PART TWO

Information Technology Infrastructure

Chapter 5
IT Infrastructure and Emerging Technologies

Chapter 6
Foundations of Business Intelligence: Databases and Information Management

Chapter 7
Telecommunications, the Internet, and Wireless Technology

Chapter 8
Securing Information Systems

Part Two provides the technical foundation for understanding information systems by examining hardware, software, database, and networking technologies along with tools and techniques for security and control. This part answers questions such as: What technologies do businesses today need to accomplish their work? What do I need to know about these technologies to make sure they enhance the performance of the firm? How are these technologies likely to change in the future? What technologies and procedures are required to ensure that systems are reliable and secure?

198

IT Infrastructure and Emerging Technologies

CHAPTER

LEARNING OBJECTIVES

After reading this chapter, you will be able to answer the following questions:

1. What is IT infrastructure and what are the stages and drivers of IT infrastructure evolution?
2. What are the components of IT infrastructure?
3. What are the current trends in computer hardware platforms?
4. What are the current trends in computer software platforms?
5. What are the challenges of managing IT infrastructure and management solutions?

MyMISLab™

Visit **mymislab.com** for simulations, tutorials, and end-of-chapter problems.

CHAPTER CASES

Toyota Motor Europe Manages with the Cloud
The Greening of the Data Center
Is It Time for Cloud Computing?
The Pleasures and Pitfalls of BYOD

VIDEO CASES

ESPN.com: Getting to eXtreme Scale On the Web
Salesforce.com: Managing by Smartphone
Acxiom's Strategic Advantage: IBM's Virtual Blade Platform

Instructional Videos:

Google and IBM Produce Cloud Computing
IBM Blue Cloud Is Ready-to-Use Computing

TOYOTA MOTOR EUROPE MANAGES WITH THE CLOUD

Toyota Motor Europe (TME) manages the sales, marketing, engineering, and dealership network for Toyota and Lexus vehicles in Europe. TME is based in Brussels, Belgium. Founded in 1963, TME sales peaked at 1.3 units in 2007, and ran about 900,000 units in 2014. TME employs about 95,000 people. T31 National Marketing and Sales Companies (NMSCs) coordinate 3,100 sales outlets or dealerships, and nine manufacturing plants.

With millions of consumers in Europe who have purchased Toyota cars in the last 50 years, providing a consistent repair and maintenance service to customers was always a challenge. Repair and maintenance manuals had been replaced by personal computers in the 1990s and 2000s. These early computerization efforts simply automated the traditional printed manuals, but had the advantage that they could be updated by distributing disks to the sales outlets. One problem: some outlets did not update their PCs in a timely fashion, and there was no way for TME management to manage the upgrade process.

TME changed its approach to maintenance as onboard computers and sensors were introduced into Toyota vehicles. These onboard computers could gather and store information from vehicle sensors. The vehicle data could be analyzed by PCs at dealerships. TME provided dealerships with over 3,500 new PCs running up-to-date software that could immediately analyze a customer's car engine performance, and provide critical maintenance and repair information and recommendations to local mechanics.



© Prodakszyn/Shutterstock

While this change in approach was a vast improvement, management did not have a solution for how to manage these 3,500 computers, or to know if local outlets were using the latest versions of the software. In fact these PCs remained outside the TME firewall, and they were essentially stand-alone computers as in the past. Thus, it was impossible for headquarters to know if dealers were providing the same quality of service to customers, or to know if dealers' PCs were free of viruses and operating appropriately. One partial solution was to send TME IT staff to the dealerships to install software updates, check for viruses, and advise mechanics on new techniques. But this was wasteful and time-consuming.

Management decided to use a cloud-based solution from Microsoft called Windows Intune. Intune allows central IT staff to run PC management tasks remotely. New software can be installed, virus protection programs run, and the status updated. Each local PC needs and Internet connection, and standard browser. The local PC downloads a client version of the program. The client sends information on the PC to central IT staff.

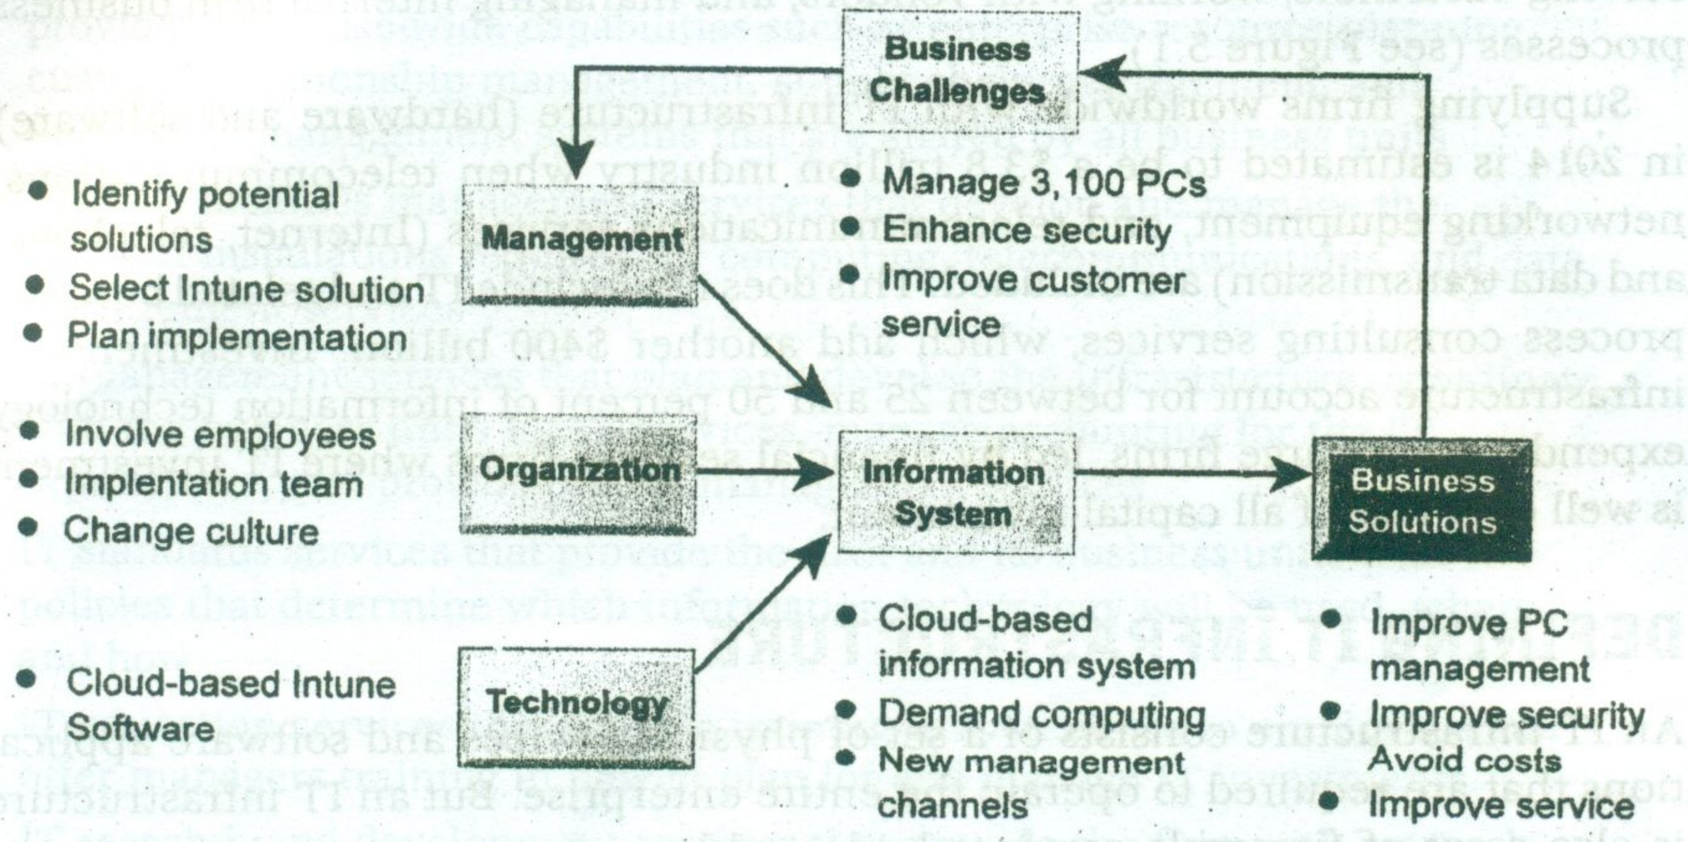
Compared to building their own system to manage its PCs, TME found the implementation process was more efficient. TME sent each dealership an installation package, and a link to the corporate TMW Web site. Dealerships would establish a link to the Web site, and download the dealership software. The Windows Intune interface is easy to use and requires only a few hours of training.

Rather than spend on building their own IT infrastructure and software to management its dealers' PCs, the Intune cloud-based model provides a far less expensive pay-as-you go, (demand computing) model, and there are no additional hardware or software costs, or maintenance costs. TME believes it saved over €1 million in infrastructure and management costs over a three-year period. The transition to tablet PCs in 2015 is expected to be problem-free because of their investment in a coordinated, cloud-based solution. Other management objectives were also achieved such as better security, customer service, and reduced operating costs.

Sources: Microsoft Corporation, "Mitchells and Butlers Boosts Service and Satisfaction with Managed Mobile Platform," www.microsoft.com/en-us/server-cloud/products/microsoft-intune/Microsoft Corporation; "Toyota Motor Europe Automotive Retailer Avoids 1.3 Million in IT Costs with Cloud-Based PC Management Tool," www.microsoft.com, November 12, 2014; Marcia Savage, "8 Cloud-based IT Management Tools," Networkcomputing.com, June 18, 2013.

The experience of Toyota Motor Europe illustrates the importance of information technology infrastructure in running a business today. The right technology at the right price can improve organizational performance. Toyota Europe was able to use contemporary cloud technology infrastructure to improve the quality of its service, enhance security, and reduce operating. But it wasn't just technology that needed updating. Toyota also changed its entire philosophy of how to maintain vehicles using embedded computers in vehicles. In addition, it needed to change how it supported its 3,100 dealers. Rather than sending expensive staff members to each dealer, they were able to rely instead on the cloud-based solution to provide support.

The chapter-opening case diagram calls attention to important points raised by this case and this chapter. In an effort to upgrade the dealer computer system, Toyota Motor Europe found that new computer infrastructure was needed in order to properly manage their 3,000 strong dealer network. Increasingly, dealer PCs were under attack from computer viruses, and dealers could often not upgrade



their software in a coordinated fashion. Customers in turn were expecting dealers to provide the most current solutions for their cars in need of repair. The existing IT infrastructure would require significant additional investment to solve these challenges.

The solution to these challenges involved changes in management, organization, and technology. Management chose to use a new approach to coordinating dealer PCs by adopting a cloud-based solution from Microsoft called Intune. Management developed an implementation plan, and made organizational changes in training, and in the culture of the organization. Rather than rely on visits from headquarters IT staff, dealers were now expected to maintain their local PCs using the company's new management software. The solution is serving important management goals: lower cost, better security, and better management of the dealer PC network.

Here are some questions to think about: How does information technology help Toyota Motor Europe solve its own business problems? How does the new cloud-based technology help the dealers solve their business problems?

5.1 WHAT IS IT INFRASTRUCTURE AND WHAT ARE THE STAGES AND DRIVERS OF IT INFRASTRUCTURE EVOLUTION?

In Chapter 1, we defined *information technology (IT) infrastructure* as the shared technology resources that provide the platform for the firm's specific information system applications. An IT infrastructure includes investment in hardware, software, and services—such as consulting, education, and training—that are shared across the entire firm or across entire business units in the firm. A firm's IT infrastructure provides the foundation for

serving customers, working with vendors, and managing internal firm business processes (see Figure 5.1).

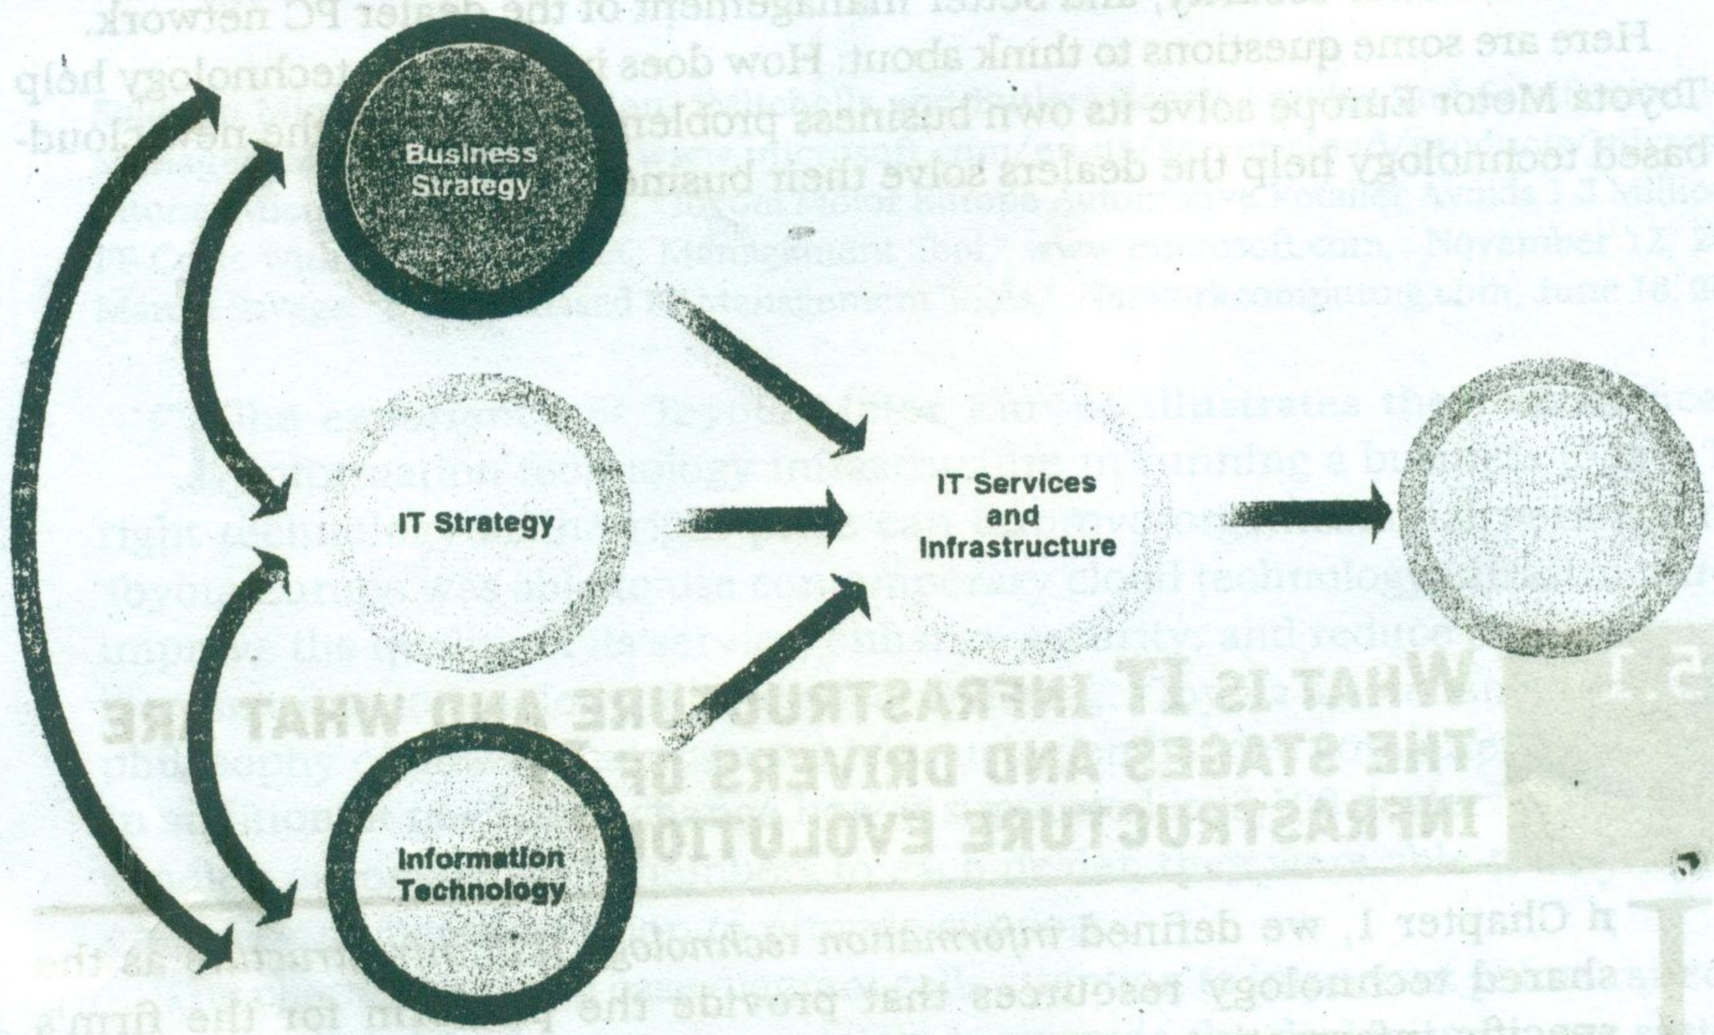
Supplying firms worldwide with IT infrastructure (hardware and software) in 2014 is estimated to be a $3.8 trillion industry when telecommunications, networking equipment, and telecommunications services (Internet, telephone, and data transmission) are included. This does not include IT and related business process consulting services, which add another $400 billion. Investments in infrastructure account for between 25 and 50 percent of information technology expenditures in large firms, led by financial services firms where IT investment is well over half of all capital investment.

DEFINING IT INFRASTRUCTURE

An IT infrastructure consists of a set of physical devices and software applications that are required to operate the entire enterprise. But an IT infrastructure is also a set of firmwide services budgeted by management and comprising both human and technical capabilities. These services include the following:

- Computing platforms used to provide computing services that connect employees, customers, and suppliers into a coherent digital environment, including large mainframes, midrange computers, desktop and laptop computers, and mobile handheld and remote cloud computing services.
- Telecommunications services that provide data, voice, and video connectivity to employees, customers, and suppliers
- Data management services that store and manage corporate data and provide capabilities for analyzing the data

FIGURE 5.1 CONNECTION BETWEEN THE FIRM, IT INFRASTRUCTURE, AND BUSINESS CAPABILITIES



The services a firm is capable of providing to its customers, suppliers, and employees are a direct function of its IT infrastructure. Ideally, this infrastructure should support the firm's business and information systems strategy. New information technologies have a powerful impact on business and IT strategies, as well as the services that can be provided to customers.

- Application software services, including online software services, that provide enterprise-wide capabilities such as enterprise resource planning, customer relationship management, supply chain management, and knowledge management systems that are shared by all business units
- Physical facilities management services that develop and manage the physical installations required for computing, telecommunications, and data management services
- IT management services that plan and develop the infrastructure, coordinate with the business units for IT services, manage accounting for the IT expenditure, and provide project management services
- IT standards services that provide the firm and its business units with policies that determine which information technology will be used, when, and how
- IT education services that provide training in system use to employees and offer managers training in how to plan for and manage IT investments
- IT research and development services that provide the firm with research on potential future IT projects and investments that could help the firm differentiate itself in the marketplace

This "service platform" perspective makes it easier to understand the business value provided by infrastructure investments. For instance, the real business value of a fully loaded personal computer operating at 3.4 gigahertz that costs about $1,000 and a high-speed Internet connection is hard to understand without knowing who will use it and how it will be used. When we look at the services provided by these tools, however, their value becomes more apparent: The new PC makes it possible for a high-cost employee making $100,000 a year to connect to all the company's major systems and the public Internet. The high-speed Internet service saves this employee about one hour per day in reduced wait time for Internet information. Without this PC and Internet connection, the value of this one employee to the firm might be cut in half.

EVOLUTION OF IT INFRASTRUCTURE

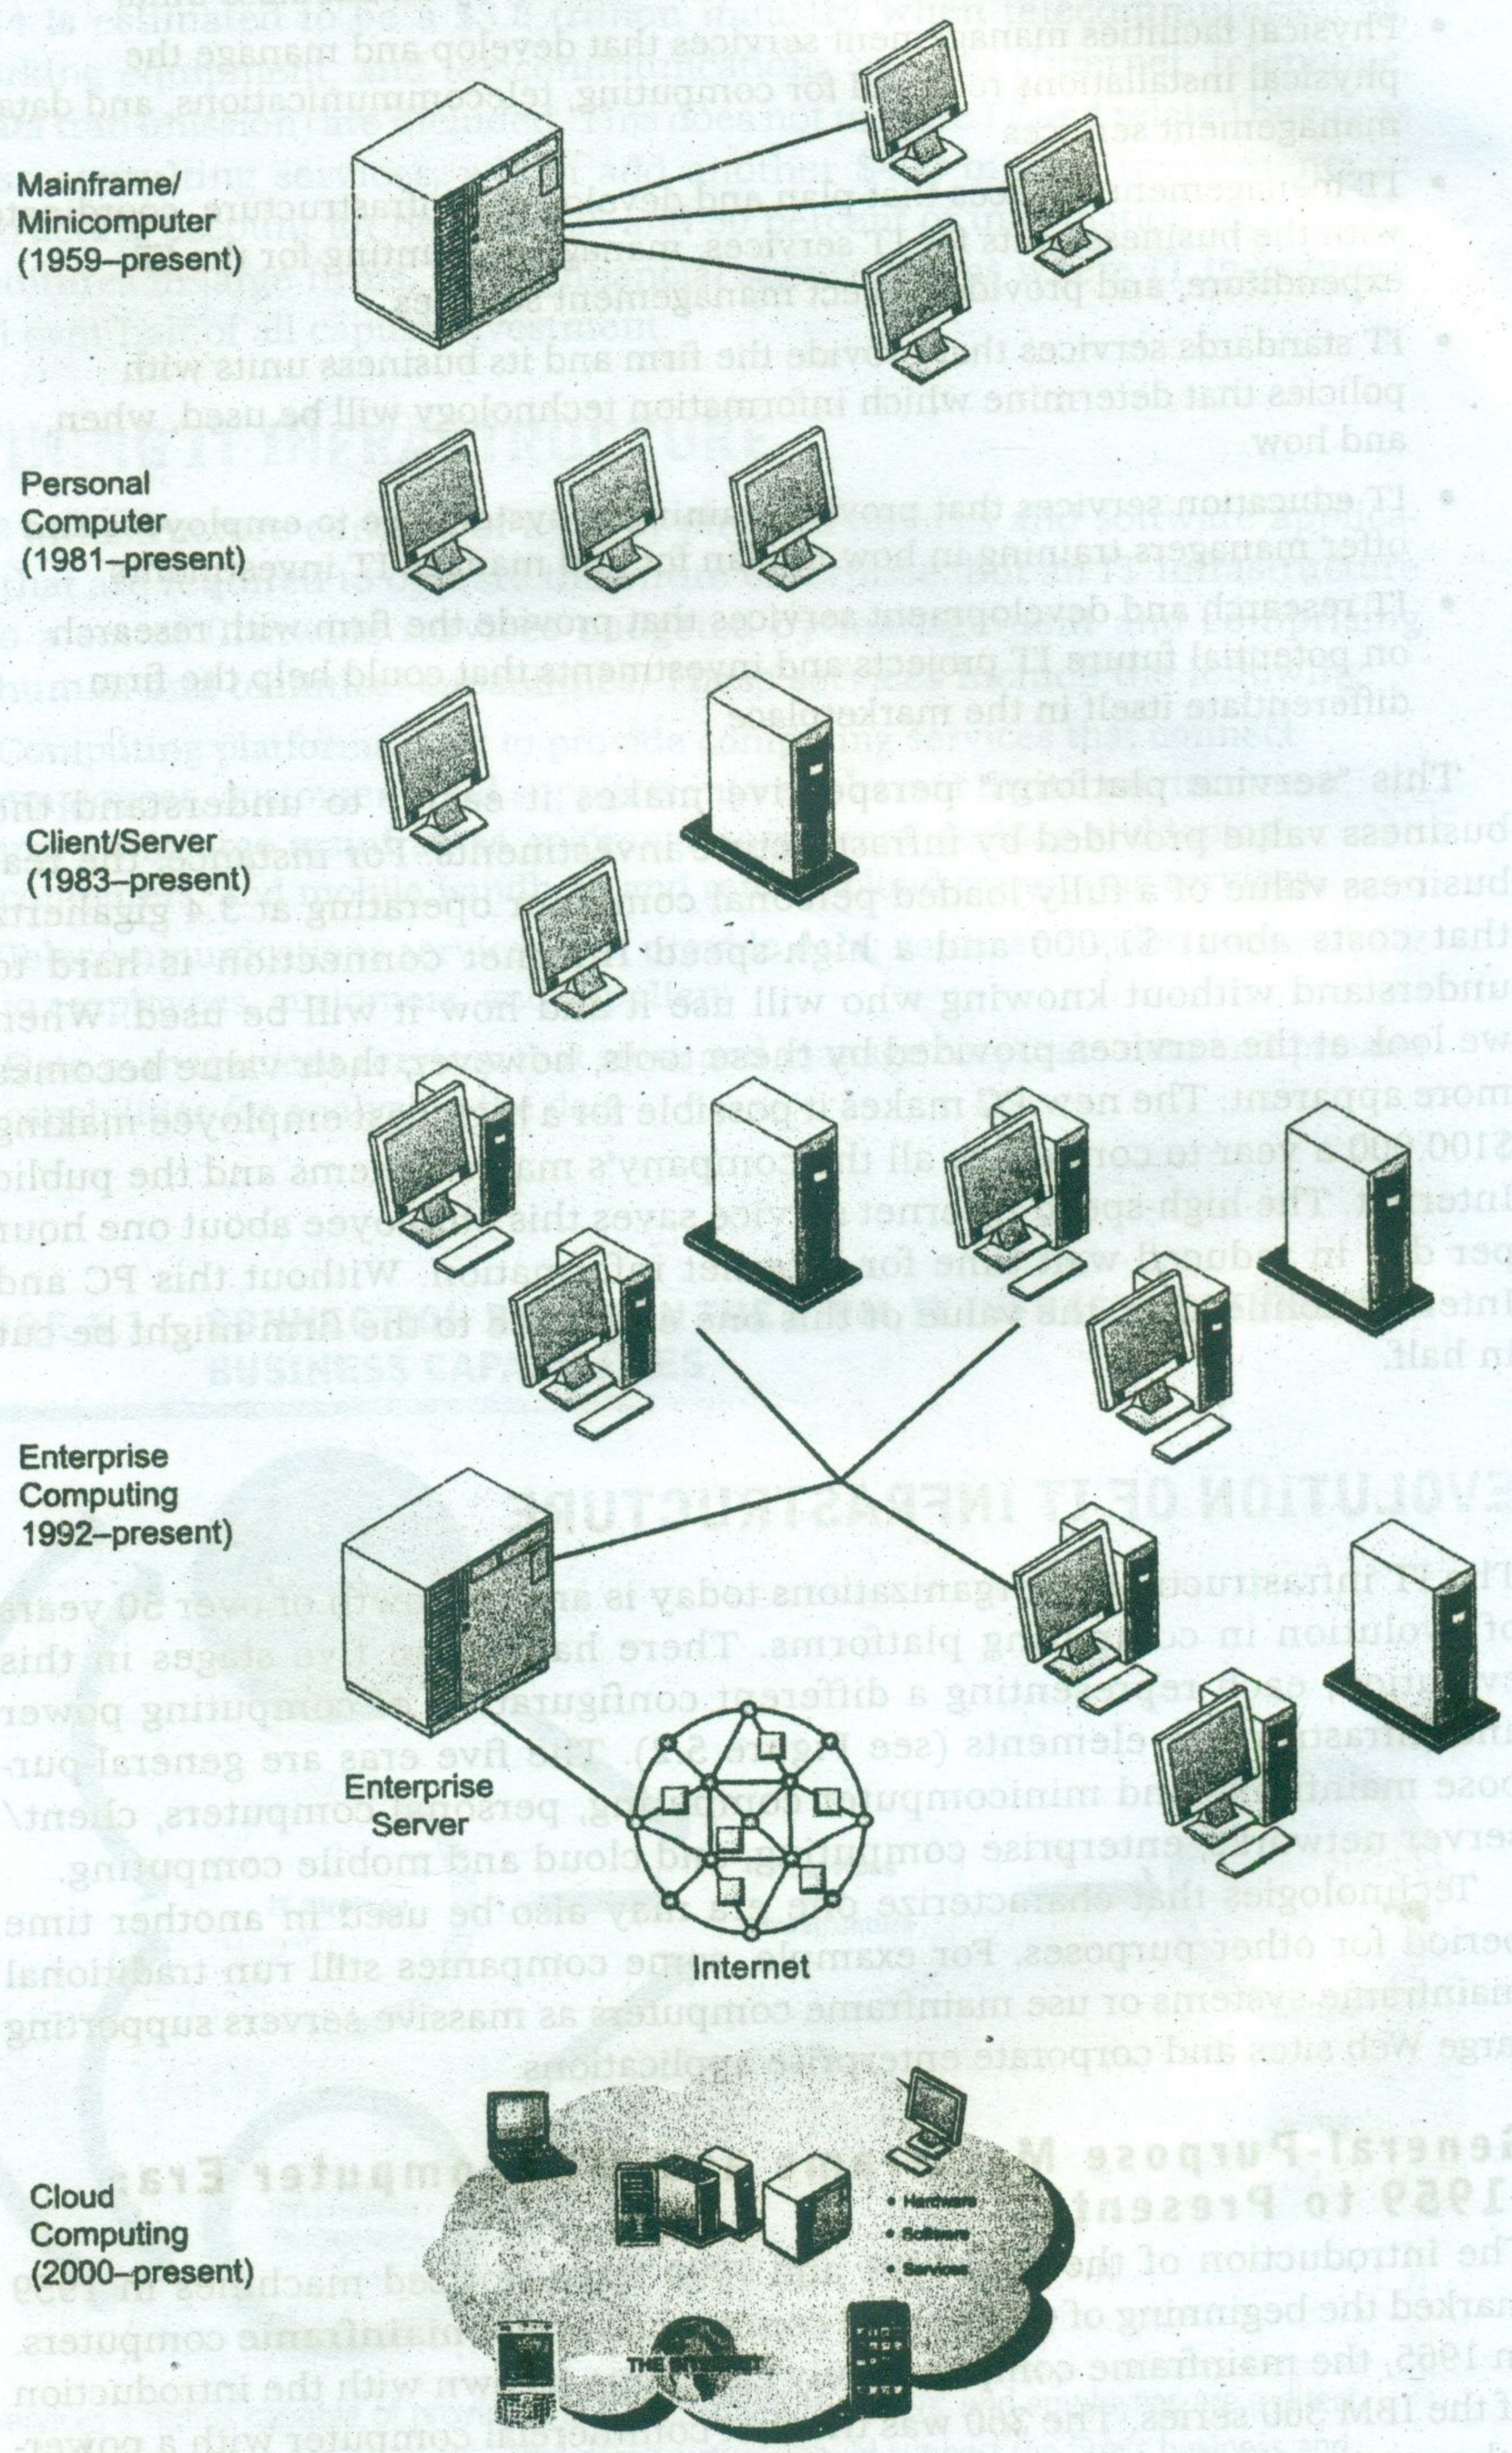
The IT infrastructure in organizations today is an outgrowth of over 50 years of evolution in computing platforms. There have been five stages in this evolution, each representing a different configuration of computing power and infrastructure elements (see Figure 5.2). The five eras are general-purpose mainframe and minicomputer computing, personal computers, client/server networks, enterprise computing, and cloud and mobile computing.

Technologies that characterize one era may also be used in another time period for other purposes. For example, some companies still run traditional mainframe systems or use mainframe computers as massive servers supporting large Web sites and corporate enterprise applications.

General-Purpose Mainframe and Minicomputer Era: (1959 to Present)

The introduction of the IBM 1401 and 7090 transistorized machines in 1959 marked the beginning of widespread commercial use of **mainframe** computers. In 1965, the mainframe computer truly came into its own with the introduction of the IBM 360 series. The 360 was the first commercial computer with a powerful operating system that could provide time sharing, multitasking, and virtual memory in more advanced models. IBM has dominated mainframe computing

FIGURE 5.2 ERAS IN IT INFRASTRUCTURE EVOLUTION



illustrated here are the typical computing configurations characterizing each of the five eras of IT infrastructure evolution.

from this point on. Mainframe computers became powerful enough to support thousands of online remote terminals connected to the centralized mainframe using proprietary communication protocols and proprietary data lines.

The mainframe era was a period of highly centralized computing under the control of professional programmers and systems operators (usually in a corporate data center), with most elements of infrastructure provided by a single vendor, the manufacturer of the hardware and the software.

This pattern began to change with the introduction of **minicomputers** produced by Digital Equipment Corporation (DEC) in 1965. DEC minicomputers (PDP-11 and later the VAX machines) offered powerful machines at far lower prices than IBM mainframes, making possible decentralized computing, customized to the specific needs of individual departments or business units rather than time sharing on a single huge mainframe. In recent years, the minicomputer has evolved into a midrange computer or midrange server and is part of a network.

Personal Computer Era: (1981 to Present)

Although the first truly personal computers (PCs) appeared in the 1970s (the Xerox Alto, the MITS Altair 8800, and the Apple I and II, to name a few), these machines had only limited distribution to computer enthusiasts. The appearance of the IBM PC in 1981 is usually considered the beginning of the PC era because this machine was the first to be widely adopted by American businesses. At first using the DOS operating system, a text-based command language, and later the Microsoft Windows operating system, the **Wintel PC** computer (Windows operating system software on a computer with an Intel microprocessor) became the standard desktop personal computer. In 2014, there are an estimated 2 billion PCs in the world, and about 276 million new PCs were sold in 2014 (Gartner, 2014). Approximately 90 percent are thought to run a version of Windows, and 10 percent run a Macintosh OS. The Wintel dominance as a computing platform is receding as iPhone and Android device sales increase. About 1.75 billion people worldwide own smartphones, and most of these users access the Internet with their mobile devices.

Proliferation of PCs in the 1980s and early 1990s launched a spate of personal desktop productivity software tools—word processors, spreadsheets, electronic presentation software, and small data management programs—that were very valuable to both home and corporate users. These PCs were stand-alone systems until PC operating system software in the 1990s made it possible to link them into networks.

Client/Server Era (1983 to Present)

In **client/server computing**, desktop or laptop computers called **clients** are networked to powerful **server** computers that provide the client computers with a variety of services and capabilities. Computer processing work is split between these two types of machines. The client is the user point of entry, whereas the server typically processes and stores shared data, serves up Web pages, or manages network activities. The term "server" refers to both the software application and the physical computer on which the network software runs. The server could be a mainframe, but today, server computers typically are more powerful versions of personal computers, based on inexpensive chips and often using multiple processors in a single computer box., or in server racks.

The simplest client/server network consists of a client computer networked to a server computer, with processing split between the two types of machines. This is called a *two-tiered client/server architecture*. Whereas simple client/server

networks can be found in small businesses, most corporations have more complex, **multitiered** (often called **N-tier**) **client/server architectures** in which the work of the entire network is balanced over several different levels of servers, depending on the kind of service being requested (see Figure 5.3).

For instance, at the first level, a **Web server** will serve a Web page to a client in response to a request for service. Web server software is responsible for locating and managing stored Web pages. If the client requests access to a corporate system (a product list or price information, for instance), the request is passed along to an **application server**. Application server software handles all application operations between a user and an organization's back-end business systems. The application server may reside on the same computer as the Web server or on its own dedicated computer. Chapters 6 and 7 provide more detail on other pieces of software that are used in multitiered client/server architectures for e-commerce and e-business.

Client/server computing enables businesses to distribute computing work across a series of smaller, inexpensive machines that cost much less than centralized mainframe systems. The result is an explosion in computing power and applications throughout the firm.

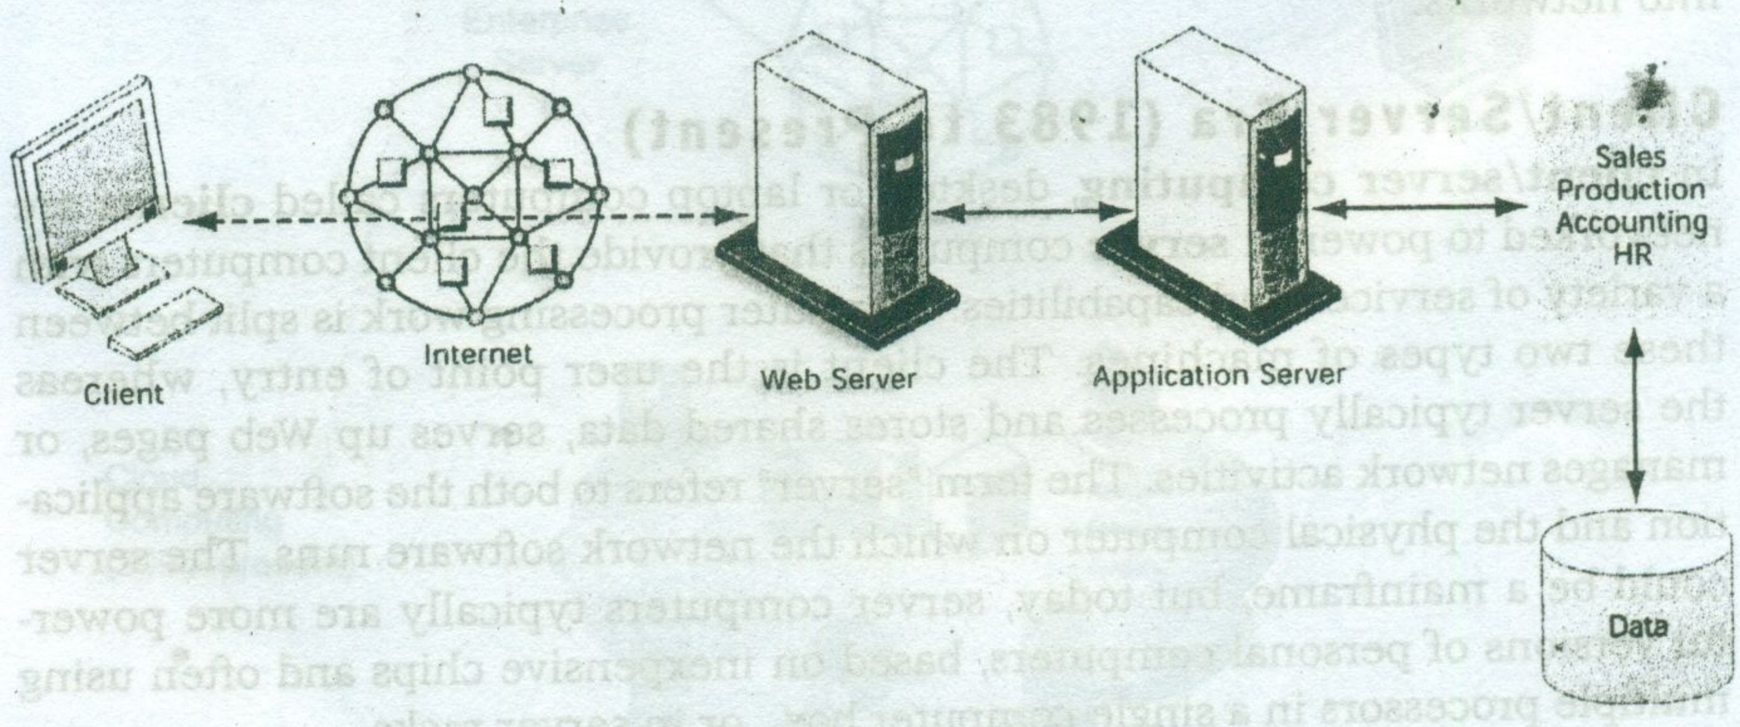
Novell NetWare was the leading technology for client/server networking at the beginning of the client/server era. Today, Microsoft is the market leader with its **Windows** operating systems (Windows Server, Windows 8, Windows 7, and Windows Vista).

Enterprise Computing Era (1992 to Present)

In the early 1990s, firms turned to networking standards and software tools that could integrate disparate networks and applications throughout the firm into an enterprise-wide infrastructure. As the Internet developed into a trusted communications environment after 1995, business firms began seriously using the *Transmission Control Protocol/Internet Protocol (TCP/IP)* networking standard to tie their disparate networks together. We discuss TCP/IP in detail in Chapter 7.

The resulting IT infrastructure links different pieces of computer hardware and smaller networks into an enterprise-wide network so that information can flow freely across the organization and between the firm and other

FIGURE 5.3 A MULTITIERED CLIENT/SERVER NETWORK (N-TIER)



In a multitiered client/server network, client requests for service are handled by different levels of servers.

organizations. It can link different types of computer hardware, including mainframes, servers, PCs, and mobile devices, and it includes public infrastructures such as the telephone system, the Internet, and public network services. The enterprise infrastructure also requires software to link disparate applications and enable data to flow freely among different parts of the business, such as enterprise applications (see Chapters 2 and 9) and Web services (discussed in Section 5.4).

Cloud and Mobile Computing Era (2000 to Present)

The growing bandwidth power of the Internet has pushed the client/server model one step further, towards what is called the "Cloud Computing Model." **Cloud computing** refers to a model of computing that provides access to a shared pool of computing resources (computers, storage, applications, and services) over a network, often the Internet. These "clouds" of computing resources can be accessed on an as-needed basis from any connected device and location. Currently, cloud computing is the fastest growing form of computing, with companies spending about $175 billion on cloud infrastructure and services in 2014 (Hamilton, 2014).

Thousands or even hundreds of thousands computers are located in cloud data centers, where they can be accessed by desktop computers, laptop computers, tablets, entertainment centers, smartphones, and other client machines linked to the Internet, with both personal and corporate computing increasingly moving to mobile platforms. Amazon, Google, IBM, and Microsoft operate huge, scalable cloud computing centers that provide computing power, data storage, and high-speed Internet connections to firms that want to maintain their IT infrastructures remotely.Firms such as Google, Microsoft, SAP, Oracle, and Salesforce.com sell software applications as services delivered over the Internet.

We discuss cloud and mobile computing in more detail in Section 5.3. The Learning Tracks include a table titled Comparing Stages in IT Infrastructure Evolution, which compares each era on the infrastructure dimensions introduced.

TECHNOLOGY DRIVERS OF INFRASTRUCTURE EVOLUTION

The changes in IT infrastructure we have just described have resulted from developments in computer processing, memory chips, storage devices, telecommunications and networking hardware and software, and software design that have exponentially increased computing power while exponentially reducing costs. Let's look at the most important developments.

Moore's Law and Microprocessing Power

In 1965, Gordon Moore, the director of Fairchild Semiconductor's Research and Development Laboratories, an early manufacturer of integrated circuits, wrote in *Electronics* magazine that since the first microprocessor chip was introduced in 1959, the number of components on a chip with the smallest manufacturing costs per component (generally transistors) had doubled each year. This assertion became the foundation of **Moore's Law**. Moore later reduced the rate of growth to a doubling every two years

This law would later be interpreted in multiple ways. There are at least three variations of Moore's Law, none of which Moore ever stated: (1) the power of microprocessors doubles every 18 months; (2) computing power doubles every 18 months; and (3) the price of computing falls by half every 18 months.

Figure 5.4 illustrates the relationship between number of transistors on a microprocessor and millions of instructions per second (MIPS), a common measure of processor power. Figure 5.5 shows the exponential decline in the cost of transistors and rise in computing power. For instance, in 2014, you can buy an Intel i7 quad-core processor chip with 2.5 billion transistors for about one ten-millionth of a dollar per transistor.

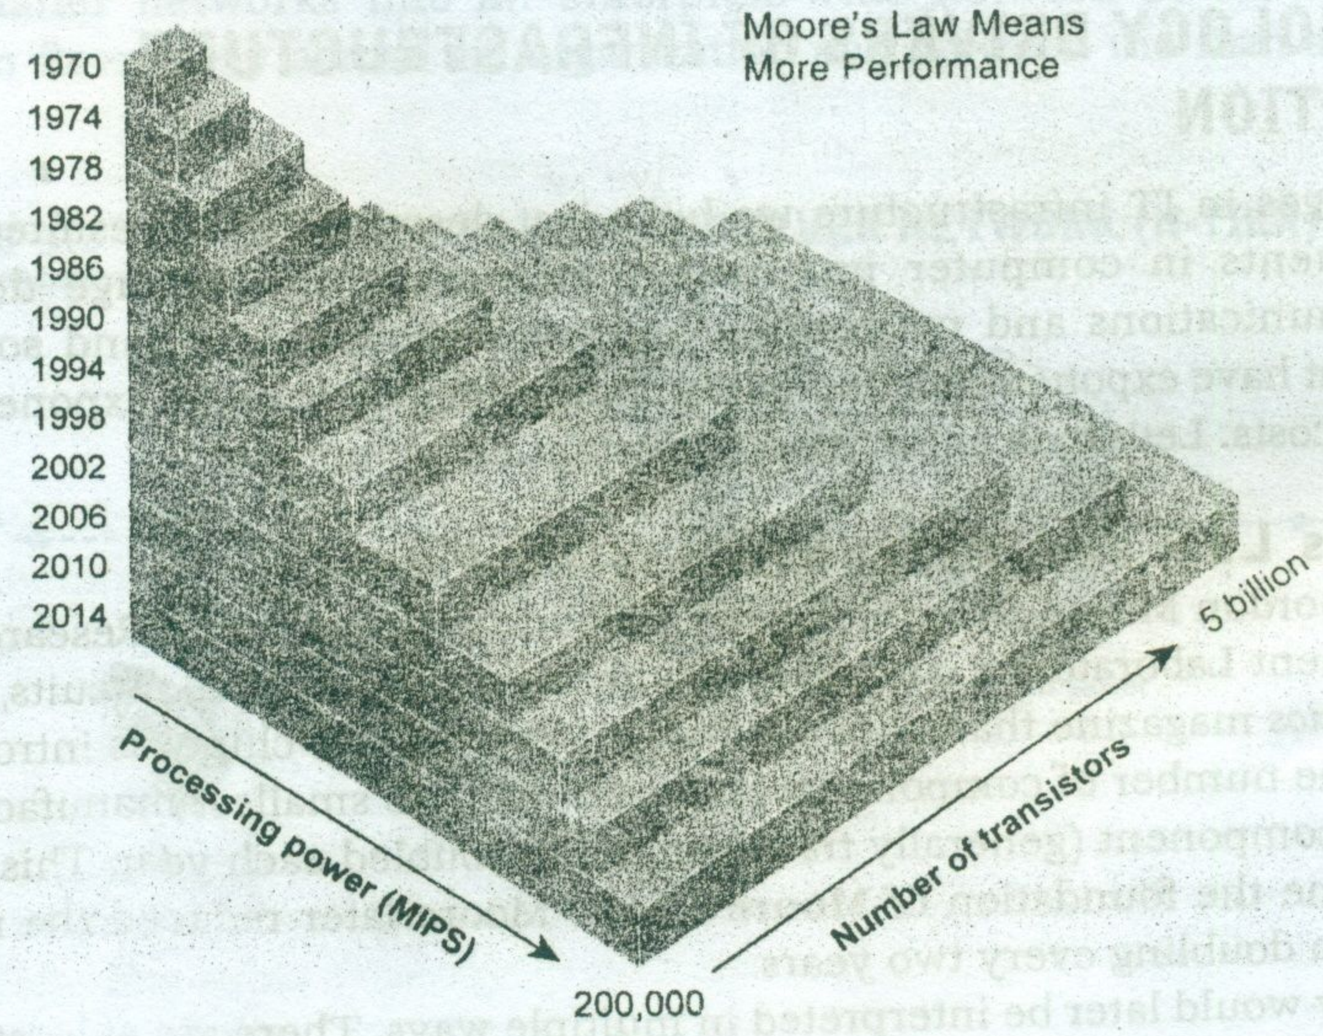
Exponential growth in the number of transistors and the power of processors coupled with an exponential decline in computing costs is likely to continue. Chip manufacturers continue to miniaturize components. Today's transistors should no longer be compared to the size of a human hair but rather to the size of a virus.

By using nanotechnology, chip manufacturers can even shrink the size of transistors down to the width of several atoms. **Nanotechnology** uses individual atoms and molecules to create computer chips and other devices that are thousands of times smaller than current technologies permit. Chip manufacturers are trying to develop a manufacturing process to produce nanotube processors economically (Figure 5.6). Stanford University scientists have built a nanotube computer.

The Law of Mass Digital Storage

A second technology driver of IT infrastructure change is the Law of Mass Digital Storage. The amount of digital information is roughly doubling every year (Gantz and Reinsel, 2011; Lyman and Varian, 2003). Fortunately, the cost of storing digital information is falling at an exponential rate of 100 percent a year. Figure 5.7 shows that the number of megabytes that can be stored on magnetic media for $1 from 1950 to the present roughly doubled every 15 months. In 2014, a 500 gigabyte hard disk drive sells at retail for about $60.

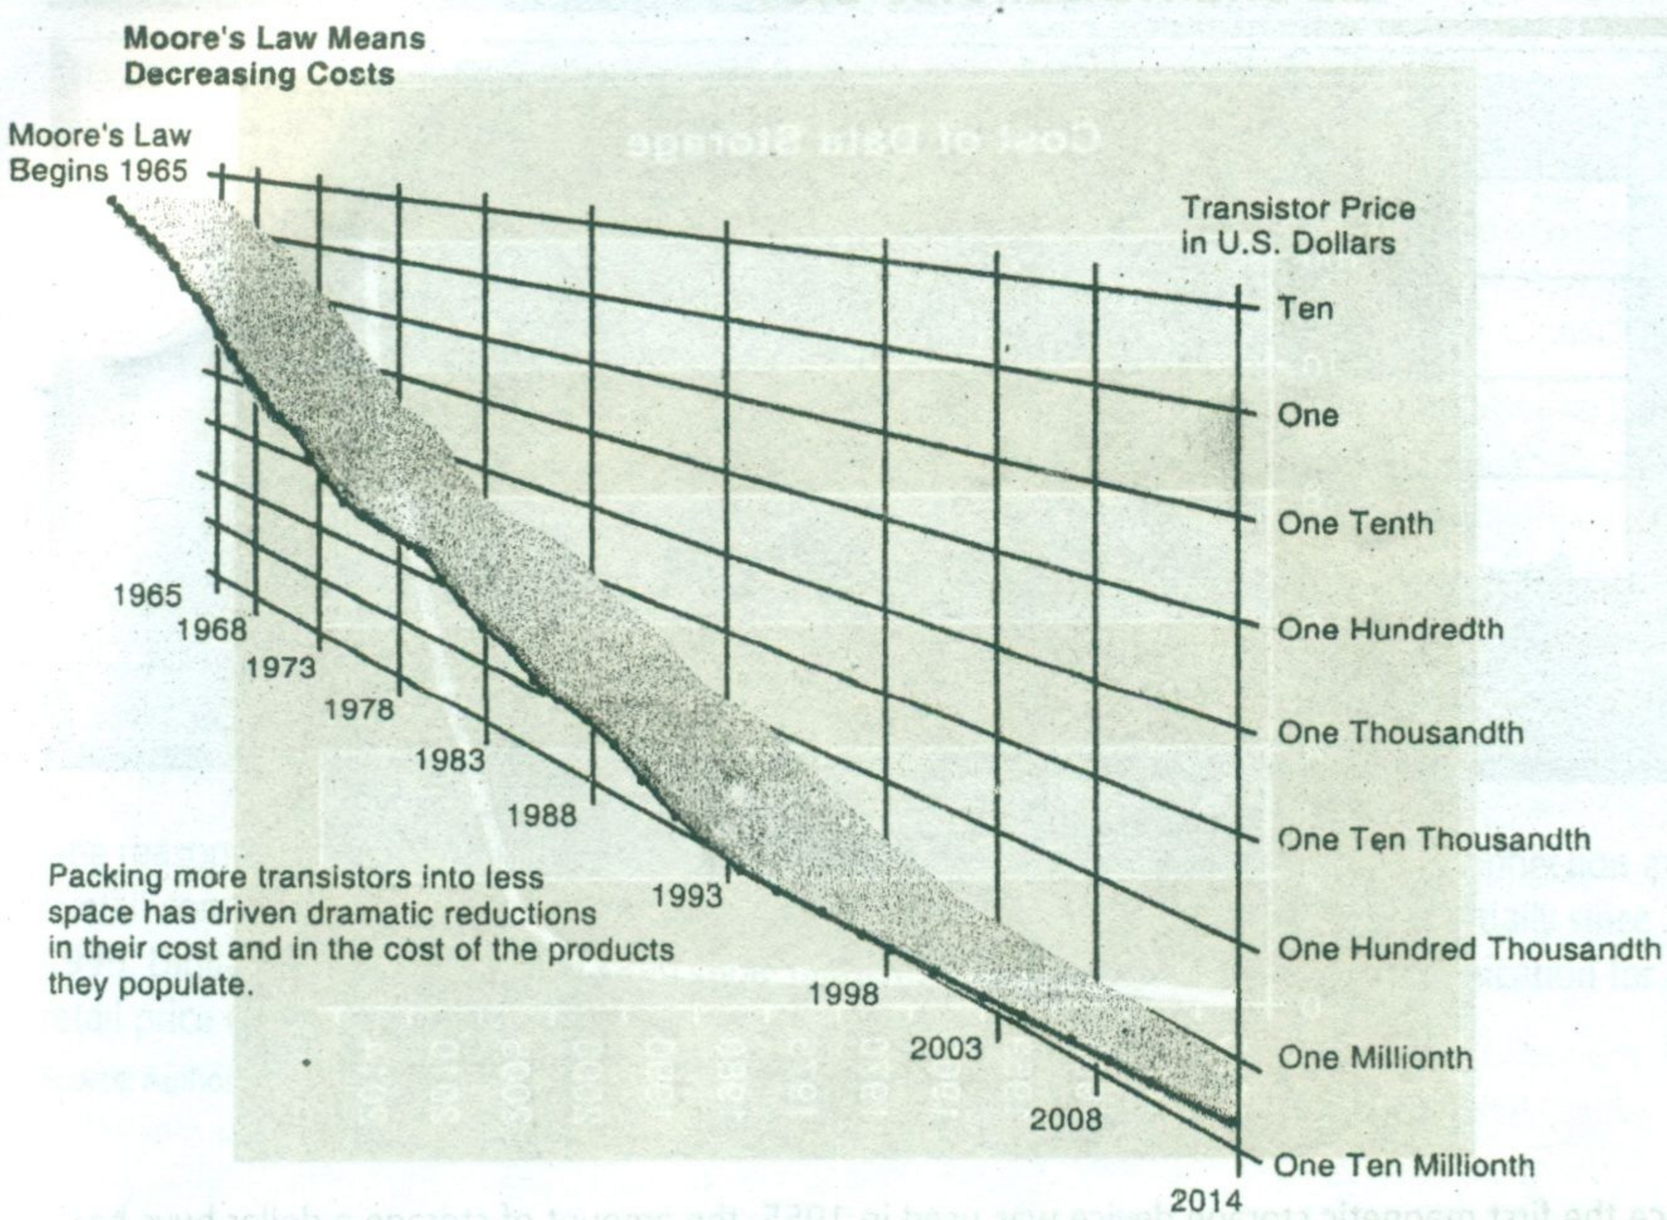
FIGURE 5.4 MOORE'S LAW AND MICROPROCESSOR PERFORMANCE



Packing over 5 billion transistors into a tiny microprocessor has exponentially increased processing power. Processing power has increased to over 200,000 MIPS (2.6 billion instructions per second).

Source: Authors' estimate.

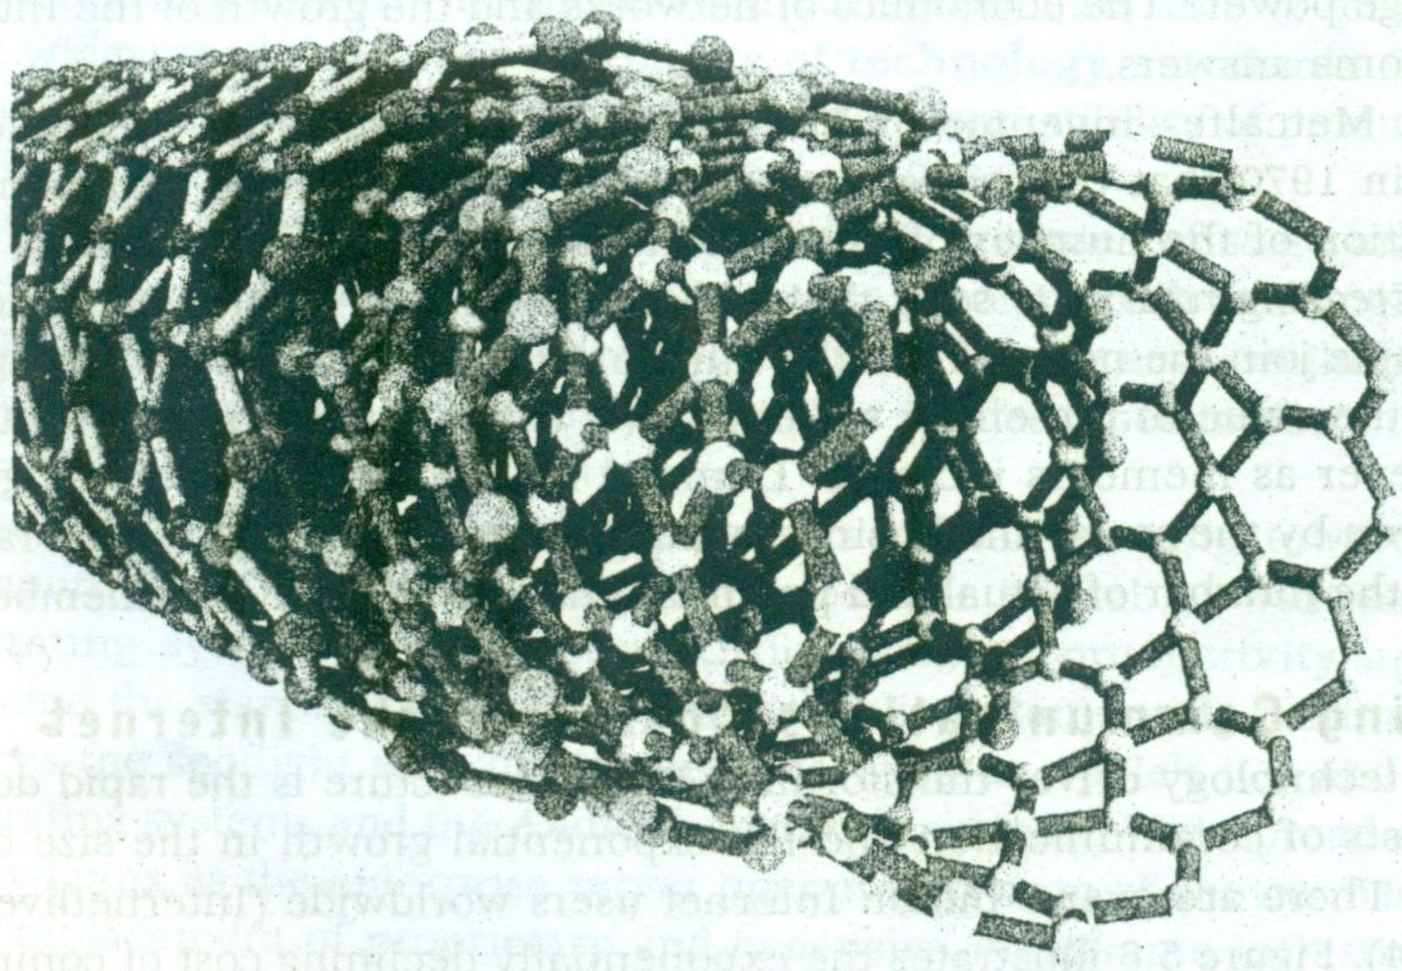
FIGURE 5.5 FALLING COST OF CHIPS



acking more transistors into less space has driven down transistor costs dramatically as well as the ost of the products in which they are used.

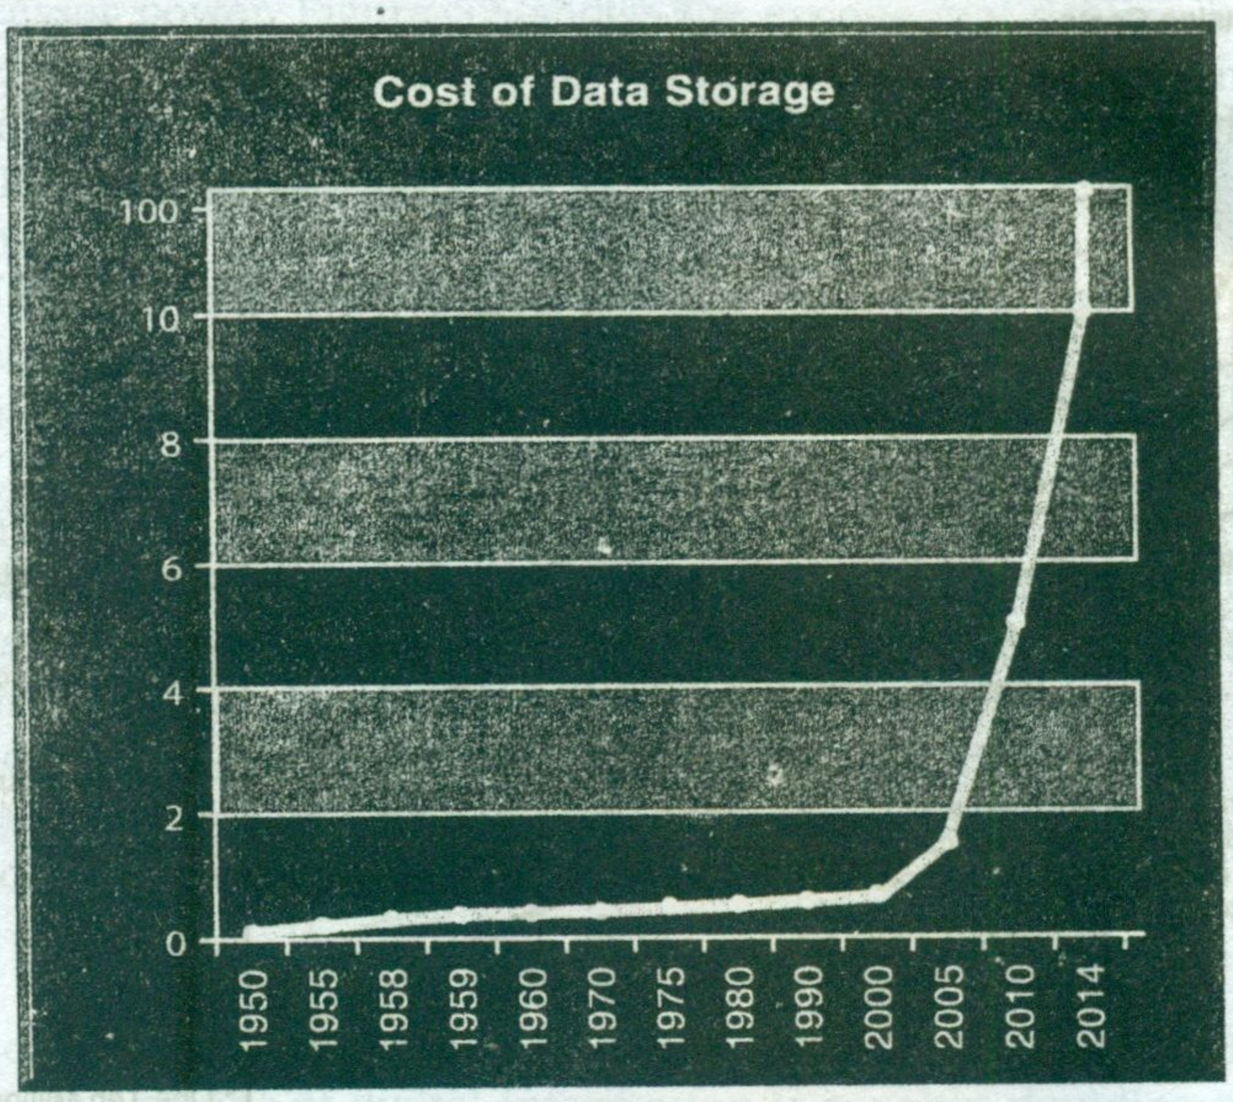
ource: Authors' estimate.

FIGURE 5.6 EXAMPLE OF NANOTUBES



Nanotubes are tiny tubes about 10,000 times thinner than a human hair. They consist of rolled-up sheets of carbon hexagons and have the potential uses as minuscule wires or in ultrasmall electronic devices and are very powerful conductors of ectrical current. © Tyler Boyes/Shutterstock.

FIGURE 5.7 THE AMOUNT OF STORAGE PER DOLLAR RISES EXPONENTIALLY, 1950–2014



Since the first magnetic storage device was used in 1955, the amount of storage a dollar buys has risen exponentially, doubling the amount of digital storage for each dollar expended every 15 months on average. Cloud storage services provide 100 gigabytes of storage for about $1.00

Source: Authors' estimates.

Metcalfe's Law and Network Economics

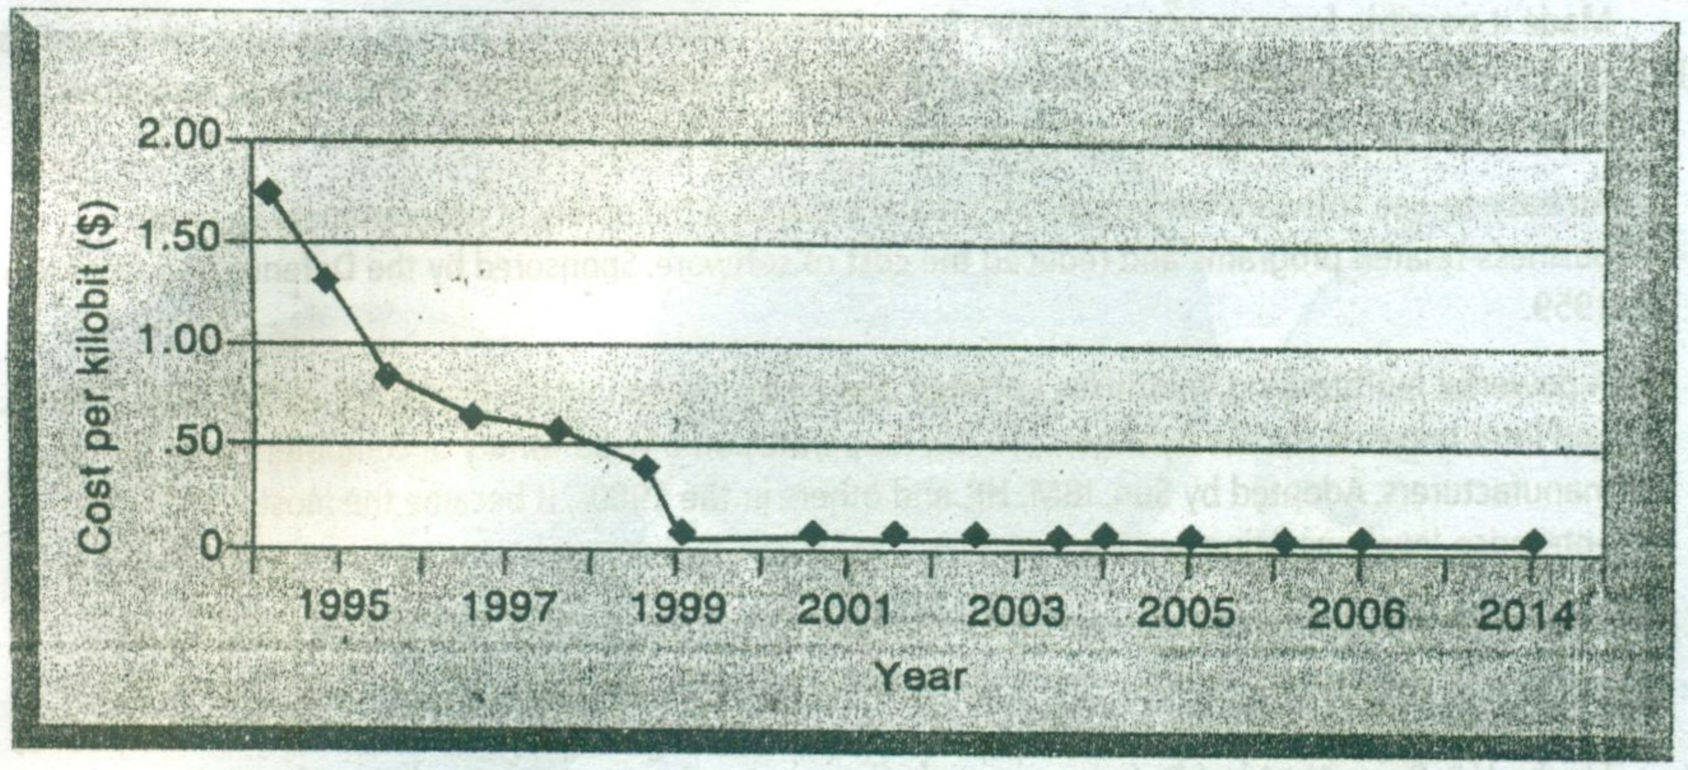
Moore's Law and the Law of Mass Storage help us understand why computing resources are now so readily available. But why do people want more computing and storage power? The economics of networks and the growth of the Internet provide some answers.

Robert Metcalfe—inventor of Ethernet local area network technology—claimed in 1970 that the value or power of a network grows exponentially as a function of the number of network members. Metcalfe and others point to the *increasing returns to scale* that network members receive as more and more people join the network. As the number of members in a network grows linearly, the value of the entire system grows exponentially and continues to grow forever as members increase. Demand for information technology has been driven by the social and business value of digital networks, which rapidly multiply the number of actual and potential links among network members.

Declining Communications Costs and the Internet

A fourth technology driver transforming IT infrastructure is the rapid decline in the costs of communication and the exponential growth in the size of the Internet. There are over 3 billion Internet users worldwide (Internetlivestats.com, 2014). Figure 5.8 illustrates the exponentially declining cost of communication both over the Internet and over telephone networks (which increasingly are based on the Internet). As communication costs fall toward a very small number and approach 0, utilization of communication and computing facilities explode

FIGURE 5.8 EXPONENTIAL DECLINES IN INTERNET COMMUNICATIONS COSTS



One reason for the growth in the Internet population is the rapid decline in Internet connection and overall communication costs. The cost per kilobit of Internet access has fallen exponentially since 1995. Digital subscriber line (DSL) and cable modems now deliver a kilobit of communication for a retail price of less than one penny.

Source: Authors.

To take advantage of the business value associated with the Internet, firms must greatly expand their Internet connections, including wireless connectivity, and greatly expand the power of their client/server networks, desktop clients, and mobile computing devices. There is every reason to believe these trends will continue.

Standards and Network Effects

Today's enterprise infrastructure and Internet computing would be impossible—both now and in the future—without agreements among manufacturers and widespread consumer acceptance of **technology standards**. Technology standards are specifications that establish the compatibility of products and the ability to communicate in a network (Stango, 2004).

Technology standards unleash powerful economies of scale and result in price declines as manufacturers focus on the products built to a single standard. Without these economies of scale, computing of any sort would be far more expensive than is currently the case. Table 5.1 describes important standards that have shaped IT infrastructure.

Beginning in the 1990s, corporations started moving toward standard computing and communications platforms. The Wintel PC with the Windows operating system and Microsoft Office desktop productivity applications became the standard desktop and mobile client computing platform. (It now shares the spotlight with other standards, such as Apple's iOS and Macintosh operating systems and the Android operating system.) Widespread adoption of Unix-Linux as the enterprise server operating system of choice made possible the replacement of proprietary and expensive mainframe infrastructures. In telecommunications, the Ethernet standard enabled PCs to connect together in small local area networks (LANs; see Chapter 7), and the TCP/IP standard enabled these LANs to be connected into firmwide networks, and ultimately, to the Internet.

TABLE 5.1 SOME IMPORTANT STANDARDS IN COMPUTING

STANDARD	SIGNIFICANCE
American Standard Code for Information Interchange (ASCII) (1958)	Made it possible for computer machines from different manufacturers to exchange data; later used as the universal language linking input and output devices such as keyboards and mice to computers. Adopted by the American National Standards Institute in 1963.
Common Business Oriented Language (COBOL) (1959)	An easy-to-use software language that greatly expanded the ability of programmers to write business-related programs and reduced the cost of software. Sponsored by the Defense Department in 1959.
Unix (1969–1975)	A powerful multitasking, multiuser, portable operating system initially developed at Bell Labs (1969) and later released for use by others (1975). It operates on a wide variety of computers from different manufacturers. Adopted by Sun, IBM, HP, and others in the 1980s, it became the most widely used enterprise-level operating system.
Transmission Control Protocol/Internet Protocol (TCP/IP) (1974)	Suite of communications protocols and a common addressing scheme that enables millions of computers to connect together in one giant global network (the Internet). Later, it was used as the default networking protocol suite for local area networks and intranets. Developed in the early 1970s for the U.S. Department of Defense.
Ethernet (1973)	A network standard for connecting desktop computers into local area networks that enabled the widespread adoption of client/server computing and local area networks, and further stimulated the adoption of personal computers.
IBM/Microsoft/Intel Personal Computer (1981)	The standard Wintel design for personal desktop computing based on standard Intel processors and other standard devices, Microsoft DOS, and later Windows software. The emergence of this standard, low-cost product laid the foundation for a 25-year period of explosive growth in computing throughout all organizations around the globe. Today, more than 1 billion PCs power business and government activities every day.
World Wide Web (1989–1993)	Standards for storing, retrieving, formatting, and displaying information as a worldwide web of electronic pages incorporating text, graphics, audio, and video enables creation of a global repository of billions of Web pages.

5.2 WHAT ARE THE COMPONENTS OF IT INFRASTRUCTURE?

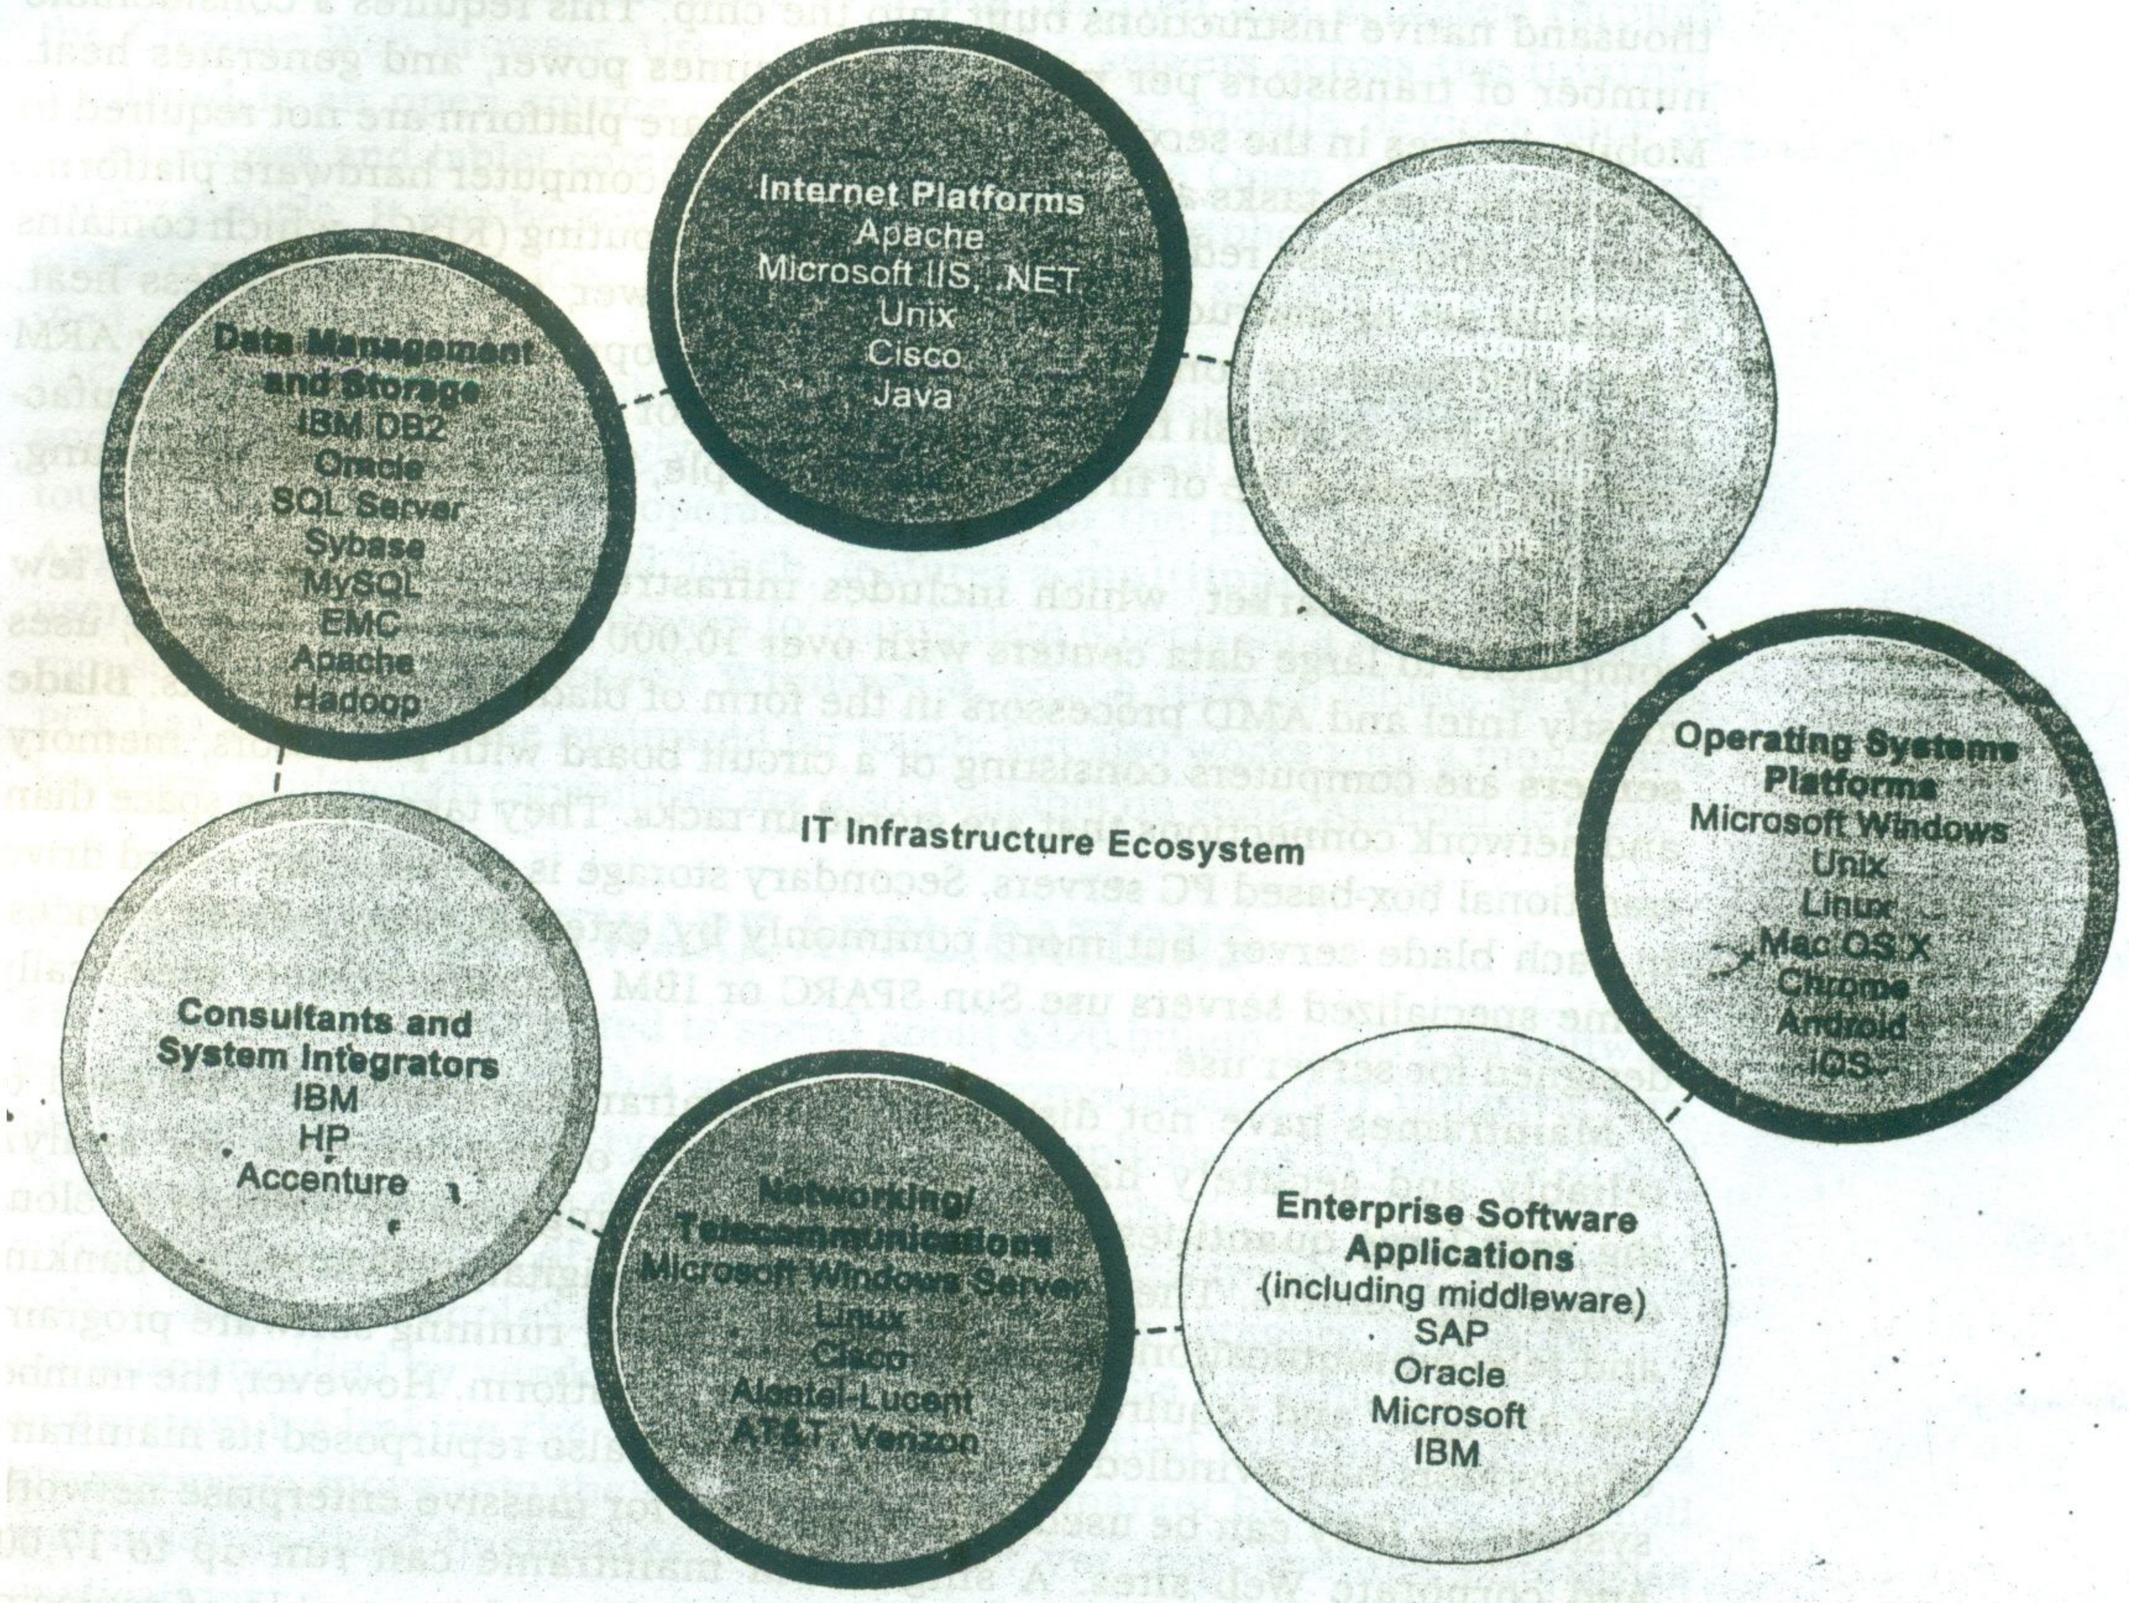
IT infrastructure today is composed of seven major components. Figure 5.9 illustrates these infrastructure components and the major vendors within each component category. These components constitute investments that must be coordinated with one another to provide the firm with a coherent infrastructure.

