals, periodical abstracts, part listings, and statistical databases of business and economic activity. Interactive multimedia applications in business, education, and entertainment are another major use of optical disks. The large storage capacities of CD and DVD disks are a natural choice for computer video games, educational videos, multimedia encyclopedias, and advertising presentations.

Radio Frequency Identification

One of the newest and most rapidly growing storage technologies is **radio frequency identification (RFID)**, a system for tagging and identifying mobile objects such as store merchandise, postal packages, and sometimes even living organisms (like pets). Using a special device called an **RFID reader**, RFID allows objects to be labeled and tracked as they move from place to place.

The RFID technology works using small (sometimes smaller than a grain of sand) pieces of hardware called **RFID chips**. These chips feature an antenna to transmit and receive radio signals. Currently, there are two general types of RFID chips: *passive* and *active*. **Passive RFID** chips do not have a power source and must derive their power from the signal sent from the reader. **Active RFID** chips are self-powered and do not need to be close to the reader to transmit their signal. Any RFID chips may be attached to objects or, in the case of some passive RFID systems, injected into objects. A recent use for RFID chips is the identification of pets such as dogs or cats. By having a tiny RFID chip injected just under their skin, they can be easily identified if they become lost. The RFID chip contains contact information about the owner of the pet. Taking this a step further, the Transportation Security Administration is considering using RFID tags embedded in airline boarding passes to keep track of passengers.

Whenever a reader within range sends appropriate signals to an object, the associated RFID chip responds with the requested information, such as an identification number or

product date. The reader, in turn, displays the response data to an operator. Readers may also forward data to a networked central computer system. Such RFID systems generally support storing information on the chips as well as simply reading data.

The RFID systems were created as an alternative to common bar codes. Relative to bar codes, RFID allows objects to be scanned from a greater distance, supports storing of data, and allows more information to be tracked per object.

Recently (as discussed in the next section), RFID has raised some privacy concerns as a result of the invisible nature of the system and its capability to transmit fairly sophisticated messages. As these types of issues are resolved, we can expect to see RFID technology used in just about every way imaginable.

RFID Privacy Issues

How would you like it if, for instance, one day you realized your underwear was reporting on your whereabouts?—California State Senator Debra Bowen, at a 2003 hearing on RFID privacy concerns.

The use of RFID technology has caused considerable controversy and even product boycotts by consumer privacy advocates who refer to RFID tags as *spy chips*. The two main privacy concerns regarding RFID are:

- Because the owner of an item will not necessarily be aware of the presence of an RFID tag, and the tag can be read at a distance without the knowledge of the individual, it becomes possible to gather sensitive data about an individual without consent.
- If a customer pays for a tagged item by credit card or in conjunction with a loyalty card, then it would be possible to deduce the identity of the purchaser indirectly by reading the globally unique ID of that item (contained in the RFID tag).

Most concerns revolve around the fact that RFID tags affixed to products remain functional even after the products have been purchased and taken home; thus, they can be used for surveillance and other purposes unrelated to their supply chain inventory functions.

Read range, however, is a function of both the reader and the tag itself. Improvements in technology may increase read ranges for tags. Having readers very close to the tags makes short-range tags readable. Generally, the read range of a tag is limited to the distance from the reader over which the tag can draw enough energy from the reader field to power the tag. Tags may be read at longer ranges by increasing reader power. The limit on read distance then becomes the signal-to-noise ratio of the signal reflected from the tag back to the reader. Researchers at two security conferences have demonstrated that passive UHF RFID tags (not the HF-type used in U.S. passports), normally read at ranges of up to 30 feet, can be read at ranges of 50–69 feet using suitable equipment. Many other types of tag signals can be intercepted from 30–35 feet away under good conditions, and the reader signal can be detected from miles away if there are no obstructions.

The potential for privacy violations with RFID was demonstrated by its use in a pilot program by the Gillette Company, which conducted a “smart shelf” test at a Tesco in Cambridge, England. They automatically photographed shoppers taking RFID-tagged safety razors off the shelf to see if the technology could be used to deter shoplifting. This trial resulted in consumer boycott against Gillette and Tesco. In another incident, uncovered by the *Chicago Sun-Times*, shelves in a Wal-Mart in Broken Arrow, Oklahoma, were equipped with readers to track the Max Factor Lipfinity lipstick containers stacked on them. Webcam images of the shelves were viewed 750 miles away by Procter & Gamble researchers in Cincinnati, Ohio, who could tell when lipsticks were removed from the shelves and observe the shoppers in action.

The controversy surrounding the use of RFID technologies was furthered by the accidental exposure of a proposed Auto-ID consortium public relations campaign that was designed to “neutralize opposition” and get consumers to “resign themselves to the inevitability of it” while merely pretending to address their concerns. During the U.N. World Summit on the Information Society (WSIS) on November 16–18, 2005,

Richard Stallman, founder of the free software movement, protested the use of RFID security cards. During the first meeting, it was agreed that future meetings would no longer use RFID cards; upon finding out this assurance was broken, he covered his card in tin foil and would uncover it only at the security stations. This protest caused the security personnel considerable concern. Some did not allow him to leave a conference room in which he had been the main speaker, and then prevented him from entering another conference room, where he was due to speak.

The Food and Drug Administration in the United States has approved the use of RFID chips in humans. Some business establishments have also started to “chip” customers, such as the Baja Beach Nightclub in Barcelona. This has provoked concerns into privacy of individuals, as they can potentially be tracked wherever they go by an identifier unique to them. There are concerns that this could lead to abuse by an authoritarian government or lead to removal of other freedoms.

In July 2006, Reuters reported that Newitz and Westhues, two hackers, showed at a conference in New York City that they could clone the RFID signal from a human-implanted RFID chip, which proved that the chip is not as secure as was previously believed.

All of these examples share a common thread, showing that whatever can be encoded can also be decoded. RFID presents the potential for enormous efficiencies and cost savings. It also presents significant challenges to privacy and security. Until these issues are worked through, much controversy will continue to surround RFID technologies.

RFID-Enabled Magazines: Tracking Reading Patterns

One of the most vexing problems for magazine publishers is trying to figure out just how many people read printed copies of magazines, rather than letting them languish in stacks of unread mail. Other questions have been raging since the dawn of the printing press, such as: How long and often do readers spend reading the pages? Do readers skip around among the articles? Do they read from front to back or from back to front? And does anybody look at the advertisements? Historically, these have been mostly unanswerable questions, left to estimates and guesswork. But a marketing research company, Mediamark Research & Intelligence (MRI), is testing radio frequency identification (RFID) technology to measure magazine readership in public waiting rooms.

The real-world testing follows up a year of laboratory testing. Jay Mattlin, senior vice president of new ventures at MRI, points out that the system needs to be tested “in a non-laboratory setting to determine how well it holds up in this important reading environment.”

The project’s objectives are to determine whether the RFID-driven passive print monitoring system “can reliably measure—in a waiting room setting—the total time spent with a specific magazine issue, the number of individual reading occasions and potentially, reader exposure to individual magazine pages,” according to an MRI statement.

For the lab testing, MRI created an “intelligent” magazine prototype—containing the passive print measuring system—that keeps track of reader activity with designated pages. “Essentially, an RFID tag attached to the magazine sends a signal to a tag reader each time the test subjects turn to one of the designated magazine pages,” notes MRI. “The system records the times of the openings and closings of designated pages, as well as the opening/closings of the magazine itself.”

Mattlin reported that the system correctly identified magazine openings and closings an average of 95 percent of the time in internal tests.

“We’ve learned a lot so far in our controlled environment,” he noted. “But considering the complexity of trying to measure a non-electronic medium, like magazines, with electronic signals, it’s going to take a while before we have a firm grip on the full potential of RFID with regard to magazine audience measurement.”

Of course, the most interesting thing to note about this story is the timing: How much value is there in solving the age-old viewership problem as print magazine readership continues to decline, and publishers have shifted most of their focus and content online?

Source: Adapted from Thomas Wailgum, “RFID Chips in Your Magazines,” *CIO Magazine*, December 12, 2007.

Predictions for the Future

If Moore’s law prevails and technology advancement continues, we can expect to see our lives change in remarkable and unimaginable ways. Although we cannot really predict the future, it is interesting and fun to read the predictions of futurists—people whose job is to think about what the future might bring. Here’s one man’s perspective on what computing technology might do to change our lives in the decades to come.

Computers Will Enable People to Live Forever

In just 15 years, we’ll begin to see the merger of human and computer intelligence that ultimately will enable people to live forever. At least that’s the prediction of author and futurist Ray Kurzweil.

Kurzweil suggests that nanobots will roam our bloodstreams, fixing diseased or aging organs, while computers will back up our human memories and rejuvenate our bodies by keeping us young in appearance and health.

The author of the book *The Singularity Is Near*, Kurzweil says that within a quarter of a century, nonbiological intelligence will match the range and subtlety of human intelligence. He predicts that it will then soar past human ability because of the continuing acceleration of information-based technologies, as well as the ability of machines to share their knowledge instantly.

Kurzweil predicts people and computers will intermix with nanobots, blood cell-sized robots, that will be integrated into everything from our clothing to our bodies and brains. People simply need to live long enough—another 15–30 years—to live forever. Think of it as replacing everyone’s “human body version 1.0” with *nanotechnology* that will repair or replace ailing or aging tissue, he says. Parts will become easily replaceable.

“A \$1,000 worth of computation in the 2020s will be 1,000 times more powerful than the human brain,” says Kurzweil, adding that in 25 years we’ll have multiplied our computational power by a billion. “Fifteen years from now, it’ll be a very different world. We’ll have cured cancer and heart disease, or at least rendered them to manageable chronic conditions that aren’t life threatening. We’ll get to the point where we can stop the aging process and stave off death.”

Actually, we’ll hit a point where human intelligence simply can’t keep up with, or even follow, the progress that computers will make, according to Kurzweil. He expects that nonbiological intelligence will have access to its own design plans and be able to improve itself rapidly. Computer, or nonbiological, intelligence created in the year 2045 will be one billion times more powerful than all human intelligence today.

“Supercomputing is behind the progress in all of these areas,” says Kurzweil, adding that a prerequisite for nonbiological intelligence is to reverse-engineer biology and the human brain. That will give scientists a “toolkit of techniques” to apply when developing intelligent computers. In a written report, he said, “We won’t experience 100 years of technological advance in the 21st century; we will witness on the order of 20,000 years of progress, or about 1,000 times greater than what was achieved in the 20th century.”

According to Kurzweil, here’s what we can expect in the not-so-distant future:

- Doctors will be doing a backup of our memories by the late 2030s.
- By the late 2020s, doctors will be sending intelligent bots, or nanobots, into our bloodstreams to keep us healthy, and into our brains to keep us young.

- In 15 years, human longevity will be greatly extended. By the 2020s, we'll be adding a year of longevity or more for every year that passes.
- In the same time frame, we'll routinely be in virtual reality environments. Instead of making a cell call, we could "meet" someone in a virtual world and take a walk on a virtual beach and chat. Business meetings and conference calls will be held in calming or inspiring virtual locations.
- When you're walking down the street and see someone you've met before, background information about that person will pop up on your glasses or in the periphery of your vision.
- Instead of spending hours in front of a desktop machine, computers will be more ingrained in our environment. For instance, computer monitors could be replaced by projections onto our retinas or on a virtual screen hovering in the air.
- Scientists will be able to rejuvenate all of someone's body tissues and organs by transforming their skin cells into youthful versions of other cell types.
- Need a little boost? Kurzweil says scientists will be able to regrow our own cells, tissues, and even whole organs, and then introduce them into our bodies, all without surgery. As part of what he calls the "emerging field of rejuvenation medicine," new tissue and organs will be built out of cells that have been made younger.
- Got heart trouble? No problem, says Kurzweil. "We'll be able to create new heart cells from your skin cells and introduce them into your system through the bloodstream. Over time, your heart cells get replaced with these new cells, and the result is a rejuvenated, young heart with your own DNA."
- One trick we'll have to master is staying ahead of the game. Kurzweil warns that terrorists could obviously use this same technology against us. For example, they could build and spread a bioengineered biological virus that's highly powerful and stealthy.

According to Kurzweil, we're not that far away from solving a medical problem that has plagued scientists and doctors for quite some time now: the common cold. He notes that though nanotechnology could go into our bloodstreams and knock it out, before we even get to that stage, biotechnology should be able to cure the cold in just 10 years.

Source: Adapted from Sharon Gaudin, "Kurzweil: Computers Will Enable People to Live Forever," *InformationWeek*, November 21, 2006.

Summary

- **Computer Systems.** Major types of computer systems are summarized in Figure 3.3. Microcomputers are used as personal computers, network computers, personal digital assistants, technical workstations, and information appliances. Midrange systems are increasingly used as powerful network servers and for many multiuser business data processing and scientific applications. Mainframe computers are larger and more powerful than most midsize systems. They are usually faster, have more memory capacity, and can support more network users and peripheral devices. They are designed to handle the information processing needs of large organizations with high volumes of transaction processing or with complex computational problems. Supercomputers are a special category of extremely powerful mainframe computer systems designed for massive computational assignments.
- **The Computer Systems Concept.** A computer is a system of information processing components that perform input, processing, output, storage, and control functions. Its hardware components include input and output devices, a central processing unit (CPU), and primary and secondary storage devices. The major functions and hardware in a computer system are summarized in Figure 3.10.
- **Peripheral Devices.** Refer to Figures 3.14 and 3.22 to review the capabilities of peripheral devices for input, output, and storage discussed in this chapter.

Key Terms and Concepts

These are the key terms and concepts of this chapter. The page number of their first explanation is given in parentheses.

- | | | |
|---|---|---|
| 1. Binary representation (108) | 16. MIPS (million instructions per second) (94) | 29. RFID (radio frequency identification) (115) |
| 2. Central processing unit (93) | 17. Moore's law (94) | 30. Secondary storage (93) |
| 3. Computer system (92) | 18. Network computer (86) | 31. Semiconductor memory (110) |
| 4. Computer terminal (85) | 19. Network server (83) | <i>a.</i> RAM (random-access memory) (110) |
| 5. Cycles per second (94) | 20. Network terminal (85) | <i>b.</i> ROM (read-only memory) (110) |
| 6. Direct access (109) | 21. Off-line (97) | 32. Sequential access (109) |
| 7. Graphical user interface (97) | 22. Online (97) | 33. Speech recognition (101) |
| 8. Information appliance (86) | 23. Optical disks (114) | 34. Storage capacity (109) |
| 9. Magnetic disks (112) | 24. Optical scanning (103) | <i>a.</i> Bit (108) |
| <i>a.</i> Floppy disk (112) | 25. Peripherals (97) | <i>b.</i> Byte (108) |
| <i>b.</i> Hard disk (113) | 26. Pointing devices (97) | <i>c.</i> Kilobyte (109) |
| <i>c.</i> RAID (redundant array of independent disks) (113) | 27. Primary storage unit (93) | <i>d.</i> Megabyte (109) |
| 10. Magnetic stripe (105) | 28. Processing speed (94) | <i>e.</i> Gigabyte (109) |
| 11. Magnetic tape (113) | <i>a.</i> Millisecond (94) | <i>f.</i> Terabyte (109) |
| 12. Mainframe system (89) | <i>b.</i> Microsecond (94) | <i>g.</i> Petabyte (109) |
| 13. Microcomputer (83) | <i>c.</i> Nanosecond (94) | 35. Supercomputer (90) |
| 14. Midrange system (87) | <i>d.</i> Picosecond (94) | 36. Volatility (110) |
| 15. Minicomputer (88) | | 37. Workstation computer (83) |

Review Quiz

Match one of the previous key terms and concepts with one of the following brief examples or definitions. Try to find the best fit for answers that seem to fit more than one term or concept. Defend your choices.

- | | |
|---|---|
| ___ 1. A computer is a combination of components that perform input, processing, output, storage, and control functions. | ___ 11. A common computer interface using a desktop metaphor and icons. |
| ___ 2. The main processing component of a computer system. | ___ 12. Can be a desktop/laptop or handheld computer. |
| ___ 3. A measure of computer speed in terms of processor cycles. | ___ 13. A computer category between microcomputers and mainframes. |
| ___ 4. Devices for consumers to access the Internet. | ___ 14. A small, portable magnetic disk encased in a thin plastic shell. |
| ___ 5. The memory of a computer. | ___ 15. A large-capacity disk typically found in computer systems. |
| ___ 6. Magnetic disks and tape and optical disks perform this function. | ___ 16. Low-cost microcomputers for use with the Internet and corporate intranets. |
| ___ 7. Input/output and secondary storage devices for a computer system. | ___ 17. A redundant array of inexpensive hard drives. |
| ___ 8. Connected to and controlled by a CPU. | ___ 18. A terminal that depends on network servers for its software and processing power. |
| ___ 9. Separate from and not controlled by a CPU. | ___ 19. A computer that manages network communications and resources. |
| ___ 10. Results from the presence or absence or change in direction of electric current, magnetic fields, or light rays in computer circuits and media. | ___ 20. The most powerful type of computer. |

- 21. A magnetic tape technology for credit cards.
- 22. One-billionth of a second.
- 23. Roughly 1 billion characters of storage.
- 24. Includes electronic mice, trackballs, pointing sticks, and touch pads.
- 25. Early midrange systems used for processing-intensive applications such as scientific research and engineering analysis.
- 26. The largest of the three main types of computers.
- 27. Processor power measured in terms of number of instructions processed.
- 28. Prediction that computer power will double approximately every 18 to 24 months.
- 29. Promises to be the easiest, most natural way to communicate with computers.
- 30. Capturing data by processing light reflected from images.
- 31. The speed of a computer.
- 32. One one-thousandth of a second.
- 33. 1,024 bytes.
- 34. A device with a keyboard and a video display networked to a computer is a typical example.
- 35. The amount of data a storage device can hold.
- 36. A personal computer used as a technical workstation.
- 37. The smallest unit of data storage.
- 38. One trillion bytes.
- 39. You cannot erase the contents of these storage circuits.
- 40. The memory of most computers consists of these storage circuits.
- 41. The property that determines whether data are lost or retained when power fails.
- 42. Each position of storage can be accessed in approximately the same time.
- 43. Each position of storage can be accessed according to a predetermined order.
- 44. Microelectronic storage circuits on silicon chips.
- 45. Uses magnetic spots on metal or plastic disks.
- 46. Uses magnetic spots on plastic tape.
- 47. Uses a laser to read microscopic points on plastic disks.
- 48. A millionth of a second.
- 49. A trillionth of a second.
- 50. A grouping of eight bits that represents one alphabetic or special character.
- 51. A short-range wireless technology most commonly used to tag, track, and identify objects.
- 52. Around a million bytes; more precisely, 2 to the 20th power.
- 53. A unit of information or computer storage equal to one quadrillion bytes, or 1,024 terabytes.

REYNOLDS, DONALD

Discussion Questions

1. What trends are occurring in the development and use of the major types of computer systems?
2. Will the convergence of PDAs, subnotebook PCs, and cell phones produce an information appliance that will make all of those categories obsolete? Why or why not?
3. Refer to the Real World Case IT asset management in the chapter. What advice would you provide to a growing company to avoid facing the issues discussed in the case?
4. Do you think that information appliances like PDAs will replace personal computers (PCs) in business applications? Explain.
5. Are networks of PCs and servers making mainframe computers obsolete? Explain.
6. Refer to the Real World Case on speech recognition in health care in the chapter. Although these and other technologies are becoming more prevalent in health care, doctors have traditionally been reluctant to adopt them. Why do you think this is the case? How would these technologies change the way doctors perform their job?
7. What are several trends that are occurring in computer peripheral devices? How do these trends affect business uses of computers?
8. What are several important computer hardware developments that you expect to happen in the next 10 years? How will these affect the business use of computers?
9. What processor, memory, magnetic disk storage, and video display capabilities would you require for a personal computer that you would use for business purposes? Explain your choices.
10. What other peripheral devices and capabilities would you want to have for your business PC? Explain your choices.

WAYNE 2833BU

Analysis Exercises

1. Hardware Costs

Purchasing Computer Systems for Your Workgroup

You have been asked to get pricing information for a potential purchase of PCs for the members of your workgroup. Go to the Internet to get prices for these units from Dell and Hewlett-Packard. Look for a high-end office desktop model.

The table below shows the specifications for the basic system you have been asked to price and potential upgrades to each feature. You will want to get a price for the basic system described below and a separate price for each of the upgrades shown.

Component	Basic Unit	Upgrade
CPU (gigahertz)	2.8	3.4
Hard drive (gigabytes)	160	500
RAM (gigabytes)	1	2
Removable media	16× DVD+R/W	48× DVD+R/W
Monitor	17-inch flat screen	19-inch flat screen

Select the standard software licenses; your IT department will install the necessary software for your workgroup. Take a two-year warranty and servicing coverage offered by each supplier. If a two-year warranty is not available, simply note any differences in the coverage with the closest match.

- Prepare a spreadsheet summarizing this pricing information and showing the cost from each supplier of the following options: (1) units with the basic configuration, (2) the incremental cost of each upgrade separately, and (3) the cost of a fully upgraded unit. If you cannot find features that exactly match the requirements, then use the next higher standard for comparison and make a note of the difference.
- Prepare a set of PowerPoint slides summarizing your results. Include a discussion of the warranty and servicing contract options offered by each supplier.

2. Price and Performance Trends for Computer Hardware

Hardware Analysis

The table below details price and capacity figures for common components of personal computers. Typical prices for microprocessors, random-access memory (RAM), and hard disk storage are displayed.

The performance of typical components has increased substantially over time, so the speed (for the microprocessor) or the capacity (for the storage devices)

is also listed for comparison purposes. Although not all improvements in these components are reflected in these capacity measures, it is interesting to examine trends in these measurable characteristics.

- Create a spreadsheet based on the figures above and include a new row for each component, showing the price per unit of capacity (cost per megahertz of speed for microprocessors and cost per megabyte of storage for RAM and hard disk devices).
- Create a set of graphs highlighting your results and illustrating trends in price per unit of performance (speed) or capacity.
- Write a short paper discussing the trends you found. How long do you expect these trends to continue? Why?
- Prepare a summary presentation outlining the points from your paper (above). Be sure to *link* your Excel chart into the PowerPoint presentation so that it automatically updates when any data change in the spreadsheet.

3. Can Computers Think Like People?

The Turing Test

The Turing test is a hypothetical test to determine whether a computer system has reached the level of artificial intelligence. If the computer can fool a person into thinking it is another person, then it has artificial intelligence. Except in very narrow areas, no computer has passed the Turing test.

Free e-mail account providers such as Hotmail or Yahoo take advantage of this fact. They need to distinguish between new account registrations generated by a person and registrations generated by spammers' software. Why? Spammers burn through thousands of e-mail accounts to send millions of e-mails. To help them, spammers need automated tools to generate these accounts. Hotmail fights this practice by requiring registrants to enter correctly an alphanumeric code hidden within an image. Spammers' programs have trouble correctly reading the code, but most humans do not. With this reverse Turing test, also called a CAPTCHA, Hotmail can distinguish between a person and a program and allow only humans to register. As a result, spammers must look elsewhere for free accounts.

- Aside from those mentioned above, in what applications might businesses find it useful to distinguish between a human and a computer?

	1991	1993	1995	1997	1999	2001	2003	2005
Processor: Speed, MHz	25	33	100	125	350	1000	3,000	3,800
Cost	\$180	\$125	\$275	\$250	\$300	\$251	\$395	\$549
RAM chip: MB per chip	1	4	4	16	64	256	512	2,000
Cost	\$55	\$140	\$120	\$97	\$125	\$90	\$59	\$149
Hard drive: GB per drive	0.105	0.250	0.540	2.0	8.0	40.0	160.0	320
Cost	\$480	\$375	\$220	\$250	\$220	\$138	\$114	\$115

REYNOLDS

DONALD WARENE

2833BU

- b. Describe a Turing test that a visually impaired person, but not a computer, might pass.
- c. Search the Internet for the term CAPTCHA and describe its strengths and weaknesses.

3. Radio Frequency Identification

Input Device or Invasion of Privacy?

Punch cards, keyboards, bar code scanners—the trend is clear. Input devices have continued to promote faster and more accurate data entry. Key to this advance is capturing data at their source, and no tool does this better than radio frequency identification (RFID) systems. An RFID transmitter sends out a coded radio signal. An RFID tag changes and reflects this signal back to an antenna. The RFID system can read the reflection's unique pattern and record it in a database. Depending on the system, this pattern may be associated with a product line, shipping palette, or even a person. Although an RFID system's range is limited to a few dozen feet, this approach enables remarkable inventory tracking that doesn't rely on a human to keyboard interaction or scan. Except for the presence of a 1-inch-square

(5-cm-square) RFID tag, humans may have no idea an RFID system is in operation.

Indeed, that may be part of the problem. Consumers have expressed concern that RFID chips attached to products they purchase may be used to track them. Others fear their government may require embedded RFID chips as a form of personal identification and tracking. What started as a new and improved input device has devolved into a matter of public policy.

- a. How would you feel if your university used RFID tags embedded in student IDs to replace the magnetic swipe strip? On a campus, RFID tags might be used to control building access, manage computer access, or even automatically track class attendance.
- b. Enter “RFID” into an Internet search engine and summarize the search results. Of the top 20 results, how many were positive, negative, or neutral?
- c. Enter “RFID” and “privacy” into an Internet search engine, select a page expressing privacy concerns, and summarize them in a brief essay. Do you find these concerns compelling?

REYNOLDS

,

DONALD

WAYNE

2833BU

REAL WORLD CASE

3

IBM, Wachovia, and PayPal: Grid Computing Makes It Easier and Cheaper

IBM researchers and a team of doctors are building a database of digital images they hope will enable oncologists to diagnose and treat cancer patients faster and with more success. Researchers at the Cancer Institute of New Jersey have digitized CAT scans, MRIs, and other images using a high-performance system and computational time on the World Community Grid, also known as the world's largest public computing grid.

"Digitizing images should enable doctors to diagnose cancers earlier and detect their growth or shrinkage more accurately during treatment," says Robin Willner, vice president of the global community initiatives at IBM. "Right now, the doctor is basically eyeballing it when he's analyzing tissues and biopsies. They're trying to figure out what type of cancer it is and if there's been progress during treatment. If you digitize the image, you're able to compare numbers because you've turned an image into bits and bytes. Now it's a much more accurate comparison."

Researchers have been using the grid to convert hundreds of thousands of images of cancerous tissues and cells into digital images. Once the images are digitized, the grid can check the accuracy of the digital information to ensure that the bits and bytes are translating into real diagnoses. The World Community Grid acts as a virtual supercomputer that is based on thousands of volunteers donating their unused computer time. "If we can improve treatment and diagnosis for cancer, that's great for everybody," said Willner. "There couldn't be a better use for the grid."

The next phase of the project is to build a database that will hold hundreds of thousands, if not millions, of these images. A \$2.5 million grant from the National Institutes of Health (NIH) will enable the Cancer Institute of New Jersey, Rutgers University, and cancer centers around the country to pool their digital images in the database. Willner said the database will enable doctors to compare patients' new images to ones already in the database to help them diagnose the cancer and figure out the best way to treat it. Doctors should be able to use the database to personalize treatments for cancer patients based on how other patients with similar protein expression signatures and cancers have reacted to various treatments.

"The overarching goal of the new NIH grant is to expand the library to include signatures for a wider range of disorders and make it, along with the decision-support technology, available to the research and clinical communities as grid-enabled deployable software," said David J. Foran, a director of the Cancer Institute of New Jersey. "We hope to deploy these technologies to other cancer research centers around the nation."

This isn't IBM's first foray into the medical arena by any means. IBM has also teamed up with the Mayo Clinic to develop a research facility to advance medical imaging. Researchers from both the Mayo Clinic and IBM are working

at the new Medical Imaging Informatics Innovation Center in Rochester, Minnesota. Bradley Erickson, chairman of radiology at the Mayo Clinic, said a joint team is already working to find ways to use the Cell chip, mostly known for running inside the PlayStation 3 video-game console, in a medical imaging system. Erickson said that the technology could either reduce work that now takes minutes to a matter of seconds, or work that now takes hours to only minutes.

Grid computing, however, is not limited to nonprofit institutions. Financial services firm Wachovia Corp. has freed some of its Java-based applications from dedicated servers and is allowing these transaction applications to draw computing power from a 10,000-CPU resource pool on servers spread across cities in the United States and in London. Wachovia is tapping into computing power that's available on other systems to perform work. That capability allows companies to avoid dedicated hardware costs and make better use of underutilized hardware.

Tony Bishop, a Wachovia senior vice president and director of product management, said that to use dedicated systems as an alternative would be "three times the cost in terms of capital and people to support it otherwise." Wachovia has eight applications running on its grid that are used in internal transactions, such as order management. The servers are in New York, Philadelphia, London, and at the company's corporate headquarters in Charlotte, North Carolina. Jamie Bernardin, chief technology officer at DataSynapse, the company that developed the technology, said that to improve transaction speeds, the transaction application running on it can grow and contract as needed.

Because the system can provide resources as needed for the applications, Bishop said performance has improved on some transactions fivefold. "This ability to speed processing means decisions and services can be made and delivered more rapidly. As things get more and more automated and more and more real time it will be IT in this business that differentiates," says Bishop.

A Linux grid is the power behind the payment system at PayPal, and it has converted a mainframe believer. Scott Thompson, the former executive VP of technology solutions at Inovant, ran the Visa subsidiary responsible for executing Visa credit card transactions worldwide. The VisaNet system was strictly based on IBM mainframes.

In February 2005, Thompson became chief technology officer at the eBay payments company, PayPal, where he confronted a young Internet organization building its entire transaction processing infrastructure on open-source Linux and low-cost servers. Hmmmm, he thought at the time. "I came from Visa, where I had responsibility for VisaNet. It was a fabulous processing system, very big and very global. I was intrigued by PayPal. How would you use Linux for processing payments and never be wrong, never lose messages, never fall behind the pace of transactions?" he wondered.

He now supervises the PayPal electronic payment processing system, which is smaller than VisaNet in volume and total dollar value of transactions, but it's growing fast. It is currently processing \$1,571 worth of transactions per second in 17 different currencies. In 2006, the online payments firm, which started out over a bakery in Palo Alto, processed a total of \$37.6 billion in transactions. It's headed toward \$50 billion very soon.

Now located in San Jose, PayPal grants its consumer members options in payment methods: credit cards, debit cards, or directly from a bank account. It has 165 million account holders worldwide, and it has recently added such businesses as Northwest Airlines, Southwest Airlines, U.S. Airways, and Overstock.com, which now permit PayPal payments on their Web sites.

Thompson supervises a payment system that operates on about 4,000 servers running Red Hat Linux in the same manner that eBay and Google conduct their business on top of a grid of Linux servers. "I have been pleasantly surprised at how much we've been able to do with this approach. It operates like a mainframe," he says.

As PayPal grows, it's much easier to grow the grid with Intel-based servers than it would be to upgrade a mainframe,

according to Thompson. "The cost to increase capacity a planned 15 or 20 percent in a mainframe environment is enormous. It could be in the tens of millions to do a step increase. In PayPal's world, we add hundreds of servers in the course of a couple of nights and the cost is in the thousands, not millions."

PayPal takes Red Hat Enterprise Linux and strips out all features unnecessary to its business, and then adds proprietary extensions around security. Another virtue of the grid is that PayPal's 800 engineers can all get a copy of that customized system on their development desktops, run tests on their raw software as they work, and develop to PayPal's needs faster because they're working in the target environment. That's harder to do when the core of the data center consists of large boxes or mainframes. It's not cheap in either case to install duplicates for developers, says Thompson.

Source: Adapted from Sharon Gaudin, "IBM Uses Grid to Advance Cancer Diagnosis and Treatment," *Computerworld*, January 28, 2008; Patrick Thibodeau, "Wachovia Uses Grid Technology to Speed Up Transaction Apps," *Computerworld*, May 15, 2006; and Charles Babcock, "PayPal Says Linux Grid Can Replace Mainframes," *InformationWeek*, November 28, 2007.

CASE STUDY QUESTIONS

1. Applications for grid computing in this case include medical diagnosis and financial transaction processing. What other areas do you think would be well suited to the use of grid computing and why? Provide several examples from organizations other than those included in the case.
2. The joint effort by IBM and the Cancer Institute of New Jersey works by digitalizing medical diagnoses on the World Community Grid (WCG). What are the advantages and disadvantages of relying on a volunteer-based network such as this? Provide examples of both. Visit the Web site of the WCG to inform your answer.
3. IBM, Wachovia, and PayPal are arguably large organizations. However, several vendors have started offering computing power for rent to smaller companies, using the principles underlying grid computing. How could small and medium companies benefit from these technologies? Search the Internet for these offerings to help you research your answer.

REAL WORLD ACTIVITIES

1. Grid computing technology is becoming increasingly popular and has recently received support from giants such as IBM, Sun, and Oracle. Visit their Web sites (www.ibm.com, www.sun.com, and www.oracle.com) and review their current offerings in this regard. How do their products compare to each other? Prepare a presentation to share your findings with the class.
2. One of the main benefits of grid computing arises from the possibility of replacing expensive hardware, such as mainframes or supercomputers, with commodity-priced servers and even personal computers. What about the cost of administering so many different servers and the power consumption associated with them? Go online to search for information that would allow you to compare grid computing to more traditional, mainframe-based alternatives. Write a report to present your findings.

REAL WORLD

CASE

4

Apple, Microsoft, IBM, and Others:
The Touch Screen Comes of Age

The WIMP human-computer interface may have an uninspiring name, but Windows, Icons, Menus, and Pointing (WIMP) devices have dominated computing for some 15 years. The keyboard, mouse, and display screen have served users extraordinarily well.

Now the hegemony of WIMP may be coming to an end, say developers of technologies based on human touch and gesture. For evidence, look no further than Apple's iPhone. From a human-interface point of view, the combined display and input capabilities of the iPhone's screen, which can be manipulated by multiple fingers in a variety of intuitive touches and gestures, is nothing short of revolutionary.

The iPhone isn't the only commercial device to take the human-computer interface to a new level. The Microsoft Surface computer puts input and output devices in a large tabletop device that can accommodate touches and gestures and even recognize physical objects laid on it. In addition, the DiamondTouch Table from Mitsubishi is a touch- and gesture-activated display that supports small-group collaboration. It can even tell who is touching it.

These devices point the way toward an upcoming era of more natural and intuitive interaction between human and machine. Robert Jacob, a computer science professor at Tufts University, says touch is just one component of a booming field of research on post-WIMP interfaces, a broad coalition of technologies he calls reality-based interaction.

Those technologies include virtual reality, context-aware computing, perceptual and affective computing, and tangible interaction, in which physical objects are recognized directly by a computer.

"What's similar about all these interfaces is that they are more like the real world, Jacob says.

For example, the iPhone "uses gestures you know how to do right away, such as touching two fingers to an image or application, then pulling them apart to zoom in or pinching them together to zoom out." These actions have also found their way into the iPod Touch and the track pad of the new MacBook Air. "Just think of the brain cells you don't have to devote to remembering the syntax of the user interface! You can devote those brain cells to the job you are trying to do." In particular, he says, the ability of the iPhone to handle multiple touches at once is a huge leap past the single-touch technology that dominates in traditional touch applications such as ATMs.

Although they have not gotten much traction in the marketplace yet, advanced touch technologies from IBM may point a way to the future. In its Everywhere Displays Project, IBM mounts projectors in one or more parts of an ordinary room and projects images of touch screens onto ordinary surfaces, such as tables, walls, or the floor. Video cameras capture images of users touching various parts of the surfaces and send that information for interpretation by a computer. The touch screens contain no electronics—

indeed, no computer parts at all—so they can be easily moved and reconfigured.

A variation on that concept has been deployed by a wine store in Germany, says Claudio Pinhanez at IBM Research. The METRO Future Store in Rheinberg has a kiosk that enables customers to get information about the wines the store stocks. "But the store's inventory was so vast customers often had trouble finding the particular wine they wanted on the shelf. They often ended up buying a low-margin wine in a nearby bin of sales specials," Pinhanez says. Now the kiosk contains a "show me" button that, when pressed, shines a spotlight on the floor in front of the chosen item.

IBM is also working on a prototype system for grocery stores that might, for example, illuminate a circle on the floor that asks, Do you want to take the first steps toward more fiber in your diet? If the customer touches "yes" with his foot, the system projects footsteps to the appropriate products, such as high-fiber cereal. "Then you could make the cereal box itself interactive," says Pinhanez. "You touch it, and the system would project information about that box on a panel above the shelf." Asked if interactive cereal boxes might be a solution in search of a problem, Pinhanez says, "The point is, with projection and camera technology you can transform any everyday object into a touch screen." He says alternatives that are often discussed (e.g., a store system that talks to customers through their handheld devices) are hard to implement because of a lack of standards for the technology.

Microsoft is working with several commercial partners, including Starwood Hotels & Resorts, which owns the prestigious Sheraton, W, Westin, and Méridien brands, among others, to introduce Surface. It will initially target leisure, entertainment, and retail applications, says Mark Bolger, director of marketing for Surface Computing. For example, he says, one could imagine a hotel guest using a virtual concierge in a Surface computer in the lobby to manipulate maps, photos, restaurant menus, and theater information.

Some researchers say that a logical extension of touch technology is gesture recognition, by which a system recognizes hand or finger movements across a screen or close to it without requiring an actual touch. "Our technology is half-way there," IBM's Pinhanez says, "because we recognize the gesture of touching rather than the occlusion of a particular area. You can go over buttons without triggering them."

Patrick Baudisch at Microsoft Research says the Microsoft prototypes can already act on finger gestures, with the system recognizing finger motions, as well as positions, and understanding the meaning of different numbers of fingers. For example, the motion of one finger is seen as equivalent to a mouse movement, a finger touch is interpreted as a click, and two fingers touching and moving is seen as a scroll command.

Touch technology in its many variations is an idea whose time has come. "It's been around a long time, but traditionally

in niche markets. The technology was more expensive, and there were ergonomic problems,” he says. “But it’s all kind of coming together right now.” The rise of mobile devices is a big catalyst, because the devices are getting smaller and their screens are getting bigger. When a screen covers the entire device, there is no room for conventional buttons, which makes it necessary to have other types of interaction (e.g., voice).

Of course, researchers and inventors have envisioned even larger touch displays, including whole interactive walls. A quick YouTube search for “multitouch wall” shows that a number of these fascinating devices have reached the prototype stage, causing multitudes at technology conferences to be entranced. Experts predict, however, that this is just the beginning.

Pradeep Khosla, professor of electrical and computer engineering and robotics at Carnegie Mellon University in Pittsburgh, says touch technology will proliferate, but not by itself. “When we talk face to face, I make eye gestures, face gestures, hand gestures, and somehow you interpret them all to understand what I am saying. I think that’s where we are headed,” he says. “There is room for all these things, and multimodal gestures will be the future.”

Bill Buxton, a researcher at Microsoft, also anticipates a fusion of different interaction technologies. “Touch now may be where the mouse was in about 1983,” Buxton says. “People now understand there is something interesting here that’s different. But I don’t think we yet know what that difference could lead to. Until just one or two years ago there was a real separation between input devices

and output devices. A display was a display and a mouse was a mouse.”

“There’s been this notion that less is more—you try to get less and less stuff to reduce complexity,” he says. “But there’s this other view that more is actually less—more of the right stuff in the right place, and complexity disappears.” In the office of the future, Buxton predicts, desktop computers might be much the same as they are today. “But you can just throw stuff, with the mouse or a gesture, up onto a wall or whiteboard and then work with it with your hands by touch and gesture standing up. Then you’ll just pull things into your mobile and have this surface in your hand. The mobile, the wall, the desktop—they are all suitable for different purposes.”

Will that be the end of the WIMP interface? Tufts University’s Jacob advises users not to discard their keyboards and mice anytime soon. “They really are extremely good,” he says. “WIMP almost completely dislodged the command-line interface. The WIMP interface was such a good invention that people just kind of stopped there, but I can’t believe it’s the end of the road forever.”

Buxton agrees. “WIMP is the standard interface going back 20-plus years, and all the applications have been built around that,” he says. “The challenge is, without throwing the baby out with the bath, how do we reap the benefits of these new approaches while preserving the best parts of the things that exist?”

Source: Adapted from Gary Anthes, “Give Your Computer the Finger: Touch-Screen Tech Comes of Age,” *Computerworld*, February 1, 2008.

CASE STUDY QUESTIONS

1. What benefits may Starwood Hotels derive from the introduction of touch-screen technology as noted in the case? What possible disruptions may occur as a result. Provide several examples of each.
2. Bill Buxton of Microsoft stated that “[t]ouch now may be where the mouse was in about 1983.” What do you make of his comments, and what do you think it would take for touch technology to displace the WIMP interface? Justify your answer.
3. Is advanced touch-screen technology really a solution in search of a problem? Do you agree with this statement? Why or why not?

REAL WORLD ACTIVITIES

1. Most of the fame attached to the iPhone has resulted from individual, end-user applications. How could companies use the iPhone as a platform for commercial use? Break into small groups and brainstorm some possible uses of the technology, as well as what benefits organizations can derive from them. Then prepare a presentation to share your ideas with the class.
2. Information technology advances rapidly, and touch screen is no exception. Go online and search for developments more recent than those mentioned in the case. What new large-scale (i.e., wall-sized) applications could you find? Prepare a report comparing new developments with the examples mentioned here.


R
E
Y
N
O
L
D
S
,

D
O
N
A
L
D

W
A
Y
N
E

2
8
3
3
B
U

CHAPTER 4



COMPUTER SOFTWARE

Chapter Highlights

Section I

Application Software: End-User Applications

Introduction to Software

Real World Case: GE, H.B. Fuller Co., and Others:
Successful Implementations of Software-as-a-Service

Business Application Software

Software Suites and Integrated Packages

Web Browsers and More

Electronic Mail, Instant Messaging, and Weblogs

Word Processing and Desktop Publishing

Electronic Spreadsheets

Presentation Graphics

Personal Information Managers

Groupware

Software Alternatives

Section II

System Software: Computer System Management

System Software Overview

Operating Systems

Real World Case: U.S. Department of Defense: Enlisting
Open-Source Applications

Other System Management Programs

Programming Languages

Web Languages and Services

Programming Software

Real World Case: Wolf Peak International: Failure and
Success in Application Software for the Small-to-Medium
Enterprise

Real World Case: Power Distribution and Law
Enforcement: Reaping the Benefits of Sharing Data
through XML

Learning Objectives

1. Describe several important trends occurring in computer software.
2. Give examples of several major types of application and system software.
3. Explain the purpose of several popular software packages for end-user productivity and collaborative computing.
4. Define and describe the functions of an operating system.
5. Describe the main uses of computer programming software, tools, and languages.
6. Describe the issues associated with open-source software.

R
E
Y
N
O
L
D
S
,
D
O
N
A
L
D
W
A
Y
N
E
2
8
3
3
B
U

SECTION I

Application Software: End-User Applications

Introduction to Software

This chapter provides an overview of the major types of software you depend on as you work with computers and access computer networks. It discusses their characteristics and purposes and gives examples of their uses. Before we begin, let's look at an example of the changing world of software in business.

Read the Real World Case discussing some innovative and successful implementations of Software-as-a-Service (SaaS). We can learn a lot about the promise of this approach to technology use from this example. See Figure 4.1.

What Is Software?

To fully appreciate the need for and value of the wide variety of software available, we should be sure we understand what software is. **Software** is the general term for various kinds of programs used to operate and manipulate computers and their peripheral devices. One common way of describing hardware and software is to say that software can be thought of as the variable part of a computer and hardware as the invariable part. There are many types and categories of software. We will focus our attention on the different types of software and its uses in this chapter.

Types of Software

Let's begin our analysis of software by looking at an overview of the major types and functions of **application software** and **system software** available to computer users, shown in Figure 4.2. This figure summarizes the major categories of system and application software we will discuss in this chapter. Of course, this figure is a conceptual illustration. The types of software you will encounter depend primarily on the types of computers and networks you use and on the specific tasks you want to accomplish. We will discuss application software in this section and the major types of system software in Section II.

Application Software for End Users

Figure 4.2 shows that **application software** includes a variety of programs that can be subdivided into general-purpose and function-specific application categories. **General-purpose application programs** are programs that perform common information processing jobs for end users. For example, word processing, spreadsheet, database management, and graphics programs are popular with microcomputer users for home, education, business, scientific, and many other purposes. Because they significantly increase the productivity of end users, they are sometimes known as *productivity packages*. Other examples include Web browsers, e-mail, and groupware, which help support communication and collaboration among workgroups and teams.

An additional common way of classifying software is based on how the software was developed. **Custom software** is the term used to identify software applications that are developed within an organization for use by that organization. In other words, the organization that writes the program code is also the organization that uses the final software application. In contrast, **COTS software** (an acronym that stands for *commercial off-the-shelf*) is developed with the intention of selling the software in multiple copies (and usually for a profit). In this case, the organization that writes the software is not the intended target audience for its use.

Several characteristics are important when describing COTS software. First, as stated in our definition, COTS software products are sold in many copies with minimal changes beyond scheduled upgrade releases. Purchasers of COTS software generally have no control over the specification, schedule, evolution, or access to either the source code or the internal documentation. A COTS product is sold, leased, or licensed to the general public, but in virtually all cases, the vendor of the product retains the intellectual property rights of the software. Custom software, in contrast,

REAL WORLD

CASE

1

GE, H.B. Fuller Co., and Others:
Successful Implementations of
Software-as-a-Service

General Electric's supply chain is not simply enormous. It's a Byzantine web of sourcing partners, touching all corners of the globe: 500,000 suppliers in more than 100 countries that cut across 14 different languages. Each year, GE spends some \$55 billion among its vast supplier base.

Long-time GE CIO Gary Reiner knows this problem all too well, since, among his other duties, he is responsible for how the \$173 billion conglomerate spends that \$55 billion, utilizing GE's Six Sigma practices and taking advantage of its hefty purchasing power. GE, for instance, buys \$150 million in desktops and laptops each year from a single supplier, Dell—"at a very low price," says Reiner.

For years, GE's Global Procurement Group faced a challenging reality: trying to accurately track and make sense of all of the supply chain interactions with half a million suppliers—contracts, compliance initiatives, certifications, and other critical data, which needed to be centrally stored, managed, and made accessible to thousands across the globe. GE was using what it called a Global Supplier Library, a homegrown system that, Reiner says, had a "rudimentary capability." Reiner and his staff knew that GE needed something better, but they

didn't want to build it. They wanted a supplier information system that was easy to use and install, could unite GE's sourcing empire into one central repository, had multilanguage capabilities, and also offered self-service functionality so that each of its suppliers could manage its own data.

The destination was obvious: To achieve one common view of its supplier base, and one version of the truth in all that data, a goal that torments nearly every company today. But to get there, Reiner and his IT and procurement teams took a different route. In 2008, GE bought the application of a little-known Software-as-a-Service (SaaS) vendor that would ultimately become the largest SaaS deployment to date.

"When we judge a solution, we are indifferent to whether it's hosted by a supplier or by us," Reiner says. "We look for the functionality of the solution and at the price." And that, he claims, has been the way they've always operated. Reiner says that his group doesn't see a big difference in cost and in capabilities between on-premise and SaaS products. "And let me emphasize," he adds, "we don't see a big difference in cost either from the point of view of the ongoing operating costs, or the transition costs." Furthermore, when looking at implementation costs, "they're largely around interfacing with existing systems, process changes and data cleansing," he says. "Those three costs exist regardless of whether GE hosts that application or whether the supplier hosts that application."

The Aravo technology platform was untested at GE's level of requirements and, with just 20 or so customers, coupled with the sheer scale of GE's needs, did not really concern Reiner. "We could have been concerned about that," he concedes. "But that would have also been a concern if we had hosted the software on our own servers. We knew Aravo could handle it." Plus, Reiner says that no other supply chain vendor offered the type of functionality that Aravo's Supplier Information Management (SIM) product offered, and Reiner and his team reasoned that it was much cheaper to buy than build. "We'd much rather work with them," he says, "than build it on our own." One GE sourcing manager told Aravo that GE's ROI on the project is not just positive, "it's massively positive."