In the past, technology vendors supplying these components were often in competition with one another, offering purchasing firms a mixture of incompatible, proprietary, partial solutions. But increasingly the vendor firms have been forced by large customers to cooperate in strategic partnerships with one another. For instance, a hardware and services provider such as IBM cooperates with all the major enterprise software providers, has strategic relationships with system integrators, and promises to work with whichever database products its client firms wish to use (even though it sells its own database management software called DB2).

COMPUTER HARDWARE PLATFORMS

Firms worldwide are expected to spend $669 billion on computer hardware devices in 2014, including mainframes, servers, PCs, tablets, and smartphones.

FIGURE 5.9 THE IT INFRASTRUCTURE ECOSYSTEM



There are seven major components that must be coordinated to provide the firm with a coherent IT infrastructure. Listed here are major technologies and suppliers for each component.

You can think of all these computers and their processors as the computer hardware platform for corporate (and personal) computing worldwide.

There are approximately 2 billion PCs in the world today, 2,000 government data centers, and an estimated 8,000 corporate data centers and cloud computing centers. Nearly all their computing takes place using microprocessor "chips" manufactured or designed by Intel Corporation, and to a lesser extent, AMD Corporation. Intel and AMD processors are often referred to as "i86" processors because the original IBM PCs used an Intel 8086 processor, and all the Intel (and AMD) chips that followed are downward compatible with this processor. (For instance, you should be able to run a software application designed ten years ago on a new PC computer you bought yesterday.) Without this commonality in i86 processors, it is unlikely that today's installed base of 2 billion PCs would exist.

The computer platform has changed dramatically in the last decade, with the introduction of mobile computing devices, from the iPod in 2001, to the iPhone in 2007, and the iPad in 2010. Worldwide, 1.7 billion people use smartphones. You can think of these devices as a second computer hardware platform, one that is consumer device-driven.

The computers with Intel microprocessors in the first computer hardware platform use complex instruction set computing (CISC), with several thousand native instructions built into the chip. This requires a considerable number of transistors per processor, consumes power, and generates heat. Mobile devices in the second computer hardware platform are not required to perform as many tasks as computers in the first computer hardware platform. They are able to use reduced instruction set computing (RISC), which contains a smaller set of instructions, consumes less power, and generates less heat. Apple and Samsung consumer devices use microprocessors designed by ARM Holdings, Inc., a British firm. RISC processors for mobile devices are manufactured by a wide range of firms, including Apple, Texas Instruments, Samsung, and Qualcomm.

The server market, which includes infrastructures ranging from a few computers to large data centers with over 10,000 individual computers, uses mostly Intel and AMD processors in the form of blade servers on racks. **Blade servers** are computers consisting of a circuit board with processors, memory and network connections that are stored in racks. They take up less space than traditional box-based PC servers. Secondary storage is provided by a hard drive in each blade server, but more commonly by external mass-storage devices. Some specialized servers use Sun SPARC or IBM microprocessors specifically designed for server use.

Mainframes have not disappeared. Mainframes continue to be used to reliably and securely handle huge volumes of transactions, for analyzing very large quantities of data, and for handling large workloads in cloud computing centers. The mainframe is still the digital workhorse for banking and telecommunications networks that are often running software programs that are older and require a specific hardware platform. However, the number of providers has dwindled to one:IBM. IBM has also repurposed its mainframe systems so they can be used as giant servers for massive enterprise networks and corporate Web sites. A single IBM mainframe can run up to 17,000 instances of Linux or Windows Server software and is capable of replacing thousands of smaller blade servers (see the discussion of virtualization in Section 5.3).

OPERATING SYSTEM PLATFORMS

Microsoft Windows Server comprises about 35 percent of the server operating system market, with 65 percent of corporate servers using some form of the **Unix** operating system or **Linux**, an inexpensive and robust open source relative of Unix. Microsoft Windows Server is capable of providing enterprise-wide operating system and network services, and appeals to organizations seeking Windows-based IT infrastructures.

Unix and Linux are scalable, reliable, and much less expensive than mainframe operating systems. They can also run on many different types of processors. The major providers of Unix operating systems are IBM, HP, and Sun, each with slightly different and partially incompatible versions.

At the client level, 90 percent of PCs use some form of the Microsoft Windows **operating system** (such as Windows 8, Windows 7, or Windows Vista) to manage the resources and activities of the computer. However, there is now a much greater variety of operating systems than in the past, with new operating systems for computing on handheld mobile digital devices or cloud-connected computers.

Google's **Chrome OS** provides a lightweight operating system for cloud computing using a Web-connected computer. Programs are not stored on the user's computer but are used over the Internet and accessed through the Chrome Web browser. User data reside on servers across the Internet. **Android** is an open source operating system for mobile devices such as smartphones and tablet computers developed by the Open Handset Alliance led by Google. It has become the most popular smartphone platform worldwide, competing with iOS, Apple's mobile operating system for the iPhone, iPad, and iPod Touch.

Conventional client operating system software is designed around the mouse and keyboard, but increasingly becoming more natural and intuitive by using touch technology. **iOS**, the operating system for the phenomenally popular Apple iPad, iPhone, and iPod Touch, features a **multitouch** interface, where users employ one or more fingers to manipulate objects on a screen without a mouse or keyboard. Microsoft's **Windows 8,** which runs on tablets as well as PCs, has a user interface optimized for touch, but also works with a mouse and keyboard. Multitouch capabilities are also available on some Android devices.

ENTERPRISE SOFTWARE APPLICATIONS

Firms worldwide are expected to spend about $320 billion in 2014 on software for enterprise applications that are treated as components of IT infrastructure. We introduced the various types of enterprise applications in Chapter 2, and Chapter 9 provides a more detailed discussion of each.

The largest providers of enterprise application software are SAP and Oracle (which acquired PeopleSoft). Also included in this category is middleware software supplied by vendors such as IBM and Oracle for achieving firmwide integration by linking the firm's existing application systems. Microsoft is attempting to move into the lower ends of this market by focusing on small and medium-sized businesses that have not yet implemented enterprise applications.

DATA MANAGEMENT AND STORAGE

Enterprise database management software is responsible for organizing and managing the firm's data so that they can be efficiently accessed and used. Chapter 6 describes this software in detail. The leading database software providers are IBM (DB2), Oracle, Microsoft (SQL Server), and Sybase (Adaptive Server Enterprise), which supply more than 90 percent of the U.S. database software marketplace. MySQL is a Linux open source relational database product now owned by Oracle Corporation, and Apache Hadoop is an open source software framework for managing massive data sets (see Chapter 6).

The physical data storage market is dominated by EMC Corporation for large-scale systems, and a small number of PC hard disk manufacturers led by Seagate and Western Digital.

Digital information is doubling every two years and the market for digital data storage devices has been growing at more than 15 percent annually over the last five years. In addition to traditional disk arrays and tape libraries, large firms are turning to network-based storage technologies. **Storage area networks (SANs)** connect multiple storage devices on a separate high-speed network dedicated to storage. The SAN creates a large central pool of storage that can be rapidly accessed and shared by multiple servers.

NETWORKING/TELECOMMUNICATIONS PLATFORMS

Companies worldwide are expected to spend $1.65 trillion for telecommunications services in 2014 (Gartner, 2014). Chapter 7 is devoted to an in-depth description of the enterprise networking environment, including the Internet. Windows Server is predominantly used as a local area network operating system, followed by Linux and Unix. Large, enterprise wide area networks use some variant of Unix. Most local area networks, as well as wide area enterprise networks, use the TCP/IP protocol suite as a standard (see Chapter 7).

The leading networking hardware providers are Cisco, Alcatel-Lucent, and Juniper Networks. Telecommunications platforms are typically provided by telecommunications/telephone services companies that offer voice and data connectivity, wide area networking, wireless services, and Internet access. Leading telecommunications service vendors include AT&T and Verizon. This market is exploding with new providers of cellular wireless, high-speed Internet, and Internet telephone services.

INTERNET PLATFORMS

Internet platforms overlap with, and must relate to, the firm's general networking infrastructure and hardware and software platforms. They include hardware, software, and management services to support a firm's Web site, including Web hosting services, routers, and cabling or wireless equipment. A **Web hosting service** maintains a large Web server, or series of servers, and provides fee-paying subscribers with space to maintain their Web sites.

The Internet revolution created a veritable explosion in server computers, with many firms collecting thousands of small servers to run their Internet operations. Since then there has been a steady push toward server consolidation, reducing the number of server computers by increasing the size and power of each and by using software tools that make it possible to run more applications on a single server. The Internet hardware server market has become increasingly concentrated in the hands of IBM, Dell, Sun (Oracle), and HP, as prices have fallen dramatically.

The major Web software application development tools and suites are supplied by Microsoft (Microsoft Visual Studio and the Microsoft .NET family of development tools), Oracle-Sun (Sun's Java is the most widely used tool for developing interactive Web applications on both the server and client sides), and a host of independent software developers, including Adobe (Creative Suite) and Real Networks (media software). Chapter 7 describes the components of the firm's Internet platform in greater detail.

CONSULTING AND SYSTEM INTEGRATION SERVICES

Today, even a large firm does not have the staff, the skills, the budget, or the necessary experience to deploy and maintain its entire IT infrastructure. Implementing a new infrastructure requires (as noted in Chapters 3 and 14) significant changes in business processes and procedures, training and education, and software integration. Leading consulting firms providing this expertise include Accenture, IBM Global Services, HP, Infosys, and Wipro Technologies.

Software integration means ensuring the new infrastructure works with the firm's older, so-called legacy systems and ensuring the new elements of the infrastructure work with one another. **Legacy systems** are

generally older transaction processing systems created for mainframe computers that continue to be used to avoid the high cost of replacing or redesigning them. Replacing these systems is cost-prohibitive and generally not necessary if these older systems can be integrated into a contemporary infrastructure.

5.3 WHAT ARE THE CURRENT TRENDS IN COMPUTER HARDWARE PLATFORMS?

The exploding power of computer hardware and networking technology has dramatically changed how businesses organize their computing power, putting more of this power on networks and mobile handheld devices. We look at seven hardware trends: the mobile digital platform, consumerization of IT, quantum computing, virtualization, cloud computing, green computing (as described in the Interactive Session on Technology), and high-performance/power-saving processors.

THE MOBILE DIGITAL PLATFORM

Chapter 1 pointed out that new mobile digital computing platforms have emerged as alternatives to PCs and larger computers. Smartphones such as the iPhone, Android, and BlackBerry smartphones have taken on many functions of PCs, including transmission of data, surfing the Web, transmitting e-mail and instant messages, displaying digital content, and exchanging data with internal corporate systems. The new mobile platform also includes small, lightweight netbooks optimized for wireless communication and Internet access, **tablet computers** such as the iPad, and digital e-book readers such as Amazon's Kindle with Web access capabilities.

Smartphones and tablet computers are becoming an important means of accessing the Internet. These devices are increasingly used for business computing as well as for consumer applications. For example, senior executives at General Motors are using smartphone applications that drill down into vehicle sales information, financial performance, manufacturing metrics, and project management status. Wearable computing devices are another recent addition to the mobile digital platform. These include smartwatches, smart glasses, smart badges, and activity trackers. Wearable computing technology is still in its infancy, but it already has business uses.

CONSUMERIZATION OF IT AND BYOD

The popularity, ease of use, and rich array of useful applications for smartphones and tablet computers have created a groundswell of interest in allowing employees to use their personal mobile devices in the workplace, a phenomenon popularly called *"bring your own device" (BYOD)*. BYOD is one aspect of the **consumerization of IT**, in which new information technology that first emerges in the consumer market spreads into business organizations. Consumerization of IT includes not only mobile personal devices but also business uses of software services that originated in the consumer marketplace as well, such as Google and Yahoo search, Gmail, Google Apps, Dropbox (see Chapter 2), and even Facebook and Twitter.

Consumerization of IT is forcing businesses, especially large enterprises, to rethink the way they obtain and manage information technology equipment

INTERACTIVE SESSION: TECHNOLOGY

THE GREENING OF THE DATA CENTER

What's too hot to handle? It might very well be your company's data center, which can easily consume more than 100 times more power than a standard office building. Data-hungry tasks such as video on demand, maintaining Web sites, or analyzing large pools of transactions or social media data require more and more power-hungry machines. Power and cooling costs for data centers have skyrocketed, with cooling a server requiring roughly the same number of kilowatts of energy as running one. All this additional power consumption has a negative impact on the environment as well as corporate operating costs.

Companies are now looking to green computing for solutions. The standard for measuring data center energy efficiency is power usage effectiveness (PUE). This metric is a ratio of the total annual power consumed by a data center divided by how much is used annually by IT equipment. The lower the ratio, the better, with a PUE of 1.0 representing a desirable target. The PUE of traditional data centers has hovered around 2.0. That means the data center is using twice the amount of electricity that's actually needed to do the computing. (The extra power is consumed by lighting, cooling, and other systems.) PUE is influenced by many factors, including hardware efficiency, data center size, the types of servers and their uses, the proficiency of monitoring software, building architecture, and the climate outside the facility. New data center designs with PUEs of 1.5 or better are emerging

Virtualization is a highly effective tool for cost-effective green computing because it reduces the number of servers and storage resources in the firm's IT infrastructure. About five years ago, Acorda, a $210 million-per-year maker of drugs to treat nervous disorders such as multiple sclerosis, found it needed more servers and was outgrowing its data center. The company invested $100,000 in virtual servers running on technology from VMware. Using virtualization, Acorda avoided spending an additional $1.5 million on more physical servers and increasing energy consumption. Moreover, when the company moved to a new building in 2012, it was able to significantly shrink the size of its data center, further lowering cooling costs.

Acorda took additional steps to boost energy efficiency at the new facility. The company installed motion sensors that shut off lights after five minutes if no movement is detected, and it invested in a more intelligent cooling system that automatically changes settings as conditions change. Current plans include replacing all of Acorda's host servers and taking advantage of a VMware feature that moves virtual servers from one cluster to another, thereby reducing the number of clusters that require power at any given time. Acorda is also preparing to test virtualized desktops, which will greatly reduce the power required to run workstations and laptops.

Other tools and techniques are also available to make data centers more energy-efficient. Google and Microsoft have built data centers that take advantage of hydroelectric power. In April 2011 Facebook publicly posted the specifications for the design of its data centers, including motherboards, power supply, server chassis, server rack, and battery cabinets, as well as data center electrical and mechanical construction specifications. Facebook hardware engineers re-thought the electric design, power distribution, and thermal design of its servers to optimize energy efficiency, reducing power usage by 13 percent. The power supply, which converts alternating current into direct current consumed by the motherboard, operates at 94.5 percent efficiency. Instead of using air conditioning or air ducts, the servers are cooled by evaporative cooling and misting machines, which flow air through grill-covered walls. The server racks are taller to provide for bigger heat sinks, and the data center's large fans can move air through the servers more efficiently. Facebook's engineers modified the programming in the servers to work with these larger fans and reduce their reliance on small, individual fans that consume more power. This data center design, which has a 1.07 PUE rating, was implemented at Facebook's Prineville, Oregon data center. All of these changes have reduced Facebook's energy consumption per unit of computing power by 38 percent and operating costs by nearly 25 percent. The Prineville data center reports that its PUE is 1.07, one of the lowest.

By using ambient air cooling techniques and running warmer than average, Google's newest data centers deliver a PUE rating of 1.16. Yahoo's new Lockport, N.Y., data center has a PUE of 1.08. Lockport's cool climate, prevailing winds, and hydropower help cool Yahoo's 120 foot by 60 foot server buildings. FedEx located its energy-efficient Colorado

data center at an elevation of 6,000 feet so that the building can be cooled using outside air instead of internal air conditioning.

In addition to lowering IT costs, using cloud computing services may save energy as well. Cloud computing centers pack in servers that have been optimized for virtualization and for supporting as many different subscribing companies as possible. Cloud vendors are willing to invest heavily in cost-lowering virtualization software and energy-conserving server hardware because those efforts can produce major savings when doing the computing for large numbers of companies. A study by the Carbon Disclosure Project predicted that by 2020, large U.S. companies with revenues of more than $1 billion that used cloud computing would be able to achieve annual energy savings of $12.3 billion and annual carbon reductions equivalent to 200 million barrels of oil—enough to power 5.7 million cars for one year.

Experts note that it's important for companies to measure their energy use and inventory and track their information technology assets both before and after they start their green initiatives. And it isn't always necessary to purchase new technologies to achieve "green" goals. Organizations can achieve sizable efficiencies by better managing the computing resources they already have. Unfortunately, many information systems departments still aren't deploying their existing technology resources efficiently or using green measurement tools.

In 2015 and beyond, electricity demands to power data centers have never been higher, thanks to the continuing shift to the mobile platform. Experts predict there will be 15 billion mobile devices in use by year's end, and other estimates suggest there may be 50 billion by 2030. Meeting this considerable demand will require continued innovation in green computing to minimize environmental impact by any possible means.

Sources: Robert Bryce, "Green Computing Can't Power the Cloud," economics21.org, May 22, 2014; Tony Kontzer, "Energy Management Revamps the Data Center," *Baseline*, January 30, 2013; Charles Babcock, "Facebook's Data Center: Where Likes Live," *Information Week*, March 6 , 2013; Doug Mohney, "The Little Guys: Survival vs. Green," *Green Data Center News*, May 15, 2013; Chris Murphy, "FedEx's Strategic Tech Shift," *Information Week*, May 20, 2013; "How Facebook's Data Center Leads by Example," *CIO Insight*, August 20, 2012; Sam Greengard, "IT Gets Greener," *Baseline*, April 11, 2012; "New Study: Cloud Computing Can Dramatically Reduce Energy Costs and Carbon Emissions," AT&T, July 21, 2011; and Kenneth Miller, "The Data Center Balancing Act," *Information Week*, May 16, 2011.

CASE STUDY QUESTIONS

1. What business and social problems does data center power consumption cause?
2. What solutions are available for these problems? Are they management, organizational, or technology solutions? Explain your answer.
3. What are the business benefits and costs of these solutions?
4. Should all firms move toward green computing? Why or why not?

and services. Historically, at least in large firms, the IT department was responsible for selecting and managing the information technology and applications used by the firm and its employees. It furnished employees with desktops or laptops that were able to access corporate systems securely. The IT department maintained control over the firm's hardware and software to ensure that the business was being protected and that information systems served the purposes of the firm and its management. Today, employees and

business departments are playing a much larger role in technology selection, in many cases demanding that employees be able to use their own personal computers, smartphones, and tablets to access the corporate network. It is more difficult for the firm to manage and control these consumer technologies, and make sure they serve the needs of the business. The chapter-ending case study explores some of these management challenges created by BYOD and IT consumerization.

QUANTUM COMPUTING

Quantum computing is an emerging technology with the potential to dramatically boost computer processing power to find answers to problems that would take conventional computers many years to solve. **Quantum computing** uses the principles of quantum physics to represent data and perform operations on these data. A quantum computer would gain enormous processing power through the ability to be in many different states at once, allowing it to perform multiple operations simultaneously and solve some scientific and business problems millions of times faster than can be done today. Researchers at IBM, MIT, and the Los Alamos National Laboratory have been working on quantum computing, and the aerospace firm Lockheed Martin has purchased a quantum computer for commercial use.

VIRTUALIZATION

Virtualization is the process of presenting a set of computing resources (such as computing power or data storage) so that they can all be accessed in ways that are not restricted by physical configuration or geographic location. Virtualization enables a single physical resource (such as a server or a storage device) to appear to the user as multiple logical resources. For example, a server or mainframe can be configured to run many instances of an operating system (or different operating systems) so that it acts like many different machines. Each virtual server "looks" like a real physical server to software programs, and multiple virtual servers can run in parallel on a single machine. Virtualization also enables multiple physical resources (such as storage devices or servers) to appear as a single logical resource, as would be the case with storage area networks. VMware is the leading virtualization software vendor for Windows and Linux servers.

Server virtualization is a common method of reducing technology costs by providing the ability to host multiple systems on a single physical machine. Most servers run at just 15-20 percent of capacity, and virtualization can boost utilization server utilization rates to 70 percent or higher. Higher utilization rates translate into fewer computers required to process the same amount of work, reduced data center space to house machines, and lower energy usage. Virtualization also facilitates centralization and consolidation of hardware administration. It is now possible for companies and individuals to perform all of their computing work using a virtualized IT infrastructure in a remote location, as is the case with cloud computing.

CLOUD COMPUTING

Cloud computing is a model of computing in which computer processing, storage, software, and other services are provided as a pool of virtualized resources over a network, primarily the Internet. These "clouds" of computing

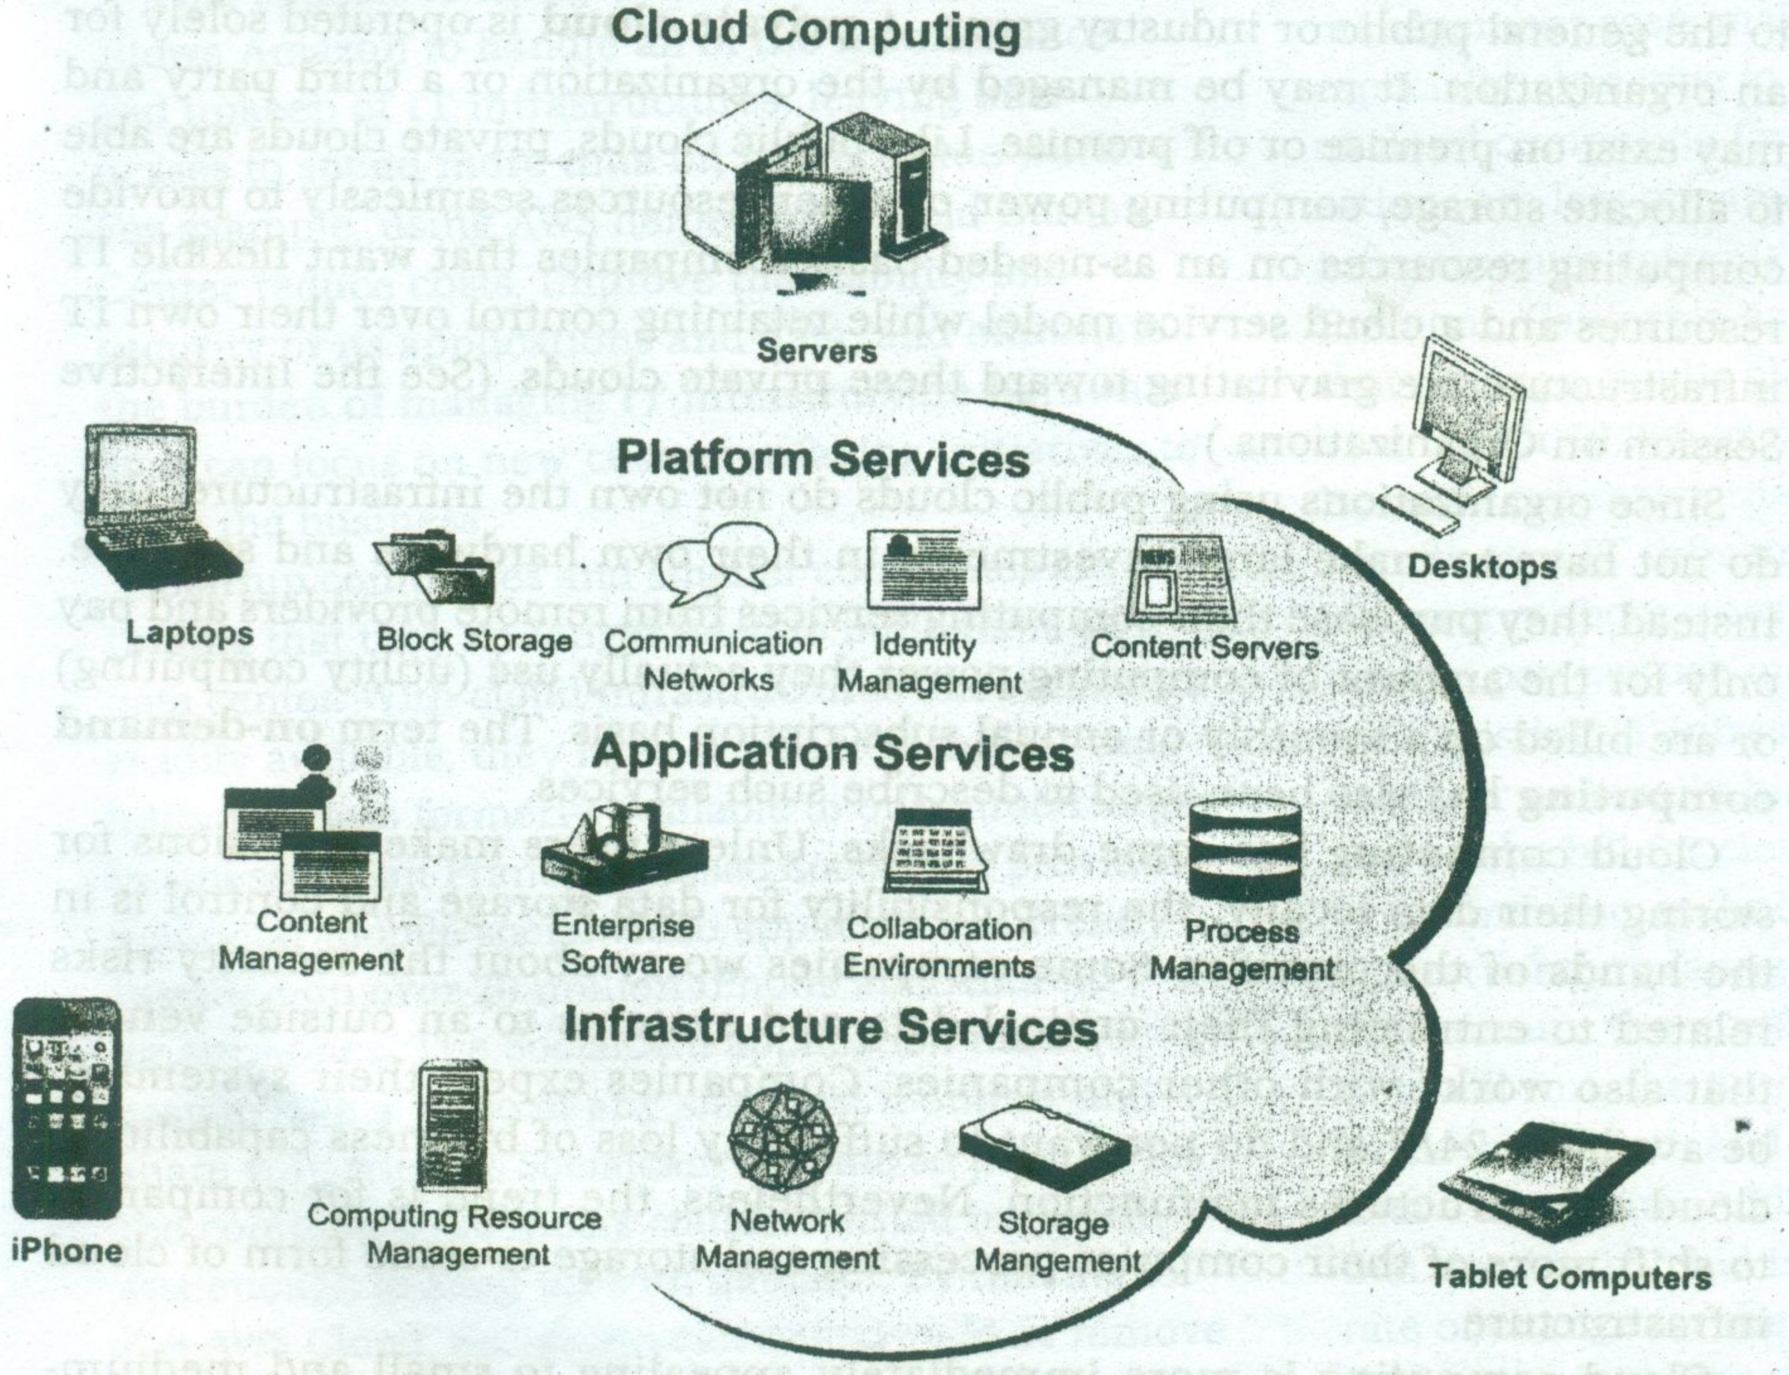
resources can be accessed on an as-needed basis from any connected device and location. Figure 5.10 illustrates the cloud computing concept.

The U.S. National Institute of Standards and Technology (NIST) defines cloud computing as having the following essential characteristics (Mell and Grance, 2009):

- **On-demand self-service:** Consumers can obtain computing capabilities such as server time or network storage as needed automatically on their own.
- **Ubiquitous network access:** Cloud resources can be accessed using standard network and Internet devices, including mobile platforms.
- **Location-independent resource pooling:** Computing resources are pooled to serve multiple users, with different virtual resources dynamically assigned according to user demand. The user generally does not know where the computing resources are located.
- **Rapid elasticity:** Computing resources can be rapidly provisioned, increased, or decreased to meet changing user demand.
- **Measured service:** Charges for cloud resources are based on amount of resources actually used.

Cloud computing consists of three different types of services:

FIGURE 5.10 CLOUD COMPUTING PLATFORM



In cloud computing, hardware and software capabilities are a pool of virtualized resources provided over a network, often the Internet. Businesses and employees have access to applications and IT infrastructure anywhere, at any time, and on any device.

- **Infrastructure as a Service (IaaS):**Customers use processing, storage, networking, and other computing resources from cloud service providers to run their information systems. For example, Amazon uses the spare capacity of its IT infrastructure to provide a broadly based cloud environment selling IT infrastructure services. These include its Simple Storage Service (S3) for storing customers' data and its Elastic Compute Cloud (EC2) service for running their applications. Users pay only for the amount of computing and storage capacity they actually use. (See the Interactive Session on Organizations).
- **Platform as a Service (PaaS):** Customers use infrastructure and programming tools supported by the cloud service provider to develop their own applications. For example, IBM offers a Smart Business Application Development & Test service for software development and testing on the IBM Cloud. Another example is Salesforce.com's Force.com, which allows developers to build applications that are hosted on its servers as a service.
- **Software as a Service (SaaS):**Customers use software hosted by the vendor on the vendor's cloud infrastructure and delivered over a network. Leading examples are Google Apps, which provides common business applications online and Salesforce.com, which also leases customer relationship management and related software services over the Internet. Both charge users an annual subscription fee, although Google provides a free version of some of its business productivity tools. Users access these applications from a Web browser, and the data and software are maintained on the providers' remote servers.

A cloud can be private or public. A **public cloud** is owned and maintained by a cloud service provider, such as Amazon Web Services, and made available to the general public or industry group. A **private cloud** is operated solely for an organization. It may be managed by the organization or a third party and may exist on premise or off premise. Like public clouds, private clouds are able to allocate storage, computing power, or other resources seamlessly to provide computing resources on an as-needed basis. Companies that want flexible IT resources and a cloud service model while retaining control over their own IT infrastructure are gravitating toward these private clouds. (See the Interactive Session on Organizations.)

Since organizations using public clouds do not own the infrastructure, they do not have to make large investments in their own hardware and software. Instead, they purchase their computing services from remote providers and pay only for the amount of computing power they actually use (utility computing) or are billed on a monthly or annual subscription basis. The term **on-demand computing** has also been used to describe such services.

Cloud computing has some drawbacks. Unless users make provisions for storing their data locally, the responsibility for data storage and control is in the hands of the provider. Some companies worry about the security risks related to entrusting their critical data and systems to an outside vendor that also works with other companies. Companies expect their systems to be available 24/7 and do not want to suffer any loss of business capability if cloud infrastructures malfunction. Nevertheless, the trend is for companies to shift more of their computer processing and storage to some form of cloud infrastructure.

Cloud computing is more immediately appealing to small and medium-sized businesses that lack resources to purchase and own their own hardware and software. However, large corporations have huge investments in complex proprietary systems supporting unique business processes, some of which give

INTERACTIVE SESSION: ORGANIZATIONS

IS IT TIME FOR CLOUD COMPUTING?

Cloud computing is taking off. The biggest players in the cloud computing marketplace include Amazon Web Services division (AWS), Microsoft, and Google. These companies have streamlined cloud computing and made it an affordable and sensible option for companies ranging from tiny Internet startups to established companies like FedEx.

For example, AWS provides subscribing companies with flexible computing power and data storage, as well as data management, messaging, payment, and other services that can be used together or individually as the business requires. Anyone with an Internet connection and a little bit of money can harness the same computing systems that Amazon itself uses to run its retail business. If customers provide the amount of server space, bandwidth, storage, and any other services they require, AWS can automatically allocate those resources. Amazon's sales pitch is that you don't pay a monthly or yearly fee to use their computing resources – instead, you pay for exactly what you use.

This appeals to many businesses because it allows Amazon to handle all of the maintenance and upkeep of IT infrastructures, leaving businesses to spend more time on higher-value work. For example, using AWS helped Merrifield Garden Center reduce costs, improve the stability and security of its applications and data, and eliminate the burden of managing IT infrastructure hardware so it can focus on new customer-facing initiatives to grow the business.

Startup companies and smaller companies are finding that they no longer need to build their own data center. With cloud infrastructures like Amazon's readily available, they have access to technical capability that was formerly available to only much larger businesses. San Francisco-based Socialcam provides a popular mobile social video application currently installed on over 20 million iPhone and Android smartphones. The Socialcam application makes it easy to take a video of any size, post it online, and share with friends. Socialcam became so popular that the company's engineers couldn't install hardware fast enough to keep up with demand. By moving to the AWS Cloud, Socialcam can quickly add or remove capacity to meet demand. Netco Sports produces the Canal+ Football app that enables viewers to replay any move from any camera angle, on any device, within 3 minutes after it happens. By using AWS, Netco Sports can scale 100 servers in under 10 minutes to support streaming for 500,000 viewers.

Until recently, banks have been reluctant to use public cloud services due to security and regulatory concerns, but shrinking profits are encouraging them to take a second look. In the meantime, some banks are using private clouds for their sensitive financial transactions. National Australia Bank (NAB), with $793 billion in assets, uses an internal private cloud based on IBM's infrastructure on demand. IBM had already been managing the bank's IT infrastructure under a seven-year contract signed in 2010. The private cloud hosts the bank's main production environment, including a new Oracle banking system, and will support short-term computing-intense projects such as marketing campaigns. NAB pays only for what it uses so that it doesn't have to make large IT capital expenditures. The equipment is all hosted in NAB's data centers, which is unusual for on-demand environments.

Although low overhead and infrastructure management costs make public cloud computing especially attractive to startups, the financial benefits of cloud computing for large and midsized organizations are less apparent. Cliff Olson, director of infrastructure systems at FP International, Inc., a Fremont, California-based packaging company, notes that paying a public cloud provider a monthly service fee for 10,000 or more employees will probably be more expensive than having the company maintain its own IT infrastructure and staff. Companies also worry about unexpected "runaway costs" from using a pay-per-use model. Integrating cloud services with existing IT infrastructures, errors, mismanagement, or unusually high volumes of Web traffic will run up the bill for cloud service users.

Gartner Inc. technology consultants advises clients contemplating public cloud services to take into account the number of machines an organization will run; the number of hours per day or per week they'll run; and the amount of storage their data will require. Additional costs include licenses that need to be paid for on a recurring basis; the rate of change for the data; and how much new data the business is expected to generate. A very large company may find it cheaper to own and manage its own data center or private cloud. But as public

clouds become more efficient and secure and the technology grows cheaper, large companies will start using more cloud resources.

A major barrier to widespread cloud adoption is concern about cloud reliability and security. Amazon's cloud experienced significant outages in April and August 2011; on June 14 and 29, 2012; on December 24, 2012; on January 31, 2013; and on August 25, 2013. Normally, cloud networks are very reliable, and often more so than private networks operated by individual companies. But when a cloud of significant size like Amazon's goes down, it sends ripples across the entire Web. The August 2013 outage was caused by a hardware failure lasting 49 minutes at Amazon's U.S.-East data center in North Virginia. It led to spiraling problems at a host of well-trafficked online services, including Instagram, Vine, AirBnB, and the mobile magazine app Flipboard. Amazon attributed the outage to glitches with a single networking device that resulted in data loss.

The outages have been proof that the vision of a cloud with 100 percent uptime is still far from reality. Nevertheless, some large cloud users such as Netflix believe that overall cloud service availability and reliability have steadily improved. A number of experts recommend that companies for whom an outage would be a major risk consider using another computing service as a backup.

Most midsized and large companies will gravitate toward a hybrid approach. For example, InterContinental Hotels revamped its IT infrastructure to include both private and public cloud usage. To improve response time for customers, InterContinental moved its core room reservation transaction system onto a private cloud within its own data center, but it moved room availability and pricing Web site applications onto public cloud data centers on the East and West coasts. Customers receive data faster if the data are located on a server that is physically close to them, and cloud computing helps InterContinental to take advantage of this.

Sources: Beth Pariseau, "Enterprises Hit Tipping Point in AWS Cloud vs. Private Cloud Costs," searchAWS.com, April 17, 2014; Penny Crosman, "Banks Pushed Toward Cloud Computing by Cost Pressures," Information Management, March 11, 2014; "Customer Success. Powered by the AWS Cloud," www.aws.com, accessed April 1, 2014; Brad Stone, "Another Amazon Outage Exposes the Cloud's Dark Lining," Bloomberg Business Week, August 26, 2013; Charles Babcock, "Cloud Implementation Costs, Complexity Surprise Companies," *Information Week*, February 6, 2013; and Penny Crossman, "How New Core, Cloud Computing Are Transforming an Aussie Bank," Information Management, January 2, 2013.

CASE STUDY QUESTIONS

1. What business benefits do cloud computing services provide? What problems do they solve?
2. What are the disadvantages of cloud computing?
3. How do the concepts of capacity planning, scalability, and TCO apply to this case? Apply these concepts both to Amazon and to subscribers of its services.
4. What kinds of businesses are most likely to benefit from using cloud computing? Why?

them strategic advantages. The cost savings from switching to cloud services are not always easy to determine for large companies that already have their own IT infrastructures in place. Corporate data centers typically work with an IT budget that accounts for a mix of operational and capital expenses. Pricing for cloud services is usually based on a per-hour or other per-use charge. Even if a company can approximate the hardware and software costs to run a specific computing task on premises, it still needs to figure in how much of the firm's network management, storage management, system administration, electricity, and real estate costs should be allocated to a single on-premises IT service. An information systems department may not have the right information to analyze those factors on a service-by-service basis.

Large firms are most likely to adopt a **hybrid cloud** computing model where they use their own infrastructure for their most essential core activities and adopt public cloud computing for less-critical systems or for additional processing capacity during peak business periods. Cloud computing will gradually shift firms from having a fixed infrastructure capacity toward a more flexible infrastructure, some of it owned by the firm, and some of it rented from giant computer centers owned by computer hardware vendors. You can find out more about cloud computing in the Learning Tracks for this chapter.

GREEN COMPUTING

By curbing hardware proliferation and power consumption, virtualization has become one of the principal technologies for promoting green computing. **Green computing** or **green IT**, refers to practices and technologies for designing, manufacturing, using, and disposing of computers, servers, and associated devices such as monitors, printers, storage devices, and networking and communications systems to minimize impact on the environment.

Reducing computer power consumption has been a very high "green" priority. Information technology is responsible for about 2 percent of total U.S. power demand and is believed to contribute about 2 percent of the world's greenhouse gases. A corporate data center can easily consume more than 100 times more power than a standard office building. All this additional power consumption has a negative impact on the environment and corporate operating costs. The chapter-opening case on Portugal Telecom illustrates some of the technologies and data center design considerations for green computing and the business and environmental benefits of cutting power consumption in the data center.

HIGH-PERFORMANCE AND POWER-SAVING PROCESSORS

Another way to reduce power requirements and hardware sprawl is to use more efficient and power-saving processors. Contemporary microprocessors now feature multiple processor cores (which perform the reading and execution of computer instructions) on a single chip. A **multicore processor** is an integrated circuit to which two or more processor cores have been attached for enhanced performance, reduced power consumption, and more efficient simultaneous processing of multiple tasks. This technology enables two or more processing engines with reduced power requirements and heat dissipation to perform tasks faster than a resource-hungry chip with a single processing core. Today you'll find PCs with dual-core, quad-core, six-core, and eight core processors and servers with 16-core processors.

Intel and other chip manufacturers have developed microprocessors that minimize power consumption, which is essential for prolonging battery life in small mobile digital devices. You will now find highly power-efficient microprocessors, such as the A6 and A7 processors used in the Apple iPhone and iPad, and Intel's Atom processor for netbooks and mobile Internet devices. The Apple micrprocessors have about one-fiftieth the power consumption of a laptop dual-core processor. Intel recently unveiled a line of ultrasmall low-power microprocessors called Quark that can be used in wearable devices, skin patches, or even swallowed to gather medical data.

5.4 WHAT ARE THE CURRENT TRENDS IN COMPUTER SOFTWARE PLATFORMS?

There are four major themes in contemporary software platform evolution:

- Linux and open source software
- Java, HTML, and HTML5
- Web services and service-oriented architecture
- Software outsourcing and cloud services

LINUX AND OPEN SOURCE SOFTWARE

Open source software is software produced by a community of several hundred thousand programmers around the world. According to the leading open source professional association, OpenSource.org, open source software is free and can be modified by users. Works derived from the original code must also be free, and the software can be redistributed by the user without additional licensing. Open source software is by definition not restricted to any specific operating system or hardware technology, although most open source software is currently based on a Linux or Unix operating system.

The open source movement has been evolving for more than 30 years and has demonstrated that it can produce commercially acceptable, high-quality software. Popular open source software tools include the Linux operating system, the Apache HTTP Web server, the Mozilla Firefox Web browser, and the Apache OpenOffice desktop productivity suite. Google's Android mobile operating system and Chrome Web browser are based on open source tools. You can find out more out more about the Open Source Definition from the Open Source Initiative and the history of open source software at the Learning Tracks for this chapter.

Linux

Perhaps the most well-known open source software is Linux, an operating system related to Unix. Linux was created by the Finnish programmer Linus Torvalds and first posted on the Internet in August 1991. Linux applications are embedded in cell phones, smartphones, tablet computers and consumer electronics. Linux is available in free versions downloadable from the Internet or in low-cost commercial versions that include tools and support from vendors such as Red Hat.

Although Linux is not used in many desktop systems, it is a leading operating system on servers, mainframe computers, and supercomputers. Linux has become the operating system of choice in the high performance computing market, powering 97 percent of the world's fastest computers. IBM, HP, Intel, Dell, and Oracle have made Linux a central part of their offerings to corporations. The popular Android operating system for mobile devices has a Linux foundation.

The rise of open source software, particularly Linux and the applications it supports, has profound implications for corporate software platforms: cost reduction, reliability and resilience, and integration, because Linux works on all the major hardware platforms from mainframes to servers to clients.

SOFTWARE FOR THE WEB: JAVA, HTML, AND HTML5

Java is an operating system-independent, processor-independent, object-oriented programming language that has become the leading interactive

environment for the Web. Java was created by James Gosling and the Green Team at Sun Microsystems in 1992. In November 13, 2006, Sun released much of Java as open source software under the terms of the GNU General Public License (GPL), completing the process on May 8, 2007.

The Java platform has migrated into cell phones, smartphones, automobiles, music players, game machines, and into set-top cable television systems serving interactive content and pay-per-view services. Java software is designed to run on any computer or computing device, regardless of the specific microprocessor or operating system the device uses. Java is the most popular development platform for mobile devices running the Android operating system. For each of the computing environments in which Java is used, Sun created a Java Virtual Machine that interprets Java programming code for that machine. In this manner, the code is written once and can be used on any machine for which there exists a Java Virtual Machine.

Java developers can create small applet programs that can be embedded in Web pages and downloaded to run on a Web browser. A **Web browser** is an easy-to-use software tool with a graphical user interface for displaying Web pages and for accessing the Web and other Internet resources. Microsoft's Internet Explorer, Mozilla Firefox, and Google Chrome browser are examples. At the enterprise level, Java is being used for more complex e-commerce and e-business applications that require communication with an organization's back-end transaction processing systems.

HTML and HTML5

HTML (Hypertext Markup Language) is a page description language for specifying how text, graphics, video, and sound are placed on a Web page and for creating dynamic links to other Web pages and objects. Using these links, a user need only point at a highlighted keyword or graphic, click on it, and immediately be transported to another document.

HTML was originally designed to create and link static documents composed largely of text. Today, however, the Web is much more social and interactive, and many Web pages have multimedia elements—images, audio, and video. Third-party plug-in applications like Flash, Silverlight, and Java have been required to integrate these rich media with Web pages. However, these add-ons require additional programming and put strains on computer processing. The next evolution of HTML, called **HTML5**, solves this problem by making it possible to embed images, audio, video, and other elements directly into a document without processor-intensive add-ons. HTML5 makes it easier for Web pages to function across different display devices, including mobile devices as well as desktops, and it will support the storage of data offline for apps that run over the Web. Other popular programming tools for Web applications include Ruby and Python. Ruby is an object-oriented programming language known for speed and ease of use in building Web applications and Python (praised for its clarity) is being used for building cloud computing applications. Major Web sites such as Google, Facebook, Amazon, and Twitter use Python and Ruby as well as Java.

WEB SERVICES AND SERVICE-ORIENTED ARCHITECTURE

Web services refer to a set of loosely coupled software components that exchange information with each other using universal Web communication

228

standards and languages. They can exchange information between two different systems regardless of the operating systems or programming languages on which the systems are based. They can be used to build open standard Web-based applications linking systems of two different organizations, and they can also be used to create applications that link disparate systems within a single company. Web services are not tied to any one operating system or programming language, and different applications can use them to communicate with each other in a standard way without time-consuming custom coding.

The foundation technology for Web services is **XML**, which stands for **Extensible Markup Language**. This language was developed in 1996 by the World Wide Web Consortium (W3C, the international body that oversees the development of the Web) as a more powerful and flexible markup language than hypertext markup language (HTML) for Web pages. Whereas HTML is limited to describing how data should be presented in the form of Web pages, XML can perform presentation, communication, and storage of data. In XML, a number is not simply a number; the XML tag specifies whether the number represents a price, a date, or a ZIP code. Table 5.2 illustrates some sample XML statements.

By tagging selected elements of the content of documents for their meanings, XML makes it possible for computers to manipulate and interpret their data automatically and perform operations on the data without human intervention. Web browsers and computer programs, such as order processing or enterprise resource planning (ERP) software, can follow programmed rules for applying and displaying the data. XML provides a standard format for data exchange, enabling Web services to pass data from one process to another.

Web services communicate through XML messages over standard Web protocols. Companies discover and locate Web services through a directory much as they would locate services in the Yellow Pages of a telephone book. Using Web protocols, a software application can connect freely to other applications without custom programming for each different application with which it wants to communicate. Everyone shares the same standards.

The collection of Web services that are used to build a firm's software systems constitutes what is known as a service-oriented architecture. A **service-oriented architecture (SOA)** is set of self-contained services that *communicate with each other to create a working software application.* Business tasks are accomplished by executing a series of these services. Software developers reuse these services in other combinations to assemble other applications as needed.

Virtually all major software vendors provide tools and entire platforms for building and integrating software applications using Web services. IBM includes Web service tools in its WebSphere e-business software platform, and Microsoft has incorporated Web services tools in its Microsoft .NET platform.

TABLE 5.2 EXAMPLES OF XML

PLAIN ENGLISH	XML
Subcompact	<AUTOMOBILETYPE="Subcompact">
4 passenger	<PASSENGERUNIT="PASS">4</PASSENGER>
$16,800	<PRICE CURRENCY="USD">$16,800</PRICE>

F-15/B

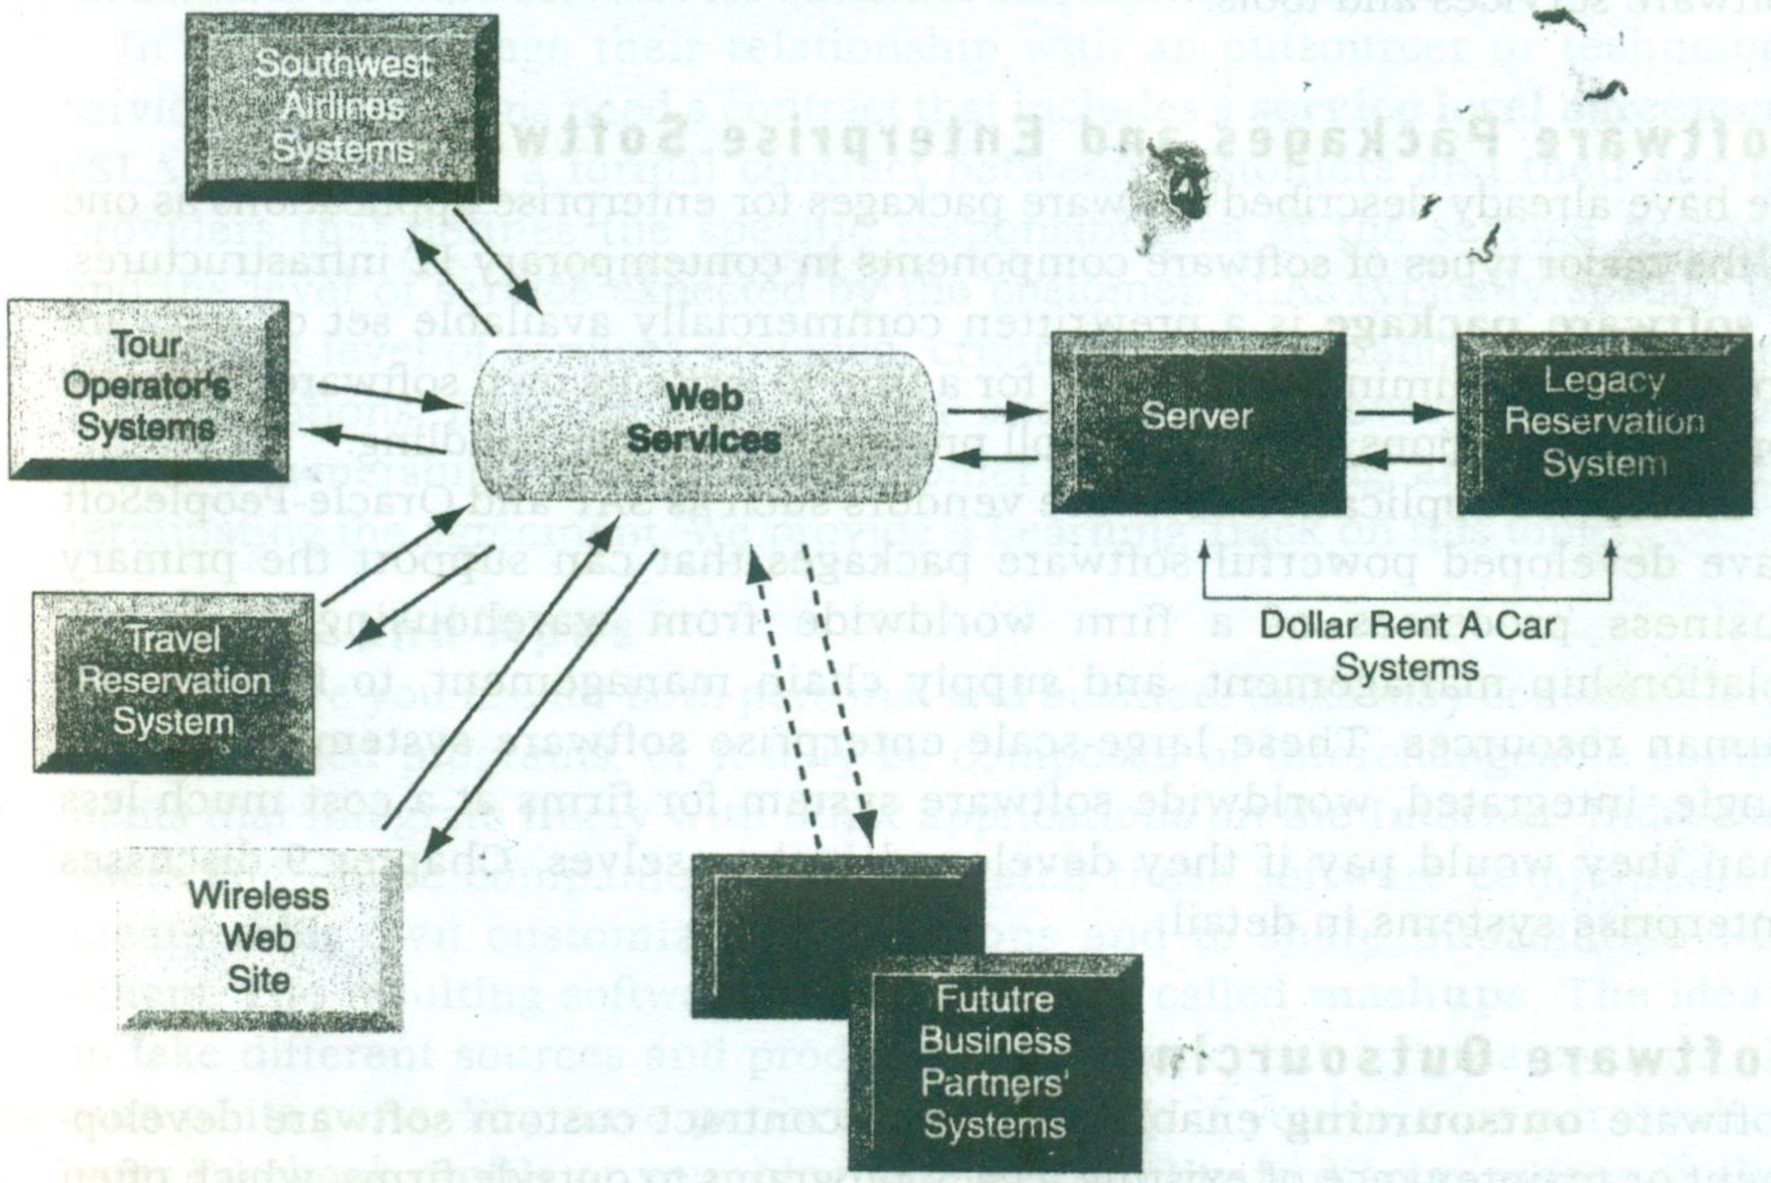
Dollar Rent A Car's systems use Web services for its online booking system with Southwest Airlines' Web site. Although both companies' systems are based on different technology platforms, a person booking a flight on Southwest.com can reserve a car from Dollar without leaving the airline's Web site. Instead of struggling to get Dollar's reservation system to share data with Southwest's information systems, Dollar used Microsoft .NET Web services technology as an intermediary. Reservations from Southwest are translated into Web services protocols, which are then translated into formats that can be understood by Dollar's computers.

Other car rental companies have linked their information systems to airline companies' Web sites before. But without Web services, these connections had to be built one at a time. Web services provide a standard way for Dollar's computers to "talk" to other companies' information systems without having to build special links to each one. Dollar is now expanding its use of Web services to link directly to the systems of a small tour operator and a large travel reservation system as well as a wireless Web site for cell phones and smartphones. It does not have to write new software code for each new partner's information systems or each new wireless device (see Figure 5.11).

SOFTWARE OUTSOURCING AND CLOUD SERVICES

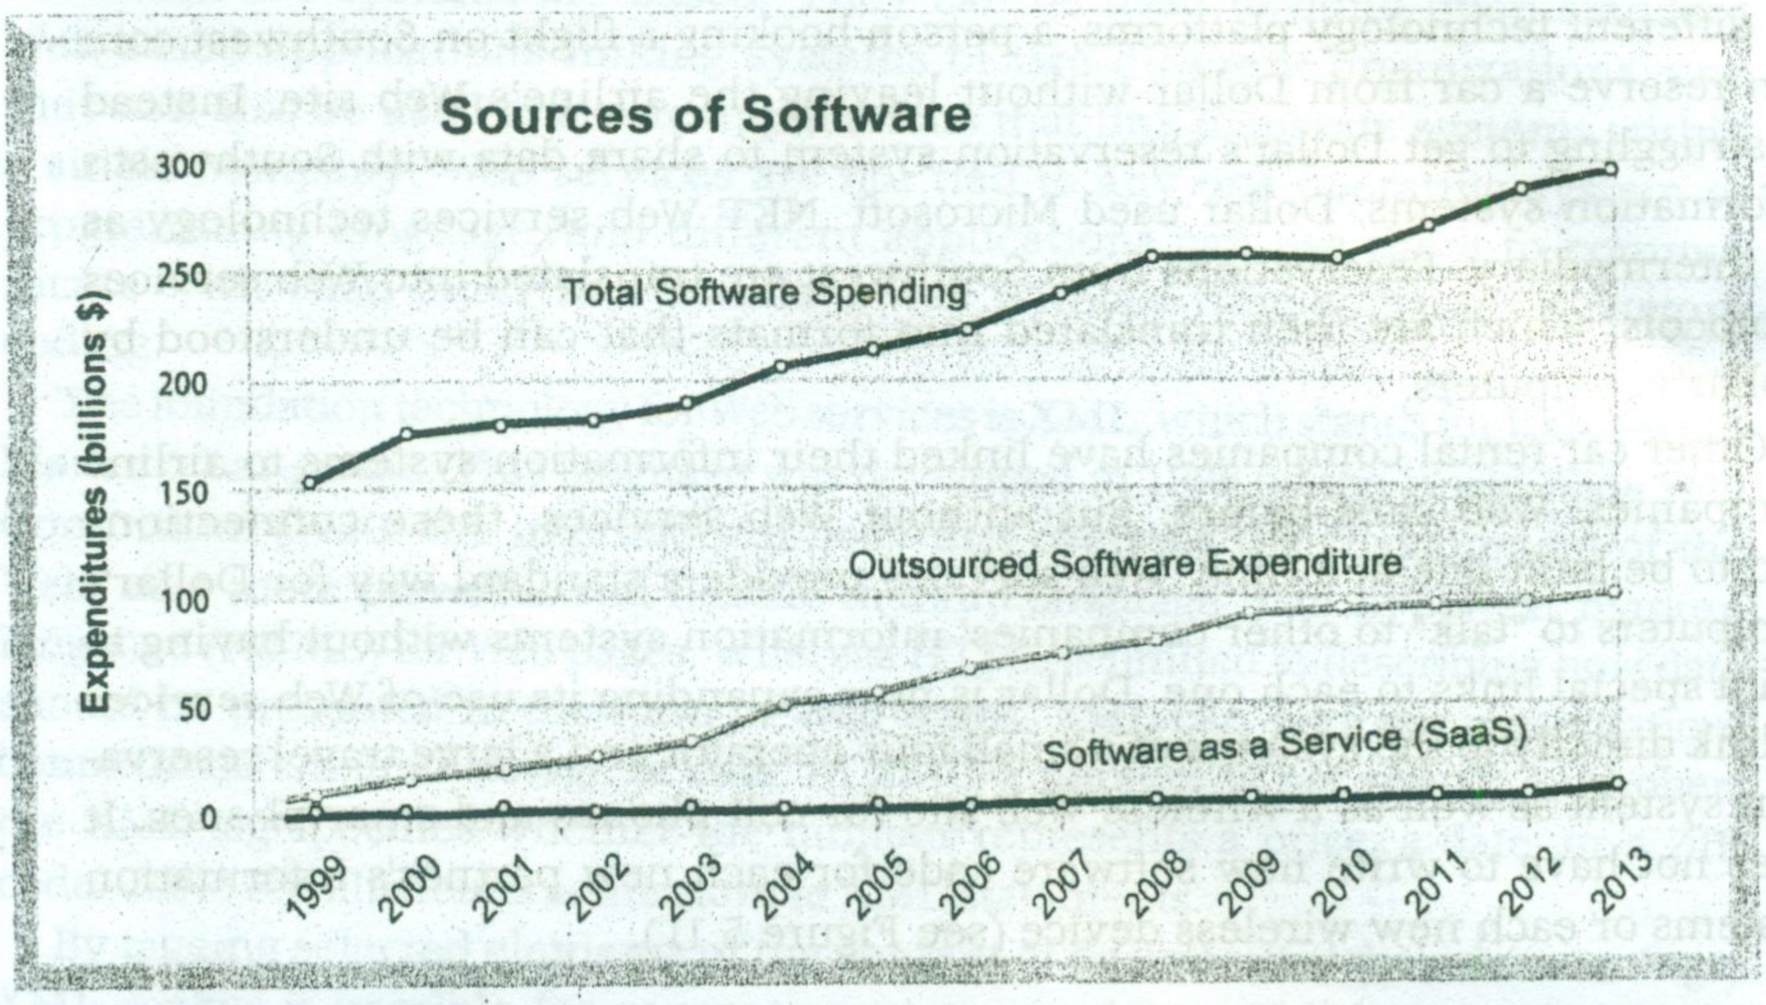
Today, many business firms continue to operate legacy systems that continue to meet a business need and that would be extremely costly to replace. But they will purchase or rent most of their new software applications from external sources. Figure 5.12 illustrates the rapid growth in external sources of software for U.S. firms.

FIGURE 5.11 HOW DOLLAR RENT A CAR USES WEB SERVICES



Dollar Rent A Car uses Web services to provide a standard intermediate layer of software to "talk" to other companies' information systems. Dollar Rent A Car can use this set of Web services to link to other companies' information systems without having to build a separate link to each firm's systems.

FIGURE 5.12 CHANGING SOURCES OF FIRM SOFTWARE



In 2014, U.S. firms will spend over $279 billion on software. About 35 percent of that will originate outside the firm, either from enterprise software vendors selling firmwide applications or individual application service providers leasing or selling software modules. Another 4 percent ($11 billion) will be provided by SaaS vendors as an online cloud-based service.

Sources: BEA National Income and Product Accounts, 2014; authors' estimates.

There are three external sources for software: software packages from a commercial software vendor, outsourcing custom application development to an external vendor, (which may or may not be offshore), and cloud-based software services and tools.

Software Packages and Enterprise Software

We have already described software packages for enterprise applications as one of the major types of software components in contemporary IT infrastructures. A **software package** is a prewritten commercially available set of software programs that eliminates the need for a firm to write its own software programs for certain functions, such as payroll processing or order handling.

Enterprise application software vendors such as SAP and Oracle-PeopleSoft have developed powerful software packages that can support the primary business processes of a firm worldwide from warehousing, customer relationship management, and supply chain management, to finance and human resources. These large-scale enterprise software systems provide a single, integrated, worldwide software system for firms at a cost much less than they would pay if they developed it themselves. Chapter 9 discusses enterprise systems in detail.

Software Outsourcing

Software **outsourcing** enables a firm to contract custom software development or maintenance of existing legacy programs to outside firms, which often operate offshore in low-wage areas of the world. According to industry analysts, spending on global IT outsourcing services was approximately $440 billion in 2014 (Kanaracus, 2014).

For example, Cemex, Mexico's largest cement manufacturer, signed a 10-year $1 billion outsourcing deal with IBM in July 2012. Under the deal, IBM responsibilities include application development and maintenance as well as IT infrastructure management at Cemex company headquarters in Monterrey, Mexico, and around the globe. IBM will take over and run Cemex's finance, accounting, and human resources systems (McDougall, 2012).

Offshore software outsourcing firms have primarily provided lower-level maintenance, data entry, and call center operations, although more sophisticated and experienced offshore firms, particularly in India, have been hired for new-program development. However, as wages offshore rise, and the costs of managing offshore projects are factored in (see Chapter 13), some work that would have been sent offshore is returning to domestic companies.

Cloud-Based Software Services and Tools

In the past, software such as Microsoft Word or Adobe Illustrator came in a box and was designed to operate on a single machine. Today, you're more likely to download the software from the vendor's Web site, or to use the software as a cloud service delivered over the Internet.

Cloud-based software and the data it uses are hosted on powerful servers in massive data centers, and can be accessed with an Internet connection and standard Web browser. In addition to free or low-cost tools for individuals and small businesses provided by Google or Yahoo, enterprise software and other complex business functions are available as services from the major commercial software vendors. Instead of buying and installing software programs, subscribing companies rent the same functions from these services, with users paying either on a subscription or per-transaction basis. Services for delivering and providing access to software remotely as a Web-based service are now referred to as **software as a service (SaaS)**. A leading example is Salesforce.com, which provides on-demand software services for customer relationship management.

In order to manage their relationship with an outsourcer or technology service provider, firms need a contract that includes a **service level agreement (SLA)**. The SLA is a formal contract between customers and their service providers that defines the specific responsibilities of the service provider and the level of service expected by the customer. SLAs typically specify the nature and level of services provided, criteria for performance measurement, support options, provisions for security and disaster recovery, hardware and software ownership and upgrades, customer support, billing, and conditions for terminating the agreement. We provide a Learning Track on this topic.

Mashups and Apps

The software you use for both personal and business tasks may consist of large self-contained programs, or it may be composed of interchangeable components that integrate freely with other applications on the Internet. Individual users and entire companies mix and match these software components to create their own customized applications and to share information with others. The resulting software applications are called **mashups**. The idea is to take different sources and produce a new work that is "greater than" the sum of its parts. You have performed a mashup if you've ever personalized your Facebook profile or your blog with a capability to display videos or slide shows.

Web mashups combine the capabilities of two or more online applications to create a kind of hybrid that provides more customer value than the original sources alone. For instance, ZipRealty uses Google Maps and data provided by

online real estate database Zillow.com to display a complete list of multiple listing service (MLS) real estate listings for any zip code specified by the user. Amazon uses mashup technologies to aggregate product descriptions with partner sites and user profiles.

Apps are small specialized software programs that run on the Internet, on your computer, or on your mobile phone or tablet and are generally delivered over the Internet. Google refers to its online services as apps, including the Google Apps suite of desktop productivity tools. But when we talk about apps today, most of the attention goes to the apps that have been developed for the mobile digital platform. It is these apps that turn smartphones and other mobile handheld devices into general-purpose computing tools.

Some downloaded apps do not access the Web but many do, providing faster access to Web content than traditional Web browsers. They provide a non-browser pathway for users to experience the Web and perform a number of tasks, ranging from reading the newspaper to shopping, searching, and buying. Because so many people are now accessing the Internet from their mobile devices, some say that apps are "the new browsers." Apps are also starting to influence the design and function of traditional Web sites as consumers are attracted to the look and feel of apps and their speed of operation.

Many apps are free or purchased for a small charge, much less than conventional software, which further adds to their appeal. There are already over 1 million apps for the Apple iPhone and iPad platform and a similar number that run on devices using Google's Android operating system. The success of these mobile platforms depends in large part on the quantity and the quality of the apps they provide. Apps tie the customer to a specific hardware platform: As the user adds more and more apps to his or her mobile phone, the cost of switching to a competing mobile platform rises.

At the moment, the most commonly downloaded apps are games, news and weather, maps/navigation, social networking, music, and video/movies. But there are also serious apps for business users that make it possible to create and edit documents, connect to corporate systems, schedule and participate in meetings, track shipments, and dictate voice messages (see the Chapter 1 Interactive Session on Management). There are also a huge number of e-commerce apps for researching and buying goods and services online.

5.5 WHAT ARE THE CHALLENGES OF MANAGING IT INFRASTRUCTURE AND MANAGEMENT SOLUTIONS?

Creating and managing a coherent IT infrastructure raises multiple challenges: dealing with platform and technology change (including cloud and mobile computing),. management and governance, and making wise infrastructure investments.

DEALING WITH PLATFORM AND INFRASTRUCTURE CHANGE

As firms grow, they often quickly outgrow their infrastructure. As firms shrink, they can get stuck with excessive infrastructure purchased in better times. How can a firm remain flexible when most of the investments in IT infrastructure are fixed-cost purchases and licenses? How well does the infrastructure scale?

Scalability refers to the ability of a computer, product, or system to expand to serve a large number of users without breaking down. New applications, mergers and acquisitions, and changes in business volume all impact computer workload and must be considered when planning hardware capacity.

Firms using mobile computing and cloud computing platforms will require new policies and procedures for managing these platforms. They will need to inventory all of their mobile devices in business use and develop policies and tools for tracking, updating, and securing them and for controlling the data and applications that run on them. Firms using cloud computing and SaaS will need to fashion new contractual arrangements with remote vendors to make sure that the hardware and software for critical applications are always available when needed and that they meet corporate standards for information security. It is up to business management to determine acceptable levels of computer response time and availability for the firm's mission-critical systems to maintain the level of business performance they expect.

MANAGEMENT AND GOVERNANCE

A long-standing issue among information system managers and CEOs has been the question of who will control and manage the firm's IT infrastructure. Chapter 2 introduced the concept of IT governance and described some issues it addresses. Other important questions about IT governance are: Should departments and divisions have the responsibility of making their own information technology decisions or should IT infrastructure be centrally controlled and managed? What is the relationship between central information systems management and business unit information systems management? How will infrastructure costs be allocated among business units? Each organization will need to arrive at answers based on its own needs.

MAKING WISE INFRASTRUCTURE INVESTMENTS

IT infrastructure is a major investment for the firm. If too much is spent on infrastructure, it lies idle and constitutes a drag on the firm's financial performance. If too little is spent, important business services cannot be delivered and the firm's competitors (who spent just the right amount) will outperform the under-investing firm. How much should the firm spend on infrastructure? This question is not easy to answer.

A related question is whether a firm should purchase and maintain its own IT infrastructure components or rent them from external suppliers, including those offering cloud services. The decision either to purchase your own IT assets or rent them from external providers is typically called the *rent-versus-buy* decision.

Cloud computing may be a low-cost way to increase scalability and flexibility, but firms should evaluate this option carefully in light of security requirements and impact on business processes and workflows. In some instances, the cost of renting software adds up to more than purchasing and maintaining an application in-house. Yet there may be benefits to using cloud services, if this allows the company to focus on core business issues instead of technology challenges.

Total Cost of Ownership of Technology Assets

The actual cost of owning technology resources includes the original cost of acquiring and installing hardware and software, as well as ongoing administration costs for hardware and software upgrades, maintenance, technical

support, training, and even utility and real estate costs for running and housing the technology. The **total cost of ownership (TCO)** model can be used to analyze these direct and indirect costs to help firms determine the actual cost of specific technology implementations. Table 5.3 describes the most important TCO components to consider in a TCO analysis.

When all these cost components are considered, the TCO for a PC might run up to three times the original purchase price of the equipment. Although the purchase price of a wireless handheld for a corporate employee may run several hundred dollars, the TCO for each device is much higher, ranging from $1,000 to $3,000, according to various consultant estimates. Gains in productivity and efficiency from equipping employees with mobile computing devices must be balanced against increased costs from integrating these devices into the firm's IT infrastructure and from providing technical support. Other cost components include fees for wireless airtime, end-user training, help desk support, and software for special applications. Costs are higher if the mobile devices run many different applications or need to be integrated into back-end systems such as enterprise applications.

Hardware and software acquisition costs account for only about 20 percent of TCO, so managers must pay close attention to administration costs to understand the full cost of the firm's hardware and software. It is possible to reduce some of these administration costs through better management. Many large firms are saddled with redundant, incompatible hardware and software because their departments and divisions have been allowed to make their own technology purchases.

In addition to switching to cloud services, these firms could reduce their TCO through greater centralization and standardization of their hardware and software resources. Companies could reduce the size of the information systems staff required to support their infrastructure if the firm minimizes the number of different computer models and pieces of software that employees are allowed to use. In a centralized infrastructure, systems can be administered from a central location and troubleshooting can be performed from that location.

TABLE 5.3 TOTAL COST OF OWNERSHIP (TCO) COST COMPONENTS

INFRASTRUCTURE COMPONENT	COST COMPONENTS
Hardware acquisition	Purchase price of computer hardware equipment, including computers, terminals, storage, and printers
Software acquisition	Purchase or license of software for each user
Installation	Cost to install computers and software
Training	Cost to provide training for information systems specialists and end users
Support	Cost to provide ongoing technical support, help desks, and so forth
Maintenance	Cost to upgrade the hardware and software
Infrastructure	Cost to acquire, maintain, and support related infrastructure, such as networks and specialized equipment (including storage backup units)
Downtime	Cost of lost productivity if hardware or software failures cause the system to be unavailable for processing and user tasks
Space and energy	Real estate and utility costs for housing and providing power for the technology

Competitive Forces Model for IT Infrastructure Investment

Figure 5.13 illustrates a competitive forces model you can use to address the question of how much your firm should spend on IT infrastructure.

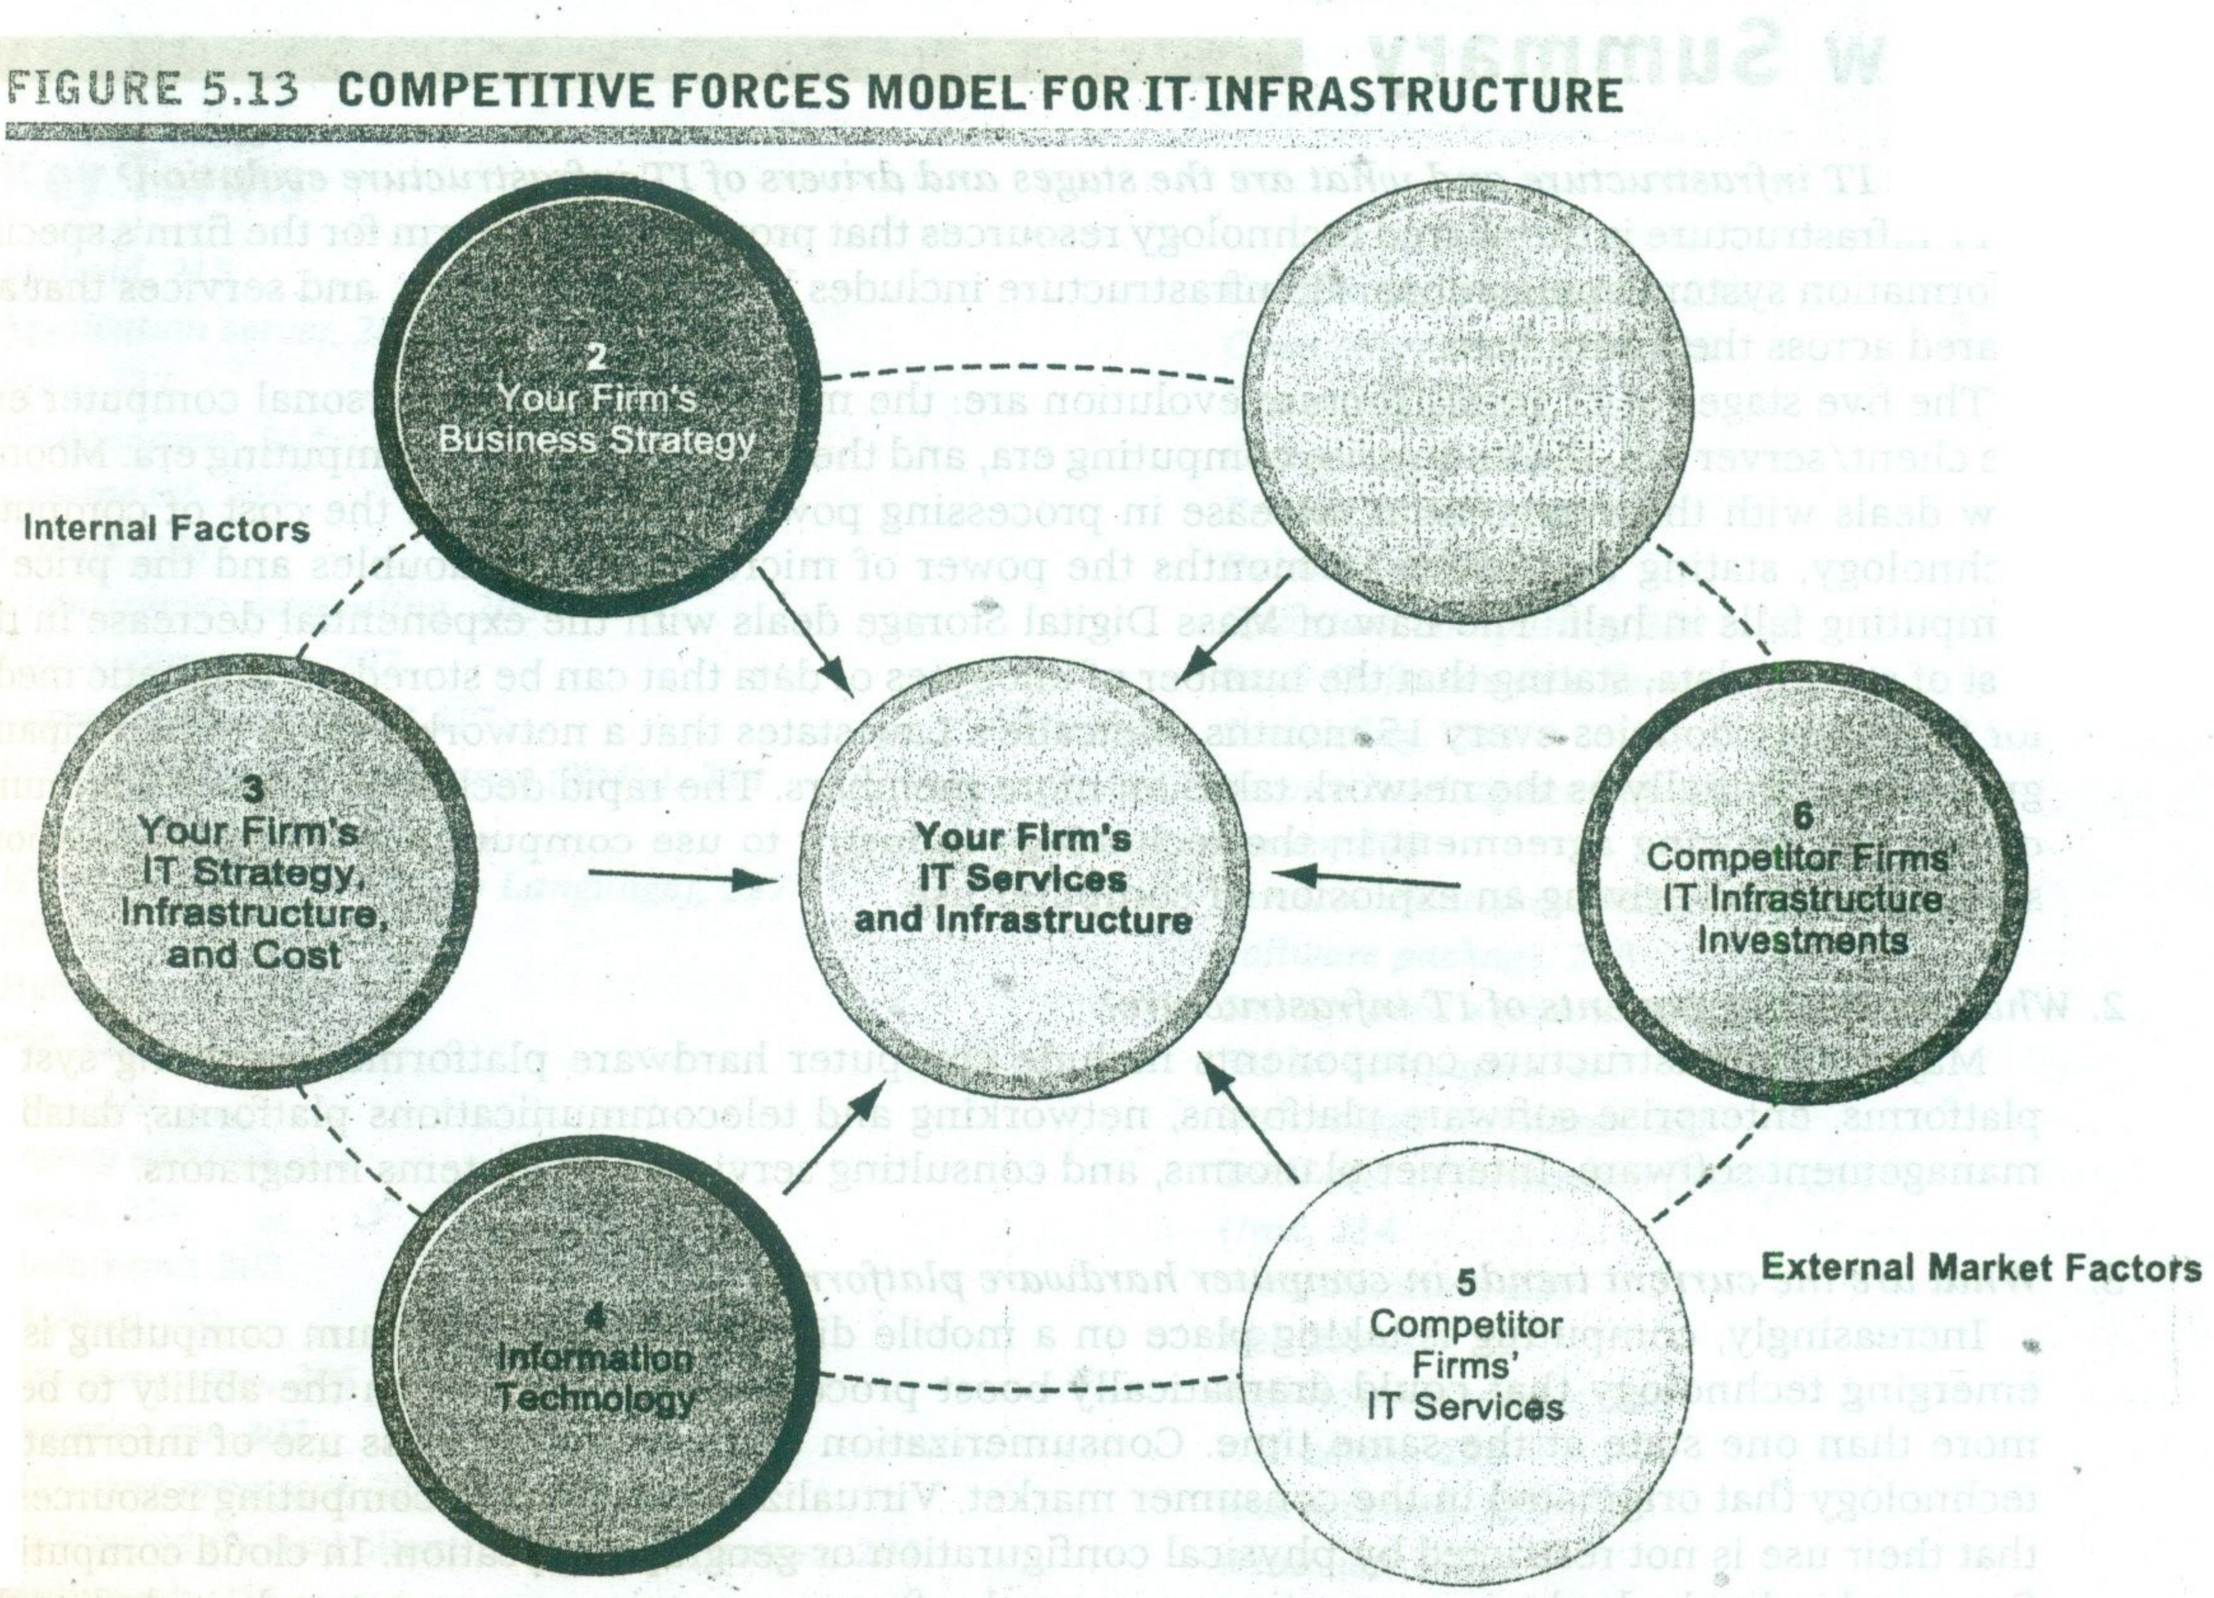
Market demand for your firm's services. Make an inventory of the services you currently provide to customers, suppliers, and employees. Survey each group, or hold focus groups to find out if the services you currently offer are meeting the needs of each group. For example, are customers complaining of slow responses to their queries about price and availability? Are employees complaining about the difficulty of finding the right information for their jobs? Are suppliers complaining about the difficulties of discovering your production requirements?