"They're using SaaS for 100,000 users and 500,000 suppliers in six languages: that's a major technology deployment shift," says Mickey North Rizza, research director at AMR Research. She says that the sheer volume of transactions, combined with the fact that GE supply chain and procurement employees around the world can now access the same sourcing partner information, all from the same central spot, is significant not only for the supply chain management space but also for the SaaS and cloud computing world. "Finally we have a very large company tackling the data transparency issue by using a SaaS product," North Rizza says. "It's a huge deal."

So far, the thorny issue of data quality in GE's supplier data has been improved, because suppliers now use the self-service capabilities in the SaaS system to manage their own data.

FIGURE 4.1



Software-as-a-Service enables one of the largest and most impressive supply chains in the world.

Source: © Ryan McVay/Getty Images.

REYNOLDS, DONALD W. ARAYNE

GE has 327,000 employees worldwide, and its sourcing systems have more than 100,000 users. There is still more work to do to the SIM platform—for example, GE sourcing employees will add more workflows and new queries to the system; more languages might be added as well (six are operational now).

Reiner says that GE is committed to working with Aravo for the long term and that the system has performed well so far. And SaaS, as an application delivery mechanism, appears to have a bright future at GE.

When Steven John took over as CIO at specialty chemical manufacturer H.B. Fuller Co., he inherited a North American payroll system implementation that was expensive and going nowhere. The business units hadn't participated in the technology decision, and the project was bogged down with customization issues and other concerns. John chose to relinquish control of payroll software and switched to SaaS.

"I wanted to do an implementation that was simple and straightforward—to configure but not customize—and see the benefits of a standard, global platform," John says. "This was a way to teach, save money and outsource a noncore system." Giving up control was an easy trade-off compared with the headaches he would face trying to fix the existing software.

"You're getting a lot more innovation," says Ray Wang, an analyst at Forrester Research Inc. "The products are a lot more configurable than what most people have in their own applications. You can change fields, rename things, and move attributes and workflows. So there's a good level of control there."

What's more, the configuration choices are more refined and well thought out, giving users a few good choices instead of myriad options. John found that configuration rather than customization allows H.B. Fuller to maintain its "lean core." "I believe that more standardization leads to more agility," John says. "SaaS allows us to say, 'This is good enough . . . for what we need.' So you don't end up with these horrible situations where you have these highly customized systems. We go with configuration option A, B or C. If one of those three doesn't meet our need, we can try to influence the next release. But in most cases, A, B or C is going to meet the need."

At H.B. Fuller, the move to SaaS for human resources tools allowed the company to empower its people. "I can do

a reorganization and have it reflected within minutes, and I don't have to call someone in HR to update everything," John says. "I can also pull up other people's organization charts and see where they are and what they're doing and better understand the organization."

When it comes to managing SaaS, neither the IT department nor the business unit using the software should be eager to relinquish control. "The buying decisions are shifting from IT to the business leaders," who often opt to charge the software as an expense rather than wait for approval through the capital budget committee, Wang says. Still, he adds, "it's very important to engage IT in these SaaS decisions because there are overall IT architectures and blueprints to consider." It becomes very costly when applications don't integrate or interoperate well with one another.

"It's good to at least have some parameters and policies in place so that people understand what type of apps will work better within the environment, what will be cheaper to share information and data with," says Wang.

One of the problems with SaaS is that if your vendor were to go bankrupt, everything would shut down. You don't own the software. It's on lease. The question is, what do you own? If the vendor doesn't have a separate on-premises deployment option, "you need the ability to take out transactional data, master file information, any kind of migration programs, just in case, so you can convert it to an on-premises alternative if they were to go down," Wang says.

In the long term, Wang envisions an IT culture where software as a service is commonplace. "We may live in a world where everything is provisioned. All our applications don't stay on premises, and business leaders are out procuring applications," he says. "IT teams are testing them to make sure they work well in the environment and there are no bugs or viruses and things integrate well, and basically the IT staff will spend a lot of time provisioning services and implementing, integrating, doing installs. That's where we envision the market in 2020."

Source: Adapted from Thomas Wailgum, "GE CIO Gets His Head in the Cloud for New SaaS Supply Chain App," *CIO Magazine*, January 22, 2009; and Stacy Collett, "SaaS Puts Focus on Functionality," *Computerworld*, March 23, 2009.

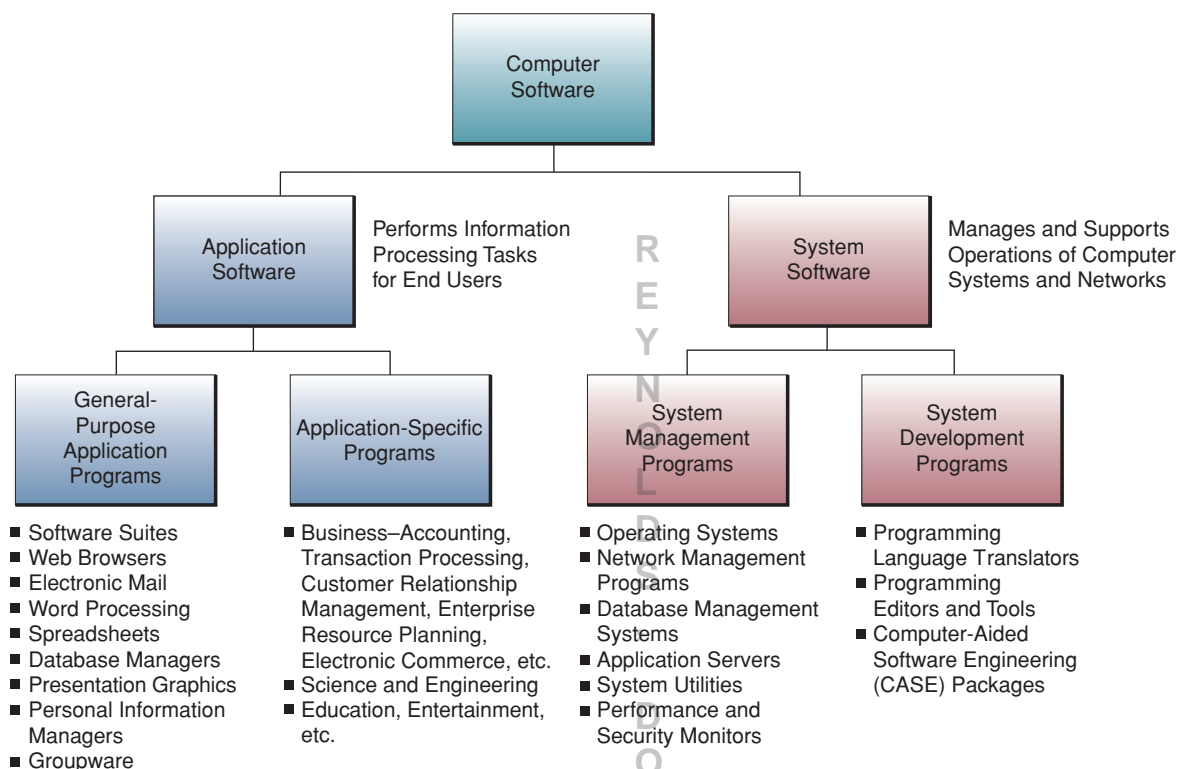
CASE STUDY QUESTIONS

1. What factors should companies take into consideration when making the decision between developing their own applications, purchasing them from a vendor, or taking the SaaS route, as discussed here? Make a list of factors and discuss their importance to this decision.
2. What risks did GE take on when they contracted with a small and less experienced vendor? What contingencies could have been put in place to prevent any problems from arising? Provide several examples.
3. What should companies do if none of the "configuration options" perfectly fits with their needs? Should they attempt to customize, or select the least-worst alternative? When would they do each?

REAL WORLD ACTIVITIES

1. The case mentions that GE's implementation of SaaS was, at the time, the largest rollout of the technology in the world. What other companies have started using SaaS extensively since? Go online and research recent implementations. How are those different from GE's experience? Prepare a report to share your findings.
2. By implementing systems based on SaaS, companies are relinquishing control over ownership of the technology and are putting access to valuable data in the hands of a third party. What are the perils of taking this approach? How could companies guard against them? Break into small groups to discuss these issues and provide some suggestions and recommendations.

FIGURE 4.2 An overview of computer software. Note the major types and examples of application and system software.



is generally owned by the organization that developed it (or that paid to have it developed), and the specifications, functionality, and ownership of the final product are controlled or retained by the developing organization.

The newest innovation in software development is called **open-source software**. In this approach, developers collaborate on the development of an application using programming standards that allow for anyone to contribute to the software. Furthermore, as each developer completes his or her project, the code for the application becomes available and free to anyone else who wishes to use it. We will discuss this new approach to software development in greater detail in Section II of this chapter.

SAP Business Suite 7: Introducing Modular Scenarios Cutting Across Organizational Functions



Germany-based SAP AG is tackling business processes in a novel way with the newest version of its Business Suite, which embeds analytics acquired from Business Objects SA and introduces industry-specific “value scenarios.” Version 7.0 of SAP Business Suite, a library of business processes, adds industry best practices through more than 30 modular value scenarios—like Superior Customer Value, Product Lifecycle Management (PLM)—designed to cross traditional organizational boundaries.

These “pre-defined end-to-end business processes” are intended to be implemented in small steps by organizations as they need it, says Jim Hagemann Snabe, SAP executive board member. The value scenarios basically illustrate interrelationships between SAP product capabilities using graphical guides and business terms, not feature and function lists. The customer can also see the impact on the associated systems, and ultimately, the specific SAP modules that would need to be activated.

Ray Wang, vice president at Cambridge, Massachusetts-based Forrester Research Inc., says customers will find the value scenarios “compelling as they align with the key business drivers users face.” But as with all best practices, Wang notes that “SAP will need to make it easy for customers to modify those scenarios, reduce the overall cost of owning SAP, and provide more frequent levels of innovation.”

One customer, Colgate-Palmolive Co., has large implementations in CRM and PLM that would benefit from the new capabilities of version 7.0, says the company’s senior vice president of IT and business services, Ed Toben. “Particularly when you look at PLM, which is newer, the processes and the enhancement pack concept of turning on pieces should make us move faster,” says Toben.

Another customer, pharmaceutical company Roche, requires the flexibility and ability to scale as the business changes in order to remain current, says chief information officer Jennifer Allerton. “IT investments . . . have got to make sense in their own right,” she says. “And, the pharmaceuticals business is one where you invest for the long term and when you make investments about IT packages, you’re not going to change your mind the next day about them.”

IBM Corp., also a customer, is focused on a number of transformation programs, including the area of operational efficiency, says Jeannette Horan, vice president of enterprise business transformation with the office of the chief information officer at IBM. To that end, the company’s strategy, says Horan, is to globally integrate the enterprise through common processes, using the Business Suite, she says, to “mix and match components of the business to go to market in new and interesting ways.”

But while companies are taking a hard look at spending and reviewing projects, “that does not mean . . . that companies do not spend, they just spend very smartly and very wisely,” says Léo Apotheker, co-CEO of SAP AG. There is a need, says Apotheker, “to provide better and faster insight, a higher level of efficiency, a need to introduce a whole new degree of flexibility in the way we do business.”

Source: Adapted from Kathleen Lau, “Industry ‘Value Scenarios’ in SAP Business Suite 7,” *Computerworld Canada*, February 12, 2009.

Business Application Software

Thousands of **function-specific application software** packages are available to support specific applications of end users in business and other fields. For example, business application software supports the reengineering and automation of business processes with strategic e-business applications like customer relationship management, enterprise resource planning, and supply chain management. Other examples are software packages that Web-enable electronic commerce applications or apply to the functional areas of business like human resource management and accounting and finance. Still other software empowers managers and business professionals with decision support tools like data mining, enterprise information portals, or knowledge management systems.

We will discuss these applications in upcoming chapters that go into more detail about these business software tools and applications. For example, data warehousing and data mining are discussed in Chapters 5 and 10; accounting, marketing, manufacturing, human resource management, and financial management applications are covered in Chapters 7 and 8. Customer relationship management, enterprise resource planning, and supply chain management are also covered in Chapter 8. Electronic commerce is the focus of Chapter 9, and decision support and data analysis applications are explored in Chapter 10. Figure 4.3 illustrates some of the many types of business application software that are available today. These particular applications are integrated in the Oracle E-Business Suite software product of Oracle Corp.

FIGURE 4.3 The business applications in Oracle's E-Business Suite software illustrate some of the many types of business application software being used today.



Source: Adapted from Oracle Corp., "E-Business Suite: Manage by Fact with Complete Automation and Complete Information," Oracle.com, 2002.

Software Suites and Integrated Packages

Let's begin our discussion of popular general-purpose application software by looking at **software suites**. The most widely used productivity packages come bundled together as software suites, such as Microsoft Office, Lotus SmartSuite, Corel WordPerfect Office, Sun's StarOffice, and their open-source product, OpenOffice. Examining their components gives us an overview of the important software tools that you can use to increase your productivity.

Figure 4.4 compares the basic programs that make up the top four software suites. Notice that each suite integrates software packages for word processing, spreadsheets, presentation graphics, database management, and personal information management. Microsoft, Lotus, Corel, and Sun bundle several other programs in each suite, depending on the version you select. Examples include programs for Internet access, e-mail, Web publishing, desktop publishing, voice recognition, financial management, and electronic encyclopedias.

A software suite costs a lot less than the total cost of buying its individual packages separately. Another advantage is that all programs use a similar *graphical user interface* (GUI) of icons, tool and status bars, menus, and so on, which gives them the same look and feel and makes them easier to learn and use. Software suites also share common tools such as spell checkers and help wizards to increase their efficiency. Another big advantage of suites is that their programs are designed to work together seamlessly and import each other's files easily, no matter which program you are using at the time. These capabilities make them more efficient and easier to use than a variety of individual package versions.

FIGURE 4.4 The basic program components of the top four software suites. Other programs may be included, depending on the suite edition selected.

Programs	Microsoft Office	Lotus SmartSuite	Corel WordPerfect Office	Sun Open Office
Word Processor	Word	WordPro	WordPerfect	Writer
Spreadsheet	Excel	1-2-3	Quattro Pro	Calc
Presentation Graphics	PowerPoint	Freelance	Presentations	Impress
Database Manager	Access	Approach	Paradox	Base
Personal Information Manager	Outlook	Organizer	Corel Central	Schedule

Of course, putting so many programs and features together in one supersize package does have some disadvantages. Industry critics argue that many software suite features are never used by most end users. The suites take up a lot of disk space (often upward of 250 megabytes), depending on which version or functions you install. Because of their size, software suites are sometimes derisively called *bloatware* by their critics. The cost of suites can vary from as low as \$100 for a competitive upgrade to more than \$700 for a full version of some editions of the suites.

These drawbacks are one reason for the continued use of **integrated packages** like Microsoft Works, Lotus eSuite WorkPlace, and AppleWorks. Integrated packages combine some of the functions of several programs—word processing, spreadsheets, presentation graphics, database management, and so on—into one software package.

Because integrated packages leave out many features and functions that are in individual packages and software suites, they are considered less powerful. Their limited functionality, however, requires a lot less disk space (often less than 10 megabytes), costs less than \$100, and is frequently preinstalled on many low-end microcomputer systems. Integrated packages offer enough functions and features for many computer users while providing some of the advantages of software suites in a smaller package.

Web Browsers and More

The most important software component for many computer users today is the once simple and limited, but now powerful and feature-rich, **Web browser**. Browsers such as Microsoft Explorer, Netscape Navigator, Mozilla Firefox, and Opera are software applications designed to support navigation through the point-and-click hyperlinked resources of the World Wide Web and the rest of the Internet, as well as corporate intranets and extranets. Once limited to surfing the Web, browsers are becoming the universal software platform from which end users launch information searches, e-mail, multimedia file transfers, discussion groups, and many other Internet-based applications.

Figure 4.5 illustrates the use of the Microsoft Internet Explorer browser to access search engines on the Netscape.com Web site. Netscape uses top-rated Google as its

FIGURE 4.5

Using the Microsoft Internet Explorer browser to access Google and other search engines on the Netscape.com Web site.



Source: Netscape and the “N” Logo are registered trademarks of Netscape Communications Corporation. Netscape content © 2010. Used with permission.

Electronic Mail, Instant Messaging, and Weblogs

default search engine but also provides links to other popular search tools including Ask Jeeves, Look Smart, Lycos, and Overture. Using search engines to find information has become an indispensable part of business and personal Internet, intranet, and extranet applications.

Industry experts predict the Web browser will be the model for how most people use networked computers in the future. Even today, whether you want to watch a video, make a phone call, download some software, hold a videoconference, check your e-mail, or work on a spreadsheet of your team's business plan, you can use your browser to launch and host such applications. That's why browsers are sometimes called the *universal client*, that is, the software component installed on all of the networked computing and communications devices of the clients (users) throughout an enterprise. As an aside, this entire book was revised and edited in a browser-based authoring program called PowerXEditor (we will learn more about PowerXEditor later in this chapter).

The first thing many people do at work, all over the world, is check their electronic mail. **E-mail** has changed the way people work and communicate. Millions of end users now depend on e-mail software to communicate with one another by sending and receiving electronic messages and file attachments via the Internet or their organizations' intranets or extranets. E-mail is stored on networked mail servers until you are ready. Whenever you want to, you can read your e-mail by displaying it on your workstation. So, with only a few minutes of effort (and a few microseconds of transmission time), a message to one or many individuals can be composed, sent, and received.

As we mentioned previously, e-mail software is now a mainstay component of top software suites and Web browsers. Free e-mail packages such as Microsoft HotMail, Yahoo! Mail, and Netscape WebMail are available to Internet users from online services and Internet service providers. Most e-mail software like Microsoft Outlook Express, Windows Mail, or Netscape Messenger can route messages to multiple end users based on predefined mailing lists and provide password security, automatic message forwarding, and remote user access. They also allow you to store messages in folders and make it easy to add documents and Web file attachments to e-mail messages. E-mail packages enable you to edit and send graphics and multimedia files, as well as text, and provide computer conferencing capabilities. In addition, your e-mail software may automatically filter and sort incoming messages (even news items from online services) and route them to appropriate user mailboxes and folders. Finally, many e-mail clients also include calendaring and contact management functions.

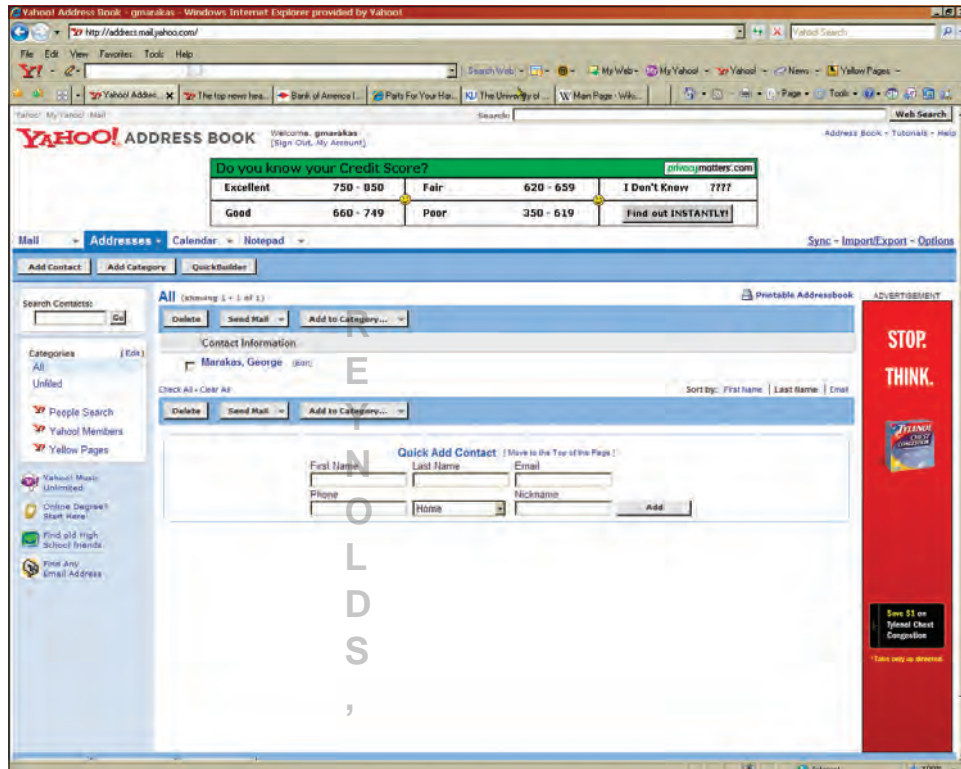
Instant messaging (IM) is an e-mail/computer-conferencing hybrid technology that has grown so rapidly that it has become a standard method of electronic messaging for millions of Internet users worldwide. By using instant messaging, groups of business professionals or friends and associates can send and receive electronic messages instantly and thus communicate and collaborate in real time in a near-conversational mode. Messages pop up instantly in an IM window on the computer screens of everyone in your business workgroup or friends on your IM "buddy list," as long as they are online, no matter what other tasks they are working on at that moment. Instant messaging software can be downloaded and IM services implemented by subscribing to many popular IM systems, including AOL's Instant Messenger and ICQ, MSN Messenger, and Yahoo Messenger. See Figure 4.6.

A **Weblog** (usually shortened to **blog** or written as "Web log" or "weblog") is a **Web site** of personal or noncommercial origin that uses a dated log format updated daily or very frequently with new information about a particular subject or range of subjects. The information can be written by the site owner, gleaned from other Web sites or other sources, or contributed by users via e-mail.

A Weblog often has the quality of being a kind of "log of our times" from a particular point of view. Generally, Weblogs are devoted to one or several subjects or themes, usually of topical interest. In general, Weblogs can be thought of as developing

FIGURE 4.6

Using the e-mail features of the Yahoo! instant messaging system.



Source: YAHOO! and the YAHOO! logo are trademarks of Yahoo! Inc. Reproduced with permission of Yahoo! Inc.

commentaries, individual or collective, on their particular themes. A Weblog may consist of the recorded ideas of an individual (a sort of diary) or be a complex collaboration open to anyone. Most of the latter are *moderated discussions*.

Because there are a number of variations on this idea and new variations can easily be invented, the meaning of this term is apt to gather additional connotations with time. As a formatting and content approach for a Web site, the Weblog seems popular because the viewer knows that something changes every day, there is a personal (rather than bland commercial) point of view, and, on some sites, there is an opportunity to collaborate with or respond to the Web site and its participants.

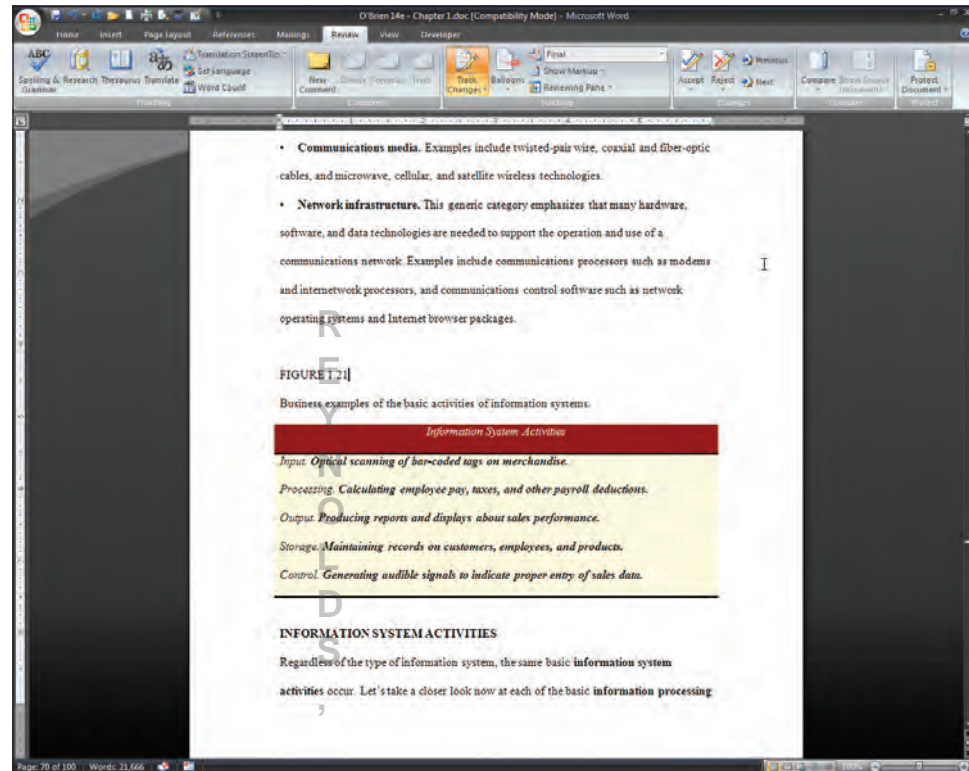
Word Processing and Desktop Publishing

Software for **word processing** has transformed the process of writing just about anything. Word processing packages computerize the creation, editing, revision, and printing of *documents* (e.g., letters, memos, reports) by electronically processing *text data* (words, phrases, sentences, and paragraphs). Top word processing packages like Microsoft Word, Lotus WordPro, Corel WordPerfect, and OpenOffice Writer can provide a wide variety of attractively printed documents with their desktop publishing capabilities. These packages can also convert documents to HTML format for publication as Web pages on corporate intranets or the World Wide Web.

Word processing packages also provide other helpful features. For example, a *spelling checker* capability can identify and correct spelling errors, and a *thesaurus* feature helps you find a better choice of words to express ideas. You can also identify and correct grammar and punctuation errors, as well as suggest possible improvements in your writing style, with grammar and style checker functions. In addition to converting documents to HTML format, you can use the top packages to design and create Web pages from scratch for an Internet or intranet Web site. See Figure 4.7.

FIGURE 4.7

Using the Microsoft Word word processing package. Note the insertion of a table in the document.



Source: Courtesy of Microsoft®.

End users and organizations can use **desktop publishing (DTP)** software to produce their own printed materials that look professionally published. That is, they can design and print their own newsletters, brochures, manuals, and books with several type styles, graphics, photos, and colors on each page. Word processing packages and desktop publishing packages like Adobe InDesign, Microsoft Publisher, and QuarkXPress are used for desktop publishing. Typically, text material and graphics can be generated by word processing and graphics packages and imported as text and graphics files. Optical scanners may be used to input text and graphics from printed material. You can also use files of *clip art*, which are predrawn graphic illustrations provided by the software package or available from other sources.

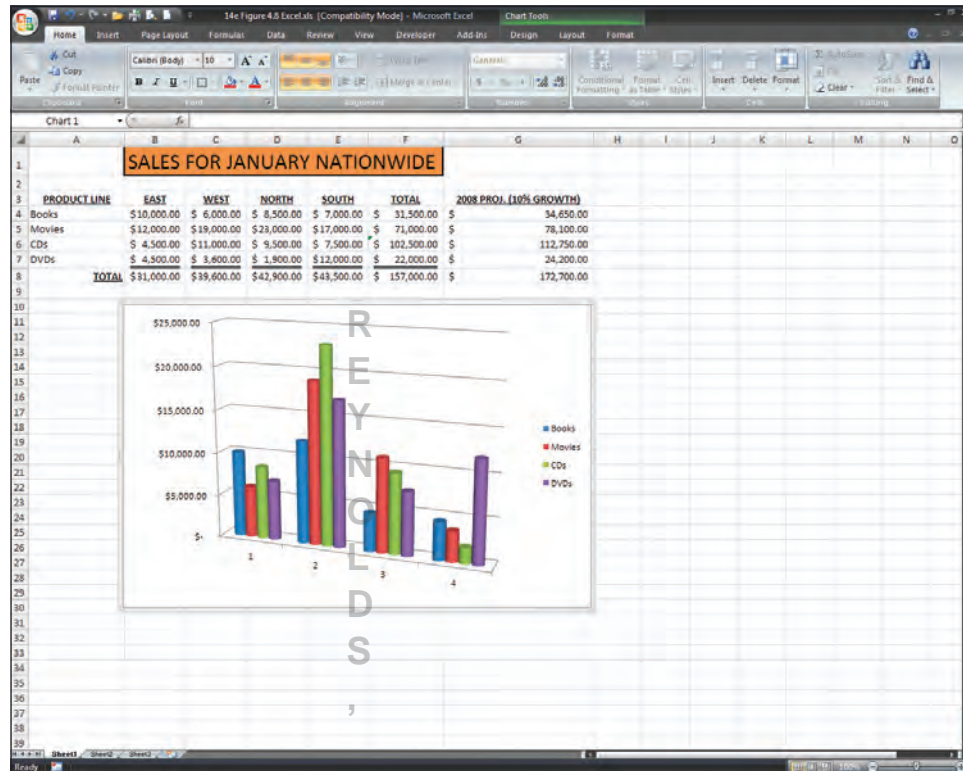
Electronic Spreadsheets

Spreadsheet packages like Lotus 1-2-3, Microsoft Excel, OpenOffice Calc, and Corel QuattroPro are used by virtually every business for analysis, planning, and modeling. They help you develop an *electronic spreadsheet*, which is a worksheet of rows and columns that can be stored on your PC or on a network server, or converted to HTML format and stored as a Web page or Web sheet on the World Wide Web. Developing a spreadsheet involves designing its format and developing the relationships (formulas) that will be used in the worksheet. In response to your input, the computer performs necessary calculations according to the formulas you defined in the spreadsheet and displays the results immediately, whether on your workstation or Web site. Most packages also help you develop charts and graphic displays of spreadsheet results. See Figure 4.8.

For example, you could develop a spreadsheet to record and analyze past and present advertising performance for a business. You could also develop hyperlinks to a similar Web sheet on your marketing team's intranet Web site. Now you have a decision support tool to help you answer *what-if questions* you may have about advertising. For example, "What would happen to market share if advertising expenses were to

FIGURE 4.8

Using an electronic spreadsheet package, Microsoft Excel. Note the use of graphics.



Source: Courtesy of Microsoft®.

increase by 10 percent?" To answer this question, you would simply change the advertising expense formula on the advertising performance worksheet you developed. The computer would recalculate the affected figures, producing new market share figures and graphics. You would then have better insight into the effect of advertising decisions on market share. Then you could share this insight with a note on the Web sheet on your team's intranet Web site.

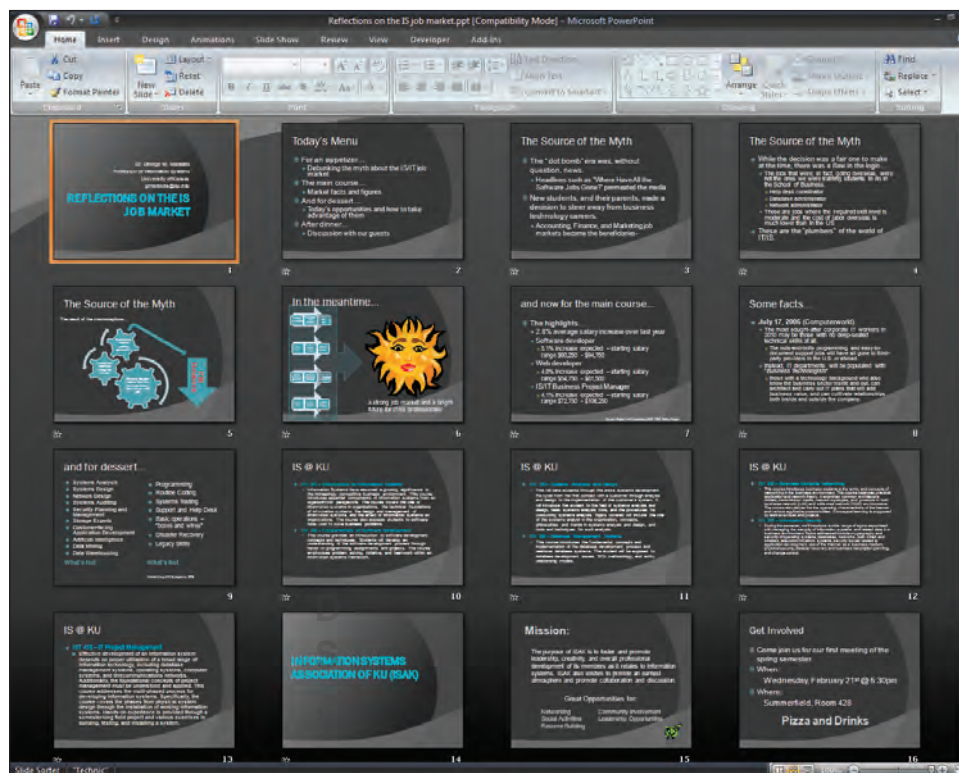
Presentation Graphics

Presentation graphics software packages help you convert numeric data into graphics displays such as line charts, bar graphs, pie charts, and many other types of graphics. Most of the top packages also help you prepare multimedia presentations of graphics, photos, animation, and video clips, including publishing to the World Wide Web. Not only are graphics and multimedia displays easier to comprehend and communicate than numeric data, but multiple-color and multiple-media displays can more easily emphasize key points, strategic differences, and important trends in the data. Presentation graphics have proved to be much more effective than tabular presentations of numeric data for reporting and communicating in advertising media, management reports, or other business presentations. See Figure 4.9.

Presentation graphics software packages like Microsoft PowerPoint, OpenOffice Impress, Lotus Freelance, or Corel Presentations give you many easy-to-use capabilities that encourage the use of graphics presentations. For example, most packages help you design and manage computer-generated and orchestrated *slide shows* containing many integrated graphics and multimedia displays. You can select from a variety of predesigned *templates* of business presentations, prepare and edit the outline and notes for a presentation, and manage the use of multimedia files of graphics, photos, sounds, and video clips. Of course, the top packages help you tailor your graphics and multimedia presentation for transfer in HTML format to Web sites on corporate intranets or the World Wide Web.

FIGURE 4.9

Using the slide preview feature of a presentation graphics package, Microsoft PowerPoint.



Source: Courtesy of Microsoft®.

Personal Information Managers

The **personal information manager (PIM)** is a popular software package for end-user productivity and collaboration, as well as a popular application for personal digital assistant (PDA) handheld devices. Various PIMs such as Lotus Organizer and Microsoft Outlook help end users store, organize, and retrieve information about customers, clients, and prospects or schedule and manage appointments, meetings, and tasks. A PIM package will organize data you enter and retrieve information in a variety of forms, depending on the style and structure of the PIM and the information you want. For example, information can be retrieved as an electronic calendar or list of appointments, meetings, or other things to do; as the timetable for a project; or as a display of key facts and financial data about customers, clients, or sales prospects. Most PIMs now include the ability to access the World Wide Web and provide e-mail capability. Also, some PIMs use Internet and e-mail features to support team collaboration by sharing information such as contact lists, task lists, and schedules with other networked PIM users. See Figure 4.10.

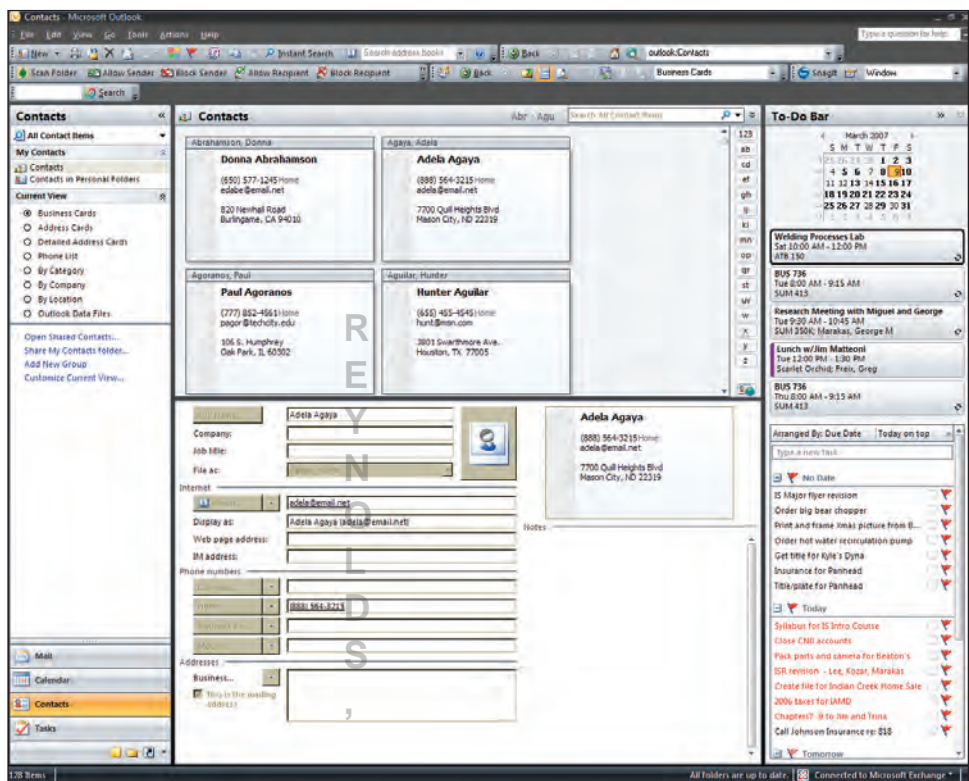
Groupware

Groupware is software that helps workgroups and teams collaborate to accomplish group assignments. Groupware is a category of general-purpose application software that combines a variety of software features and functions to facilitate collaboration. For example, groupware products like Lotus Notes, Novell GroupWise, and Microsoft Exchange support collaboration through e-mail, discussion groups and databases, scheduling, task management, data, audio and videoconferencing, and so on.

Groupware products rely on the Internet and corporate intranets and extranets to make collaboration possible on a global scale by *virtual teams* located anywhere in the world. For example, team members might use the Internet for global e-mail, project discussion forums, and joint Web page development. Or they might use corporate intranets to publish project news and progress reports and work jointly on documents stored on Web servers. See Figure 4.11.

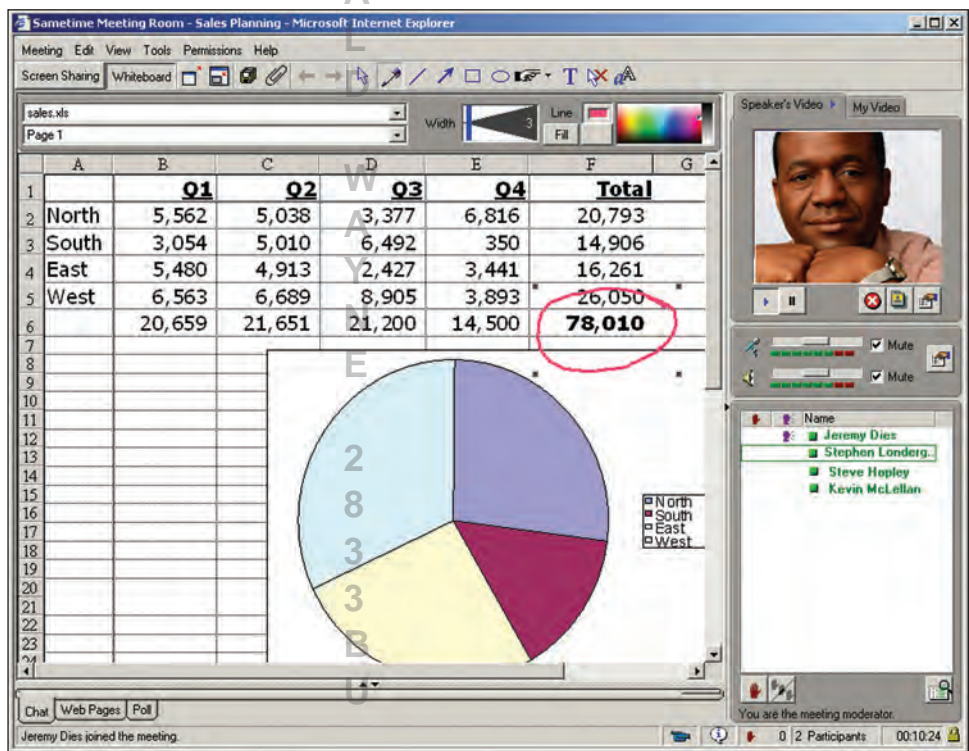
Collaborative capabilities are also being added to other software to give it groupware-like features. For example, in the Microsoft Office software suite, Microsoft Word

FIGURE 4.10
Using a personal
information manager
(PIM): Microsoft Outlook.



Source: Courtesy of Microsoft®.

FIGURE 4.11
Lotus Sametime enables
workgroups and project
teams to share spreadsheets
and other work documents
in an interactive online
collaboration process.



Source: Courtesy of International Business Machines Corporation. Unauthorized use not permitted.

keeps track of who made revisions to each document, Excel tracks all changes made to a spreadsheet, and Outlook lets you keep track of tasks you delegate to other team members. Recently, the Microsoft Office suite has included functions that allow multiple people to work on and edit the same document at the same time. Using this feature, any changes made by one team member will become visible to all team members as they are being made.

Two recent additions to the collaborative software marketplace are Microsoft's Windows SharePoint Services and IBM's WebSphere. Both products allow teams to create sophisticated Web sites for information sharing and document collaboration quickly. Furthermore, businesses can use these products as a platform for application development to facilitate the efficient creation of Web-based business portals and transaction processing applications. Web sites built with collaborative development tools can integrate a wide variety of individual applications that can help increase both individual and team productivity.

Software Alternatives

Many businesses are finding alternatives to acquiring, installing, and maintaining business application software purchased from software vendors or developing and maintaining their own software in-house with their own software developer employees. For example, as we will discuss further in Chapter 14, many large companies are *outsourcing* the development and maintenance of software they need to *contract programming* firms and other software development companies, including the use of *offshore* software developers in foreign countries, and employing the Internet to communicate, collaborate, and manage their software development projects.

Application Service Providers

A large and fast-growing number of companies are turning to **application service providers (ASPs)**, instead of developing or purchasing the application software they need to run their businesses. Application service providers are companies that own, operate, and maintain application software and the computer system resources (servers, system software, networks, and IT personnel) required to offer the use of the application software for a fee as a service over the Internet. The ASP can bill their customers on a per-use basis or on a monthly or annual fee basis.

Businesses are using an ASP instead of owning and maintaining their own software for many reasons. One of the biggest advantages is the low cost of initial investment, and in many cases, the short time needed to get the Web-based application set up and running. The ASP's pay-as-you-go fee structure is usually significantly less expensive than the cost of developing or purchasing, as well as running and maintaining, application software. In addition, using an ASP eliminates or drastically reduces the need for much of the IT infrastructure (servers, system software, and IT personnel) that would be needed to acquire and support application software, including the continual challenges of distributing and managing companywide software patches and upgrades. Consequently, the use of ASPs by businesses and other organizations is expected to accelerate in the coming years. See Figure 4.12.

McAfee Inc.: Security under a Software-as-a- Service Model

Santa Clara, California-based security vendor McAfee Inc. released a software-as-a-service Web security tool for protecting a distributed workforce from Web threats, while rendering IT departments fewer upfront costs in light of current budgetary constraints.

Especially in tough economic times, a SaaS model of software delivery, like the McAfee Web Protection Service, saves cash-strapped organizations money because IT staff don't have to spend valuable time managing on-site equipment, says Mark Campbell, senior product marketing manager with McAfee Inc. "They get the advantages of having a tool that is always on, always up-to-date and with uptime guarantees," says Campbell. One challenge with on-premise tools, he continues, is

that when vendors issue a feature update, a period of time can elapse before the enhancements are up and running in the environment, says Campbell. That problem goes away when the software is hosted centrally.

Features of the hosted security offering include reputation-based filtering, based on McAfee's reputation system TrustedSource, to block constantly morphing threats. There's flexible policy manager for setting policies for certain employee groups like access to certain social networking sites by contract employees versus executive management. Users have the ability to run reports and use dashboards to reap insight into an organization's Web usage. "Are employees spending all day on Facebook and does this align with our appropriate usage of Web tools?" says Campbell. Other features include malware protection, remote office and user support, and transparent user authentication.

James Quin, senior research analyst with London, Ontario-based Info-Tech Research Group Ltd., can't yet say if McAfee's SaaS offering will be cheaper in the long run given monthly recurring costs for the service. That said, in this climate of eliminated capital budgets, Quin says "a solution like this offers them value there." Small organizations in particular, says Quin, will benefit from not having to retain as much security expertise.

Offering a SaaS option for malware technology that is "pretty commoditized" is certainly a move by McAfee to differentiate itself in a crowded space, says Quin. "And it puts them out front first because they're not going to be the last ones to offer this kind of service," he says.

Campbell thinks customers' perception of hosted security products has changed for the better, helped along by the successful adoption of hosted CRM tools like salesforce.com. "More and more IT departments are beginning to accept and really realize the benefits of it," says Campbell.

Source: Adapted from Kathleen Lau, "SaaS Web Security a Cheaper Option, McAfee Says," *CIO.com*, April 30, 2009.

FIGURE 4.12

Salesforce.com is a leading application service provider of Web-based sales management and customer relationship management services to both large and small businesses.



Source: Courtesy of Salesforce.com.

Cloud Computing

One of the most recent advances in computing and software delivery is called **cloud computing**. Cloud computing is a style of computing in which software and, in some cases, virtualized hardware resources are provided as a service over the Internet. Users need not have knowledge of, expertise in, or control over the technology infrastructure “in the cloud” that supports them. The term cloud is used as a metaphor for the Internet, based on how the Internet is often depicted in computer network diagrams.

The concept incorporates technology trends that have the common theme of reliance on the Internet for satisfying the computing needs of the users. Examples of vendors providing cloud services include SAP Business ByDesign, MidlandHR’s “iTrent as a Service,” Salesforce.com, and Google Apps, which provide common business applications online that are accessed from a Web browser, while the software and data are stored on the servers.

Cloud computing is often confused with grid computing (recall the concept from Chapter 3 where the CPU power of multiple computers is harnessed to act like one big computer when necessary). Indeed, many cloud computing deployments depend on grids, but cloud computing can be seen as a natural next step from the grid model. The majority of cloud computing infrastructure consists of reliable services delivered through data centers and built on servers with different levels of virtualization technologies. The services are accessible anywhere that has access to networking infrastructure. The cloud appears as a single point of access for all the computing needs of consumers.

As many computer software users generally do not own the infrastructure around them, they can avoid capital expenditure and consume resources as a service, paying instead for what they use. If this sounds a lot like how you pay for your electricity or natural gas, it is because the same basic model has been adopted. Many cloud computing offerings have adopted the utility computing model, which is analogous to how traditional utilities like electricity are consumed, while others are billed on a subscription basis. Sharing “perishable and intangible” computing power among multiple users or enterprises can improve utilization rates, as servers are left idle less often because more people are accessing and using the computing resources. Through this approach, significant reductions in costs can be realized while increasing the overall speed of application development. A side effect of this approach is that a given user’s or enterprise’s computing capacity can be scaled upward almost instantly as needed without having to own an infrastructure that is engineered to be ready for short-term peak loads. Cloud computing has been enabled by large increases in available commercial bandwidth, which makes it possible to receive the same response times from centralized infrastructure at other sites.

The real benefit to the organization comes from the cost savings. Cloud computing users can avoid capital expenditure on hardware, software, and services, by simply paying a provider only for what they use. As stated above, consumption is billed on a utility (e.g. resources consumed, like electricity) or subscription (e.g., time based, like a newspaper) basis, with little or no upfront cost. Other benefits of this time-sharing style approach are low barriers to entry, shared infrastructure and costs, low management overhead, and immediate access to a broad range of applications. Users can generally terminate the contract at any time, and the services are often covered by service level agreements with financial penalties in the event the agreed-upon service levels are not delivered. It is predicted that someday, everyone will compute “in the cloud.”

Software Licensing

Regardless of whether a software application is purchased COTS or accessed via an ASP, the software must be licensed for use. Software licensing is a complex topic that involves considerations of the special characteristics of software in the context of the underlying intellectual property rights, including copyright, trademark, and trade secrets, as well as traditional contract law, including the Uniform Commercial Code (UCC).