Your firm's business strategy. Analyze your firm's five-year business strategy and try to assess what new services and capabilities will be required to achieve strategic goals.

Your firm's IT strategy, infrastructure, and cost. Examine your firm's information technology plans for the next five years and assess its alignment with the firm's business plans. Determine the total IT infrastructure costs. You will want to perform a TCO analysis. If your firm has no IT strategy, you will need to devise one that takes into account the firm's five-year strategic plan.

Information technology assessment. Is your firm behind the technology curve or at the bleeding edge of information technology? Both situations are to be avoided. It is usually not desirable to spend resources on advanced technolo-

FIGURE 5.13 COMPETITIVE FORCES MODEL FOR IT INFRASTRUCTURE



There are six factors you can use to answer the question, "How much should our firm spend on IT infrastructure?"

gies that are still experimental, often expensive, and sometimes unreliable. You want to spend on technologies for which standards have been established and IT vendors are competing on cost, not design, and where there are multiple suppliers. However, you do not want to put off investment in new technologies or allow competitors to develop new business models and capabilities based on the new technologies.

Competitor firm services. Try to assess what technology services competitors offer to customers, suppliers, and employees. Establish quantitative and qualitative measures to compare them to those of your firm. If your firm's service levels fall short, your company is at a competitive disadvantage. Look for ways your firm can excel at service levels.

Competitor firm IT infrastructure investments. Benchmark your expenditures for IT infrastructure against your competitors. Many companies are quite public about their innovative expenditures on IT. If competing firms try to keep IT expenditures secret, you may be able to find IT investment information in public companies' SEC Form 10-K annual reports to the federal government when those expenditures impact a firm's financial results.

Your firm does not necessarily need to spend as much as, or more than, your competitors. Perhaps it has discovered much less-expensive ways of providing services, and this can lead to a cost advantage. Alternatively, your firm may be spending far less than competitors and experiencing commensurate poor performance and losing market share.

Review Summary

1. *What is IT infrastructure and what are the stages and drivers of IT infrastructure evolution?*

IT infrastructure is the shared technology resources that provide the platform for the firm's specific information system applications. IT infrastructure includes hardware, software, and services that are shared across the entire firm.

The five stages of IT infrastructure evolution are: the mainframe era, the personal computer era, the client/server era, the enterprise computing era, and the cloud and mobile computing era. Moore's Law deals with the exponential increase in processing power and decline in the cost of computer technology, stating that every 18 months the power of microprocessors doubles and the price of computing falls in half. The Law of Mass Digital Storage deals with the exponential decrease in the cost of storing data, stating that the number of kilobytes of data that can be stored on magnetic media for $1 roughly doubles every 15 months. Metcalfe's Law states that a network's value to participants grows exponentially as the network takes on more members. The rapid decline in costs of communication and growing agreement in the technology industry to use computing and communications standards is also driving an explosion of computer use.

2. *What are the components of IT infrastructure?*

Major IT infrastructure components include computer hardware platforms, operating system platforms, enterprise software platforms, networking and telecommunications platforms, database management software, Internet platforms, and consulting services and systems integrators.

3. *What are the current trends in computer hardware platforms?*

Increasingly, computing is taking place on a mobile digital platform. Quantum computing is an emerging technology that could dramatically boost processing power through the ability to be in more than one state at the same time. Consumerization of IT is the business use of information technology that originated in the consumer market. Virtualization organizes computing resources so that their use is not restricted by physical configuration or geographic location. In cloud computing, firms and individuals obtain computing power and software as services over a network, including the Internet, rather than purchasing and installing the hardware and software on their own computers. A

multicore processor is a microprocessor to which two or more processing cores have been attached for enhanced performance. Green computing includes practices and technologies for producing, using, and disposing of information technology hardware to minimize negative impact on the environment.

4. *What are the current trends in computer software platforms?*

Open source software is produced and maintained by a global community of programmers and is often downloadable for free. Linux is a powerful, resilient open source operating system that can run on multiple hardware platforms and is used widely to run Web servers. Java is an operating-system– and hardware-independent programming language that is the leading interactive programming environment for the Web. HTML5 makes it possible to embed images, audio, and video directly into a Web document without add-on programs. Web services are loosely coupled software components based on open Web standards that work with any application software and operating system. They can be used as components of Web-based applications linking the systems of two different organizations or to link disparate systems of a single company. Companies are purchasing their new software applications from outside sources, including software packages, by outsourcing custom application development to an external vendor (that may be offshore), or by renting online software services (SaaS). Mashups combine two different software services to create new software applications and services. Apps are small pieces of software that run on the Internet, on a computer, or on a mobile phone and are generally delivered over the Internet.

5. *What are the challenges of managing IT infrastructure and management solutions?*

Major challenges include dealing with platform and infrastructure change, infrastructure management and governance, and making wise infrastructure investments. Solution guidelines include using a competitive forces model to determine how much to spend on IT infrastructure and where to make strategic infrastructure investments, and establishing the total cost of ownership (TCO) of information technology assets. The total cost of owning technology resources includes not only the original cost of computer hardware and software but also costs for hardware and software upgrades, maintenance, technical support, and training.

Key Terms

Android, 215
Application server, 206
Apps, 232
Blade servers, 214
Chrome OS, 215
Clients, 205
Client/server computing, 205
Cloud computing, 207
Consumerization of IT, 217
Extensible Markup Language (XML), 228
Green computing, 225
HTML (Hypertext Markup Language), 227
HTML5, 227
Hybrid cloud, 225
iOS, 215
Java, 226
Legacy systems, 216
Linux, 214
Mainframe, 203
Mashup, 231
Minicomputers, 205
Moore's Law, 207
Multicore processor, 225
Multitiered (N-tier) client/server architecture, 206
Multitouch, 215
Nanotechnology, 208
On-demand computing, 222
Open source software, 226
Operating system, 214
Outsourcing, 230
Private cloud, 222
Public cloud, 222
Quantum computing, 220
SaaS (Software as a Service), 231
Scalability, 233
Service level agreement (SLA),231
Server, 205
Service-oriented architecture (SOA), 228
Software package, 230
Storage area network (SAN), 215
Tablet computers, 217
Technology standards, 211
Total cost of ownership (TCO), 234
Unix, 214
Virtualization, 220
Web browser, 227
Web hosting service, 216
Web server, 206
Web services, 227
Windows, 206
Windows 8, 215
Wintel PC, 205

MyMISLab

Go to **mymislab.com** to complete the problems marked with this icon ✪.

Review Questions

5-1 What is IT infrastructure and what are the stages and drivers of IT infrastructure evolution?

- Define IT infrastructure from both a technology and a services perspective.
- List each of the eras in IT infrastructure evolution and describe its distinguishing characteristics.
- Define and describe the following: Web server, application server, multitiered client/server architecture.
- Describe Moore's Law and the Law of Mass Digital Storage.
- Describe how network economics, declining communications costs, and technology standards affect IT infrastructure.

5-2 What are the components of IT infrastructure?

- List and describe the components of IT infrastructure that firms need to manage.

5-3 What are the current trends in computer hardware platforms?

- Describe the evolving mobile platform, consumerization of IT, and cloud computing.
- Explain how businesses can benefit from virtualization, green computing, and multicore processors.
- List and describe some popular programming languages for business applications.
- List the essential characteristics of cloud computing and distinguish between a public cloud and a private cloud.

5-4 What are the current trends in computer software platforms?

- Define and describe open source software and Linux and explain their business benefits.
- Define HTML and explain how HTML5 provides additional functionality.
- Define and describe Web services and the role played by XML.
- Name and describe the three external sources for software.
- Define and describe software mashups and apps.

5-5 What are the challenges of managing IT infrastructure and management solutions?

- Name and describe the management challenges posed by IT infrastructure.
- Describe some methods that firms can use to reduce the TCO of technology assets.

Discussion Questions

✪ **5-6** Why is selecting computer hardware and software for the organization an important management decision? What management, organization, and technology issues should be considered when selecting computer hardware and software?

✪ **5-7** Should organizations use software service providers for all their software needs? Why or why not? What management, organization, and technology factors should be considered when making this decision?

✪ **5-8** What are the advantages and disadvantages of the BYOD movement?

Hands-On MIS Projects

The projects in this section give you hands-on experience in developing solutions for managing IT infrastructures and IT outsourcing, using spreadsheet software to evaluate alternative desktop systems, and using Web research to budget for a sales conference.

Management Decision Problems

5-9 Hischornklinik Group is a leading private medical clinic group in Germany. It relies on information systems to operate 14 hospitals, as well as hundreds of specialist institutes. Demand for additional servers and storage technology is growing by 20 percent each year. Hischornklinik was setting up a separate server for every application, and its servers and other computers were running a number of different operating systems, including several versions of UNIX and Windows. Hischornklinik had to manage technologies from many different vendors, including Hewlett-Packard (HP), Sun Microsystems, Microsoft, and IBM. Assess the impact of this situation on business performance. What factors and management decisions must be considered when developing a solution to this problem?

5-10 Qantas Airways, Australia's leading airline, faces cost pressures from high fuel prices and lower levels of global airline traffic. To remain competitive, the airline must find ways to keep costs low while providing a high level of customer service. Qantas had a 30-year-old data center. Management had to decide whether to replace its IT infrastructure with newer technology or outsource it. What factors should be considered by Qantas management when deciding whether to outsource? If Qantas decides to outsource, list and describe points that should be addressed in a service level agreement.

Improving Decision Making: Using a Spreadsheet to Evaluate Hardware and Software Options

Software skills: Spreadsheet formulas
Business skills: Technology pricing

5-11 In this exercise, you'll use software at various online travel sites to obtain pricing for total travel and lodging costs for a sales conference.

EuroPlastiques is a leading EU plastics company. EuroPlastiques is planning a two-day sales conference for June 19–20, starting with a reception on the evening of June 18. The conference consists of all-day meetings that the entire sales force, numbering 100 sales representatives and their 15 managers, must attend. Each sales representative requires his or her own room, and the company needs two common meeting rooms, one large enough to hold the entire sales force plus a few visitors (150 total) and the other able to hold half the force. The company would like to hold the conference in either Zurich, Switzerland, or Milan, Italy, at a Marriott- or Novotel-owned hotel. Use the Marriott and Novotel Web sites to select a hotel in whichever of these cities would enable the company to hold its sales conference within its budget and meet its sales conference requirements. Then locate flights arriving the afternoon prior to the conference.

Your attendees will be coming from Paris (44), London (22), Amsterdam (20), Stockholm (18), and Dublin (11). Determine costs of each airline ticket from these cities. When you are finished, create a budget for the conference. The budget will include the cost of each airline ticket, the room cost, and 60 euros per attendee per day for food

Improving Decision Making: Using Web Research to Budget for a Sales Conference

Software skills: Internet-based software
Business skills: Researching transportation and lodging costs

5-12 The Foremost Composite Materials Company is planning a two-day sales conference for October 19–20, starting with a reception on the evening of October 18. The conference consists of all-day meetings that the entire sales force, numbering 120 sales representatives and their 16 managers, must

attend. Each sales representative requires his or her own room, and the company needs two common meeting rooms, one large enough to hold the entire sales force plus a few visitors (200) and the other able to hold half the force. Management has set a budget of $175,000 for the representatives' room rentals. The company would like to hold the conference in either Miami or Marco Island, Florida, at a Hilton- or Marriott-owned hotel.

Use the Hilton and Marriott Web sites to select a hotel in whichever of these cities that would enable the company to hold its sales conference within its budget and meet its sales conference requirements. Then locate flights arriving the afternoon prior to the conference. Your attendees will be coming from Los Angeles (54), San Francisco (32), Seattle (22), Chicago (19), and Pittsburgh (14). Determine costs of each airline ticket from these cities. When you are finished, create a budget for the conference. The budget will include the cost of each airline ticket, the room cost, and $70 per attendee per day for food.

Collaboration and Teamwork Project

5-13 In MyMISLab, you will find a Collaboration and Teamwork Project dealing with the concepts in this chapter. You will be able to use Google Drive, Google Docs, Google Sites, Google+, or other open-source collaboration tools to complete the assignment.

The Pleasures and Pitfalls of BYOD

CASE STUDY

Just about everyone who has a smartphone wants to be able to bring it to work and use it on the job. And why not? Employees using their own smartphones would allow companies to enjoy all of the same benefits of a mobile workforce without spending their own money to purchase these devices. Smaller companies are able to go mobile without making large investments in devices and mobile services. One IBM-sponsored study by Forrester Consulting found that a BYOD program using mobile enterprise services from IBM achieved a 108 percent return on investment and payback within one month. "Anywhere/anytime" access to computing tools increased workplace productivity and raised effective employee work time by 45–60 minutes per week. According to Gartner Inc., by 2017, 50 percent of employers will require employees to supply their own mobile devices for the workplace. BYOD is becoming the "new normal."

But...wait a minute. Nearly three out of five enterprises believe that BYOD represents a growing problem for their organizations, according to a survey of 162 enterprises conducted by Osterman Research on behalf of Dell Inc. Although BYOD can improve employee job satisfaction and productivity, it also can cause a number of problems if not managed properly: support for personally owned devices is more difficult than it is for company-supplied devices, the cost of managing mobile devices can increase, and protecting corporate data and networks becomes more difficult. Research conducted by the Aberdeen Group found that on average, an enterprise with 1,000 mobile devices spends an extra $170,000 per year when it allows BYOD. So it's not that simple.

BYOD requires a significant portion of corporate IT resources dedicated to managing and maintaining a large number of devices within the organization. In the past, companies tried to limit business smartphone use to a single platform. This made it easier to keep track of each mobile device and to roll out software upgrades or fixes, because all employees were using the same devices, or at the very least, the same operating system. The most popular employer-issued smartphone used to be Research in Motion's BlackBerry, because it was considered the "most secure" mobile platform available. (BlackBerry mobile devices access corporate e-mail and data using a proprietary software and networking platform that is company-controlled and protected from outsiders.)

Today, the mobile digital landscape is much more complicated, with a variety of devices and operating systems on the market that do not have well-developed tools for administration and security. Android has over 79 percent of the worldwide smartphone market, but it is more difficult to use for corporate work than Apple mobile devices using the iOS operating system. IOS is considered a closed system and runs only on a limited number of different Apple mobile devices. In contrast, Android's fragmentation makes it more difficult and costly for corporate IT to manage. As of July 2013, there were at least 11,868 different Android-based devices were available from more than 1,700 different brands, according to a report by OpenSignal, which researches wireless networks and devices. Android's huge consumer market share attracts many hackers. Android is also vulnerable because it has an open-source architectureand comes in multiple versions.

If employees are allowed to work with more than one type of mobile device and operating system, companies need an effective way to keep track of all the devices employees are using. To access company information, the company's networks must be configured to receive connections from that device. When employees make changes to their personal phone, such as switching cellular carriers, changing their phone number, or buying a new mobile device altogether, companies will need to quickly and flexibly ensure that their employees are still able to remain productive. Firms need an efficient inventory management system that keeps track of which devices employees are using, where the device is located, whether it is being used, and what software it is equipped with. For unprepared companies, keeping track of who gets access to what data could be a nightmare.

With the large variety of phones and operating systems available, providing adequate technical support for every employee could be difficult. When employees are not able to access critical data or encounter other problems with their mobile devices, they will need assistance from the information systems department. Companies that rely on desktop computers tend to have many of the same computers

with the same specs and operating systems, making tech support that much easier. Mobility introduces a new layer of variety and complexity to tech support that companies need to be prepared to handle.

There are significant concerns with securing company information accessed with mobile devices. If a device is stolen or compromised, companies need ways to ensure that sensitive or confidential information isn't freely available to anyone. Mobility puts assets and data at greater risk than if they were only located within company walls and on company machines. Companies often use technologies that allow them to wipe data from devices remotely, or encrypt data so that if it is stolen, it cannot be used. You'll find a detailed discussion of mobile security issues in Chapter 8.

IBM's CIO Jeanette Horan believes that BYOD may cause as many problems as it solves. BYOD was not saving IBM any money and had actually created new challenges for the IT department because employees' devices are full of software that IBM doesn't control. IBM provides secure BlackBerrys for about 40,000 of its 400,000 workers while allowing 80,000 more employees to use their own smartphones or tablets to access IBM networks.

The IBM IT department found it had no grasp of which apps and services employees were using on their personal devices, and employees themselves were "blissfully unaware" of the security risks posed by popular apps. IBM decided to ban the use of such popular services as the Dropbox cloud-based cyberlocker, fearing that employees would put IBM-sensitive information in their personal Dropbox accounts, forward internal email to public Web mail services, or use their smartphones as mobile Wi-Fi hotspots. According to research by the International Data Company (IDC), 20 percent of corporate employees using personal cloud storage services admitted to using them to store enterprise data, so this is becoming a serious problem.

IBM will not allow an employee to access its corporate networks with his or her personal device unless it secures the device. The IT department configures the device so that its memory can be erased remotely if it is lost or stolen. The IT group also disables public file-transfer programs like Apple's iCloud; instead, employees use an IBM-hosted version called MyMobileHub. IBM even turns off Siri, the voice-activated personal assistant, on employees' iPhones because the spoken queries are uploaded to Apple servers.

Each employee's device is treated differently, depending on the model and the job responsibilities of the person using it. Some people are only allowed to receive IBM e-mail, calendars, and contacts on their portable devices, while others can access internal IBM applications and files (see Chapter 8). IBM equips the mobile devices of the latter category of employees with additional software, such as programs that encrypt information as it travels to and from corporate networks.

One company that has successfully implemented BYOD is Intel Corporation, the giant semiconductor company. About 70 percent of the 39,000 devices registered on its network are personal devices. Intel approached in BYOD in a positive manner, trying to find ways to make it work rather than to defeat it. Diane Bryant, then Intel's CIO, didn't want to be dependent on a single mobile vendor or device.

Intel hammered out a BYOD strategy and created an end-user service-level agreement that clarified that end users were voluntarily using BYOD rather than being mandated by Intel. The company developed different policies, rules, and access limits for each type of device-smartphone, tablet, or laptop—with multiple levels of controls in place. Intel maintains a list of approved devices. If a device does not meet its requirements, it is blocked from the network. Intel's BYOD program today offers 40 proprietary applications, including travel tools to help schedule a flight and conference room finders. The company has an internal "app store" and uses a variety of software and security tools, including mobile device management (MDM) software and mobile app management (MAM) software.

Intel's goal for BYOD is not to save money but to make employees happier and more productive. Employees like being able to use their own device and apps alongside specialized Intel apps. On average, Intel workers report that bringing their own devices saves them about 57 minutes per day, which amounts to 5 million hours annually company-wide.

Canadian Tire decided not to allow BYOD at all and issued new BlackBarry Q10 and Z10 smartphones to its 3,000 corporate employees. (Canadian Tire is one of Canada's largest companies, with an online e-commerce store and 1,200 retail outlets sell ing automotive, sports, leisure, home products, and apparel; petroleum outlets; and financial services.) The company felt that for its purposes, the bring-your-own-device model was not sufficiently secure. Canadian Tire's chief technology officer (CTO) Eugene Roman worries that an email could sent a virus into the company's core infrastructure. At present, Canadian Tire's management thinks BYOD

is interesting but is not yet ready for the company's mainstream business applications.

In order to successfully deploy mobile devices, companies need to carefully examine their business processes and determine whether or not mobility makes sense for them. Not every firm will benefit from mobility to the same degree. Without a clear idea of how exactly mobile devices fit into the long term plans for the firm, companies will end up wasting their money on unnecessary devices and programs.

Sources: Dennis McCafferty, "Surprising Facts About Mobility and BYOD," Baseline, January 29, 2014; Beatrice Piquer-Durand, "BYOD and BYOA: Dangers and Complications," Techradar Pro, March 24, 2014; Tam Harbert, "Android Invades the Enterprise," Computerworld BYOD Consumerization of IT," November 2013; Forrester Consulting, "The Total Economic Impact of IBM Managed Mobility for BYOD," May 2013; Fred Donovan, "The Growing BYOD Problem," FierceMobileIT, February 13, 2013; Brian Bergstein, "IBM Faces the Perils of 'Bring Your Own Device'," MIT Technology Review, May 21, 2013; and Matt Hamblen, "Canadian Tire forgoes BYOD, Issues BlackBerries to Workers," Computerworld, May 20, 2013.

CASE STUDY QUESTIONS

5-14 What are the advantages and disadvantages of allowing employees to use their personal smartphones for work?

5-15 What management, organization, and technology factors should be addressed when deciding whether to allow employees to use their personal smartphones for work?

5-16 Compare the BYOD experiences of IBM and Intel. Why did BYOD at Intel work so well?

5-17 Allowing employees use their own smartphones for work will save the company money. Do you agree? Why or why not?

MyMISLab

Go to **mymislab.com** for the following Assisted-graded writing questions.

5-18 What are the distinguishing characteristics of cloud computing and what are the three types of cloud services?

5-19 What is the total cost of ownership of technology assets and what are its cost components?

Chapter 5 References

Andersson, Henrik, James Kaplan, and Brent Smolinski. "Capturing Value from IT Infrastructure Innovation." McKinsey Quarterly (October 2012).

Babcock, Charles. "Cloud's Thorniest Question: Does It Pay Off?" Information Week (June 4, 2012).

Benlian , Alexander, Marios Koufaris and Thomas Hess. "Service Quality in Software-as-a-Service: Developing the SaaS-Qual Measure and Examining Its Role in Usage Continuance." Journal of Management Information Systems, 28, No. 3 (Winter 2012).

Carr, Nicholas. *The Big Switch*. New York: Norton (2008).

Clark, Don. "Intel Unveils Tiny Quark Chips for Wearable Devices." *Wall Street Journal* (September 10, 2013).

Choi, Jae, Derek L. Nazareth, and Hemant K. Jain. "Implementing Service-Oriented Architecture in Organizations." *Journal of Management Information Systems* 26, No. 4 (Spring 2010).

David, Julie Smith, David Schuff, and Robert St. Louis. "Managing Your IT Total Cost of Ownership." *Communications of the ACM* 45, No. 1 (January 2002).

EMarketer. "Smartphone Users Worldwide Will Total 1.75 Billion in 2014." (Jan 16, 2014).

Evangelista, Michelle. "The Total Impact of IBM Managed Mobility for BYOD." IBM (2013).

Fitzgerald, Brian. "The Transformation of Open Source Software." *MIS Quarterly* 30, No. 3 (September 2006).

Gartner, Inc. "Gartner Says Worldwide IT Spending on Pace to Reach $3.8 Trillion in 2014." (January 6, 2014).

Gartner, Inc. "Gartner Says Worldwide Traditional PC, Tablet, Ultramobile and Mobile Phone Shipments to Grow 4.2 Percent in 2014." (July 7, 2014).

Grossman, Lev. "Quantum Leap." *Time* (February 17, 2014).

Hagel III, John and John Seeley Brown. "Your Next IT Strategy." *Harvard Business Review* (October, 2001).

Hamilton, David. "Enterprise Cloud IT Spending to Grow 20% in 2014, Reaching $174.2B: IHS Research." The Whir.com (February 19, 2014).

Hamilton, David. "Smartphone Users Worldwide Will Total 1.75 Billion in 2014". TheWhir.com (February 19, 2014).

Hardy, Quentin. "The Era of Cloud Computing." *New York Times* (June 11, 2014).

"Internet Users." Internetlivestats.com (April 1, 2014).

Kanaracus, Chris. "Global IT Spending Outlook 'Better but Subpar' for 2014, Forrester Says." *CIO* (January 2, 2014).

Kauffman, Robert J. and Julianna Tsai. "The Unified Procurement Strategy for Enterprise Software: A Test of the 'Move to the Middle' Hypothesis." *Journal of Management Information Systems* 26, No. 2 (Fall 2009).

King, John. "Centralized vs. Decentralized Computing: Organizational Considerations and Management Options." *Computing Surveys* (October 1984).

Lyman, Peter and Hal R. Varian. "How Much Information 2003?" University of California at Berkeley School of Information Management and Systems (2003).

McAfee, Andrew. "What Every CEO Needs to Know about the Cloud." Harvard Business Review (November 2011).

McCafferty, Dennis. "Eight Interesting Facts About Java." *CIO Insight* (June 16, 2014).

Mell, Peter and Tim Grance. "The NIST Definition of Cloud Computing" Version 15. NIST (October 17, 2009).

Moore, Gordon. "Cramming More Components Onto Integrated Circuits," *Electronics* 38, Number 8 (April 19, 1965).

Mueller, Benjamin, Goetz Viering, Christine Legner, and Gerold Riempp. "Understanding the Economic Potential of Service-Oriented Architecture." *Journal of Management Information Systems* 26, No. 4 (Spring 2010).

Schuff, David and Robert St. Louis. "Centralization vs. Decentralization of Application Software." *Communications of the ACM* 44, No. 6 (June 2001).

Stango, Victor. "The Economics of Standards Wars." *Review of Network Economics* 3, Issue 1 (March 2004).

Susarla, Anjana, Anitesh Barua, and Andrew B. Whinston. "A Transaction Cost Perspective of the 'Software as a Service' Business Model. " *Journal of Management Information Systems* 26, No. 2 (Fall 2009).

Taft, Darryl K. "Application Development: Java Death Debunked: 10 Reasons Why It's Still Hot." eWeek (February 22, 2012).

Torode, Christine, Linda Tucci and Karen Goulart. "Managing the Next-Generation Data Center." *Modern Infrastructure CIO Edition* (January 2013).

Weill, Peter, and Marianne Broadbent. *Leveraging the New Infrastructure*. Cambridge, MA: Harvard Business School Press (1998).

Weitzel, Tim. *Economics of Standards in Information Networks*. Springer (2004).

Foundations of Business Intelligence: Databases and Information Management

CHAPTER 6

LEARNING OBJECTIVES

After reading this chapter, you will be able to answer the following questions:

1. What are the problems of managing data resources in a traditional file environment?
2. What are the major capabilities of database management systems (DBMS) and why is a relational DBMS so powerful?
3. What are the principal tools and technologies for accessing information from databases to improve business performance and decision making?
4. Why are information policy, data administration, and data quality assurance essential for managing the firm's data resources?

MyMISLab™

Visit **mymislab.com** for simulations, tutorials, and end-of-chapter problems.

CHAPTER CASES

BAE Systems
Driving ARI Fleet Management with Real-Time Analytics
American Water Keeps Data Flowing
Lego: Embracing Change by Combining BI with a Flexible Information System

VIDEO CASES

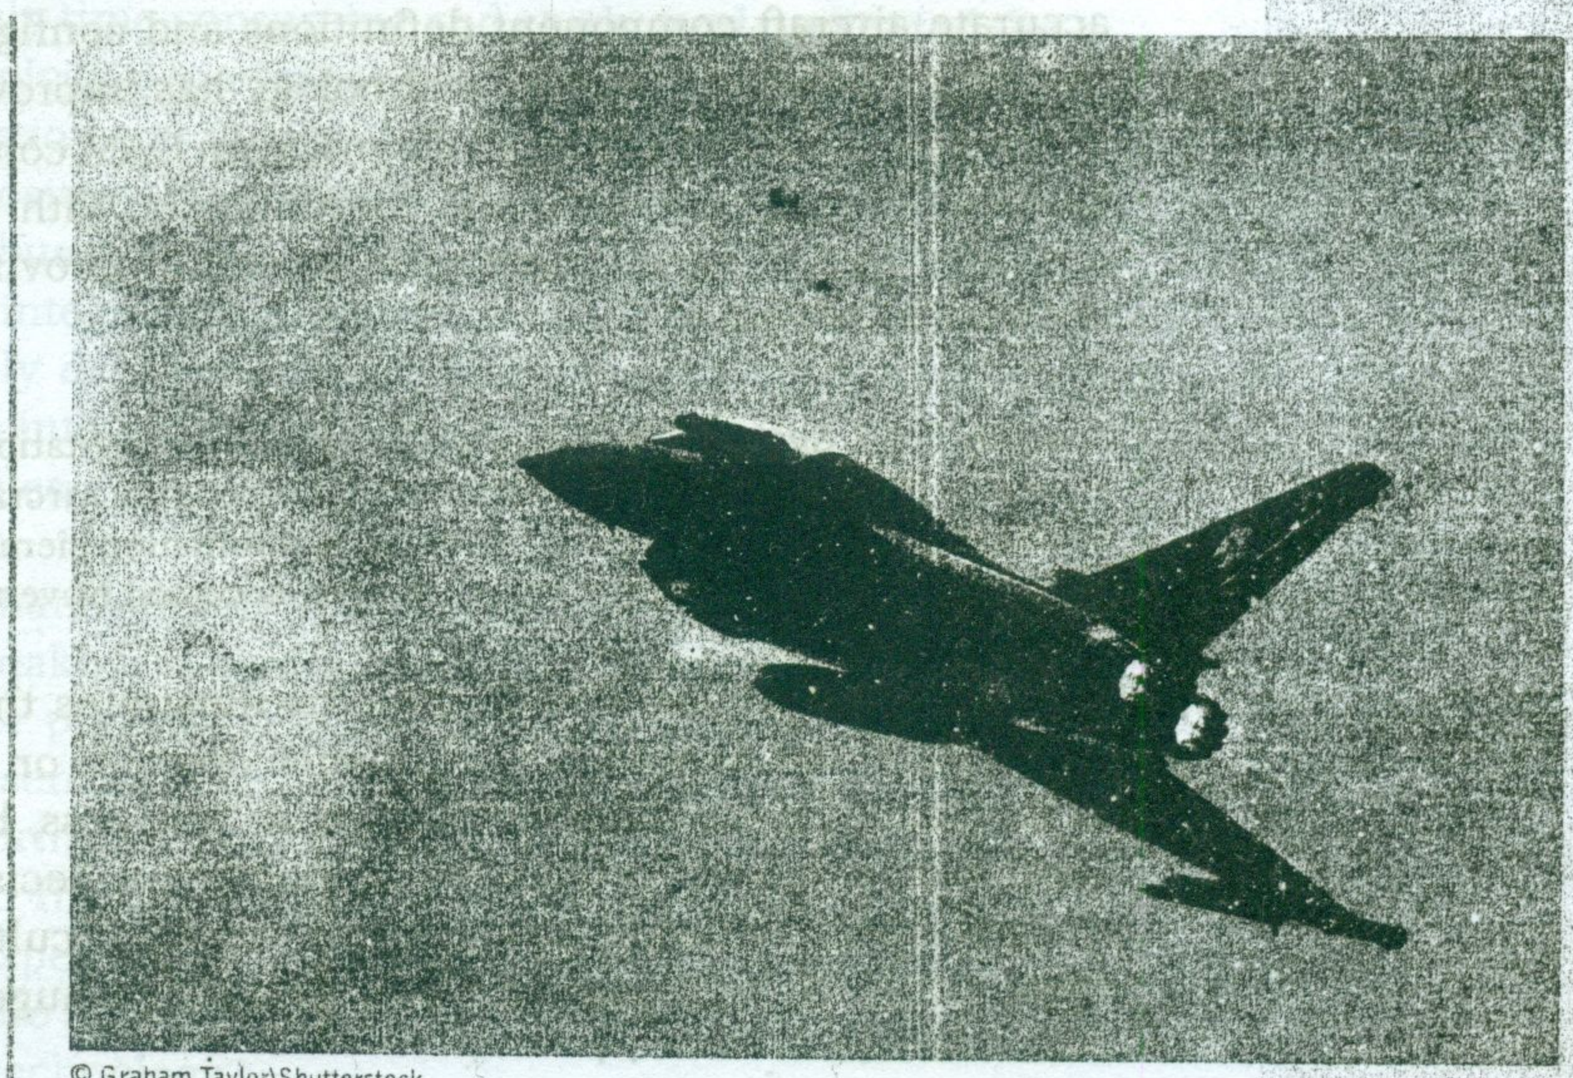
Dubuque Uses Cloud Computing and Sensors to Build a Smarter City
Data Warehousing at REI: Understanding the Customer
Maruti Suzuki Business Intelligence and Enterprise Databases

BAE SYSTEMS

BAE Systems (BAE) is the United Kingdom's (UK) largest manufacturing company and one of the largest commercial aerospace and defence organizations in Europe. Its high-technology, information-driven products and services range from one of the world's most capable multi-role combat fighters, the Eurofighter Typhoon, to the Jetstream family of commercial aircraft, to the provision of information technology (IT) and information systems (IS) for e-business to develop and implement logistics, IT, and e-capability services. With sales, manufacturing, and support sites throughout the world, including the UK, Europe, the United States, and Australia, BAE employs 88,000 people and generates more than US$30 billion in annual revenue.

Although BAE has consolidated its competitive position in established markets, and continues to expand into new markets in the Middle East and Asia, its performance in the aircraft part of the business was being impeded by legacy information systems that support the computer-aided design (CAD) and computer-aided manufacturing (CAM) of its aircraft. The distributed nature of BAE's design and manufacturing sites meant that storing and analyzing accurate sets of operational data describing the complex components of the various aircraft types to produce aircraft assembly reports for the production lines became increasingly challenging and resource-consuming. Data describing the same aircraft component parts might need resolution as in the case of various part-naming conventions and codes.

Accessing the data from the many systems was a complex task involving many technical challenges. As the aircraft business of BAE grew, so did the likelihood of delays in producing the aircraft assembly reports and other operations



© Graham Taylor\Shutterstock

data sets necessary for aircraft production management decision-making. In the worst case, the production of aircraft on the assembly line would stop until accurate information was available, with consequent schedule and cost implications. BAE's CAD/CAM staff were storing and analyzing data sets sourced from five major aircraft design and manufacturing sites spread throughout the UK, each host to thousands of staff involved in the design and manufacturing process, so that assembly reports and other operations data could be produced. Although the data that the legacy systems processed were held principally in computer files, there were numerous occasions when paper drawings with annotations containing component design and manufacturing information were used to reconcile ambiguities and inconsistencies in the assembly reports. When these data ambiguities and inconsistencies occurred, they gave rise to a sense of uncertainty in the assembly reports produced.

What BAE needed was a single repository for CAD/CAM data that would also facilitate the integration of data held in its legacy systems. The company decided to replace its legacy systems with an enterprise-wide knowledge management system that would bring the design and manufacturing data into a single database that could be concurrently accessed by the design and manufacturing engineers. BAE implemented Siemens' Teamcenter product lifecycle management software and Dassault Systemes' CATIA CAD/CAM software. Teamcenter can also be configured to take advantage of recent developments in cloud computing using Microsoft's Azure, IBM's SmartCloud Enterprise+, and Amazon Web Services.

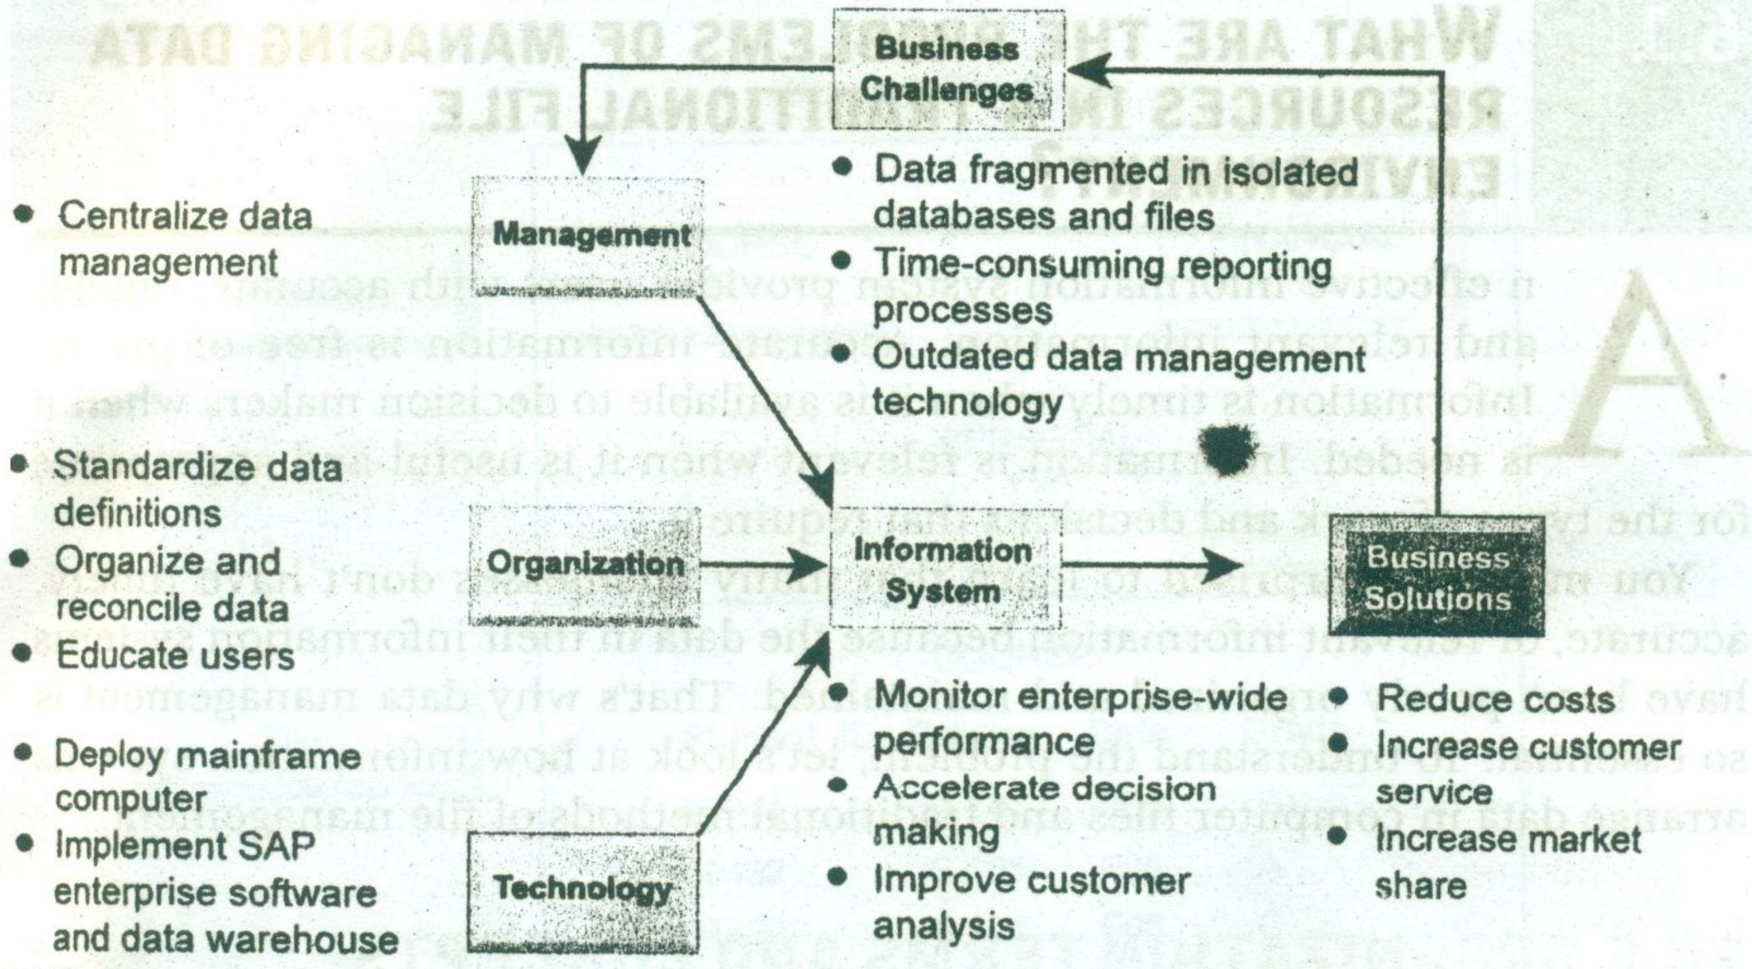
Bringing together Siemens' Teamcenter and Dassault Systemes' CATIA has given BAE Systems powerful integrated data management tools. The Teamcenter database includes tools for component markup and rollup capabilities, allowing users to visualize the effect of component design changes and configuration selections in real-time.

The new solution has produced significant cost savings at BAE in terms of its design and manufacturing data management and storage while boosting performance. With fewer legacy systems and data files to manage, BAE has been able to meet quality, time, and cost requirements by being able to produce complete and accurate aircraft component definitions and configurations. BAE's new design and manufacturing database technology has improved speed-to-market by synchronizing upstream CAD and downstream CAM component definitions, thereby enabling better cross-discipline coordination. With these savings, the company has been able to spend more resources on improving data management across the entire enterprise.

Sources: "BAE Systems Half-Yearly Report and Presentation 2012," www.baesystems.com, accessed November 8, 2012; "Teamcenter Supports Aircraft through 50-year Cycle: BAE Systems Military Air Solutions," www.plm.automation.siemens.com, accessed November 8, 2012; "CATIA V5 Fact Sheet," www.3ds.com, accessed November 8, 2012.

The experience of BAE Systems illustrates the importance of data management. Business performance depends on the accuracy and reliability of its data. The company has grown its business, but, both operational CAD/CAM efficiency and production management decision making were impeded by data stored in legacy systems that were difficult to access. How businesses store, organize, and manage their data has a huge impact on organizational effectiveness.

The chapter-opening diagram calls attention to important points raised by this case and this chapter. BAE Systems management decided that the firm needed to improve the management of its data. Pieces of data about design



components, manufactured components, and their final assembly had been stored in many large legacy systems that made it extremely difficult for the data to be retrieved and correctly unified so that it could be used in the production line assembly of aircraft components. The data were often redundant and inconsistent, limiting their usefulness. Management was unable to obtain an enterprise-view of the company.

In the past, BAE Systems had used manual paper processes to reconcile its inconsistent and redundant data and to assemble data for management reporting. This solution was extremely time-consuming and costly, and prevented the company's information technology department from performing higher-value work. A more appropriate solution was to install new hardware and software to create an enterprise-wide repository for business information that would support a more streamlined set of business applications. The new software included enterprise software that was integrated with an up-to-date database management system that could supply data for enterprise-wide reporting. The company had to reorganize its data into a standard company-wide format, eliminate redundancies, and establish rules, responsibilities, and procedures for updating and using the data.

A state-of-the-art database management system suite of software helps BAE Systems boost efficiency by making it easier to locate and assemble data for management reporting and for processing day-to-day CAD/CAM transactions for final aircraft component assembly. The data are more accurate and reliable, and costs for managing and storing the data have been considerably reduced.

Here are some questions to think about: What kinds of data management problems did BAE Systems experience in its legacy database environment? What work had to be done before the company could effectively take advantage of the new data management technology?

6.1 WHAT ARE THE PROBLEMS OF MANAGING DATA RESOURCES IN A TRADITIONAL FILE ENVIRONMENT?

An effective information system provides users with accurate, timely, and relevant information. Accurate information is free of errors. Information is timely when it is available to decision makers when it is needed. Information is relevant when it is useful and appropriate for the types of work and decisions that require it.

You might be surprised to learn that many businesses don't have timely, accurate, or relevant information because the data in their information systems have been poorly organized and maintained. That's why data management is so essential. To understand the problem, let's look at how information systems arrange data in computer files and traditional methods of file management.

FILE ORGANIZATION TERMS AND CONCEPTS

A computer system organizes data in a hierarchy that starts with bits and bytes and progresses to fields, records, files, and databases (see Figure 6.1). A **bit** represents the smallest unit of data a computer can handle. A group of bits, called a **byte**, represents a single character, which can be a letter, a number, or another symbol. A grouping of characters into a word, a group of words, or a complete number (such as a person's name or age) is called a **field**. A group of related fields, such as the student's name, the course taken, the date, and the grade, comprises a **record**; a group of records of the same type is called a **file**.

For example, the records in Figure 6.1 could constitute a student course file. A group of related files makes up a database. The student course file illustrated in Figure 6.1 could be grouped with files on students' personal histories and financial backgrounds to create a student database.

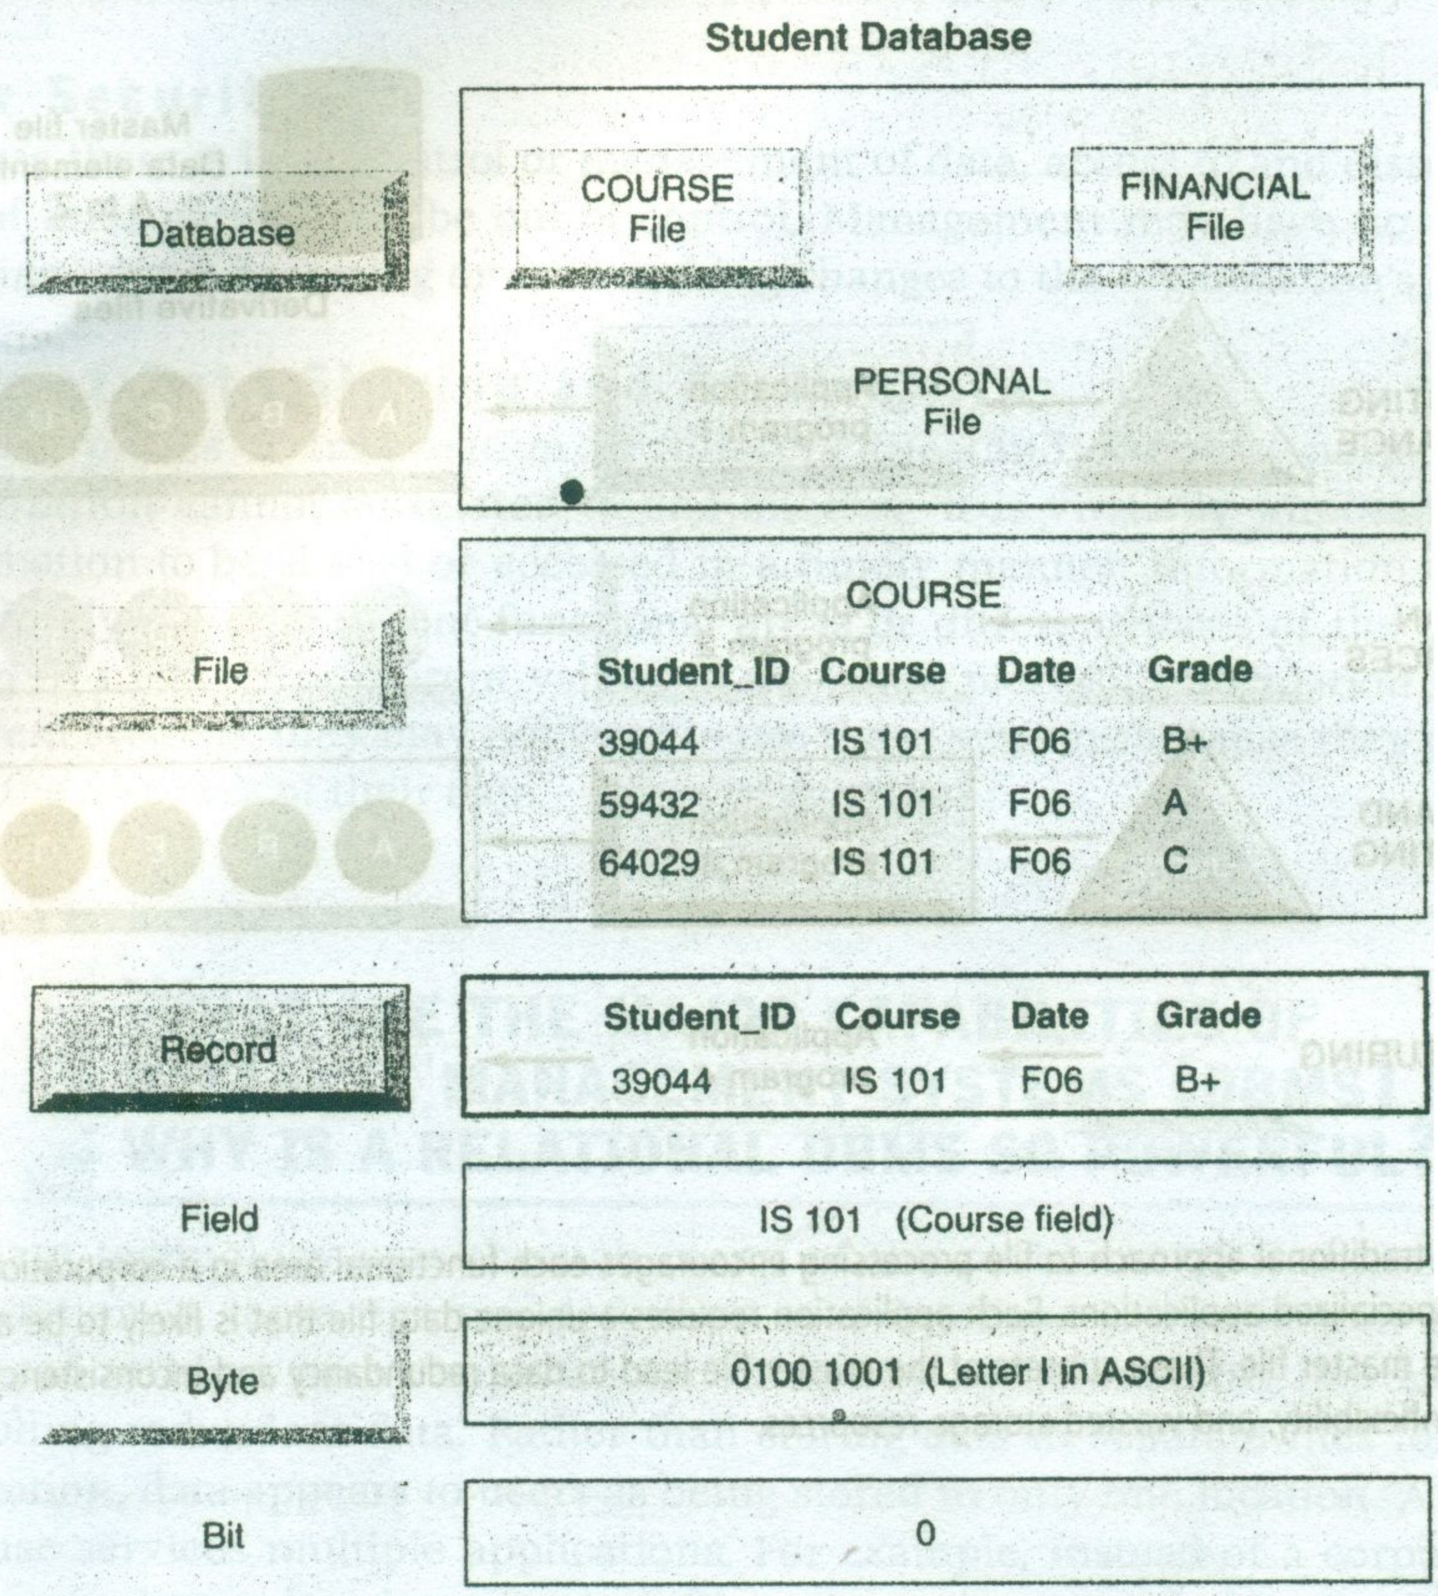
A record describes an entity. An **entity** is a person, place, thing, or event on which we store and maintain information. Each characteristic or quality describing a particular entity is called an **attribute**. For example, Student_ID, Course, Date, and Grade are attributes of the entity COURSE. The specific values that these attributes can have are found in the fields of the record describing the entity COURSE.

PROBLEMS WITH THE TRADITIONAL FILE ENVIRONMENT

In most organizations, systems tended to grow independently without a company-wide plan. Accounting, finance, manufacturing, human resources, and sales and marketing all developed their own systems and data files. Figure 6.2 illustrates the traditional approach to information processing.

Each application, of course, required its own files and its own computer program to operate. For example, the human resources functional area might have a personnel master file, a payroll file, a medical insurance file, a pension file, a mailing list file, and so forth until tens, perhaps hundreds, of files and programs existed. In the company as a whole, this process led to multiple master files created, maintained, and operated by separate divisions or departments. As this process goes on for 5 or 10 years, the organization is saddled with hundreds of programs and applications that are very difficult to maintain

FIGURE 6.1 **THE DATA HIERARCHY**



A computer system organizes data in a hierarchy that starts with the bit, which represents either a 0 or a 1. Bits can be grouped to form a byte to represent one character, number, or symbol. Bytes can be grouped to form a field, and related fields can be grouped to form a record. Related records can be collected to form a file, and related files can be organized into a database.

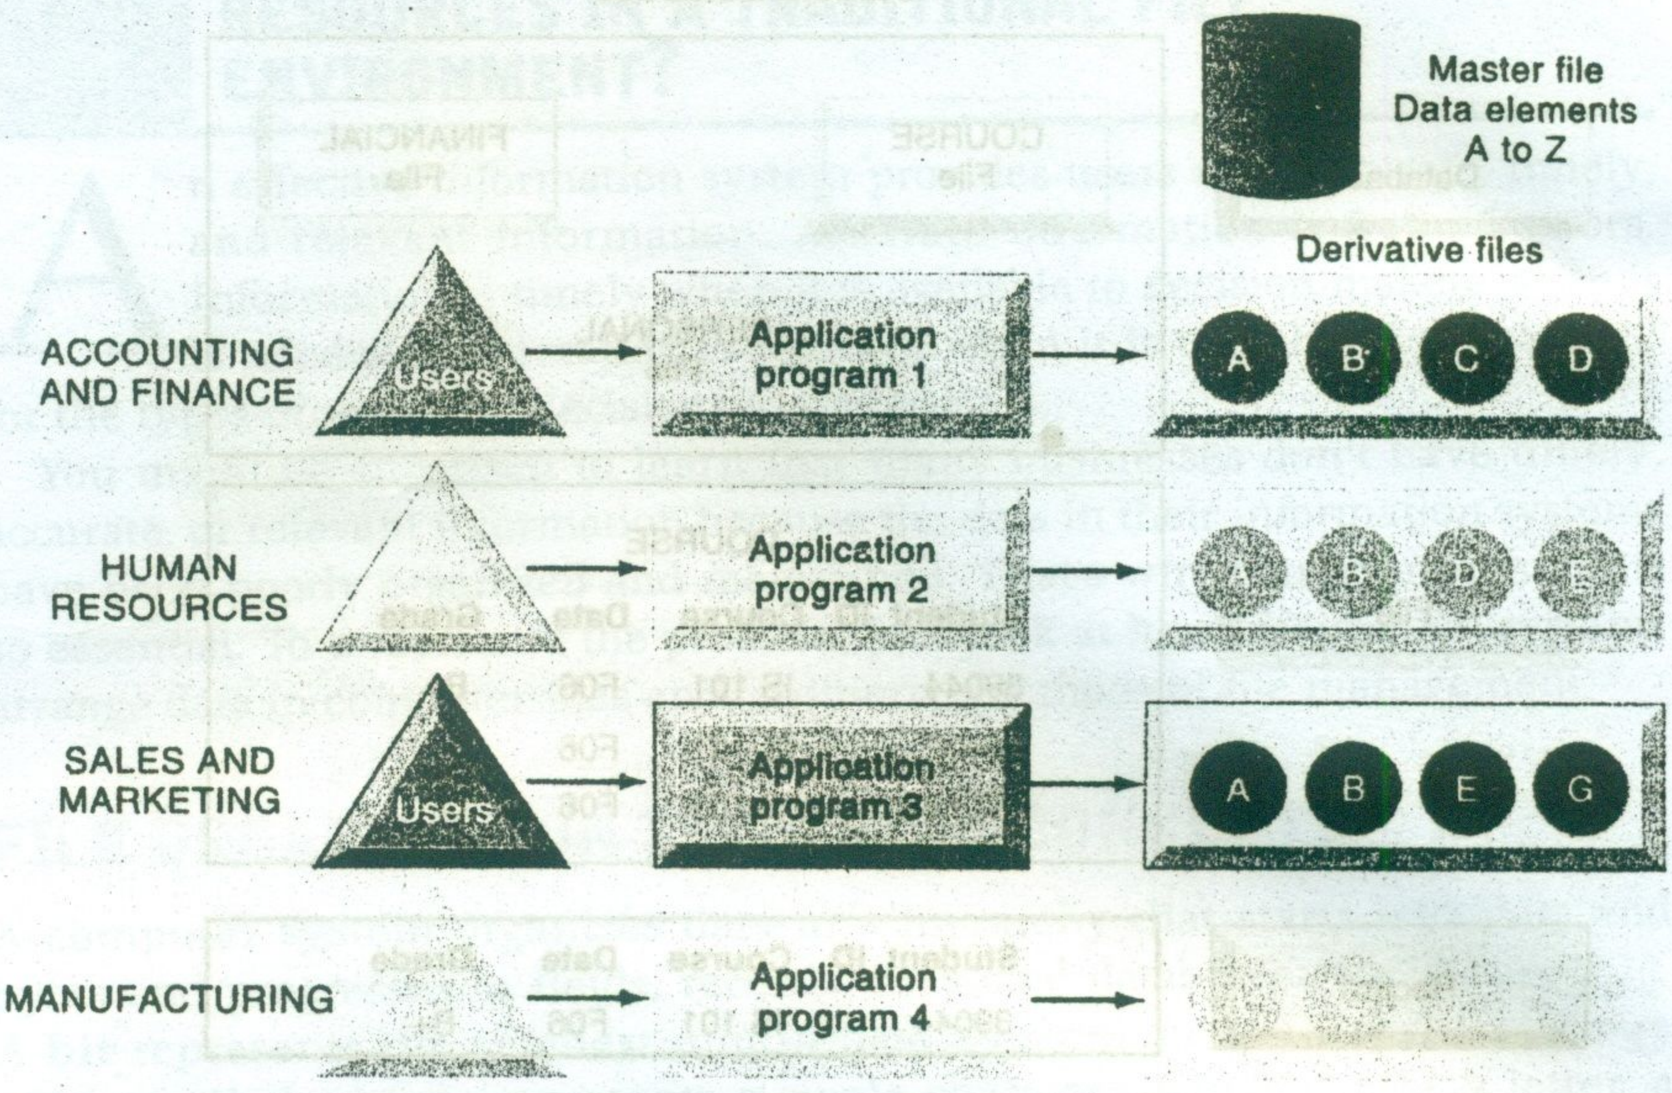
and manage. The resulting problems are data redundancy and inconsistency program-data dependence, inflexibility, poor data security, and an inability to share data among applications.

Data Redundancy and Inconsistency

Data redundancy is the presence of duplicate data in multiple data files so that the same data are stored in more than one place or location. Data redundancy occurs when different groups in an organization independently collect the same piece of data and store it independently of each other. Data redundancy wastes storage resources and also leads to **data inconsistency**, where the same attribute may have different values. For example, in instances of the entity COURSE illustrated in Figure 6.1, the Date may be updated in some systems but not in others. The same attribute, Student_ID, may also have different names in different systems throughout the organization. Some systems might use Student_ID and others might use ID, for example.

Additional confusion might result from using different coding systems to represent values for an attribute. For instance, the sales, inventory, and

FIGURE 6.2 TRADITIONAL FILE PROCESSING



The use of a traditional approach to file processing encourages each functional area in a corporation to develop specialized applications. Each application requires a unique data file that is likely to be a subset of the master file. These subsets of the master file lead to data redundancy and inconsistency, processing inflexibility, and wasted storage resources.

manufacturing systems of a clothing retailer might use different codes to represent clothing size. One system might represent clothing size as "extra large," whereas another might use the code "XL" for the same purpose. The resulting confusion would make it difficult for companies to create customer relationship management, supply chain management, or enterprise systems that integrate data from different sources.

Program-Data Dependence

Program-data dependence refers to the coupling of data stored in files and the specific programs required to update and maintain those files such that changes in programs require changes to the data. Every traditional computer program has to describe the location and nature of the data with which it works. In a traditional file environment, any change in a software program could require a change in the data accessed by that program. One program might be modified from a five-digit to a nine-digit zip code. If the original data file were changed from five-digit to nine-digit zip codes, then other programs that required the five-digit zip code would no longer work properly. Such changes could cost millions of dollars to implement properly.

Lack of Flexibility

A traditional file system can deliver routine scheduled reports after extensive programming efforts, but it cannot deliver ad hoc reports or respond to unanticipated information requirements in a timely fashion. The information required by ad hoc requests is somewhere in the system but may be too expensive to

retrieve. Several programmers might have to work for weeks to put together the required data items in a new file.

Poor Security

Because there is little control or management of data, access to and dissemination of information may be out of control. Management may have no way of knowing who is accessing or even making changes to the organization's data.

Lack of Data Sharing and Availability

Because pieces of information in different files and different parts of the organization cannot be related to one another, it is virtually impossible for information to be shared or accessed in a timely manner. Information cannot flow freely across different functional areas or different parts of the organization. If users find different values of the same piece of information in two different systems, they may not want to use these systems because they cannot trust the accuracy of their data.

6.2 WHAT ARE THE MAJOR CAPABILITIES OF DATABASE MANAGEMENT SYSTEMS (DBMS) AND WHY IS A RELATIONAL DBMS SO POWERFUL?

Database technology cuts through many of the problems of traditional file organization. A more rigorous definition of a **database** is a collection of data organized to serve many applications efficiently by centralizing the data and controlling redundant data. Rather than storing data in separate files for each application, data appears to users as being stored in only one location. A single database services multiple applications. For example, instead of a corporation storing employee data in separate information systems and separate files for personnel, payroll, and benefits, the corporation could create a single common human resources database.

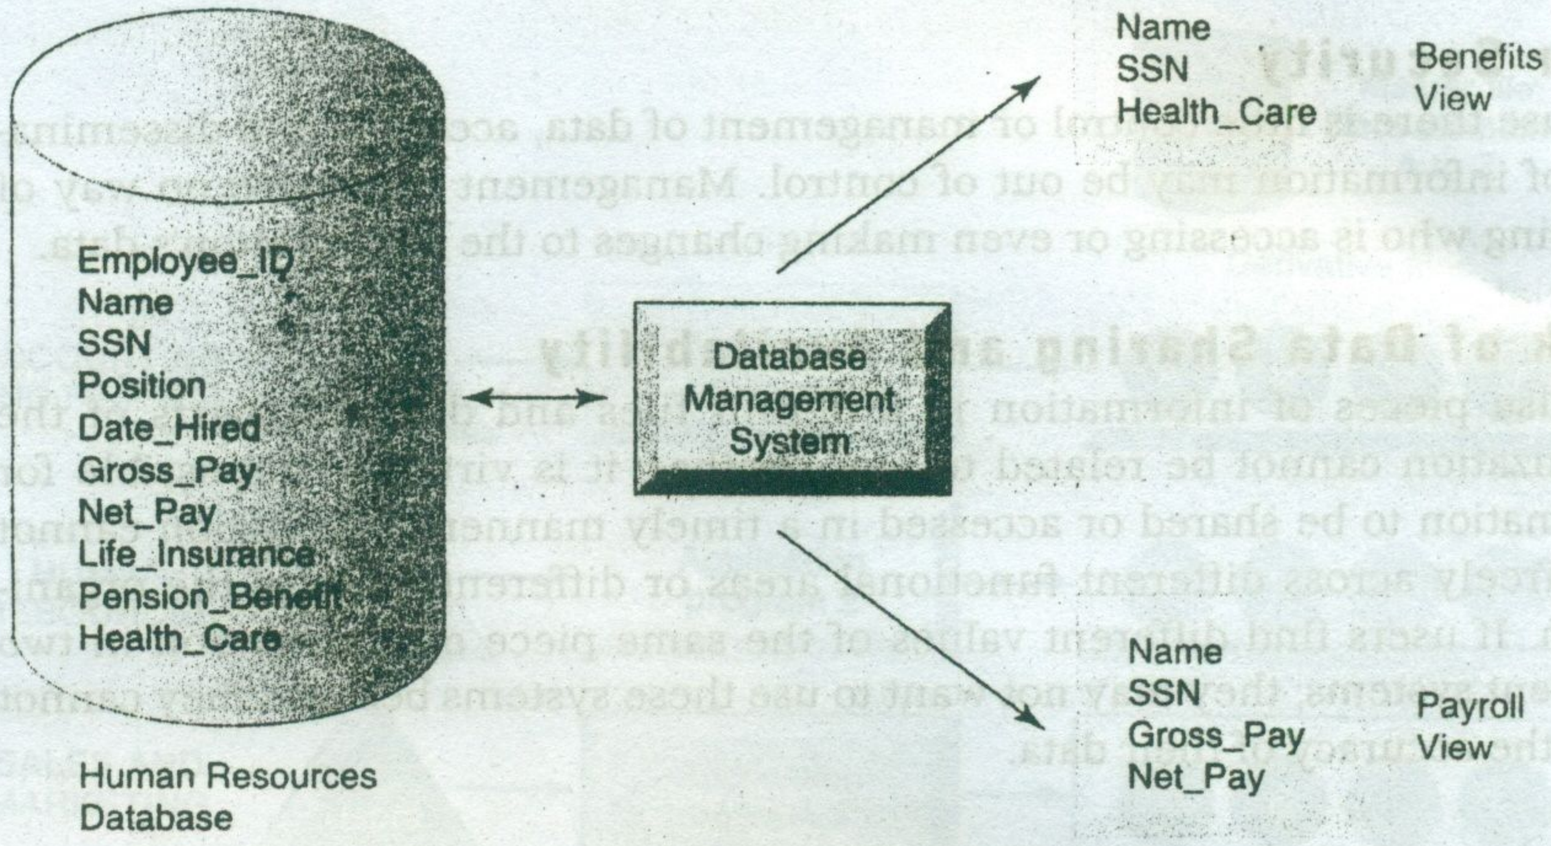
DATABASE MANAGEMENT SYSTEMS

A **database management system (DBMS)** is software that permits an organization to centralize data, manage them efficiently, and provide access to the stored data by application programs. The DBMS acts as an interface between application programs and the physical data files. When the application program calls for a data item, such as gross pay, the DBMS finds this item in the database and presents it to the application program. Using traditional data files, the programmer would have to specify the size and format of each data element used in the program and then tell the computer where they were located.

The DBMS relieves the programmer or end user from the task of understanding where and how the data are actually stored by separating the logical and physical views of the data. The *logical view* presents data as they would be perceived by end users or business specialists, whereas the *physical view* shows how data are actually organized and structured on physical storage media.

The database management software makes the physical database available for different logical views required by users. For example, for the human resources database illustrated in Figure 6.3, a benefits specialist might require a view consisting of the employee's name, social security number, and health

FIGURE 6.3 HUMAN RESOURCES DATABASE WITH MULTIPLE VIEWS



A single human resources database provides many different views of data, depending on the information requirements of the user. Illustrated here are two possible views, one of interest to a benefits specialist and one of interest to a member of the company's payroll department.

insurance coverage. A payroll department member might need data such as the employee's name, social security number, gross pay, and net pay. The data for all these views are stored in a single database, where they can be more easily managed by the organization.

How a DBMS Solves the Problems of the Traditional File Environment

A DBMS reduces data redundancy and inconsistency by minimizing isolated files in which the same data are repeated. The DBMS may not enable the organization to eliminate data redundancy entirely, but it can help control redundancy. Even if the organization maintains some redundant data, using a DBMS eliminates data inconsistency because the DBMS can help the organization ensure that every occurrence of redundant data has the same values. The DBMS uncouples programs and data, enabling data to stand on their own. The description of the data used by the program does not have to be specified in detail each time a different program is written. Access and availability of information will be increased and program development and maintenance costs reduced because users and programmers can perform ad hoc queries of the database for many simple applications without having to write complicated programs. The DBMS enables the organization to centrally manage data, their use, and security. Data-sharing throughout the organization is easier because the data are presented to users as being in a single location rather than fragmented in many different systems and files.

Relational DBMS

Contemporary DBMS use different database models to keep track of entities, attributes, and relationships. The most popular type of DBMS today for PCs as well as for larger computers and mainframes is the **relational DBMS**. Relational databases represent data as two-dimensional tables (called relations). Tables may be referred to as files. Each table contains data on an entity and its attributes. Microsoft Access is a relational DBMS for desktop systems, whereas

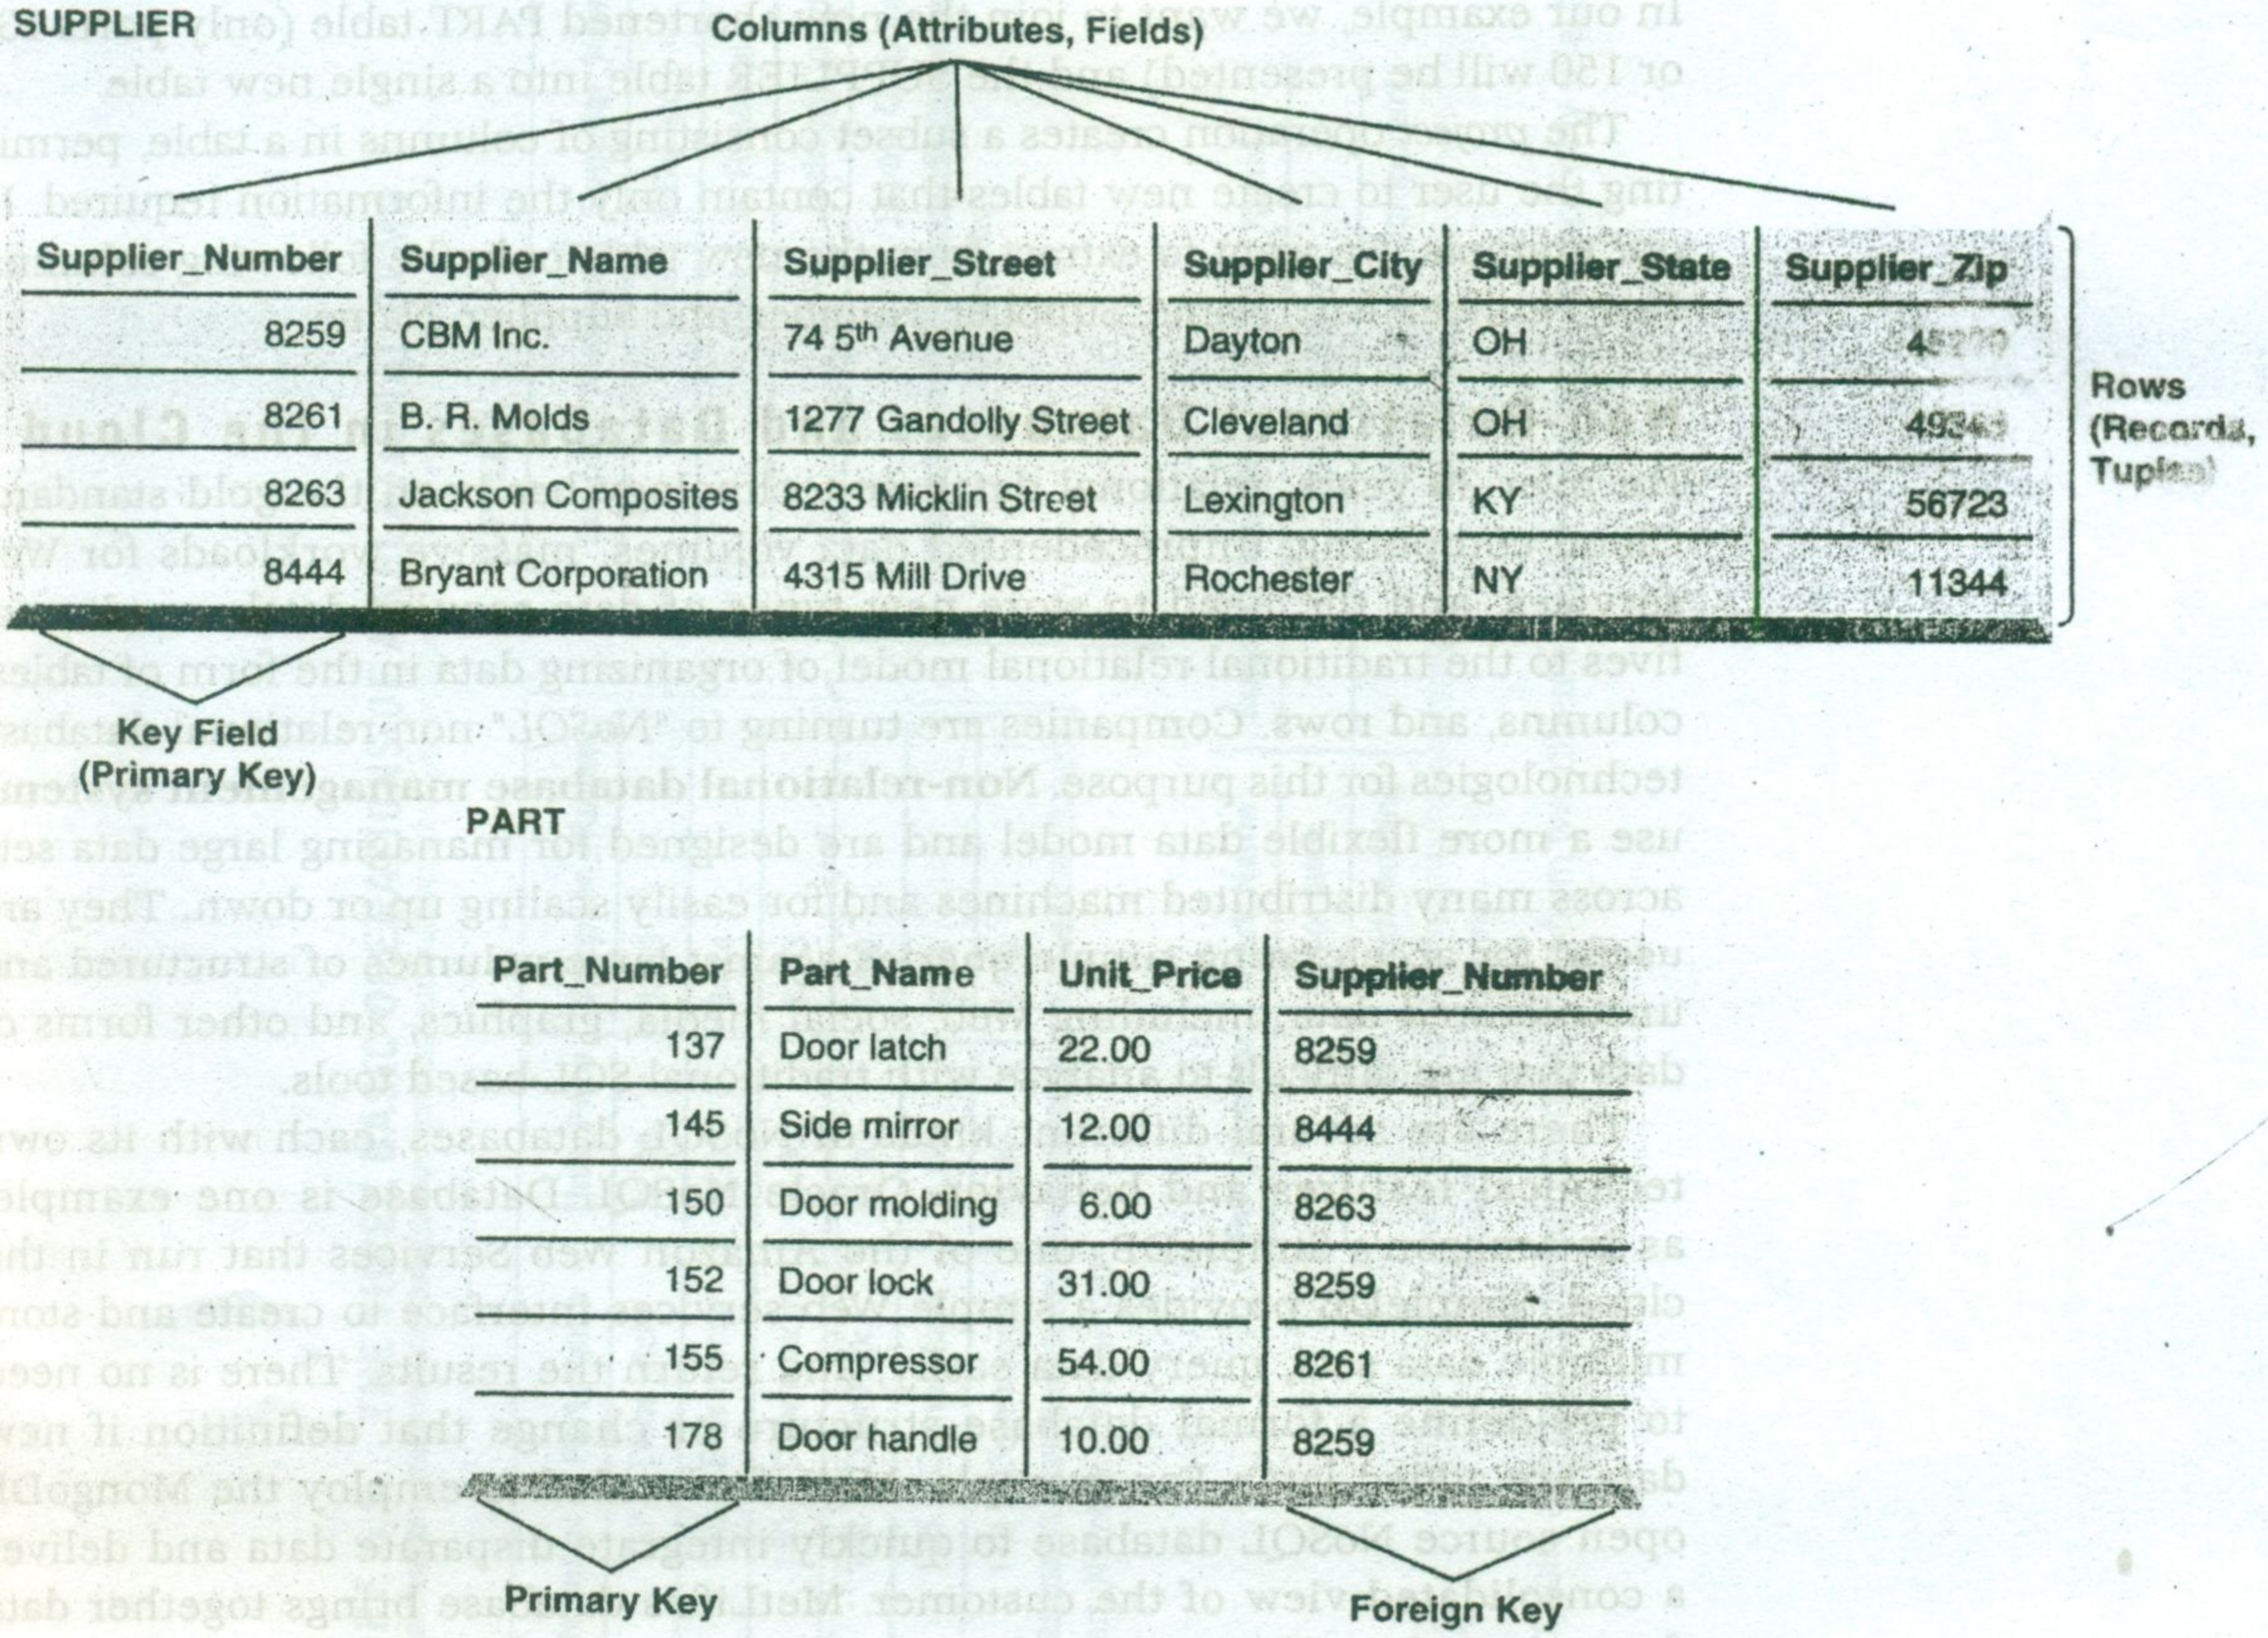
DB2, Oracle Database, and Microsoft SQL Server are relational DBMS for large mainframes and midrange computers. MySQL is a popular open source DBMS.