Contrary to what many believe, when an individual or company buys a software application, they have not purchased rights of ownership. Rather, they have purchased a license to use the software under the terms of the software licensing agreement. Software is generally licensed to better protect the vendor's intellectual property rights. The license often prohibits reverse engineering, modifying, disclosing, or transferring the software. In most cases, the license also gives the purchaser permission to sell or dispose of the rights provided by the license but not to duplicate or resell multiple copies of the software.

The requirement for licensing does not disappear when use of the software is obtained through an ASP. In this case, the license to dispense use of the software is granted to the ASP by the various software vendors, and in return, the ASP agrees to pay the software vendor a royalty based on the number of user accounts to which the ASP resells the rights.

Software vendors are working hard to provide easy licensing and access to their products while simultaneously preventing software piracy, which serves only to raise the ultimate cost of the product.

In the next section, we will learn about an entirely new approach to software licensing: open-source code.

SECTION II

System Software: Computer System Management

System Software Overview

System software consists of programs that manage and support a computer system and its information processing activities. For example, operating systems and network management programs serve as a vital *software interface* between computer networks and hardware and the application programs of end users.

Read the Real World Case on the use of open-source software by the U.S. Department of Defense. We can learn a lot about new trends regarding the use of open-source applications from this example. See Figure 4.13.

Overview

We can group system software into two major categories (see Figure 4.14):

- **System Management Programs.** Programs that manage the hardware, software, network, and data resources of computer systems during the execution of the various information processing jobs of users. Examples of important system management programs are operating systems, network management programs, database management systems, and system utilities.
- **System Development Programs.** Programs that help users develop information system programs and procedures and prepare user programs for computer processing. Major software development programs are programming language translators and editors, and a variety of CASE (computer-aided software engineering) and other programming tools. We will take a closer look at CASE tools later in this chapter.

Operating Systems

The most important system software package for any computer is its operating system. An **operating system** is an integrated system of programs that manages the operations of the CPU, controls the input/output and storage resources and activities of the computer system, and provides various support services as the computer executes the application programs of users.

The primary purpose of an operating system is to maximize the productivity of a computer system by operating it in the most efficient manner. An operating system minimizes the amount of human intervention required during processing. It helps your application programs perform common operations such as accessing a network, entering data, saving and retrieving files, and printing or displaying output. If you have any hands-on experience with a computer, you know that the operating system must be loaded and activated before you can accomplish other tasks. This requirement emphasizes that operating systems are the most indispensable components of the software interface between users and the hardware of their computer systems.

Operating Systems Functions

An operating system performs five basic functions in the operation of a computer system: providing a user interface, resource management, task management, file management, and utilities and support services. See Figure 4.15.

The User Interface. The **user interface** is the part of the operating system that allows you to communicate with it so you can load programs, access files, and accomplish other tasks. Three main types of user interfaces are the *command-driven*, *menu-driven*, and *graphical user interfaces*. The trend in user interfaces for operating systems and other software is moving away from the entry of brief end-user commands, or even the selection of choices from menus of options. Instead, most software provides an easy-to-use graphical user interface (GUI) that uses icons, bars, buttons, boxes, and other images. These GUIs rely on pointing devices like the electronic mouse or touch pad to make selections that help you get things done. Currently, the most common and widely recognized GUI is the Microsoft Windows desktop.

REAL WORLD

CASE

2

U.S. Department of Defense:
Enlisting Open-Source Applications

The U.S. Defense Department is enlisting an open-source approach to software development, which is an about-face for such a historically top-down organization. The Department of Defense (DoD) says open-source software is equal to commercial software in almost all cases and by law should be considered by the agency when making technology purchase decisions.

In terms of guidance, the DoD says open-source software (OSS) meets the definition of “commercial computer software” and thus executive agencies are required to include open source when evaluating software that meets their computing needs. OSS is defined as “software for which the human-readable source code is available for use, study, reuse, modification, enhancement, and redistribution by the users of that software.”

In addition, it lays out a list of open-source positives, including broad peer-review that helps eliminate defects, modification rights that help speed changes when needed, a reduction in the reliance on proprietary vendors, a licensing model that facilitates quick provisioning, cost reduction in some cases, reduction in maintenance and ownership costs, and favorable characteristics for rapid prototyping and experimentation. “The continuous and broad peer-review enabled by publicly available source code supports software reliability and security efforts through the identification and elimination of defects that might otherwise go unrecognized by a more limited core development team,” states deputy CIO David Wennergren in a memo to top military officials.

“I would consider this a milestone day,” says John Scott, director of open-source software and open integration for

Mercury Federal Systems, a technology consultancy to the U.S. government. Scott helped draft some of the open-source guidance contained in the memo, which took about 18 months to draft. “The 2003 policy study was OK to use, but this one goes a bit further in expanding on what open source is and why you would want to use it. But it is not just about usage, it is also about helping create OSS by submitting changes back out to the public.”

Scott says he believes this is the first time guidance has been issued about sharing the government’s own open-source changes with the public.

Taken together, two developments show how the Defense Department is trying to take advantage of Web-based communities to speed up software development and reduce its costs. Dave Mihelcic, CTO of the Defense Information Systems Agency, says the military believes in the core Web 2.0 philosophy of the power of collaboration.

The military has launched a collaborative platform called Forge.mil for its developers to share software, systems components, and network services. The agency also signed an agreement with the Open Source Software Institute (OSSSI) to allow 50 internally developed workforce management applications to be licensed to other government agencies, universities, and companies.

“The Web is a platform for harvesting collective intelligence,” Mihelcic says. He points to “remixable data sources, services in perpetual beta and lightweight programming models” as some of the aspects of open-source software development that are applicable to the Defense Department.

One example of the Defense Department’s new community-based approach to software development is Forge.mil, which was made generally available for unclassified use within the department in April 2009. Forge.mil is powered by CollabNet Team Forge, a commercial lifecycle management platform for distributed software development teams, and is modeled after the popular SourceForge.net.

The Defense Information Systems Agency (DISA) has issued version two of SoftwareForge (software that runs on the Forge.mil site to enable sharing and collaborative development of open-source software) after a three-month trial that grew to 1,300 users. SoftwareForge provides software version control, bug tracking, requirements management, and release packaging for software developers, along with collaboration tools such as wikis, discussion forums, and document repositories, DISA says.

DISA also says it will deploy a cloud computing-based version of the SoftwareForge tools for classified environments. DISA also plans to add software testing and certification services to Forge.mil. Mihelcic says Forge.mil is similar to the “Web 2.0 paradigm of putting services on the Web and making them accessible to a large number of users to increase the adoption of capabilities. We’re using the same collaboration approach to speed the development of DOD systems.”

FIGURE 4.13



The U.S. Department of Defense is becoming both an adopter and a provider of open-source software.

Source: Andrea Comas/Reuters/Landov.

Meanwhile, DISA has licensed its Corporate Management Information System (CMIS) to the OSSSI to develop an open-source version of the 50-odd applications that DISA uses to manage its workforce. The CMIS applications support human resources, training, payroll, and other personnel management functions that meet federal regulations. DISA, which provides IT services to the Department of Defense, made the decision to share its applications after other agencies expressed interest in them, says Richard Nelson, chief of personnel systems support at DISA's manpower, personnel, and security directorate. "Federal agencies discovered that the applications we have could be of benefit more widely," he says. Interest is coming from states and counties, as well.

DISA worked with the nonprofit OSSSI, which promotes the use of open source in government and academia. OSSSI copyrighted the software stack and licensed it back to DISA, making it available at no cost to government agencies under the Open Software License 3.0. "It's already paid for because the taxpayer paid for us to build it," Nelson says.

OSSSI wanted to create a process that could be repeated with other government-built applications. "The opportunity was more than the product," executive director John Weatherby says. "One of the key things was to set up a system, a process that can be replicated by other government agencies."

CMIS comprises more than 50 Web applications, including workforce management, automated workflow, learning management, balanced scorecard, and telework management. CMIS has 16,000 users, including DISA employees and military contractors. Originally written in 1997, CMIS was revamped in January 2006 using the latest Web-based tools, including an Adobe Cold Fusion front-end and a Microsoft SQL Server 2005 back-end.

Nelson says CMIS is easy to use because it takes advantage of modern Web-based interfaces, including drop-down lists for data input. "We've been able to cut down on help desk support so substantially," Nelson says. "With the old version, we were running anywhere from 75 to 100 help desk calls and e-mails a day. Now our average is less than five

e-mails and calls. It's not because people are using it less but because it has fewer problems."

Nelson says a key driver for CMIS is that it needs to be so intuitive that users don't need training. "If the customer requires instruction on the product, we have failed and we will do it over," Nelson says. "The reason that we're able to do that so successfully is that we take a somewhat different approach to the way most software is designed. Most software is designed so that business logic and processes need to follow software logic and process. Therefore it requires substantial training. We do it exactly opposite."

The OSSSI will make CMIS available in two different licenses: a regular open-source license for government agencies and companies, and a free license for academia. Nelson says CMIS has a cutting-edge approach to learning management, handling everything from training course sign-up to approvals and payment. Another unusual feature of CMIS is its telework management application.

Nelson says he hopes many organizations will license CMIS and start adding new capabilities so DISA can take advantage of a vibrant CMIS community of developers. Within three years, "I would hope that a number of others inside government and beyond are using it," Nelson says. "I'm hoping we all have ready access to qualified developers. I'm hoping that DISA gets access to a substantial number of additional applications . . . without having to build them ourselves."

Going forward, DISA wants to encourage use of and training in Adobe Cold Fusion, which it used to build OSCMIS, to increase the talent pool of OSCMIS developers. "We would even like to start with kids in high school to get them interested in software development as a career," Nelson says.

Source: Adapted from Carolyn Duffy Marsan, "Military Enlists Open Source Community," *Network World*, April 27, 2009; John Fontana, "DoD: Open Source as Good as Proprietary Software," *Network World*, October 27, 2009; J. Nicholas Hoover, "Defense CIO Touts Benefits of Open Source," *InformationWeek*, October 28, 2009; and J. Nicholas Hoover, "Defense Info Agency Open-Sources Its Web Apps," *InformationWeek*, August 21, 2009.

CASE STUDY QUESTIONS

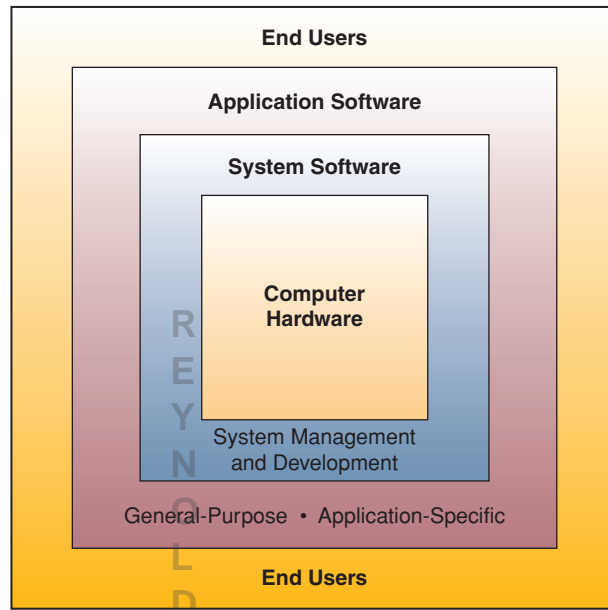
1. Given the critical nature of defense activities, security in this environment is a primary concern. How do the agencies discussed in the case address this issue? Can you think of anything else they could be doing? Provide some recommendations.
2. The U.S. Department of Defense is arguably one of the largest organizations in the world. Managing technology for such an organization is certainly a major endeavor. Does the shift toward open-source initiatives help in this regard? Does it hurt? Discuss the advantages and disadvantages of adopting open-source applications in large organizations.
3. After reading the case, do you think the shift to open-source software involved a major cultural change for the Department of Defense? Would you expect the same to be the case for large companies? Justify your answer.

REAL WORLD ACTIVITIES

1. Small open-source applications have been around for quite some time, but large-scale open-source systems have begun to emerge. Go online and search the Internet for examples of businesses adopting open-source technologies for major organizational systems. Prepare a presentation to highlight several examples from your research.
2. How does the open-source model of application development and distribution differ from the more common, proprietary approach? Do open-source applications present a legitimate threat to commercial software development, or will they remain niche applications? Break into small groups to discuss various reasons that the companies may or may not want to adopt open-source technologies.

FIGURE 4.14

The system and application software interface between end users and computer hardware.



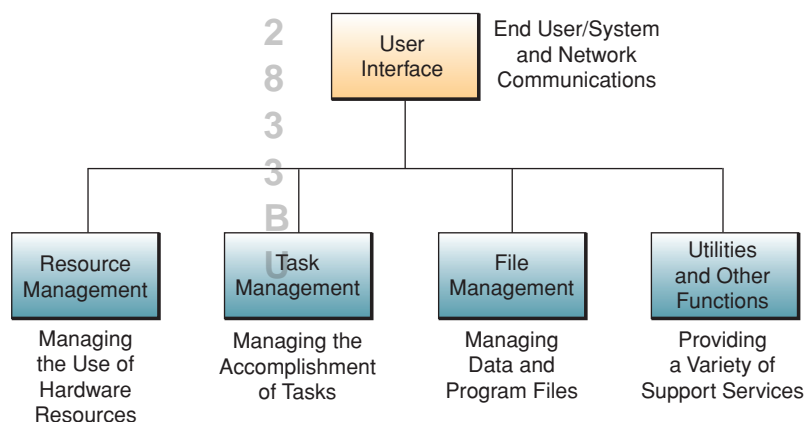
Resource Management. An operating system uses a variety of resource management programs to manage the hardware and networking resources of a computer system, including its CPU, memory, secondary storage devices, telecommunications processors, and input/output peripherals. For example, memory management programs keep track of where data and programs are stored. They may also subdivide memory into a number of sections and swap parts of programs and data between memory and magnetic disks or other secondary storage devices. This process can provide a computer system with a **virtual memory** capability that is significantly larger than the real memory capacity of its primary storage circuits. So, a computer with a virtual memory capability can process large programs and greater amounts of data than the capacity of its memory chips would normally allow.

File Management. An operating system contains **file management** programs that control the creation, deletion, and access of files of data and programs. File management also involves keeping track of the physical location of files on magnetic disks and other secondary storage devices. So operating systems maintain directories of information about the location and characteristics of files stored on a computer system's secondary storage devices.

Task Management. The **task management** programs of an operating system help accomplish the computing tasks of end users. The programs control which task gets

FIGURE 4.15

The basic functions of an operating system include a user interface, resource management, task management, file management, and utilities and other functions.



access to the CPU and for how much time. The task management functions can allocate a specific slice of CPU time to a particular task and interrupt the CPU at any time to substitute a higher priority task. Several different approaches to task management may be taken, each with advantages in certain situations.

Multitasking (sometimes referred to as *multiprogramming* or *time-sharing*) is a task management approach that allows for several computing tasks to be performed in a seemingly simultaneous fashion. In reality, multitasking assigns only one task at a time to the CPU, but it switches from one program to another so quickly that it gives the appearance of executing all of the programs at the same time. There are two basic types of multitasking: *preemptive* and *cooperative*. In preemptive multitasking, the task management functions parcel out CPU *time slices* to each program. In contrast, cooperative multitasking allows each program to control the CPU for as long as it needs it. If a program is not using the CPU, however, it can allow another program to use it temporarily. Most Windows and UNIX-based operating systems use the preemptive approach, whereas most Macintosh-style platforms use cooperative multitasking. Although the terms *multitasking* and *multiprocessing* are often used interchangeably, they are actually different concepts based on the number of CPUs being used. In multiprocessing, more than one CPU is being accessed, but in multitasking, only one CPU is in operation.

Most computers make use of some sort of multitasking. On modern microcomputers, multitasking is made possible by the development of powerful processors and their ability to address much larger memory capacities directly. This capability allows primary storage to be subdivided into several large partitions, each of which is used by a different software application.

In effect, a single computer can act as if it were several computers, or *virtual machines*, because each application program runs independently at the same time. The number of programs that can be run concurrently depends on the amount of memory that is available and the amount of processing each job demands. That's because a microprocessor (or CPU) can become overloaded with too many jobs and provide unacceptably slow response times. However, if memory and processing capacities are adequate, multitasking allows end users to switch easily from one application to another, share data files among applications, and process some applications in a *background* mode. Typically, background tasks include large printing jobs, extensive mathematical computations, or unattended telecommunications sessions.

Microsoft Windows

For many years, MS-DOS (Microsoft Disk Operating System) was the most widely used microcomputer operating system. It is a single-user, single-tasking operating system but was given a graphical user interface and limited multitasking capabilities by combining it with Microsoft **Windows**. Microsoft began replacing its DOS/Windows combination in 1995 with the Windows 95 operating system, featuring a graphical user interface, true multitasking, networking, multimedia, and many other capabilities. Microsoft introduced an enhanced Windows 98 version during 1998, and a Windows Me (Millennium Edition) consumer PC system in 2000.

Microsoft introduced its **Windows NT** (New Technology) operating system in 1995. Windows NT is a powerful, multitasking, multiuser operating system that was installed on many network servers to manage PCs with high-performance computing requirements. New Server and Workstation versions were introduced in 1997. Microsoft substantially enhanced its Windows NT products with the **Windows 2000** operating system during the year 2000.

Late in 2001, Microsoft introduced **Windows XP** Home Edition and Professional versions, and thus formally merged its two Windows operating system lines for consumer and business users, uniting them around the Windows NT and Windows 2000 code base. With Windows XP, consumers and home users finally received an enhanced Windows operating system with the performance and stability features that business users had in Windows 2000 and continue to have in Windows XP Professional. Microsoft also introduced four new **Windows Server 2003** versions in 2003, which are summarized and compared in Figure 4.16.

FIGURE 4.16 Comparing the purposes of the four versions of the Microsoft Windows Server 2008 operating system.

Microsoft Windows Server 2008 Comparisons	
• Windows Server 2008, Standard Edition	For smaller server applications, including file and print sharing, Internet and intranet connectivity, and centralized desktop application deployment.
• Windows Server 2008, Enterprise Edition	For larger business applications, XML Web services, enterprise collaboration, and enterprise network support.
• Windows Server 2008, Datacenter Edition	For business-critical and mission-critical applications demanding the highest levels of scalability and availability.
• Windows Server 2008, Web Edition	For Web serving and hosting, providing a platform for developing and deploying Web services and applications.

In 2006, Microsoft released their next operating system called Vista. Vista contains hundreds of new features; some of the most significant include an updated graphical user interface and visual style dubbed Windows Aero, improved search features, new multimedia creation tools such as Windows DVD Maker, and completely redesigned networking, audio, print, and display subsystems. Vista also aims to increase the level of communication between machines on a home network using peer-to-peer technology, making it easier to share files and digital media between computers and devices.

For developers, Vista introduced version 3.0 of the .NET Framework, which aims to make it significantly easier for developers to write high-quality applications than with the previous versions of Windows.

Microsoft's primary stated objective with Vista, however, was to improve the state of security in the Windows operating system. One of the most common criticisms of Windows XP and its predecessors has been their commonly exploited security vulnerabilities and overall susceptibility to malware, viruses, and buffer overflows. In light of these complaints, then-Microsoft chairman Bill Gates announced in early 2002 a companywide "Trustworthy Computing Initiative" to incorporate security work into every aspect of software development at the company. Microsoft claimed that it prioritized improving the security of Windows XP and Windows Server 2003 rather than finishing Windows Vista, significantly delaying its completion.

During 2008, a new server product, entitled (appropriately enough) Windows Server 2008, has emerged. Windows Server 2008 is built from the same code base as Windows Vista; therefore, it shares much of the same architecture and functionality. Because the code base is common, it automatically comes with most of the technical, security, management, and administrative features new to Windows Vista such as the rewritten networking processes (native IPv6, native wireless, speed, and security improvements); improved image-based installation, deployment, and recovery; improved diagnostics, monitoring, event logging, and reporting tools; new security features; improved Windows Firewall with secure default configuration; .NET Framework 3.0 technologies; and the core kernel, memory, and file system improvements. Processors and memory devices are modeled as Plug and Play devices, to allow hot-plugging of these devices.

Windows Server 2008 is already in release 2 as several performance and security enhancements required a major upgrade.

In 2009, Microsoft released their newest operating system, Windows 7. Unlike its predecessor, Vista, which introduced a large number of new features, Windows 7 was intended to be a more focused and incremental upgrade with the goal of being fully compatible with applications and hardware with which Vista was already compatible. Windows 7 has been very well received and is rapidly replacing the installed base of Vista without receiving any of the complaints and struggles encountered by Vista adopters and users.

UNIX

Originally developed by AT&T, **UNIX** now is also offered by other vendors, including Solaris by Sun Microsystems and AIX by IBM. UNIX is a multitasking, multiuser, network-managing operating system whose portability allows it to run on mainframes, midrange computers, and microcomputers. UNIX is still a popular choice for Web and other network servers.

Linux

Linux is a low-cost, powerful, and reliable UNIX-like operating system that is rapidly gaining market share from UNIX and Windows servers as a high-performance operating system for network servers and Web servers in both small and large networks. Linux was developed as free or low-cost *shareware* or *open-source software* over the Internet in the 1990s by Linus Torvald of Finland and millions of programmers around the world. Linux is still being enhanced in this way but is sold with extra features and support services by software vendors such as Red Hat, Caldera, and SUSE Linux. PC versions, which support office software suites, Web browsers, and other application software, are also available.

Open-Source Software

The concept of **open-source software** (OSS) is growing far beyond the Linux operating system. The basic idea behind open source is very simple: When programmers can read, redistribute, and modify the source code for a piece of software, the software evolves. People improve it, people adapt it, people fix bugs. This development can happen at a speed that, if one is accustomed to the slow pace of conventional software development, seems astonishing. The open-source community of software developers has learned that this rapid evolutionary process produces better software than the traditional commercial (closed) model, in which only a very few programmers can see the source. The concept of open source, admittedly, runs counter to the highly commercial (and proprietary) world of traditional software development. Nonetheless, an increasingly large number of developers have embraced the open-source concept and come to realize that the proprietary approach to software development has hidden costs that can often outweigh its benefits.

Since 1998, the OSS movement has become a revolution in software development. This revolution, however, can actually trace its roots back more than 30 years. Typically, in the PC era, computer software had been sold only as a finished product, otherwise called a *precompiled binary*, which is installed on a user's computer by copying files to appropriate directories or folders. Moving to a new computer platform (Windows to Macintosh, for example) usually required the purchase of a new license. If the company went out of business or discontinued support of a product, users of that product had no recourse. Bug fixes were completely dependent on the organization that sold the software. In contrast, OSS is software that is licensed to guarantee free access to the programming behind the precompiled binary, otherwise called the *source code*. This access allows the user to install the software on a new platform without an additional purchase and to get support (or create a support consortium with other like-minded users) for a product whose creator no longer supports it. Those who are technically inclined can fix bugs themselves rather than waiting for someone else to do so. Generally, there is a central distribution mechanism that allows users to obtain the source code, as well as precompiled binaries in some cases. There are also mechanisms by which users may pay a fee to obtain the software, such as on a CD-ROM or DVD, which may also include some technical support. A variety of licenses are used to ensure that the source code will remain available, wherever the code is actually used.

To be clear, there are several things open source is not: It is not shareware, public-domain software, freeware, or software viewers and readers made freely available without access to source code. Shareware, whether or not the user registers it and pays the registration fee, typically allows no access to the underlying source code. Unlike freeware and public-domain software, OSS is copyrighted and distributed with license terms designed to ensure that the source code will always be available. Although a fee may be charged for the software's packaging, distribution, or support, the complete package needed to create files is included, not simply a portion needed to view files created elsewhere.