Let's look at how a relational database organizes data about suppliers and parts (see Figure 6.4). The database has a separate table for the entity SUPPLIER and a table for the entity PART. Each table consists of a grid of columns and rows of data. Each individual element of data for each entity is stored as a separate field, and each field represents an attribute for that entity. Fields in a relational database are also called columns. For the entity SUPPLIER, the supplier identification number, name, street, city, state, and zip code are stored as separate fields within the SUPPLIER table and each field represents an attribute for the entity SUPPLIER.

The actual information about a single supplier that resides in a table is called a row. Rows are commonly referred to as records, or in very technical terms, as **tuples**. Data for the entity PART have their own separate table.

The field for Supplier_Number in the SUPPLIER table uniquely identifies each record so that the record can be retrieved, updated, or sorted. It is called a **key field**. Each table in a relational database has one field that is designated as its **primary key**. This key field is the unique identifier for all the information

FIGURE 6.4 RELATIONAL DATABASE TABLES



A relational database organizes data in the form of two-dimensional tables. Illustrated here are tables for the entities SUPPLIER and PART showing how they represent each entity and its attributes. Supplier_Number is a primary key for the SUPPLIER table and a foreign key for the PART table.

in any row of the table and this primary key cannot be duplicated. Supplier_Number is the primary key for the SUPPLIER table and Part_Number is the primary key for the PART table. Note that Supplier_Number appears in both the SUPPLIER and PART tables. In the SUPPLIER table, Supplier_Number is the primary key. When the field Supplier_Number appears in the PART table, it is called a **foreign key** and is essentially a lookup field to look up data about the supplier of a specific part.

Operations of a Relational DBMS

Relational database tables can be combined easily to deliver data required by users, provided that any two tables share a common data element. Suppose we wanted to find in this database the names of suppliers who could provide us with part number 137 or part number 150. We would need information from two tables: the SUPPLIER table and the PART table. Note that these two files have a shared data element: Supplier_Number.

In a relational database, three basic operations, as shown in Figure 6.5, are used to develop useful sets of data: select, join, and project. The *select* operation creates a subset consisting of all records in the file that meet stated criteria. Select creates, in other words, a subset of rows that meet certain criteria. In our example, we want to select records (rows) from the PART table where the Part_Number equals 137 or 150. The *join* operation combines relational tables to provide the user with more information than is available in individual tables: In our example, we want to join the now-shortened PART table (only parts 137 or 150 will be presented) and the SUPPLIER table into a single new table.

The *project* operation creates a subset consisting of columns in a table, permitting the user to create new tables that contain only the information required. In our example, we want to extract from the new table only the following columns: Part_Number, Part_Name, Supplier_Number, and Supplier_Name.

Non-Relational Databases and Databases in the Cloud

For over 30 years, relational database technology has been the gold standard. Cloud computing, unprecedented data volumes, massive workloads for Web services, and the need to store new types of data require database alternatives to the traditional relational model of organizing data in the form of tables, columns, and rows. Companies are turning to *"NoSQL"* non-relational database technologies for this purpose. **Non-relational database management systems** use a more flexible data model and are designed for managing large data sets across many distributed machines and for easily scaling up or down. They are useful for accelerating simple queries against large volumes of structured and unstructured data, including Web, social media, graphics, and other forms of data that are difficult to analyze with traditional SQL-based tools.

There are several different kinds of NoSQL databases, each with its own technical features and behavior. Oracle NoSQL Database is one example, as is Amazon's SimpleDB, one of the Amazon Web Services that run in the cloud. SimpleDB provides a simple Web services interface to create and store multiple data sets, query data easily, and return the results. There is no need to pre-define a formal database structure or change that definition if new data are added later. For example, MetLife decided to employ the MongoDB open source NoSQL database to quickly integrate disparate data and deliver a consolidated view of the customer. MetLife's database brings together data from more than 70 separate administrative systems, claims systems and other data sources, including semi-structured and unstructured data, such as images of health records and death certificates. The NoSQLdatabase is able to ingest

FIGURE 6.5 THE THREE BASIC OPERATIONS OF A RELATIONAL DBMS

PART

Part_Number	Part_Name	Unit_Price	Supplier_Number
137	Door latch	22.00	8259
145	Side mirror	12.00	8444
150	Door molding	6.00	8263
152	Door lock	31.00	8259
155	Compressor	54.00	8261
178	Door handle	10.00	8259

Select Part_Number = 137 or 150

SUPPLIER

Supplier_Number	Supplier_Name	Supplier_Street	Supplier_City	Supplier_State	Supplier_Zip
8259	CBM Inc.	74 5th Avenue	Dayton	OH	45220
8261	B. R. Molds	1277 Gandolly Street	Cleveland	OH	49345
8263	Jackson Components	8233 Micklin Street	Lexington	KY	58723
8444	Bryant Corporation	4315 Mill Drive	Rochester	[illegible]	11344

Join by Supplier_Number

Part_Number	Part_Name	Supplier_Number	Supplier_Name
137	Door latch	8259	CBM Inc.
150	Door molding	8263	Jackson Components

Project selected columns

The select, join, and project operations enable data from two different tables to be combined and only selected attributes to be displayed.

structured, semi-structured and unstructured information without requiring tedious, expensive and time-consuming database-mapping (Henschen, 2013).

Amazon and other cloud computing vendors provide relational database services as well. Amazon Relational Database Service (Amazon RDS) offers MySQL, SQL Server, or Oracle Database as database engines. Pricing is based on usage. Oracle has its own database cloud service using its relational Oracle Database and Microsoft SQL Azure Database is a cloud-based relational database service based on Microsoft's SQL Server DBMS. Cloud-based data management services have special appeal for Web-focused start-ups or small to medium-sized businesses seeking database capabilities at a lower price than in-house database products.

In addition to public cloud-based data management services, companies now have the option of using databases in private clouds. For example, Sabre Holdings, the world's largest software as a service (SaaS) provider for the aviation industry, has a private database cloud that supports more than 100 projects and 700 users. A consolidated database spanning a pool of standardized servers running Oracle Database provides database services for multiple applications.

CAPABILITIES OF DATABASE MANAGEMENT SYSTEMS

A DBMS includes capabilities and tools for organizing, managing, and accessing the data in the database. The most important are its data definition language, data dictionary, and data manipulation language.

DBMS have a **data definition** capability to specify the structure of the content of the database. It would be used to create database tables and to define the characteristics of the fields in each table. This information about the database would be documented in a data dictionary. A **data dictionary** is an automated or manual file that stores definitions of data elements and their characteristics.

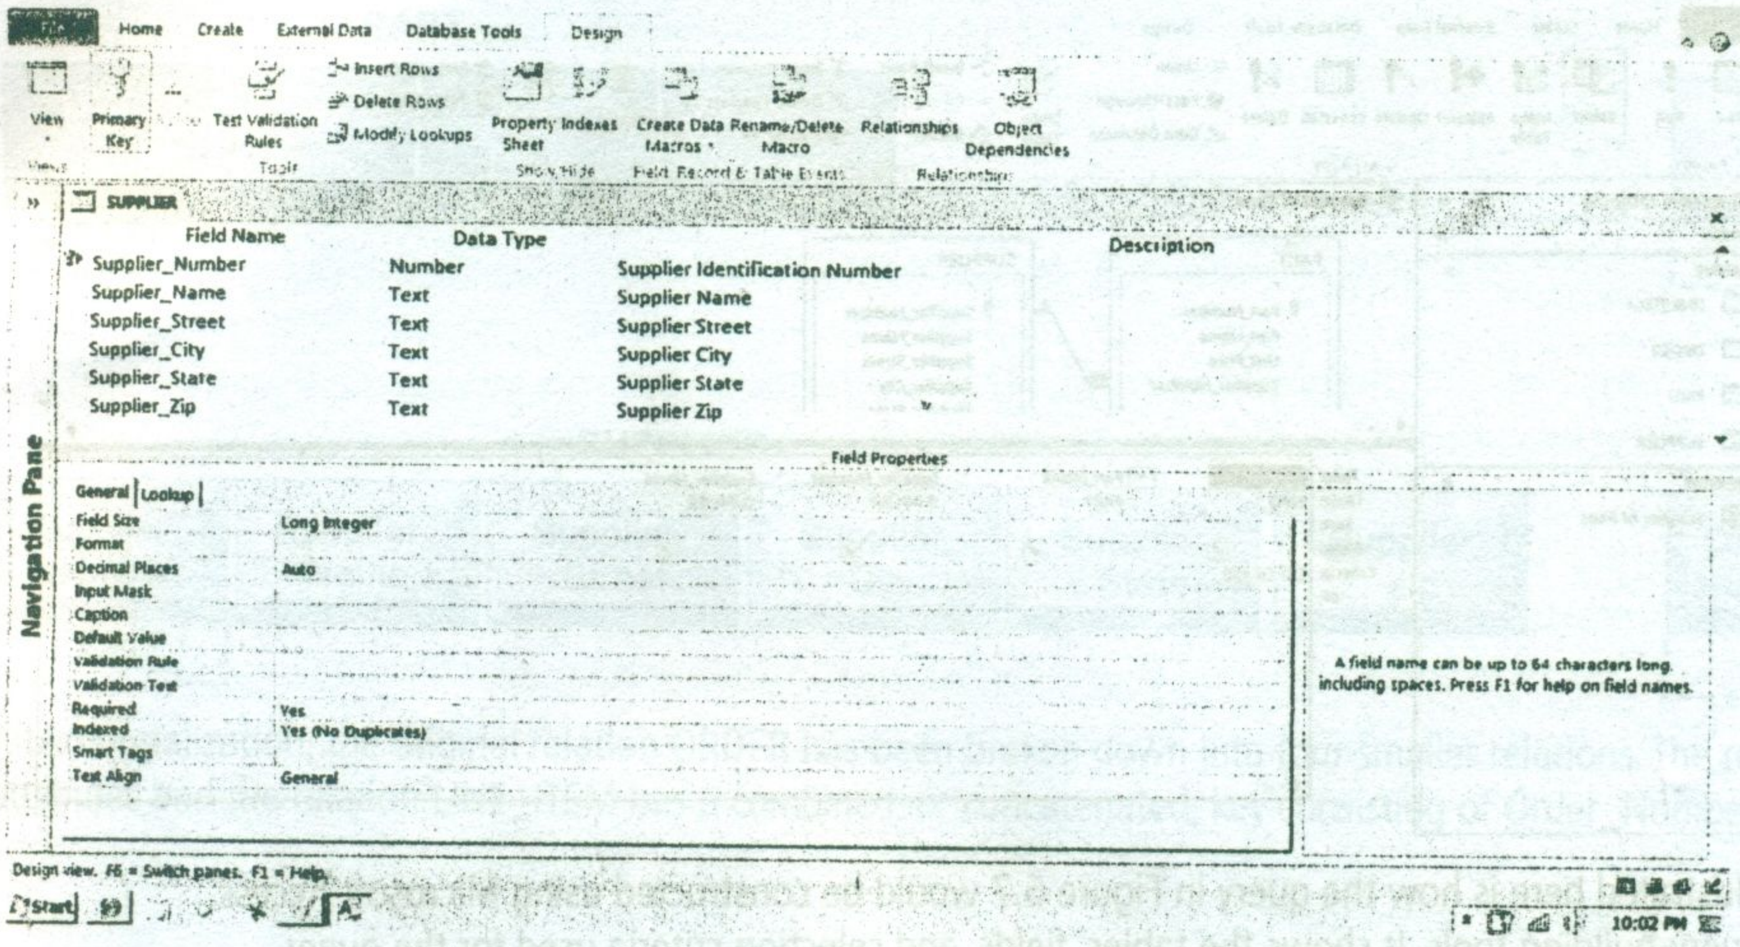
Microsoft Access has a rudimentary data dictionary capability that displays information about the name, description, size, type, format, and other properties of each field in a table (see Figure 6.6). Data dictionaries for large corporate databases may capture additional information, such as usage, ownership (who in the organization is responsible for maintaining the data), authorization, security, and the individuals, business functions, programs, and reports that use each data element.

Querying and Reporting

DBMS includes tools for accessing and manipulating information in databases. Most DBMS have a specialized language called a **data manipulation language** that is used to add, change, delete, and retrieve the data in the database. This language contains commands that permit end users and programming specialists to extract data from the database to satisfy information requests and develop applications. The most prominent data manipulation language today is **Structured Query Language**, or **SQL**. Figure 6.7 illustrates the SQL query that would produce the new resultant table in Figure 6.5. You can find out more about how to perform SQL queries in our Learning Tracks for this chapter.

Users of DBMS for large and midrange computers, such as DB2, Oracle, or SQL Server, would employ SQL to retrieve information they needed from the database. Microsoft Access also uses SQL, but it provides its own set of user-friendly tools for querying databases and for organizing data from databases into more polished reports.

FIGURE 6.6 ACCESS DATA DICTIONARY FEATURES



Microsoft Access has a rudimentary data dictionary capability that displays information about the size, format, and other characteristics of each field in a database. Displayed here is the information maintained in the SUPPLIER table. The small key icon to the left of Supplier_Number indicates that it is a key field.

In Microsoft Access, you will find features that enable users to create queries by identifying the tables and fields they want and the results, and then selecting the rows from the database that meet particular criteria. These actions in turn are translated into SQL commands. Figure 6.8 illustrates how the same query as the SQL query to select parts and suppliers would be constructed using the Microsoft query-building tools.

Microsoft Access and other DBMS include capabilities for report generation so that the data of interest can be displayed in a more structured and polished format than would be possible just by querying. Crystal Reports is a popular report generator for large corporate DBMS, although it can also be used with Access. Access also has capabilities for developing desktop system applications. These include tools for creating data entry screens, reports, and developing the logic for processing transactions.

DESIGNING DATABASES

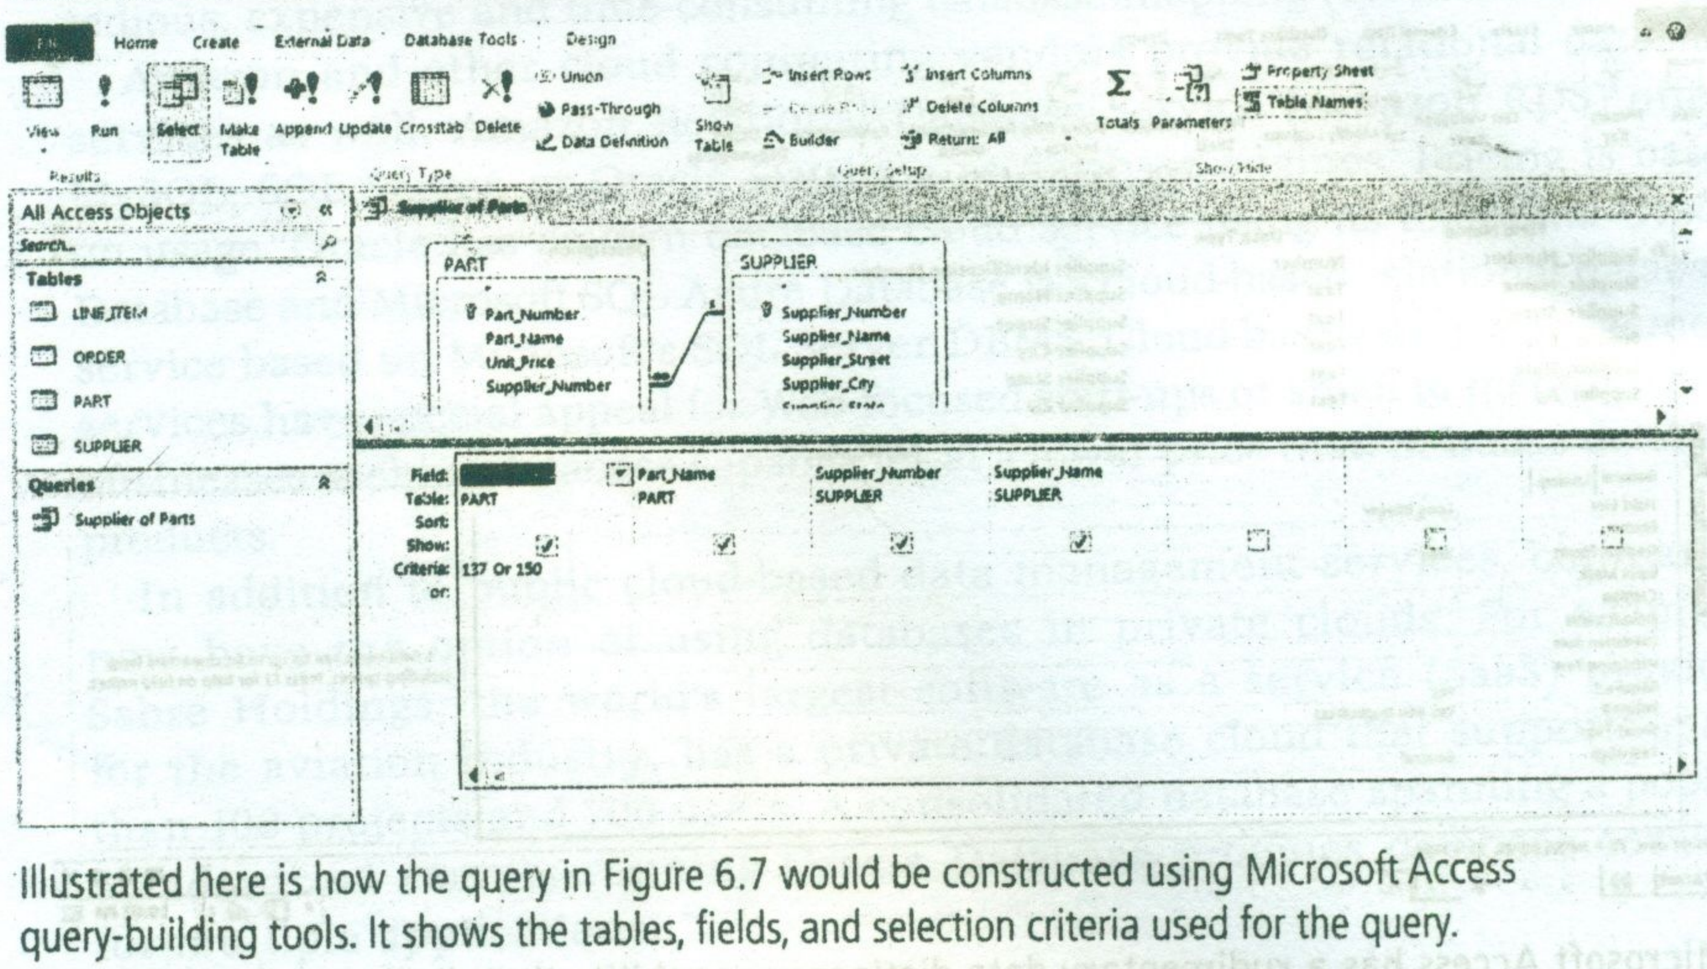
To create a database, you must understand the relationships among the data, the type of data that will be maintained in the database, how the data will be

FIGURE 6.7 EXAMPLE OF AN SQL QUERY

```
SELECT PART.Part_Number, PART.Part_Name, SUPPLIER.Supplier_Number,
SUPPLIER.Supplier_Name
FROM PART, SUPPLIER
WHERE PART.Supplier_Number = SUPPLIER.Supplier_Number AND
Part_Number = 137 OR Part_Number = 150;
```

Illustrated here are the SQL statements for a query to select suppliers for parts 137 or 150. They produce a list with the same results as Figure 6.5.

FIGURE 6.8 AN ACCESS QUERY



Illustrated here is how the query in Figure 6.7 would be constructed using Microsoft Access query-building tools. It shows the tables, fields, and selection criteria used for the query.

used, and how the organization will need to change to manage data from a company-wide perspective. The database requires both a conceptual design and a physical design. The conceptual, or logical, design of a database is an abstract model of the database from a business perspective, whereas the physical design shows how the database is actually arranged on direct-access storage devices.

Normalization and Entity-Relationship Diagrams

The conceptual database design describes how the data elements in the database are to be grouped. The design process identifies relationships among data elements and the most efficient way of grouping data elements together to meet business information requirements. The process also identifies redundant data elements and the groupings of data elements required for specific application programs. Groups of data are organized, refined, and streamlined until an overall logical view of the relationships among all the data in the database emerges.

To use a relational database model effectively, complex groupings of data must be streamlined to minimize redundant data elements and awkward many-to-many relationships. The process of creating small, stable, yet flexible and adaptive data structures from complex groups of data is called **normalization**. Figures 6.9 and 6.10 illustrate this process.

In the particular business modeled here, an order can have more than one part but each part is provided by only one supplier. If we build a relation called ORDER with all the fields included here, we would have to repeat the name

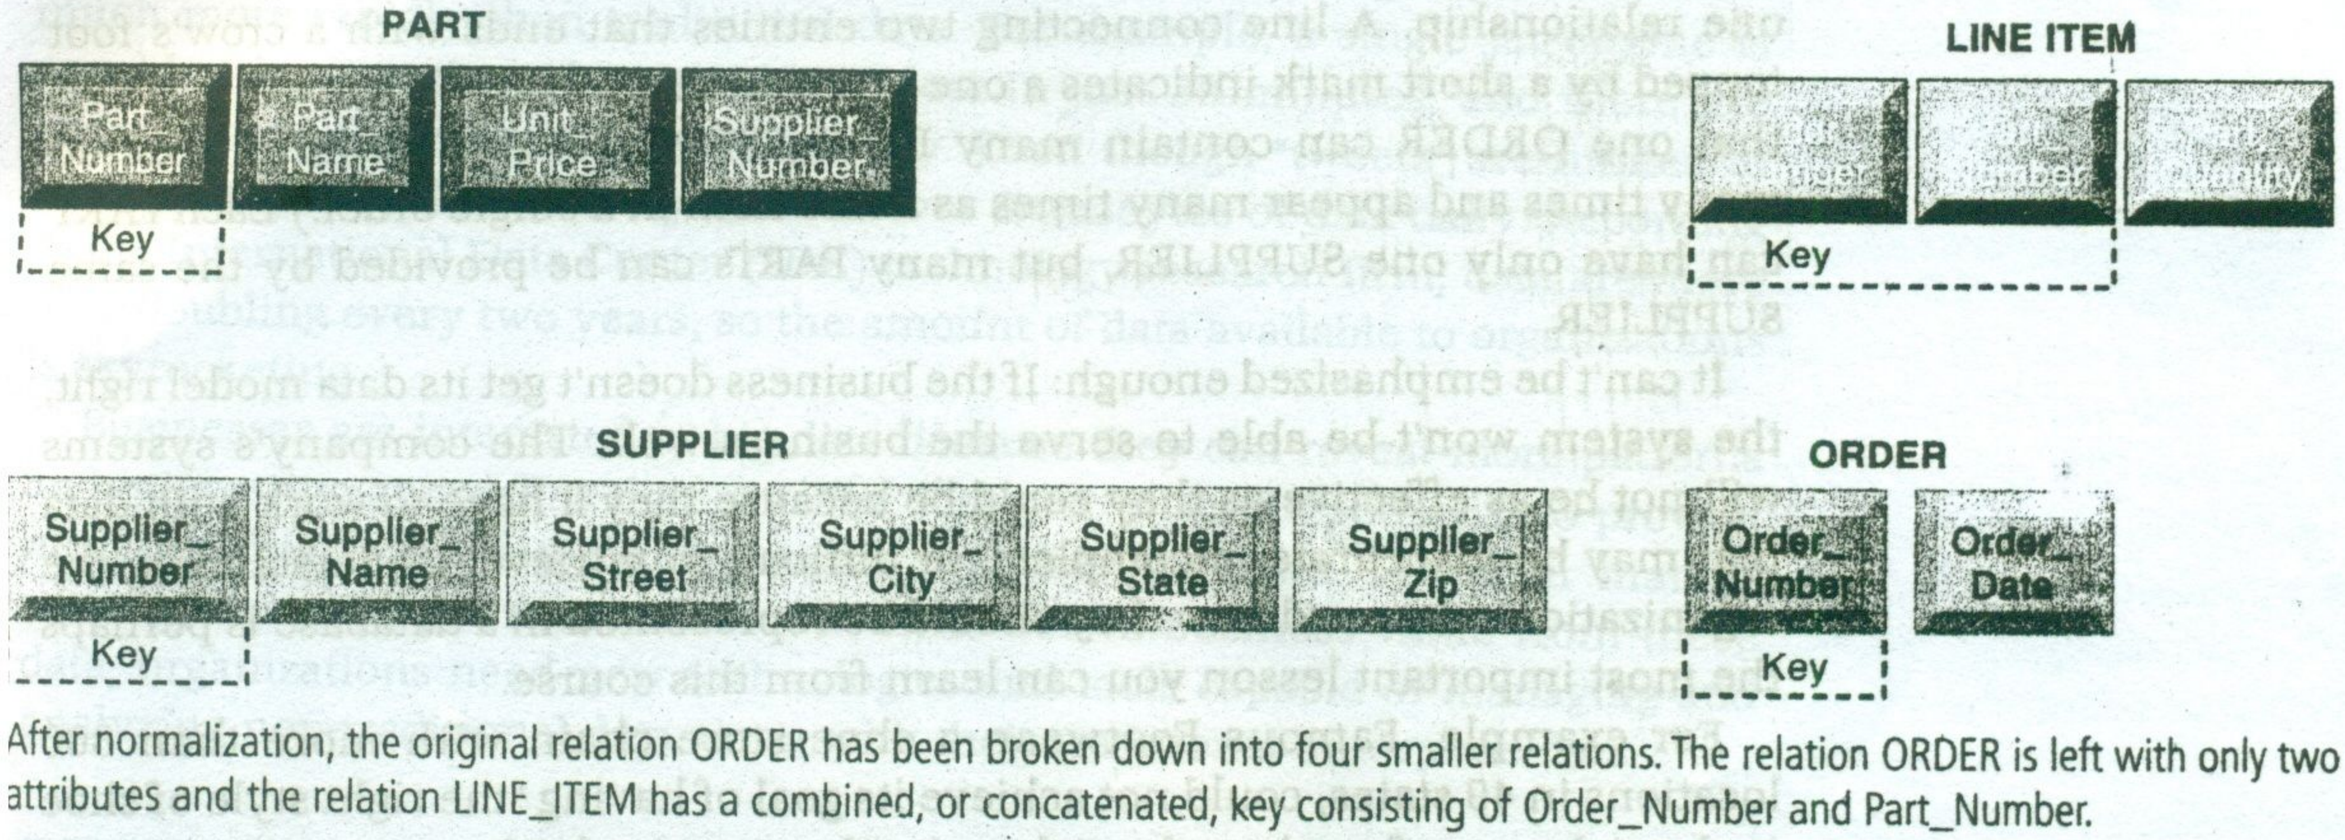
FIGURE 6.9 AN UNNORMALIZED RELATION FOR ORDER

ORDER (Before Normalization)

Order_ Number	Order_ Date	Part_ Number	Part_ Name	Unit_ Price	Part_ Quantity	Supplier_ Number	Supplier_ Name	Supplier_ Street	Supplier_ City	Supplier_ State	Supplier_ Zip

An unnormalized relation contains repeating groups. For example, there can be many parts and suppliers for each order. There is only a one-to-one correspondence between Order_Number and Order_Date.

FIGURE 6.10 NORMALIZED TABLES CREATED FROM ORDER



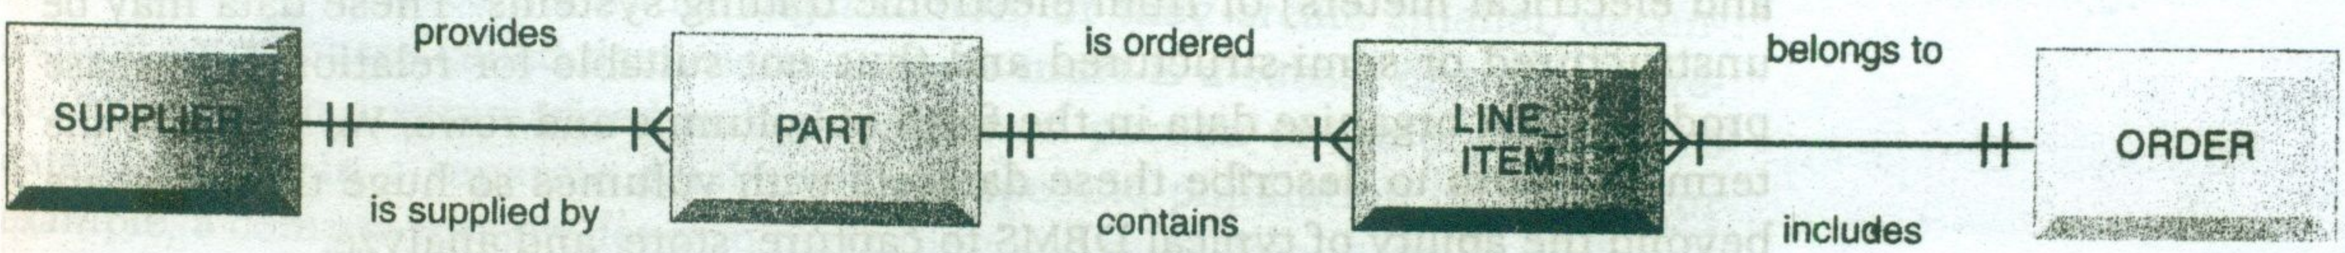
After normalization, the original relation ORDER has been broken down into four smaller relations. The relation ORDER is left with only two attributes and the relation LINE_ITEM has a combined, or concatenated, key consisting of Order_Number and Part_Number.

and address of the supplier for every part on the order, even though the order is for parts from a single supplier. This relationship contains what are called repeating data groups because there can be many parts on a single order to a given supplier. A more efficient way to arrange the data is to break down ORDER into smaller relations, each of which describes a single entity. If we go step by step and normalize the relation ORDER, we emerge with the relations illustrated in Figure 6.10. You can find out more about normalization, entity-relationship diagramming, and database design in the Learning Tracks for this chapter.

Relational database systems try to enforce **referential integrity** rules to ensure that relationships between coupled tables remain consistent. When one table has a foreign key that points to another table, you may not add a record to the table with the foreign key unless there is a corresponding record in the linked table. In the database we examined earlier in this chapter, the foreign key Supplier_Number links the PART table to the SUPPLIER table. We may not add a new record to the PART table for a part with Supplier_Number 8266 unless there is a corresponding record in the SUPPLIER table for Supplier_Number 8266. We must also delete the corresponding record in the PART table if we delete the record in the SUPPLIER table for Supplier_Number 8266. In other words, we shouldn't have parts from nonexistent suppliers!

Database designers document their data model with an **entity-relationship diagram**, illustrated in Figure 6.11. This diagram illustrates the relationship between the entities SUPPLIER, PART, LINE_ITEM, and ORDER. The boxes

FIGURE 6.11 AN ENTITY-RELATIONSHIP DIAGRAM



This diagram shows the relationships between the entities SUPPLIER, PART, LINE_ITEM, and ORDER that might be used to model the database in Figure 6.10.

represent entities. The lines connecting the boxes represent relationships. A line connecting two entities that ends in two short marks designates a one-to-one relationship. A line connecting two entities that ends with a crow's foot topped by a short mark indicates a one-to-many relationship. Figure 6.11 shows that one ORDER can contain many LINE_ITEMs. (A PART can be ordered many times and appear many times as a line item in a single order.) Each PART can have only one SUPPLIER, but many PARTs can be provided by the same SUPPLIER.

It can't be emphasized enough: If the business doesn't get its data model right, the system won't be able to serve the business well. The company's systems will not be as effective as they could be because they'll have to work with data that may be inaccurate, incomplete, or difficult to retrieve. Understanding the organization's data and how they should be represented in a database is perhaps the most important lesson you can learn from this course.

For example, Famous Footwear, a shoe store chain with more than 800 locations in 49 states, could not achieve its goal of having "the right style of shoe in the right store for sale at the right price" because its database was not properly designed for rapidly adjusting store inventory. The company had an Oracle relational database running on a midrange computer, but the database was designed primarily for producing standard reports for management rather than for reacting to marketplace changes. Management could not obtain precise data on specific items in inventory in each of its stores. The company had to work around this problem by building a new database where the sales and inventory data could be better organized for analysis and inventory management.

6.3 WHAT ARE THE PRINCIPAL TOOLS AND TECHNOLOGIES FOR ACCESSING INFORMATION FROM DATABASES TO IMPROVE BUSINESS PERFORMANCE AND DECISION MAKING?

Businesses use their databases to keep track of basic transactions, such as paying suppliers, processing orders, keeping track of customers, and paying employees. But they also need databases to provide information that will help the company run the business more efficiently, and help managers and employees make better decisions. If a company wants to know which product is the most popular or who is its most profitable customer, the answer lies in the data.

THE CHALLENGE OF BIG DATA

Most data collected by organizations used to be transaction data that could easily fit into rows and columns of relational database management systems. We are now witnessing an explosion of data from Web traffic, e-mail messages, and social media content (tweets, status messages), as well as machine-generated data from sensors (used in smart meters, manufacturing sensors, and electrical meters) or from electronic trading systems. These data may be unstructured or semi-structured and thus not suitable for relational database products that organize data in the form of columns and rows. We now use the term **big data** to describe these datasets with volumes so huge that they are beyond the ability of typical DBMS to capture, store, and analyze.

Big data doesn't refer to any specific quantity, but usually refers to data in the petabyte and exabyte range—in other words, billions to trillions of records,

all from different sources. Big data are produced in much larger quantities and much more rapidly than traditional data. For example, a single jet engine is capable of generating 10 terabytes of data in just 30 minutes, and there are more than 25,000 airline flights each day. Even though "tweets" are limited to 140 characters each, Twitter generates over 8 terabytes of data daily. According to the International Data Center (IDC) technology research firm, data are more than doubling every two years, so the amount of data available to organizations is skyrocketing.

Businesses are interested in big data because they can reveal more patterns and interesting anomalies than smaller data sets, with the potential to provide new insights into customer behavior, weather patterns, financial market activity, or other phenomena. However, to derive business value from these data, organizations need new technologies and tools capable of managing and analyzing non-traditional data along with their traditional enterprise data.

BUSINESS INTELLIGENCE INFRASTRUCTURE

Suppose you wanted concise, reliable information about current operations, trends, and changes across the entire company. If you worked in a large company, the data you need might have to be pieced together from separate systems, such as sales, manufacturing, and accounting, and even from external sources, such as demographic or competitor data. Increasingly, you might need to use big data. A contemporary infrastructure for business intelligence has an array of tools for obtaining useful information from all the different types of data used by businesses today, including semi-structured and unstructured big data in vast quantities. These capabilities include data warehouses and data marts, Hadoop, in-memory computing, and analytical platforms. Some of these capabilities are available as cloud services.

Data Warehouses and Data Marts

The traditional tool for analyzing corporate data for the past two decades has been the data warehouse. A **data warehouse** is a database that stores current and historical data of potential interest to decision makers throughout the company. The data originate in many core operational transaction systems, such as systems for sales, customer accounts, and manufacturing, and may include data from Web site transactions. The data warehouse extracts current and historical data from multiple operational systems inside the organization. These data are combined with data from external sources and transformed by correcting inaccurate and incomplete data and restructuring the data for management reporting and analysis before being loaded into the data warehouse.

The data warehouse makes the data available for anyone to access as needed, but it cannot be altered. A data warehouse system also provides a range of ad hoc and standardized query tools, analytical tools, and graphical reporting facilities .

Companies often build enterprise-wide data warehouses, where a central data warehouse serves the entire organization, or they create smaller, decentralized warehouses called data marts. A **data mart** is a subset of a data warehouse in which a summarized or highly focused portion of the organization's data is placed in a separate database for a specific population of users. For example, a company might develop marketing and sales data marts to deal with customer information. Bookseller Barnes & Noble used to maintain a series of data marts—one for point-of-sale data in retail stores, another for college bookstore sales, and a third for online sales.

Hadoop

Relational DBMS and data warehouse products are not well-suited for organizing and analyzing big data or data that do not easily fit into columns and rows used in their data models. For handling unstructured and semi-structured data in vast quantities, as well as structured data, organizations are using **Hadoop**. Hadoop is an open source software framework managed by the Apache Software Foundation that enables distributed parallel processing of huge amounts of data across inexpensive computers. It breaks a big data problem down into sub-problems, distributes them among up to thousands of inexpensive computer processing nodes, and then combines the result into a smaller data set that is easier to analyze. You've probably used Hadoop to find the best airfare on the Internet, get directions to a restaurant, do a search on Google, or connect with a friend on Facebook.

Hadoop consists of several key services: the Hadoop Distributed File System (HDFS) for data storage and MapReduce for high-performance parallel data processing. HDFS links together the file systems on the numerous nodes in a Hadoop cluster to turn them into one big file system. Hadoop's MapReduce was inspired by Google's MapReduce system for breaking down processing of huge datasets and assigning work to the various nodes in a cluster. HBase, Hadoop's non-relational database, provides rapid access to the data stored on HDFS and a transactional platform for running high-scale real-time applications.

Hadoop can process large quantities of any kind of data, including structured transactional data, loosely structured data such as Facebook and Twitter feeds, complex data such as Web server log files, and unstructured audio and video data. Hadoop runs on a cluster of inexpensive servers, and processors can be added or removed as needed. Companies use Hadoop for analyzing very large volumes of data as well as for a staging area for unstructured and semi-structured data before they are loaded into a data warehouse. Facebook stores much of its data on its massive Hadoop cluster, which holds an estimated 100 petabytes, about 10,000 times more information than the Library of Congress. Yahoo uses Hadoop to track user behavior so it can modify its home page to fit their interests. Life sciences research firm NextBio uses Hadoop and HBase to process data for pharmaceutical companies conducting genomic research. Top database vendors such as IBM, Hewlett-Packard, Oracle, and Microsoft have their own Hadoop software distributions. Other vendors offer tools for moving data into and out of Hadoop or for analyzing data within Hadoop.

In-Memory Computing

Another way of facilitating big data analysis is to use **in-memory computing**, which relies primarily on a computer's main memory (RAM) for data storage. (Conventional DBMS use disk storage systems.) Users access data stored in system primary memory, thereby eliminating bottlenecks from retrieving and reading data in a traditional, disk-based database and dramatically shortening query response times. In-memory processing makes it possible for very large sets of data, amounting to the size of a data mart or small data warehouse, to reside entirely in memory. Complex business calculations that used to take hours or days are able to be completed within seconds, and this can even be accomplished on handheld devices. (See the Interactive Session on Technology.)

The previous chapter describes some of the advances in contemporary computer hardware technology that make in-memory processing possible, such as powerful high-speed processors, multicore processing, and falling computer memory prices. These technologies help companies optimize the use of memory and accelerate processing performance while lowering costs.

INTERACTIVE SESSION: TECHNOLOGY

DRIVING ARI FLEET MANAGEMENT WITH REAL-TIME ANALYTICS

Automotive Resources International®, better known as simply ARI®, is the world's largest privately-held company for vehicle fleet management services. ARI is headquartered in Mt. Laurel, New Jersey and has 2,500 employees and offices throughout North America, Europe, the UK, and Hong Kong. The company manages more than 1,000,000 vehicles in the U.S., Canada, Mexico, Puerto Rico and Europe.

Businesses that need vehicles for shipments (trucks, vans, cars, ships, and rail cars) may choose to manage their own fleet of vehicles or they may outsource fleet management to companies such as ARI which specialize in these services. ARI manages the entire life cycle and operation of a fleet of vehicles for its customers, from up-front specification and acquisition to resale, including financing, maintenance, fuel management, and risk management services such as driver safety training and accident management. ARI also maintains six call centers in North America that operate 24/7, 365 days a year to support customers' fleet operations by providing assistance regarding repairs, breakdowns, accident response, preventive maintenance, and other driver needs. These call centers handle about 3.5 million calls per year from customers, drivers, and suppliers who expect access to real-time actionable information.

Providing this information has become increasingly challenging. Operating a single large commercial vehicle fleet generates high volumes of complex data, such as data on fuel consumption, maintenance, licensing, and compliance. A fuel transaction, for example, requires data on state taxes paid, fuel grade, total sale, amount sold, and time and place of purchase. A simple brake job and preventive maintenance checkup generates dozens of records for each component that is serviced. Each part and service performed on a vehicle is tracked using American Trucking Association codes. ARI collects and analyzes over 14,000 pieces of data per vehicle. Then multiply the data by hundreds of fleets, some with up to 10,000 vehicles, all operating simultaneously throughout the globe, and you'll have an idea of the enormous volume of data ARI needs to manage, both for itself and for its customers.

ARI provided its customers with detailed information about their fleet operations, but the type of information it could deliver was very limited. For example, ARI could generate detailed reports on line-item expenditures, vehicle purchases, maintenance records, and other operational information presented as simple spreadsheets, charts, or graphs, but it was not possible to analyze all the data to spot trends and make recommendations. ARI was able to analyze data customer by customer, but it was not able to aggregate data across its entire customer base. For instance, if ARI was managing a pharmaceutical company's vehicle fleet, its information systems could not benchmark that fleet's performance against others in the industry. That type of problem required too much manual work and time, and still didn't deliver the level of insight management thought was possible.

What's more, in order to create reports, ARI had to go through internal subject matter experts in various aspects of fleet operations, who were called "reporting power users." Every request for information was passed to these power users. A request for a report would take 5 days to fill. If the report was unsatisfactory, it would go back to the report writer to make changes. ARI's process for analyzing its data was extremely drawn out.

In mid-2011, ARI implemented SAP BusinessObjects Explorer to give customers the enhanced ability to access data and run their own reports. SAP BusinessObjects Explorer is a business intelligence tool that enables business users to view, sort and analyze business intelligence data. Users search through data sources using an iTunes-like interface. They do not have to create queries to search the data and results are shown with a chart that indicates the best information match. The graphical representation of results changes as the user asks further questions of the data.

In early 2012, ARI integrated SAP BusinessObjects Explorer with HANA, SAP's in-memory computing platform that is deployable as an on-premise appliance (hardware and software) or in the cloud. HANA is optimized for performing real-time analytics and handling very high volumes of operational and transactional data in real time. HANA's in-memory analytics queries data stored in random access memory (RAM) instead of on a hard disk or flash storage.

Things started happening quickly after that. When ARI's controller wanted an impact analysis of the company's top 10 customers, SAP HANA produced the result in 3 to 3 ½ seconds. In ARI's old systems

environment, this task would have been assigned to a power user versed in using reporting tools, specifications would have to be drawn up and a program designed for that specific query, a process that would have taken about 36 hours.

Using HANA, ARI is now able to quickly mine its vast data resources and generate predictions based on the results. For example, the company can produce precise figures on what it costs to operate a fleet of a certain size over a particular route across specific industries during a certain type of weather and predict what the impact of changes in any of these variables. And it can do so nearly as easily as providing customers with a simple history of their expenditures on fuel. With such helpful information ARI provides more value to its customers.

HANA has also reduced the time required for each transaction handled by ARI's call centers—from the time a call center staffer takes a call to retrieving and delivering the requested information—by 5 percent. Since call center staff account for 40 percent of ARI's direct overhead, that time reduction translates into major cost savings.

ARI plans to make some of these real-time reporting and analytic capabilities available on mobile devices, which will enable customers to instantly approve a variety of operational procedures, such as authorizing maintenance repairs. Customers will also be able to use the mobile tools for instant insight into their fleet operations, down to a level of detail such as a specific vehicle's tire history.

Sources: "Driving 2 Million Vehicles with SAP Data," www.sap.com, accessed February 1, 2014; www.arifleet.com, accessed February 1, 2014; and "ARI Fleet Management Drives Real-Time Analytics to Customers," SAP InsiderPROFILES, April 1, 2013.

CASE STUDY QUESTIONS

1. Why was data management so problematic at ARI?
2. Describe ARI's earlier capabilities for data analysis and reporting and their impact on the business.
3. Was SAP HANA a good solution for ARI? Why or why not?
4. Describe the changes in the business as a result of adopting HANA.

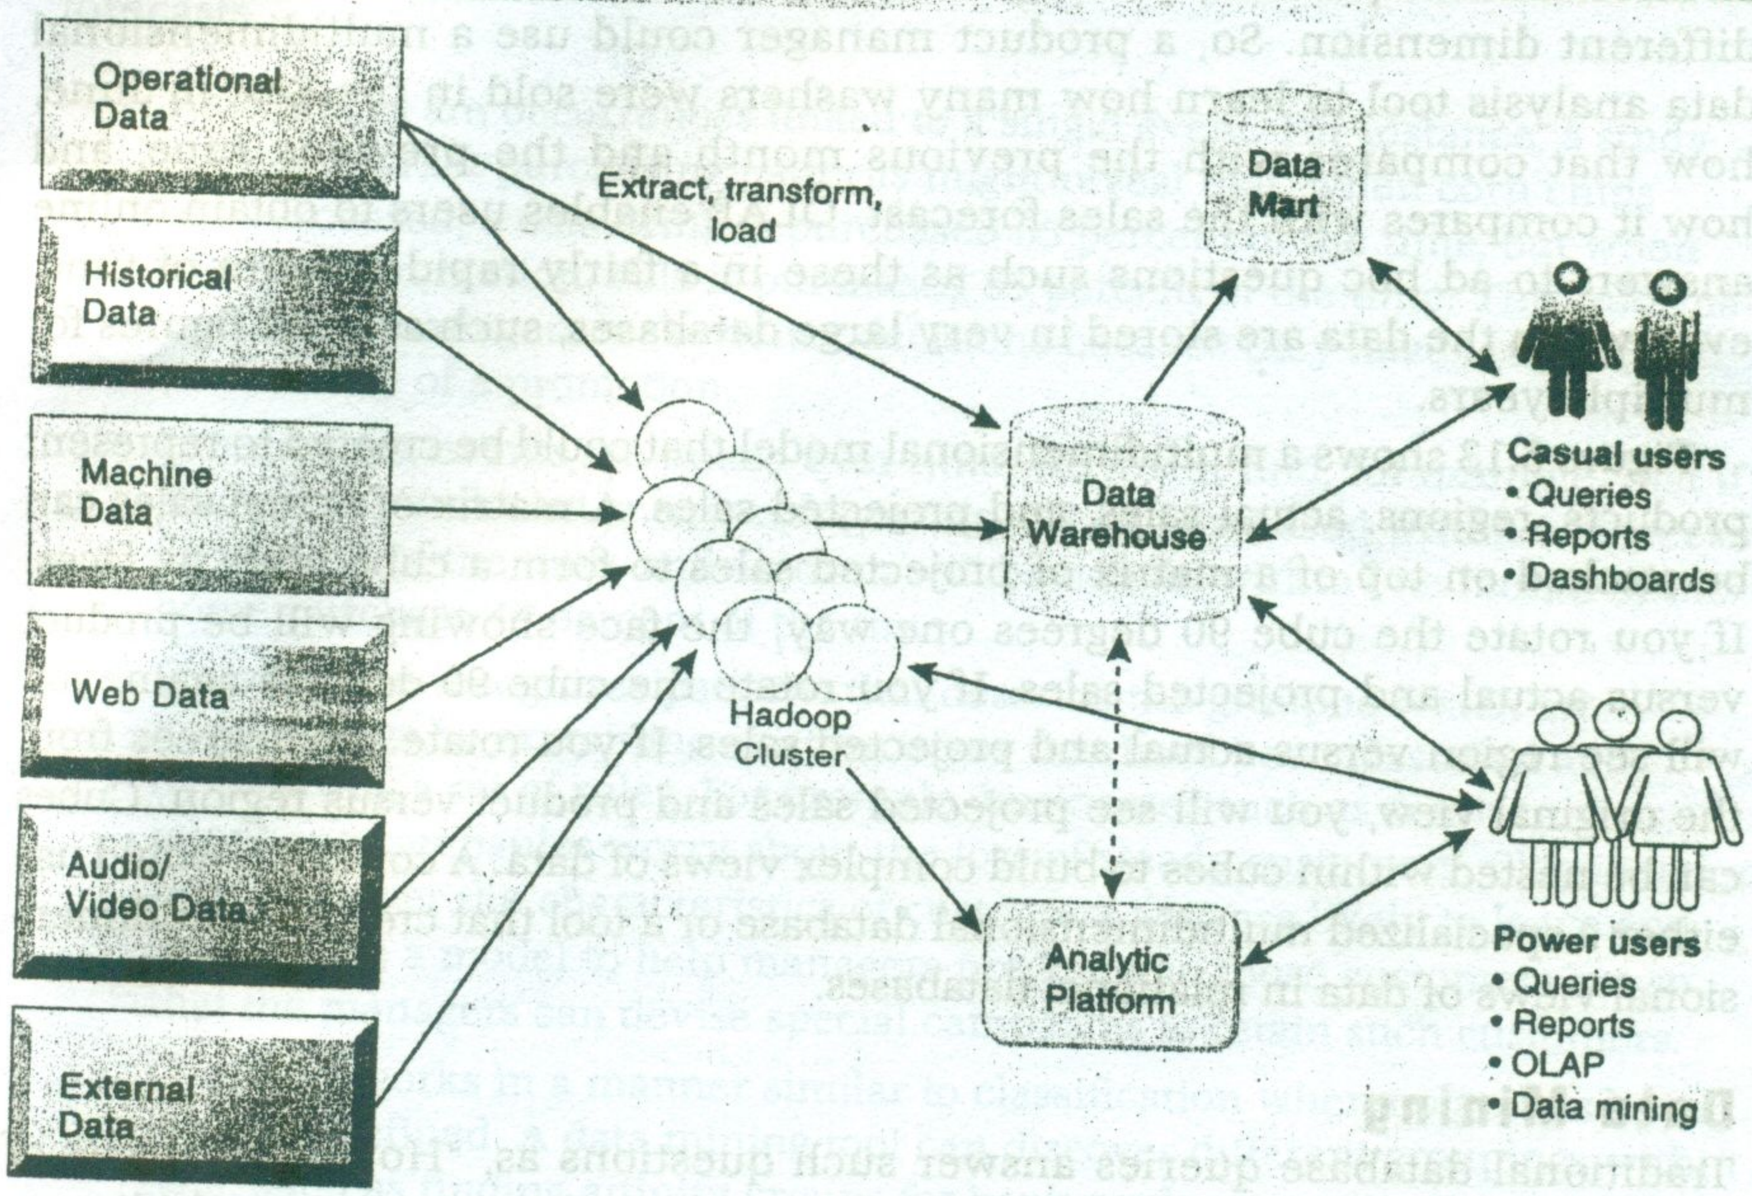
Leading commercial products for in-memory computing include SAP's High Performance Analytics Appliance (HANA) and Oracle Exalytics. Each provides a set of integrated software components, including in-memory database software and specialized analytics software, that run on hardware optimized for in-memory computing work.

Analytic Platforms

Commercial database vendors have developed specialized high-speed **analytic platforms** using both relational and non-relational technology that are optimized for analyzing large datasets. Analytic platforms such as IBM Netezza and Oracle Exadata feature preconfigured hardware-software systems that are specifically designed for query processing and analytics. For example, IBM Netezza features tightly integrated database, server, and storage components that handle complex analytic queries 10 to 100 times faster than traditional systems. Analytic platforms also include in-memory systems and NoSQL non-relational database management systems. Analytic platforms are now available as cloud services.

Figure 6.12 illustrates a contemporary business intelligence infrastructure using the technologies we have just described. Current and historical data are extracted from multiple operational systems along with Web data, machine-generated data, unstructured audio/visual data, and data from external sources that's been restructured and reorganized for reporting and analysis. Hadoop clusters pre-process big data for use in the data warehouse, data marts, or an

FIGURE 6.12 **CONTEMPORARY BUSINESS INTELLIGENCE INFRASTRUCTURE**



A contemporary business intelligence infrastructure features capabilities and tools to manage and analyze large quantities and different types of data from multiple sources. Easy-to-use query and reporting tools for casual business users and more sophisticated analytical toolsets for power users are included.

analytic platform, or for direct querying by power users. Outputs include reports and dashboards as well as query results. Chapter 12 discusses the various types of BI users and BI reporting in greater detail.

ANALYTICAL TOOLS: RELATIONSHIPS, PATTERNS, TRENDS

Once data have been captured and organized using the business intelligence technologies we have just described, they are available for further analysis using software for database querying and reporting, multidimensional data analysis (OLAP), and data mining. This section will introduce you to these tools, with more detail about business intelligence analytics and applications in Chapter 12.

Online Analytical Processing (OLAP)

Suppose your company sells four different products—nuts, bolts, washers, and screws—in the East, West, and Central regions. If you wanted to ask a fairly straightforward question, such as how many washers sold during the past quarter, you could easily find the answer by querying your sales database. But what if you wanted to know how many washers sold in each of your sales regions and compare actual results with projected sales?

To obtain the answer, you would need **online analytical processing (OLAP)**. OLAP supports multidimensional data analysis, enabling users to view the same data in different ways using multiple dimensions. Each aspect

of information—product, pricing, cost, region, or time period—represents a different dimension. So, a product manager could use a multidimensional data analysis tool to learn how many washers were sold in the East in June, how that compares with the previous month and the previous June, and how it compares with the sales forecast. OLAP enables users to obtain online answers to ad hoc questions such as these in a fairly rapid amount of time, even when the data are stored in very large databases, such as sales figures for multiple years.

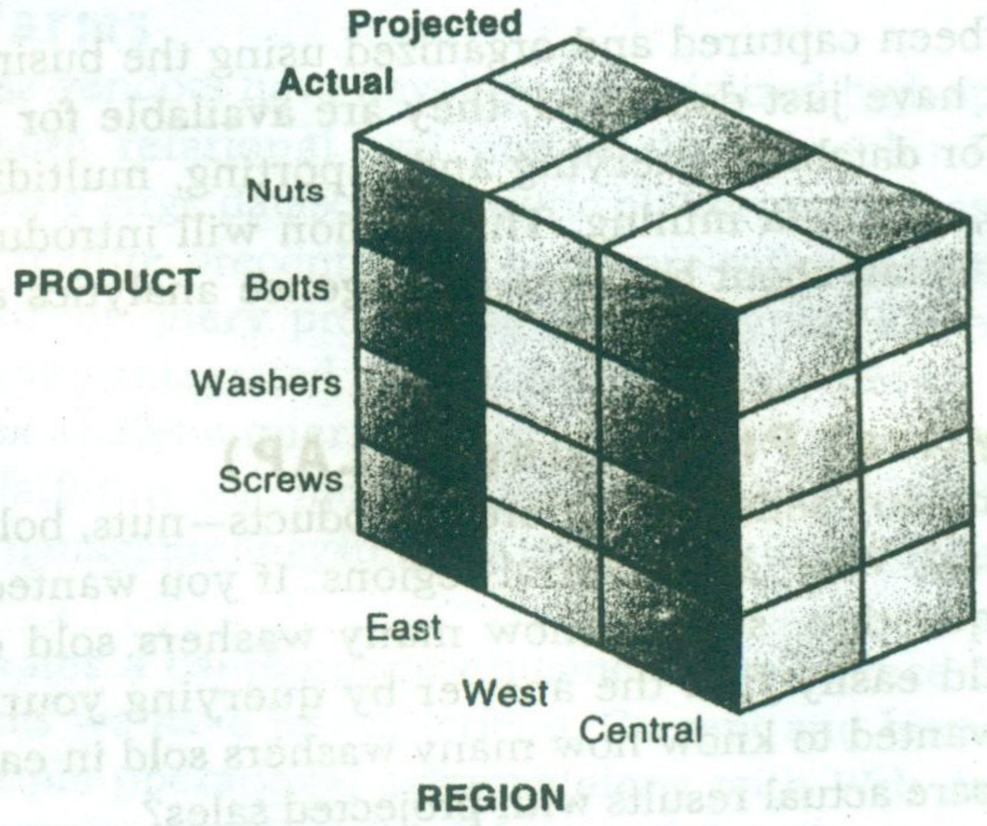
Figure 6.13 shows a multidimensional model that could be created to represent products, regions, actual sales, and projected sales. A matrix of actual sales can be stacked on top of a matrix of projected sales to form a cube with six faces. If you rotate the cube 90 degrees one way, the face showing will be product versus actual and projected sales. If you rotate the cube 90 degrees again, you will see region versus actual and projected sales. If you rotate 180 degrees from the original view, you will see projected sales and product versus region. Cubes can be nested within cubes to build complex views of data. A company would use either a specialized multidimensional database or a tool that creates multidimensional views of data in relational databases.

Data Mining

Traditional database queries answer such questions as, "How many units of product number 403 were shipped in February 2013?" OLAP, or multidimensional analysis, supports much more complex requests for information, such as, "Compare sales of product 403 relative to plan by quarter and sales region for the past two years." With OLAP and query-oriented data analysis, users need to have a good idea about the information for which they are looking.

Data mining is more discovery-driven. Data mining provides insights into corporate data that cannot be obtained with OLAP by finding hidden patterns and relationships in large databases and inferring rules from them to predict future behavior. The patterns and rules are used to guide decision making and forecast the effect of those decisions. The types of information obtainable

FIGURE 6.13 MULTIDIMENSIONAL DATA MODEL



This view shows product versus region. If you rotate the cube 90 degrees, the face that will show is product versus actual and projected sales. If you rotate the cube 90 degrees again, you will see region versus actual and projected sales. Other views are possible.

from data mining include associations, sequences, classifications, clusters, and forecasts.

- *Associations* are occurrences linked to a single event. For instance, a study of supermarket purchasing patterns might reveal that, when corn chips are purchased, a cola drink is purchased 65 percent of the time, but when there is a promotion, cola is purchased 85 percent of the time. This information helps managers make better decisions because they have learned the profitability of a promotion.
- In *sequences*, events are linked over time. We might find, for example, that if a house is purchased, a new refrigerator will be purchased within two weeks 65 percent of the time, and an oven will be bought within one month of the home purchase 45 percent of the time.
- *Classification* recognizes patterns that describe the group to which an item belongs by examining existing items that have been classified and by inferring a set of rules. For example, businesses such as credit card or telephone companies worry about the loss of steady customers. Classification helps discover the characteristics of customers who are likely to leave and can provide a model to help managers predict who those customers are so that the managers can devise special campaigns to retain such customers.
- *Clustering* works in a manner similar to classification when no groups have yet been defined. A data mining tool can discover different groupings within data, such as finding affinity groups for bank cards or partitioning a database into groups of customers based on demographics and types of personal investments.
- Although these applications involve predictions, *forecasting* uses predictions in a different way. It uses a series of existing values to forecast what other values will be. For example, forecasting might find patterns in data to help managers estimate the future value of continuous variables, such as sales figures.

These systems perform high-level analyses of patterns or trends, but they can also drill down to provide more detail when needed. There are data mining applications for all the functional areas of business, and for government and scientific work. One popular use for data mining is to provide detailed analyses of patterns in customer data for one-to-one marketing campaigns or for identifying profitable customers.

Entertainment, formerly known as Harrah's Entertainment, is the largest gaming company in the world. It continually analyzes data about its customers gathered when people play its slot machines or use its casinos and hotels. The corporate marketing department uses this information to build a detailed gambling profile, based on a particular customer's ongoing value to the company. For instance, data mining lets Caesars know the favorite gaming experience of a regular customer at one of its riverboat casinos, along with that person's preferences for room accommodations, restaurants, and entertainment. This information guides management decisions about how to cultivate the most profitable customers, encourage those customers to spend more, and attract more customers with high revenue-generating potential. Business intelligence improved Caesars's profits so much that it became the centerpiece of the firm's business strategy.

Text Mining and Web Mining

Unstructured data, most in the form of text files, is believed to account for over 80 percent of useful organizational information and is one of the major sources of

big data that firms want to analyze. E-mail, memos, call center transcripts, survey responses, legal cases, patent descriptions, and service reports are all valuable for finding patterns and trends that will help employees make better business decisions. **Text mining** tools are now available to help businesses analyze these data. These tools are able to extract key elements from unstructured big data sets, discover patterns and relationships, and summarize the information.

Businesses might turn to text mining to analyze transcripts of calls to customer service centers to identify major service and repair issues or to measure customer sentiment about their company. **Sentiment analysis** software is able to mine text comments in an e-mail message, blog, social media conversation, or survey form to detect favorable and unfavorable opinions about specific subjects.

For example, the discount broker Charles Schwab uses Attensity Analyze software to analyze hundreds of thousands of its customer interactions each month. The software analyzes Schwab's customer service notes, e-mails, survey responses, and online discussions to discover signs of dissatisfaction that might cause a customer to stop using the company's services. Attensity is able to automatically identify the various "voices" customers use to express their feedback (such as a positive, negative, or conditional voice) to pinpoint a person's intent to buy, intent to leave, or reaction to a specific product or marketing message. Schwab uses this information to take corrective actions such as stepping up direct broker communication with the customer and trying to quickly resolve the problems that are making the customer unhappy.

The Web is another rich source of unstructured big data for revealing patterns, trends, and insights into customer behavior. The discovery and analysis of useful patterns and information from the World Wide Web is called **Web mining**. Businesses might turn to Web mining to help them understand customer behavior, evaluate the effectiveness of a particular Web site, or quantify the success of a marketing campaign. For instance, marketers use the Google Trends and Google Insights for Search services, which track the popularity of various words and phrases used in Google search queries, to learn what people are interested in and what they are interested in buying.

Web mining looks for patterns in data through content mining, structure mining, and usage mining. Web content mining is the process of extracting knowledge from the content of Web pages, which may include text, image, audio, and video data. Web structure mining examines data related to the structure of a particular Web site. For example, links pointing to a document indicate the popularity of the document, while links coming out of a document indicate the richness or perhaps the variety of topics covered in the document. Web usage mining examines user interaction data recorded by a Web server whenever requests for a Web site's resources are received. The usage data records the user's behavior when the user browses or makes transactions on the Web site and collects the data in a server log. Analyzing such data can help companies determine the value of particular customers, cross marketing strategies across products, and the effectiveness of promotional campaigns.

The chapter-ending case describes organizations' experiences as they use the analytical tools and business intelligence technologies we have described to grapple with "big data" challenges.

DATABASES AND THE WEB

Have you ever tried to use the Web to place an order or view a product catalog? If so, you were probably using a Web site linked to an internal

corporate database. Many companies now use the Web to make some of the information in their internal databases available to customers and business partners.

Suppose, for example, a customer with a Web browser wants to search an online retailer's database for pricing information. Figure 6.14 illustrates how that customer might access the retailer's internal database over the Web. The user accesses the retailer's Web site over the Internet using Web browser software on his or her client PC. The user's Web browser software requests data from the organization's database, using HTML commands to communicate with the Web server.

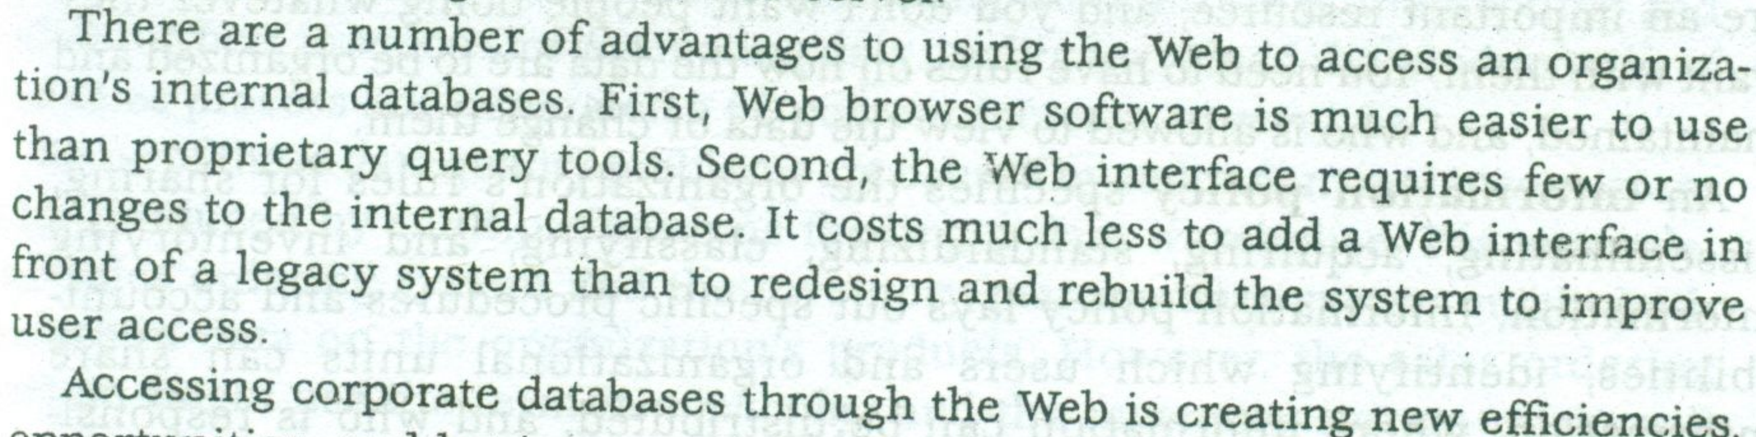
Because many back-end databases cannot interpret commands written in HTML, the Web server passes these requests for data to software that translates HTML commands into SQL so the commands can be processed by the DBMS working with the database. In a client/server environment, the DBMS resides on a dedicated computer called a **database server**. The DBMS receives the SQL requests and provides the required data. Middleware transfers information from the organization's internal database back to the Web server for delivery in the form of a Web page to the user.

Figure 6.14 shows that the middleware working between the Web server and the DBMS is an application server running on its own dedicated computer (see Chapter 5). The application server software handles all application operations, including transaction processing and data access, between browser-based computers and a company's back-end business applications or databases. The application server takes requests from the Web server, runs the business logic to process transactions based on those requests, and provides connectivity to the organization's back-end systems or databases. Alternatively, the software for handling these operations could be a custom program or a CGI script. A CGI script is a compact program using the *Common Gateway Interface (CGI)* specification for processing data on a Web server.

There are a number of advantages to using the Web to access an organization's internal databases. First, Web browser software is much easier to use than proprietary query tools. Second, the Web interface requires few or no changes to the internal database. It costs much less to add a Web interface in front of a legacy system than to redesign and rebuild the system to improve user access.

Accessing corporate databases through the Web is creating new efficiencies, opportunities, and business models. ThomasNet.com provides an up-to-date online directory of more than 700,000 suppliers of industrial products, such as chemicals, metals, plastics, rubber, and automotive equipment. Formerly called

FIGURE 6.14 LINKING INTERNAL DATABASES TO THE WEB



Users access an organization's internal database through the Web using their desktop PCs and Web browser software.

Thomas Register, the company used to send out huge paper catalogs with this information. Now it provides this information to users online via its Web site and has become a smaller, leaner company.

Other companies have created entirely new businesses based on access to large databases through the Web. One is the social networking service Facebook, which helps users stay connected with each other and meet new people. Facebook features "profiles" with information on more than 1.3 billion active users with information about themselves, including interests, friends, photos, and groups with which they are affiliated. Facebook maintains a massive database to house and manage all of this content. There are also many Web-enabled databases in the public sector to help consumers and citizens access helpful information.

6.4 WHY ARE INFORMATION POLICY, DATA ADMINISTRATION, AND DATA QUALITY ASSURANCE ESSENTIAL FOR MANAGING THE FIRM'S DATA RESOURCES?

Setting up a database is only a start. In order to make sure that the data for your business remain accurate, reliable, and readily available to those who need it, your business will need special policies and procedures for data management.

ESTABLISHING AN INFORMATION POLICY

Every business, large and small, needs an information policy. Your firm's data are an important resource, and you don't want people doing whatever they want with them. You need to have rules on how the data are to be organized and maintained, and who is allowed to view the data or change them.

An **information policy** specifies the organization's rules for sharing, disseminating, acquiring, standardizing, classifying, and inventorying information. Information policy lays out specific procedures and accountabilities, identifying which users and organizational units can share information, where information can be distributed, and who is responsible for updating and maintaining the information. For example, a typical information policy would specify that only selected members of the payroll and human resources department would have the right to change and view sensitive employee data, such as an employee's salary or social security number, and that these departments are responsible for making sure that such employee data are accurate.

If you are in a small business, the information policy would be established and implemented by the owners or managers. In a large organization, managing and planning for information as a corporate resource often requires a formal data administration function. **Data administration** is responsible for the specific policies and procedures through which data can be managed as an organizational resource. These responsibilities include developing information policy, planning for data, overseeing logical database design and data dictionary development, and monitoring how information systems specialists and end-user groups use data.

You may hear the term **data governance** used to describe many of these activities. Promoted by IBM, data governance deals with the policies and

processes for managing the availability, usability, integrity, and security of the data employed in an enterprise, with special emphasis on promoting privacy, security, data quality, and compliance with government regulations.

A large organization will also have a database design and management group within the corporate information systems division that is responsible for defining and organizing the structure and content of the database, and maintaining the database. In close cooperation with users, the design group establishes the physical database, the logical relations among elements, and the access rules and security procedures. The functions it performs are called **database administration**.

ENSURING DATA QUALITY

A well-designed database and information policy will go a long way toward ensuring that the business has the information it needs. However, additional steps must be taken to ensure that the data in organizational databases are accurate and remain reliable.

What would happen if a customer's telephone number or account balance were incorrect? What would be the impact if the database had the wrong price for the product you sold or your sales system and inventory system showed different prices for the same product? Data that are inaccurate, untimely, or inconsistent with other sources of information lead to incorrect decisions, product recalls, and financial losses. Gartner Inc. reported that more than 25 percent of the critical data in large Fortune 1000 companies' databases is inaccurate or incomplete, including bad product codes and product descriptions, faulty inventory descriptions, erroneous financial data, incorrect supplier information, and incorrect employee data. A Sirius Decisions study on "The Impact of Bad Data on Demand Creation" found that 10 to 25 percent of customer and prospect records contain critical data errors. Correcting these errors at their source and following best practices for promoting data quality increased the productivity of the sales process and generated a 66 percent increase in revenue.

Some of these data quality problems are caused by redundant and inconsistent data produced by multiple systems feeding a data warehouse. For example, the sales ordering system and the inventory management system might both maintain data on the organization's products. However, the sales ordering system might use the term *Item Number* and the inventory system might call the same attribute *Product Number*. The sales, inventory, or manufacturing systems of a clothing retailer might use different codes to represent values for an attribute. One system might represent clothing size as "extra large," whereas the other system might use the code "XL" for the same purpose. During the design process for the warehouse database, data describing entities, such as a customer, product, or order, should be named and defined consistently for all business areas using the database.

Think of all the times you've received several pieces of the same direct mail advertising on the same day. This is very likely the result of having your name maintained multiple times in a database. Your name may have been misspelled or you used your middle initial on one occasion and not on another or the information was initially entered onto a paper form and not scanned properly into the system. Because of these inconsistencies, the database would treat you as different people! We often receive redundant mail addressed to Laudon, Lavdon, Lauden, or Landon.

If a database is properly designed and enterprise-wide data standards established, duplicate or inconsistent data elements should be minimal.

Most data quality problems, however, such as misspelled names, transposed numbers, or incorrect or missing codes, stem from errors during data input. The incidence of such errors is rising as companies move their businesses to the Web and allow customers and suppliers to enter data into their Web sites that directly update internal systems.

Before a new database is in place, organizations need to identify and correct their faulty data and establish better routines for editing data once their database is in operation. Analysis of data quality often begins with a **data quality audit**, which is a structured survey of the accuracy and level of completeness of the data in an information system. Data quality audits can be performed by surveying entire data files, surveying samples from data files, or surveying end users for their perceptions of data quality.

Data cleansing, also known as *data scrubbing*, consists of activities for detecting and correcting data in a database that are incorrect, incomplete, improperly formatted, or redundant. Data cleansing not only corrects errors but also enforces consistency among different sets of data that originated in separate information systems. Specialized data-cleansing software is available to automatically survey data files, correct errors in the data, and integrate the data in a consistent company-wide format.

Data quality problems are not just business problems. They also pose serious problems for individuals, affecting their financial condition and even their jobs. For example, inaccurate or outdated data about consumers' credit histories maintained by credit bureaus can prevent creditworthy individuals from obtaining loans or lower their chances of finding or keeping a job.

The Interactive Session on Management illustrates American Water's experience with managing data as a resource. As you read this case, try to identify the policies, procedures, and technologies that were required to improve data management at this company.

INTERACTIVE SESSION: MANAGEMENT

AMERICAN WATER KEEPS DATA FLOWING

American Water, founded in 1886, is the largest public water utility in the United States. Headquartered in Voorhees, N.J., the company employs more than 7,000 dedicated professionals who provide drinking water, wastewater and other related services to approximately 16 million people in 35 states, as well as Ontario and Manitoba, Canada. Most of American Water's services support locally-managed utility subsidiaries that are regulated by the U.S. state in which each operates as well as the federal government. American Water also owns subsidiaries that manage municipal drinking water and wastewater systems under contract and others that supply businesses and residential communities with water management products and services.

Until recently, American water's systems and business processes were very localized, and many of these processes were manual. Over time, this information environment became increasingly difficult to manage. Many systems were not integrated, so that running any type of report that had to provide information about more than one region was a heavily manual process. Data had to be extracted from the systems supporting each region and then combined manually to create the desired output. When the company was preparing to hold an initial public offering of its stock in 2006, its software systems could not handle the required regulatory controls, so roughly 80 percent of this work had to be performed manually. It was close to a nightmare.

Management wanted to change the company from a decentralized group of independent regional businesses into a more centralized organization with standard company-wide business processes and enterprise-wide reporting. The first step toward achieving this goal was to implement an enterprise resource planning (ERP) system designed to replace disparate systems with a single integrated software platform. The company selected SAP as its ERP system vendor.

An important step of this project was to migrate the data from American Water's old systems to the new platform. The company's data resided in many different systems in various formats. Each regional business maintained some of its own data in its own systems, and a portion of these data were redundant and inconsistent. For example, there were duplicate pieces of materials master data because a material might be called one thing in the company's Missouri operation and another in its New Jersey business. These names had to be standardized so that the same name for a piece of data was used by every business unit. American Water's business users had to buy into this new company-wide view of data.

Data migration entails much more than just transferring data between old and new systems. Business users need to know that data are not just a responsibility of the information systems department: the business "owns" the data. It is business needs that determine the rules and standards for managing the data. Therefore, it is up to business users to inventory and review all the pieces of data in their systems to determine precisely which pieces of data from the old system will be used in the new system and which data do not need to be brought over. The data also need to be reviewed to make sure they are accurate and consistent and that redundant data are eliminated.

Most likely some type of data cleansing will be required. For example, American Water had data on more than 70,000 vendors in its vendor master data file. Andrew Clarkson, American Water's Business Intelligence Lead, asked business users to define an active vendor and to use that definition to identify which data to migrate. He also worked with various functional groups to standardize how to present address data.

One of the objectives of American Water's data management work was to support an enterprise-wide business intelligence program based on a single view of the business. An analytical system and data warehouse would be able to combine data from the SAP ERP System with data from other sources, including new customer information and enterprise asset management systems. That meant that American Water's business users had to do a lot of thinking about the kinds of reports they wanted. The company had originally planned to have the system provide 200 reports, but later reduced that number by half. Business users were trained to generate these reports and customize them. Most financial users initially tried to create their reports using Microsoft Excel spreadsheet software. Over time, however, they learned to do the same thing using SAP Business Objects Web Intelligence tools that came with the system. SAP Business Objects Web

Intelligence is a set of tools that enables business users to view, sort, and analyze business intelligence data. It includes tools for generating queries, reports and interactive dashboards.

At present, American Water is focusing on promoting the idea that data must be "clean" to be effective and has poured an incredible amount of effort into its data cleansing work—identifying incomplete, incorrect, inaccurate, and irrelevant pieces of data and then replacing, modifying, or deleting the "dirty" data. According to Clarkson, just as water treatment plants have measurements and meters to check water quality as its being treated, data management needs to ensure the quality of data at every step to make sure the final product will be genuinely useful for the company.

Sources: "SAP to Deliver Software Solution to American Water," www.sap.com, accessed January 31, 2014; David Hannon, "Clean Smooth-Flowing Data at American Water," SAP Insider Profiles, January–March 2013 and www.amwater.com, accessed February 2, 2014.

CASE STUDY QUESTIONS

1. Discuss the role of information policy, data administration, and efforts to ensure data quality in improving data management at American Water.
2. Describe roles played by information systems specialists and end users in American Water's systems transformation project.
3. Why was the participation of business users so important? If they didn't play this role, what would have happened?
4. How did implementing a data warehouse help American Water move toward a more centralized organization?
5. Give some examples of problems that would have occurred at American Water if its data were not "clean"?
6. How did American Water's data warehouse improve operations and management decision making?

Review Summary

1. *What are the problems of managing data resources in a traditional file environment?*

Traditional file management techniques make it difficult for organizations to keep track of all of the pieces of data they use in a systematic way and to organize these data so that they can be easily accessed. Different functional areas and groups were allowed to develop their own files independently. Over time, this traditional file management environment creates problems such as data redundancy and inconsistency, program-data dependence, inflexibility, poor security, and lack of data sharing and availability. A database management system (DBMS) solves these problems with software that permits centralization of data and data management so that businesses have a single consistent source for all their data needs. Using a DBMS minimizes redundant and inconsistent files.

2. *What are the major capabilities of DBMS and why is a relational DBMS so powerful?*

The principal capabilities of a DBMS includes a data definition capability, a data dictionary capability, and a data manipulation language. The data definition capability specifies the structure and content of the database. The data dictionary is an automated or manual file that stores information about the data in the database, including names, definitions, formats, and descriptions of data elements. The data manipulation language, such as SQL, is a specialized language for accessing and manipulating the data in the database.

The relational database has been the primary method for organizing and maintaining data in information systems because it is so flexible and accessible. It organizes data in two-dimensional tables called relations with rows and columns. Each table contains data about an entity and its attributes. Each row represents a record and each column represents an attribute or field. Each table also contains a key field to uniquely identify each record for retrieval or manipulation. Relational database tables can be combined easily to deliver data required by users, provided that any two tables share a common

data element. Non-relational databases are becoming popular for managing types of data that can't be handled easily by the relational data model. Both relational and non-relational database products are available as cloud computing services.