The philosophy of open source is based on a variety of models that sometimes conflict; indeed, it often seems there are as many philosophies and models for developing and managing OSS as there are major products. In 1998, a small group of open-source enthusiasts decided it was time to formalize some things about open source. The newly formed group registered themselves on the Internet as www.open-source.org and began the process of defining exactly what is, and what is not, open-source software. As it stands today, open-source licensing is defined by the following characteristics:

- The license shall not restrict any party from selling or giving away the software as a component of an aggregate software distribution containing programs from several different sources.
- The program must include source code and must allow distribution in source code, as well as compiled form.
- The license must allow modifications and derived works and must allow them to be distributed under the same terms as the license of the original software.
- The license may restrict source code from being distributed in modified form only if the license allows the distribution of patch files with the source code for the purpose of modifying the program at build time.
- The license must not discriminate against any person or group of persons.
- The license must not restrict anyone from making use of the program in a specific field of endeavor.
- The rights attached to the program must apply to all to whom the program is re-distributed without the need for execution of an additional license by those parties.
- The license must not be specific to a product.
- The license must not contaminate other software by placing restrictions on any software distributed along with the licensed software.

This radical approach to software development and distribution is not without its detractors—most notably Microsoft. Nonetheless, the open-source movement is flourishing and stands to continue to revolutionize the way we think about software development.

OpenOffice.org 3

A relative newcomer to the open-source arena is an entire office suite offered by Sun Microsystems called OpenOffice.org 3. This product, built under the open-source standards described above, is a complete integrated office suite that provides all the common applications including word processing, spreadsheet, presentation graphics, and database management. It can store and retrieve files in a wide variety of data formats, including all of the file formats associated with the other major office suite applications on the market.

Best of all, OpenOffice.org 3 can be downloaded and used *entirely free of any license fees*. OpenOffice.org 3 is released under the LGPL license. This means you may use it for any purpose: domestic, commercial, educational, or public administration. You may install it on as many computers as you like, and you may make copies and give them away to family, friends, students, employees—anyone you like.

Mac OS X

Actually based on a form of UNIX, the **Mac OS X** (pronounced MAC OS 10) is the latest operating system from Apple for the iMac and other Macintosh microcomputers. The Mac OS X version 10.2 Jaguar has an advanced graphical user interface and multi-tasking and multimedia capabilities, along with an integrated Web browser, e-mail, instant messaging, search engine, digital media player, and many other features.

Mac OS X was a radical departure from previous Macintosh operating systems; its underlying code base is completely different from previous versions. Its core, named Darwin, is an open source, UNIX-like operating system. Apple layered over Darwin a

number of proprietary components, including the Aqua interface and the Finder, to complete the GUI-based operating system that is Mac OS X.

Mac OS X also included a number of features intended to make the operating system more stable and reliable than Apple's previous operating systems. Preemptive multitasking and memory protection, for example, improved the ability of the operating system to run multiple applications simultaneously that don't interrupt or corrupt each other.

The most visible change was the Aqua theme. The use of soft edges, translucent colors, and pinstripes—similar to the hardware design of the first iMacs—brought more texture and color to the interface than OS 9's "Platinum" appearance had offered. Numerous users of the older versions of the operating system decried the new look as "cutesy" and lacking in professional polish. However, Aqua also has been called a bold and innovative step forward at a time when user interfaces were seen as "dull and boring." Despite the controversy, the look was instantly recognizable, and even before the first version of Mac OS X was released, third-party developers started producing skins (look-and-feel colors and styles for application interfaces) for customizable applications that mimicked the Aqua appearance.

Mac OS X also includes its own software development tools, most prominently an integrated development environment called Xcode. Xcode provides interfaces to compilers that support several programming languages including C, C++, Objective-C, and Java. For the Apple Intel Transition, it was modified so that developers could easily create an operating system to remain compatible with both the Intel-based and PowerPC-based Macintosh.

Application Virtualization

Consider all of the various types of software applications we discussed in the first section of this chapter along with the multiple operating systems we just discussed. What happens when a user who has a machine running Windows needs to run an application designed specifically for a machine running Mac OS X? The answer used to be "Borrow someone's Mac." Through the development of application virtualization, a much more useful and productive answer exists. *Application virtualization* is an umbrella term that describes software technologies that improve portability, manageability, and compatibility of applications by insulating them from the underlying operating system on which they are executed. A fully virtualized application is not installed in the traditional sense; it is just executed as if it is. The application is fooled into believing that it is directly interfacing with the original operating system and all the resources managed by it, when in reality it is not. Application virtualization is just an extension of operating system virtualization where the same basic concepts fool the whole operating system into thinking it is running on a particular type of hardware when it is, in fact, not.

The concept of virtualization is not a recent development. The use of a virtual machine was a common practice during the mainframe era where extremely large machines were partitioned into smaller, separate virtual machines or domains to allow multiple users to run unique sets of applications and processes simultaneously. Each user constituency used a portion of the total available machine resources and the virtualization approach made it appear that each domain was an entirely separate machine from all the rest. If you have ever set up a new PC and created a partition on the hard drive, you have taken advantage of virtualization. You have taken one physical drive and created two virtual drives—one for each partition.

Application virtualization is a logical next step from these early roots. The benefits to the enterprise range from the cost savings associated with not having to have multiple platforms for multiple applications, to the energy savings associated with not having a multitude of servers running at low capacity while eating up electricity and generating heat.

A thorough discussion of virtualization is well beyond the scope of this text but suffice to say it is rapidly blurring the boundaries between machines and operating systems and operating systems and applications. Add this to the cloud computing concept and we have the makings of an anytime, anywhere, any machine, any application world.

Toronto's Hospital for Sick Children: Challenges in Making Virtualization Work



Toronto's Hospital for Sick Children has learned the hard way that virtualization efforts won't be successful if vendors aren't ready to support you, according to its director of technology, Ana Andreasian. The hospital (usually referred to as "Sick Kids") has already consolidated a considerable amount of its server infrastructure, which now includes 300 physical and 60 virtual machines. Sick Kids employs about 110 IT staff members who serve more than 5,000 employees.

Andreasian said the biggest issue she's experienced so far has come from vendors who do not properly test their applications before offering them to virtualization customers. "They'll say, 'Give me one CPU, one gig of memory, and I'm good,'" she says. "Then you'll find they need four CPUs and four gigs of RAM. You wind up having a never-ending discussion on how to solve the performance problems."

Another challenge has been vendors who say they're willing to support virtual environments, but not fully. "Some vendors have a condition: if you have a problem, you have to move (the application) out of a virtual environment," she says. "That's just not practical."

Sick Kids Hospital is somewhat unusual in that it started its virtualization journey by focusing on storage systems rather than servers. Andreasian explained that the organization currently manages some 150 terabytes of data, which is always on the increase. Devices to handle that data, meanwhile, always end up going out of support. "We were facing the question: How do you migrate that data? It's a huge cost," she says, adding that no one wants to experience any downtime associated with such a migration. And all this has to happen in such a way that's transparent to the user.

The hospital has also turned to Citrix for application virtualization in order to allow remote support, which is important in a hospital situation where many clinicians may need to work from home. Sick Kids is now using VMware to deal with the more common issues around managing server fleets, such as lack of real estate, power costs, and the need to provision (that is, set up) machines more quickly.

"In the physical world, if you have good planning and processes in place, that will help you with virtualization," says Dennis Corning, HP's worldwide senior manager of product marketing for virtualization.

Andreasian agrees. "Provisioning (a virtual server) is easy. De-provisioning once the business user no longer needs it is where it's difficult," she says. "They might not tell you it's no longer necessary. You need governance and monitoring and process."

Source: Adapted from Shane Schick, "Hospital CTO Identifies Virtualization Gotchas," *CIO.com*, January 28, 2010.

Other System Management Programs

There are many other types of important system management software besides operating systems. These include *database management systems*, which we will cover in Chapter 5, and *network management programs*, which we will cover in Chapter 6. Figure 4.17 compares several types of system software offered by IBM and its competitors.

Several other types of system management software are marketed as separate programs or included as part of an operating system. Utility programs, or **utilities**, are an important example. Programs like Norton Utilities perform miscellaneous housekeeping and file conversion functions. Examples include data backup, data recovery, virus protection, data compression, and file defragmentation. Most operating systems also provide many utilities that perform a variety of helpful chores for computer users.

FIGURE 4.17 Comparing system software offered by IBM and its main competitors.

Software Category	What It Does	IBM Product	Customers	Main Competitor	Customers
Network management	Monitors networks to keep them up and running.	Tivoli	T. Rowe Price uses it to safeguard customer records.	HP OpenView	Amazon.com uses it to monitor its servers.
Application server	Shuttles data between business apps and the Web.	WebSphere	REI uses it to serve up its Web site and distribute data.	BEA WebLogic	Washingtonpost.com builds news pages with it.
Database manager	Provides digital storehouses for business data.	DB2	Mikasa uses it to help customers find its products online.	Oracle 11g	It runs Southwest Airlines' frequent-flyer program.
Collaboration tools	Powers everything from e-mail to electronic calendars.	Lotus	Retailer Sephora uses it to coordinate store maintenance.	Microsoft Exchange	Time Inc. uses it to provide e-mail to its employees.
Development tools	Allows programmers to craft software code quickly.	Rational	Merrill Lynch used it to build code for online trading.	Microsoft Visual Studio .NET	Used to develop management system.

Other examples of system support programs include performance monitors and security monitors. **Performance monitors** are programs that monitor and adjust the performance and usage of one or more computer systems to keep them running efficiently. **Security monitors** are packages that monitor and control the use of computer systems and provide warning messages and record evidence of unauthorized use of computer resources. A recent trend is to merge both types of programs into operating systems like Microsoft's Windows 2008 Datacenter Server or into system management software like Computer Associates' CA-Unicenter, which can manage both mainframe systems and servers in a data center.

Another important software trend is the use of system software known as **application servers**, which provide a *middleware* interface between an operating system and the application programs of users. **Middleware** is software that helps diverse software applications and networked computer systems exchange data and work together more efficiently. Examples include application servers, Web servers, and enterprise application integration (EAI) software. Thus, for example, application servers like BEA's WebLogic and IBM's WebSphere help Web-based e-business and e-commerce applications run much faster and more efficiently on computers using Windows, UNIX, and other operating systems.

Programming Languages

To understand computer software, you need a basic knowledge of the role that programming languages play in the development of computer programs. A **programming language** allows a programmer to develop the sets of instructions that constitute a computer program. Many different programming languages have been developed, each with its own unique vocabulary, grammar, and uses.

Machine Languages

Machine languages (or *first-generation languages*) are the most basic level of programming languages. In the early stages of computer development, all program instructions had to be written using binary codes unique to each computer. This type of programming involves the difficult task of writing instructions in the form of strings of binary digits (ones and zeros) or other number systems. Programmers must have a detailed knowledge of the internal operations of the specific type of CPU they are using. They must write long series of detailed instructions to accomplish even simple processing tasks. Programming in machine language requires specifying the storage

FIGURE 4.18

Examples of four levels of programming languages. These programming language instructions might be used to compute the sum of two numbers as expressed by the formula $X = Y + Z$.

Four Levels of Programming Languages	
<ul style="list-style-type: none">• Machine Languages: Use binary coded instructions 1010 11001 1011 11010 1100 11011	<ul style="list-style-type: none">• High-Level Languages: Use brief statements or arithmetic notations BASIC: $X = Y + Z$ COBOL: COMPUTE $X = Y + Z$
<ul style="list-style-type: none">• Assembler Languages: Use symbolic coded instructions LOD Y ADD Z STR X	<ul style="list-style-type: none">• Fourth-Generation Languages: Use natural and nonprocedural statements SUM THE FOLLOWING NUMBERS

locations for every instruction and item of data used. Instructions must be included for every switch and indicator used by the program. These requirements make machine language programming a difficult and error-prone task. A machine language program to add two numbers together in the CPU of a specific computer and store the result might take the form shown in Figure 4.18.

Assembler Languages

Assembler languages (or *second-generation languages*) are the next level of programming languages. They were developed to reduce the difficulties in writing machine language programs. The use of assembler languages requires language translator programs called *assemblers* that allow a computer to convert the instructions of such language into machine instructions. Assembler languages are frequently called symbolic languages because symbols are used to represent operation codes and storage locations. Convenient alphabetic abbreviations called *mnemonics* (memory aids) and other symbols represent operation codes, storage locations, and data elements. For example, the computation $X = Y + Z$ in an assembler language might take the form shown in Figure 4.18.

Assembler languages are still used as a method of programming a computer in a machine-oriented language. Most computer manufacturers provide an assembler language that reflects the unique machine language instruction set of a particular line of computers. This feature is particularly desirable to *system programmers*, who program system software (as opposed to application programmers, who program application software), because it provides them with greater control and flexibility in designing a program for a particular computer. They can then produce more efficient software—that is, programs that require a minimum of instructions, storage, and CPU time to perform a specific processing assignment.

High-Level Languages

High-level languages (or *third-generation languages*) use instructions, which are called *statements*, that include brief statements or arithmetic expressions. Individual high-level language statements are actually *macroinstructions*; that is, each individual statement generates several machine instructions when translated into machine language by high-level language translator programs called *compilers* or *interpreters*. High-level language statements resemble the phrases or mathematical expressions required to express the problem or procedure being programmed. The *syntax* (vocabulary, punctuation, and grammatical rules) and *semantics* (meanings) of such statements do not reflect the internal code of any particular computer. For example, the computation $X = Y + Z$ would be programmed in the high-level languages of BASIC and COBOL as shown in Figure 4.18.

High-level languages like BASIC, COBOL, and FORTRAN are easier to learn and program than an assembler language because they have less rigid rules, forms, and syntaxes. However, high-level language programs are usually less efficient than assembler language programs and require a greater amount of computer time for translation into machine instructions. Because most high-level languages are machine-independent,

programs written in a high-level language do not have to be reprogrammed when a new computer is installed, and programmers do not have to learn a different language for each type of computer.

Fourth-Generation Languages

The term **fourth-generation language** describes a variety of programming languages that are more nonprocedural and *conversational* than prior languages. These languages are called fourth-generation languages (4GLs) to differentiate them from machine languages (first generation), assembler languages (second generation), and high-level languages (third generation).

Most fourth-generation languages are *nonprocedural languages* that encourage users and programmers to specify the results they want, while the computer determines the sequence of instructions that will accomplish those results. Thus, fourth-generation languages have helped simplify the programming process. **Natural languages** are sometimes considered *fifth-generation languages* (5GLs) and are very close to English or other human languages. Research and development activity in artificial intelligence (AI) is developing programming languages that are as easy to use as ordinary conversation in one's native tongue. For example, INTELLECT, a natural language, would use a statement like, "What are the average exam scores in MIS 200?" to program a simple average exam score task.

In the early days of 4GLs, results suggested that high-volume transaction processing environments were not in the range of a 4GL's capabilities. Although 4GLs were characterized by their ease of use, they were also viewed as less flexible than their predecessors, primarily due to their increased storage and processing speed requirements. In today's large data volume environment, 4GLs are widely used and no longer viewed as a trade-off between ease of use and flexibility.

Modern (and Automatic?) Code Generation

Twenty years ago, software engineer Fred Brooks famously observed that there was no silver bullet that could slay "the monster of missed schedules, blown budgets and flawed products." Today, the creation of software might seem as expensive, trouble-prone, and difficult as ever—and yet progress is being made. Although no silver bullet is in sight, an array of new techniques promises to further boost a programmer's productivity, at least in some application domains.

The techniques span a broad spectrum of methods and results, but all are aimed at generating software automatically. Typically, they generate code from high-level, machine-readable designs or from domain-specific languages—assisted by advanced compilers—that sometimes can be used by nonprogrammers.

Gordon Novak, a computer science professor at the University of Texas at Austin and a member of the school's Laboratory for Artificial Intelligence, is working on "automatic programming"—using libraries of generic versions of programs, such as algorithms—to sort or find items in a list. Unlike traditional subroutines, which have simple but rigid interfaces and are invoked by other lines of program code, his technique works at a higher level and is therefore more flexible and easier to use.

Novak's users construct "views" that describe application data and principles and then connect the views by arrows in diagrams that show the relationships among the data. The diagrams are, in essence, very high-level flowcharts of the desired program. They get compiled in a way that customizes the stored generic algorithms for the user's specific problem, and the result is ordinary source code such as C, C++, or Java.

Novak says he was able to generate 250 lines of source code for an indexing program in 90 seconds with his system. That's equivalent to a week of productivity for an average programmer using a traditional language. "You are describing your program at a higher level," he says. "And what my program is saying is, 'I can tailor the algorithm for your application for free.'"

Douglas Smith, principal scientist at Kestrel Institute, a nonprofit computer science research firm in Palo Alto, California, is developing tools to “automate knowledge and get it into the computer.” A programmer starts with Kestrel’s Specware, which is a general-purpose, fifth-generation language that specifies a program’s functions without regard to the ultimate programming language, system architecture, algorithms, data structures, and so on. Specware draws on a library of components, but the components aren’t code. They are at a higher level and include design knowledge and principles about algorithms, data structures, and so on. Smith calls them “abstract templates.”

In addition, Specware can produce proofs that the working code is “correct”—that is, that it conforms to the requirements put in by the user (which, of course, may contain errors). “Some customers want that for very-high-assurance applications, with no security flaws,” Smith says. Kestrel does work for NASA and U.S. military and security agencies.

“It’s a language for writing down problem requirements, a high-level statement of what a solution should be, without saying how to solve the problem,” Smith says. “We think it’s the ultimate frontier in software engineering. It’s what systems analysts do.”

Source: Adapted from Gary Anthes, “In the Labs: Automatic Code Generators,” *Computerworld*, March 20, 2006.

Object-Oriented Languages

Object-oriented languages like Visual Basic, C++, and Java are also considered fifth-generation languages and have become major tools of software development. Briefly, whereas most programming languages separate data elements from the procedures or actions that will be performed on them, object-oriented languages tie them together into **objects**. Thus, an object consists of data and the actions that can be performed on the data. For example, an object could be a set of data about a bank customer’s savings account and the operations (e.g., interest calculations) that might be performed on the data. An object also could be data in graphic form, such as a video display window plus the display actions that might be used on it. See Figure 4.19.

In procedural languages, a program consists of procedures to perform actions on each data element. However, in object-oriented systems, objects tell other objects to perform actions on themselves. For example, to open a window on a computer video display, a beginning menu object could send a window object a message to open, and

FIGURE 4.19

An example of a bank savings account object. This object consists of data about a customer’s account balance and the basic operations that can be performed on those data.


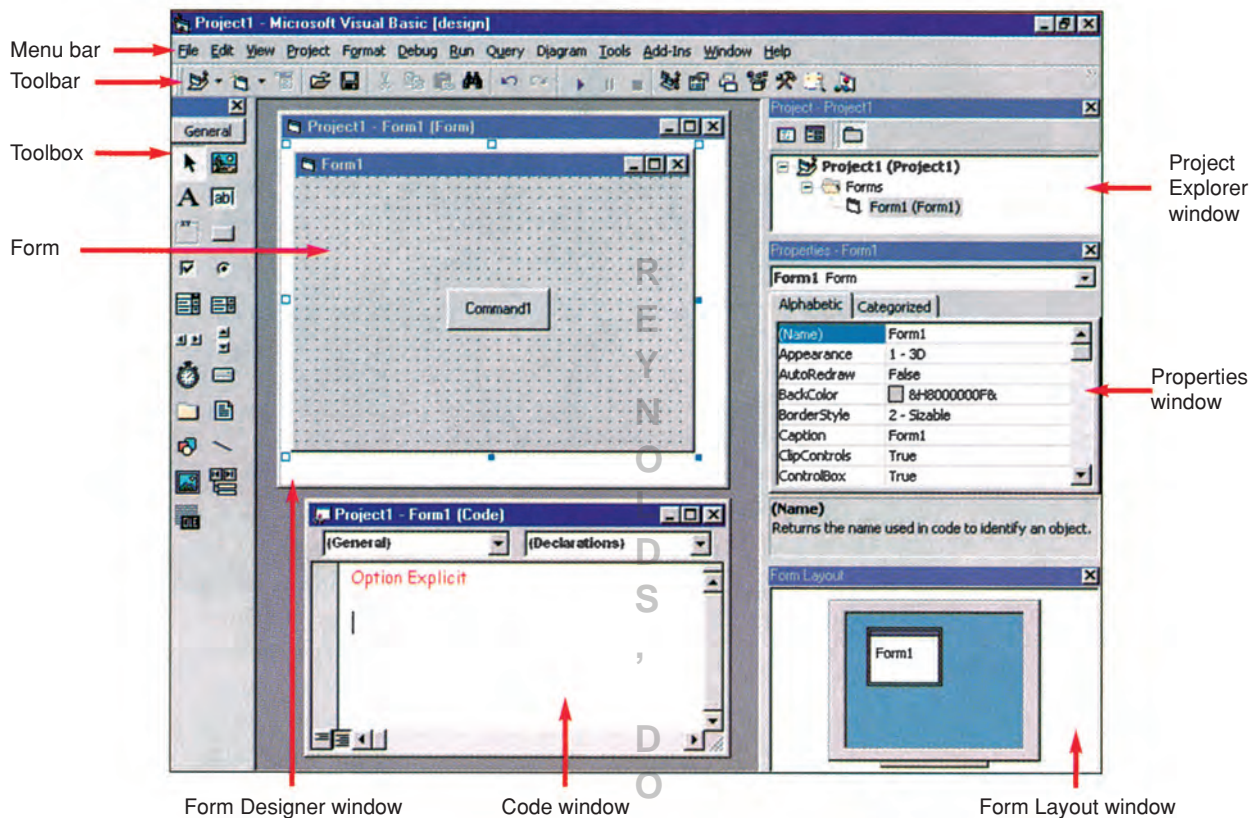


FIGURE 4.20 The Visual Basic object-oriented programming environment.



Source: Courtesy of Microsoft®.

a window would appear on the screen. That's because the window object contains the program code for opening itself.

Object-oriented languages are easier to use and more efficient for programming the graphics-oriented user interfaces required by many applications. Therefore, they are the most widely used programming languages for software development today. Also, once objects are programmed, they are reusable. Therefore, reusability of objects is a major benefit of object-oriented programming. For example, programmers can construct a user interface for a new program by assembling standard objects such as windows, bars, boxes, buttons, and icons. Therefore, most object-oriented programming packages provide a GUI that supports a point-and-click, drag-and-drop visual assembly of objects known as *visual programming*. Figure 4.20 shows a display of the Visual Basic object-oriented programming environment. Object-oriented technology is discussed further in the coverage of object-oriented databases in Chapter 5.

Web Languages and Services

HTML

HTML, XML, and Java are three programming languages that are important tools for building multimedia Web pages, Web sites, and Web-based applications. In addition, XML and Java have become strategic components of the software technologies that support many Web services initiatives in business.

HTML (Hypertext Markup Language) is a page description language that creates hypertext or hypermedia documents. HTML inserts control codes within a document at points you can specify that create links (*hyperlinks*) to other parts of the document or to other documents anywhere on the World Wide Web. HTML embeds control codes in the ASCII text of a document that designate titles, headings, graphics, and multimedia components, as well as hyperlinks within the document.

As we mentioned previously, several of the programs in the top software suites automatically convert documents into HTML formats. These include Web browsers, word processing and spreadsheet programs, database managers, and presentation graphics packages. These and other specialized *Web publishing* programs like Microsoft FrontPage, Lotus FastSite, and Macromedia's DreamWeaver provide a range of features to help you design and create multimedia Web pages without formal HTML programming.

XML

XML (eXtensible Markup Language) is not a Web page format description language like HTML. Instead, XML describes the contents of Web pages (including business documents designed for use on the Web) by applying identifying tags or *contextual labels* to the data in Web documents. For example, a travel agency Web page with airline names and flight times would use hidden XML tags like "airline name" and "flight time" to categorize each of the airline flight times on that page. Or product inventory data available at a Web site could be labeled with tags like "brand," "price," and "size." By classifying data in this way, XML makes Web site information much more searchable, easier to sort, and easier to analyze.

For example, XML-enabled search software could easily find the exact product you specify if the product data on the Web site had been labeled with identifying XML tags. A Web site that uses XML could also more easily determine which Web page features its customers use and which products they investigate. Thus, XML promises to make electronic business and commerce processes a lot easier and more efficient by supporting the automatic electronic exchange of business data between companies and their customers, suppliers, and other business partners.

As mentioned at the beginning of the chapter, this entire textbook was revised and edited for the current edition using an XML-based application called PowerXEditor by Aptara. Let's focus our attention on this unique application of XML intended to create efficiencies in the publishing industry.

Aptara, Inc.: Revolutionizing the Publishing Industry through XML

The publishing industry has experienced an upheaval in the past decade or so. The "long tail" of sales of existing books via Web sellers such as Amazon and the improvement in software and hardware technologies that can replicate the experience of reading a book or magazine means publishing houses are printing and selling fewer new books. As a result, many of these companies are venturing into digital publishing.

"All the publishers are shifting from print to digital," said Dev Ganesan, president and CEO of Aptara, which specializes in content transformation. "That's a huge change. What that means for software companies is that they need to develop platforms for content creation that meet the needs of every customer. At the same time, customers are looking at publishing in terms of handling content in terms of authors, editors, and production employees. On top of that, they're trying to automate parts of the production process. And companies must be willing to market products using traditional and new media to reach the widest possible audience. So there are a lot of challenges, but a lot of opportunities, too."

The upshot of all this is that learning professionals now can deliver content more flexibly and at a lower cost. They can make static content dynamic by taking a body of knowledge in print—such as a book—and converting it to a digital format. They can then chunk that content into smaller sizes and organize those nuggets of information according to learners' needs. Moreover, they can get content published and distributed much more quickly via digital, online media. This is critical in an industry such as health care, which faces rapid changes due to technological innovation and regulation, said another Aptara source.

"In addition to the cost savings, they want to turn it around much faster," he said. "Time to market is becoming paramount because there's so much innovation going on. If they don't have their print products out faster, they fall behind."

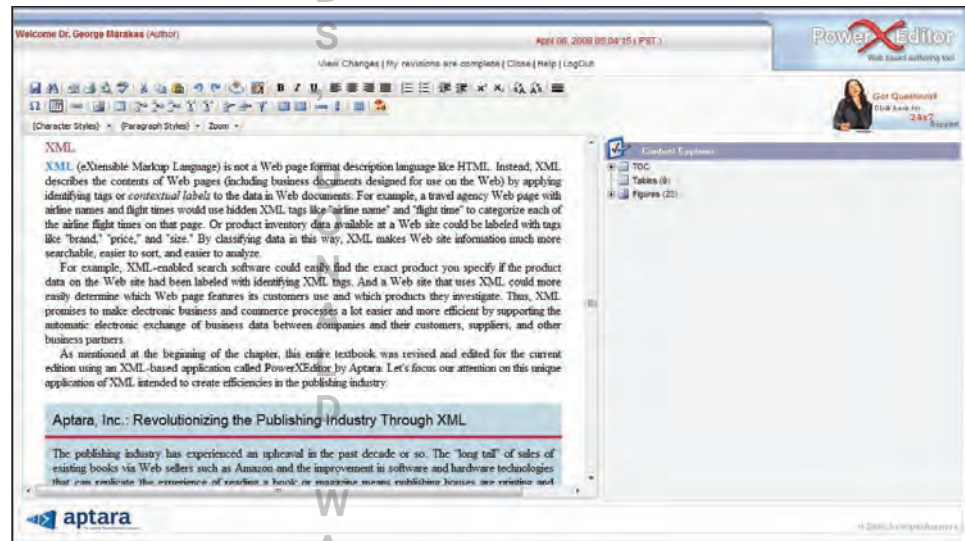
A breakthrough product from Aptara is called PowerXEditor (PXE). An XML-based application, PXE allows a publisher to upload an existing book layout; edit or revise all elements of the book, including text look and feel, figures, tables, and other elements unique to that book; and output the book to a paging program that sets the book up for final printing. The important issue is that all of this is done in a digital format instead of the previously common method of tear pages and cut and paste of figures and tables. Because the PXE content is XML-based, the application can be accessed via the Internet using any conventional Web browser. This means all of the contributors to a textbook can have access to the various chapters and elements no matter where they are. Add in the workflow management aspects of PXE, and all phases of the textbook revising, copyediting, and proofing processes can be handled with ease.

Figure 4.21 shows a typical PXE screen. You might notice that it is in the process of editing the page you are currently reading. Figure 4.22 shows the XML code for the same page.

Source: Adapted from Brian Summerfield, "Executive Briefings: Balancing Print and Digital Media," *Chief Learning Officer*, March 2008. <http://www.clomedia.com/includes/printcontent.php?aid=2133>

FIGURE 4.21

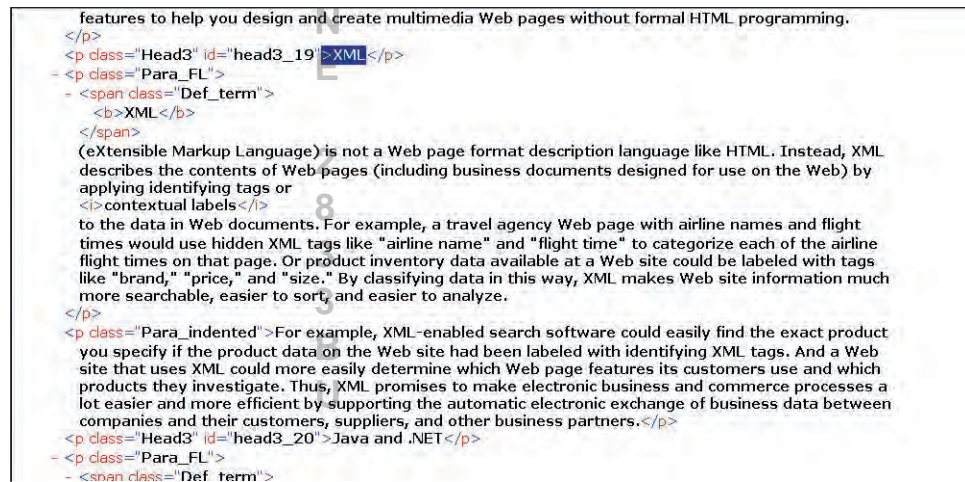
The XML-based PowerXEditor allows all the collaborators on a book project to access the elements of the book via a common Web browser. Here is a screenshot of PXE on the page you are currently reading.



Source: Courtesy of Aptara.

FIGURE 4.22

This is a section of the XML code from the page you are currently reading. While XML looks similar to HTML source code, it is far more powerful and complex.



Source: Courtesy of Aptara.

Java and .NET

Java is an object-oriented programming language created by Sun Microsystems that is revolutionizing the programming of applications for the World Wide Web and corporate intranets and extranets. Java is related to the C++ and Objective C programming languages but is much simpler and more secure and is computing-platform independent. Java is also specifically designed for real-time, interactive, Web-based network applications. Java applications consisting of small application programs, called *applets*, can be executed by any computer and any operating system anywhere in a network.

The ease of creating Java applets and distributing them from network servers to client PCs and network computers is one of the major reasons for Java's popularity. Applets can be small, special-purpose application programs or small modules of larger Java application programs. Java programs are platform-independent, too—they can run on Windows, UNIX, and Macintosh systems without modification.

Microsoft's **.NET** is a collection of programming support for what are known as Web services, the ability to use the Web rather than your own computer for various services (see Figure 4.23). .NET is intended to provide individual and business users with a seamlessly interoperable and Web-enabled interface for applications and computing devices and to make computing activities increasingly Web browser-oriented. The .NET platform includes servers, building-block services such as Web-based data storage, and device software. It also includes Passport, Microsoft's fill-in-the-form-only-once identity verification service.

The .NET platform is expected to enable the entire range of computing devices to work together and have user information automatically updated and synchronized on all of them. In addition, it will provide a premium online subscription service. The service will feature customized access to and delivery of products and services from a central starting point for the management of various applications (e.g., e-mail) or software (e.g., Office .NET). For developers, .NET offers the ability to create reusable modules, which should increase productivity and reduce the number of programming errors.

FIGURE 4.23 The benefits and limitations of the Java Enterprise Edition 6 (Java EE 6) and Microsoft .NET software development platforms.

Java EE 5		.NET	
PROS	CONS	PROS	CONS
<ul style="list-style-type: none"> • Runs on any operating system and application server (may need adjustments). • Handles complex, high-volume, high-transaction applications. • Has more enterprise features for session management, fail-over, load balancing, and application integration. • Is favored by experienced enterprise vendors such as IBM, BEA, SAP, and Oracle. • Offers a wide range of vendor choices for tools and application servers. • Has a proven track record. 	<ul style="list-style-type: none"> • Has a complex application development environment. • Tools can be difficult to use. • Java Swing environment's ability to build graphical user interfaces has limitations. • May cost more to build, deploy, and manage applications. • Lacks built-in support for Web services standards. • Is difficult to use for quick-turnaround, low-cost, and mass-market projects. 	<ul style="list-style-type: none"> • Easy-to-use tools may increase programmer productivity. • Has a strong framework for building rich graphical user interfaces. • Gives developers choice of working in more than 20 programming languages. • Is tightly integrated with Microsoft's operating system and enterprise server software. • May cost less, due in part to built-in application server in Windows, unified management, and less expensive tools. • Has built-in support for Web service standards. 	<ul style="list-style-type: none"> • Framework runs only on Windows, restricting vendor choice. • Users of prior Microsoft tools and technology face a potentially steep learning curve. • New run-time infrastructure lacks maturity. • Questions persist about the scalability and transaction capability of the Windows platform. • Choice of integrated development environments is limited. • Getting older applications to run in new .NET environment may require effort.

Source: Carol Silwa, ".NET vs. Java," *Computerworld*, May 20, 2002, p. 31.

The full release of .NET is expected to take several years to complete, with intermittent releases of products such as a personal security service and new versions of Windows and Office that implement the .NET strategy coming on the market separately. Visual Studio .NET is a development environment that is now available, and Windows XP supports certain .NET capabilities.

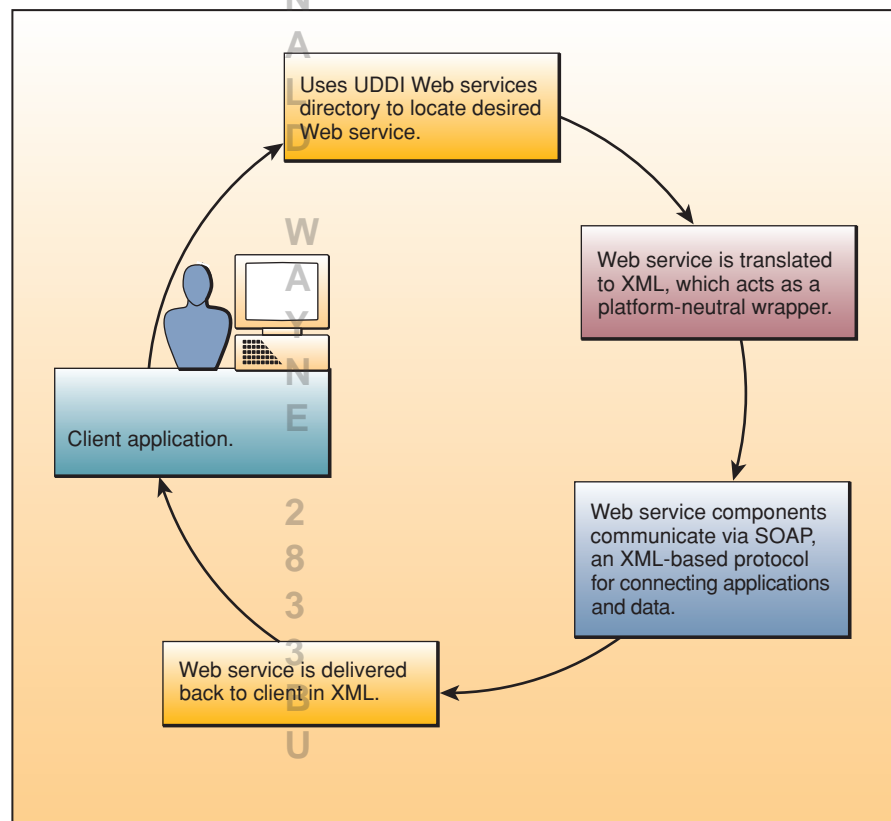
The latest version of Java is Java Enterprise Edition 6 (Java EE 6), which has become the primary alternative to Microsoft's .NET software development platform for many organizations intent on capitalizing on the business potential of Web-based applications and Web services. Figure 4.23 compares the pros and cons of using Java EE 6 and .NET for software development.

Web Services

Web services are software components that are based on a framework of Web and object-oriented standards and technologies for using the Web that electronically link the applications of different users and different computing platforms. Thus, Web services can link key business functions for the exchange of data in real time within the Web-based applications that a business might share with its customers, suppliers, and other business partners. For example, Web services would enable the purchasing application of a business to use the Web to check the inventory of a supplier before placing a large order, while the sales application of the supplier could use Web services to automatically check the credit rating of the business with a credit-reporting agency before approving the purchase. Therefore, among both business and IT professionals, the term *Web services* is commonly used to describe the Web-based business and computing functions or services accomplished by Web services software technologies and standards.

Figure 4.24 illustrates how Web services work and identifies some of the key technologies and standards that are involved. The XML language is one of the key technologies that enable Web services to make applications work between different computing

FIGURE 4.24
The basic steps in
accomplishing a Web
services application.



Source: Adapted from Bala Iyer, Jim Freedman, Mark Gaynor, and George Wyner, "Web Services: Enabling Dynamic Business Networks," *Communications of the Association for Information Systems* 11 (2003), p. 543.

platforms. Also important are **UDDI** (Universal Description, Discovery, and Integration), the “yellow pages” directory of all Web services and how to locate and use them, and **SOAP** (Simple Object Access Protocol), an XML-based protocol of specifications for connecting applications to the data that they need.

Web services promise to be the key software technology for automating access to data and application functions between a business and its trading partners. As companies increasingly move to doing business over the Web, Web services will become essential for the development of the easy and efficient e-business and e-commerce applications that will be required. The flexibility and interoperability of Web services will also be essential for coping with the fast-changing relationships between a company and its business partners that are commonplace in today’s dynamic global business environment.

Airbus: Flying on SAP and Web Services



European aircraft builder Airbus has implemented a Web services–based travel management application from SAP as a first step in a planned groupwide migration to a service-oriented architecture (SOA). The airplane manufacturer is installing the travel management component of SAP’s ERP software, mySAP, which uses SOA technology. “The new system replaces a homegrown system at the company’s plant in France, a Lotus-based system in its Spanish operations, and earlier SAP versions at facilities in Germany and the United Kingdom,” says James Westgarth, manager of travel technology procurement at Airbus.

“We like the idea of an open architecture, which SOA enables,” Westgarth says. “We like the idea of being able to manage everything internally and to cherry-pick for the best solution in every class.” “Additional components, such as online booking, could also come from SAP—if the software vendor has a superior product for that application,” says Westgarth.

The decision to deploy a new Web services–based travel management system was driven in large part by a need to reduce administration costs and improve business processes.

Airbus has a travel budget of 250 million euros, which is used to help pay for more than 180,000 trips annually. The company aims to reduce costs by eliminating the current paper-based reimbursement process, which consumes time and labor, with a system that enables employees to process their own travel expenses online from their desktops or mobile devices.

A key benefit for employees: Reimbursement time will be reduced to 3 days from about 10 days. In addition, the new system allows Airbus to integrate new service providers more easily into its operations, notes Westgarth. The manufacturer has outsourced its valued-added tax reclaim activities to a third party specializing in this service. With the help of application link enablers, Westgarth and his team are able to link their travel management system into the company’s other SAP applications, including finance and human resources. Airbus has a strategy to eventually migrate to the mySAP ERP across multiple systems and countries over a number of years.

“The company chose travel management to pilot mySAP ERP,” says Westgarth. There have been some issues with the rollout of the travel management application, Westgarth concedes. “Because we’re the first big company to implement this technology, we’ve had difficulty finding enough skilled people on the market,” he said. “And some work was required to integrate the Web interface into our portal.”

But Airbus employees, Westgarth said, like the Web-based application’s new user interface, the single sign-on and the step-by-step guidance. And the company likes the flexibility. “No one was talking about low-cost carriers five years ago,” he said. “We need to adapt to the market and to changing needs.”

Source: Adapted from John Blau, “Airbus Flies on Web Services With SAP,” *IDG News Service/CIO Magazine*, June 8, 2006.

Programming Software

Various software packages are available to help programmers develop computer programs. For example, *programming language translators* are programs that translate other programs into machine language instruction codes that computers can execute. Other software packages, such as programming language editors, are called *programming tools* because they help programmers write programs by providing a variety of program creation and editing capabilities. See Figure 4.25.

Language Translator Programs

Computer programs consist of sets of instructions written in programming languages that must be translated by a **language translator** into the computer's own machine language before they can be processed, or executed, by the CPU. Programming language translator programs (or *language processors*) are known by a variety of names. An **assembler** translates the symbolic instruction codes of programs written in an assembler language into machine language instructions, whereas a **compiler** translates high-level language statements.

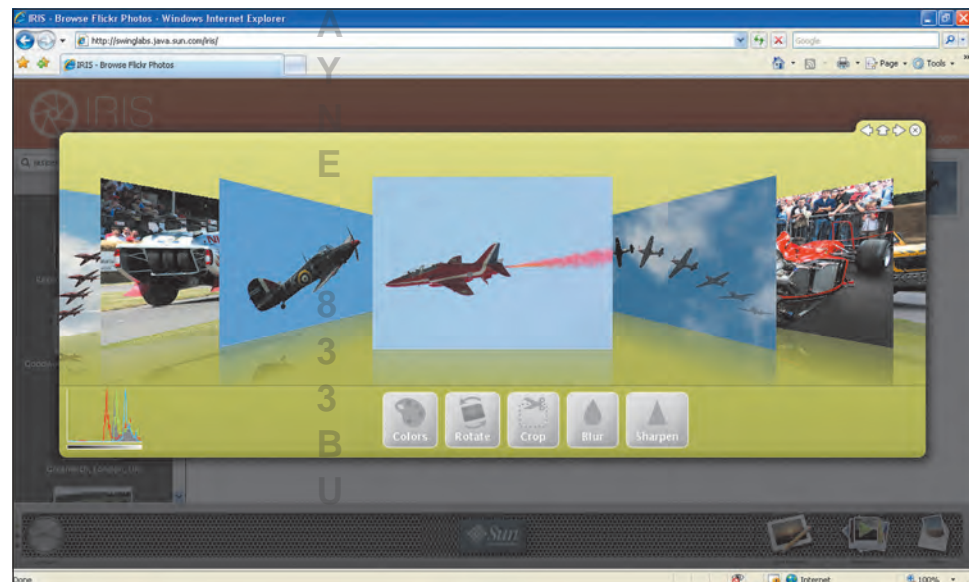
An **interpreter** is a special type of compiler that translates and executes each statement in a program one at a time, instead of first producing a complete machine language program, as compilers and assemblers do. Java is an example of an interpreted language. Thus, the program instructions in Java applets are interpreted and executed *on the fly* as the applet is being executed by a client PC.

Programming Tools

Software development and the computer programming process have been enhanced by adding *graphical programming interfaces* and a variety of built-in development capabilities. Language translators have always provided some editing and diagnostic capabilities to identify programming errors or *bugs*. However, most software development programs now include powerful graphics-oriented *programming editors* and *debuggers*. These **programming tools** help programmers identify and minimize errors while they are programming. Such programming tools provide a computer-aided programming environment, which decreases the drudgery of programming while increasing the efficiency and productivity of software developers. Other programming tools include diagramming packages, code generators, libraries of reusable objects and program code, and prototyping tools. All of these programming tools are an essential part of widely used programming languages like Visual Basic, C++, and Java.

FIGURE 4.25

Using the graphical programming interface of a Java programming tool, Forte for Java, by Sun Microsystems.



Source: Courtesy of Sun Microsystems.

CASE Tools

Since the early days of programming, software developers have needed automated tools. Initially the concentration was on program support tools such as translators, compilers, assemblers, macroprocessors, and linkers and loaders. However, as computers became more powerful and the software that ran on them grew larger and more complex, the range of support tools began to expand. In particular, the use of interactive time-sharing systems for software development encouraged the development of program editors, debuggers, and code analyzers.

As the range of support tools expanded, manufacturers began to integrate them into a single application using a common interface. Such tools were referred to as **CASE tools** (computer-aided software engineering).

CASE tools can take a number of forms and be applied at different stages of the software development process. Those CASE tools that support activities early in the life cycle of a software project (e.g., requirements, design support tools) are sometimes called *front-end* or *upper* CASE tools. Those that are used later in the life cycle (e.g., compilers, test support tools) are called *back-end* or *lower* CASE tools.

Exploring the details of CASE tools is beyond the scope of this text, and you will encounter them again when you study systems analysis and design. For now, remember that CASE is an important part of resolving the problems of complex application development and maintenance of software applications.

Summary

- **Software.** Computer software consists of two major types of programs: (1) application software that directs the performance of a particular use, or application, of computers to meet the information processing needs of users and (2) system software that controls and supports the operations of a computer system as it performs various information processing tasks. Refer to Figure 4.2 for an overview of the major types of software.
- **Application Software.** Application software includes a variety of programs that can be segregated into general-purpose and application-specific categories. General-purpose application programs perform common information processing jobs for end users. Examples are word processing, electronic spreadsheet, and presentation graphics programs. Application-specific programs accomplish information processing tasks that support specific business functions or processes, scientific or engineering applications, and other computer applications in society.
- **System Software.** System software can be subdivided into system management programs and system development programs. System management programs manage the hardware, software, network, and data resources of a computer system during its execution of information processing jobs. Examples of system management programs are operating systems, network management programs, database management systems, system utilities, application servers, and performance and security monitors. Network management programs support and manage telecommunications activities and network performance telecommunications networks. Database management systems control the development, integration, and maintenance of databases. Utilities are programs that perform routine computing functions, such as backing up data or copying files, as part of an operating system or as a separate package. System development programs like language translators and programming editors help IS specialists develop computer programs to support business processes.
- **Operating Systems.** An operating system is an integrated system of programs that supervises the operation of the CPU, controls the input/output storage functions of the computer system, and provides various support services. An operating system performs five basic functions: (1) a user interface for system and network communications with users, (2) resource management for managing the hardware resources of a computer system, (3) file management for managing files of data and programs, (4) task management for managing the tasks a computer must accomplish, and (5) utilities and other functions that provide miscellaneous support services.
- **Programming Languages.** Programming languages are a major category of system software. They require the use of a variety of programming packages to help programmers develop computer programs and language translator programs to convert programming language instructions into machine language instruction codes. The five major levels of programming languages are

machine languages, assembler languages, high-level languages, fourth-generation languages, and object-oriented languages. Object-oriented languages like Java

and special-purpose languages like HTML and XML are being widely used for Web-based business applications and services.

Key Terms and Concepts

These are the key terms and concepts of this chapter. The page number of their first explanation is given in parentheses.

- | | | |
|--|--|--|
| 1. Application service provider (ASP) (143) | 13. Groupware (141) | 27. Presentation graphics software (140) |
| 2. Application software (130) | 14. High-level language (158) | 28. Programming language (157) |
| 3. Assembler language (158) | 15. HTML (161) | 29. Software suites (135) |
| 4. CASE tools (168) | 16. Instant messaging (IM) (137) | 30. Spreadsheet package (139) |
| 5. Cloud computing (145) | 17. Integrated package (136) | 31. System software (147) |
| 6. COTS software (130) | 18. Java (164) | 32. User interface (147) |
| 7. Custom software (130) | 19. Language translator (167) | 33. Utilities (156) |
| 8. Desktop publishing (DTP) (139) | 20. Machine language (157) | 34. Virtual memory (150) |
| 9. E-mail (137) | 21. Middleware (157) | 35. Web browser (136) |
| 10. Fourth-generation language (159) | 22. Multitasking (151) | 36. Web services (165) |
| 11. Function-specific application software (134) | 23. Natural language (159) | 37. Word processing software (138) |
| 12. General-purpose application programs (130) | 24. Object-oriented language (160) | 38. XML (162) |
| | 25. Operating system (147) | |
| | 26. Personal information manager (PIM) (141) | |

Review Quiz

Match one of the previous key terms and concepts with one of the brief examples or definitions that follow. Try to find the best fit for answers that seem to fit more than one term or concept. Defend your choices.