Designing a database requires both a logical design and a physical design. The logical design models the database from a business perspective. The organization's data model should reflect its key business processes and decision-making requirements. The process of creating small, stable, flexible, and adaptive data structures from complex groups of data when designing a relational database is termed normalization. A well-designed relational database will not have many-to-many relationships, and all attributes for a specific entity will only apply to that entity. It will try to enforce referential integrity rules to ensure that relationships between coupled tables remain consistent. An entity-relationship diagram graphically depicts the relationship between entities (tables) in a relational database.

3. *What are the principal tools and technologies for accessing information from databases to improve business performance and decision making?*

Contemporary data management technology has an array of tools for obtaining useful information from all the different types of data used by businesses today, including semi-structured and unstructured big data in vast quantities. These capabilities include data warehouses and data marts, Hadoop, in-memory computing, and analytical platforms. OLAP represents relationships among data as a multidimensional structure, which can be visualized as cubes of data and cubes within cubes of data, enabling more sophisticated data analysis. Data mining analyzes large pools of data, including the contents of data warehouses, to find patterns and rules that can be used to predict future behavior and guide decision making. Text mining tools help businesses analyze large unstructured data sets consisting of text. Web mining tools focus on analysis of useful patterns and information from the World Wide Web, examining the structure of Web sites and activities of Web site users as well as the contents of Web pages. Conventional databases can be linked via middleware to the Web or a Web interface to facilitate user access to an organization's internal data.

4. *Why are information policy, data administration, and data quality assurance essential for managing the firm's data resources?*

Developing a database environment requires policies and procedures for managing organizational data as well as a good data model and database technology. A formal information policy governs the maintenance, distribution, and use of information in the organization. In large corporations, a formal data administration function is responsible for information policy, as well as for data planning, data dictionary development, and monitoring data usage in the firm.

Data that are inaccurate, incomplete, or inconsistent create serious operational and financial problems for businesses because they may create inaccuracies in product pricing, customer accounts, and inventory data, and lead to inaccurate decisions about the actions that should be taken by the firm. Firms must take special steps to make sure they have a high level of data quality. These include using enterprise-wide data standards, databases designed to minimize inconsistent and redundant data, data quality audits, and data cleansing software.

Key Terms

Analytic platform, 266
Attribute, 250
Big data, 262
Bit, 250
Byte, 250
Data administration, 272
Data cleansing, 274
Data definition, 258
Data dictionary, 258
Data governance, 272
Data inconsistency, 251
Data manipulation language, 258
Data mart, 263
Data mining, 268
Data quality audit, 274
Data redundancy, 251
Data warehouse, 263
Database, 253
Database administration, 273
Database management system (DBMS), 253
Database server, 271
Entity, 250
Entity-relationship diagram, 261
Field, 250
File, 250
Foreign key, 256

Hadoop, 264
In-memory computing, 264
Information policy, 272
Key field, 255
Non-relational database management systems, 256
Normalization, 260
Online analytical processing (OLAP), 267
Primary key, 255
Program-data dependence, 252
Record, 250
Referential integrity, 261
Relational DBMS, 254
Sentiment analysis, 270
Structured Query Language (SQL), 258
Text mining, 270
Tuple, 255
Web mining, 270

MyMISLab

Go to **mymislab.com** to complete the problems marked with this icon.

Review Questions

6-1 What are the problems of managing data resources in a traditional file environment?

- List and describe each of the components in the data hierarchy.
- Define and explain the significance of entities, attributes, and key fields.
- List and describe the problems of the traditional file environment.

6-2 What are the major capabilities of database management systems (DBMS) and why is a relational DBMS so powerful?

- Define a database and a database management system.
- Name and briefly describe the capabilities of a DBMS.
- Define a relational DBMS and explain how it organizes data.
- List and describe the three operations of a relational DBMS.
- Explain why non-relational databases are useful.
- Define referential integrity and explain its role in relational database systems.
- Define and describe an entity-relationship diagram and explain its role in database design.

6-3 What are the principal tools and technologies for accessing information from databases to improve business performance and decision making?

- Define big data and describe the technologies for managing and analyzing it.
- List and describe the components of a contemporary business intelligence infrastructure.
- Describe the capabilities of online analytical processing (OLAP).
- Define data mining, describing how it differs from OLAP and the types of information it provides.
- Explain how text mining and Web mining differ from conventional data mining.
- Define data warehouse and explain how it differs from a data mart.

6-4 Why are information policy, data administration, and data quality assurance essential for managing the firm's data resources?

- Describe the roles of information policy and data administration in information management.
- Explain why an organization needs an information policy.
- Explain why data quality audits and data cleansing are essential.

Discussion Questions

6-5 It has been said there is no bad data, just bad management. Discuss the implications of this statement.

6-6 To what extent should end users be involved in the selection of a database management system and database design?

6-7 Why is big data so interesting to businesses? What challenges does big data present?

Hands-On MIS Projects

The projects in this section give you hands-on experience in analyzing data quality problems, establishing company-wide data standards, creating a database for inventory management, and using the Web to search online databases for overseas business resources.

Management Decision Problems

6-8 Iko Instruments Group, a global supplier of measurement, analytical, and monitoring instruments and services based in the Netherlands, had a new data warehouse designed to analyze customer activity to improve service and marketing. However, the data warehouse was full of inaccurate and redundant data. The data in the warehouse came from numerous transaction processing systems in the United States, Europe, Asia, and other locations around the world. The team that designed the warehouse had assumed that sales groups in all these areas would enter customer names, telephone numbers, and addresses the same way. In fact, companies in different countries were using multiple ways of entering quote, billing, shipping, contact information and other data. Assess the potential business impact of these data quality problems. What decisions have to be made and steps taken to reach a solution?

6-9 Your industrial supply company wants to create a data warehouse where management can obtain a single corporate-wide view of critical sales information to identify bestselling products, key customers, and sales trends. Your sales and product information are stored in several different systems: a divisional sales system running on a Unix server and a corporate sales system running on an IBM mainframe. You would like to create a single standard format that consolidates these data from both systems. In MyMISLab, you can review the proposed format, along with sample files from the two systems that would supply the data for the data warehouse. Then answer the following questions:

- What business problems are created by not having these data in a single standard format?
- How easy would it be to create a database with a single standard format that could store the data from both systems? Identify the problems that would have to be addressed.
- Should the problems be solved by database specialists or general business managers? Explain.
- Who should have the authority to finalize a single company-wide format for this information in the data warehouse?

Achieving Operational Excellence: Building a Relational Database for Inventory Management

Software skills: Database design, querying, and reporting
Business skills: Inventory management

6-10 In this exercise, you will use database software to design a database for managing inventory for a small business. Sylvester's Bike Shop, located in San Francisco, California, sells road, mountain, hybrid, leisure, and children's bicycles. Currently, Sylvester's purchases bikes from three suppliers, but plans to add new suppliers in the near future. Using the information found in the tables in MyMISLab, build a simple relational database to manage information about Sylvester's suppliers and products. Once you have built the database, perform the following activities.

- Prepare a report that identifies the five most expensive bicycles. The report should list the bicycles in descending order from most expensive to least expensive, the quantity on hand for each, and the markup percentage for each.
- Prepare a report that lists each supplier, its products, the quantities on hand, and associated reorder levels. The report should be sorted alphabetically by supplier. For each supplier, the products should be sorted alphabetically.
- Prepare a report listing only the bicycles that are low in stock and need to be reordered. The report should provide supplier information for the items identified.
- Write a brief description of how the database could be enhanced to further improve management of the business. What tables or fields should be added? What additional reports would be useful?

Improving Decision Making: Searching Online Databases for Overseas Business Resources

Software skills: Online databases
Business skills: Researching services for overseas operations

6-11 This project develops skills in searching Web-enabled databases with information about products and services in faraway locations.

Your company, Caledonian Furniture, is located in Cumbernauld, Scotland, and manufactures office furniture of various types. You are considering opening a facility to manufacture and sell your products in Australia. You would like to contact organizations that offer many services necessary for you to open your Australian office and manufacturing facility, including lawyers, accountants, import-export experts, and telecommunications equipment and support firms. Access the following online databases to locate companies that you would like to meet with during your upcoming trip: Australian Business Register, AustraliaTrade Now (australiatradenow.com), and the Nationwide Business Directory of Australia (www.nationwide.com.au). If necessary, use search engines such as Yahoo! and Google.

- List the companies you would contact on your trip to determine whether they can help you with these and any other functions you think are vital to establishing your office.
- Rate the databases you used for accuracy of name, completeness, ease of use, and general helpfulness.

Collaboration and Teamwork Project

6-12 In MyMISLab, you will find a Collaboration and Teamwork Project dealing with the concepts in this chapter. You will be able to use Google Drive, Google Docs, Google Sites, Google+, or other open source collaboration tools to complete the assignment.

Lego: Embracing Change by Combining BI with a Flexible Information System

CASE STUDY

The Lego Group, which is headquartered in Billund, Denmark, is one of the largest toy manufacturers in the world. Lego's main products have been the bricks and figures that children have played with for generations. The Danish company has experienced sustained growth since its founding in 1932, and for most of its history its major manufacturing facilities were located in Denmark.

In 2003, Lego was facing tough competition from imitators and manufacturers of electronic toys. In an effort to reduce costs, the group decided to initiate a gradual restructuring process that continues today. In 2006, the company announced that a large part of its production would be outsourced to the electronics manufacturing service company Flextronics, which has plants in Mexico, Hungary, and the Czech Republic. The decision to outsource production came as a direct consequence of an analysis of Lego's total supply chain. To reduce labor costs, manually intensive processes were outsourced, keeping only the highly skilled workers in Billund. Lego's workforce was gradually reduced from 8,300 employees in 2003 to approximately 4,200 in 2010. Additionally, production had to be relocated to places closer to its natural markets. As a consequence of all these changes, Lego transformed itself from a manufacturing firm to a market-oriented company that is capable of reacting fast to changing global demand.

Lego's restructuring process, coupled with double-digit sales growth in the past few years, has led to the company's expansion abroad and made its workforce more international. These changes presented supply chain and human resources challenges to the company. The supply chain had to be reengineered to simplify production without reducing quality. Improved logistics planning allowed Lego to work more closely with retailers, suppliers, and the new outsourcing companies. At the same time, the human resources (HR) department needed to play a more strategic role inside the company. HR was now responsible for implementing effective policies aimed at retaining and recruiting the most qualified employees from a diversity of cultural backgrounds.

Adapting company operations to these changes required a flexible and robust IT infrastructure with business intelligence capabilities that could help management perform better forecasting and planning. As part of the solution, Lego chose to move to SAP business suite software. SAP AG, a German company that specializes in enterprise software solutions, is one of the leading software companies in the world. SAP's software products include a variety of applications designed to efficiently support all of a company's essential functions and operations. Lego chose to implement SAP's Supply Chain Management (SCM), Product Lifecycle Management (PLM), and Enterprise Resources Planning (ERP) modules.

The SCM module includes essential features such as supply chain monitoring and analysis as well as forecasting, planning, and inventory optimization. The PLM module enables managers to optimize development processes and systems. The ERP module includes, among other applications, the Human Capital Management (HCM) application for personnel administration and development.

SAP's business suite is based on a flexible three-tier client-server architecture that can easily be adapted to the new service-oriented architecture (SOA) available in the latest versions of the software. In the first tier, a client interface—a browser-type graphical user interface (GUI) running on either a laptop, desktop, or mobile device—submits users' requests to the application servers. The applications servers (the second tier in the system) receive and process clients' requests. In turn, these application servers send the processed requests to the database system (the third tier) which consists of one or more relational databases. SAP's business suite supports databases from different vendors, including those offered by Oracle, Microsoft, MySQL, and others. The relational databases contain tables that store data on Lego's products, daily operations, the supply chain, and thousands of employees. Managers can easily use the SAP query tool to obtain reports from the databases because it does not require any technical skill. Additionally, the distributed architecture enables authorized personnel to have direct access to the database system from the company's various locations, including those in Europe, North America, and Asia.

SAP's ERP-HCM module includes advanced features such as "Talent Manager" as well as those for handling employee administration, reporting, and travel and time management. These features allow

Lego's HR personnel to select the best candidates, schedule their training, and create a stimulus plan to retain them. It is also possible to include performance measurements and get real-time insight into HR trends. Using these advanced features, together with tools from other software vendors, Lego's managers are able to track employees' leadership potential, develop their careers, and forecast the recruiting of new employees with certain skills.

The investments that The Lego Group has made in information systems and business re-design have paid off handsomely. In 2013 the Group increased sales by 11% to €3,403 million against €3,103 million the year before. Operating profit increased 10% to €1,118. Full-time employees increased to 11,755 as the company expanded production in Asia. In the first half of 2014, revenue increased 15 percent compared with the same period last year measured, and profits increased by 12% when compared to the similar period in 2013.

Reflecting its growing emphasis on developing a global company, and its substantial investment in global information systems both in the supply chain and the distribution chain, The Lego Group in 2014 is showing strong, long-term growth in all regions. In Europe, America, and Asia, sales growth has been in the double digits for over five years despite the fact that the Global Great Recession (2008 to 2013) led to flat sales of toys worldwide. In the Asian region, growth in Lego sales varied from market to market. China's growth in consumer sales of more than 50% was the most significant in the region. This supports The Lego Group's ambitions to further globalize the company and make Asia a significant contributor to future growth.

During April 2014 The Lego Group opened its first factory in China, located in Jiaxing, and a new office in Shanghai, which is one of five main offices globally for The Lego Group. The executives of Lego believe there is huge potential in Asia, and they are moving to be located close to their Asian customers. Lego executives believe that there is tremendous potential in Asia, and have decided to learn more about the Asian market and build capabilities in the region. The new factory and office represent a significant expansion of the Lego physical presence in the region. According to executives, in combination with their existing office in Singapore, the Shanghai office and the new factory enable strategically important functions to be located close to their customers as well as children and parents in China and Asia.

The decision to place a Lego factory in China is a direct consequence of The Lego Group's ambition to have production placed close to core markets. This same philosophy has led to expansions of the Lego factory in the Czech Republic, and an entirely new factory was opened in Nyiregyhaza, Hungary, in March 2014. These factories, along with the parent factory in Denmark, serve the European markets. To serve the Americas faster and with customized products, the company expanded its Lego factory in Monterrey, Mexico.

Executives believe the global approach to information systems and production facilities enables the company to deliver Lego products to retailers and ultimately to children all over the world very fast, offering a world class service to consumers. In 2014, in addition to its growth across a variety of markets, *The LEGO Movie* was also released to overwhelmingly positive reviews, bolstering the company's brand and allowing it to develop a new array of products based on the movie's themes.

The Lego Group is primed to continue its growth throughout 2015 and beyond using its organizational flexibility and the concepts it has honed for years. The company is responding to its customers and releasing new versions of some of its most popular sets of toys, including its Bionicle series of block sets. The company also launched a new office with 120 employees and an additional 80 more to come in 2015, complete with an innovative work environment without assigned desk spaces and emphasizing creativity over rigid routines. Like their expansion into Asia, the new office represents an effort to establish a global presence with offices and employees around the world to reach as many potential families and children as possible. So far, Lego has built an impressive worldwide presence, block by block.

Sources: Roar Trangbaek, "New London Office Supports Lego Group Strategy to Reach Children Globally, Newsroom, www.lego.com, November 2014; Roar Trangbaek, "Global Lego Sales Up 15 Percent in First Half of 2014," Newsroom, www.lego.com, 2014; Jens Hansegard, "Oh, Snap! Lego's Sales Surpass Mattel," *Wall Street Journal*, September 4, 2014, "How Lego Became World's Hottest Toy Company," *Economist*, March 9, 2013; "Business 2010: Embracing the Challenge of Change," *Economist Intelligence Unit*, February 2005, http://graphics. eiu.com/files/ad_pdfs/Business%202010_Global_ FINAL. pdf, accessed November 16, 2010; "Lego Creates Model Business Success with SAP and IBM," *IBM Global Financing*, May 19, 2010, www-01.ibm.com/software/

success/cssdb.nsf/CS/STRD- 85KGS6?OpenDocument, October 20, 2010; "Human Resources as an Exponent of Good Governance" (in Danish), www.sat.com, October 20, 2010; "Lego, the Toy of the Century Had to Reinvent the Supply-Chain to Save the Company," *Supply Chain Digest*, September 25, 2007, www.scdigest.com/assets/on_target/07-09-25- 7.php?cid = 1237, accessed November 16, 2010; G. W. Anderson, T. Rhodes, J. Davis, and J. Dobbins, *SAMS Teach Yourself SAP in 24 hours* (Indianapolis, IN: SAMS, 2008).

CASE STUDY QUESTIONS

6-13 Explain the role of the database in SAP's three-tier system.

6-14 Explain why distributed architectures are flexible.

6-15 Identify some of the business intelligence features included in SAP's business software suite.

6-16 What are the main advantages and disadvantages of having multiple databases in a distributed architecture? Explain.

Case contributed by Daniel Ortiz Arroyo, Aalborg University.

MyMISLab

Go to **mymislab.com** for the following Assisted-graded writing questions:

6-17 Identify the five problems of a traditional file environment and explain how a database management system solves them.

6-18 Discuss how the following facilitate the management of big data: Hadoop, in-memory computing, analytic platforms.

Chapter 6 References

Aiken, Peter, Mark Gillenson, Xihui Zhang, and David Rafner. "Data Management and Data Administration. Assessing 25 Years of Practice." *Journal of Database Management* (July-September 2011).

Barth, Paul S. "Managing Big Data: What Every CIO Needs to Know." *CIO Insight* (January 12, 2012).

Barton, Dominic and David Court. "Making Advanced Analytics Work for You." *Harvard Business Review* (October 2012).

Baum, David. "Flying High with a Private Database Cloud." *Oracle Magazine* (November/December 2011).

Beath, Cynthia, , Irma Becerra-Fernandez, Jeanne Ross and James Short. "Finding Value in the Information Explosion." *MIT Sloan Management Review* 53, No. 4 (Summer 2012).

Bughin, Jacques, John Livingston, and Sam Marwaha."Seizing the Potential for Big Data." The McKinsey Quarterly (October 2011).

Clifford, James, Albert Croker, and Alex Tuzhilin. "On Data Representation and Use in a Temporal Relational DBMS." *Information Systems Research* 7, No. 3 (September 1996).

Davenport, Thomas H. and D.J. Patil. "Data Scientist: The Sexiest Job of the 21st Century." *Harvard Business Review* (October 2012).

Davenport, Thomas H. Big Data at Work: *Dispelling the Myths, Uncovering the Opportunities.* Harvard Business Press (2014).

Eckerson, Wayne W. "Analytics in the Era of Big Data: Exploring a Vast New Ecosystem." TechTarget (2012).

_____. "Data Quality and the Bottom Line." The Data Warehousing Institute (2002).

Greengard, Samuel. "Big Data Unlocks Business Value." Baseline (January 2012).

Henschen, Doug. " MetLife Uses NoSQL for Customer Service Breakthrough." *Information Week* (May 13, 2013).

Hoffer, Jeffrey A., Ramesh Venkataraman, and Heikki Toppi. *Modern Database Management*, 11th ed. Upper Saddle River, NJ: Prentice-Hall (2013).

Jinesh Radadia. "Breaking the Bad Data Bottlenecks." *Information Management* (May/June 2010).

Jordan, John. "The Risks of Big Data for Companies." *Wall Street Journal* (October 20, 2013).

Kajepeeta, Sreedhar. "How Hadoop Tames Enterprises' Big Data." *Information Week* (February 2012).

Kroenke, David M. and David Auer. *Database Processing: Fundamentals, Design, and Implementation*, 13e. Upper Saddle River, NJ: Prentice-Hall (2014).

Lee, Yang W., and Diane M. Strong. "Knowing-Why about Data Processes and Data Quality." *Journal of Management Information Systems* 20, No. 3 (Winter 2004).

Lohr, Steve. "The Age of Big Data." *New York Times* (February 11, 2012).

Loveman, Gary. "Diamonds in the Datamine." *Harvard Business Review* (May 2003)

Marcus, Gary and Ernest Davis. "Eight (No, Nine!) Problems With Big Data." *New York Times* (April 6, 2014).

Martens, David and Foster Provost. "Explaining Data-Driven Document Classifications." *MIS Quarterly* 38, No. 1 (March 2014).

McAfee, Andrew and Erik Brynjolfsson. "Big Data: The Management Revolution." *Harvard Business Review* (October 2012).

McKinsey Global Institute. "Big Data: The Next Frontier for Innovation, Competition, and Productivity." McKinsey & Company (2011).

Morrison, Todd and Mark Fontecchio, "In-memory Technology Pushes Analytics Boundaries, Boosts BI Speeds," SearchBusinessAnalytics.techtarget.com, accessed May 17, 2013.

Morrow, Rich. "Apache Hadoop: The Swiss Army Knife of IT." Global Knowledge (2013).

Mulani, Narendra. "In-Memory Technology: Keeping Pace with Your Data." *Information Management* (February 27, 2013).

Redman, Thomas. *Data Driven: Profiting from Your Most Important Business Asset*. Boston: Harvard Business Press (2008).

Redman, Thomas C. "Data's Credibility Problem" *Harvard Business Review* (December 2013).

Rosenbush, Steven and Michael Totty. "How Big Data Is Transforming Business." *Wall Street Journal* (March 10, 2013).

Ross, Jeanne W., Cynthia M. Beath, and Anne Quaadgras. "You May Not Need Big Data After All." *Harvard Business Review* (December 2013).

Wallace, David J. "How Caesar's Entertainment Sustains a Data-Driven Culture." *DataInformed* (December 14, 2012).

Telecommunications, the Internet, and Wireless Technology

CHAPTER 7

LEARNING OBJECTIVES

After reading this chapter, you will be able to answer the following questions:

1. What are the principal components of telecommunications networks and key networking technologies?
2. What are the different types of networks?
3. How do the Internet and Internet technology work and how do they support communication and e-business?
4. What are the principal technologies and standards for wireless networking, communication, and Internet access?

MYMISLAB™

Visit **mymislab.com** for simulations, tutorials, and end-of-chapter problems.

CHAPTER CASES

RFID and Wireless Technology Speed Up Production at Continental Tires

The Battle over Net Neutrality

Monitoring Employees on Networks: Unethical or Good Business?

RFID Propels the Angkasa Library Management System

VIDEO CASES

Telepresence Moves Out of the Boardroom and Into the Field

Virtual Collaboration with Lotus Sametime

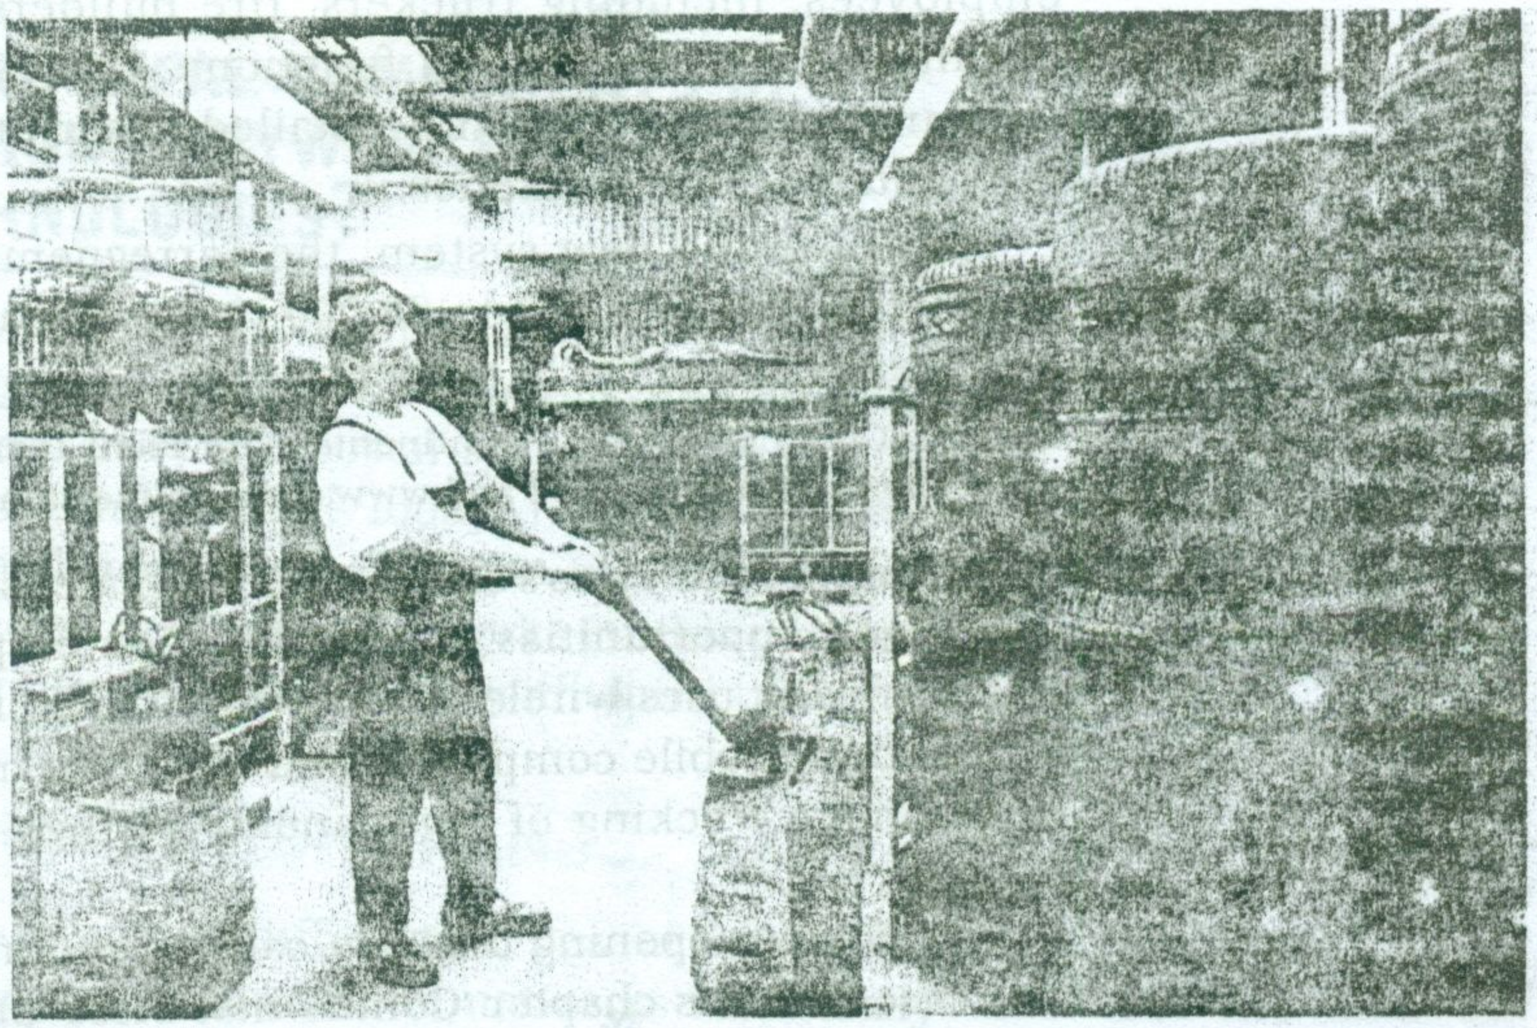
RFID AND WIRELESS TECHNOLOGY SPEED UP PRODUCTION AT CONTINENTAL TIRES

Continental AG, headquartered in Hanover, Germany, is a global auto and truck parts manufacturing company, with 164,000 employees in 46 countries. It is also the world's fourth largest tire manufacturer and one of the top five automotive suppliers in the world.

One of the factories for Continental's Tire Division is located in Sarreguemines, France. This facility produces 1,000 different kinds of tires and encompasses nearly 1.5 million square feet. The production process requires large wheeled carts loaded with sheets of rubber or other components to be transported from storage to workstations as tires are being built. Until recently, if a carrier was not in its expected location, a worker had to look for it manually. Manual tracking was time-consuming and inaccurate, and the plant often lost track of tire components altogether.

Missing materials created bottlenecks and production delays at a time when business was growing and the company needed to increase production capacity. Continental found a solution in a new real-time location system based on a Wi-Fi wireless network using radio frequency identification (RFID) tags, AeroScout MobileView software, mobile computers, and Global Data Sciences' material inventory tracking system software.

The Sarreguemines plant mounted AeroScout T2-EB Industrial RFID tags on the sides of 1,100 of its carriers. As the carriers move from one manufacturing or storage station to another, location information about the cart is transmitted to nearby nodes of a Cisco Wi-Fi wireless network. AeroScout's MobileView software picks up the location and represents the carrier as an icon on a map of the facility displayed on computer screens. Fifteen Honeywell Dolphin 6500 and Motorola Solutions MC9190 handheld



© Caro / Alamy

computers are used to confirm that a carrier has been loaded with components or has arrived at a specific workstation.

Seven of the plant's tuggers, which are small trucks for hauling the carriers around the plant, are equipped with DLOG mobile vehicle-mounted computers. When a tugger driver is looking for a specific component, he or she can use the mobile device to access the MobileView system, pull up a map of the facility, and see an icon indicating where that component's carrier is located. The location tracking system provides a real-time snapshot of all the components used in the factory.

A bar code label is attached to each component and carrier, and the system starts tracking that component as soon as it is placed in a carrier. Plant workers use one of the Motorola or Honeywell handhelds and the MobileView software to scan the bar code labels on both the component and its carrier, which is associated with the ID number transmitted by an RFID tag mounted on the carrier. The scanned bar code data are stored in a material inventory tracking system. The MobileView software tracks the carrier's location as it is being transported to a storage area, and also the location where it is placed in storage.

When components are needed for manufacturing, a tugger driver uses the DLOG mobile computer to identify the location of the carrier with those specific components, and then goes to that location. After the carrier has been retrieved and taken to a workstation, its bar code is scanned by an employee at that station using one of the handheld computers. This updates the system to show that the required components have been received.

By enabling tugger drivers to quickly locate components, the new system has increased productivity and ensures that materials are not overlooked or misplaced. Fewer materials are thrown away because they expired and were not used when they were needed. The system is able to send alerts of materials that have been sitting too long in one spot.

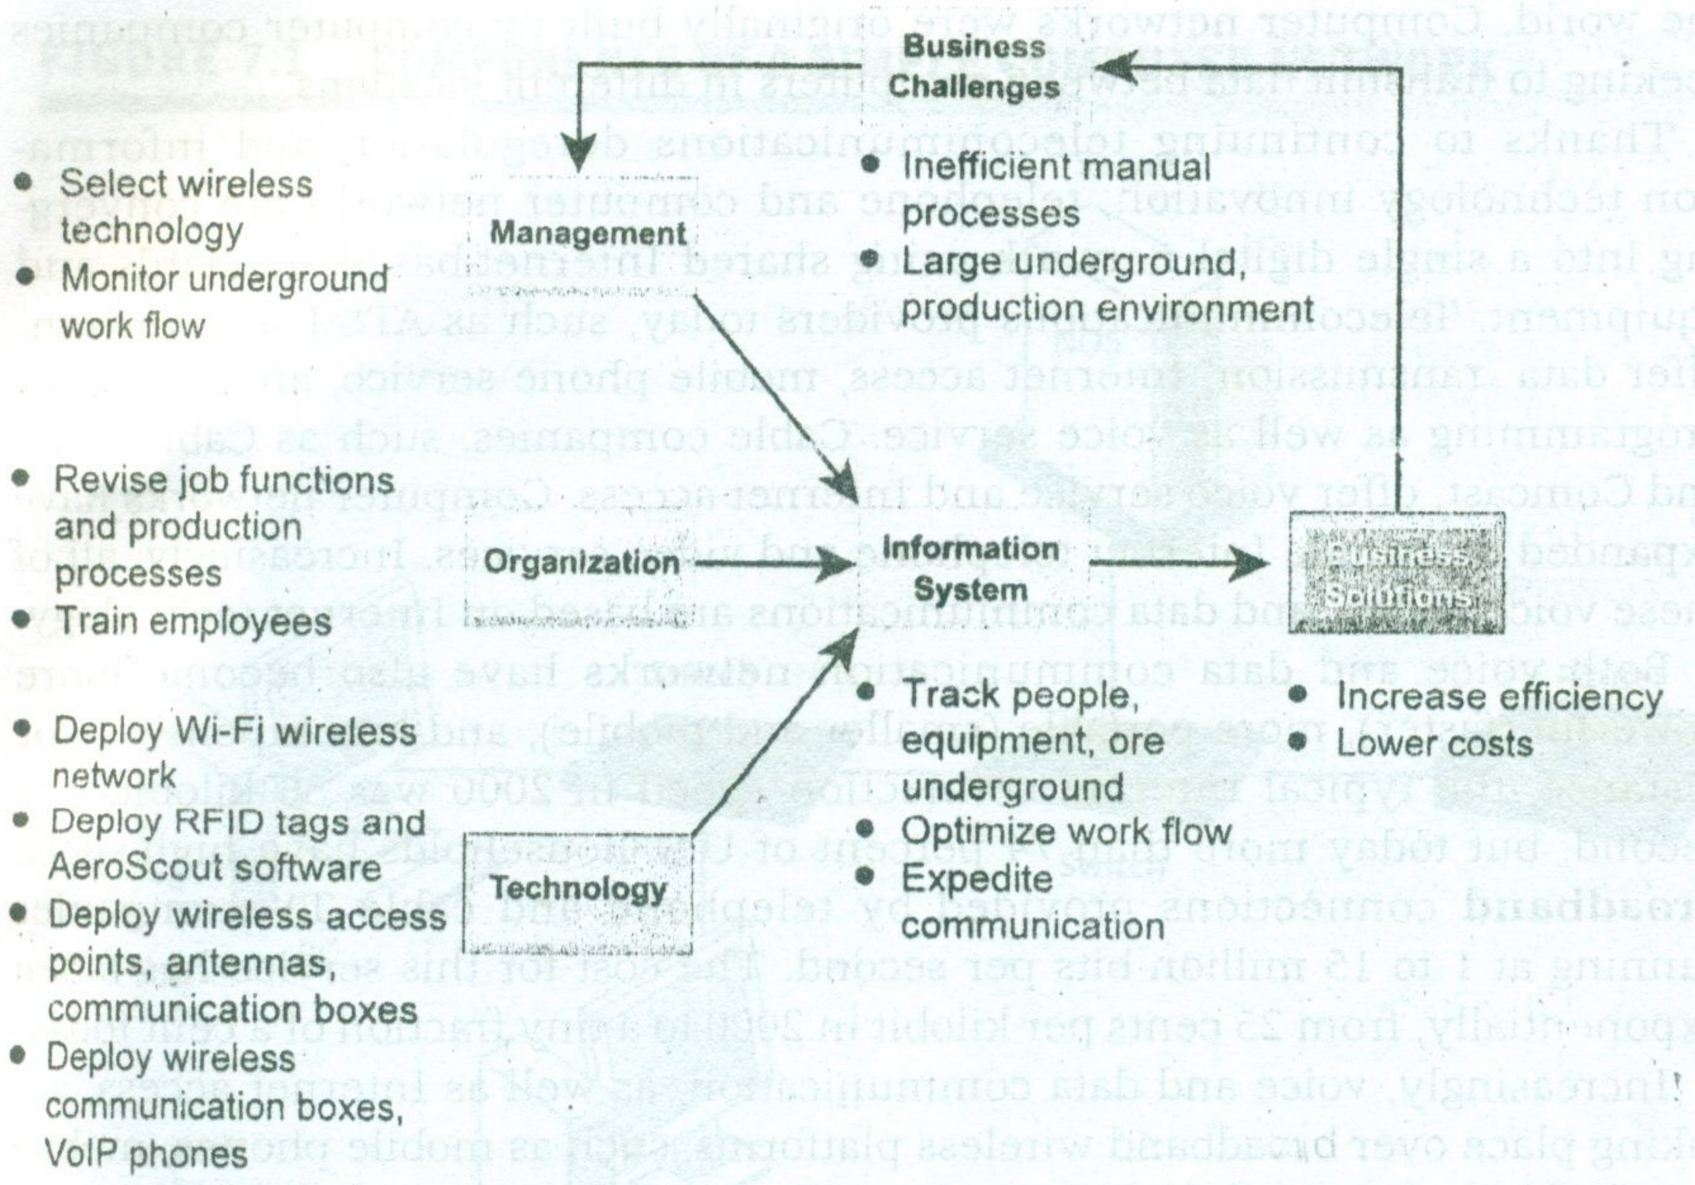
When AeroScout and the new material inventory tracking system were implemented in September 2011, Continental made sure all production employees, including truckers, tire builders, and management, received training in the new system functions. The company also provided workers with instruction cards with detailed descriptions of system functions that they could use for reference.

Thanks to the new system, the Sarreguemines tire factory has increased production from 33,000 to 38,000 tires per day. Wastage of tire components has been reduced by 20 percent.

Sources: Claire Swedberg, "Continental Tire Plant Increases Productivity, Reduces Waste," *RFID Journal*, April 25, 2012 and www.conti-online.com, accessed May 2, 2012.

Continental Tires's experience illustrates some of the powerful capabilities and opportunities provided by contemporary networking technology. The company uses wireless networking, radio frequency identification (RFID) technology, mobile computers, and materials inventory management software to automate tracking of components as they move through the production process.

The chapter-opening diagram calls attention to important points raised by this case and this chapter. Continental Tires' production environment extends over a very large area, and requires intensive oversight and coordination to make sure that components are available when and where they are needed in the production process. Tracking components manually was very slow and cumbersome, increasing the possibility that components would be overlooked or lost.



Management decided that wireless technology and RFID tagging provided a solution and arranged for the deployment of a wireless RFID network throughout the entire Sarreguemines production facility. The network made it much easier to track components and to optimize tugger truck movements. Continental Tires had to redesign its production and other work processes and train employees in the new system to take advantage of the new technology.

Here are some questions to think about: How did Continental's real-time location system transform operations? Why was training so important?

7.1 WHAT ARE THE PRINCIPAL COMPONENTS OF TELECOMMUNICATIONS NETWORKS AND KEY NETWORKING TECHNOLOGIES?

If you run or work in a business, you can't do without networks. You need to communicate rapidly with your customers, suppliers, and employees. Until about 1990, businesses used the postal system or telephone system with voice or fax for communication. Today, however, you and your employees use computers, e-mail, text messaging, the Internet, mobile phones, and mobile computers connected to wireless networks for this purpose. Networking and the Internet are now nearly synonymous with doing business.

NETWORKING AND COMMUNICATION TRENDS

Firms in the past used two fundamentally different types of networks: telephone networks and computer networks. Telephone networks historically handled voice communication, and computer networks handled data traffic. Telephone networks were built by telephone companies throughout the twentieth century using voice transmission technologies (hardware and software), and these companies almost always operated as regulated monopolies throughout

the world. Computer networks were originally built by computer companies seeking to transmit data between computers in different locations.

Thanks to continuing telecommunications deregulation and information technology innovation, telephone and computer networks are converging into a single digital network using shared Internet-based standards and equipment. Telecommunications providers today, such as AT&T and Verizon, offer data transmission, Internet access, mobile phone service, and television programming as well as voice service. Cable companies, such as Cablevision and Comcast, offer voice service and Internet access. Computer networks have expanded to include Internet telephone and video services. Increasingly, all of these voice, video, and data communications are based on Internet technology.

Both voice and data communication networks have also become more powerful (faster), more portable (smaller and mobile), and less expensive. For instance, the typical Internet connection speed in 2000 was 56 kilobits per second, but today more than 74 percent of U.S. households have high-speed **broadband** connections provided by telephone and cable TV companies running at 1 to 15 million bits per second. The cost for this service has fallen exponentially, from 25 cents per kilobit in 2000 to a tiny fraction of a cent today.

Increasingly, voice and data communication, as well as Internet access, are taking place over broadband wireless platforms, such as mobile phones, mobile handheld devices, and PCs in wireless networks. More than half the Internet users in the United States use smartphones and tablets to access the Internet.

WHAT IS A COMPUTER NETWORK?

If you had to connect the computers for two or more employees together in the same office, you would need a computer network. Exactly what is a network? In its simplest form, a network consists of two or more connected computers. Figure 7.1 illustrates the major hardware, software, and transmission components used in a simple network: a client computer and a dedicated server computer, network interfaces, a connection medium, network operating system software, and either a hub or a switch.

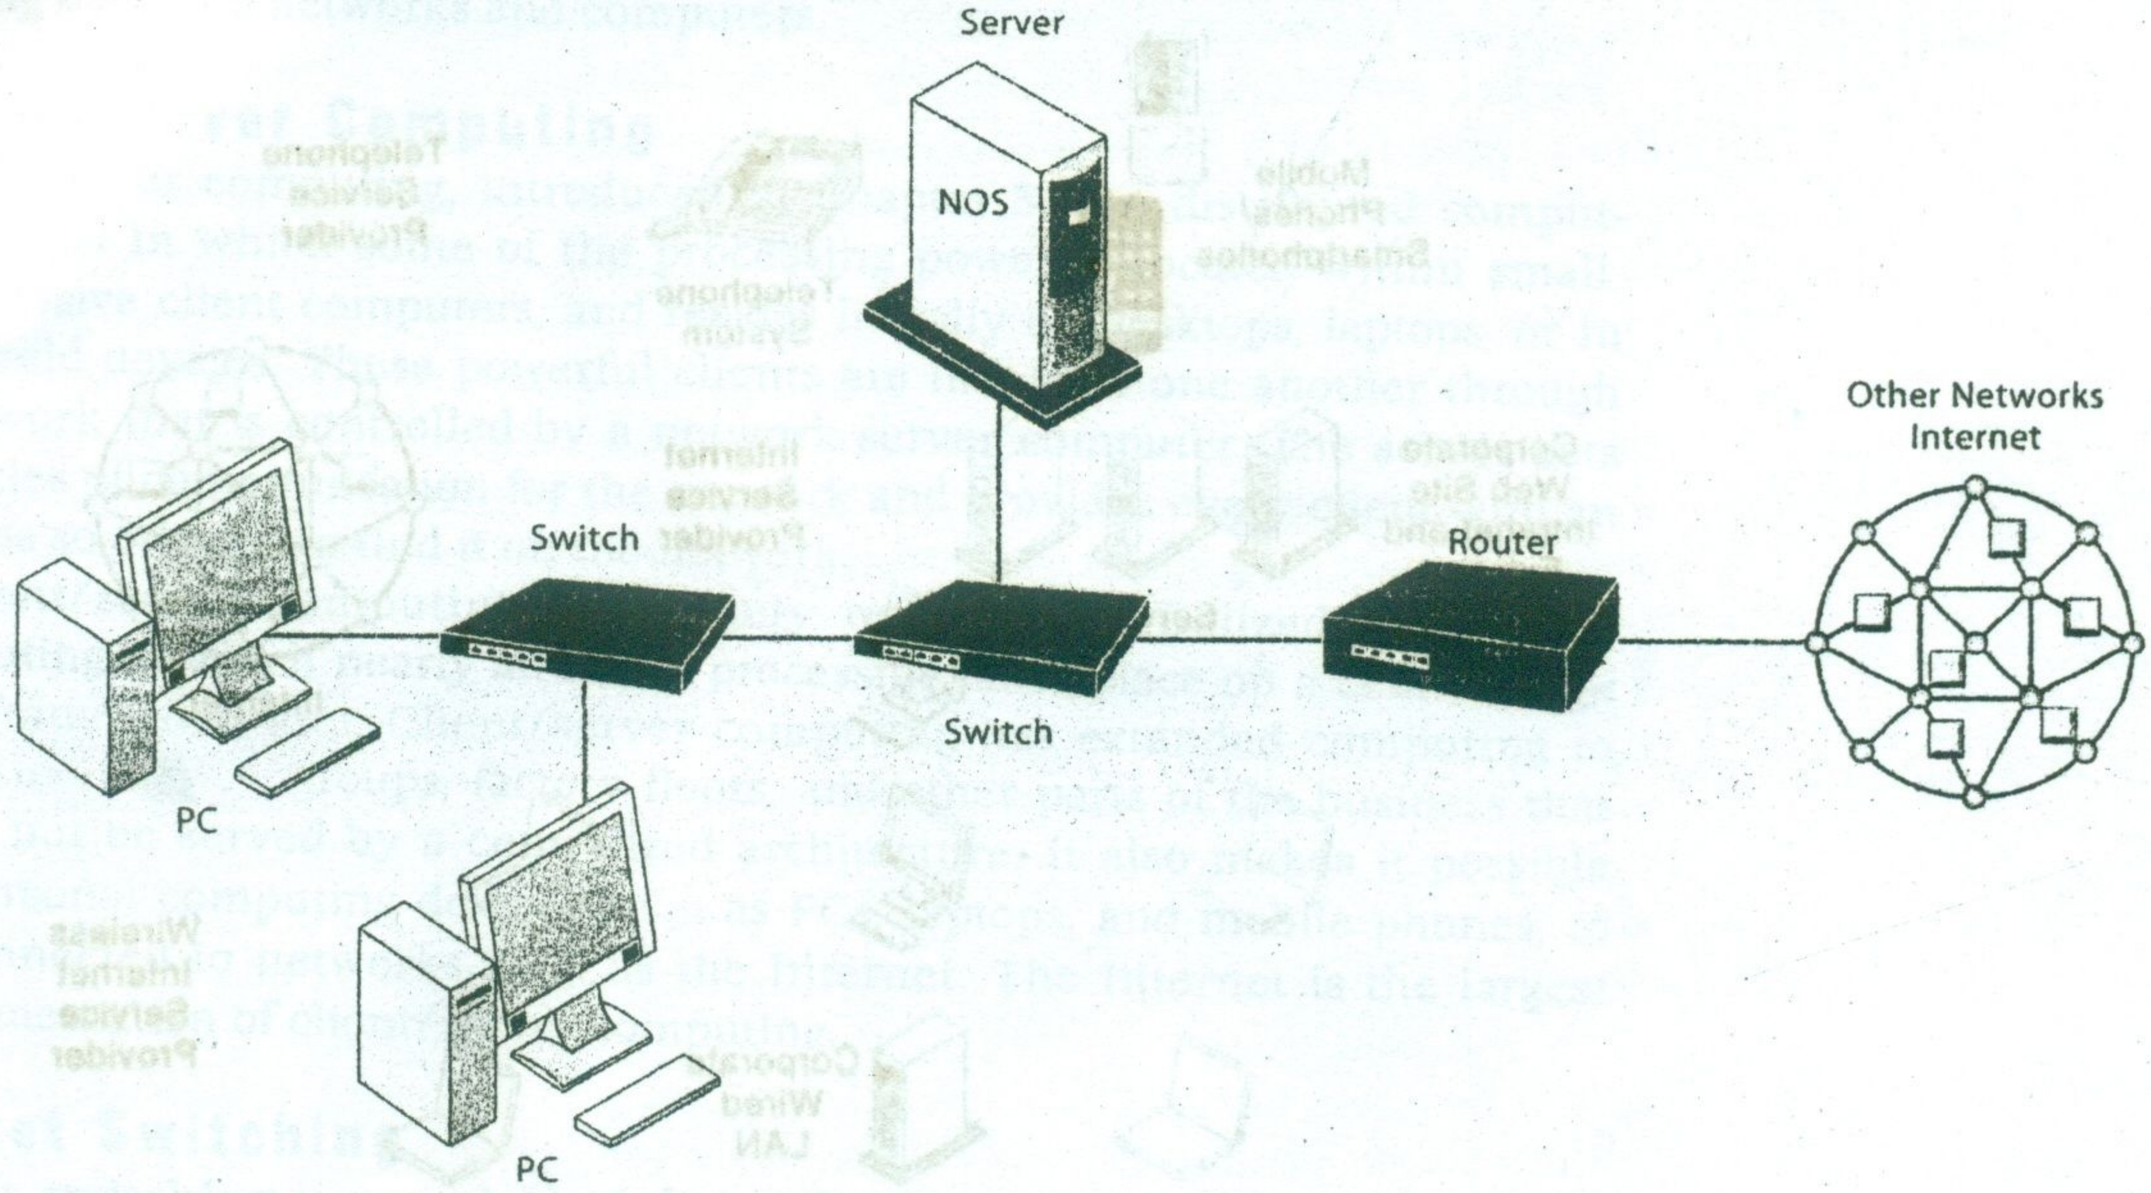
Each computer on the network contains a network interface device to link the computer to the network. The connection medium for linking network components can be a telephone wire, coaxial cable, or radio signal in the case of cell phone and wireless local area networks (Wi-Fi networks).

The network operating system (NOS) routes and manages communications on the network and coordinates network resources. It can reside on every computer in the network, or it can reside primarily on a dedicated server computer for all the applications on the network. A server computer is a computer on a network that performs important network functions for client computers, such as serving up Web pages, storing data, and storing the network operating system (and hence controlling the network). Microsoft Windows Server, Linux, and Novell Open Enterprise Server are the most widely used network operating systems.

Most networks also contain a switch or a hub acting as a connection point between the computers. **Hubs** are very simple devices that connect network components, sending a packet of data to all other connected devices. A **switch** has more intelligence than a hub and can filter and forward data to a specified destination on the network.

What if you want to communicate with another network, such as the Internet? You would need a router. A **router** is a communications processor used to route packets of data through different networks, ensuring that the data sent gets to the correct address.

FIGURE 7.1 COMPONENTS OF A SIMPLE COMPUTER NETWORK



Illustrated here is a very simple computer network, consisting of computers, a network operating system (NOS) residing on a dedicated server computer, cable (wiring) connecting the devices, switches, and a router.

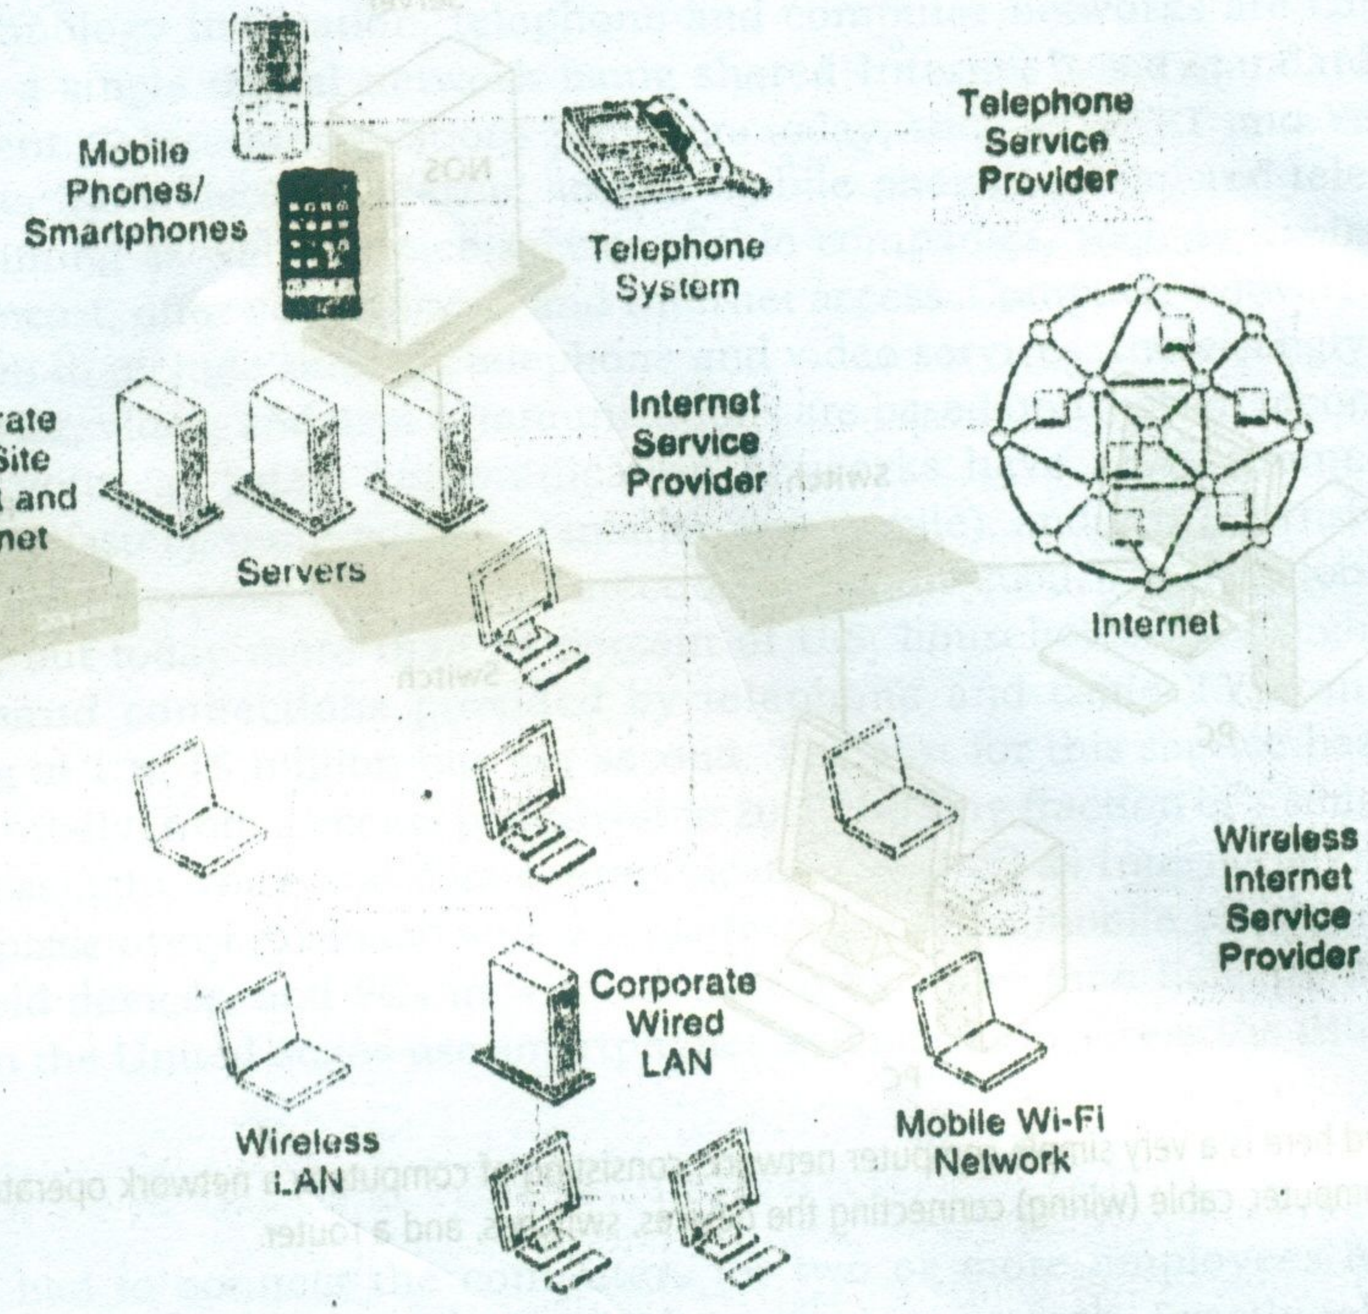
Network switches and routers have proprietary software built into their hardware for directing the movement of data on the network. This can create network bottlenecks and makes the process of configuring a network more complicated and time-consuming. **Software-defined networking (SDN)** is a new networking approach in which many of these control functions are managed by one central program, which can run on inexpensive commodity servers that are separate from the network devices themselves. This is especially helpful in a cloud computing environment with many different pieces of hardware because it allows a network administrator to manage traffic loads in a flexible and more efficient manner.

Networks in Large Companies

The network we've just described might be suitable for a small business. But what about large companies with many different locations and thousands of employees? As a firm grows, and collects hundreds of small local area networks, these networks can be tied together into a corporate-wide networking infrastructure. The network infrastructure for a large corporation consists of a large number of these small local area networks linked to other local area networks and to firmwide corporate networks. A number of powerful servers support a corporate Web site, a corporate intranet, and perhaps an extranet. Some of these servers link to other large computers supporting back-end systems.

Figure 7.2 provides an illustration of these more complex, larger scale corporate-wide networks. Here you can see that the corporate network infrastructure supports a mobile sales force using mobile phones and smartphones, mobile employees linking to the company Web site, internal

FIGURE 7.2 CORPORATE NETWORK INFRASTRUCTURE



Today's corporate network infrastructure is a collection of many different networks from the public switched telephone network, to the Internet, to corporate local area networks linking workgroups, departments, or office floors.

company networks using mobile wireless local area networks (Wi-Fi networks), and a videoconferencing system to support managers across the world. In addition to these computer networks, the firm's infrastructure may include a separate telephone network that handles most voice data. Many firms are dispensing with their traditional telephone networks and using Internet telephones that run on their existing data networks (described later).

As you can see from this figure, a large corporate network infrastructure uses a wide variety of technologies—everything from ordinary telephone service and corporate data networks to Internet service, wireless Internet, and mobile phones. One of the major problems facing corporations today is how to integrate all the different communication networks and channels into a coherent system that enables information to flow from one part of the corporation to another, and from one system to another. As more and more communication networks become digital, and based on Internet technologies, it will become easier to integrate them.

KEY DIGITAL NETWORKING TECHNOLOGIES

Contemporary digital networks and the Internet are based on three key technologies: client/server computing, the use of packet switching, and the development of widely used communications standards (the most important

of which is Transmission Control Protocol/Internet Protocol, or TCP/IP) for linking disparate networks and computers.

Client/Server Computing

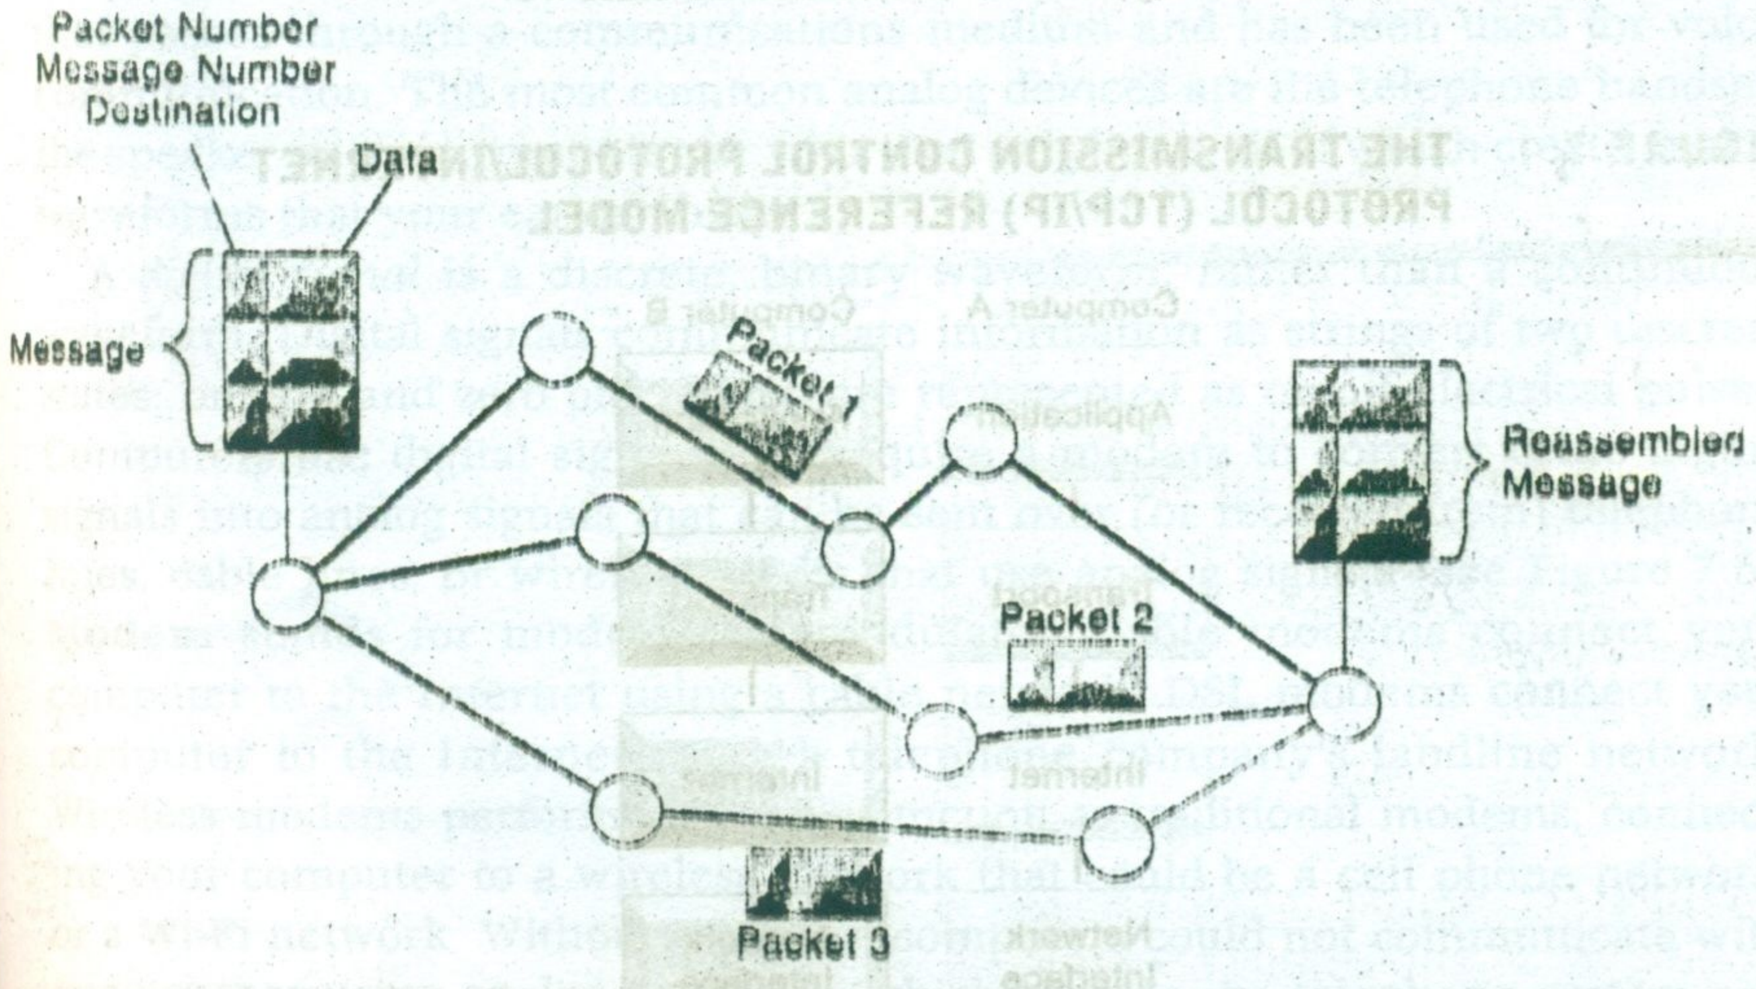
Client/server computing, introduced in Chapter 5, is a distributed computing model in which some of the processing power is located within small, inexpensive client computers, and resides literally on desktops, laptops, or in handheld devices. These powerful clients are linked to one another through a network that is controlled by a network server computer. The server sets the rules of communication for the network and provides every client with an address so others can find it on the network.

Client/server computing has largely replaced centralized mainframe computing in which nearly all of the processing takes place on a central large mainframe computer. Client/server computing has extended computing to departments, workgroups, factory floors, and other parts of the business that could not be served by a centralized architecture. It also makes it possible for personal computing devices such as PCs, laptops, and mobile phones, to be connected to networks such as the Internet. The Internet is the largest implementation of client/server computing.

Packet Switching

Packet switching is a method of slicing digital messages into parcels called packets, sending the packets along different communication paths as they become available, and then reassembling the packets once they arrive at their destinations (see Figure 7.3). Prior to the development of packet switching, computer networks used leased, dedicated telephone circuits to communicate with other computers in remote locations. In circuit-switched networks, such as the telephone system, a complete point-to-point circuit is assembled, and then communication can proceed. These dedicated circuit-switching techniques

FIGURE 7.3 **PACKET-SWITCHED NETWORKS AND PACKET COMMUNICATIONS**



Data are grouped into small packets, which are transmitted independently over various communications channels and reassembled at their final destination.

were expensive and wasted available communications capacity—the circuit was maintained regardless of whether any data were being sent.

Packet switching makes much more efficient use of the communications capacity of a network. In packet-switched networks, messages are first broken down into small fixed bundles of data called packets. The packets include information for directing the packet to the right address and for checking transmission errors along with the data. The packets are transmitted over various communications channels using routers, each packet traveling independently. Packets of data originating at one source will be routed through many different paths and networks before being reassembled into the original message when they reach their destinations.

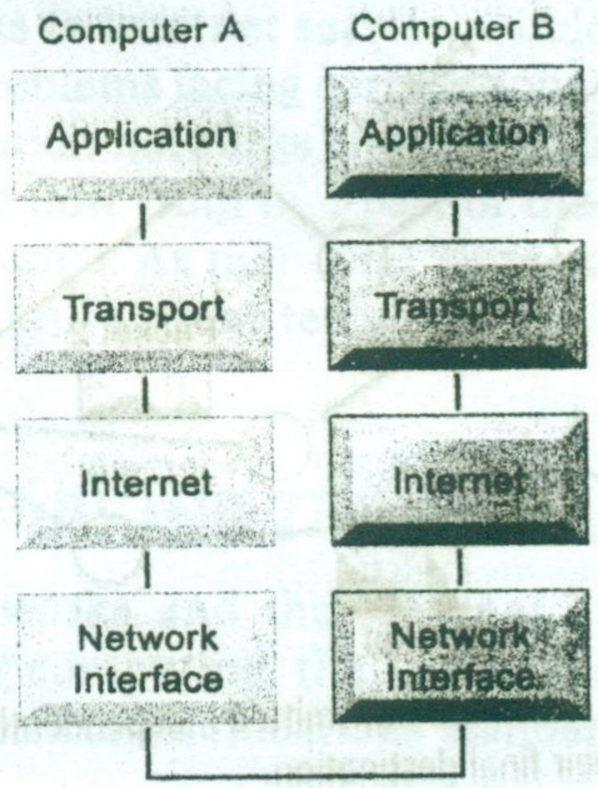
TCP/IP and Connectivity

In a typical telecommunications network, diverse hardware and software components need to work together to transmit information. Different components in a network communicate with each other only by adhering to a common set of rules called protocols. A **protocol** is a set of rules and procedures governing transmission of information between two points in a network.

In the past, many diverse proprietary and incompatible protocols often forced business firms to purchase computing and communications equipment from a single vendor. But today, corporate networks are increasingly using a single, common, worldwide standard called **Transmission Control Protocol/Internet Protocol (TCP/IP). TCP/IP** was developed during the early 1970s to support U.S. Department of Defense Advanced Research Projects Agency (DARPA) efforts to help scientists transmit data among different types of computers over long distances.

TCP/IP uses a suite of protocols, the main ones being TCP and IP. TCP refers to the Transmission Control Protocol, which handles the movement of data between computers. TCP establishes a connection between the computers, sequences the transfer of packets, and acknowledges the packets sent. IP refers to the Internet Protocol (IP), which is responsible for the delivery of packets and includes the disassembling and reassembling of packets during transmission. Figure 7.4 illustrates the four-layered Department of Defense reference model for TCP/IP, and the layers are described as follows:

FIGURE 7.4 THE TRANSMISSION CONTROL PROTOCOL/INTERNET PROTOCOL (TCP/IP) REFERENCE MODEL



This figure illustrates the four layers of the TCP/IP reference model for communications.

1. *Application layer.* The Application layer enables client application programs to access the other layers and defines the protocols that applications use to exchange data. One of these application protocols is the Hypertext Transfer Protocol (HTTP), which is used to transfer Web page files.
2. *Transport layer.* The Transport layer is responsible for providing the Application layer with communication and packet services. This layer includes TCP and other protocols.
3. *Internet layer.* The Internet layer is responsible for addressing, routing, and packaging data packets called IP datagrams. The Internet Protocol is one of the protocols used in this layer.
4. *Network Interface layer.* At the bottom of the reference model, the Network Interface layer is responsible for placing packets on and receiving them from the network medium, which could be any networking technology.

Two computers using TCP/IP are able to communicate even if they are based on different hardware and software platforms. Data sent from one computer to the other passes downward through all four layers, starting with the sending computer's Application layer and passing through the Network Interface layer. After the data reach the recipient host computer, they travel up the layers and are reassembled into a format the receiving computer can use. If the receiving computer finds a damaged packet, it asks the sending computer to retransmit it. This process is reversed when the receiving computer responds.

7.2 WHAT ARE THE DIFFERENT TYPES OF NETWORKS?

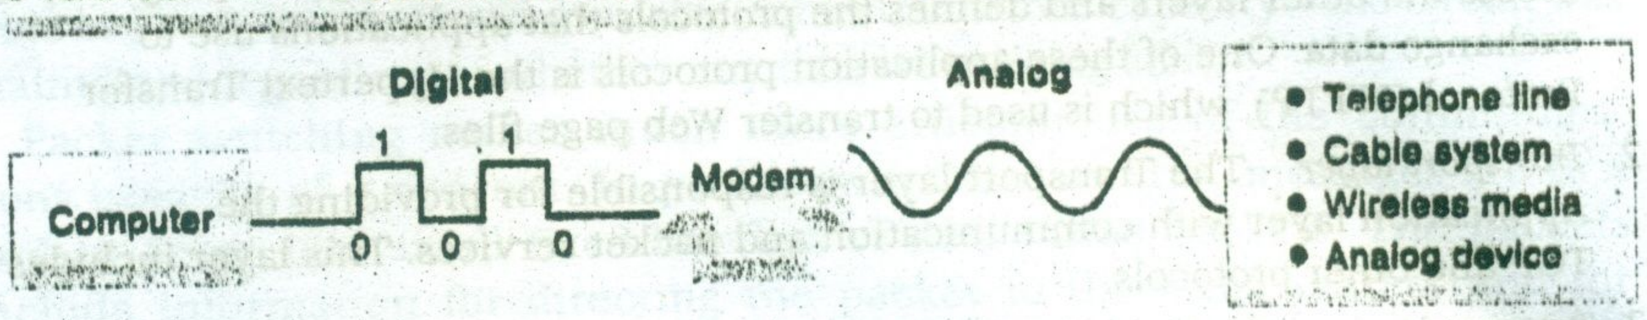
Let's look more closely at alternative networking technologies available to businesses.

SIGNALS: DIGITAL VS. ANALOG

There are two ways to communicate a message in a network: an analog signal or a digital signal. An *analog signal* is represented by a continuous waveform that passes through a communications medium and has been used for voice communication. The most common analog devices are the telephone handset, the speaker on your computer, or your iPod earphone, all of which create analog waveforms that your ear can hear.

A *digital signal* is a discrete, binary waveform, rather than a continuous waveform. Digital signals communicate information as strings of two discrete states: one bit and zero bits, which are represented as on-off electrical pulses. Computers use digital signals and require a modem to convert these digital signals into analog signals that can be sent over (or received from) telephone lines, cable lines, or wireless media that use analog signals (see Figure 7.5). **Modem** stands for modulator-demodulator. Cable modems connect your computer to the Internet using a cable network. DSL modems connect your computer to the Internet using a telephone company's landline network. Wireless modems perform the same function as traditional modems, connecting your computer to a wireless network that could be a cell phone network, or a Wi-Fi network. Without modems, computers could not communicate with one another using analog networks (which include the telephone system and cable networks).

FIGURE 7.5 FUNCTIONS OF THE MODEM



A modem is a device that translates digital signals into analog form (and vice versa) so that computers can transmit data over analog networks such as telephone and cable networks.

There are many different kinds of networks and ways of classifying them. One way of looking at networks is in terms of their geographic scope (see Table 7.1).

Local Area Networks

If you work in a business that uses networking, you are probably connecting to other employees and groups via a local area network. A **local area network (LAN)** is designed to connect personal computers and other digital devices within a half-mile or 500-meter radius. LANs typically connect a few computers in a small office, all the computers in one building, or all the computers in several buildings in close proximity. LANs also are used to link to long-distance wide area networks (WANs, described later in this section) and other networks around the world using the Internet.

Review Figure 7.1, which could serve as a model for a small LAN that might be used in an office. One computer is a dedicated network file server, providing users with access to shared computing resources in the network, including software programs and data files.

The server determines who gets access to what and in which sequence. The router connects the LAN to other networks, which could be the Internet or another corporate network, so that the LAN can exchange information with networks external to it. The most common LAN operating systems are Windows, Linux, and Novell. Each of these network operating systems supports TCP/IP as its default networking protocol.

Ethernet is the dominant LAN standard at the physical network level, specifying the physical medium to carry signals between computers, access control rules, and a standardized set of bits used to carry data over the system.

TABLE 7.1 TYPES OF NETWORKS

TYPE	AREA
Local area network (LAN)	Up to 500 meters (half a mile); an office or floor of a building
Campus area network (CAN)	Up to 1,000 meters (a mile); a college campus or corporate facility
Metropolitan area network (MAN)	A city or metropolitan area
Wide area network (WAN)	A transcontinental or global area

Originally, Ethernet supported a data transfer rate of 10 megabits per second (Mbps). Newer versions, such as Gigabit Ethernet, support a data transfer rate of 1 gigabit per second (Gbps).

The LAN illustrated in Figure 7.1 uses a client/server architecture where the network operating system resides primarily on a single file server, and the server provides much of the control and resources for the network. Alternatively, LANs may use a peer-to-peer architecture. A peer-to-peer network treats all processors equally and is used primarily in small networks with 10 or fewer users. The various computers on the network can exchange data by direct access and can share peripheral devices without going through a separate server.

In LANs using the Windows Server family of operating systems, the **peer-to-peer** architecture is called the *workgroup network model*, in which a small group of computers can share resources, such as files, folders, and printers, over the network without a dedicated server. The *Windows domain network model*, in contrast, uses a dedicated server to manage the computers in the network.

Larger LANs have many clients and multiple servers, with separate servers for specific services, such as storing and managing files and databases (file servers or database servers), managing printers (print servers), storing and managing e-mail (mail servers), or storing and managing Web pages (Web servers).

Metropolitan and Wide Area Networks

Wide area networks (WANs) span broad geographical distances—entire regions, states, continents, or the entire globe. The most universal and powerful WAN is the Internet. Computers connect to a WAN through public networks, such as the telephone system or private cable systems, or through leased lines or satellites. A **metropolitan area network (MAN)** is a network that spans a metropolitan area, usually a city and its major suburbs. Its geographic scope falls between a WAN and a LAN.

TRANSMISSION MEDIA AND TRANSMISSION SPEED

Networks use different kinds of physical transmission media, including twisted pair wire, coaxial cable, fiber-optic cable, and media for wireless transmission. Each has advantages and limitations. A wide range of speeds is possible for any given medium depending on the software and hardware configuration. Table 7.2 compares these media.

Bandwidth: Transmission Speed

The total amount of digital information that can be transmitted through any telecommunications medium is measured in bits per second (bps). One signal change, or cycle, is required to transmit one or several bits; therefore, the transmission capacity of each type of telecommunications medium is a function of its frequency. The number of cycles per second that can be sent through that medium is measured in **hertz**—one hertz is equal to one cycle of the medium.

The range of frequencies that can be accommodated on a particular telecommunications channel is called its **bandwidth**. The bandwidth is the difference between the highest and lowest frequencies that can be accommodated on a single channel. The greater the range of frequencies, the greater the bandwidth and the greater the channel's transmission capacity.

TABLE 7.2 PHYSICAL TRANSMISSION MEDIA

TRANSMISSION MEDIUM	DESCRIPTION	SPEED
Twisted pair wire (CAT 5)	Strands of copper wire twisted in pairs for voice and data communications. CAT 5 is the most common 10 Mbps LAN cable. Maximum recommended run of 100 meters.	10–100+ Mbps
Coaxial cable	Thickly insulated copper wire, which is capable of high-speed data transmission and less subject to interference than twisted wire. Currently used for cable TV and for networks with longer runs (more than 100 meters).	Up to 1 Gbps
Fiber-optic cable	Strands of clear glass fiber, transmitting data as pulses of light generated by lasers. Useful for high-speed transmission of large quantities of data. More expensive than other physical transmission media and harder to install; often used for network backbone.	15 Mbps to 6+ Tbps
Wireless transmission media	Based on radio signals of various frequencies and includes both terrestrial and satellite microwave systems and cellular networks. Used for long-distance, wireless communication and Internet access.	Up to 600+ Mbps

7.3 How do the Internet and Internet technology work and how do they support communication and e-business?

We all use the Internet, and many of us can't do without it. It's become an indispensable personal and business tool. But what exactly is the Internet? How does it work, and what does Internet technology have to offer for business? Let's look at the most important Internet features.

WHAT IS THE INTERNET?

The Internet has become the world's most extensive, public communication system. It's also the world's largest implementation of client/server computing and internetworking, linking millions of individual networks all over the world. This global network of networks began in the early 1970s as a U.S. Department of Defense network to link scientists and university professors around the world.

Most homes and small businesses connect to the Internet by subscribing to an Internet service provider. **An Internet service provider (ISP)** is a commercial organization with a permanent connection to the Internet that sells temporary connections to retail subscribers. EarthLink, NetZero, AT&T, and Time Warner are ISPs. Individuals also connect to the Internet through their business firms, universities, or research centers that have designated Internet domains.

There are a variety of services for ISP Internet connections. Connecting via a traditional telephone line and modem, at a speed of 56.6 kilobits per second (Kbps), used to be the most common form of connection worldwide, but it has been largely replaced by broadband connections. Digital subscriber line, cable, satellite Internet connections, and T lines provide these broadband services.

Digital subscriber line (DSL) technologies operate over existing telephone lines to carry voice, data, and video at transmission rates ranging from 385 Kbps

all the way up to 40 Mbps, depending on usage patterns and distance. **Cable Internet connections** provided by cable television vendors use digital cable coaxial lines to deliver high-speed Internet access to homes and businesses. They can provide high-speed access to the Internet of up to 50 Mbps, although most providers offer service ranging from 1 Mbps to 6 Mbps. In areas where DSL and cable services are unavailable, it is possible to access the Internet via satellite, although some satellite Internet connections have slower upload speeds than other broadband services.

T1 and T3 are international telephone standards for digital communication. They are leased, dedicated lines suitable for businesses or government agencies requiring high-speed guaranteed service levels. **T1 lines** offer guaranteed delivery at 1.54 Mbps, and T3 lines offer delivery at 45 Mbps. The Internet does not provide similar guaranteed service levels, but simply "best effort."

INTERNET ADDRESSING AND ARCHITECTURE

The Internet is based on the TCP/IP networking protocol suite described earlier in this chapter. Every computer on the Internet is assigned a unique **Internet Protocol (IP) address**, which currently is a 32-bit number represented by four strings of numbers ranging from 0 to 255 separated by periods. For instance, the IP address of www.microsoft.com is 207.46.250.119.

When a user sends a message to another user on the Internet, the message is first decomposed into packets using the TCP protocol. Each packet contains its destination address. The packets are then sent from the client to the network server and from there on to as many other servers as necessary to arrive at a specific computer with a known address. At the destination address, the packets are reassembled into the original message.