- | | |
|---|---|
| ___ 1. An approach to computing where tasks are assigned to a combination of connections, software, and services accessed over a network. | ___ 10. The ability to do several computing tasks concurrently. |
| ___ 2. Programs that direct the performance of a specific use of computers. | ___ 11. Converts numeric data into graphic displays. |
| ___ 3. A system of programs that manages the operations of a computer system. | ___ 12. Translates high-level instructions into machine language instructions. |
| ___ 4. Companies that own, operate, and maintain application software for a fee as a service over the Internet. | ___ 13. Performs housekeeping chores for a computer system. |
| ___ 5. Integrated software tool that supports the development of software applications. | ___ 14. A category of application software that performs common information processing tasks for end users. |
| ___ 6. Software designed in-house for use by a specific organization or set of users. | ___ 15. Software available for the specific applications of end users in business, science, and other fields. |
| ___ 7. The function that provides a means of communication between end users and an operating system. | ___ 16. Helps you surf the Web. |
| ___ 8. Acronym meaning commercial off-the-shelf. | ___ 17. Uses your networked computer to send and receive messages. |
| ___ 9. Provides a greater memory capability than a computer's actual memory capacity. | ___ 18. Creates and displays a worksheet for analysis. |
| | ___ 19. Allows you to create and edit documents. |

- 20. Enables you to produce your own brochures and newsletters.
- 21. Helps you keep track of appointments and tasks.
- 22. A program that performs several general-purpose applications.
- 23. A combination of individual general-purpose application packages that work easily together.
- 24. Software to support the collaboration of teams and workgroups.
- 25. Uses instructions in the form of coded strings of ones and zeros.
- 26. Uses instructions consisting of symbols representing operation codes and storage locations.
- 27. Uses instructions in the form of brief statements or the standard notation of mathematics.
- 28. Might take the form of query languages and report generators.
- 29. Languages that tie together data and the actions that will be performed on the data.
- 30. As easy to use as one's native tongue.
- 31. Includes programming editors, debuggers, and code generators.
- 32. Produces hyperlinked multimedia documents for the Web.
- 33. A Web document content description language.
- 34. A popular object-oriented language for Web-based applications.
- 35. Windows, Linux, and Mac OS are common examples.
- 36. Software that helps diverse applications work together.
- 37. Enables you to communicate and collaborate in real time with the online associates in your workgroup.
- 38. Links business functions within applications for the exchange of data between companies via the Web.

Discussion Questions

1. What major trends are occurring in software? What capabilities do you expect to see in future software packages?
2. How do the different roles of system software and application software affect you as a business end user? How do you see this changing in the future?
3. Refer to the Real World Case on Software-as-a-Service (SaaS) in the chapter. Do you think GE would have been better off developing a system specifically customized to their needs, given that GE's supply chain is like nothing else in the world?
4. Why is an operating system necessary? That is, why can't an end user just load an application program into a computer and start computing?
5. Should a Web browser be integrated into an operating system? Why or why not?
6. Refer to the Real World Case about the U.S. Department of Defense and its adoption of open-source software in the chapter. Would such an approach work for a commercial organization, or is it limited to government entities? What would be the most important differences in each case, if any?
7. Are software suites, Web browsers, and groupware merging together? What are the implications for a business and its end users?
8. How are HTML, XML, and Java affecting business applications on the Web?
9. Do you think Linux will surpass, in adoption and use, other operating systems for network and Web servers? Why or why not?
10. Which application software packages are the most important for a business end user to know how to use? Explain the reasons for your choices.

Analysis Exercises

Complete the following exercises as individual or group projects that apply chapter concepts to real-world business situations.

1. Desktop Application Recognition Tool Selection

ABC Department Stores would like to acquire software to do the following tasks. Identify which software packages they need.

- a. Surf the Web and their intranets and extranets.
- b. Send messages to one another's computer workstations.
- c. Help employees work together in teams.

- d. Use a group of productivity packages that work together easily.
- e. Help sales reps keep track of meetings and sales calls.
- f. Type correspondence and reports.
- g. Analyze rows and columns of sales figures.
- h. Develop a variety of graphical presentations.

2. Y2K Revisited

The End of Time

Decades ago, programmers trying to conserve valuable storage space shortened year values to two digits. This shortcut created what became known as the “Y2K” problem or “millennium bug” at the turn of the century. Programmers needed to review billions of lines of code to ensure important programs would continue to operate correctly. The Y2K problem merged with the dot-com boom and created a tremendous demand for information technology employees. Information system users spent billions of dollars fixing or replacing old software. The IT industry is only now beginning to recover from the postboom slump. Could such hysteria happen again? It can and, very likely, it will.

Today, most programs use several different schemes to record dates. One scheme, POSIX time, widely employed on UNIX-based systems, requires a signed 32-bit integer to store a number representing the number of seconds since January 1, 1970. “0” represents midnight on January 1, “10” represents 10 seconds after midnight, and “-10” represents 10 seconds before midnight. A simple program then converts these data into any number of international date formats for display. This scheme works well because it allows programmers to subtract one date/time from another date/time and directly determine the interval between them. It also requires only 4 bytes of storage space. But 32 bits still calculates to a finite number, whereas time is infinite. As a business manager, you will need to be aware of this new threat and steer your organization away from repeating history. The following questions will help you evaluate the situation and learn from history.

- a. If 1 represents 1 second and 2 represents 2 seconds, how many seconds can be represented in a binary number 32 bits long? Use a spreadsheet to show your calculations.
- b. Given that POSIX time starts at midnight, January 1, 1970, in what year will time “run out”? Remember that half the available numbers represent dates before 1970. Use a spreadsheet to show your calculations.
- c. As a business manager, what can you do to minimize this problem for your organization?

3. Tracking Project Work

Queries and Reports

You are responsible for managing information systems development projects at AAA Systems. To better track

progress in completing projects, you have decided to maintain a simple database table to track the time your employees spend on various tasks and the projects with which they are associated. It will also allow you to keep track of employees’ billable hours each week. The table below provides a sample data set.

- a. Build a database table to store the data shown and enter the records as a set of sample data.
- b. Create a query that will list the hours worked for all workers who worked more than 40 hours during production week 20.
- c. Create a report grouped by project that will show the number of hours devoted to each task on the project and the subtotal number of hours devoted to each project, as well as a grand total of all hours worked.
- d. Create a report grouped by employee that will show each employee’s hours worked on each task and total hours worked. The user should be able to select a production week and find data for just that week presented.

4. Matching Training to Software Use

3-D Graphing

You have the responsibility to manage software training for Sales, Accounting, and Operations Department workers in your organization. You have surveyed the workers to get a feel for the amounts of time spent using various packages, and the results are shown below. The values shown are the total number of workers in each department and the total weekly hours the department’s workers spend using each software package. You have been asked to prepare a spreadsheet summarizing these data and comparing the use of the various packages across departments.

Department	Employees	Spreadsheet	Database	Presentations
Sales	225	410	1,100	650
Operations	75	710	520	405
Accounting	30	310	405	50

- a. Create a spreadsheet illustrating each application’s average use per department. To do this, you will first enter the data shown above. Then compute the average weekly spreadsheet use by dividing spreadsheet hours by the number of Sales workers. Do this for each department. Repeat these three calculations for both database and presentation use. Round results to the nearest 1/100th.
- b. Create a three-dimensional bar graph illustrating the averages by department and software package.
- c. A committee has been formed to plan software training classes at your company. Prepare a slide presentation with four slides illustrating your findings. The first slide should serve as an introduction to the data; the second slide should

contain a copy of the original data table (without the averages); the third slide should contain a copy of the three-dimensional bar graph from the previous answer; and the fourth slide should contain

your conclusions regarding key applications per department. Use professional labels, formatting, and backgrounds.

Project_Name	Task_Name	Employee_ID	Production_Week	Hours_Worked
Fin-Goods-Inv	App. Devel.	456	21	42
Fin-Goods-Inv	DB Design	345	20	20
Fin-Goods-Inv	UI Design	234	20	16
HR	Analysis	234	21	24
HR	Analysis	456	20	48
HR	UI Design	123	20	8
HR	UI Design	123	21	40
HR	UI Design	234	21	32
Shipmt-Tracking	DB Design	345	20	24
Shipmt-Tracking	DB Design	345	21	16
Shipmt-Tracking	DB Development	345	21	20
Shipmt-Tracking	UI Design	123	20	32
Shipmt-Tracking	UI Design	234	20	24

REYNOLDS, DONALD WAYNE 2833BU

REAL WORLD

CASE

3

Wolf Peak International: Failure and Success in Application Software for the Small-to-Medium Enterprise

One of the hazards of a growing small business is a software upgrade. If you pick the wrong horse, you may find yourself riding in the wrong direction. Correcting your course may mean not only writing off your first upgrade selection but then going through the agonizing process of finding a better software solution for your company. That's what happened to Wolf Peak International of Layton, Utah, which designs and manufactures eyewear for the safety, sporting, driving, and fashion industries. Founded in 1998, the privately held small to midsize enterprise (SME) also specializes in overseas production, sourcing, importing, and promotional distribution services.

In Wolf Peak's early days, founder-owner Kurt Daems was happy using QuickBooks to handle accounting chores. The package is user friendly and allowed him to drill down to view transaction details or combine data in a variety of ways to create desired reports. As the company prospered, however, it quickly outgrew the capabilities of QuickBooks.

"As Wolf Peak got bigger, the owner felt the need to get into a more sophisticated accounting system," says Ron Schwab, CFO at Wolf Peak International. "There were no financial people in-house at the time the decision was made to purchase a replacement for QuickBooks, and the decision was made without a finance person in place to review it."

Wolf Peak selected one of several accounting software packages promoted to growing SMEs. By the time Schwab joined the company, the package had been installed for six months, following an implementation period that lasted a full year. "The biggest difficulty for QuickBooks users is to go from a very friendly user interface and the ability to find information easily to a more sophisticated, secured, batch-oriented accounting system that became an absolute nightmare to get data out of," notes Schwab. "So the company paid a lot of money to have this new accounting system, but nobody knew how to go in and extract financial or operational data used to make critical business decisions."

There were other problems. Developing reusable reports was difficult, time-consuming, and expensive. The company paid IT consultants to develop reports for specific needs, some of which still had not been delivered, months after they were commissioned. Ad hoc reporting was similarly intractable. Furthermore, the company's prior-year history in QuickBooks could not be converted into the new accounting package. A situation like this creates serious problems.

Accustomed to keeping close tabs on the company's operations, Daems found that he simply could not get the information he wanted. He began to lose track of his business. "He got so fed up he finally came to me and said he was ready to look at a SAP software alternative he'd heard about," Schwab recalls. "He wasn't ready to buy it, though, because he'd just sunk a lot of money into the new accounting package."

One year after Wolf Peak had switched over to the new accounting software, Schwab called the offices of JourneyTEAM, a local SAP services partner, and asked their software consultants to demonstrate the SAP Business One software suite.

SAP Business One is an integrated business management software package designed specifically for SMEs like Wolf Peak: The application automates critical operations including sales, finance, purchasing, inventory, and manufacturing and delivers an accurate, up-to-the-minute view of the business. Its relative affordability promises a rapid return on investment, and its simplicity means users have a consistent, intuitive environment that they can learn quickly and use effectively.

"We had a wish list from various company employees asking for a variety of capabilities," recalls Schwab. "The JourneyTEAM people came in and demonstrated all those functionalities and more. They even generated four or five reports that we had spent several thousand dollars and several months trying to get from our other software consultants and had not yet received. Based on our data that they had input into Business One, JourneyTEAM put those reports together in an afternoon."

Daems still had a few reservations: He needed the buy-in of his VP of sales and was concerned about cost. He still wasn't ready to write off the recently installed accounting software.

JourneyTEAM came in and gave another presentation for the Wolf Peak sales team and, following that, came back with an acceptable quote. With some pain, but also considerable relief, Daems wrote off the existing accounting package. "We felt the benefits of SAP Business One far outweighed the costs and time already invested in that software system," Daems says.

Implementation of Business One took just seven weeks from the day of the initial sales presentation. "We implemented SAP Business One during our busiest period of the year with no disruptions," notes Schwab. "It went better than I expected, in particular the cutover and conversion to Business One. JourneyTEAM did an amazing job of getting all our old records converted with no real problems at all. We met our June 30 deadline and cutover during the succeeding long weekend without incident."

Schwab's enthusiasm for SAP Business One is high. "This is the best accounting program I've ever worked with," he says. "I can drill down to anything I want. And with the XL Reporter tool, I can build reports on the fly."

Business One includes a seamlessly integrated reporting and financial analysis tool called XL Reporter that works with Microsoft Excel to provide instant access to financial and operational data. It reports on live data drawn from a variety of sources including general ledger, receivables, payables, sales, purchasing, and inventory software. "Now we're

REYNOLDS, DONALD W. AYNÉ 2833 BU

building the reports we want,” says Schwab. “To have a program like XL Reporter that lets us build custom reports, preset regular updates, and then work within Microsoft Excel—that’s hugely valuable to us. Nobody else offers the ability to do ad hoc queries so easily. Even people who aren’t serious programmers can go in and create the documents they need within the limits of their authorizations. So I highly recommend it.”

For years, Daems had been running an open receivables report that presents, for example, all the invoices that are 15 days past due and greater than \$450. Unfortunately, he simply could not run a report like that with the software package he bought to replace his old QuickBooks program. That situation has now changed.

“With SAP Business One, we can go in there and ask for those parameters and then sort it by oldest, biggest amount, or customer,” says Schwab. “And it’s paperless. The accounts receivable person doesn’t have to print anything out and then write a bunch of notes on it and type them into the system for someone else to find. It’s all right there.”

Wolf Peak also requires a very complicated commissions report, used to generate the checks that go out to the company’s commissioned sales representatives, who receive individualized reports as well. The previous consultants were unable to deliver this set of reports. JourneyTEAM was able to develop it on Business One in an afternoon.

Wolf Peak is already expanding its use of SAP Business One into other areas. The company has applied the software to warehouse management, where it enables Wolf Peak to manage inventory, receiving, warehouse delivery, shipping, and all the other aspects of the warehousing task. Inventory is one of the company’s biggest assets, and it has to be managed well. “We have an audit report that lists all of the

inventory, the current on-hand quantity, and the demands on it through sales orders or outstanding purchase orders,” Schwab says.

This report then lists the value of that inventory and allows Schwab to look at the activity against any inventory item during any period. Beyond that, it enables him to drill down to the actual invoices that affect that inventory item. “We want to minimize what we have on hand,” he says, “but we always have to be sure we have enough to meet our customers’ needs. Business One lets us do that.”

Wolf Peak’s management has also begun using the customer relationship management (CRM) functionality within Business One to assist with its collection of receivables. The company’s plan is to extend its use of the software to develop and track sales opportunities as well. Three months following its installation, Wolf Peak is quite happy with its decision to go with SAP’s Business One software. “Reports that used to take months to create—if we could get them at all—can now be created in minutes,” says Schwab.

A less tangible but no less important benefit is the renewed confidence Business One brings to management. “A company’s greatest untapped asset is its own financial information,” says Schwab. “SAP Business One creates an environment where the decision makers get the information they want on a timely basis, in a format they can use. It’s amazing what happens when management begins to see what is really happening inside the enterprise. Business One delivers useful information to help make good business decisions—and that’s really the bottom line. This is a business management tool.”

Source: Adapted from SAP America, “Wolf Peak: Making the Best Choice to Support Growth,” *SAP BusinessInsights*, March 2007; JourneyTEAM, “Wolf Peak Success Story—SAP Business One,” ABCComputer.com, March 2007.

CASE STUDY QUESTIONS

1. What problems occurred when Wolf Peak upgraded from QuickBooks to a new accounting software package? How could these problems have been avoided?
2. Why did SAP’s Business One prove to be a better choice for Wolf Peak’s management than the new accounting software? Give several examples to illustrate your answer.
3. Should most SMEs use an integrated business software suite like SAP Business One instead of specialized accounting and other business software packages? Why or why not?

REAL WORLD ACTIVITIES

1. This case demonstrates failure and success in the software research, selection, and installation process, as well as some major differences among business application software packages in capabilities, such as ease of use and information access for employees and management. Search the Internet to find several more examples of such success and failure for software suites like SAP Business One or Oracle E-Business Suite and specialized business packages like QuickBooks or Great Plains Accounting.
2. Break into small groups with your classmates to discuss several key differences you have found on the basis of your Internet research. Then make recommendations to the class for how these differences should shape the business application software selection decision for an SME.

REAL WORLD

CASE

4

Power Distribution and Law Enforcement: Reaping the Benefits of Sharing Data through XML

A power consortium that distributes a mix of “green” and conventional electricity is implementing an XML-based settlements system that drives costs out of power distribution. The Northern California Power Agency (NCPA) is one of several state-chartered coordinators in California that schedules that delivery of power to the California power grid and then settles the payment due to suppliers. NCPA sells the power generated by the cities of Palo Alto and Santa Clara, as well as hydro- and geothermal sources farther north.

Power settlements are a highly regulated and complicated process. Each settlement statement contains how much power a particular supplier delivered and how much was used by commercial vs. residential customers, and the two have different rates of payment. The settlements are complicated by the fact that electricity meters are read only once every 90 days; many settlements must be based on an estimate of consumption that gets revised as meter readings come in.

On behalf of a supplier, NCPA can protest that fees for transmission usage weren’t calculated correctly, and the dispute requires a review of all relevant data. Getting one or more of these factors wrong is commonplace. “Power settlements are never completely settled,” says Bob Caracristi, manager of power settlements for NCPA. “Negotiations over details may still be going on a year or two after the power has been delivered.”

Furthermore, “the enormity of the data” has in the past required a specialist vendor that creates software to analyze the massive settlement statements produced by the grid’s manager, the California Independent System Operator. NCPA sought these vendor bids three years ago and received quotes that were “several hundred thousand dollars a year in licensing fees and ongoing maintenance,” remarks Caracristi. The need for services from these customized systems adds to the cost of power consumption for every California consumer.

Faced with such a large annual expense, NCPA sought instead to develop the in-house expertise to deal with the statements. Senior programmer analyst Carlo Tiu and his team at NCPA used Oracle’s XML-handling capabilities to develop a schema to handle the data and a configuration file that contained the rules for determining supplier payment from the data. That file can be regularly updated, without needing to modify the XML data themselves. In doing so, the NCPA gained a step on the rest of the industry, as the California Independent System Operator started requiring all of its vendors to provide power distribution and billing data as XML files. NCPA has already tested its ability to process XML settlement statements automatically and has scaled out its Oracle system to 10 times its needs “without seeing any bottlenecks,” says Tiu.

Being able to process the Independent System Operator statements automatically will represent huge cost savings to

NCPA, according to IS manager Tom Breckon. “When settlement statements come in,” Breckon says, “NCPA has eight working days to determine where mistakes may have been made. If we fail to get back to [the California Independent System Operator], we lose our chance to reclaim the monies from corrections.” Yet, he acknowledges, “we can’t inspect that volume of data on a manual basis.”

Gaining the expertise to deal with settlements as XML data over the past three years has cost NCPA the equivalent of one year’s expense of a manager’s salary. Meanwhile, NCPA has positioned itself to become its own statement processor and analyzer, submit disputes to the California Independent System Operator for corrections, and collect more of those corrected payments for members on a timely basis. “In my opinion,” says Breckon, “everybody will be doing it this way five years from now. It would reduce costs for all rate payers.”

In the state of Ohio, almost 1,000 police departments have found critical new crime-fighting tools by gaining access to the digital records kept by neighboring law enforcement agencies. The Ohio Law Enforcement Gateway Search Engine is an Internet-based tool that can securely comb through numerous crime databases using a single log-in and query, making it easier to use than separate crime databases. For police officers, searching for information on a suspect or a rash of crimes used to require manually logging into several separate crime databases, which could take hours. Now, officers in even the smallest communities can log in just once and quickly gain access to criminal information.

The project, which began in 2003, faced a major hurdle: finding a way to get the disparate crime information systems to interoperate with each other. “Everybody wants to share, but nobody wants to use the same product,” says Chief Gary Vest of the Powell, Ohio, Police Department, near Columbus. In a major metropolitan area in Ohio, there can be 30 different police departments, each using different products that aren’t linked, he says. “That made it difficult for local departments to link suspects and crimes in neighboring jurisdictions.”

To make the systems compatible, crime records management vendors rewrote their software so that data from participating departments could be converted into the gateway format for easier data sharing. The vendors used a special object-oriented Global Justice XML Data Model and interoperability standards developed by the U.S. Department of Justice for such purposes. What makes this project different from other fledgling police interoperability programs in the United States is that it’s a standards-based system. “You don’t have to throw out your vendor to play,” notes Vest.

So far, Ohio police can’t search on criminal “M.O.’s,” but that capability is being worked on. By combing local police records, officers can search for a suspect’s name even

REYNOLDS, DONALD W. AYNIE 2833BU

before it's in the national databases or other larger data repositories, says Vest. "You're a step earlier." Other regional police interoperability projects are in progress around the nation, but this is believed to be the first statewide effort.

In San Diego County, police agencies have been sharing crime data for 25 years using a custom program called the Automated Regional Justice Information System (ARJIS). Barbara Montgomery, project manager for ARJIS, says it differs from the Ohio initiative because it is mainframe-based and all police agencies have to use the same software to access information. Such data-sharing programs are not widespread in the United States because of their cost, especially for smaller police departments, she says. In fact, ARJIS was made possible only after a number of departments pooled their money.

"No single police department could afford to buy [the hardware and the skills of] a bunch of computer programmers so it was truly a 'united we stand, divided we fall' approach," Montgomery says. "The next generation of ARJIS is being planned now, with the system likely to evolve over the next few years from its mainframe roots to a server-based enterprise architecture for more flexibility," says Montgomery.

Along the same lines, the Florida Department of Law Enforcement will begin work on a \$15 million project to integrate the back-end systems of 500 law enforcement

organizations across the state. In many cases, investigators in Florida law enforcement offices now gather information from other departments in the state via telephone or e-mail. The Florida Law Enforcement Exchange project promises to provide access to statewide law enforcement data with a single query, says state's CIO Brenda Owens, whose IT unit is overseeing the project.

"Our goal is to provide seamless access to data across the state," says Owens. "An operator sitting at a PC in a police department doesn't know or care what the data look like; they can put the inquiry in and get the information back."

Large integration projects such as this often derail because it's difficult to get different groups to agree on metadata types. "The metadata management or understanding the common elements is a huge part of [an integration project]," notes Ken Vollmer, an analyst at Forrester Research. "Trying to combine information from two agencies—that is hard enough. In Florida, you're talking 500 agencies, and they have to have some software to help them determine what the common data elements are."

Source: Adapted from Charles Babcock, "Electricity Costs Attacked through XML," *InformationWeek*, December 26, 2007; Todd Weiss, "Ohio Police Use Specialized Software to Track Data (and Bad Guys)," *Computerworld*, June 23, 2006; and Heather Havenstein, "Florida Begins Linking Its Law Enforcement Agencies," *Computerworld*, February 13, 2006.

CASE STUDY QUESTIONS

1. What is the business value of XML to the organizations described in the case? How are they able to achieve such large returns on investment?
2. What are other ways in which XML could be used by organizations to create value and share data? Look for examples involving for-profit organizations to gain a more complete perspective on the issue.
3. What seem to be important elements in the success of projects relying on extensive use of XML across organizations, and why? Research the concept of metadata to inform your answer.

REAL WORLD ACTIVITIES

1. XBRL stands for eXtensible Business Reporting Language, and it is one of the family of XML languages that is becoming standard for business communication across companies. Among other uses, the Securities and Exchange Commission has run a voluntary XBRL filing program since 2005. Go online and research the current status of XBRL implementation and adoption, including examples of companies that are already using it for business purposes. Prepare a report to share your findings.
2. Investigate other large-scale, systemwide implementations of XML such as the one described in the case involving the California Independent System Operator. Prepare a presentation with the proposed or realized costs and benefits of those efforts and share your findings with the class.