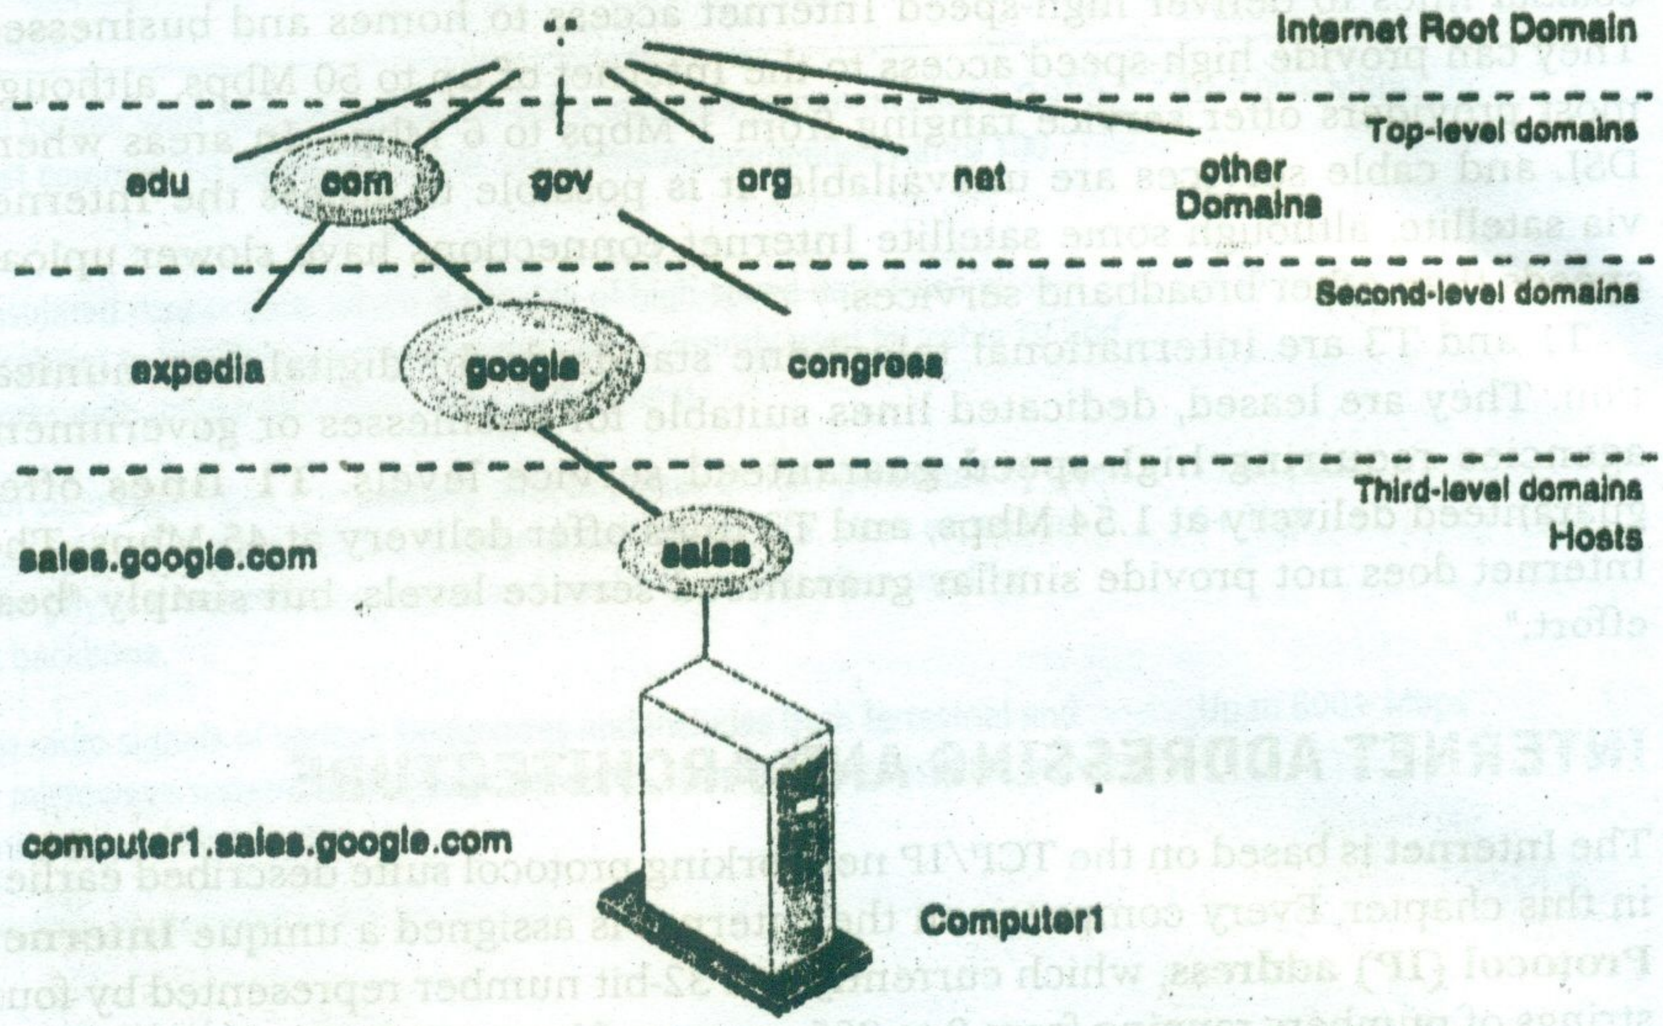
The Domain Name System

Because it would be incredibly difficult for Internet users to remember strings of 12 numbers, the **Domain Name System (DNS)** converts domain names to IP addresses. The **domain name** is the English-like name that corresponds to the unique 32-bit numeric IP address for each computer connected to the Internet. DNS servers maintain a database containing IP addresses mapped to their corresponding domain names. To access a computer on the Internet, users need only specify its domain name.

DNS has a hierarchical structure (see Figure 7.6). At the top of the DNS hierarchy is the root domain. The child domain of the root is called a top-level domain, and the child domain of a top-level domain is called a second-level domain. Top-level domains are two- and three-character names you are familiar with from surfing the Web, for example, .com, .edu, .gov, and the various country codes such as .ca for Canada or .it for Italy. Second-level domains have two parts, designating a top-level name and a second-level name—such as buy.com, nyu.edu, or amazon.ca. A host name at the bottom of the hierarchy designates a specific computer on either the Internet or a private network.

The most common domain extensions currently available and officially approved are shown in the following list. Countries also have domain names such as .uk, .au, and .fr (United Kingdom, Australia, and France, respectively), and there is a new class of "internationalized" top-level domains that use non-English characters. In the future, this list will expand to include many more types of organizations and industries.

FIGURE 7.6 THE DOMAIN NAME SYSTEM



Domain Name System is a hierarchical system with a root domain, top-level domains, second-level domains, and host computers at the third level.

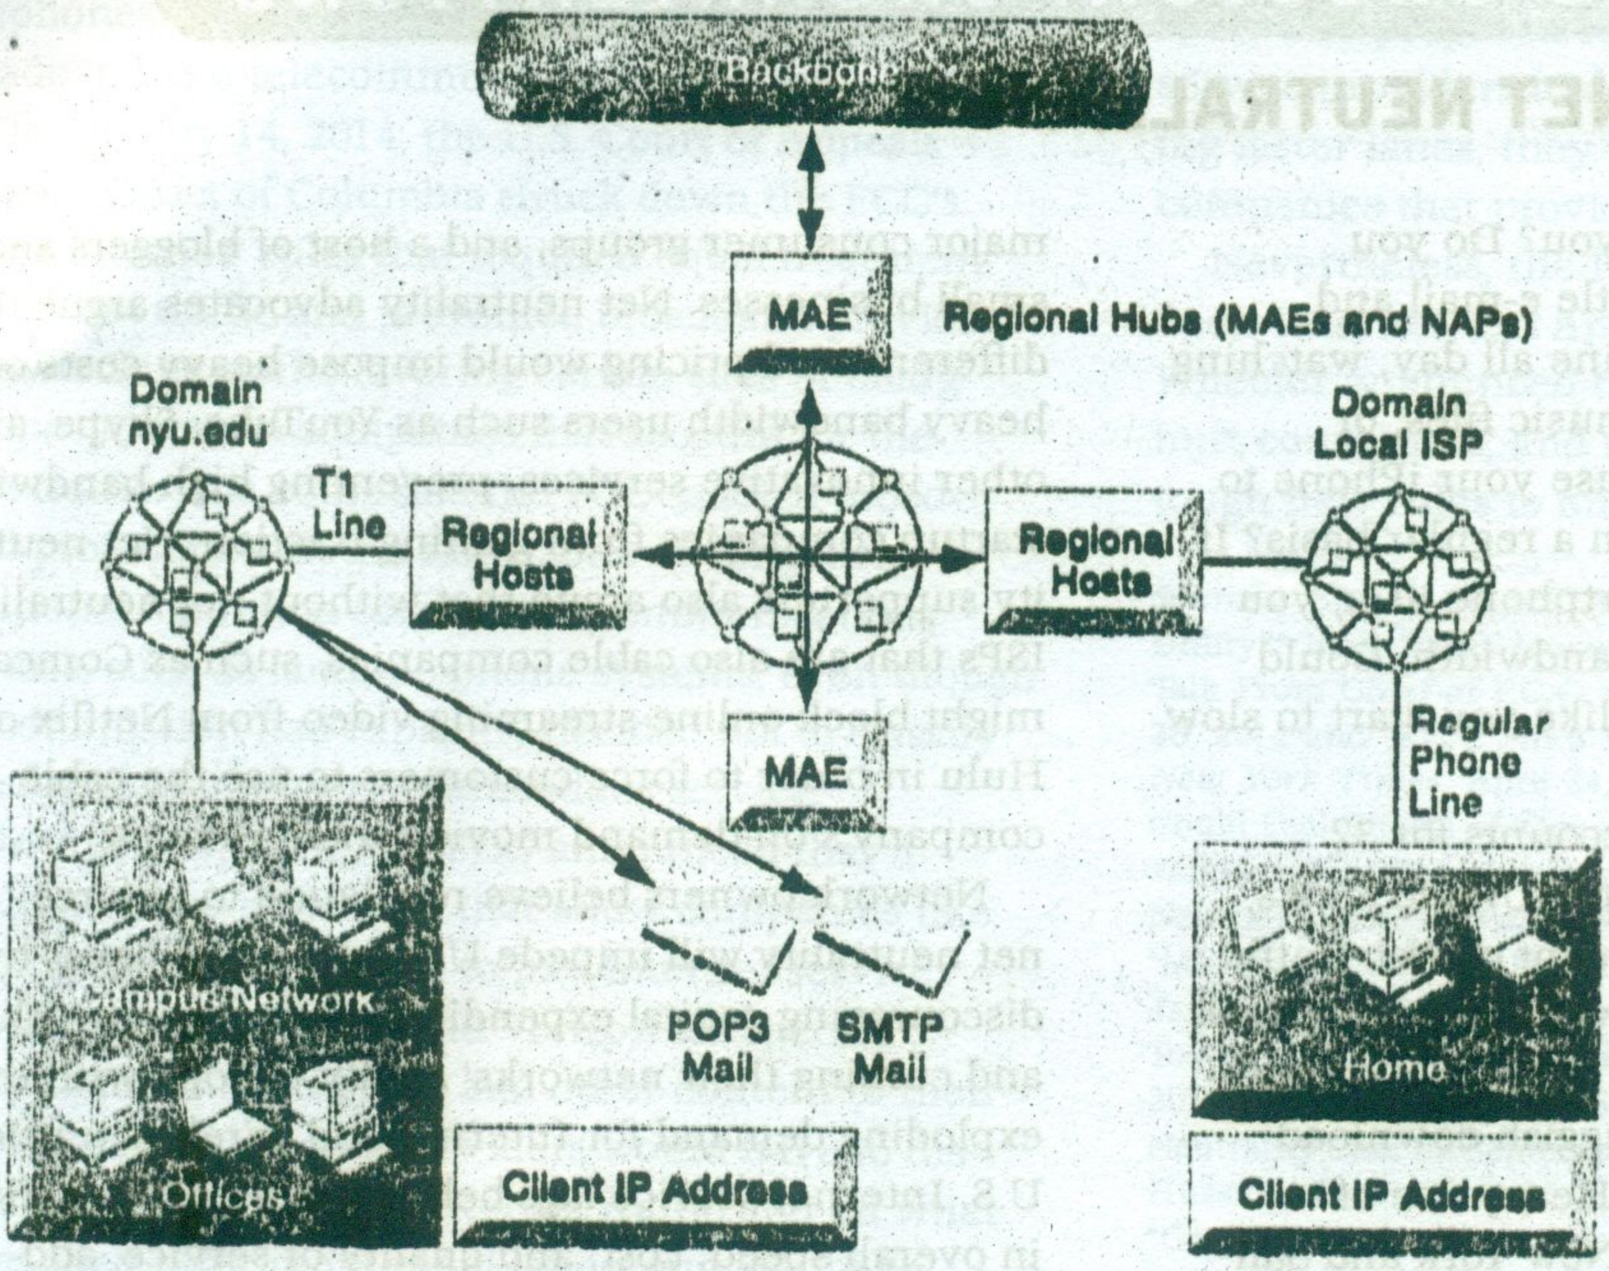
.com	Commercial organizations/businesses
.edu	Educational institutions
.gov	U.S. government agencies
.mil	U.S. military
.net	Network computers
.org	Nonprofit organizations and foundations
.biz	Business firms
.info	Information providers

Internet Architecture and Governance

Internet data traffic is carried over transcontinental high-speed backbone networks that generally operate in the range of 45 Mbps to 2.5 Gbps (see Figure 7.7). These trunk lines are typically owned by long-distance telephone companies (called *network service providers*) or by national governments. Local connection lines are owned by regional telephone and cable television companies in the United States that connect retail users in homes and businesses to the Internet. The regional networks lease access to ISPs, private companies, and government institutions.

Each organization pays for its own networks and its own local Internet connection services, a part of which is paid to the long-distance trunk line owners. Individual Internet users pay ISPs for using their service, and they generally pay a flat subscription fee, no matter how much or how little they use the Internet. A debate is now raging on whether this arrangement should continue or whether heavy Internet users who download large video and music files should pay more for the bandwidth they consume. The Interactive Session on Organizations explores this topic, by examining the pros and cons of net neutrality.

FIGURE 7.7 INTERNET NETWORK ARCHITECTURE



The Internet backbone connects to regional networks, which in turn provide access to Internet service providers, large firms, and government institutions. Network access points (NAPs) and metropolitan area exchanges (MAEs) are hubs where the backbone intersects regional and local networks and where backbone owners connect with one another.

No one "owns" the Internet, and it has no formal management. However, worldwide Internet policies are established by a number of professional organizations and government bodies, including the Internet Architecture Board (IAB), which helps define the overall structure of the Internet; the Internet Corporation for Assigned Names and Numbers (ICANN), which assigns IP addresses; and the World Wide Web Consortium (W3C), which sets Hypertext Markup Language and other programming standards for the Web.

These organizations influence government agencies, network owners, ISPs, and software developers with the goal of keeping the Internet operating as efficiently as possible. The Internet must also conform to the laws of the sovereign nation-states in which it operates, as well as the technical infrastructures that exist within the nation-states. Although in the early years of the Internet and the Web there was very little legislative or executive interference, this situation is changing as the Internet plays a growing role in the distribution of information and knowledge, including content that some find objectionable.

ICANN was created by the Department of Commerce in 1986 to manage the domain name system, and the central core servers of the Internet domains located in the United States. Emerging nations and others have long called for the U.S. dominance of the Internet to end, and instead support a multi-national controlling agency. After the terrorist attacks on the United States on September 11, 2001, the U.S. refused to consider such a change. This opinion changed again in March 2014 when the Commerce Department announced its intention to transfer the domain name functions of the Internet

INTERACTIVE SESSION: ORGANIZATIONS

THE BATTLE OVER NET NEUTRALITY

What kind of Internet user are you? Do you primarily use the Net to do a little e-mail and online banking? Or are you online all day, watching YouTube videos, downloading music files, or playing online games? Do you use your iPhone to stream TV shows and movies on a regular basis? If you're a power Internet or smartphone user, you are consuming a great deal of bandwidth. Could hundreds of millions of people like you start to slow the Internet down?

Video streaming on Netflix accounts for 32 percent of all bandwidth use in the United States, and Google's YouTube for 19 percent of Web traffic at peak hours. If user demand overwhelms network capacity, the Internet might not come to a screeching halt, but users could face sluggish download speeds and video transmission. Heavy use of iPhones in urban areas such as New York and San Francisco has already degraded service on the AT&T wireless network. AT&T reported that 3 percent of its subscriber base accounted for 40 percent of its data traffic.

Internet service providers (ISPs) assert that network congestion is a serious problem and that expanding their networks would require passing on burdensome costs to consumers. These companies believe differential pricing methods, which include data caps and metered use—charging based on the amount of bandwidth consumed—are the fairest way to finance necessary investments in their network infrastructures. But metering Internet use is not widely accepted, because of an ongoing debate about net neutrality.

Net neutrality is the idea that Internet service providers must allow customers equal access to content and applications, regardless of the source or nature of the content. Presently, the Internet is neutral: all Internet traffic is treated equally on a first-come, first-served basis by Internet backbone owners. However, this arrangement prevents telecommunications and cable companies from charging differentiated prices based on the amount of bandwidth consumed by the content being delivered over the Internet.

The strange alliance of net neutrality advocates includes MoveOn.org; the Christian Coalition; the American Library Association; data-intensive Web businesses such as Netflix, Amazon, and Google; major consumer groups; and a host of bloggers and small businesses. Net neutrality advocates argue that differentiated pricing would impose heavy costs on heavy bandwidth users such as YouTube, Skype, and other innovative services, preventing high-bandwidth startup companies from gaining traction. Net neutrality supporters also argue that without net neutrality, ISPs that are also cable companies, such as Comcast, might block online streaming video from Netflix or Hulu in order to force customers to use the cable company's on-demand movie rental services.

Network owners believe regulation to enforce net neutrality will impede U.S. competitiveness by discouraging capital expenditure for new networks and curbing their networks' ability to cope with the exploding demand for Internet and wireless traffic. U.S. Internet service lags behind many other nations in overall speed, cost, and quality of service, adding credibility to this argument. And with enough options for Internet access, dissatisfied consumers could simply switch to providers who enforce net neutrality and allow unlimited Internet use.

The wireless industry had been largely exempted from net neutrality rules, because the government determined it was a less mature network and companies should be allowed more freedom to manage traffic. Wireless providers already have tiered plans that charge heavy bandwidth users larger service fees.

A December 2012 report by the non-profit, non-partisan, public policy institute, New America Foundation (NAF), disputes these claims. Like personal computers, the processing capacity of the routers and switches in wired broadband networks has vastly expanded while the price has declined. Although total U.S. Internet data consumption rose 120% in 2012, the cost to transport the data decreased at a faster pace. The net cost to carriers was at worst flat and for the most part, down. The NAF report further asserts that lack of competition has enabled wired broadband carriers to charge higher rates, institute data caps, and spend less on the capital expenditures needed to upgrade and maintain their networks than they have in the past.

The courts have maintained that the Federal Communications Commission (FCC) has no authority to dictate how the Internet operates. The Communications Act of 1996 forbids the agency from managing the Internet as a "common carrier," the

regulatory approach the commission took toward telephones, and the FCC itself decided not to classify broadband as a telecommunications service.

On January 14, 2014, the U.S. Court of Appeals for the District of Columbia struck down the FCC's "Open Internet" rules that required equal treatment of Internet traffic and prevented broadband providers from blocking traffic favoring certain sites or charging special fees to companies that account for the most traffic. The court said the FCC saddled broadband providers with the same sorts of obligations as traditional "common carrier" telecommunications services, such as landline phone systems, even though the commission had explicitly decided not to classify broadband as a telecommunications service.

On April 24, 2014, the FCC announced that it would propose new rules that allow companies like Disney, Google or Netflix to pay Internet service providers like Comcast and Verizon for special, faster lanes to send video and other content to their customers. Broadband providers would have to disclose how they treat all Internet traffic and on what terms they offer more rapid lanes, and would be required to act in a "commercially reasonable manner." Providers would not be allowed to block Web sites. The proposed rules would also require Internet service providers to disclose whether, in assigning faster lanes, they had favored their affiliated companies that provide content.

Nevertheless, the FCC continues to push for an open Internet. On April 30, 2014, FCC chairman Tom Wheeler announced that lack of competition has hurt consumers, and that the FCC planned to write tough new rules to enforce net neutrality.

Sources: "Should the U.S. Regulate Broadband Internet Access as a Utility?" *Wall Street Journal*, May 11, 2014; Edward Wyatt, "Stern Talk From Chief of F.C.C. on Open Net," *New York Times*, April 30, 2014 and "F.C.C., in a Shift, Backs Fast Lane for Web Traffic," *New York Times*, April 24, 2014; Amol Sharma, "Netflix, YouTube Could Feel Effects of 'Open Internet' Ruling," *Wall Street Journal*, January 14, 2014; Gautham Nagesh, "FCC to Propose New 'Net Neutrality' Rules," *Wall Street Journal*, April 23, 2014; Shira Ovide, "Moving Beyond the Net Neutrality Debate," *Wall Street Journal*, January 14, 2014; Gautham Nagesh and Amol Sharma, "Court Tosses Rules of Road for Internet," *Wall Street Journal*, January 4, 2014; UpdAlina Selyukh," S. Court to Hear Oral Arguments in Net Neutrality Case on September 9," Reuters, June 25, 2013; and Hibah Hussain, Danielle Kehl, Benjamin Lennett, and Patrick Lucey, "Capping the Nation's Broadband Future? Dwindling Competition Is Fueling the Rise of Increasingly Costly and Restrictive Internet Usage Caps," New America Foundation, December 17, 2012.

CASE STUDY QUESTIONS

1. What is net neutrality? Why has the Internet operated under net neutrality up to this point in time?
2. Who's in favor of net neutrality? Who's opposed? Why?
3. What would be the impact on individual users, businesses, and government if Internet providers switched to a tiered service model for transmission over land lines as well as wireless?
4. It has been said that net neutrality is the most important issue facing the Internet since the advent of the Internet. Discuss the implications of this statement.
5. Are you in favor of legislation enforcing network neutrality? Why or why not?

over to a global, multi-stakeholder community which will be determined in 2015. Multi-stakeholder means that the leadership of the global Internet would follow the pattern of ICANN and would be composed of representatives from academia, business, governments, and public interest groups rather than a government led or an inter-governmental body. Until this body is formed, the Internet domain name system will remain under the control of the Department of Commerce. The announcement came in part as a response to widespread global hostility to U.S. control over the DNS amidst the revelations of Edward Snowden, describing how U.S. intelligence agencies used the Internet to conduct surveillance over individuals and groups around the world even though such surveillance had nothing to do with the operations of ICANN or the Department of Commerce, but instead were enabled by other

technical means (NTIA, 2014; Wyatt, 2014). The hope is that the transition to a multinational body will not disrupt the orderly operation of the Internet.

The Future Internet: IPv6 and Internet2

The Internet was not originally designed to handle the transmission of massive quantities of data and billions of users. Because of sheer Internet population growth, the world is about to run out of available IP addresses using the old addressing convention. The old addressing system is being replaced by a new version of the IP addressing schema called **IPv6** (Internet Protocol version 6), which contains 128-bit addresses (2 to the power of 128), or more than a quadrillion possible unique addresses. IPv6 is not compatible with the existing Internet addressing system, so the transition to the new standard will take years.

Internet2 is an advanced networking consortium representing over 350 U.S. universities, private businesses, and government agencies working with 66,000 institutions across the United States and international networking partners from more than 100 countries. To connect these communities, Internet2 developed a high-capacity 100 Gbps network that serves as a testbed for leading-edge technologies that may eventually migrate to the public Internet, including telemedicine, distance learning, and other advanced applications not possible with consumer-grade Internet services. The fourth generation of this network is being rolled out to provide 8.8 terabits of capacity.

INTERNET SERVICES AND COMMUNICATION TOOLS

The Internet is based on client/server technology. Individuals using the Internet control what they do through client applications on their computers, such as Web browser software. The data, including e-mail messages and Web pages, are stored on servers. A client uses the Internet to request information from a particular Web server on a distant computer, and the server sends the requested information back to the client over the Internet. Chapters 5 and 6 describe how Web servers work with application servers and database servers to access information from an organization's internal information systems applications and their associated databases. Client platforms today include not only PCs and other computers but also smartphones and tablets.

Internet Services

A client computer connecting to the Internet has access to a variety of services. These services include e-mail, chatting and instant messaging, electronic discussion groups, **Telnet**, **File Transfer Protocol (FTP)**, and the Web. Table 7.3 provides a brief description of these services.

Each Internet service is implemented by one or more software programs. All of the services may run on a single server computer, or different services may be allocated to different machines. Figure 7.8 illustrates one way that these services can be arranged in a multitiered client/server architecture.

E-mail enables messages to be exchanged from computer to computer, with capabilities for routing messages to multiple recipients, forwarding messages, and attaching text documents or multimedia files to messages. Most e-mail today is sent through the Internet. The cost of e-mail is far lower than equivalent voice, postal, or overnight delivery costs, and e-mail messages arrive anywhere in the world in a matter of seconds.

Nearly 90 percent of U.S. workplaces have employees communicating interactively using **chat** or instant messaging tools. Chatting enables two or more people who are simultaneously connected to the Internet to hold live,

TABLE 7.3 MAJOR INTERNET SERVICES

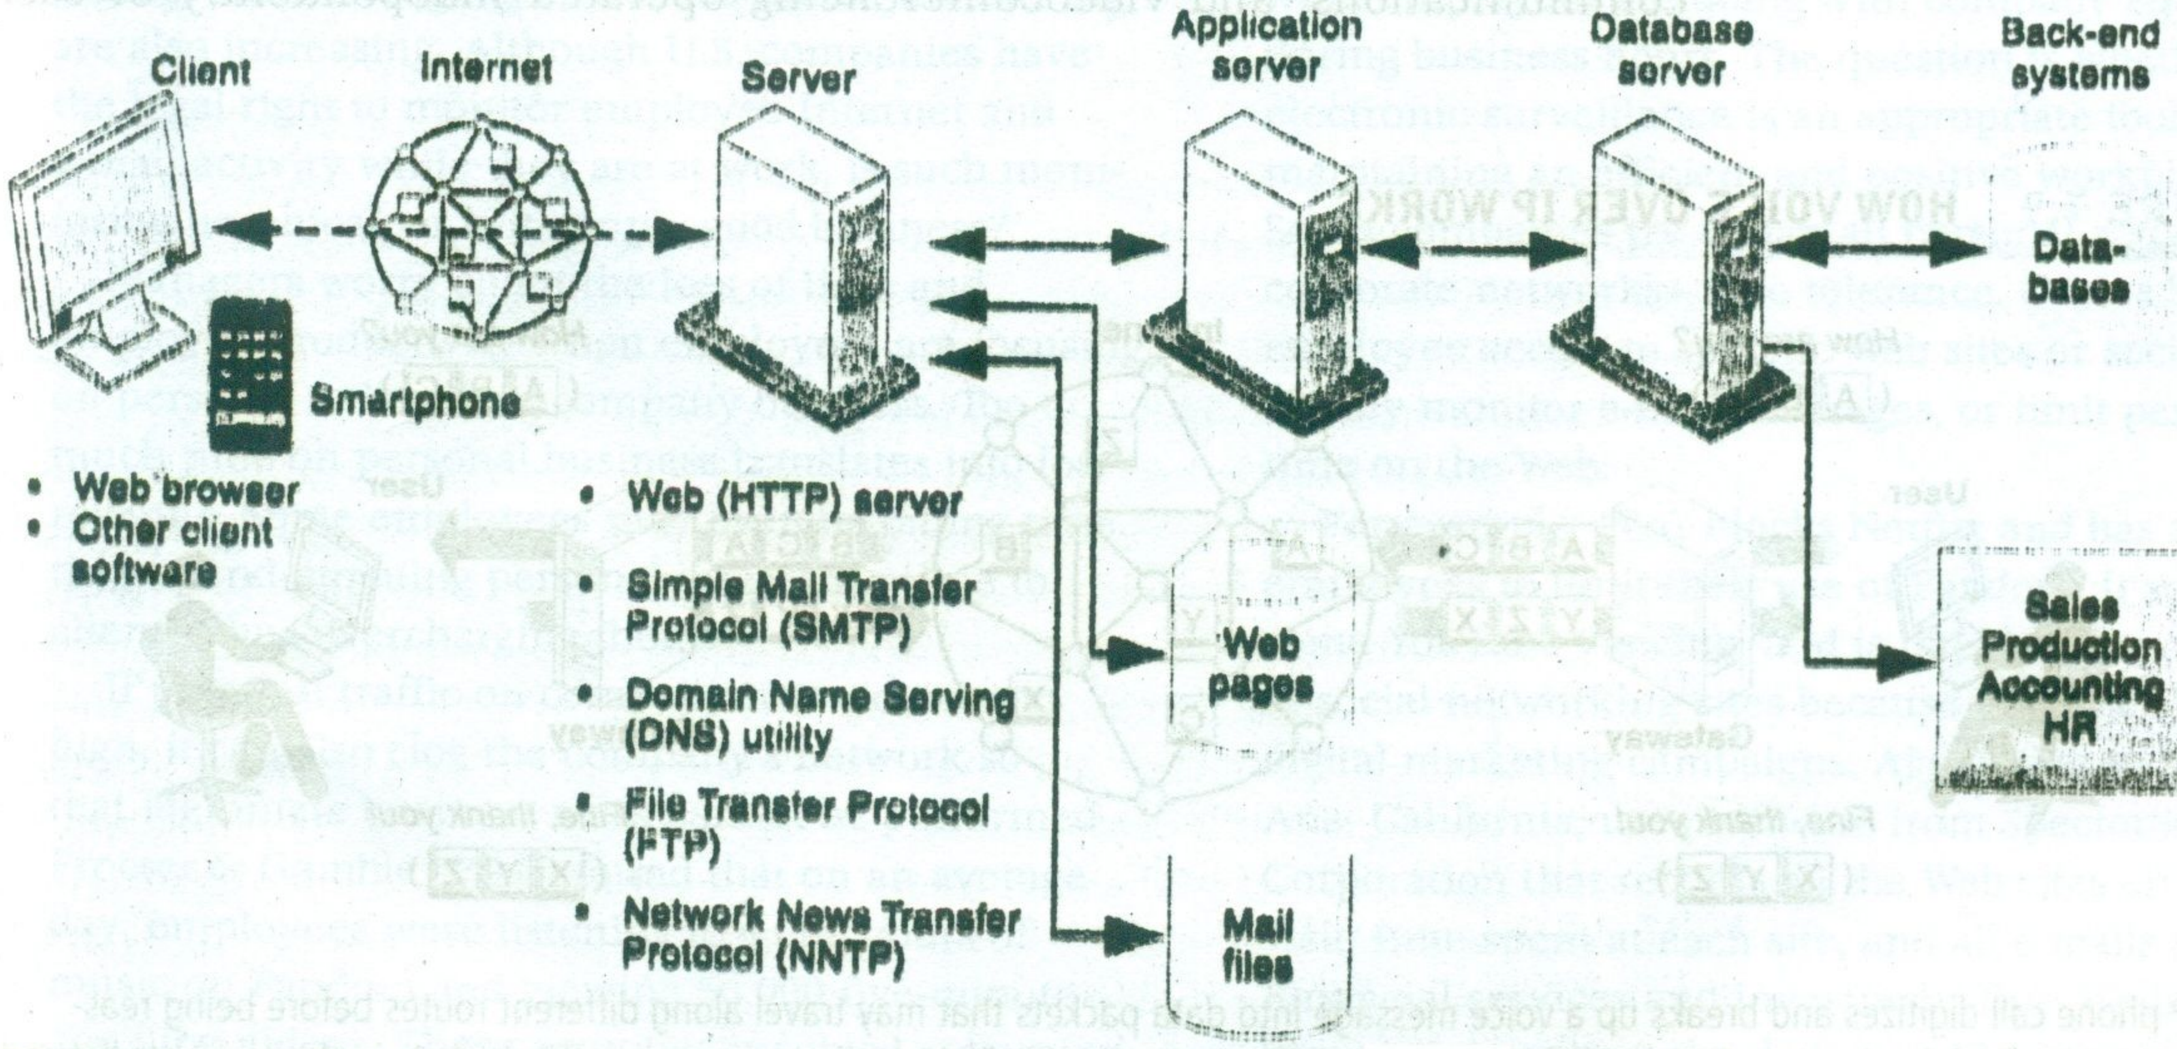
CAPABILITY	FUNCTIONS SUPPORTED
E-mail	Person-to-person messaging; document sharing
Chatting and instant messaging	Interactive conversations
Newsgroups	Discussion groups on electronic bulletin boards
Telnet	Logging on to one computer system and doing work on another
File Transfer Protocol (FTP)	Transferring files from computer to computer
World Wide Web	Retrieving, formatting, and displaying information (including text, audio, graphics, and video) using hypertext links

interactive conversations. Chat systems now support voice and video chat as well as written conversations. Many online retail businesses offer chat services on their Web sites to attract visitors, to encourage repeat purchases, and to improve customer service.

Instant messaging is a type of chat service that enables participants to create their own private chat channels. The instant messaging system alerts the user whenever someone on his or her private list is online so that the user can initiate a chat session with other individuals. Instant messaging systems for consumers include Yahoo! Messenger, Google Talk, AOL Instant Messenger, and Facebook Chat. Companies concerned with security use proprietary communications and messaging systems such as IBM Sametime.

Newsgroups are worldwide discussion groups posted on Internet electronic bulletin boards on which people share information and ideas on a defined topic, such as radiology or rock bands. Anyone can post messages on these bulletin boards for others to read. Many thousands of groups exist that discuss almost all conceivable topics.

FIGURE 7.8 CLIENT/SERVER COMPUTING ON THE INTERNET



Client computers running Web browsers and other software can access an array of servers over the Internet. These services may all run on a single server or on multiple specialized servers.

Employee use of e-mail, instant messaging, and the Internet is supposed to increase worker productivity, but the accompanying Interactive Session on Management shows that this may not always be the case. Many company managers now believe they need to monitor and even regulate their employees' online activity. But is this ethical? Although there are some strong business reasons why companies may need to monitor their employees' e-mail and Web activities, what does this mean for employee privacy?

Voice over IP

The Internet has also become a popular platform for voice transmission and corporate networking. **Voice over IP (VoIP)** technology delivers voice information in digital form using packet switching, avoiding the tolls charged by local and long-distance telephone networks (see Figure 7.9). Calls that would ordinarily be transmitted over public telephone networks travel over the corporate network based on the Internet Protocol, or the public Internet. Voice calls can be made and received with a computer equipped with a microphone and speakers or with a VoIP-enabled telephone.

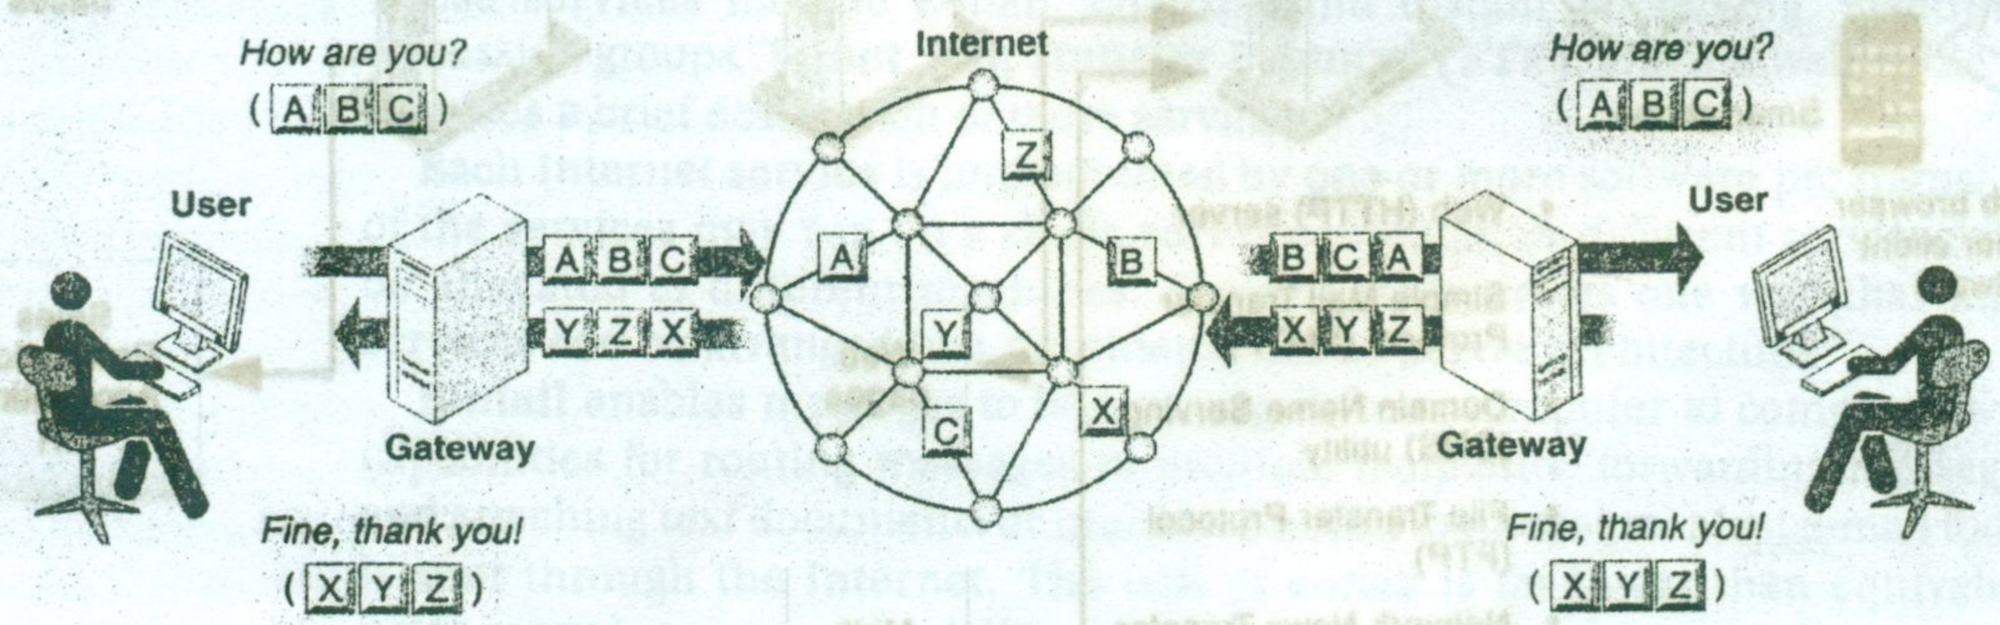
Cable firms such as Time Warner and Cablevision provide VoIP service bundled with their high-speed Internet and cable offerings. Skype offers free VoIP worldwide using a peer-to-peer network, and Google has its own free VoIP service.

Although there are up-front investments required for an IP phone system, VoIP can reduce communication and network management costs by 20 to 30 percent. For example, VoIP saves Virgin Entertainment Group $700,000 per year in long-distance bills. In addition to lowering long-distance costs and eliminating monthly fees for private lines, an IP network provides a single voice-data infrastructure for both telecommunications and computing services. Companies no longer have to maintain separate networks or provide support services and personnel for each different type of network.

Unified Communications

In the past, each of the firm's networks for wired and wireless data, voice communications, and videoconferencing operated independently of each

FIGURE 7.9 HOW VOICE OVER IP WORKS



A VoIP phone call digitizes and breaks up a voice message into data packets that may travel along different routes before being reassembled at the final destination. A processor nearest the call's destination, called a gateway, arranges the packets in the proper order and directs them to the telephone number of the receiver or the IP address of the receiving computer.

INTERACTIVE SESSION: MANAGEMENT

MONITORING EMPLOYEES ON NETWORKS: UNETHICAL OR GOOD BUSINESS?

The Internet has become an extremely valuable business tool, but it's also a huge distraction for workers on the job. Employees are wasting valuable company time by surfing inappropriate Web sites (Facebook, shopping, sports, etc.), sending and receiving personal email, talking to friends via online chat, and downloading videos and music. According to IT research firm Gartner Inc., non-work-related Internet surfing results in an estimated 40% productivity loss each year for American businesses. A recent Gallup Poll found that the average employee spends over 75 minutes per day using office computers for non-business related activity. That translates into an annual loss of $6,250 per year, per employee. An average mid-size company of 500 employees could be expected to lose $3.25 million in lost productivity due to Internet misuse.

Many companies have begun monitoring employee use of e-mail and the Internet, sometimes without their knowledge. Many tools are now available for this purpose, including SONAR, Spector CNE Investigator, iSafe, OsMonitor, IMonitor, Work Examiner, Net Spy, Activity Monitor, Mobistealth, and Spytech. These products enable companies to record online searches, monitor file downloads and uploads, record keystrokes, keep tabs on emails, create transcripts of chats, or take certain screenshots of images displayed on computer screens. Instant messaging, text messaging, and social media monitoring are also increasing. Although U.S. companies have the legal right to monitor employee Internet and e-mail activity while they are at work, is such monitoring unethical, or is it simply good business?

Managers worry about the loss of time and employee productivity when employees are focusing on personal rather than company business. Too much time on personal business translates into lost revenue. Some employees may even be billing time they spend pursuing personal interests online to clients, thus overcharging them.

If personal traffic on company networks is too high, it can also clog the company's network so that legitimate business work cannot be performed. Procter & Gamble (P&G) found that on an average day, employees were listening to 4,000 hours of music on Pandora and viewing 50,000 five-minute YouTube videos. These activities involved streaming huge quantities of data, which slowed down P&G's Internet connection.

When employees use e-mail or the Web (including social networks) at employer facilities or with employer equipment, anything they do, including anything illegal, carries the company's name. Therefore, the employer can be traced and held liable. Management in many firms fear that racist, sexually explicit, or other potentially offensive material accessed or traded by their employees could result in adverse publicity and even lawsuits for the firm. An estimated 27 percent of Fortune 500 organizations have had to defend themselves against claims of sexual harassment stemming from inappropriate email. Even if the company is found not to be liable, responding to lawsuits could run up huge legal bills. Symantec's 2011 Social Media Protection Flash Poll found that the average litigation cost for companies with social media incidents ran over $650,000.

Companies also fear leakage of confidential information and trade secrets through e-mail or social networks. Another survey conducted by the American Management Association and the ePolicy Institute found that 14 percent of the employees polled admitted they had sent confidential or potentially embarrassing company e-mails to outsiders.

U.S. companies have the legal right to monitor what employees are doing with company equipment during business hours. The question is whether electronic surveillance is an appropriate tool for maintaining an efficient and positive workplace. Some companies try to ban all personal activities on corporate networks—zero tolerance. Others block employee access to specific Web sites or social sites, closely monitor e-mail messages, or limit personal time on the Web.

For example, P&G blocks Netflix and has asked employees to limit their use of Pandora. It still allows some YouTube viewing, and is not blocking access to social networking sites because staff use them for digital marketing campaigns. Ajax Boiler in Santa Ana, California, uses software from SpectorSoft Corporation that records all the Web sites employees visit, time spent at each site, and all e-mails sent. Financial services and investment firm Wedbush Securities monitors the daily e-mails, instant messag-

ing, and social networking activity of its 1,000-plus employees. The firm's e-mail monitoring software flags certain types of messages and keywords within messages for further investigation.

A number of firms have fired employees who have stepped out of bounds. A Proofpoint survey found that one in five large U.S. companies fired an employee for violating e-mail policies in the past year. Among managers who fired employees for Internet misuse, the majority did so because the employees' e-mail contained sensitive, confidential, or embarrassing information.

No solution is problem-free, but many consultants believe companies should write corporate policies on employee e-mail, social media, and Web use. The policies should include explicit ground rules that state, by position or level, under what circumstances employees can use company facilities for e-mail, blogging, or Web surfing. The policies should also inform employees whether these activities are monitored and explain why.

IBM now has "social computing guidelines" that cover employee activity on sites such as Facebook and Twitter. The guidelines urge employees not to conceal their identities, to remember that they are personally responsible for what they publish, and to refrain from discussing controversial topics that are not related to their IBM role.

The rules should be tailored to specific business needs and organizational cultures. For example, investment firms will need to allow many of their employees access to other investment sites. A company dependent on widespread information sharing, innovation, and independence could very well find that monitoring creates more problems than it solves.

Sources: "Should Companies Monitor Their Employees' Social Media?" *Wall Street Journal*, May 11, 2014; Rhodri Marsden, "Workplace monitoring mania may be risky business," Brisbane Times, March 30, 2014; Donna Iadipaolo, "Invading Your Privacy Is Now the Norm in the Workplace," Philly.com, April 28, 2014; "Office Slacker Stats," www.staffmonitoring.com, accessed May 1, 2014; "Office Productivity Loss," Staffmonitoring.com, accessed May 1, 2014; "Workplace Privacy and Employee Monitoring," Privacy Rights Clearinghouse, June 2013; Samuel Greengard, "How Smartphone Addiction Hurts Productivity," *CIO Insight*, March 11, 2013; Emily Glazer, "P&G Curbs Employees' Internet Use," *The Wall Street Journal*, April 4, 2012; and David L. Barron, "Social Media: Frontier for Employee Disputes," Baseline, January 19, 2012.

CASE STUDY QUESTIONS

1. Should managers monitor employee e-mail and Internet usage? Why or why not?
2. Describe an effective e-mail and Web use policy for a company.
3. Should managers inform employees that their Web behavior is being monitored? Or should managers monitor secretly? Why or why not?

other and had to be managed separately by the information systems department. Now, however, firms are able to merge disparate communications modes into a single universally accessible service using unified communications technology. **Unified communications** integrates disparate channels for voice communications, data communications, instant messaging, e-mail, and electronic conferencing into a single experience where users can seamlessly switch back and forth between different communication modes. Presence technology shows whether a person is available to receive a call. Companies will need to examine how work flows and business processes will be altered by this technology in order to gauge its value.

CenterPoint Properties, a major Chicago area industrial real estate company, used unified communications technology to create collaborative Web sites for each of its real estate deals. Each Web site provides a single point for accessing structured and unstructured data. Integrated presence technology lets team members e-mail, instant message, call, or videoconference with one click.

Virtual Private Networks

What if you had a marketing group charged with developing new products and services for your firm with members spread across the United States? You would want them to be able to e-mail each other and communicate with the home office without any chance that outsiders could intercept the communications. In the past, one answer to this problem was to work with large private networking firms who offered secure, private, dedicated networks to customers. But this was an expensive solution. A much less-expensive solution is to create a virtual private network within the public Internet.

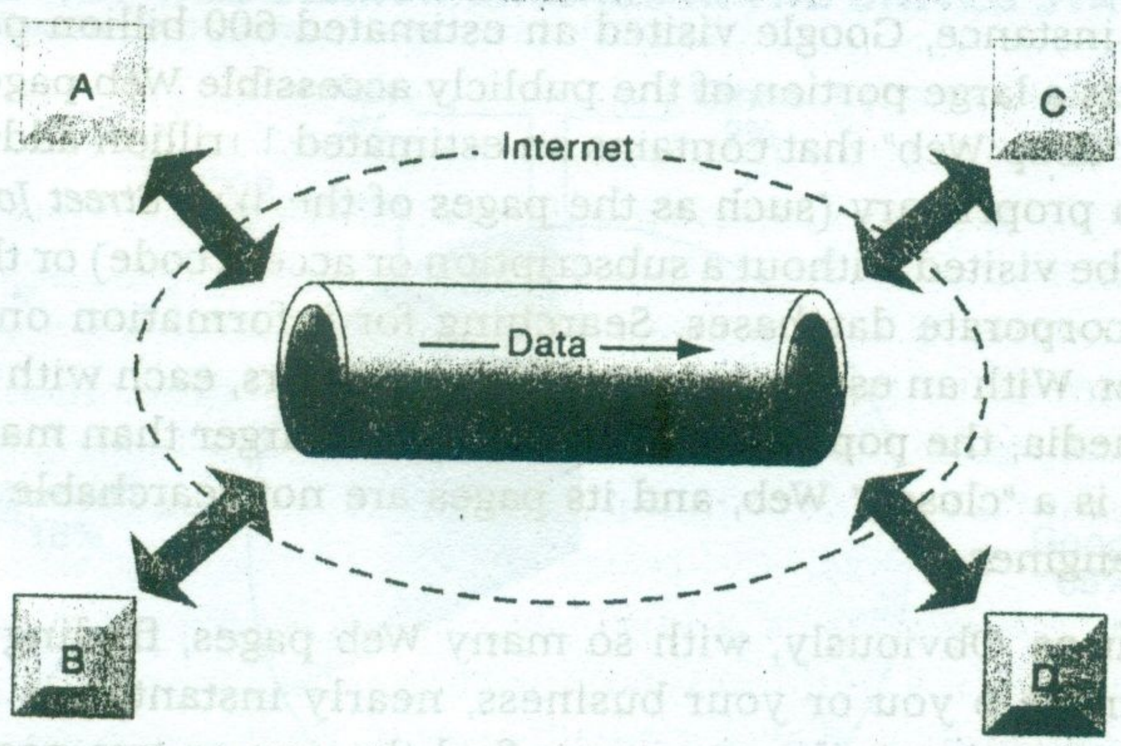
A **virtual private network (VPN)** is a secure, encrypted, private network that has been configured within a public network to take advantage of the economies of scale and management facilities of large networks, such as the Internet (see Figure 7.10). A VPN provides your firm with secure, encrypted communications at a much lower cost than the same capabilities offered by traditional non-Internet providers who use their private networks to secure communications. VPNs also provide a network infrastructure for combining voice and data networks.

Several competing protocols are used to protect data transmitted over the public Internet, including *Point-to-Point Tunneling Protocol (PPTP)*. In a process called tunneling, packets of data are encrypted and wrapped inside IP packets. By adding this wrapper around a network message to hide its content, business firms create a private connection that travels through the public Internet.

THE WEB

The Web is the most popular Internet service. It's a system with universally accepted standards for storing, retrieving, formatting, and displaying information using a client/server architecture. Web pages are formatted using hypertext with embedded links that connect documents to one another and that also link pages to other objects, such as sound, video, or animation files. When you click a graphic and a video clip plays, you have clicked a hyperlink. A typical **Web site** is a collection of Web pages linked to a home page.

FIGURE 7.10 A VIRTUAL PRIVATE NETWORK USING THE INTERNET



This VPN is a private network of computers linked using a secure "tunnel" connection over the Internet. It protects data transmitted over the public Internet by encoding the data and "wrapping" them within the Internet Protocol (IP). By adding a wrapper around a network message to hide its content, organizations can create a private connection that travels through the public Internet.

310

Hypertext

Web pages are based on a standard Hypertext Markup Language (HTML), which formats documents and incorporates dynamic links to other documents and pictures stored in the same or remote computers (see Chapter 5). Web pages are accessible through the Internet because Web browser software operating your computer can request Web pages stored on an Internet host server using the **Hypertext Transfer Protocol (HTTP)**. HTTP is the communications standard used to transfer pages on the Web. For example, when you type a Web address in your browser, such as http://www.sec.gov, your browser sends an HTTP request to the sec.gov server requesting the home page of sec.gov.

HTTP is the first set of letters at the start of every Web address, followed by the domain name, which specifies the organization's server computer that is storing the document. Most companies have a domain name that is the same as or closely related to their official corporate name. The directory path and document name are two more pieces of information within the Web address that help the browser track down the requested page. Together, the address is called a **uniform resource locator (URL)**. When typed into a browser, a URL tells the browser software exactly where to look for the information. For example, in the URL *http://www.megacorp.com/content/features/082610.html*, *http* names the protocol used to display Web pages, www.megacorp.com is the domain name, content/features is the directory path that identifies where on the domain Web server the page is stored, and 082610.html is the document name and the name of the format it is in (it is an HTML page).

Web Servers

A Web server is software for locating and managing stored Web pages. It locates the Web pages requested by a user on the computer where they are stored and delivers the Web pages to the user's computer. Server applications usually run on dedicated computers, although they can all reside on a single computer in small organizations.

The most common Web server in use today is Apache HTTP Server, followed by Microsoft Internet Information Services (IIS). Apache is an open source product that is free of charge and can be downloaded from the Web.

Searching for Information on the Web

No one knows for sure how many Web pages there really are. The surface Web is the part of the Web that search engines visit and about which information is recorded. For instance, Google visited an estimated 600 billion pages in 2013, and this reflects a large portion of the publicly accessible Web page population. But there is a "deep Web" that contains an estimated 1 trillion additional pages, many of them proprietary (such as the pages of the *Wall Street Journal Online*, which cannot be visited without a subscription or access code) or that are stored in protected corporate databases. Searching for information on Facebook is another matter. With an estimated 1.3 billion members, each with pages of text, photos, and media, the population of Web pages is larger than many estimates. But Facebook is a "closed" Web, and its pages are not searchable by Google or other search engines.

Search Engines Obviously, with so many Web pages, finding specific Web pages that can help you or your business, nearly instantly, is an important problem. The question is, how can you find the one or two pages you really want and need out of billions of indexed Web pages? **Search engines** attempt to solve the problem of finding useful information on the Web nearly instantly,

and, arguably, they are the "killer app" of the Internet era. Today's search engines can sift through HTML files, files of Microsoft Office applications, PDF files, as well as audio, video, and image files. There are hundreds of different search engines in the world, but the vast majority of search results are supplied by Google, Yahoo!, and Microsoft's Bing (see Figure 7.11).

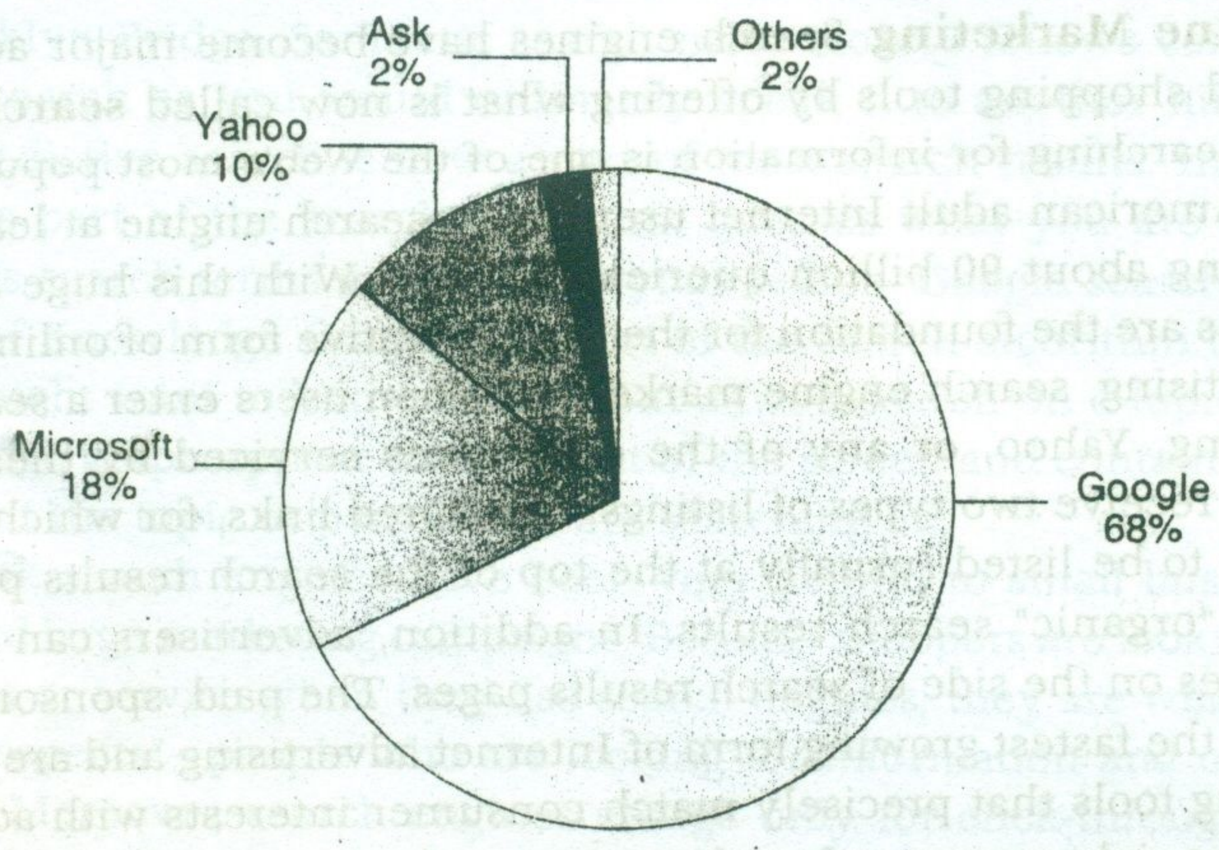
Web search engines started out in the early 1990s as relatively simple software programs that roamed the nascent Web, visiting pages and gathering information about the content of each page. The first search engines were simple keyword indexes of all the pages they visited, leaving the user with lists of pages that may not have been truly relevant to their search.

In 1994, Stanford University computer science students David Filo and Jerry Yang created a hand-selected list of their favorite Web pages and called it "Yet Another Hierarchical Officious Oracle," or Yahoo. Yahoo was not initially a search engine but rather an edited selection of Web sites organized by categories the editors found useful. Currently Yahoo relies on Microsoft's Bing for search results.

In 1998, Larry Page and Sergey Brin, two other Stanford computer science students, released their first version of Google. This search engine was different: Not only did it index each Web page's words but it also ranked search results based on the relevance of each page. Page patented the idea of a page ranking system (called PageRank System), which essentially measures the popularity of a Web page by calculating the number of sites that link to that page as well as the number of pages to which it links. The premise is that really popular Web pages are more "relevant" to users. Brin contributed a unique Web crawler program that indexed not only keywords on a page but also combinations of words (such as authors and the titles of their articles). These two ideas became the foundation for the Google search engine. Figure 7.12 illustrates how Google works.

Mobile Search With the growth of mobile smartphones and tablet computers, and with about 167 million Americans accessing the Internet via mobile devices, the nature of e-commerce and search is changing. Mobile search from smartphones and tablets made up about 50 percent of all searches in 2014, and

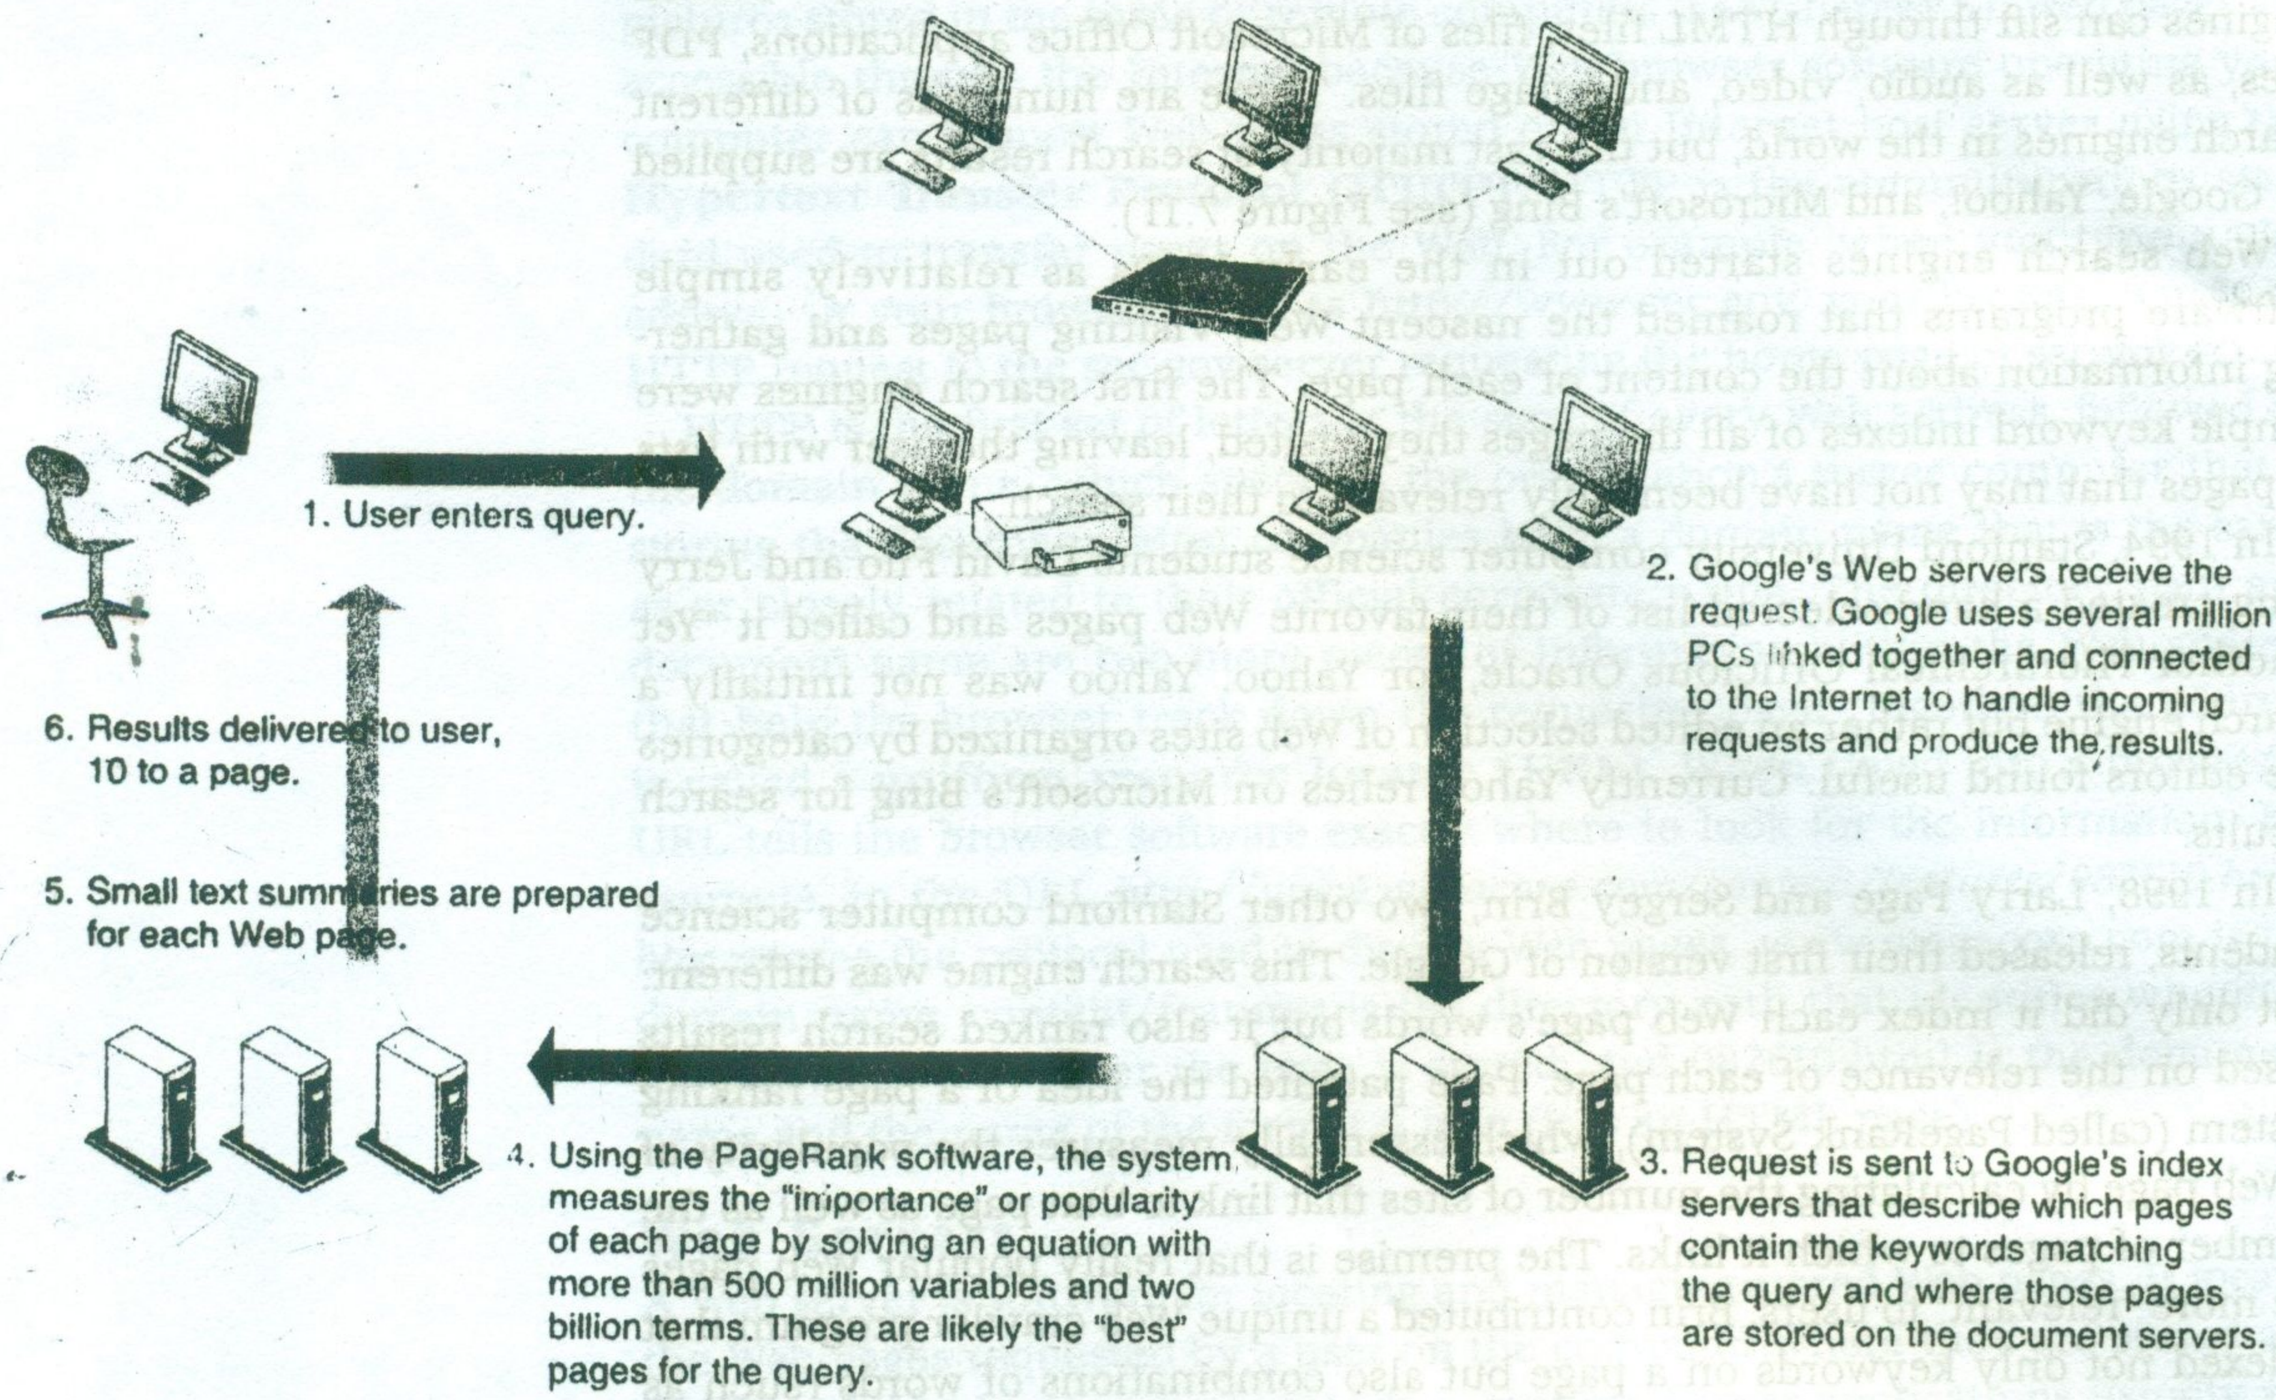
FIGURE 7.11 TOP WEB SEARCH ENGINES IN THE UNITED STATES



Google is the most popular search engine, handling nearly 70 percent of Web searches in the United States and around 90% in Europe..

Sources: Based on data from comScore Inc., February 2014.

FIGURE 7.12 HOW GOOGLE WORKS



The Google search engine is continuously crawling the Web, indexing the content of each page, calculating its popularity, and storing the pages so that it can respond quickly to user requests to see a page. The entire process takes about one-half second.

according to Google will expand rapidly in the next few years. Both Google and Yahoo have developed new search interfaces to make searching and shopping from smartphones more convenient. Amazon, for instance, sold over $1 billion in goods in 2013 through mobile searches of its store (Search Agency, 2013). While smartphones are widely used to shop, actual purchases typically take place on laptops or desktops, followed by tablets.

Search Engine Marketing Search engines have become major advertising platforms and shopping tools by offering what is now called **search engine marketing**. Searching for information is one of the Web's most popular activities: 60% of American adult Internet users use a search engine at least once a day, generating about 90 billion queries a month. With this huge audience, search engines are the foundation for the most lucrative form of online marketing and advertising, search engine marketing. When users enter a search term at Google, Bing, Yahoo, or any of the other sites serviced by these search engines, they receive two types of listings: sponsored links, for which advertisers have paid to be listed (usually at the top of the search results page), and unsponsored "organic" search results. In addition, advertisers can purchase small text boxes on the side of search results pages. The paid, sponsored advertisements are the fastest growing form of Internet advertising and are powerful new marketing tools that precisely match consumer interests with advertising messages at the right moment. Search engine marketing monetizes the value of the search process. In 2014, search engine marketing is expected to generate $22.8 billion in revenue, nearly half of all online advertising ($51 billion).

Google will for over 38% of all online advertising in 2014. About 97% of Google's revenue of $60 billion in 2013 came from online advertising, and 95% of the ad revenue came from search engine marketing (Google, 2014).

Because search engine marketing is so effective (it has the highest click-through rate and the highest return on ad investment), companies seek to optimize their Web sites for search engine recognition. The better optimized the page is, the higher a ranking it will achieve in search engine result listings. **Search engine optimization (SEO)** is the process of improving the quality and volume of Web traffic to a Web site by employing a series of techniques that help a Web site achieve a higher ranking with the major search engines when certain keywords and phrases are put into the search field. One technique is to make sure that the keywords used in the Web site description match the keywords likely to be used as search terms by prospective customers. For example, your Web site is more likely to be among the first ranked by search engines if it uses the keyword "lighting" rather than "lamps" if most prospective customers are searching for "lighting." It is also advantageous to link your Web site to as many other Web sites as possible because search engines evaluate such links to determine the popularity of a Web page and how it is linked to other content on the Web. Search engines can be gamed by scammers who create thousands of phony Web site pages and link them altogether, or link them to a single retailer's site in an attempt to fool Google's search engine. Firms can also pay so-called "link farms" to link to their site. Google changed its search algorithm in 2012. Codenamed "Penguin," the new algorithm examines the quality of links more carefully with the intent of down ranking sites that have a suspicious pattern of sites linking to them. Penguin is updated annually and published.

Google and other search engine firms are attempting to refine search engine algorithms to capture more of what the user intended, and more the "meaning" of a search. Google introduced Hummingbird, its new search algorithm in September 2013. Rather than evaluate each word separately in a search, Google's semantically informed Hummingbird will try to evaluate an entire sentence. So, if your search is a long sentence like "Google annual report selected financial data 2013," Hummingbird should be able to figure out that you really want the SEC Form 10k report filed with the Securities and Exchange Commission on March 31, 2014. How about "Italian restaurant Brooklyn Bridge"? This will return the name and location of a number of Italian restaurants in vicinity of the Brooklyn Bridge. Semantic search more closely follows conversational search, or search as you would ordinarily speak it to another human being. Google's predictive search is now a part of most search results. In **predictive search**, this part of the search algorithm guesses what you are looking for, and suggests search terms as you enter your search. Google searches also take advantage of Knowledge Graph, an effort of the search algorithm to anticipate what you might want to know more about as you search on a topic. Results of the knowledge graph appear on the right of the screen and contain more information about the topic or person you are searching on.

In general, search engines have been very helpful to small businesses that cannot afford large marketing campaigns. Because shoppers are looking for a specific product or service when they use search engines, they are what marketers call "hot prospects"—people who are looking for information and often intending to buy. Moreover, search engines charge only for click-throughs to a site. Merchants do not have to pay for ads that don't work, only for ads that receive a click. Consumers benefit from search engine marketing because ads for merchants appear only when consumers are looking for a specific product. There

are no pop-ups, Flash animations, videos, interstitials, e-mails, or other irrelevant communications to deal with. Thus, search engine marketing saves consumers cognitive energy and reduces search costs (including the cost of transportation needed to physically search for products). One study estimated the global value of search to both merchants and consumers to be more than $800 billion, with about 65 percent of the benefit going to consumers in the form of lower search costs and lower prices (McKinsey, 2011). Google and Microsoft face challenges ahead as desktop PC search growth slows, and revenues decline because the price of search engine ads is declining slightly. The growth in mobile search does not make up for the loss of desktop search revenues because mobile ads sell for generally half as much as desktop search ads.

Social Search One problem with Google and mechanical search engines is that they are so thorough: Enter a search for "ultra computers" and in .2 seconds you will receive over 300 million reponses! Search engines are not very discriminating. **Social search** is an effort to provide fewer, more relevant, and trustworthy search results based on a person's network of social contacts. In contrast to the top search engines that use a mathematical algorithm to find pages that satisfy your query, a social search Web site would review your friends' recommendations (and their friends'), their past Web visits, and their use of "Like" buttons.

In January 2013 Facebook launched Graph Search, a social network search engine that responds to user search queries with information from the user's social network of friends and connections. Graph Search relies upon the huge amount of data on Facebook that is, or can be, linked to individuals and organizations. You might use Graph Search to search for Boston restaurants that your friends like, alumni from the University of South Carolina who like Lady Gaga, or pictures of your friends before 2010. Google has developed Google +1 as a social layer on top of its existing search engine. Users can place a +1 next to the Web sites they found helpful, and their friends will be notified automatically. Subsequent searches by their friends would list the +1 sites recommended by friends higher up on the page. One problem with social search is that your close friends may not have intimate knowledge of topics you are exploring, or they may have tastes you don't appreciate. It's also possible your close friends don't have any knowledge about what you are searching for.

Semantic Search Another way for search engines to become more discriminating and helpful is to make search engines that could understand what it is we are really looking for. Called "semantic search" the goal is to build a search engine that could really understand human language and behavior. For instance, in 2012 Google's search engine began delivering more than millions of links. It started to give users more facts and direct answers, and to provide more relevant links to sites based on the search engine's estimation of what the user intended, and even on the user's past search behavior. Google's search engine is trying to understand what people are most likely thinking about when they search for something. Google hopes to use its massive database of objects (people, places, things), and smart software to provide users better results than just millions of hits. For instance, do a search on "Lake Tahoe" and the search engine will return basic facts about Tahoe (altitude, average temperature, and local fish), a map, and hotel accommodations (Efrati, 2012).

Although search engines were originally designed to search text documents, the explosion of photos and videos on the Internet created a demand for searching and classifying these visual objects. Facial recognition software can create a digital version of a human face. In 2012 Facebook introduced

its facial recognition software and combined it with tagging, to create a new feature called Tag Suggest. The software creates a digital facial print, similar to a finger print. Users can put their own tagged photo on their timeline, and their friend's timelines. Once a person's photo is tagged, Facebook can pick that person out of a group photo, and identify for others who is in the photo. You can also search for people on Facebook using their digital image to find and identify them.

Intelligent Agent Shopping Bots Chapter 11 describes the capabilities of software agents with built-in intelligence that can gather or filter information and perform other tasks to assist users. **Shopping bots** use intelligent agent software for searching the Internet for shopping information. Shopping bots such as MySimon or PriceGrabber can help people interested in making a purchase filter and retrieve information about products of interest, evaluate competing products according to criteria the users have established, and negotiate with vendors for price and delivery terms. Many of these shopping agents search the Web for pricing and availability of products specified by the user and return a list of sites that sell the item along with pricing information and a purchase link.

Web 2.0

Today's Web sites don't just contain static content—they enable people to collaborate, share information, and create new services and content online. These second-generation interactive Internet-based services are referred to as **Web 2.0**. If you have shared photos over the Internet at Flickr or another photo site, pinned a photo on Pinterest, posted a video to YouTube, created a blog, or added an app to your Facebook page, you've used some of these Web 2.0 services.

Web 2.0 has four defining features: interactivity, real-time user control, social participation (sharing), and user-generated content. The technologies and services behind these features include cloud computing, software mashups and apps, blogs, RSS, wikis, and social networks.

Mashups, which we introduced in Chapter 5, are software services that enable users and system developers to mix and match content or software components to create something entirely new. For example, Yahoo's photo storage and sharing site Flickr combines photos with other information about the images provided by users and tools to make it usable within other programming environments. Web 2.0 tools and services have fueled the creation of social networks and other online communities where people can interact with one another in the manner of their choosing.

A **blog**, the popular term for a Weblog, is a personal Web site that typically contains a series of chronological entries (newest to oldest) by its author, and links to related Web pages. Blogging is a major activity for U.S. Internet users: 74 million read blogs, and 22 million write blogs or post to blogs. The blog may include a *blogroll* (a collection of links to other blogs) and *trackbacks* (a list of entries in other blogs that refer to a post on the first blog). Most blogs allow readers to post comments on the blog entries as well. The act of creating a blog is often referred to as "blogging." Blogs can be hosted by a third-party service such as Blogger.com, TypePad.com, and Xanga.com, and blogging features have been incorporated into social networks such as Facebook and collaboration platforms such as Lotus Notes. WordPress is a leading open source blogging tool and content management system. **Microblogging**, used in Twitter, is a type of blogging that features short posts of 140 characters or less.

Blog pages are usually variations on templates provided by the blogging service or software. Therefore, millions of people without HTML skills of any kind can post their own Web pages and share content with others. The totality of blog-related Web sites is often referred to as the **blogosphere**. Although blogs have become popular personal publishing tools, they also have business uses (see Chapters 2 and 10).

If you're an avid blog reader, you might use RSS to keep up with your favorite blogs without constantly checking them for updates. **RSS**, which stands for Really Simple Syndication or Rich Site Summary, pulls specified content from Web sites and feeds it automatically to users' computers. RSS reader software gathers material from the Web sites or blogs that you tell it to scan and brings new information from those sites to you. RSS readers are available through Web sites such as Google and Yahoo, and they have been incorporated into the major Web browsers and e-mail programs.

Blogs allow visitors to add comments to the original content, but they do not allow visitors to change the original posted material. **Wikis**, in contrast, are collaborative Web sites where visitors can add, delete, or modify content on the site, including the work of previous authors. Wiki comes from the Hawaiian word for "quick."

Wiki software typically provides a template that defines layout and elements common to all pages, displays user-editable software program code, and then renders the content into an HTML-based page for display in a Web browser. Some wiki software allows only basic text formatting, whereas other tools allow the use of tables, images, or even interactive elements, such as polls or games. Most wikis provide capabilities for monitoring the work of other users and correcting mistakes.

Because wikis make information sharing so easy, they have many business uses. The U.S. Department of Homeland Security's National Cyber Security Center (NCSC) deployed a wiki to facilitate collaboration among federal agencies on cybersecurity. NCSC and other agencies use the wiki for real-time information sharing on threats, attacks, and responses and as a repository for technical and standards information. Pixar Wiki is a collaborative community wiki for publicizing the work of Pixar Animation Studios. The wiki format allows anyone to create or edit an article about a Pixar film.

Social networking sites enable users to build communities of friends and professional colleagues. Members typically create a "profile," a Web page for posting photos, videos, MP3 files, and text, and then share these profiles with others on the service identified as their "friends" or contacts. Social networking sites are highly interactive, offer real-time user control, rely on user-generated content, and are broadly based on social participation and sharing of content and opinions. Leading social networking sites include Facebook, Twitter (with 1.3 billion and 270 million monthly active users respectively in 2014), and LinkedIn (for professional contacts).

For many, social networking sites are the defining Web 2.0 application, and one that has radically changed how people spend their time online; how people communicate and with whom; how business people stay in touch with customers, suppliers, and employees; how providers of goods and services learn about their customers; and how advertisers reach potential customers. The large social networking sites are also morphing into application development platforms where members can create and sell software applications to other members of the community. Facebook alone has over 1 million developers who created over 550,000 applications for gaming, video sharing, and communicating with friends and family. We talk more about business

applications of social networking in Chapters 2 and 10, and you can find social networking discussions in many other chapters of this book. You can also find a more detailed discussion of Web 2.0 in our Learning Tracks.

Web 3.0 and the Future Web

Americans conducted about 19 billion searches in January 2014 (comScore, 2014). How many of these produced a meaningful result (a useful answer in the first three listings)? Arguably, fewer than half. Google, Yahoo, Microsoft, and Amazon are all trying to increase the odds of people finding meaningful answers to search engine queries. But with over 500 billion Web pages indexed, the means available for finding the information you really want are quite primitive, based on the words used on the pages, and the relative popularity of the page among people who use those same search terms. In other words, it's hit or miss.

To a large extent, the future of the Web involves developing techniques to make searching the 500 billion public Web pages more productive and meaningful for ordinary people. Web 1.0 solved the problem of obtaining access to information. Web 2.0 solved the problem of sharing that information with others and building new Web experiences. **Web 3.0** is the promise of a future Web where all this digital information, all these contacts, can be woven together into a single meaningful experience.

Sometimes this is referred to as the **Semantic Web**. "Semantic" refers to meaning. Most of the Web's content today is designed for humans to read and for computers to display, not for computer programs to analyze and manipulate. Semantic Search, described above, is a subset of a larger effort to make the Web more intelligent, more humanlike (W3C, 2012). Search engines can discover when a particular term or keyword appears in a Web document, but they do not really understand its meaning or how it relates to other information on the Web. You can check this out on Google by entering two searches. First, enter "Paris Hilton". Next, enter "Hilton in Paris". Because Google does not understand ordinary English, it has no idea that you are interested in the Hilton Hotel in Paris in the second search. Because it cannot understand the meaning of pages it has indexed, Google's search engine returns the most popular pages for those queries where "Hilton" and "Paris" appear on the pages.

First described in a 2001 *Scientific* American article, the Semantic Web is a collaborative effort led by the World Wide Web Consortium to add a layer of meaning atop the existing Web to reduce the amount of human involvement in searching for and processing Web information. For instance, the New York Times launched a semantic application called Longitude which provides a graphical interface to access the Times content. You can ask for stories about Germany in the last 24 hours, or a city in the United States, to retrieve all recent stories in the Times.

Views on the future of the Web vary, but they generally focus on ways to make the Web more "intelligent," with machine-facilitated understanding of information promoting a more intuitive and effective user experience. For instance, let's say you want to set up a party with your tennis buddies at a local restaurant Friday night after work. One problem is that you are already scheduled to go to a movie with another friend. In a Semantic Web 3.0 environment, you would be able to coordinate this change in plans with the schedules of your tennis buddies and the schedule of your movie friend, and make a reservation at the restaurant all with a single set of commands issued as text or voice to your handheld smartphone. Right now, this capability is beyond our grasp.

Work proceeds slowly on making the Web a more intelligent experience, in large part because it is difficult to make machines, including software programs, that are truly intelligent like humans. But there are other views of the future Web. Some see a 3-D Web where you can walk through pages in a 3-D environment. Others point to the idea of a pervasive Web that controls everything from a city's traffic lights and water usage, to the lights in your living room, to your car's rear view mirror, not to mention managing your calendar and appointments. This is referred to as the "Internet of Things."

The Internet of Things includes the widespread use and distribution of sensors. Firms like IBM, HP, and Oracle are exploring how to build smart machines, factories, and cities through extensive use of remote sensors and fast cloud computing. We provide more detail on this topic in the following section.

The "App Internet" is another element in the future Web. The growth of apps within the mobile platform is astounding: Over 80% of mobile minutes in the United States are generated through apps, only 20% using browsers. Apps give users direct access to content and are much faster than loading a browser and searching for content.

The **visual Web** is another part of the future Web. The "visual Web" refers to Web sites like Pinterest where pictures replace text documents, where users search on pictures, and where pictures of products replace display ads for products. Pinterest is a social networking site that provides users (as well as brands) with an online board to which they can "pin" interesting pictures. Looking for a blue dress, or black dress shirt? Google will deliver thousands of links to sites that sell these items. Pinterest will deliver a much smaller collection of magazine quality photos linked subtly to vendor Web sites. Considered the fastest growing Web site in history, Pinterest has 70 million monthly users and was the 35th most popular Web destination in 2014. The Instagram app is another example of the visual Web. Instagram is a photo and video sharing site that allows users to take pictures, enhance them, and share them with friends on other social sites like Facebook, Twitter, Tumblr, and Google+. In 2014 Instagram had 220 million monthly active users.

Other complementary trends leading toward a future Web 3.0 include more widespread use of cloud computing and software as a service (SaaS) business models, ubiquitous connectivity among mobile platforms and Internet access devices, and the transformation of the Web from a network of separate siloed applications and content into a more seamless and interoperable whole. These more modest visions of the future Web 3.0 are more likely to be realized in the near term.

7.4 What are the principal technologies and standards for wireless networking, communication, and Internet access?

Welcome to the wireless revolution! Cell phones, smartphones, tablets, and wireless-enabled personal computers have morphed into portable media and computing platforms that let you perform many of the computing tasks you used to do at your desk, and a whole lot more. We introduced smartphones in our discussions of the mobile digital platform in Chapters 1 and 5. **Smartphones** such as the iPhone, Android phones, and BlackBerry combine the functionality of a cell phone with that of a mobile laptop computer with Wi-Fi capability. This makes it possible to combine music, video, Internet

access, and telephone service in one device. A large part of the Internet is becoming a mobile, access-anywhere, broadband service for the delivery of video, music, and Web search.

CELLULAR SYSTEMS

In 2013, over 1.8 billion cell phones were sold worldwide. In the United States, there are 365 million cell phone subscriptions, and 164 million people have smartphones. About 167 million people access the Web using their phone (eMarketer, 2014). In a few years, smartphones will be the predominant source of searches, not the desktop PC. Digital cellular service uses several competing standards. In Europe and much of the rest of the world outside the United Sates, the standard is Global System for Mobile Communications (GSM). GSM's strength is its international roaming capability. There are GSM cell phone systems in the United States, including T-Mobile and AT&T.

A competing standard in the United States is Code Division Multiple Access (CDMA), which is the system used by Verizon and Sprint. CDMA was developed by the military during World War II. It transmits over several frequencies, occupies the entire spectrum, and randomly assigns users to a range of frequencies over time, making it more efficient than GSM.

Earlier generations of cellular systems were designed primarily for voice and limited data transmission in the form of short text messages. Today wireless carriers offer 3G and 4G networks. **3G networks**, with transmission speeds ranging from 144 Kbps for mobile users in, say, a car, to more than 2 Mbps for stationary users, offer fair transmission speeds for e-mail, browsing the Web, and online shopping, but are too slow for videos. **4G networks** have much higher speeds: 100 megabits/second download, and 50 megabits upload speed, with more than enough capacity for watching high definition video on your smartphone. Long Term Evolution (LTE) and mobile Worldwide Interoperability for Microwave Access (WiMax—see the following section) are the current 4G standards.

WIRELESS COMPUTER NETWORKS AND INTERNET ACCESS

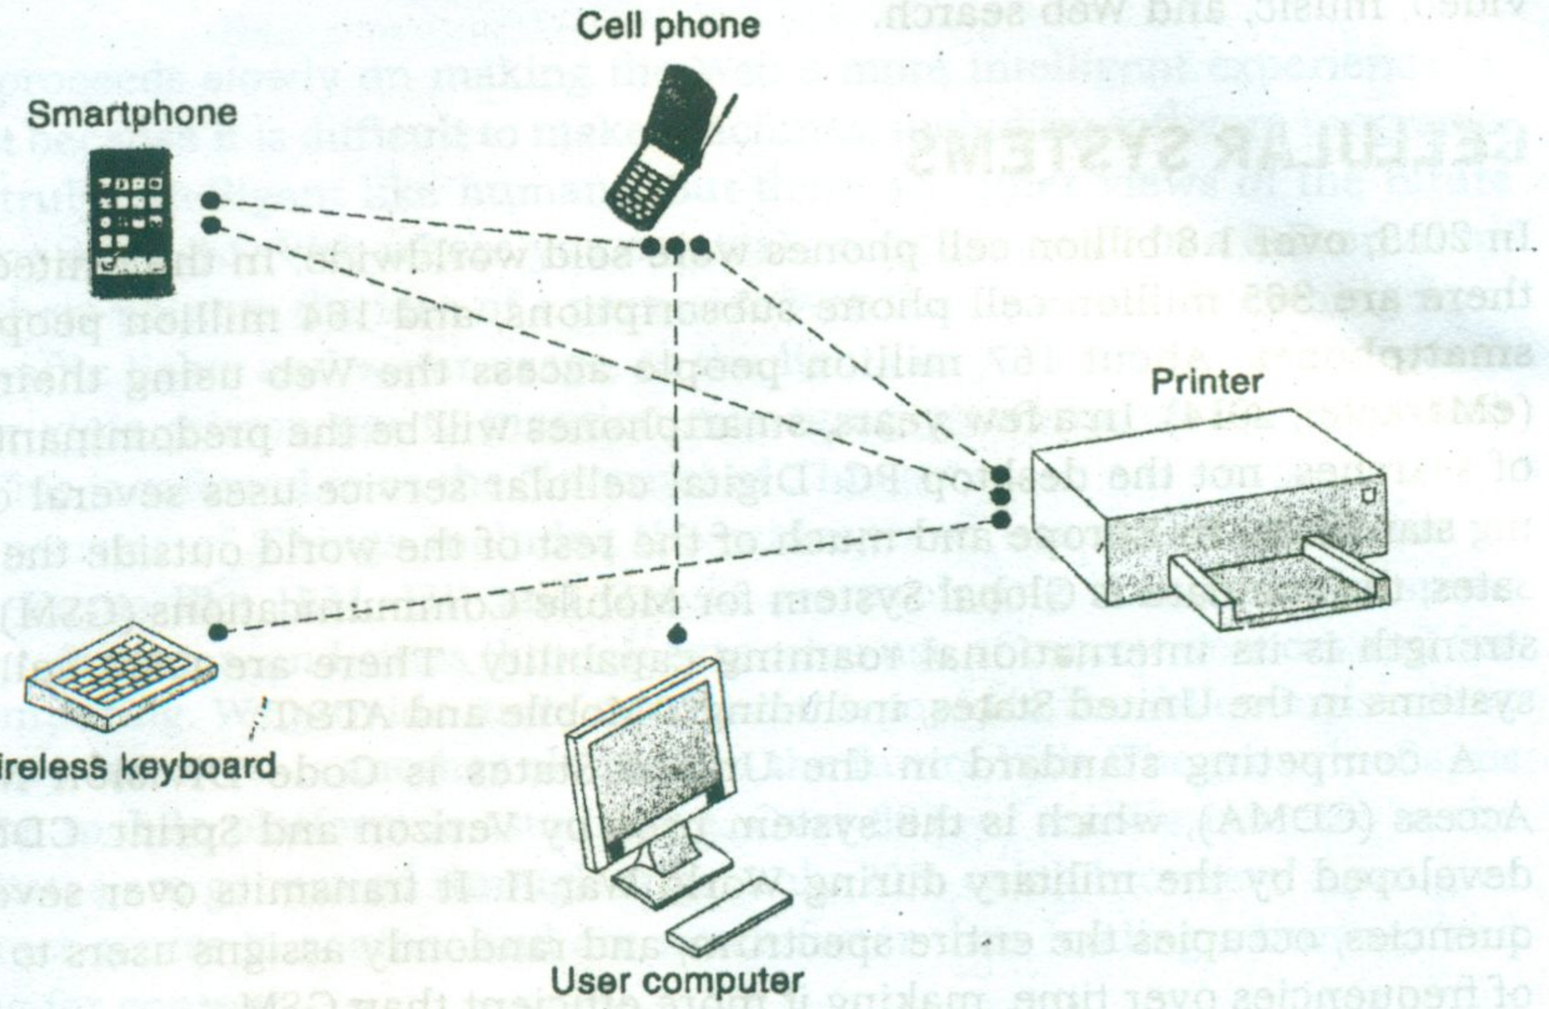
An array of technologies provide high-speed wireless access to the Internet for PCs and mobile devices. These new high-speed services have extended Internet access to numerous locations that could not be covered by traditional wired Internet services, and have made ubiquitous computing, anywhere, anytime, a reality.

Bluetooth

Bluetooth is the popular name for the 802.15 wireless networking standard, which is useful for creating small **personal area networks (PANs)**. It links up to eight devices within a 10-meter area using low-power, radio-based communication and can transmit up to 722 Kbps in the 2.4-GHz band.

Wireless phones, pagers, computers, printers, and computing devices using Bluetooth communicate with each other and even operate each other without direct user intervention (see Figure 7.13). For example, a person could direct a notebook computer to send a document file wirelessly to a printer. Bluetooth connects wireless keyboards and mice to PCs or cell phones to earpieces without wires. Bluetooth has low power requirements, making it appropriate for battery-powered handheld computers or cell phones.

FIGURE 7.13 A BLUETOOTH NETWORK (PAN)



Bluetooth enables a variety of devices, including cell phones, smartphones, wireless keyboards and mice, PCs, and printers, to interact wirelessly with each other within a small 30-foot (10-meter) area. In addition to the links shown, Bluetooth can be used to network similar devices to send data from one PC to another, for example.

Although Bluetooth lends itself to personal networking, it has uses in large corporations. For example, FedEx drivers use Bluetooth to transmit the delivery data captured by their handheld PowerPad computers to cellular transmitters, which forward the data to corporate computers. Drivers no longer need to spend time docking their handheld units physically in the transmitters, and Bluetooth has saved FedEx $20 million per year.

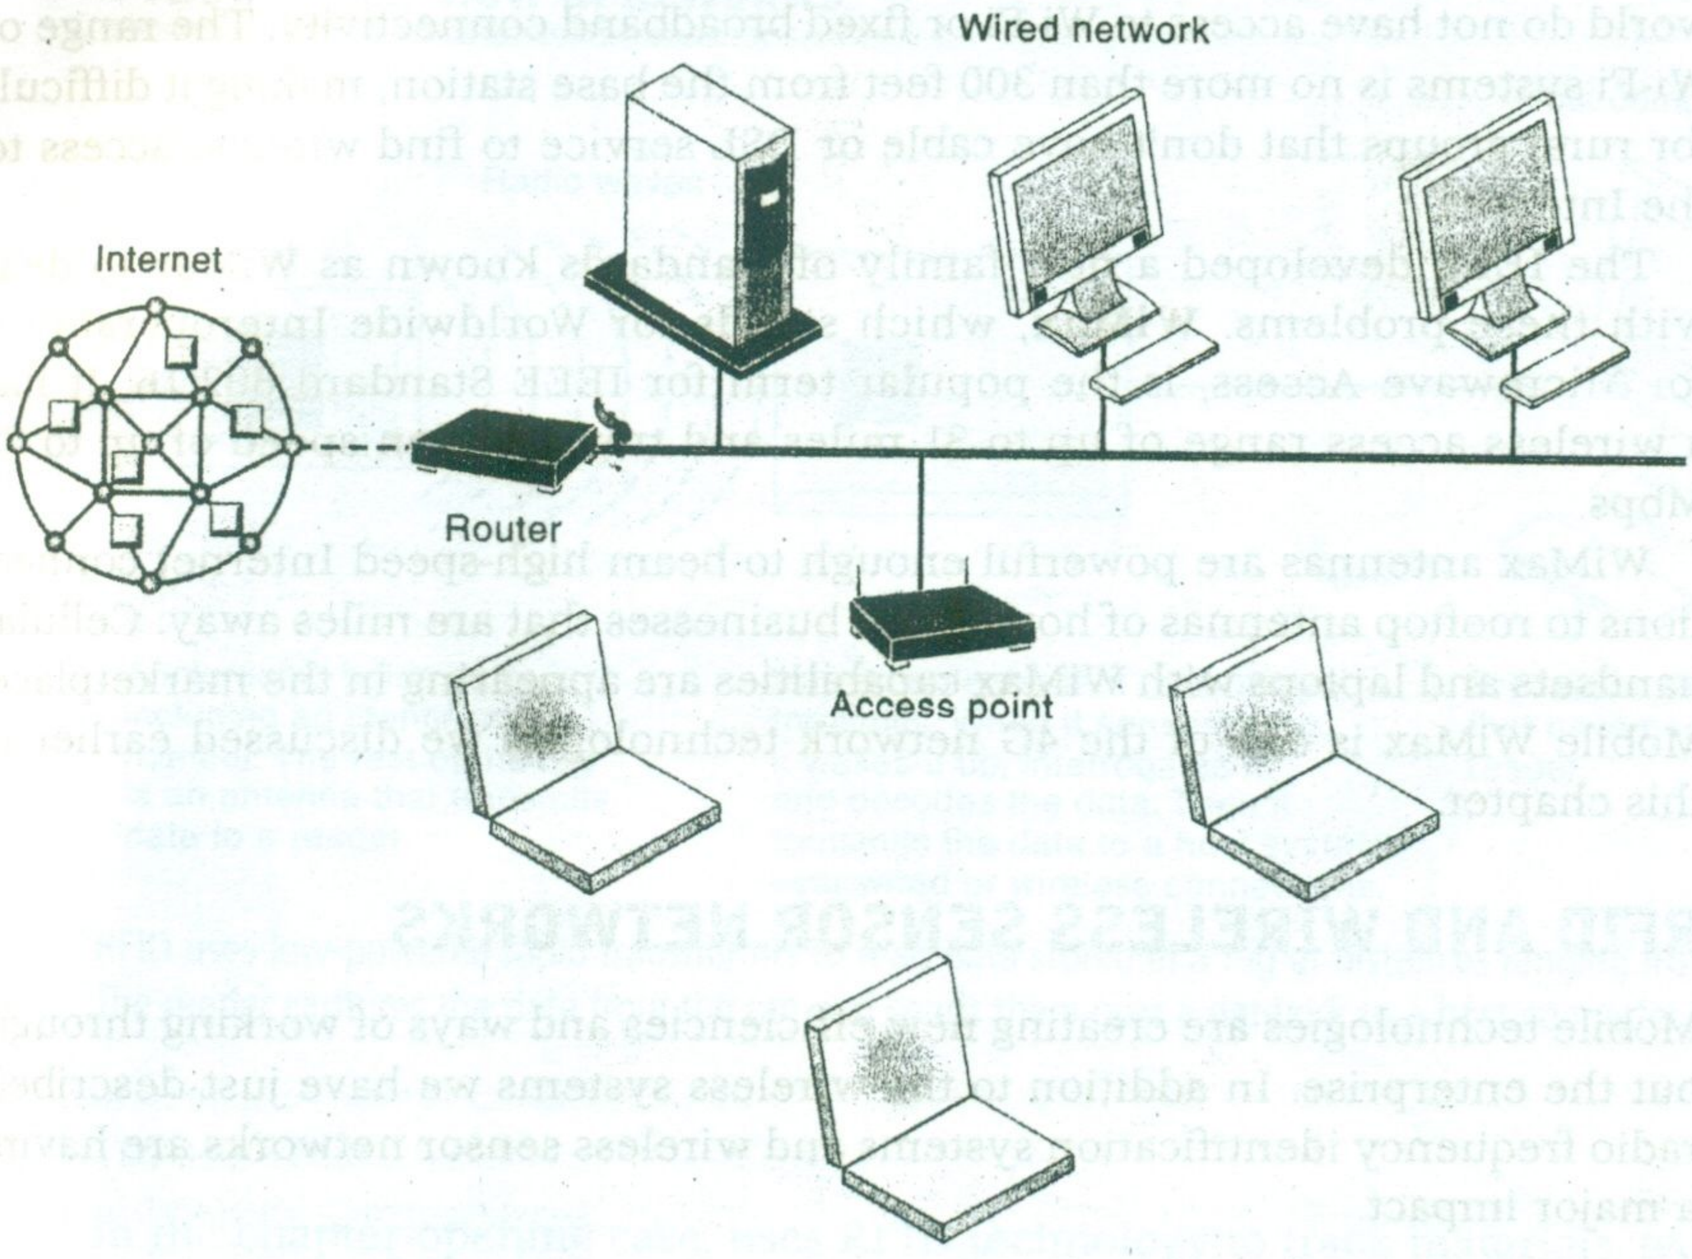
Wi-Fi and Wireless Internet Access

The 802.11 set of standards for wireless LANs and wireless Internet access is also known as **Wi-Fi**. The first of these standards to be widely adopted was 802.11b, which can transmit up to 11 Mbps in the unlicensed 2.4-GHz band and has an effective distance of 30 to 50 meters. The 802.11g standard can transmit up to 54 Mbps in the 2.4-GHz range. 802.11n is capable of transmitting over 100 Mbps. Today's PCs and netbooks have built-in support for Wi-Fi, as do the iPhone, iPad, and other smartphones.

In most Wi-Fi communication, wireless devices communicate with a wired LAN using access points. An access point is a box consisting of a radio receiver/transmitter and antennas that links to a wired network, router, or hub. Mobile access points such as Verizon's Mobile Hotspots use the existing cellular network to create Wi-Fi connections.

Figure 7.14 illustrates an 802.11 wireless LAN that connects a small number of mobile devices to a larger wired LAN and to the Internet. Most wireless devices are client machines. The servers that the mobile client stations need to use are on the wired LAN. The access point controls the wireless stations and acts as a bridge between the main wired LAN and the wireless LAN. (A bridge connects two LANs based on different technologies.) The access point also controls the wireless stations.

FIGURE 7.14 AN 802.11 WIRELESS LAN



Mobile laptop computers equipped with network interface cards link to the wired LAN by communicating with the access point. The access point uses radio waves to transmit network signals from the wired network to the client adapters, which convert them into data that the mobile device can understand. The client adapter then transmits the data from the mobile device back to the access point, which forwards the data to the wired network.

The most popular use for Wi-Fi today is for high-speed wireless Internet service. In this instance, the access point plugs into an Internet connection, which could come from a cable service or DSL telephone service. Computers within range of the access point use it to link wirelessly to the Internet.

Hotspots typically consist of one or more access points providing wireless Internet access in a public place. Some hotspots are free or do not require any additional software to use; others may require activation and the establishment of a user account by providing a credit card number over the Web.

Businesses of all sizes are using Wi-Fi networks to provide low-cost wireless LANs and Internet access. Wi-Fi hotspots can be found in hotels, airport lounges, libraries, cafes, and college campuses to provide mobile access to the Internet. Dartmouth College is one of many campuses where students now use Wi-Fi for research, course work, and entertainment.

Wi-Fi technology poses several challenges, however. One is Wi-Fi's security features, which make these wireless networks vulnerable to intruders. We provide more detail about Wi-Fi security issues in Chapter 8.

Another drawback of Wi-Fi networks is susceptibility to interference from nearby systems operating in the same spectrum, such as wireless phones, microwave ovens, or other wireless LANs. However, wireless networks based on the 802.11n standard are able to solve this problem by using multiple wireless antennas in tandem to transmit and receive data and technology called *MIMO* (multiple input multiple output) to coordinate multiple simultaneous radio signals.

WiMax

A surprisingly large number of areas in the United States and throughout the world do not have access to Wi-Fi or fixed broadband connectivity. The range of Wi-Fi systems is no more than 300 feet from the base station, making it difficult for rural groups that don't have cable or DSL service to find wireless access to the Internet.

The IEEE developed a new family of standards known as WiMax to deal with these problems. **WiMax**, which stands for Worldwide Interoperability for Microwave Access, is the popular term for IEEE Standard 802.16. It has a wireless access range of up to 31 miles and transmission speed of up to 75 Mbps.

WiMax antennas are powerful enough to beam high-speed Internet connections to rooftop antennas of homes and businesses that are miles away. Cellular handsets and laptops with WiMax capabilities are appearing in the marketplace. Mobile WiMax is one of the 4G network technologies we discussed earlier in this chapter.

RFID AND WIRELESS SENSOR NETWORKS

Mobile technologies are creating new efficiencies and ways of working throughout the enterprise. In addition to the wireless systems we have just described, radio frequency identification systems and wireless sensor networks are having a major impact.

Radio Frequency Identification (RFID)

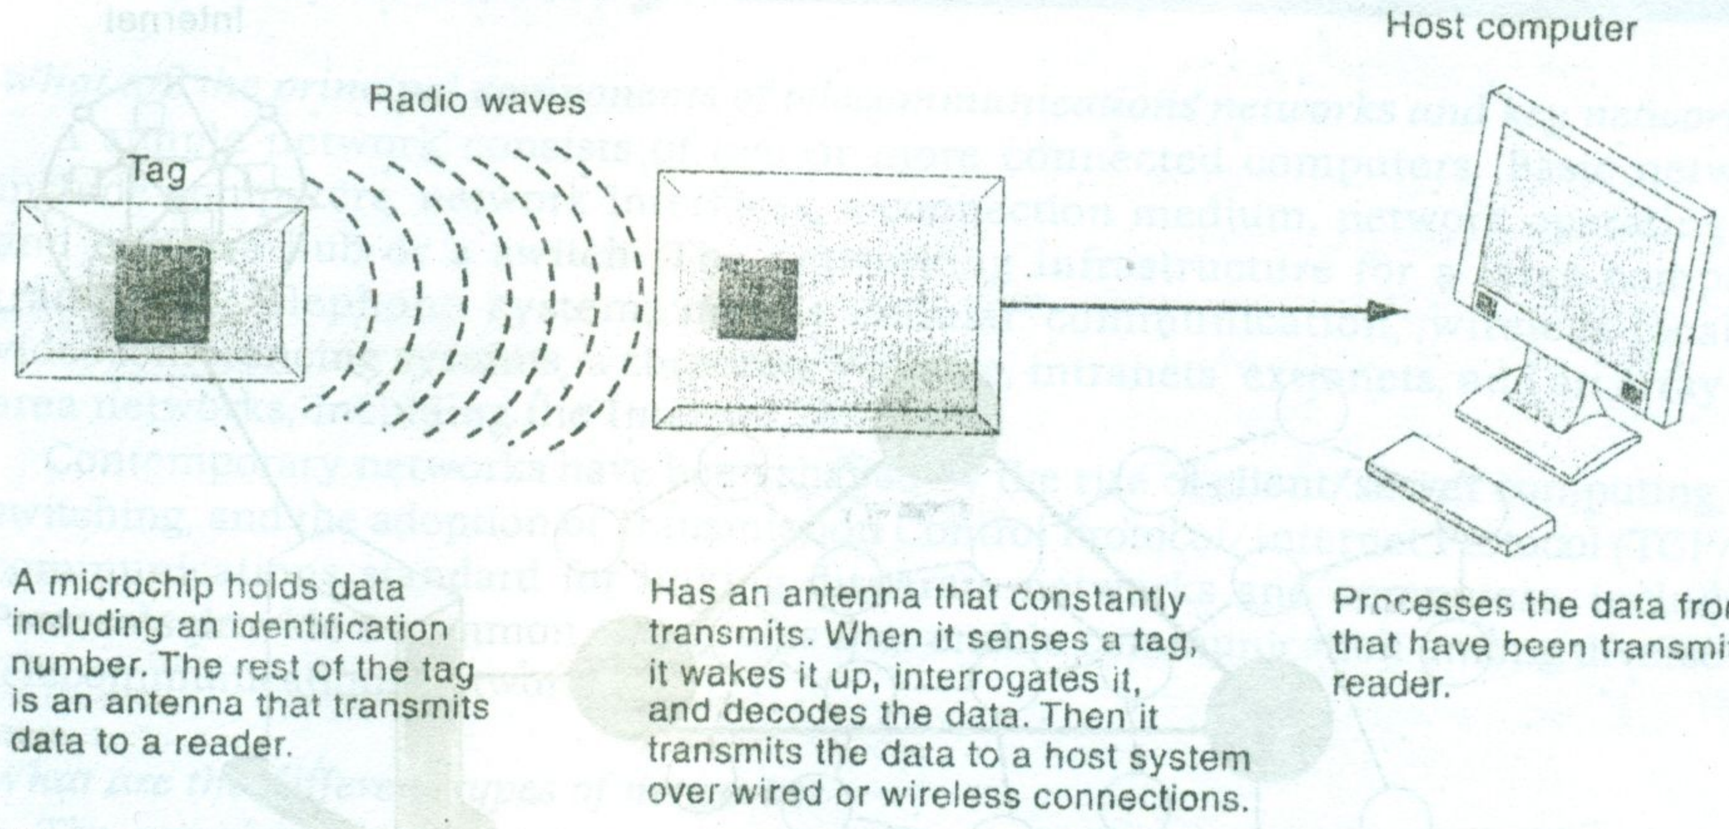
Radio frequency identification (RFID) systems provide a powerful technology for tracking the movement of goods throughout the supply chain. RFID systems use tiny tags with embedded microchips containing data about an item and its location to transmit radio signals over a short distance to RFID readers. The RFID readers then pass the data over a network to a computer for processing. Unlike bar codes, RFID tags do not need line-of-sight contact to be read.

The RFID tag is electronically programmed with information that can uniquely identify an item plus other information about the item, such as its location, where and when it was made, or its status during production. Embedded in the tag is a microchip for storing the data. The rest of the tag is an antenna that transmits data to the reader.

The reader unit consists of an antenna and radio transmitter with a decoding capability attached to a stationary or handheld device. The reader emits radio waves in ranges anywhere from 1 inch to 100 feet, depending on its power output, the radio frequency employed, and surrounding environmental conditions. When a RFID tag comes within the range of the reader, the tag is activated and starts sending data. The reader captures these data, decodes them, and sends them back over a wired or wireless network to a host computer for further processing (see Figure 7.15). Both RFID tags and antennas come in a variety of shapes and sizes.

In inventory control and supply chain management, RFID systems capture and manage more detailed information about items in warehouses or in production than bar coding systems. If a large number of items are shipped together, RFID systems track each pallet, lot, or even unit item in the shipment. This technology may help companies such as Walmart improve receiving and storage operations by improving their ability to "see" exactly what stock is stored in warehouses or on retail store shelves. Continental Tires, described

FIGURE 7.15 HOW RFID WORKS



RFID uses low-powered radio transmitters to read data stored in a tag at distances ranging from 1 inch to 100 feet. The reader captures the data from the tag and sends them over a network to a host computer for processing.

in the chapter-opening case, uses RFID technology to track materials, workers, and vehicles at its manufacturing facilities.

Walmart has installed RFID readers at store receiving docks to record the arrival of pallets and cases of goods shipped with RFID tags. The RFID reader reads the tags a second time just as the cases are brought onto the sales floor from backroom storage areas. Software combines sales data from Walmart's point-of-sale systems and the RFID data regarding the number of cases brought out to the sales floor. The program determines which items will soon be depleted and automatically generates a list of items to pick in the warehouse to replenish store shelves before they run out. This information helps Walmart reduce out-of-stock items, increase sales, and further shrink its costs.

The cost of RFID tags used to be too high for widespread use, but now it starts at around 7 cents per tag in the United States. As the price decreases, RFID is starting to become cost-effective for many applications.

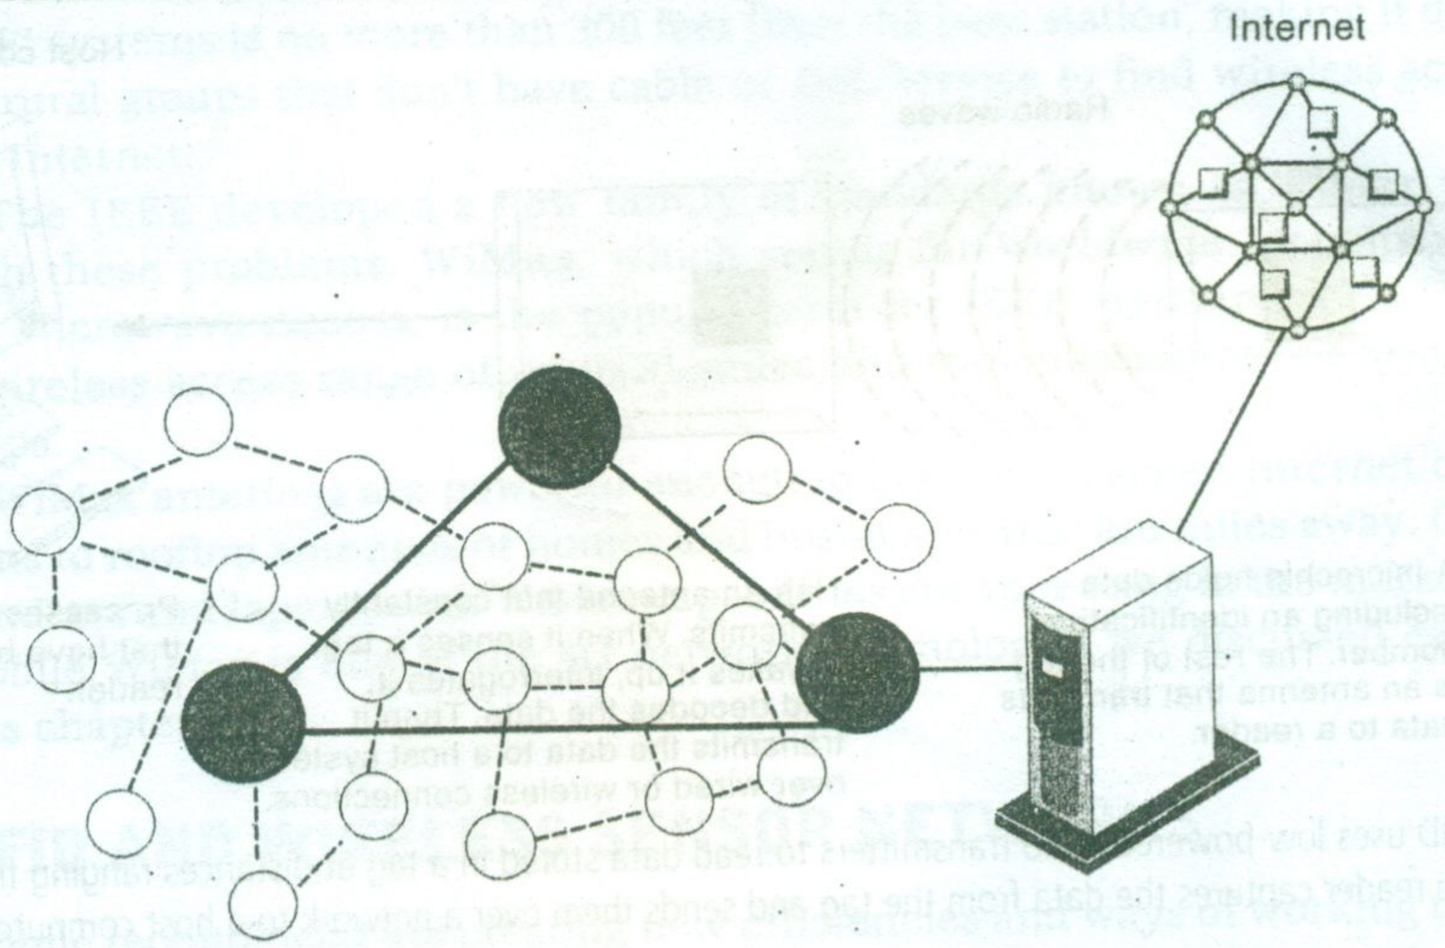
In addition to installing RFID readers and tagging systems, companies may need to upgrade their hardware and software to process the massive amounts of data produced by RFID systems—transactions that could add up to tens or hundreds of terabytes.

Software is used to filter, aggregate, and prevent RFID data from overloading business networks and system applications. Applications often need to be redesigned to accept large volumes of frequently generated RFID data and to share those data with other applications. Major enterprise software vendors, including SAP and Oracle PeopleSoft, now offer RFID-ready versions of their supply chain management applications.

Wireless Sensor Networks

If your company wanted state-of-the art technology to monitor building security or detect hazardous substances in the air, it might deploy a wireless sensor network. **Wireless sensor networks (WSNs)** are networks of interconnected wireless devices that are embedded into the physical environment to provide

FIGURE 7.16 A WIRELESS SENSOR NETWORK



The small circles represent lower-level nodes and the larger circles represent high-end nodes. Lower-level nodes forward data to each other or to higher-level nodes, which transmit data more rapidly and speed up network performance.

measurements of many points over large spaces. These devices have built-in processing, storage, and radio frequency sensors and antennas. They are linked into an interconnected network that routes the data they capture to a computer for analysis.

These networks range from hundreds to thousands of nodes. Because wireless sensor devices are placed in the field for years at a time without any maintenance or human intervention, they must have very low power requirements and batteries capable of lasting for years.

Figure 7.16 illustrates one type of wireless sensor network, with data from individual nodes flowing across the network to a server with greater processing power. The server acts as a gateway to a network based on Internet technology.

Wireless sensor networks are valuable in areas such as monitoring environmental changes, monitoring traffic or military activity, protecting property, efficiently operating and managing machinery and vehicles, establishing security perimeters, monitoring supply chain management, or detecting chemical, biological, or radiological material.

RFID systems and wireless sensor networks are major sources of "Big Data" that organizations are starting to analyze to improve their operations and decision making. Output from these systems is fueling what is called the Industrial Internet, also known as the **Internet of Things**, in which machines such as jet engines, power plant turbines, or agricultural sensors constantly gather data and send the data over the Internet for analysis. The data might signal the need to take action, such as replacing a part that's close to wearing out, restocking a product on a store shelf, starting the watering system for a soybean field, or slowing down a turbine. Over time, more and more everyday physical objects will be connected to the Internet and will be able to identify themselves to other devices, creating networks that can sense and respond as data changes. Another example is the "smart city," described in the Chapter 12 Interactive Session on Organizations. You'll see more examples of the Internet of Things in Chapter 12.

Review Summary

1. *What are the principal components of telecommunications networks and key networking technologies?*

A simple network consists of two or more connected computers. Basic network components include computers, network interfaces, a connection medium, network operating system software, and either a hub or a switch. The networking infrastructure for a large company includes the traditional telephone system, mobile cellular communication, wireless local area networks, videoconferencing systems, a corporate Web site, intranets, extranets, and an array of local and wide area networks, including the Internet.

Contemporary networks have been shaped by the rise of client/server computing, the use of packet switching, and the adoption of Transmission Control Protocol/Internet Protocol (TCP/IP) as a universal communications standard for linking disparate networks and computers, including the Internet. Protocols provide a common set of rules that enable communication among diverse components in a telecommunications network.

2. *What are the different types of networks?*

The principal physical transmission media are twisted copper telephone wire, coaxial copper cable, fiber-optic cable, and wireless transmission.

Local area networks (LANs) connect PCs and other digital devices together within a 500-meter radius and are used today for many corporate computing tasks. Wide area networks (WANs) span broad geographical distances, ranging from several miles to continents, and are often private networks that are independently managed. Metropolitan area networks (MANs) span a single urban area.

Digital subscriber line (DSL) technologies, cable Internet connections, and T1 lines are often used for high-capacity Internet connections.

3. *How do the Internet and Internet technology work, and how do they support communication and e-business?*

The Internet is a worldwide network of networks that uses the client/server model of computing and the TCP/IP network reference model. Every computer on the Internet is assigned a unique numeric IP address. The Domain Name System (DNS) converts IP addresses to more user-friendly domain names. Worldwide Internet policies are established by organizations and government bodies, such as the Internet Architecture Board (IAB) and the World Wide Web Consortium (W3C).

Major Internet services include e-mail, newsgroups, chatting, instant messaging, Telnet, FTP, and the Web. Web pages are based on Hypertext Markup Language (HTML) and can display text, graphics, video, and audio. Web site directories, search engines, and RSS technology help users locate the information they need on the Web. RSS, blogs, social networking, and wikis are features of Web 2.0.

Firms are also starting to realize economies by using VoIP technology for voice transmission and by using virtual private networks (VPNs) as low-cost alternatives to private WANs.

4. *What are the principal technologies and standards for wireless networking, communication, and Internet access?*

Cellular networks are evolving toward high-speed, high-bandwidth, digital packet-switched transmission. Broadband 3G networks are capable of transmitting data at speeds ranging from 144 Kbps to more than 2 Mbps. 4G networks capable of transmission speeds of 100 Mbps are starting to be rolled out.

Major cellular standards include Code Division Multiple Access (CDMA), which is used primarily in the United States, and Global System for Mobile Communications (GSM), which is the standard in Europe and much of the rest of the world.

Standards for wireless computer networks include Bluetooth (802.15) for small personal area networks (PANs), Wi-Fi (802.11) for local area networks (LANs), and WiMax (802.16) for metropolitan area networks (MANs).

Radio frequency identification (RFID) systems provide a powerful technology for tracking the movement of goods by using tiny tags with embedded data about an item and its location. RFID readers read the radio signals transmitted by these tags and pass the data over a network to a computer for processing. Wireless sensor networks (WSNs) are networks of interconnected wireless sensing and transmitting devices that are embedded into the physical environment to provide measurements of many points over large spaces.

Key Terms

3G Networks, 319
4G networks, 319
Bandwidth, 297
Blog, 315
Blogosphere, 316
Bluetooth, 319
Broadband, 290
Cable Internet connections, 299
Chat, 304
Digital subscriber line (DSL), 298
Domain name, 299
Domain Name System (DNS), 299
E-mail, 304
File Transfer Protocol (FTP), 304
Hertz, 297
Hotspots, 321
Hubs, 290
Hypertext Transfer Protocol (HTTP), 310
Instant messaging, 305
Internet of Things, 324
Internet Protocol (IP) address, 299
Internet service provider (ISP), 299
Internet2, 304
IPv6, 304
Local area network (LAN), 296
Metropolitan area network (MAN), 297
Microblogging, 315
Modem, 295
Network operating system (NOS), 290
Packet switching, 293
Peer-to-peer, 297
Personal area networks (PANs), 319
Predictive search, 313
Protocol, 294
Radio frequency identification (RFID), 322
Router, 290
RSS, 316
Search engines, 310
Search engine marketing, 312
Search engine optimization (SEO), 313
Semantic Web, 317
Shopping bots, 315
Smartphones, 318
Social networking, 316
Social search, 314
Software-defined networking (SDN), 291
Switch, 290
T1 lines, 299
Telnet, 304
Transmission Control Protocol/Internet Protocol (TCP/IP), 294
Unified communications, 308
Uniform resource locator (URL), 310
Virtual private network (VPN), 309
Visual Web, 318
Voice over IP (VoIP), 306
Web 2.0, 315
Web 3.0, 317
Web site, 309
Wide area networks (WANs), 297
Wi-Fi, 320
Wiki, 316
WiMax, 322
Wireless sensor networks (WSNs), 323

MyMISLab

Go to **mymislab.com** to complete the problems marked with this icon ✪.

Review Questions

7-1 What are the principal components of telecommunications networks and key networking technologies?

- Describe the features of a simple network and the network infrastructure for a large company.
- Describe one of the major problems facing corporations today with respect to telecommunications.
- Name and describe the principal technologies and trends that have shaped contemporary telecommunications systems.

7-2 What are the different types of networks?

- List and describe the different types of physical transmission media.
- Distinguish between a LAN, MAN, and WAN.

7-3 How do the Internet and Internet technology work, and how do they support communication and e-business?

- Define the Internet, describe how it works, and explain how it provides business value.
- Explain how the Domain Name System (DNS) and IP addressing system work.
- List and describe the principal Internet services.

- Define and describe VoIP and virtual private networks, and explain how they provide value to businesses.
- Define unified communications and explain how it provides value to businesses.
- Describe how online search technologies are used for marketing.

7-4 What are the principal technologies and standards for wireless networking, communications, and Internet access?

- Define Bluetooth, Wi-Fi, WiMax, and 3G and 4G networks.
- Describe the capabilities of each and for which types of applications each is best suited.
- Define RFID, explain how it works, and describe how it provides value to businesses.
- Define WSNs, explain how they work, and describe the kinds of applications that use them.

Discussion Questions

7-5 It has been said that within the next few years, smartphones will become the single most important digital device we own. Discuss the implications of this statement.

7-6 Does anyone own the Internet? Discuss how the Internet is currently governed and consider whether countries should be able to impose its own rules on the Internet within their own borders.

7-7 Compare Wi-Fi and high-speed cellular systems for accessing the Internet. What are the advantages and disadvantages of each?

Hands-On MIS Projects

The projects in this section give you hands-on experience evaluating and selecting communications technology, using spreadsheet software to improve selection of telecommunications services, and using Web search engines for business research.

Management Decision Problems

7-8 Your company supplies ceramic floor tiles to Leroy Merlin, Bricorama, and other home improvement stores in France. You have been asked to start using radio frequency identification tags on each case of tiles you ship to help your customers improve the management of your products and those of other suppliers in their warehouses. Use the Web to identify the cost of hardware, software, and networking components for an RFID system for your company. What factors should be considered? What are the key decisions that have to be made in determining whether your firm should adopt this technology?

7-9 BD sells medical devices and instrument systems to hospitals, health clinics, medical offices, and laboratories in over 50 countries. The company employs over 30,000 people around the world, including account managers, customer service and support representatives, and warehouse staff. Management is considering adopting a system for unified communications. What factors should be considered? What are the key decisions that have to be made in determining whether to adopt this technology? Use the Web, if necessary, to find out more about unified communications and its costs.

Improving Decision Making: Using Spreadsheet Software to Evaluate Wireless Services

Software skills: Spreadsheet formulas, formatting
Business skills: Analyzing telecommunications services and costs

7-10 In this project, you'll use the Web to research alternative wireless services and use spreadsheet software to calculate wireless service costs for a sales force.

You would like to equip your sales force of 35, based in Dublin, Ireland, with mobile phones that have capabilities for voice transmission, text messaging, and taking and sending photos. Use the Web to select

a wireless service provider that provides international service as well as good service in your home area. Examine the features of the mobile handsets and wireless plans offered by each of these vendors. Assume that each of the 35 salespeople will need to spend three hours per weekday between 8 a.m. and 6 p.m. on mobile voice communication, send 30 text messages per weekday, use 1 gigabyte of data per month, and send five photos per week. Use your spreadsheet software to determine the wireless service and handset that will offer the best pricing per user over a two-year period. For the purposes of this exercise, you do not need to consider corporate discounts.

Achieving Operational Excellence: Using Web Search Engines for Business Research

Software skills: Web search tools
Business skills: Researching new technologies

7-11 This project will help develop your Internet skills in using Web search engines for business research.

Use Google and Bing to obtain information about the Internet of Things. If you wish, try some other search engines as well. Compare the volume and quality of information you find with each search tool. Which tool is the easiest to use? Which produced the best results for your research? Why?

Collaboration and Teamwork Project

7-12 In MyMISLab, you will find a Collaboration and Teamwork Project dealing with the concepts in this chapter. You will be able to use Google Drive, Google Docs, Google Sites, Google +, or other open-source collaboration tools to complete the assignment.

RFID Propels the Angkasa Library Management System

CASE STUDY

RFID (Radio-frequency Identification) can play a key role in library management. The major advantage of using RFID is that it ensures traceability and security. In addition, RFID simplifies transactional processes at the library and can help to cut costs and save time. Across the globe, libraries are setting up the infrastructure for RFID as a replacement for manual management or barcodes. RFID tags can be embedded within a book, and unlike other forms of labeling, these tags can store additional information, like author and title. Using this system also speeds up the process of checking books in and out, prevents theft and is of use in inventory management. The Allianze University College of Medical Sciences (AUCMS) has successfully set up RFID tagging as part of its library management system.

AUCMS is a premier institute offering courses in medicine and applied health sciences in Penang, Malaysia. The institution has a vast array of resources, including books, journals, newspapers, and e-books, and the library caters to a large population of students and staff. AUCMS partnered with the Sains Group, which is a globally renowned software solutions and IT services provider, to install and implement the the company's Angkasa Library Management system. This system manages many of the core functions of the library, including acquisition, cataloguing, circulation, subscriptions to journals, and the management of an open-access catalogue available to members on the Web.

The system is primarily useful in terms of being a repository of all library records. It enables the automatization of library processes that would earlier have been manual and paper-based. The Acquisition Module manages the process of acquiring the books, magazines, journals and other materials for the library. This automates the entire process of acquisition, from the approval of orders to placing them, the recording of receipts and invoices, and even manages expenditure and maintains budgets. In a typical instance, lists of required reading material are sent over to the library heads, a final list is determined, all of the important information is recorded on the system (invoice numbers, supplier names, and so on). The system provides a printable order form, which is then signed off on by the library head and sent to suppliers.

Information on newly acquired books is entered into the system through the Cataloguing Module, which assigns classification numbers to each book. Data is entered in the standard machine-readable format AACR2, which stands for Anglo-American Cataloguing Rules. Based on this standard format, the technical department enters details into the system, which then generates an index card in the same format. This card serves as an index for searching, and makes import/export of bibliographic data in standard exchange formats possible. The Circulation Module performs all the functions related to circulation, providing suitable checks at every stage. Since books may be circulated for multiple reasons such as book binding and display, in addition to being issued to members, this module records the current status of a circulated library item. The Serials Management Module controls the library's subscriptions to periodicals and monitors the scheduled arrival of individual issues, in addition to recording budget information.

An important part of the library management system is the Online Public Access Catalogue (OPAC), which provides online access to library resources. The OPAC module supports keyword searches in many languages, and page navigation is user-friendly. Since the catalogue is Web-enabled, this makes it very easy to keep information on the availability of materials in the library up-to-date. Another member-centric module is the Member Management module, which stores profiles for all library members and includes a photograph and a registration number. This module sends reminders to members to return their books on time, and also records fines if items are not returned.

The first step towards setting this comprehensive system up at the AUCMS library was the tagging of the books. The RFID tag is the most important aspect of the system because the tag establishes a unique identitity for each library asset, from books to magazines, and functions as the item's identity in the system. The tag contains electronically stored information that may be read from up to several meters. Unlike a bar code, the RFID tag does not need to be in the line of sight of a barcode reader. All library issuing and returns counters have a touch screen with an RFID glass sensor. A handheld reader is also required to scan membership cards.

Borrowers simply place books and other material they want to borrow on the touch screen, and hand over their cards to library staff. The screen reads RFID tags, while the card reader opens up a member's profile on the system. The staff member assigns material to the borrower in question, and an issuing receipt is generated, with the date of return clearly marked. This system eliminates the need for physical records, stamping, and other time-consuming activities involved in the process of checking a book out of a library. In addition, the system also minimizes human error.

Returning books is also much easier now, with members simply dropping their books at any time into a specially designed "drop box" machine in the library. Members drop the books to be returned in the box one at a time. A scanner within the drop box scans the RFID tag of the dropped book and marks it as returned in the system. Member accounts are updated accordingly and a proof-of-return slip is issued automatically. The wl le process is accomplished through a link between the RFID system and the library's host computer. This eliminates the need for staff who would earlier have had to process all book returns, and is also convenient for members, since they can return their books even when the library is closed. The status of returned items is updated instantly, which means other members who are waiting for a particular item can acquire it at once.

Another major task that has always occupied library staff is shelf management. Since a large number of users access the library and do not put books back on shelves in the right order, shelves can become difficult to manage. Books that have been returned also have to be returned to specific shelves. The shelf management system in place at the AUCMS library consists of a portable scanner and a base station. The solution is designed to cover three main operations: searching for individual books, inventory checks of library stock, and locating and replacing for books that have been shelved incorrectly. Each book has a shelf ID, which is the location identification code (this code enables section-wise and rack-wise identification of books). This information is saved in a central database against the book's information and RFID tag, as well as being linked to OPAC. The portable scanner fetches the information stored in a book's RFID tag, and using the shelving information also stored in the database, library staff can move books back into their places. Similarly, missing books can also be identified, with the portable scanner pulling up the records of all the books meant to be on a specific shelf. The collected information is compared against the library database to generate a report of missing books. This system has improved accuracy and enabled faster shelving (20 books per minute). It also directs OPAC users to the exact physical locations of the books they are looking for. In total, the automatization of all of these library processes has resulted in a significant decrease in the amount of manpower required to staff the library.

Finally, one of the most pervasive problems in any library is theft. RFID tagging is extremely useful in preventing theft through the installation of Electronic Article Surveillance (EAS) gates in libraries. Just like alarm systems in stores, a library member attempting to leave with an unissued item would trigger an alarm, because the gates can sense the book's RFID tag within a range of 1 meter without interference of magnetic items. The EAS gates are linked to the library's surveillance station. When someone passes through with an unissued item, an alarm is sounded and the gates are raised, an alarm is sounded, and the camera at the gates takes a photograph of the person and sends it to the surveillance station. Theft detection is a stand-alone technology that operates independently of the library database, but it is an integral part of managing the library at AUCMS.

Sources: "Integrated Library Management Solution," SAINs Corporation, Sarawak, Malaysia, http://www.sains.com.my/solutions.php?id=47, February 11, 2015; http://www.allianzeunicollege.edu.my/; "Library Management Solution for Penang Allianze University College of Medical Sciences," *Sarawak Information Systems Newsletter*, 2012.

CASE STUDY QUESTIONS

7-13 How can RFID technology simplify basic library processes like the lending and returns of books?

7-14 What kind of technology does your school or university library use? Does IT play a crucial role in managing operations?

7-15 How is RFID technology helping the Allianze University College of Medical Sciences library to detect and prevent the theft of books?

Case contributed by Sahil Raj, Punjabi University.

MyMISLab

Go to **mymislab.com** for the following Assisted-graded writing questions.

7-16 Compare Web 2.0 and Web 3.0.

7-17 How do social search, semantic search, and mobile search differ from searching for information on the Web using conventional search engines?

Chapter 7 References

Boutin, Paul. "Search Tool on Facebook Puts Network to Work." *New York Times* (March 20, 2013).

comScore. "comScore Releases June 2014 U.S. Search Engine Rankings." (July 21, 2014).

Efrati, Amir. "Google's Search Revamp: A Step Closer to AI." *Wall Street Journal* (March 14, 2012).

eMarketer. "US Mobile Users: 2014 Complete Forecast." (Alison M McCarthy, eMarketer Report (April 2014).

Google, Inc. "SEC Form 10k for the Fiscal Year Ending December 30, 2013" Google, Inc. March 31, 2014.

Holmes, Sam and Jeffrey A. Trachtenberg. "Web Addresses Enter New Era." *Wall Street Journal* (June 21, 2011).

ICANN. "ICANN Policy Update." 10, No. 9 (September 2010).

Lahiri, Atanu, I. "The Disruptive Effect of Open Platforms on Markets for Wireless Services." *Journal of Management Information Systems* 27, No. 3 (Winter 2011).

Marin Software, Inc. "The State of Mobile Search Advertising in the US: How the Emergence of Smartphones and Tablets Changes Paid Search." Marin Software Inc. (2012).

McKinsey&Company. "The Impact of Internet Technologies: Search (July 2011).

Miller, Claire Cain. "Google, a Giant in Mobile Search, Seeks New Ways to Make It Pay," *New York Times* (April 24, 2011).

Murphy, Chris. "The Internet of Things." *Information Week* (August 13, 2012).

National Telecommunications and Information Agency. "NTIA Announces Intent to Transition Key Internet Domain Name Functions." Press Release. March 14, 2014;

Panko, Raymond R. and Julia Panko. *Business Data Networks and Telecommunications* 8e. Upper Saddle River, NJ: Prentice-Hall (2011).

SearchAgency.com. "Mobile Drives Increased Spend and Clicks on Both Google and Bing ." (July 15, 2014).

Shaw, Tony. "Innovation Web 3.0." Baseline (March/April 2011).

Simonite, Tom. "Social Indexing." Technology Review (May/June 2011).

"Software Defined Networking." Global Knowledge (2014).

SupplyChainBrain. "RFID's Role in Today's Supply Chain (November 4, 2013).

"The Internet of Things." *McKinsey Quarterly* (March 2010).

Winkler, Rolfe. "As Google Builds Out Own Content, Some Advertisers Feel Pushed Aside." *Wall Street.Journal* (August 18, 2014).

Winkler, Rolfe. "Getting More than Just Words in a Google Search Result." *Wall Street Journal* (Aug. 18, 2014).

Wittman, Art. "Here Comes the Internet of Things" *Information Week* (July 22, 2013).

Worldwide Web Consortium, " Semantic Web." w3.org/standards/semanticweb (October 18, 2012).

Worthen, Ben and Cari Tuna. "Web Running Out of Addresses." *Wall Street Journal* (Feb 1, 2011).

Wyatt, Edward. "U.S. to Cede Its Oversight of Addresses on Internet." New York Times, March 14, 2014

CHAPTER 8

Securing Information Systems

Securing Information Systems

CHAPTER 8

LEARNING OBJECTIVES

After reading this chapter, you will be able to answer the following questions:

1. Why are information systems vulnerable to destruction, error, and abuse?
2. What is the business value of security and control?
3. What are the components of an organizational framework for security and control?
4. What are the most important tools and technologies for safeguarding information resources?

MyMISLab™

Visit **mymislab.com** for simulations, tutorials, and end-of-chapter problems.

CHAPTER CASES

"MiniDuke" Exposes EU Cybersecurity Gaps

Stuxnet and the Changing Face of Cyberwarfare

MWEB Business: Hacked

Information Security Threats and Policies in Europe

VIDEO CASES

Stuxnet and Cyberwarfare

Cyberespionage: The Chinese Threat

IBM Zone Trusted Information Channel (ZTIC)

Instructional Videos:

Sony PlayStation Hacked; Data Stolen from 77 Million Users

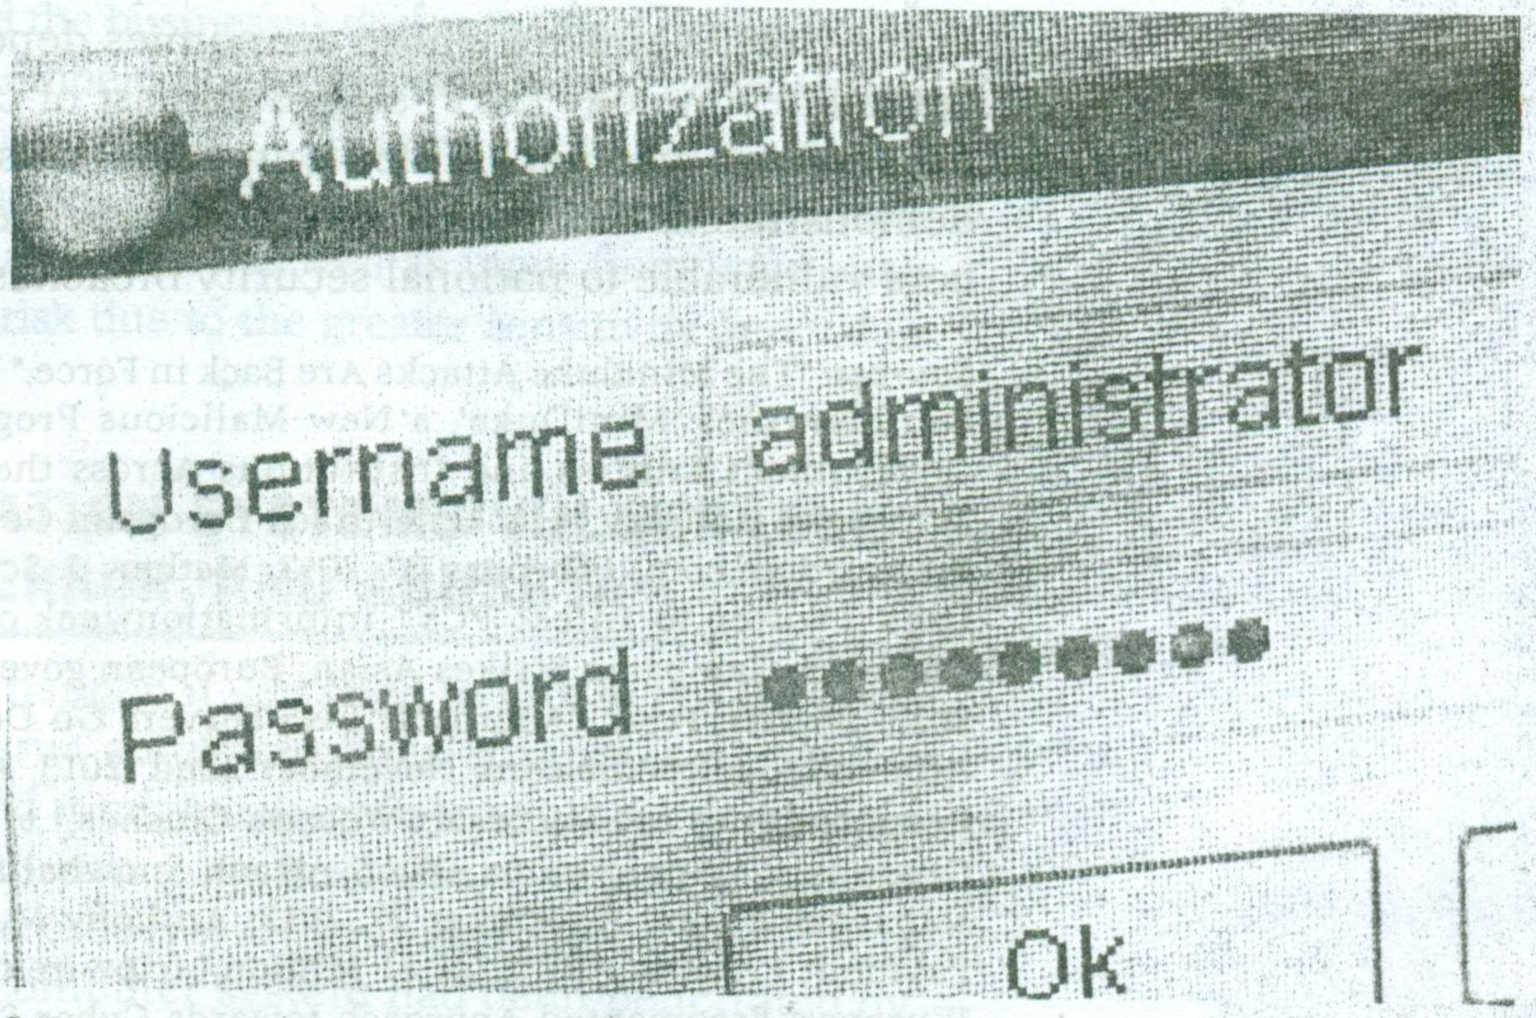
Zappos Working to Correct Online Security Breach

Meet the Hackers: Anonymous Statement on Hacking SONY

'MINIDUKE' EXPOSES EU CYBERSECURITY GAPS

When over 20 European countries, including Ukraine, Belgium, Portugal, Romania, the Czech Republic, and Ireland were infected with an unusual malware agent that combined malware writing techniques from the turn of the century with novel, new sandbox-evading techniques, security experts at Kaspersky Lab and CrySys Lab sounded the alarm. Dubbed MiniDuke, the social engineering exploit gained initial access by inducing targets to open infected PDF attachments that appeared legitimate and highly relevant to the recipient's work.

Sandboxing was developed as an alternative to traditional signature-based virus detection to find zero-day malware—unknown malware for which antivirus software signatures are not yet available. MiniDuke used a now patched Adobe Reader bug to avoid the sandbox, allowing it to install a dropper (small downloader) that provided the perpetrators with a customized backdoor unique to each system. Normal authentication procedures could then be bypassed and affected machines remotely accessed. What's more, the computer's unique fingerprint was then also used to encrypt communications back to the attackers. In hostile environments, MiniDuke goes dormant rather than proceeding to its next step—using Twitter to locate specific tweets from premade accounts created by its Command and Control (C2) operators. These tweets contain tags that identify encrypted URLs for the C2s. The C2s in turn transmit encrypted files containing commands and infected GIF files to create a bigger backdoor on the system. In the absence of Twitter, MiniDuke can use Google Search to locate the C2s.



© ID1974\Shutterstock

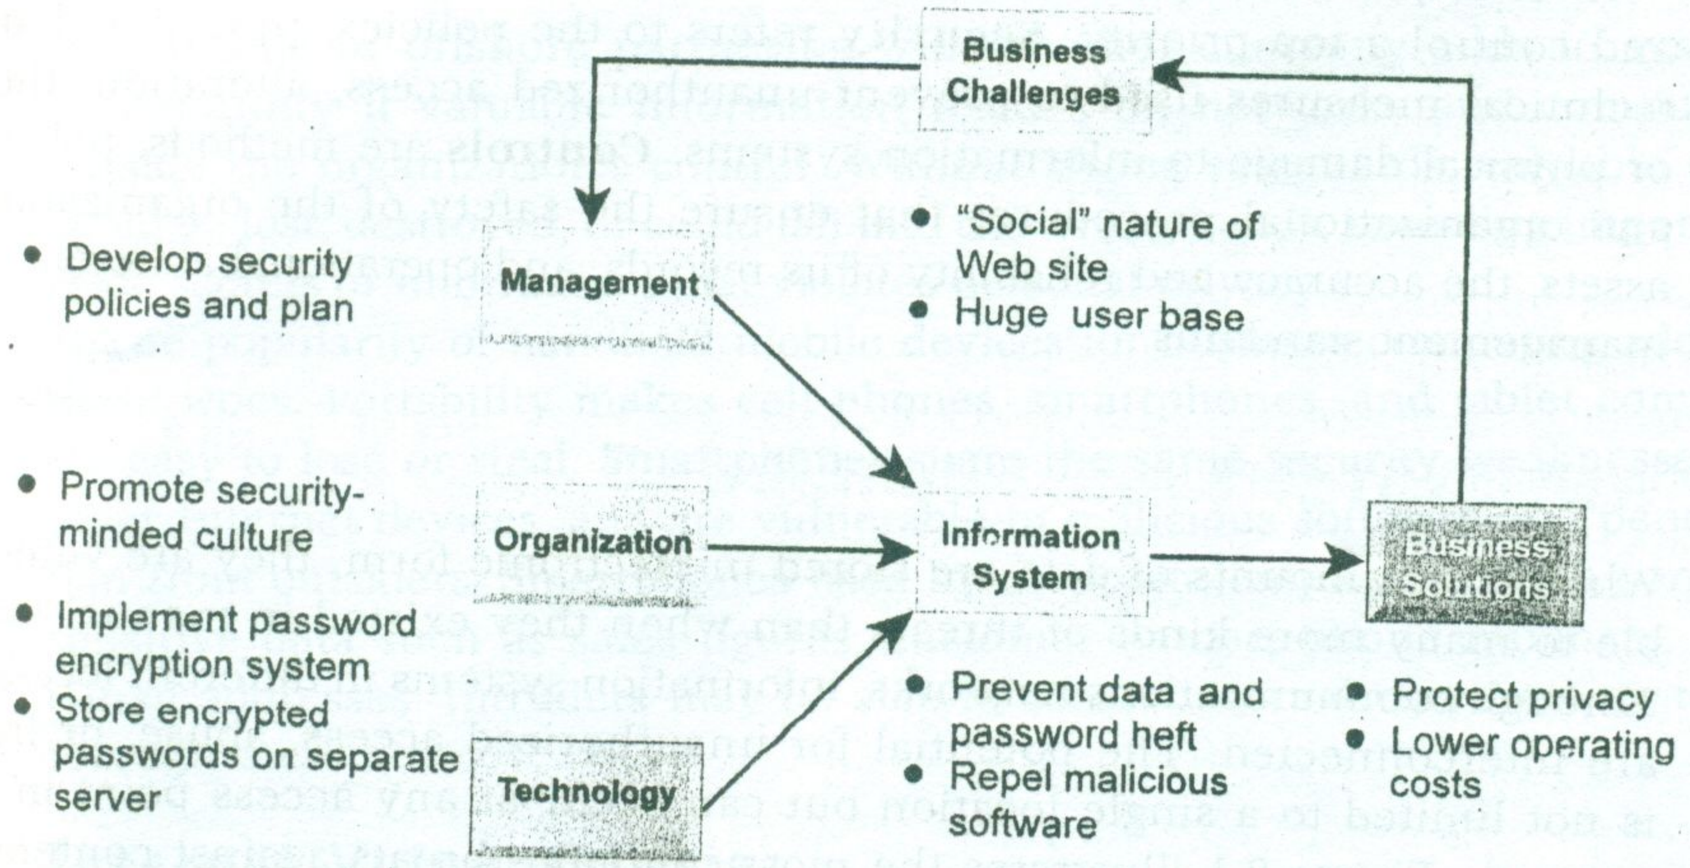
With geopolitical intelligence as the goal, victims of the February 2013 attacks included a research institute, two think tanks, a prominent research foundation, and multiple government entities and institutions. But even though security researchers identified one command server in Panama and another in Turkey, the perpetrators, and the precise information they were seeking, remained unknown.

As 2013 continued, attacks on European government agencies became a recurring theme. In July, Trend Micro discovered that members of the diplomatic community in 16 European and several Asian countries had been targeted with malware-laden e-mail attachments suspected to have originated from China's defense ministry. Like MiniDuke, the malware exploited a vulnerability in widely-used business software (Microsoft Office 2003-2010), installed a dropper, and created a backdoor. It then collected e-mail user names and passwords from Outlook and Internet Explorer along with Web site login credentials and transmitted the data to two now-defunct URLs in Hong Kong. Throughout 2014, MiniDuke continued to exert influence, as law enforcement agencies reported its usage in attacks against governments and other organizations.

With the sophistication, efficacy, frequency, and geopolitical consequences of cyberattacks on the rise, ENISA, the European Union Agency for Network and Information Security, established in 2005, urged implementation of a common cybersecurity strategy. However, the fledgling European Cybercrime Centre (EC3) created in January 2013 was not even able to stipulate a clear definition for cyber security, let alone induce members to craft and implement coordinated policies. In July, its cybersecurity strategy presentation to the European Parliament merely defined the role of governments in combatting cybercrime, bolstering national defense, and effecting international cooperation along with setting general guidelines for fundamental rights, responsibilities, and access.

The 28 members are simply not willing to cede any part of their national security policy to the EU, and the EC3 did not challenge this entrenched outlook. Lack of a centralized and cohesive approach to cybersecurity makes for disjointed and disparate policies and procedures and unequal levels of protection. While members whose economies depend on information and communications technology (ICT) stayed abreast of cybersecurity threats and defenses, others have not yet implemented the most basic protections. As MiniDuke demonstrated, having 28 different approaches to cybersecurity leaves all members vulnerable to national security breaches.

Sources: "The Miniduke Attacks Are Back in Force," Kaspersky.com, July 3, 2014; "Kaspersky Lab Identifies 'MiniDuke', a New Malicious Program Designed for Spying on Multiple Government Entities and Institutions Across the World," Kaspersky.com, February 27, 2013; Josh Halliday, "Hackers Attack European Governments Using 'MiniDuke' Malware," theguardian.com, February 27, 2013; Mathew J. Schwartz, "MiniDuke Espionage Malware Uses Twitter to Infect PCs," informationweek.com, February 28, 2013; Jeremy Kirk, "Malware Campaign Strikes Asian, European governments," computerworld.com, July 16, 2013; Bogdan Popa, "Microsoft DNS Servers Go Down for Some, Windows, Hotmail Sites Affected," softpedia.com, November 22nd, 2013; Anthony M. Freed, "Anonymous Claims Responsibility for Microsoft Website Crashes," tripwire.com, November 27, 2013; Adam Greenberg, "Anonymous DDoS Attack Snowballs, Affects Several Microsoft Services," scmagazine.com, November 27, 2013; Anthony M. Freed, "Hackers Hit European Bitcoin Payment Processor BIPS for $1 Million," tripwire.com, November 26, 2013; Karine e Silva, "Europe's Fragmented Approach towards Cyber Security," policyreview.info, October 10, 2013; Medha Basu, "EU Cyber Security Agency Issues Cloud Deployment Guide," futuregov.asia, November 25, 2013.



The problems created by malware such as MiniDuke and DDoS attacks illustrate why the EU and businesses operating there need to pay special attention to information system security. The fragmented approach to cybersecurity in the EU does not allocate responsibility among the stakeholders, and with no cohesive strategy, members are not encouraged to work together. Neither governments nor businesses are adequately protected.

In November 2013, ENISA, recognizing the switch to cloud computing, urged adoption of a unified strategy for public sector cloud security. Standard procedures for security certification of both services and providers, a shared framework for service-level agreements (SLAs), and a common certification and accreditation process were recommended as well as cloud deployment security measures to protect the integrity of the cloud supply chain.

The chapter-opening diagram calls attention to important points raised by this case and this chapter. While Web servers can never be completely secure from hackers who can steal sensitive public and private sector data, the EU must clearly do more to protect its member nations and the businesses that serve its citizens. While establishment of the EC3 was a start, member nations must be convinced to participate in a consistent, common cybersecurity strategy. This strategy must include the expert knowledge collected by private sector security companies such as Kaspersky Lab McAfee, CrySyS Lab, Trend Micro and others as well as input from large business entities which are at the greatest risk due to the greater sensitivity and value of their stored data.

8.1 WHY ARE INFORMATION SYSTEMS VULNERABLE TO DESTRUCTION, ERROR, AND ABUSE?

Can you imagine what would happen if you tried to link to the Internet without a firewall or antivirus software? Your computer would be disabled in a few seconds, and it might take you many days to recover. If you used the computer to run your business, you might not be able to sell to your customers or place orders with your suppliers while it was down. And you might find that your computer system had been penetrated by outsiders, who perhaps stole or destroyed valuable data, including confidential payment data from your customers. If too much data were destroyed or divulged, your business might never be able to recover!

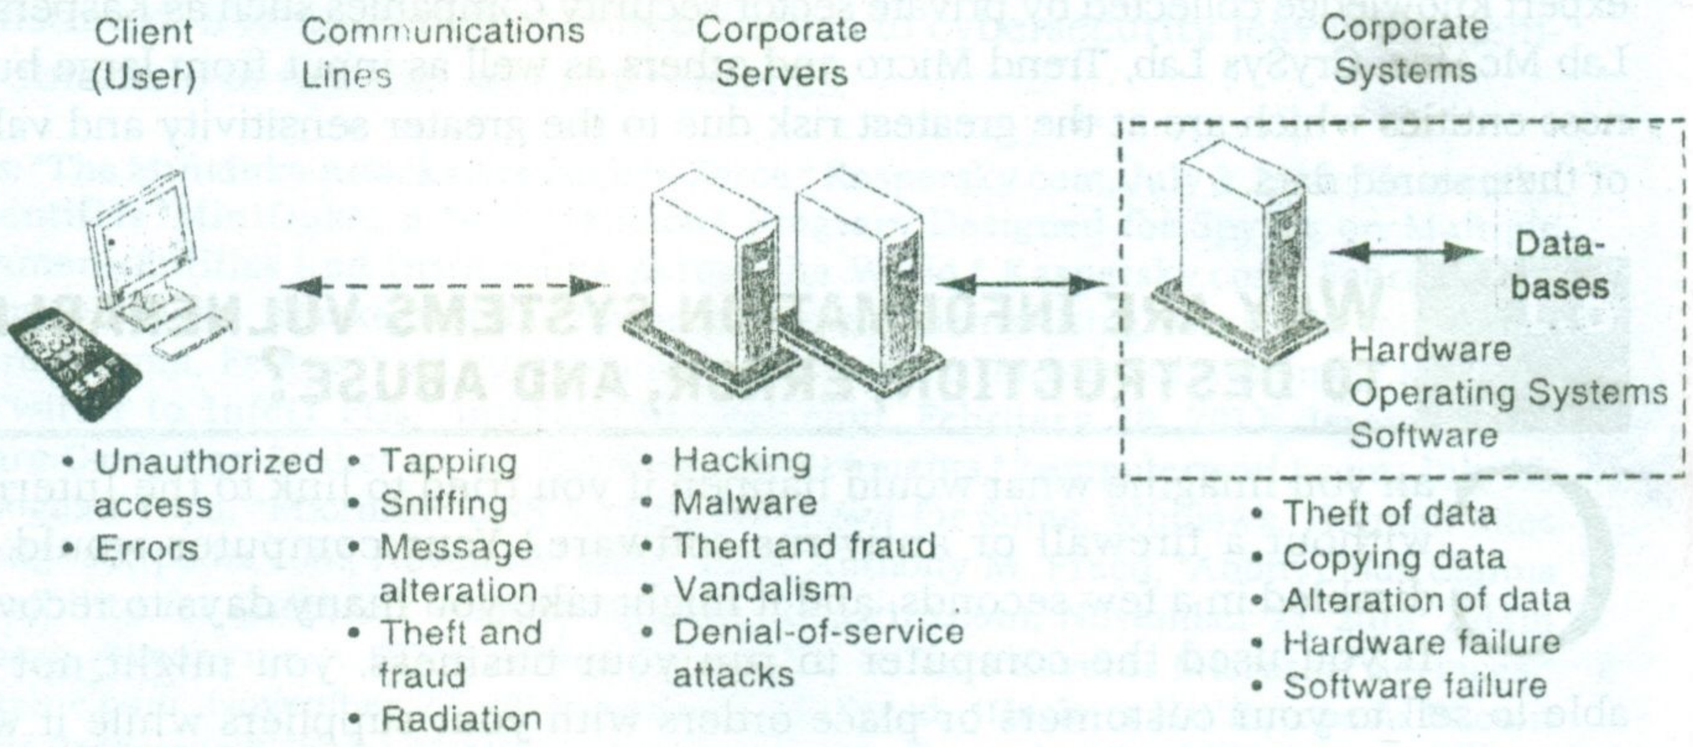
In short, if you operate a business today, you need to make security and control a top priority. **Security** refers to the policies, procedures, and technical measures used to prevent unauthorized access, alteration, theft, or physical damage to information systems. **Controls** are methods, policies, and organizational procedures that ensure the safety of the organization's assets, the accuracy and reliability of its records, and operational adherence to management standards.

WHY SYSTEMS ARE VULNERABLE

When large amounts of data are stored in electronic form, they are vulnerable to many more kinds of threats than when they existed in manual form. Through communications networks, information systems in different locations are interconnected. The potential for unauthorized access, abuse, or fraud is not limited to a single location but can occur at any access point in the network. Figure 8.1 illustrates the most common threats against contemporary information systems. They can stem from technical, organizational, and environmental factors compounded by poor management decisions. In the multi-tier client/server computing environment illustrated here, vulnerabilities exist at each layer and in the communications between the layers. Users at the client layer can cause harm by introducing errors or by accessing systems without authorization. It is possible to access data flowing over networks, steal valuable data during transmission, or alter messages without authorization. Radiation may disrupt a network at various points as well. Intruders can launch denial-of-service attacks or malicious software to disrupt the operation of Web sites. Those capable of penetrating corporate systems can steal, destroy, or alter corporate data stored in databases or files.

Systems malfunction if computer hardware breaks down, is not configured properly, or is damaged by improper use or criminal acts. Errors in programming, improper installation, or unauthorized changes cause computer software

FIGURE 8.1 CONTEMPORARY SECURITY CHALLENGES AND VULNERABILITIES



The architecture of a Web-based application typically includes a Web client, a server, and corporate information systems linked to databases. Each of these components presents security challenges and vulnerabilities. Floods, fires, power failures, and other electrical problems can cause disruptions at any point in the network.

to fail. Power failures, floods, fires, or other natural disasters can also disrupt computer systems.

Domestic or offshore partnering with another company adds to system vulnerability if valuable information resides on networks and computers outside the organization's control. Without strong safeguards, valuable data could be lost, destroyed, or could fall into the wrong hands, revealing important trade secrets or information that violates personal privacy.

The popularity of handheld mobile devices for business computing adds to these woes. Portability makes cell phones, smartphones, and tablet computers easy to lose or steal. Smartphones share the same security weaknesses as other Internet devices, and are vulnerable to malicious software and penetration from outsiders. Smartphones used by corporate employees often contain sensitive data such as sales figures, customer names, phone numbers, and e-mail addresses. Intruders may be able to access internal corporate systems through these devices.

Internet Vulnerabilities

Large public networks, such as the Internet, are more vulnerable than internal networks because they are virtually open to anyone. The Internet is so huge that when abuses do occur, they can have an enormously widespread impact. When the Internet becomes part of the corporate network, the organization's information systems are even more vulnerable to actions from outsiders.

Telephone service based on Internet technology (see Chapter 7) is more vulnerable than the switched voice network if it does not run over a secure private network. Most Voice over IP (VoIP) traffic over the public Internet is not encrypted, so anyone with a network can listen in on conversations. Hackers can intercept conversations or shut down voice service by flooding servers supporting VoIP with bogus traffic.

Vulnerability has also increased from widespread use of e-mail, instant messaging (IM), and peer-to-peer file-sharing programs. E-mail may contain attachments that serve as springboards for malicious software or unauthorized access to internal corporate systems. Employees may use e-mail messages to transmit valuable trade secrets, financial data, or confidential customer information to unauthorized recipients. Popular IM applications for consumers do not use a secure layer for text messages, so they can be intercepted and read by outsiders during transmission over the public Internet. Instant messaging activity over the Internet can in some cases be used as a back door to an otherwise secure network. Sharing files over peer-to-peer (P2P) networks, such as those for illegal music sharing, may also transmit malicious software or expose information on either individual or corporate computers to outsiders.

Wireless Security Challenges

Is it safe to log onto a wireless network at an airport, library, or other public location? It depends on how vigilant you are. Even the wireless network in your home is vulnerable because radio frequency bands are easy to scan. Both Bluetooth and Wi-Fi networks are susceptible to hacking by eavesdroppers. Local area networks (LANs) using the 802.11 standard can be easily penetrated by outsiders armed with laptops, wireless cards, external antennae, and hacking software. Hackers use these tools to detect unprotected networks, monitor network traffic, and, in some cases, gain access to the Internet or to corporate networks.

Wi-Fi transmission technology was designed to make it easy for stations to find and hear one another. The *service set identifiers* (SSIDs) that identify the access points in a Wi-Fi network are broadcast multiple times and can be

picked up fairly easily by intruders' sniffer programs (see Figure 8.2). Wireless networks in many locations do not have basic protections against **war driving**, in which eavesdroppers drive by buildings or park outside and try to intercept wireless network traffic.

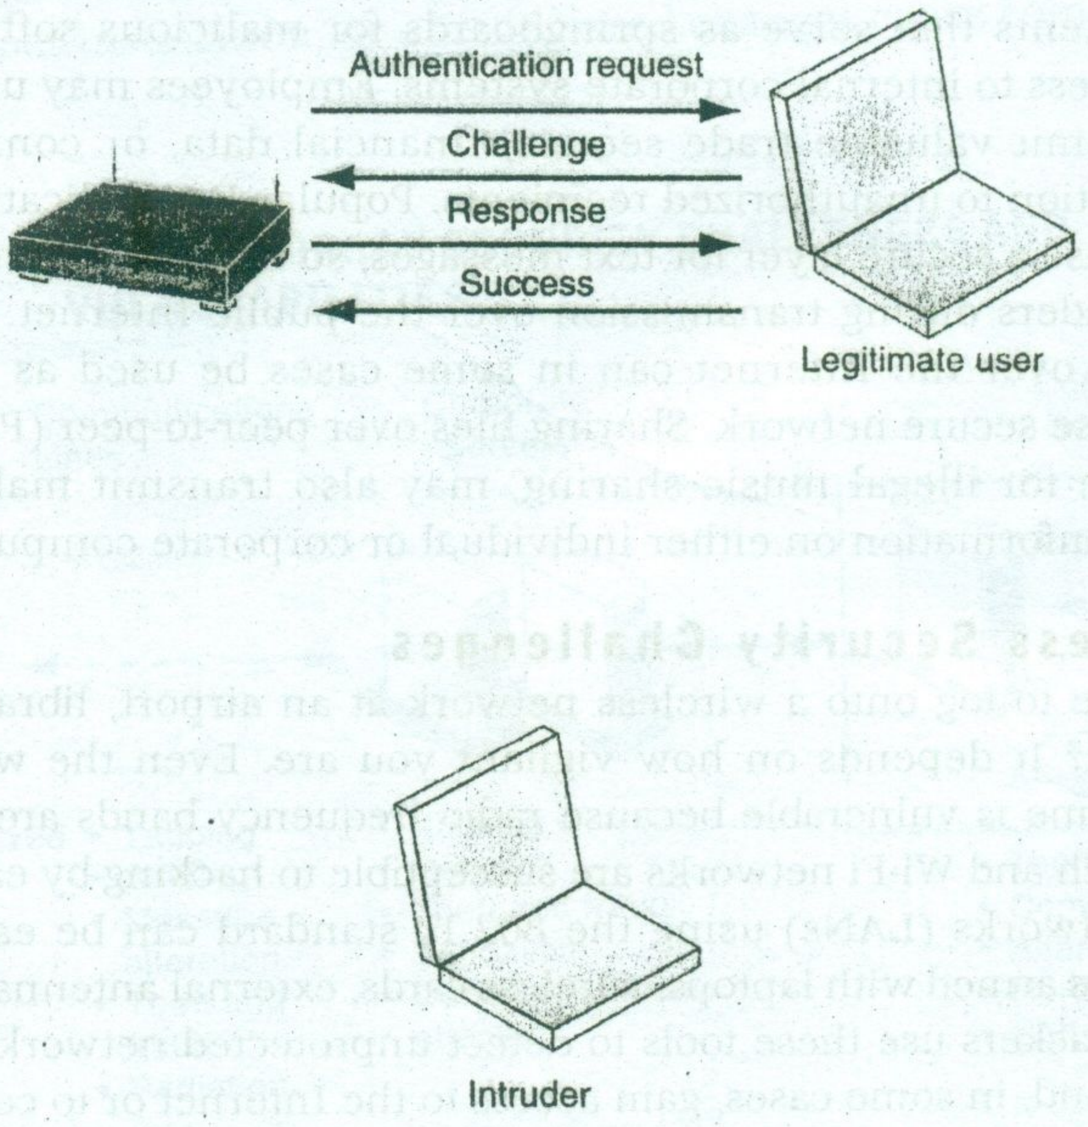
An intruder that has associated with an access point by using the correct SSID is capable of accessing other resources on the network. For example, the intruder could use the Windows operating system to determine which other users are connected to the network, access their computer hard drives, and open or copy their files.

Intruders also use the information they have gleaned to set up rogue access points on a different radio channel in physical locations close to users to force a user's radio network interface controller (NIC) to associate with the rogue access point. Once this association occurs, hackers using the rogue access point can capture the names and passwords of unsuspecting users.

MALICIOUS SOFTWARE: VIRUSES, WORMS, TROJAN HORSES, AND SPYWARE

Malicious software programs are referred to as **malware** and include a variety of threats, such as computer viruses, worms, and Trojan horses. A **computer virus** is a rogue software program that attaches itself to other software programs or data files in order to be executed, usually without user knowledge or permission. Most computer viruses deliver a "payload." The payload may be relatively

FIGURE 8.2 WI-FI SECURITY CHALLENGES



Many Wi-Fi networks can be penetrated easily by intruders using sniffer programs to obtain an address to access the resources of a network without authorization.

benign, such as instructions to display a message or image, or it may be highly destructive—destroying programs or data, clogging computer memory, reformatting a computer's hard drive, or causing programs to run improperly. Viruses typically spread from computer to computer when humans take an action, such as sending an e-mail attachment or copying an infected file.

Most recent attacks have come from **worms**, which are independent computer programs that copy themselves from one computer to other computers over a network. Unlike viruses, worms can operate on their own without attaching to other computer program files and rely less on human behavior in order to spread from computer to computer. This explains why computer worms spread much more rapidly than computer viruses. Worms destroy data and programs as well as disrupt or even halt the operation of computer networks.

Worms and viruses are often spread over the Internet from files of downloaded software, from files attached to e-mail transmissions, or from compromised e-mail messages, online ads, or instant messaging. Viruses have also invaded computerized information systems from "infected" disks or infected machines. Especially prevalent today are **drive-by downloads**, consisting of malware that comes with a downloaded file that a user intentionally or unintentionally requests.

Hackers can do to a smartphone just about anything they can do to any Internet device: request malicious files without user intervention, delete files, transmit files, install programs running in the background to monitor user actions, and potentially convert the smartphone into a robot in a botnet to send e-mail and text messages to anyone. With smartphones starting to outsell PCs, and smartphones increasingly used as payment devices, they are becoming a major avenue for malware.

According to McAfee and other IT security experts, mobile devices now pose the greatest security risks, outpacing those from larger computers. Android, which is the world's leading mobile operating system, not Windows, is the platform targeted by most hackers. (McAfee, 2014; Reisinger, 2014). Mobile device viruses pose serious threats to enterprise computing because so many wireless devices are now linked to corporate information systems, as the trend towards employees using their own devices for work continues.

Blogs, wikis, and social networking sites such as Facebook, Twitter, and LinkedIn have emerged as new conduits for malware or spyware. These applications allow users to post software code as part of the permissible content, and such code can be launched automatically as soon as a Web page is viewed. In the spring of 2014, some programs purporting to be Instagram "image viewers," or "image and video downloaders" to transfer Instagram photos to desktop PCs were bundled with malware. The malware caused Internet slowdown, unwanted redirection to other sites, and sometimes installation of other programs without the user's consent (Scharr, 2014).

Panda Security reported that in 2013, there were 30 million new malware strains in circulation, averaging 82,000 per day. Additionally 20 percent of all malware that has ever existed was created in that one year alone (Panda, 2014). Table 8.1 describes the characteristics of some of the most harmful worms and viruses that have appeared to date.

Over 70 percent of the infections Panda found were Trojan Horses. A **Trojan horse** is a software program that appears to be benign but then does something other than expected. The Trojan horse is not itself a virus because it does not replicate, but it is often a way for viruses or other malicious code to be introduced into a computer system. The term *Trojan horse* is based on the huge wooden horse used by the Greeks to trick the Trojans into opening the gates

TABLE 8.1 **EXAMPLES OF MALICIOUS CODE**

NAME	TYPE	DESCRIPTION
Conficker (aka Downadup, Downup)	Worm	First detected in November 2008 and still prevalent. Uses flaws in Windows software to take over machines and link them into a virtual computer that can be commanded remotely. Had more than 5 million computers worldwide under its control. Difficult to eradicate.
Storm	Worm/ Trojan horse	First identified in January 2007. Spreads via e-mail spam with a fake attachment. Infected up to 10 million computers, causing them to join its zombie network of computers engaged in criminal activity.
Sasser.ftp	Worm	First appeared in May 2004. Spread over the Internet by attacking random IP addresses. Causes computers to continually crash and reboot, and infected computers to search for more victims. Affected millions of computers worldwide, disrupting British Airways flight check-ins, operations of British coast guard stations, Hong Kong hospitals, Taiwan post office branches, and Australia's Westpac Bank. Sasser and its variants caused an estimated $14.8 billion to $18.6 billion in damages worldwide.
MyDoom.A	Worm	First appeared on January 26, 2004. Spreads as an e-mail attachment. Sends e-mail to addresses harvested from infected machines, forging the sender's address. At its peak, this worm lowered global Internet performance by 10 percent and Web page loading times by as much as 50 percent. Was programmed to stop spreading after February 12, 2004.
Sobig.F	Worm	First detected on August 19, 2003. Spreads via e-mail attachments and sends massive amounts of mail with forged sender information. Deactivated itself on September 10, 2003, after infecting more than 1 million PCs and doing $5 to $10 billion in damage.
ILOVEYOU	Virus	First detected on May 3, 2000. Script virus written in Visual Basic script and transmitted as an attachment to e-mail with the subject line ILOVEYOU. Overwrites music, image, and other files with a copy of itself and did an estimated $10 billion to $15 billion in damage.
Melissa	Macro virus/worm	First appeared in March 1999. Word macro script mailing infected Word file to first 50 entries in user's Microsoft Outlook address book. Infected 15 to 29 percent of all business PCs, causing $300 million to $600 million in damage.

to their fortified city during the Trojan War. Once inside the city walls, Greek soldiers hidden in the horse revealed themselves and captured the city.

An example of a modern-day Trojan horse is the Zeus Trojan which runs on computers with the Microsoft Windows operating system. It is often used to steal login credentials for banking by surreptitiously capturing peoples' keystrokes as they use their computers. Zeus is spread mainly through drive-by downloads and phishing.

SQL injection attacks have become a major malware threat. SQL injection attacks take advantage of vulnerabilities in poorly coded Web application software to introduce malicious program code into a company's systems and networks. These vulnerabilities occur when a Web application fails to properly validate or filter data entered by a user on a Web page, which might occur when ordering something online. An attacker uses this input validation error to send a rogue SQL query to the underlying database to access the database, plant malicious code, or access other systems on the network. Large Web applications have hundreds of places for inputting user data, each of which creates an opportunity for an SQL injection attack.

Malware known as **ransomware** is proliferating on both desktop and mobile devices. Ransomware tries to extort money from users by taking control of

their computers or displaying annoying pop-up messages. One nasty example, CryptoLocker, encrypts an infected computer's files, forcing users to pay hundreds of dollars to regain access. You can get ransomware from donwloading an infected attachment, clicking a link inside an email or by visiting the wrong Web site.

Some types of **spyware** also act as malicious software. These small programs install themselves surreptitiously on computers to monitor user Web surfing activity and serve up advertising. Thousands of forms of spyware have been documented.

Many users find such spyware annoying, and some critics worry about its infringement on computer users' privacy. Some forms of spyware are especially nefarious. **Keyloggers** record every keystroke made on a computer to steal serial numbers for software, to launch Internet attacks, to gain access to e-mail accounts, to obtain passwords to protected computer systems, or to pick up personal information such as credit card and or bank account numbers. The Zeus Trojan described earlier uses keylogging. Other spyware programs reset Web browser home pages, redirect search requests, or slow performance by taking up too much memory.

HACKERS AND COMPUTER CRIME

A **hacker** is an individual who intends to gain unauthorized access to a computer system. Within the hacking community, the term *cracker* is typically used to denote a hacker with criminal intent, although in the public press, the terms hacker and cracker are used interchangeably. Hackers and crackers gain unauthorized access by finding weaknesses in the security protections employed by Web sites and computer systems, often taking advantage of various features of the Internet that make it an open system and easy to use.

Hacker activities have broadened beyond mere system intrusion to include theft of goods and information, as well as system damage and **cybervandalism**, the intentional disruption, defacement, or even destruction of a Web site or corporate information system. For example, a group of pro-Syrian regime hackers called the Syrian Electronic Army hacked Skype's Twitter account, blog, and Facebook page in early January 2014, publishing a fake message which read, "Don't use Microsoft emails...They are monitoring your accounts and selling the data to the governments" (Ribeiro, 2014).

Spoofing and Sniffing

Hackers attempting to hide their true identities often spoof, or misrepresent, themselves by using fake e-mail addresses or masquerading as someone else. **Spoofing** may also involve redirecting a Web link to an address different from the intended one, with the site masquerading as the intended destination. For example, if hackers redirect customers to a fake Web site that looks almost exactly like the true site, they can then collect and process orders, effectively stealing business as well as sensitive customer information from the true site. We will provide more detail on other forms of spoofing in our discussion of computer crime.

A **sniffer** is a type of eavesdropping program that monitors information traveling over a network. When used legitimately, sniffers help identify potential network trouble spots or criminal activity on networks, but when used for criminal purposes, they can be damaging and very difficult to detect. Sniffers enable hackers to steal proprietary information from anywhere on a network, including e-mail messages, company files, and confidential reports.

Denial-of-Service Attacks

In a **denial-of-service (DoS) attack**, hackers flood a network server or Web server with many thousands of false communications or requests for services to crash the network. The network receives so many queries that it cannot keep up with them and is thus unavailable to service legitimate requests. A **distributed denial-of-service (DDoS)** attack uses numerous computers to inundate and overwhelm the network from numerous launch points.

Although DoS attacks do not destroy information or access restricted areas of a company's information systems, they often cause a Web site to shut down, making it impossible for legitimate users to access the site. For example, the blogging service Typepad was subjected to a series of DDoS attacks in April 2014 which lasted five days. During that time period, the site was flooded with so much illicit traffic that it was knocked offline, preventing Typepad users from acessing their blogs and applications (Perez, 2014). For busy e-commerce sites, these attacks are costly; while the site is shut down, customers cannot make purchases. Especially vulnerable are small and midsize businesses whose networks tend to be less protected than those of large corporations.

Perpetrators of DDoS attacks often use thousands of "zombie" PCs infected with malicious software without their owners' knowledge and organized into a **botnet**. Hackers create these botnets by infecting other people's computers with bot malware that opens a back door through which an attacker can give instructions. The infected computer then becomes a slave, or zombie, serving a master computer belonging to someone else. Once hackers infect enough computers, they can use the amassed resources of the botnet to launch DDoS attacks, phishing campaigns, or unsolicited "spam" e-mail.

Ninety percent of the world's spam and 80 percent of the world's malware are delivered via botnets. For example, the Grum botnet, once the world's third-largest botnet, was reportedly responsible for 18% of worldwide spam traffic (amounting to 18 billion spam messages per day) when it was shut down on July 19, 2012. At one point Grum had infected and controlled 560,000–840,000 computers.

Computer Crime

Most hacker activities are criminal offenses, and the vulnerabilities of systems we have just described make them targets for other types of **computer crime** as well. Computer crime is defined by the U.S. Department of Justice as "any violations of criminal law that involve a knowledge of computer technology for their perpetration, investigation, or prosecution." Table 8.2 provides examples of the computer as both a target and an instrument of crime.

No one knows the magnitude of the computer crime problem—how many systems are invaded, how many people engage in the practice, or the total economic damage. According to the Ponemon Institute's 2013 Annual Cost of Cyber Crime Study sponsored by HP Enterprise Security, the average annualized cost of cybercrime for the organizations in the study was $11.6 million per year (Ponemon Institute, 2013). Many companies are reluctant to report computer crimes because the crimes may involve employees, or the company fears that publicizing its vulnerability will hurt its reputation. The most economically damaging kinds of computer crime are denial of service attacks, activities of malicious insiders, and Web-based attacks.

Identity Theft

With the growth of the Internet and electronic commerce, identity theft has become especially troubling. **Identity theft** is a crime in which an imposter

TABLE 8.2 EXAMPLES OF COMPUTER CRIME

COMPUTERS AS TARGETS OF CRIME
Breaching the confidentiality of protected computerized data
Accessing a computer system without authority
Knowingly accessing a protected computer to commit fraud
Intentionally accessing a protected computer and causing damage, negligently or deliberately
Knowingly transmitting a program, program code, or command that intentionally causes damage to a protected computer
Threatening to cause damage to a protected computer
COMPUTERS AS INSTRUMENTS OF CRIME
Theft of trade secrets
Unauthorized copying of software or copyrighted intellectual property, such as articles, books, music, and video
Schemes to defraud
Using e-mail for threats or harassment
Intentionally attempting to intercept electronic communication
Illegally accessing stored electronic communications, including e-mail and voice mail
Transmitting or possessing child pornography using a computer

obtains key pieces of personal information, such as social security identification numbers, driver's license numbers, or credit card numbers, to impersonate someone else. The information may be used to obtain credit, merchandise, or services in the name of the victim or to provide the thief with false credentials.

Identity theft has flourished on the Internet, with credit card files a major target of Web site hackers. According to the Identity Fraud Study by Javelin Strategy & Research, identity fraud affected 13.1 million consumers in 2013. The total dolllar losses from identity theft increased to $18 billion (Javelin 2014). Moreover, e-commerce sites are wonderful sources of customer personal information—name, address, and phone number. Armed with this information, criminals are able to assume new identities and establish new credit for their own purposes.

One increasingly popular tactic is a form of spoofing called **phishing**. Phishing involves setting up fake Web sites or sending e-mail messages that look like those of legitimate businesses to ask users for confidential personal data. The e-mail message instructs recipients to update or confirm records by providing social security numbers, bank and credit card information, and other confidential data either by responding to the e-mail message, by entering the information at a bogus Web site, or by calling a telephone number. EBay, PayPal, Amazon.com, Walmart, and a variety of banks are among the top spoofed companies. In a more targeted form of phishing called *spear phishing*, messages appear to come from a trusted source, such as an individual within the recipient's own company or a friend.

Phishing techniques called evil twins and pharming are harder to detect. **Evil twins** are wireless networks that pretend to offer trustworthy Wi-Fi connections to the Internet, such as those in airport lounges, hotels, or coffee shops. The bogus network looks identical to a legitimate public network. Fraudsters try to capture passwords or credit card numbers of unwitting users who log on to the network.

Pharming redirects users to a bogus Web page, even when the individual types the correct Web page address into his or her browser. This is possible if pharming perpetrators gain access to the Internet address information stored by Internet service providers to speed up Web browsing and the ISP companies have flawed software on their servers that allows the fraudsters to hack in and change those addresses.

According to the Ponemon Institute's 2014 Cost of a Data Breach Study, the average cost of a breach to a company was $3.5 million (Ponemon, 2014). Moreover, brand damage can be significant, albeit hard to quantify. Table 8.3 describes other major data breaches.

The U.S. Congress addressed the threat of computer crime in 1986 with the Computer Fraud and Abuse Act, which makes it illegal to access a computer system without authorization. Most states have similar laws, and nations in Europe have comparable legislation. Congress passed the National Information Infrastructure Protection Act in 1996 to make malware distribution and hacker attacks to disable Web sites federal crimes.

U.S. legislation, such as the Wiretap Act, Wire Fraud Act, Economic Espionage Act, Electronic Communications Privacy Act, E-Mail Threats and Harassment Act, and Child Pornography Act, covers computer crimes involving intercepting electronic communication, using electronic communication to defraud, stealing trade secrets, illegally accessing stored electronic communications, using e-mail for threats or harassment, and transmitting or possessing child pornography. A proposed federal Data Security and Breach Notification Act would mandate organizations that possess personal information to put in place "reasonable" security procedures to keep the data secure and to notify anyone affected by a data breach, but it has not been enacted.

TABLE 8.3 **MAJOR DATA BREACHES**

DATA BREACH	DESCRIPTION
EBay	Cyberattack on eBay servers during February and March 2014 compromises database containing customer names, encrypted passwords, email addresses, physical addresses, phone numbers, and birthdates. No financial data were accessed, but the information is useful for identity theft.
Heartland Payment Systems	In 2008, criminals led by Miami hacker Albert Gonzales installed spying software on the computer network of Heartland Payment Systems, a payment processor based in Princeton, NJ, and stole the numbers of as many as 100 million credit and debit cards. Gonzales was sentenced in 2010 to 20 years in federal prison, and Heartland paid about $140 million in fines and settlements.
TJX	A 2007 data breach at TJX, the retailer that owns national chains including TJ Maxx and Marshalls, cost at least $250 million. Cyber criminals took more than 45 million credit and debit card numbers, some of which were used later to buy millions of dollars in electronics from Walmart and elsewhere. Albert Gonzales, who played a major role in the Heartland hack, was linked to this cyberattack as well.
Epsilon	In March 2011, hackers stole millions of names and e-mail addresses from the Epsilon e-mail marketing firm, which handles e-mail lists for major retailers and banks like Best Buy, JPMorgan, TiVo, and Walgreens. Costs could range from $100 million to $4 billion, depending on what happens to the stolen data, with most of the costs from losing customers due to a damaged reputation.
Sony	In April 2011, hackers obtained personal information, including credit, debit, and bank account numbers, from over 100 million PlayStation Network users and Sony Online Entertainment users. The breach could cost Sony and credit card issuers up to a total of $2 billion.

INTERACTIVE SESSION: MANAGEMENT

STUXNET AND THE CHANGING FACE OF CYBERWARFARE

In July 2010, reports surfaced about a Stuxnet worm that had been targeting Iran's nuclear facilities. In November of that year, Iran's President Mahmoud Ahmadinejad publicly acknowledged that malicious software had infected the Iranian nuclear facilities and disrupted the nuclear program by disabling the facilities' centrifuges. Stuxnet had earned its place in history as the first visible example of industrial cyberwarfare.

To date, Stuxnet is the most sophisticated cyberweapon ever deployed. Stuxnet's mission was to activate only computers that ran Supervisory Control and Data Acquisition (SCADA) software used in Siemens centrifuges to enrich uranium. The Windows-based worm had a "dual warhead." One part was designed to lay dormant for long periods, then speed up Iran's nuclear centrifuges so that they spun wildly out of control. Another secretly recorded what normal operations at the nuclear plant looked like and then played those recordings back to plant operators so it would appear that the centrifuges were operating normally when they were actually tearing themselves apart.

The worm's sophistication indicated the work of highly skilled professionals. Michael Assante, president and CEO at the National Board of Information Security Examiners, views Stuxnet as a weapons delivery system like the B-2 Bomber. The software program code was highly modular, so that it could be easily changed to attack different systems. Stuxnet only became active when it encountered a specific configuration of controllers, running a set of processes limited to centrifuge plants.

Over 60 percent of Stuxet-infected computers are in Iran, and digital security company Kaspersky Labs speculates that the worm was launched with nation-state support (probably from Israel and the United States) with the intention of disabling some or all of Iran's uranium enrichment program. Stuxnet wiped out about one-fifth of Iran's nuclear centrifuges by causing them to spin at too high a velocity. The damage was irreparable and is believed to have delayed Iran's ability to make nuclear arms by as much as five years. And no one is certain that the Stuxnet attacks are over. Some experts who examined the Stuxnet software code believe it contains the seeds for more versions and attacks.

According to a Tofino Security report, Stuxnet is capable of infecting even well-secured computer systems that follow industry best practices. Companies' need for interonnectivity between control systems make it nearly impossible to defend against a well-constructed, multi-pronged attack such as Stuxnet.

And Stuxnet is not the only cyberweapon currently at work. The Flame virus, released about five years ago, has been infecting computers in Iran, Lebanon, Sudan, Saudi Arabia, Egypt, Syria, and Israel. While researchers are still analyzing the program, the attack's main goal is stealing information and espionage. Flame is able to grab images of users' computer screens, record their instant messaging chats, collect passwords, remotely turn on their microphones to record audio conversations, scan disks for specific files, and monitor their keystrokes and network traffic. The software also records Skype conversations and can turn infected computers into Bluetooth beacons which attempt to download contact information from nearby Bluetooth-enabled devices. These data, along with locally stored documents, can be sent to one of several command and control servers that are scattered around the world. The program then awaits further instructions from these servers.

The Duqu worm, discovered in September 2011, also aims to steal information by scanning systems. Duqu infects a very small number of very specific systems around the world, but may use completely different modules for infiltrating those separate systems. One of Duqu's actions is to steal digital certificates used for authentication from attacked computers to help future viruses appear as secure software. It is going largely undetected. Security researchers believe Duqu was created by the same group of programmers behind Stuxnet.

The real worry for security experts and government officials is an act of cyberwarfare against a critical resource, such as the electric grid, financial systems, or communications systems. (In April 2009, cyberspies infiltrated the U.S. electrical grid, using weak points where computers on the grid are connected to the Internet, and left behind software programs whose purpose is unclear, but which presumably could be used to disrupt the system.)

The U.S. has no clear strategy about how the country would respond to that level of cyberattack, and the

effects of such an attack would likely be devastating. Mike McConnell, the former director of national intelligence, stated that if even a single large American bank were successfully attacked, it would have an order-of-magnitude greater impact on the global economy than the World Trade Center attacks, and that the ability to threaten the U.S. money supply is the financial equivalent of a nuclear weapon.

Many security experts believe that U.S. cybersecurity is not well-organized. Several different agencies, including the Pentagon and the National Security Agency (NSA), have their sights on being the leading agency in the ongoing efforts to combat cyberwarfare. The first headquarters designed to coordinate government cybersecurity efforts, called Cybercom, was activated in May 2010 in the hope of resolving this organizational tangle. In May 2011 President Barack Obama signed executive orders weaving cyber capabilities into U.S. military strategy, but these capabilities are still evolving.

In 2014, a virus similar to Stuxnet called Energetic Bear was found to have attacked energy companies in the U.S. and Europe, lending credence to fears the the energy grid is vulnerable to these kinds of attacks. It's one thing to develop a next-generation computer virus, but another one to develop methods of defending established computer systems from them. Will the United States and other nations be ready when the next Stuxnet appears?

Sources: Michael B. Kelley, "A Stuxnet-Like Virus Has Infected Hundreds of US and European Energy Companies," Businessinsider.com, July 1, 2014; Brian Royer, "Stuxnet, the Nation's Power Grid, and the Law of Unintended Consequences," *Dark Reading, March 12, 2012*; Thomas Erdbrink, "Iran Confirms Attack by Virus That Collects Information," *New York Times*, May 29, 2012; Nicole Perlroth, "Virus Infects Computers across Middle East," *New York Times*, May 28, 2012; Thom Shanker and Elisabeth Bumiller, "After Suffering Damaging Cyberattack, the Pentagon Takes Defensive Action," *New York Times*, July 15, 2011; Robert Leos, "Secure Best Practices No Proof Against Stuxnet," CSO, March 3, 2011; Lolita C. Baldor, "Pentagon Gets Cyberwar Guidelines," Associated Press, June 22, 2011; William J. Broad, John Markoff, and David E. Sanger, "Israel Tests on Worm Called Crucial in Iran Nuclear Delay," *New York Times*, January 15, 2011; George V. Hulme, "SCADA Insecurity" and Michael S. Mimoso, "Cyberspace Has Gone Offensive," *Information Security's Essential Guide to Threat Management* (June 14, 2011); and Sibhan Gorman and Julian A. Barnes, "Cyber Combat: Act of War," *Wall Street Journal*, May 31, 2011.

CASE STUDY QUESTIONS

1. Is cyberwarfare a serious problem? Why or why not?
2. Assess the management, organization, and technology factors that have created this problem.
3. What makes Stuxnet different from other cyberwarfare attacks? How serious a threat is this technology?
4. What solutions have been proposed for this problem? Do you think they will be effective? Why or why not?

Click Fraud

When you click on an ad displayed by a search engine, the advertiser typically pays a fee for each click, which is supposed to direct potential buyers to its products. **Click fraud** occurs when an individual or computer program fraudulently clicks on an online ad without any intention of learning more about the advertiser or making a purchase. Click fraud has become a serious problem at Google and other Web sites that feature pay-per-click online advertising.

Some companies hire third parties (typically from low-wage countries) to fraudulently click on a competitor's ads to weaken them by driving up their marketing costs. Click fraud can also be perpetrated with software programs doing the clicking, and botnets are often used for this purpose. Search engines such as Google attempt to monitor click fraud but have been reluctant to publicize their efforts to deal with the problem.

Global Threats: Cyberterrorism and Cyberwarfare

The cyber criminal activities we have described—launching malware, denial-of-service attacks, and phishing probes—are borderless. Attack servers for malware are now hosted in 206 countries and territories, according to the latest FireEye Advanced Threat Report. The most popular sources of malware attacks are the United States, Germany, South Korea, China, Netherlands, United Kingdom, and Russia (Karlovsky, 2014). The global nature of the Internet makes it possible for cybercriminals to operate—and to do harm—anywhere in the world.

Internet vulnerabilities have also turned individuals and even entire nation states into easy targets for politically-motivated hacking to conduct sabotage and espionage. **Cyberwarfare** is a state-sponsored activity designed to cripple and defeat another state or nation by penetrating its computers or networks for the purposes of causing damage and disruption.

In general, cyberwarfare attacks have become much more widespread, sophisticated, and potentially devastating. There are 250,000 probes trying to find their way into the U.S. Department of Defense networks every hour, and cyberattacks on U.S. federal agencies have increased 150 percent since 2008. Over the years, hackers have stolen plans for missile tracking systems, satellite navigation devices, surveillance drones, and leading-edge jet fighters.

Cyberwarfare poses a serious threat to the infrastructure of modern societies, since their major financial, health, government, and industrial institutions rely on the Internet for daily operations. Cyberwarfare also involves defending against these types of attacks. The chapter-ending case discusses this topic in greater detail.

INTERNAL THREATS: EMPLOYEES

We tend to think the security threats to a business originate outside the organization. In fact, company insiders pose serious security problems. Employees have access to privileged information, and in the presence of sloppy internal security procedures, they are often able to roam throughout an organization's systems without leaving a trace.

Studies have found that user lack of knowledge is the single greatest cause of network security breaches. Many employees forget their passwords to access computer systems or allow co-workers to use them, which compromises the system. Malicious intruders seeking system access sometimes trick employees into revealing their passwords by pretending to be legitimate members of the company in need of information. This practice is called **social engineering**.

Both end users and information systems specialists are also a major source of errors introduced into information systems. End users introduce errors by entering faulty data or by not following the proper instructions for processing data and using computer equipment. Information systems specialists may create software errors as they design and develop new software or maintain existing programs.

SOFTWARE VULNERABILITY

Software errors pose a constant threat to information systems, causing untold losses in productivity, and sometimes endangering people who use or depend on systems. Growing complexity and size of software programs, coupled with demands for timely delivery to markets, have contributed to an increase in software flaws or vulnerabilities. On April 16, 2013, American Airlines had to cancel or delay 1,950 flights due to a faulty software patch from a vendor or

internal software changes that were not properly tested. A systemwide network outage downed the airline's primary systems that manage airline operations, as well as backup systems that kick in when the primary systems fail (Boulton, 2013).

A major problem with software is the presence of hidden **bugs** or program code defects. Studies have shown that it is virtually impossible to eliminate all bugs from large programs. The main source of bugs is the complexity of decision-making code. A relatively small program of several hundred lines will contain tens of decisions leading to hundreds or even thousands of different paths. Important programs within most corporations are usually much larger, containing tens of thousands or even millions of lines of code, each with many times the choices and paths of the smaller programs.

Zero defects cannot be achieved in larger programs. Complete testing simply is not possible. Fully testing programs that contain thousands of choices and millions of paths would require thousands of years. Even with rigorous testing, you would not know for sure that a piece of software was dependable until the product proved itself after much operational use.

Flaws in commercial software not only impede performance but also create security vulnerabilities that open networks to intruders. Each year security firms identify thousands of software vulnerabilities in Internet and PC software. A recent example is the Heartbleed bug, which is a flaw in OpenSSL, an open-source encryption technology that is used by an estimated two-thirds of Web servers. Hackers could exploit the bug to access visitors' personal data as well as a site's encryption keys, which can be used to collect even more "protected" data.

To correct software flaws once they are identified, the software vendor creates small pieces of software called **patches** to repair the flaws without disturbing the proper operation of the software. An example is Microsoft's Windows 7 Service Pack 1, which features security, performance, and stability updates for Windows 7. It is up to users of the software to track these vulnerabilities, test, and apply all patches. This process is called *patch management*.

Because a company's IT infrastructure is typically laden with multiple business applications, operating system installations, and other system services, maintaining patches on all devices and services used by a company is often time-consuming and costly. Malware is being created so rapidly that companies have very little time to respond between the time a vulnerability and a patch are announced and the time malicious software appears to exploit the vulnerability.

8.2 What is the business value of security and control?

Many firms are reluctant to spend heavily on security because it is not directly related to sales revenue. However, protecting information systems is so critical to the operation of the business that it deserves a second look.

Companies have very valuable information assets to protect. Systems often house confidential information about individuals' taxes, financial assets, medical records, and job performance reviews. They also can contain information on corporate operations, including trade secrets, new product development plans, and marketing strategies. Government systems may store information on weapons systems, intelligence operations, and military targets. These information assets have tremendous value, and the repercussions can be devastating if they

are lost, destroyed, or placed in the wrong hands. Systems that are unable to function because of security breaches, disasters, or malfunctioning technology can permanently impact a company's financial health. Some experts believe that 40 percent of all businesses will not recover from application or data losses that are not repaired within three days.

Inadequate security and control may result in serious legal liability. Businesses must protect not only their own information assets but also those of customers, employees, and business partners. Failure to do so may open the firm to costly litigation for data exposure or theft. An organization can be held liable for needless risk and harm created if the organization fails to take appropriate protective action to prevent loss of confidential information, data corruption, or breach of privacy. For example, BJ's Wholesale Club was sued by the U.S. Federal Trade Commission for allowing hackers to access its systems and steal credit and debit card data for fraudulent purchases. Banks that issued the cards with the stolen data sought $13 million from BJ's to compensate them for reimbursing card holders for the fraudulent purchases. A sound security and control framework that protects business information assets can thus produce a high return on investment. Strong security and control also increase employee productivity and lower operational costs.

LEGAL AND REGULATORY REQUIREMENTS FOR ELECTRONIC RECORDS MANAGEMENT

Recent U.S. government regulations are forcing companies to take security and control more seriously by mandating the protection of data from abuse, exposure, and unauthorized access. Firms face new legal obligations for the retention and storage of electronic records as well as for privacy protection.

If you work in the health care industry, your firm will need to comply with the Health Insurance Portability and Accountability Act (HIPAA) of 1996. **HIPAA** outlines medical security and privacy rules and procedures for simplifying the administration of health care billing and automating the transfer of health care data between health care providers, payers, and plans. It requires members of the health care industry to retain patient information for six years and ensure the confidentiality of those records. It specifies privacy, security, and electronic transaction standards for health care providers handling patient information, providing penalties for breaches of medical privacy, disclosure of patient records by e-mail, or unauthorized network access.

If you work in a firm providing financial services, your firm will need to comply with the Financial Services Modernization Act of 1999, better known as the **Gramm-Leach-Bliley Act** after its congressional sponsors. This act requires financial institutions to ensure the security and confidentiality of customer data. Data must be stored on a secure medium, and special security measures must be enforced to protect such data on storage media and during transmittal.

If you work in a publicly traded company, your company will need to comply with the Public Company Accounting Reform and Investor Protection Act of 2002, better known as the **Sarbanes-Oxley Act** after its sponsors Senator Paul Sarbanes of Maryland and Representative Michael Oxley of Ohio. This Act was designed to protect investors after the financial scandals at Enron, WorldCom, and other public companies. It imposes responsibility on companies and their management to safeguard the accuracy and integrity of financial information that is used internally and released externally. One of the Learning Tracks for this chapter discusses Sarbanes-Oxley in detail.

Sarbanes-Oxley is fundamentally about ensuring that internal controls are in place to govern the creation and documentation of information in financial statements. Because information systems are used to generate, store, and transport such data, the legislation requires firms to consider information systems security and other controls required to ensure the integrity, confidentiality, and accuracy of their data. Each system application that deals with critical financial reporting data requires controls to make sure the data are accurate. Controls to secure the corporate network, prevent unauthorized access to systems and data, and ensure data integrity and availability in the event of disaster or other disruption of service are essential as well.

ELECTRONIC EVIDENCE AND COMPUTER FORENSICS

Security, control, and electronic records management have become essential for responding to legal actions. Much of the evidence today for stock fraud, embezzlement, theft of company trade secrets, computer crime, and many civil cases is in digital form. In addition to information from printed or typewritten pages, legal cases today increasingly rely on evidence represented as digital data stored on portable storage devices, CDs, and computer hard disk drives, as well as in e-mail, instant messages, and e-commerce transactions over the Internet. E-mail is currently the most common type of electronic evidence.

In a legal action, a firm is obligated to respond to a discovery request for access to information that may be used as evidence, and the company is required by law to produce those data. The cost of responding to a discovery request can be enormous if the company has trouble assembling the required data or the data have been corrupted or destroyed. Courts now impose severe financial and even criminal penalties for improper destruction of electronic documents.

An effective electronic document retention policy ensures that electronic documents, e-mail, and other records are well organized, accessible, and neither retained too long nor discarded too soon. It also reflects an awareness of how to preserve potential evidence for computer forensics. **Computer forensics** is the scientific collection, examination, authentication, preservation, and analysis of data held on or retrieved from computer storage media in such a way that the information can be used as evidence in a court of law. It deals with the following problems:

- Recovering data from computers while preserving evidential integrity
- Securely storing and handling recovered electronic data
- Finding significant information in a large volume of electronic data
- Presenting the information to a court of law

Electronic evidence may reside on computer storage media in the form of computer files and as *ambient data*, which are not visible to the average user. An example might be a file that has been deleted on a PC hard drive. Data that a computer user may have deleted on computer storage media can be recovered through various techniques. Computer forensics experts try to recover such hidden data for presentation as evidence.

An awareness of computer forensics should be incorporated into a firm's contingency planning process. The CIO, security specialists, information systems staff, and corporate legal counsel should all work together to have a plan in place that can be executed if a legal need arises. You can find out more about computer forensics in the Learning Tracks for this chapter.

8.3 WHAT ARE THE COMPONENTS OF AN ORGANIZATIONAL FRAMEWORK FOR SECURITY AND CONTROL?

Even with the best security tools, your information systems won't be reliable and secure unless you know how and where to deploy them. You'll need to know where your company is at risk and what controls you must have in place to protect your information systems. You'll also need to develop a security policy and plans for keeping your business running if your information systems aren't operational.

INFORMATION SYSTEMS CONTROLS

Information systems controls are both manual and automated and consist of general and application controls. **General controls** govern the design, security, and use of computer programs and the security of data files in general throughout the organization's information technology infrastructure. On the whole, general controls apply to all computerized applications and consist of a combination of hardware, software, and manual procedures that create an overall control environment.

General controls include software controls, physical hardware controls, computer operations controls, data security controls, controls over implementation of system processes, and administrative controls. Table 8.4 describes the functions of each of these controls.

Application controls are specific controls unique to each computerized application, such as payroll or order processing. They include both automated and manual procedures that ensure that only authorized data are completely and accurately processed by that application. Application controls can be classified as (1) input controls, (2) processing controls, and (3) output controls.

TABLE 8.4 GENERAL CONTROLS

TYPE OF GENERAL CONTROL	DESCRIPTION
Software controls	Monitor the use of system software and prevent unauthorized access of software programs, system software, and computer programs.
Hardware controls	Ensure that computer hardware is physically secure, and check for equipment malfunction. Organizations that are critically dependent on their computers also must make provisions for backup or continued operation to maintain constant service.
Computer operations controls	Oversee the work of the computer department to ensure that programmed procedures are consistently and correctly applied to the storage and processing of data. They include controls over the setup of computer processing jobs and backup and recovery procedures for processing that ends abnormally.
Data security controls	Ensure that valuable business data files on either disk or tape are not subject to unauthorized access, change, or destruction while they are in use or in storage.
Implementation controls	Audit the systems development process at various points to ensure that the process is properly controlled and managed.
Administrative controls	Formalize standards, rules, procedures, and control disciplines to ensure that the organization's general and application controls are properly executed and enforced.

Input controls check data for accuracy and completeness when they enter the system. There are specific input controls for input authorization, data conversion, data editing, and error handling. *Processing controls* establish that data are complete and accurate during updating. Output *controls ensure* that the results of computer processing are accurate, complete, and properly distributed. You can find more detail about application and general controls in our Learning Tracks.

RISK ASSESSMENT

Before your company commits resources to security and information systems controls, it must know which assets require protection and the extent to which these assets are vulnerable. A risk assessment helps answer these questions and determine the most cost-effective set of controls for protecting assets.

A risk assessment determines the level of risk to the firm if a specific activity or process is not properly controlled. Not all risks can be anticipated and measured, but most businesses will be able to acquire some understanding of the risks they face. Business managers working with information systems specialists should try to determine the value of information assets, points of vulnerability, the likely frequency of a problem, and the potential for damage. For example, if an event is likely to occur no more than once a year, with a maximum of a $1,000 loss to the organization, it is not wise to spend $20,000 on the design and maintenance of a control to protect against that event. However, if that same event could occur at least once a day, with a potential loss of more than $300,000 a year, $100,000 spent on a control might be entirely appropriate.

Table 8.5 illustrates sample results of a risk assessment for an online order processing system that processes 30,000 orders per day. The likelihood of each exposure occurring over a one-year period is expressed as a percentage. The next column shows the highest and lowest possible loss that could be expected each time the exposure occurred and an average loss calculated by adding the highest and lowest figures together and dividing by two. The expected annual loss for each exposure can be determined by multiplying the average loss by its probability of occurrence.

This risk assessment shows that the probability of a power failure occurring in a one-year period is 30 percent. Loss of order transactions while power is down could range from $5,000 to $200,000 (averaging $102,500) for each occurrence, depending on how long processing is halted. The probability of embezzlement occurring over a yearly period is about 5 percent, with potential losses ranging from $1,000 to $50,000 (and averaging $25,500) for each occurrence. User errors have a 98 percent chance of occurring over a yearly period, with losses ranging from $200 to $40,000 (and averaging $20,100) for each occurrence.

TABLE 8.5 ONLINE ORDER PROCESSING RISK ASSESSMENT

EXPOSURE	PROBABILITY OF OCCURRENCE (%)	LOSS RANGE/ AVERAGE ($)	EXPECTED ANNUAL LOSS ($)
Power failure	30%	$5,000–$200,000 ($102,500)	$30,750
Embezzlement	5%	$1,000–$50,000 ($25,500)	$1,275
User error	98%	$200–$40,000 ($20,100)	$19,698

Once the risks have been assessed, system builders will concentrate on the control points with the greatest vulnerability and potential for loss. In this case, controls should focus on ways to minimize the risk of power failures and user errors because anticipated annual losses are highest for these areas.

SECURITY POLICY

Once you've identified the main risks to your systems, your company will need to develop a security policy for protecting the company's assets. A **security policy** consists of statements ranking information risks, identifying acceptable security goals, and identifying the mechanisms for achieving these goals. What are the firm's most important information assets? Who generates and controls this information in the firm? What existing security policies are in place to protect the information? What level of risk is management willing to accept for each of these assets? Is it willing, for instance, to lose customer credit data once every 10 years? Or will it build a security system for credit card data that can withstand the once-in-a-hundred-year disaster? Management must estimate how much it will cost to achieve this level of acceptable risk.

The security policy drives other policies determining acceptable use of the firm's information resources and which members of the company have access to its information assets. An **acceptable use policy (AUP)** defines acceptable uses of the firm's information resources and computing equipment, including desktop and laptop computers, wireless devices, telephones, and the Internet. The policy should clarify company policy regarding privacy, user responsibility, and personal use of company equipment and networks. A good AUP defines unacceptable and acceptable actions for every user and specifies consequences for noncompliance. For example, security policy at Unilever, the giant multinational consumer goods company, requires every employee to use a company-specified device and employ a password or other method of identification when logging onto the corporate network.

Security policy also includes provisions for identity management. **Identity management** consists of business processes and software tools for identifying the valid users of a system and controlling their access to system resources. It includes policies for identifying and authorizing different categories of system users, specifying what systems or portions of systems each user is allowed to access, and the processes and technologies for authenticating users and protecting their identities.

Figure 8.3 is one example of how an identity management system might capture the access rules for different levels of users in the human resources function. It specifies what portions of a human resource database each user is permitted to access, based on the information required to perform that person's job. The database contains sensitive personal information such as employees' salaries, benefits, and medical histories.

The access rules illustrated here are for two sets of users. One set of users consists of all employees who perform clerical functions, such as inputting employee data into the system. All individuals with this type of profile can update the system but can neither read nor update sensitive fields, such as salary, medical history, or earnings data. Another profile applies to a divisional manager, who cannot update the system but who can read all employee data fields for his or her division, including medical history and salary. We provide more detail on the technologies for user authentication later on in this chapter.

FIGURE 8.3 ACCESS RULES FOR A PERSONNEL SYSTEM

SECURITY PROFILE 1	
User: Personnel Dept. Clerk	
Location: Division 1	
Employee Identification Codes with This Profile:	00753, 27834, 37665, 44116
Data Field Restrictions	**Type of Access**
All employee data for Division 1 only	Read and Update
• Medical history data	None
• Salary	None
• Pensionable earnings	None

SECURITY PROFILE 2	
User: Divisional Personnel Manager	
Location: Division 1	
Employee Identification Codes with This Profile:	27321
Data Field Restrictions	**Type of Access**
All employee data for Division 1 only	Read Only

These two examples represent two security profiles or data security patterns that might be found in a personnel system. Depending on the security profile, a user would have certain restrictions on access to various systems, locations, or data in an organization.

DISASTER RECOVERY PLANNING AND BUSINESS CONTINUITY PLANNING

If you run a business, you need to plan for events, such as power outages, floods, earthquakes, or terrorist attacks that will prevent your information systems and your business from operating. **Disaster recovery planning** devises plans for the restoration of computing and communications services after they have been disrupted. Disaster recovery plans focus primarily on the technical issues involved in keeping systems up and running, such as which files to back up and the maintenance of backup computer systems or disaster recovery services.

For example, MasterCard maintains a duplicate computer center in Kansas City, Missouri, to serve as an emergency backup to its primary computer center in St. Louis. Rather than build their own backup facilities, many firms contract with disaster recovery firms, such as Comdisco Disaster Recovery Services and SunGard Availability Services. These disaster recovery firms provide hot sites housing spare computers at locations around the country where subscribing firms can run their critical applications in an emergency. For example, Champion Technologies, which supplies chemicals used in oil and gas operations, is able to switch its enterprise systems from Houston to a SunGard data center in Scottsdale, Arizona, in two hours.

Business continuity planning focuses on how the company can restore business operations after a disaster strikes. The business continuity plan

identifies critical business processes and determines action plans for handling mission-critical functions if systems go down. For example, Deutsche Bank, which provides investment banking and asset management services in 74 different countries, has a well-developed business continuity plan that it continually updates and refines. It maintains full-time teams in Singapore, Hong Kong, Japan, India, and Australia to coordinate plans addressing loss of facilities, personnel, or critical systems so that the company can continue to operate when a catastrophic event occurs. Deutsche Bank's plan distinguishes between processes critical for business survival and those critical to crisis support and is coordinated with the company's disaster recovery planning for its computer centers.

Business managers and information technology specialists need to work together on both types of plans to determine which systems and business processes are most critical to the company. They must conduct a business impact analysis to identify the firm's most critical systems and the impact a systems outage would have on the business. Management must determine the maximum amount of time the business can survive with its systems down and which parts of the business must be restored first.

THE ROLE OF AUDITING

How does management know that information systems security and controls are effective? To answer this question, organizations must conduct comprehensive and systematic audits. An **information systems audit** examines the firm's overall security environment as well as controls governing individual information systems. The auditor should trace the flow of sample transactions through the system and perform tests, using, if appropriate, automated audit software. The information systems audit may also examine data quality.

Security audits review technologies, procedures, documentation, training, and personnel. A thorough audit will even simulate an attack or disaster to test the response of the technology, information systems staff, and business employees.

The audit lists and ranks all control weaknesses and estimates the probability of their occurrence. It then assesses the financial and organizational impact of each threat. Figure 8.4 is a sample auditor's listing of control weaknesses for a loan system. It includes a section for notifying management of such weaknesses and for management's response. Management is expected to devise a plan for countering significant weaknesses in controls.

8.4 WHAT ARE THE MOST IMPORTANT TOOLS AND TECHNOLOGIES FOR SAFEGUARDING INFORMATION RESOURCES?

Businesses have an array of technologies for protecting their information resources. They include tools for managing user identities, preventing unauthorized access to systems and data, ensuring system availability, and ensuring software quality.

IDENTITY MANAGEMENT AND AUTHENTICATION

Midsize and large companies have complex IT infrastructures and many different systems, each with its own set of users. Identity management

FIGURE 8.4 SAMPLE AUDITOR'S LIST OF CONTROL WEAKNESSES

Function: Loans **Prepared by: J. Ericson** **Received by: T. Benson**
Location: Peoria, IL **Date: June 16, 2015** **Review date: June 28, 2015**

Nature of Weakness and Impact	Chance for Error/Abuse		Notification to Management	
	Yes/ No	Justification	Report date	Management response
User accounts with missing passwords	Yes	Leaves system open to unauthorized outsiders or attackers	5/10/15	Eliminate accounts without passwords
Network configured to allow some sharing of system files	Yes	Exposes critical system files to hostile parties connected to the network	5/10/15	Ensure only required directories are shared and that they are protected with strong passwords
Software patches can update production programs without final approval from Standards and Controls group	No	All production programs require management approval; Standards and Controls group assigns such cases to a temporary production status		

This chart is a sample page from a list of control weaknesses that an auditor might find in a loan system in a local commercial bank. This form helps auditors record and evaluate control weaknesses and shows the results of discussing those weaknesses with management, as well as any corrective actions taken by management.

software automates the process of keeping track of all these users and their system privileges, assigning each user a unique digital identity for accessing each system. It also includes tools for authenticating users, protecting user identities, and controlling access to system resources.

To gain access to a system, a user must be authorized and authenticated. **Authentication** refers to the ability to know that a person is who he or she claims to be. Authentication is often established by using **passwords** known only to authorized users. An end user uses a password to log on to a computer system and may also use passwords for accessing specific systems and files. However, users often forget passwords, share them, or choose poor passwords that are easy to guess, which compromises security. Password systems that are too rigorous hinder employee productivity. When employees must change complex passwords frequently, they often take shortcuts, such as choosing passwords that are easy to guess or keeping their passwords at their workstations in plain view. Passwords can also be "sniffed" if transmitted over a network or stolen through social engineering.

New authentication technologies, such as tokens, smart cards, and biometric authentication, overcome some of these problems. A **token** is a physical device, similar to an identification card, that is designed to prove the identity of a single user. Tokens are small gadgets that typically fit on key rings and display passcodes that change frequently. A **smart card** is a device about the size of a credit card that contains a chip formatted with access permission and other data. (Smart cards are also used in electronic payment systems.) A reader device interprets the data on the smart card and allows or denies access.



© Aleksey Boldin/Alamy

This smartphone has a biometric fingerprint reader for fast yet secure access to files and networks. New models of PCs and smartphones are starting to use biometric identification to authenticate users.

Biometric authentication uses systems that read and interpret individual human traits, such as fingerprints, irises, and voices, in order to grant or deny access. Biometric authentication is based on the measurement of a physical or behavioral trait that makes each individual unique. It compares a person's unique characteristics, such as the fingerprints, face, or retinal image, against a stored profile of these characteristics to determine whether there are any differences between these characteristics and the stored profile. If the two profiles match, access is granted. Fingerprint and facial recognition technologies are just beginning to be used for security applications, with many PC laptops (and some smartphones) equipped with fingerprint identification devices and several models with built-in webcams and face recognition software.

The steady stream of incidents in which hackers have been able to access traditional passwords highlights the need for more secure means of authentication. **Two-factor authentication** increases security by validating users with a multi-step process. To be authenticated, a user must provide two means of identification, one of which is typically a physical token, such as a smartcard or chip-enabled bank card, and the other of which is typically data, such as a password or PIN (personal identification number). Biometric data, such as fingerprints, iris prints, or voice prints, can also be used as one of the authenticating mechanisms. A common example of two-factor authentication is a bank card: the card itself is the physical item and the PIN is the data that go with it.

FIREWALLS, INTRUSION DETECTION SYSTEMS, AND ANTIVIRUS SOFTWARE

Without protection against malware and intruders, connecting to the Internet would be very dangerous. Firewalls, intrusion detection systems, and antivirus software have become essential business tools.

Firewalls

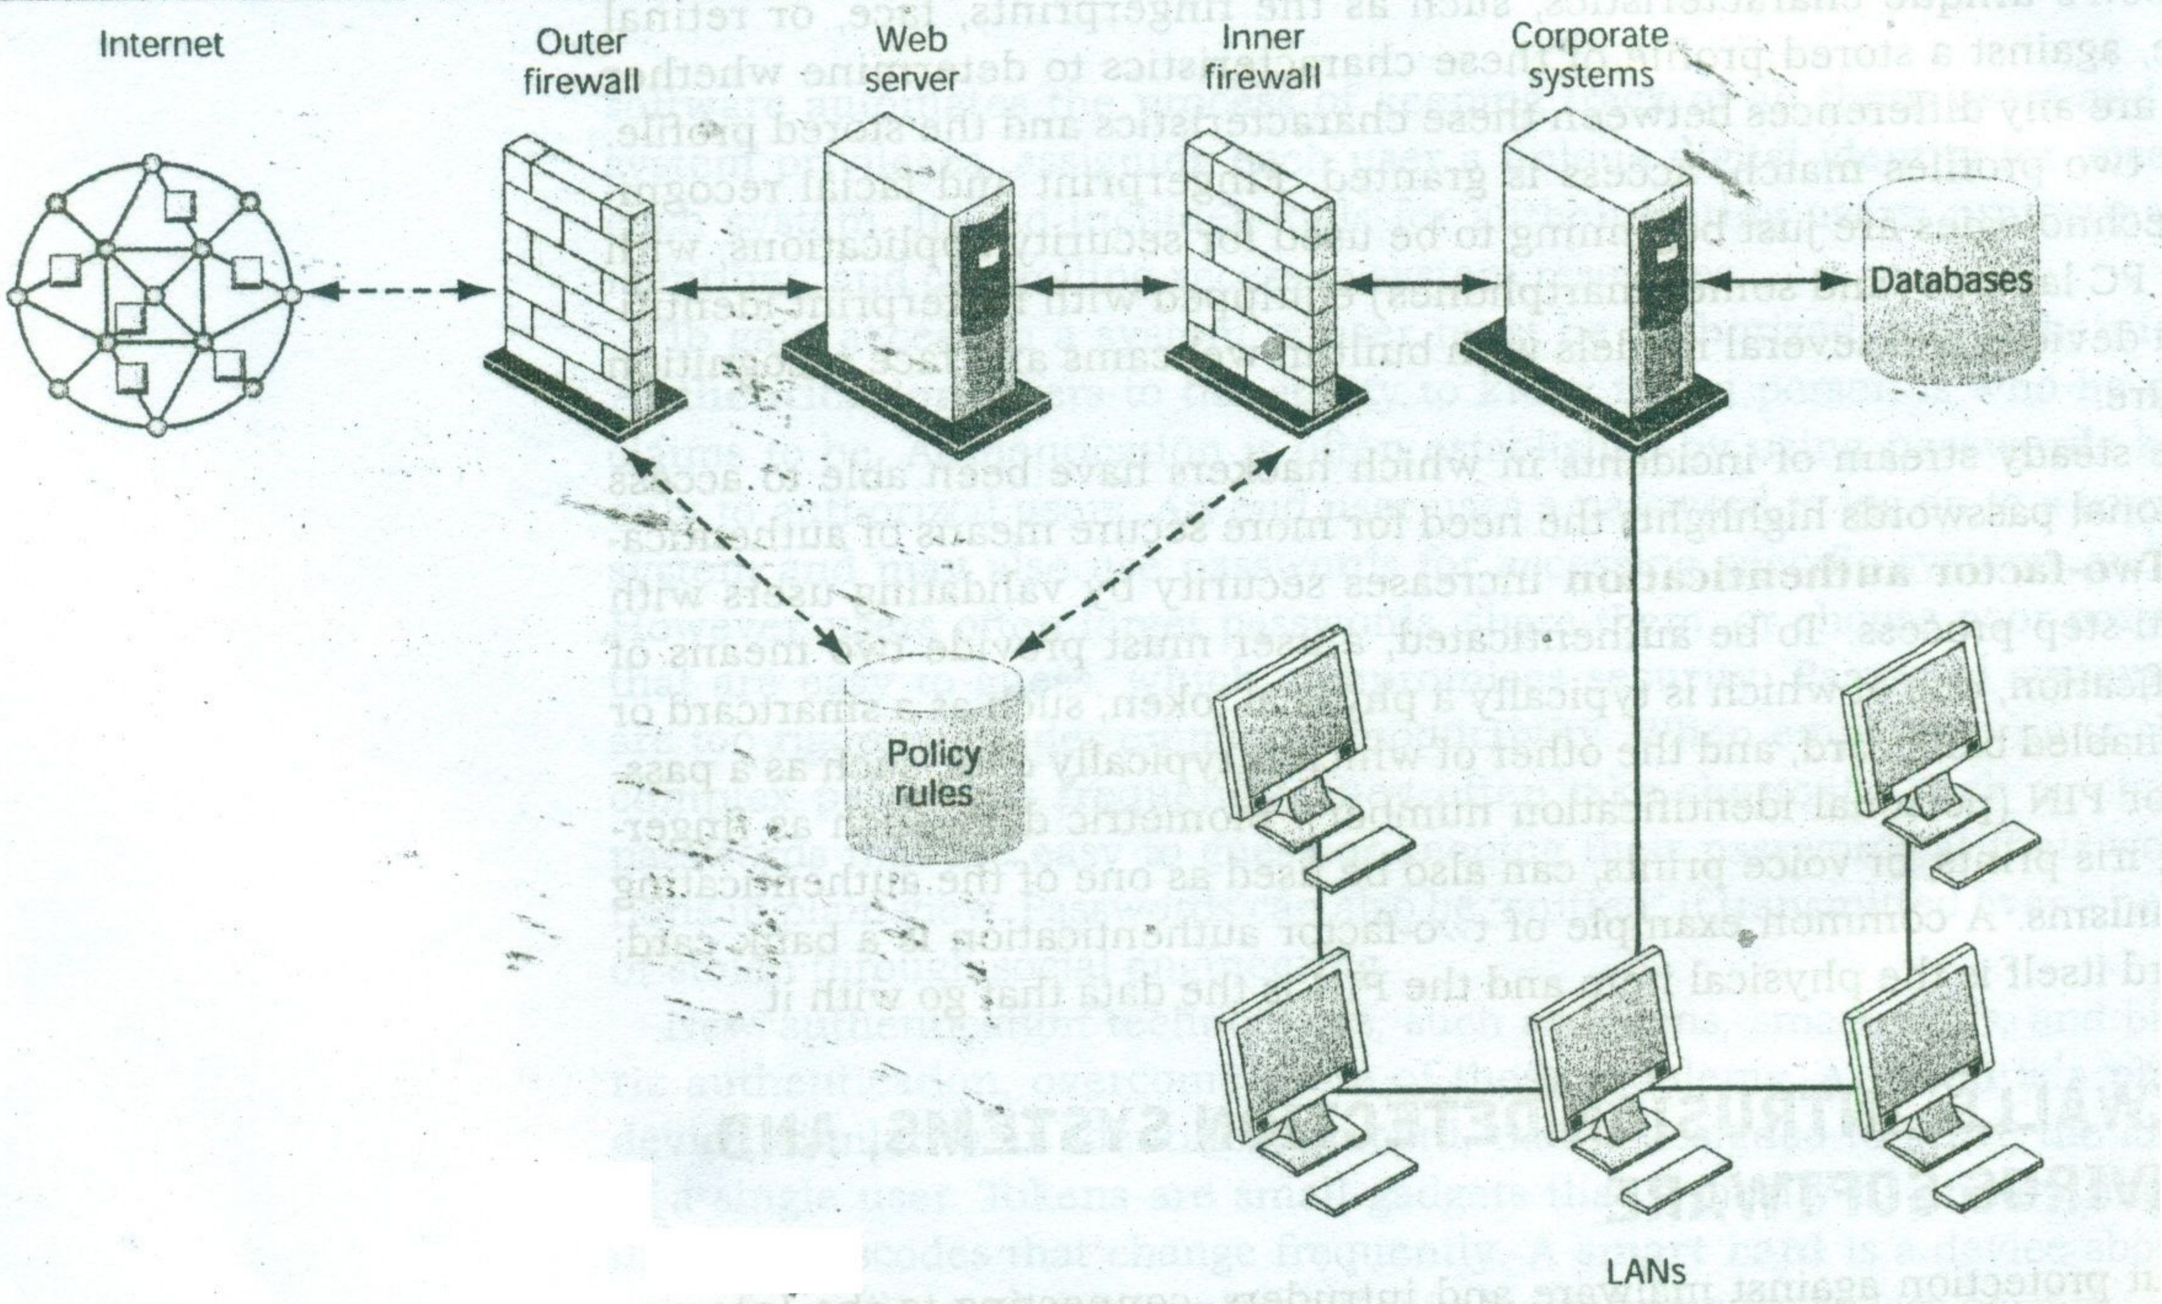
Firewalls prevent unauthorized users from accessing private networks. A firewall is a combination of hardware and software that controls the flow of incoming and outgoing network traffic. It is generally placed between the organization's private internal networks and distrusted external networks, such as the Internet, although firewalls can also be used to protect one part of a company's network from the rest of the network (see Figure 8.5).

The firewall acts like a gatekeeper who examines each user's credentials before access is granted to a network. The firewall identifies names, IP addresses, applications, and other characteristics of incoming traffic. It checks this information against the access rules that have been programmed into the system by the network administrator. The firewall prevents unauthorized communication into and out of the network.

In large organizations, the firewall often resides on a specially designated computer separate from the rest of the network, so no incoming request directly accesses private network resources. There are a number of firewall screening technologies, including static packet filtering, stateful inspection, Network Address Translation, and application proxy filtering. They are frequently used in combination to provide firewall protection.

Packet filtering examines selected fields in the headers of data packets flowing back and forth between the trusted network and the Internet, examining individual packets in isolation. This filtering technology can miss many types of attacks. *Stateful inspection* provides additional security by determining whether

FIGURE 8.5 A CORPORATE FIREWALL



The firewall is placed between the firm's private network and the public Internet or another distrusted network to protect against unauthorized traffic.

packets are part of an ongoing dialogue between a sender and a receiver. It sets up state tables to track information over multiple packets. Packets are accepted or rejected based on whether they are part of an approved conversation or whether they are attempting to establish a legitimate connection.

Network Address Translation (NAT) can provide another layer of protection when static packet filtering and stateful inspection are employed. NAT conceals the IP addresses of the organization's internal host computer(s) to prevent sniffer programs outside the firewall from ascertaining them and using that information to penetrate internal systems.

Application proxy filtering examines the application content of packets. A proxy server stops data packets originating outside the organization, inspects them, and passes a proxy to the other side of the firewall. If a user outside the company wants to communicate with a user inside the organization, the outside user first "talks" to the proxy application and the proxy application communicates with the firm's internal computer. Likewise, a computer user inside the organization goes through the proxy to talk with computers on the outside.

To create a good firewall, an administrator must maintain detailed internal rules identifying the people, applications, or addresses that are allowed or rejected. Firewalls can deter, but not completely prevent, network penetration by outsiders and should be viewed as one element in an overall security plan.

Intrusion Detection Systems

In addition to firewalls, commercial security vendors now provide intrusion detection tools and services to protect against suspicious network traffic and attempts to access files and databases. **Intrusion detection systems** feature full-time monitoring tools placed at the most vulnerable points or "hot spots" of corporate networks to detect and deter intruders continually. The system generates an alarm if it finds a suspicious or anomalous event. Scanning software looks for patterns indicative of known methods of computer attacks, such as bad passwords, checks to see if important files have been removed or modified, and sends warnings of vandalism or system administration errors. Monitoring software examines events as they are happening to discover security attacks in progress. The intrusion detection tool can also be customized to shut down a particularly sensitive part of a network if it receives unauthorized traffic.

Antivirus and Antispyware Software

Defensive technology plans for both individuals and businesses must include anti-malware protection for every computer. **Antivirus software** prevents, detects, and removes malware, including computer viruses, computer worms, Trojan horses, spyware, and adware. However, most antivirus software is effective only against malware already known when the software was written. To remain effective, the antivirus software must be continually updated, and even then it is not always effective. According to a report by Solutionary Security Engineering Research Team (SERT), 54 percent of malware evades anti-virus detection. Organizations need to use additional malware detection tools for better protection (Solutionary, 2013).

Unified Threat Management Systems

To help businesses reduce costs and improve manageability, security vendors have combined into a single appliance various security tools, including firewalls, virtual private networks, intrusion detection systems, and Web content

filtering and antispam software. These comprehensive security management products are called **unified threat management (UTM)** systems. Although initially aimed at small and medium-sized businesses, UTM products are available for all sizes of networks. Leading UTM vendors include Blue Coat, Fortinent, and Check Point, and networking vendors such as Cisco Systems and Juniper Networks provide some UTM capabilities in their products.

SECURING WIRELESS NETWORKS

The initial security standard developed for Wi-Fi, called Wired Equivalent Privacy (WEP), is not very effective because its encryption keys are relatively easy to crack. WEP provides some margin of security, however, if users remember to enable it. Corporations can further improve Wi-Fi security by using it in conjunction with virtual private network (VPN) technology when accessing internal corporate data.

In June 2004, the Wi-Fi Alliance industry trade group finalized the 802.11i specification (also referred to as Wi-Fi Protected Access 2 or WPA2) that replaces WEP with stronger security standards. Instead of the static encryption keys used in WEP, the new standard uses much longer keys that continually change, making them harder to crack. It also employs an encrypted authentication system with a central authentication server to ensure that only authorized users access the network.

ENCRYPTION AND PUBLIC KEY INFRASTRUCTURE

Many businesses use encryption to protect digital information that they store, physically transfer, or send over the Internet. **Encryption** is the process of transforming plain text or data into cipher text that cannot be read by anyone other than the sender and the intended receiver. Data are encrypted by using a secret numerical code, called an encryption key, that transforms plain data into cipher text. The message must be decrypted by the receiver.

Two methods for encrypting network traffic on the Web are SSL and S-HTTP. **Secure Sockets Layer (SSL)** and its successor Transport Layer Security (TLS) enable client and server computers to manage encryption and decryption activities as they communicate with each other during a secure Web session. **Secure Hypertext Transfer Protocol (S-HTTP)** is another protocol used for encrypting data flowing over the Internet, but it is limited to individual messages, whereas SSL and TLS are designed to establish a secure connection between two computers.

The capability to generate secure sessions is built into Internet client browser software and servers. The client and the server negotiate what key and what level of security to use. Once a secure session is established between the client and the server, all messages in that session are encrypted.

Two methods of encryption are symmetric key encryption and public key encryption. In symmetric key encryption, the sender and receiver establish a secure Internet session by creating a single encryption key and sending it to the receiver so both the sender and receiver share the same key. The strength of the encryption key is measured by its bit length. Today, a typical key will be 128 bits long (a string of 128 binary digits).

The problem with all symmetric encryption schemes is that the key itself must be shared somehow among the senders and receivers, which exposes the key to outsiders who might just be able to intercept and decrypt the key.

FIGURE 8.6 PUBLIC KEY ENCRYPTION

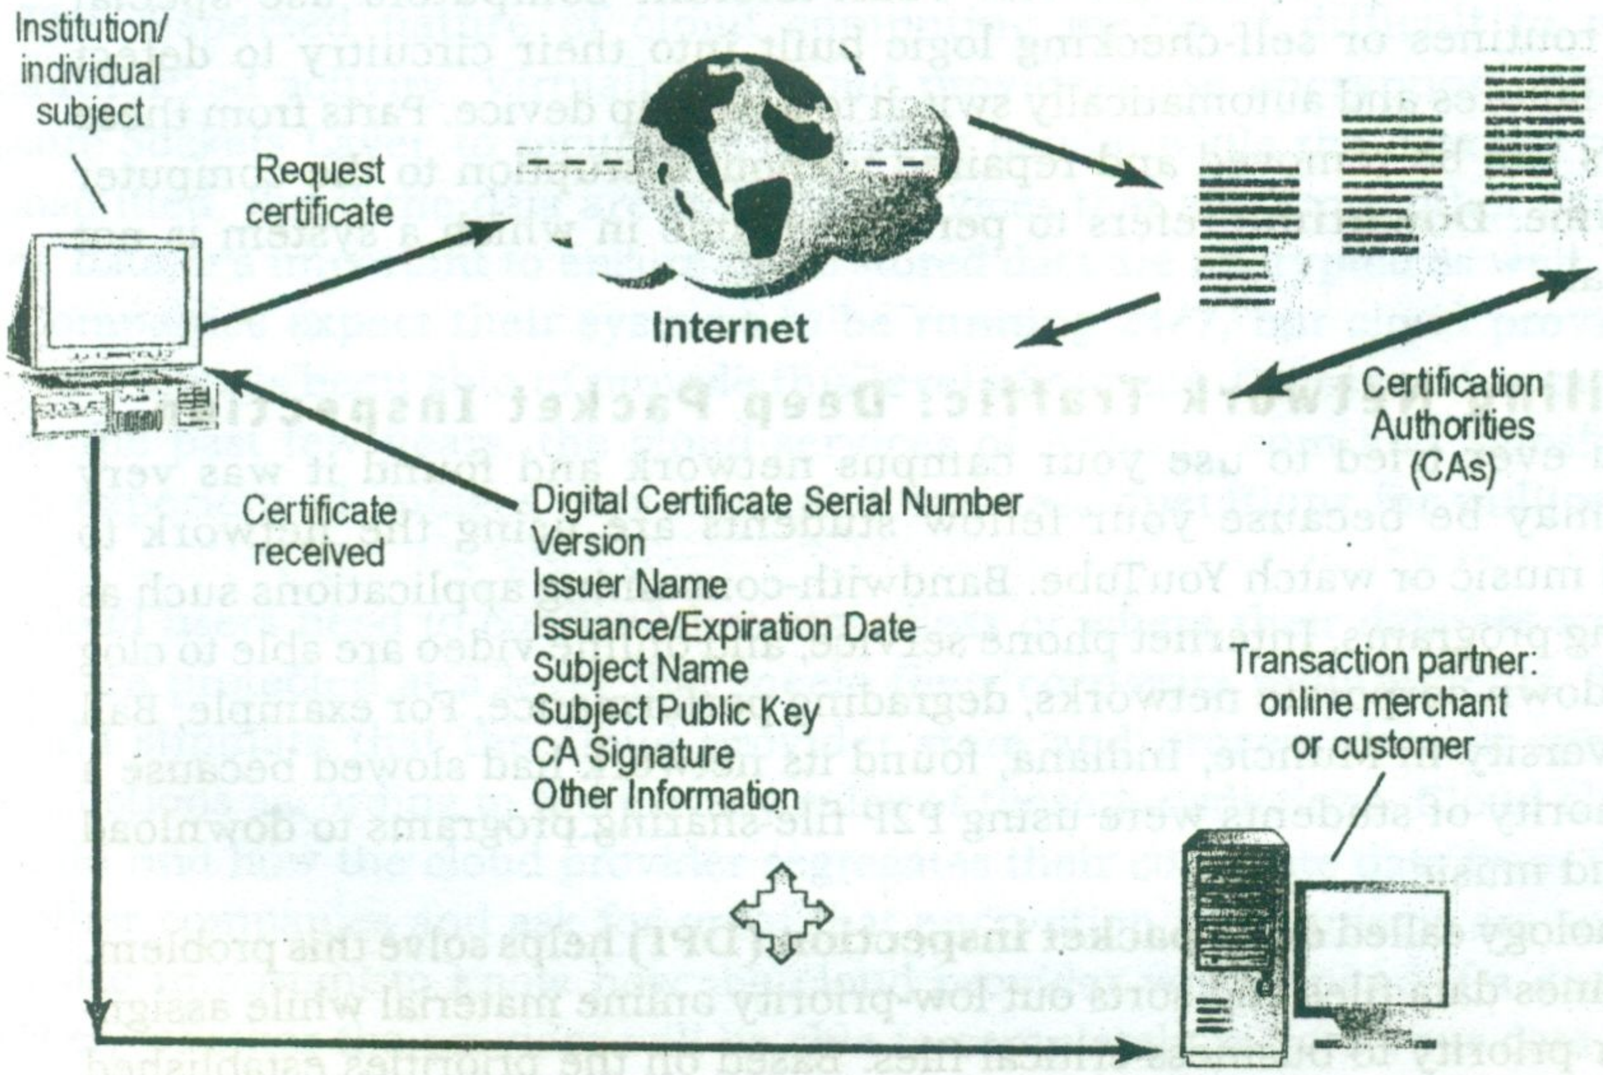


A public key encryption system can be viewed as a series of public and private keys that lock data when they are transmitted and unlock the data when they are received. The sender locates the recipient's public key in a directory and uses it to encrypt a message. The message is sent in encrypted form over the Internet or a private network. When the encrypted message arrives, the recipient uses his or her private key to decrypt the data and read the message.

A more secure form of encryption called **public key encryption** uses two keys: one shared (or public) and one totally private as shown in Figure 8.6. The keys are mathematically related so that data encrypted with one key can be decrypted using only the other key. To send and receive messages, communicators first create separate pairs of private and public keys. The public key is kept in a directory and the private key must be kept secret. The sender encrypts a message with the recipient's public key. On receiving the message, the recipient uses his or her private key to decrypt it.

Digital certificates are data files used to establish the identity of users and electronic assets for protection of online transactions (see Figure 8.7). A digital certificate system uses a trusted third party, known as a certificate authority

FIGURE 8.7 DIGITAL CERTIFICATES



Digital certificates help establish the identity of people or electronic assets. They protect online transactions by providing secure, encrypted, online communication.

(CA, or certification authority), to validate a user's identity. There are many CAs in the United States and around the world, including Symantec, GoDaddy, and Comodo.

The CA verifies a digital certificate user's identity offline. This information is put into a CA server, which generates an encrypted digital certificate containing owner identification information and a copy of the owner's public key. The certificate authenticates that the public key belongs to the designated owner. The CA makes its own public key available either in print or perhaps on the Internet. The recipient of an encrypted message uses the CA's public key to decode the digital certificate attached to the message, verifies it was issued by the CA, and then obtains the sender's public key and identification information contained in the certificate. Using this information, the recipient can send an encrypted reply. The digital certificate system would enable, for example, a credit card user and a merchant to validate that their digital certificates were issued by an authorized and trusted third party before they exchange data. **Public key infrastructure (PKI)**, the use of public key cryptography working with a CA, is now widely used in e-commerce.

ENSURING SYSTEM AVAILABILITY

As companies increasingly rely on digital networks for revenue and operations, they need to take additional steps to ensure that their systems and applications are always available. Firms such as those in the airline and financial services industries with critical applications requiring online transaction processing have traditionally used fault-tolerant computer systems for many years to ensure 100 percent availability. In **online transaction processing**, transactions entered online are immediately processed by the computer. Multitudinous changes to databases, reporting, and requests for information occur each instant.

Fault-tolerant computer systems contain redundant hardware, software, and power supply components that create an environment that provides continuous, uninterrupted service. Fault-tolerant computers use special software routines or self-checking logic built into their circuitry to detect hardware failures and automatically switch to a backup device. Parts from these computers can be removed and repaired without disruption to the computer or downtime. **Downtime** refers to periods of time in which a system is not operational.

Controlling Network Traffic: Deep Packet Inspection

Have you ever tried to use your campus network and found it was very slow? It may be because your fellow students are using the network to download music or watch YouTube. Bandwith-consuming applications such as file-sharing programs, Internet phone service, and online video are able to clog and slow down corporate networks, degrading performance. For example, Ball State University in Muncie, Indiana, found its network had slowed because a small minority of students were using P2P file-sharing programs to download movies and music.

A technology called **deep packet inspection (DPI)** helps solve this problem. DPI examines data files and sorts out low-priority online material while assigning higher priority to business-critical files. Based on the priorities established by a network's operators, it decides whether a specific data packet can continue to its destination or should be blocked or delayed while more important traffic

proceeds. Using a DPI system from Allot Communications, Ball State was able to cap the amount of file-sharing traffic and assign it a much lower priority. Ball State's preferred network traffic speeded up.

Security Outsourcing

Many companies, especially small businesses, lack the resources or expertise to provide a secure high-availability computing environment on their own. They can outsource many security functions to **managed security service providers (MSSPs)** that monitor network activity and perform vulnerability testing and intrusion detection. SecureWorks, BT Managed Security Solutions Group, and Symantec are leading providers of MSSP services.

SECURITY ISSUES FOR CLOUD COMPUTING AND THE MOBILE DIGITAL PLATFORM

Although cloud computing and the emerging mobile digital platform have the potential to deliver powerful benefits, they pose new challenges to system security and reliability. We now describe some of these challenges and how they should be addressed.

Security in the Cloud

When processing takes place in the cloud, accountability and responsibility for protection of sensitive data still reside with the company owning that data. Understanding how the cloud computing provider organizes its services and manages the data is critical.

Cloud computing is highly distributed. Cloud applications reside in large remote data centers and server farms that supply business services and data management for multiple corporate clients. To save money and keep costs low, cloud computing providers often distribute work to data centers around the globe where work can be accomplished most efficiently. When you use the cloud, you may not know precisely where your data are being hosted.

The dispersed nature of cloud computing makes it difficult to track unauthorized activity. Virtually all cloud providers use encryption, such as Secure Sockets Layer, to secure the data they handle while the data are being transmitted. But if the data are stored on devices that also store other companies' data, it's important to ensure these stored data are encrypted as well.

Companies expect their systems to be running 24/7, but cloud providers haven't always been able to provide this level of service. On several occasions over the past few years, the cloud services of Amazon.com and Salesforce.com experienced outages that disrupted business operations for millions of users.

Cloud users need to confirm that regardless of where their data are stored, they are protected at a level that meets their corporate requirements. They should stipulate that the cloud provider store and process data in specific jurisdictions according to the privacy rules of those jurisdictions. Cloud clients should find how the cloud provider segregates their corporate data from those of other companies and ask for proof that encryption mechanisms are sound. It's also important to know how the cloud provider will respond if a disaster strikes, whether the provider will be able to completely restore your data, and how long this should take. Cloud users should also ask whether cloud providers will submit to external audits and security certifications. These kinds of controls

can be written into the service level agreement (SLA) before signing with a cloud provider.

Securing Mobile Platforms

If mobile devices are performing many of the functions of computers, they need to be secured like desktops and laptops against malware, theft, accidental loss, unauthorized access, and hacking attempts.

Mobile devices accessing corporate systems and data require special protection. Companies should make sure that their corporate security policy includes mobile devices, with additional details on how mobile devices should be supported, protected, and used. They will need mobile device management tools to authorize all devices in use; to maintain accurate inventory records on all mobile devices, users, and applications; to control updates to applications; and to lock down or erase lost or stolen devices so they can't be compromised. Data loss prevention technology can identify where critical data are saved, who is accessing the data, how data are leaving the company, and where the data are going. Firms should develop guidelines stipulating approved mobile platforms and software applications as well as the required software and procedures for remote access of corporate systems. The organization's mobile security policy should forbid employees from using unsecure, consumer-based applications for transferring and storing corporate documents and files, or sending such documents and files to oneself via e-mail without encryption.

Companies should encrypt communication whenever possible. All mobile device users should be required to use the password feature found in every smartphone. Mobile security products are available from Kaspersky, Symantec, Trend Micro, and McAfee.

Some companies insist that employees use only company-issued smartphones. BlackBerry devices are considered the most secure because they run within their own secure system. But, increasingly, companies are allowing employees to use their own devices, including iPhones, iPads, and Android phones, for work, to make employees more available and productive (see the Chapter 5 discussion of BYOD). Protective software products, such as the tools from Good Technology, are now available for segregating corporate data housed within personally owned mobile devices from the device's personal content.

ENSURING SOFTWARE QUALITY

In addition to implementing effective security and controls, organizations can improve system quality and reliability by employing software metrics and rigorous software testing. Software metrics are objective assessments of the system in the form of quantified measurements. Ongoing use of metrics allows the information systems department and end users to jointly measure the performance of the system and identify problems as they occur. Examples of software metrics include the number of transactions that can be processed in a specified unit of time, online response time, the number of payroll checks printed per hour, and the number of known bugs per hundred lines of program code. For metrics to be successful, they must be carefully designed, formal, objective, and used consistently.

INTERACTIVE SESSION: TECHNOLOGY

MWEB BUSINESS: HACKED

BMWEB, launched in 1997, became South Africa's leading ISP in 1998. It has established itself as a company that provides a cutting-edge network and service infrastructure and outstanding customer service. Currently, MWEB's customer base of 320,000 includes home users; small, medium, and large business customers; and corporate clients. MWEB won the ISP of the Year award at the MyBroadband Conference in Johannesburg in 2010. The award was based on the performance of its various broadband services as well as on customer satisfaction.

Its business division, MWEB Business, was founded in January 1998. MWEB Business prides itself as being a business partner that is perfectly positioned to leverage the power of Web-based technologies in all areas of an organization. MWEB Business helps companies:

- Manage business data in ways that add real value and insight to their operations.
- Integrate existing systems with the Internet so as to close the gap between technology, strategy, and the organization's bottom line.
- Develop, manage, and maintain solutions that include all aspects of Internet connectivity, Web site development and hosting, broadband and wireless applications, e-commerce, and consulttancy services .
- Manage internal information among employees, as well as among business partners and suppliers.

MWEB has moved forward in publicizing its plans for the South African Internet market. According to MWEB CEO Rudi Jansen, the company needs to improve the quality of their network, which is not only an MWEB problem, but also a Telkom network problem. Despite having a less-than-ideal network infrastructure, MWEB uses AVG Internet Security to offer its customers the best possible security while online. AVG Internet Security offers MWEB customers the following features:

- Identity protection for safe banking and shopping
- LinkScanner for safe surfing and searching
- WebShield for safe social networking, chatting, and downloading
- Antiphishing and antispam for a safe uncluttered inbox
- High-speed antivirus/antispyware software with automatic updates
- An enhanced firewall

In addition, MWEB automatically protects customers against junk email and viruses that are sent via email. Its virus filter ensures that only virus-free email is delivered to clients' inboxes by automatically cleaning e-mails from recognized malware sources. MWEB advises its customers to keep their ADSL connections safe from bandwidth theft and account abuse by blocking unsolicited incoming connections to network ports commonly used by hackers.

Despite the multitude of security services offered by MWEB, a number of MWEB Business subscribers' account details were compromised when their login and password details were published on the Internet by hackers. Initial reports indicated that as many as 2,390 users of MWEB's business digital subscriber lines were affected. The company disclosed the security breach on October 25, 2010. It appears that hackers gained access to Internet Solutions' self-service management system, which MWEB Business uses to provide and manage business accounts that have not yet been migrated to the MWEB network.

Historically, MWEB Business was a reseller of Internet Solutions' Uncapped & Fixed IP ADSL services, which were provisioned and managed by MWEB using a Web-based management interface provided by Internet Solutions. All new Business ADSL services provided after April 2010, as well as the bulk of legacy services already migrated, used MWEB's internal authentication systems, which were completely unaffected by this incident.

MWEB responded quickly to the hacking incident. According to Jansen, about 1,000 clients on the Internet Solutions network needed to be migrated from the old server that was attacked by hackers. Although the network was quickly secured, most customers had recently been moved to MWEB's IPC network. MWEB also contacted these customers to reset their passwords as an added security measure. Jansen was quick to note that no personal information was lost and that none of MWEB's clients suffered any losses as their usernames and passwords had been recreated and changed. He further added that MWEB successfully repels 5,000 attacks a day.

Andre Joubert, general manager of MWEB Business, emphasized that only ADSL authentica-

tion usernames and passwords had been compromised. The integrity of the personal or private data related to the accounts remained intact, as did the access credentials for each customer's bundled onsite router. Joubert did acknowledge the seriousness of the hack, apologizing for any inconvenience the breach may have caused to MWEB's customers. As soon as the breach was identified, MWEB took immediate action to evaluate the extent of the breach and to limit any damage. In MWEB's defense, Jansen said that MWEB constantly advises its customers to be vigilant regarding their online data and security. In addition, MWEB was working closely with Internet Solutions to investigate the nature and source of the breach to ensure that it does not happen again.

Sources: "2010 MyBroadband Awards: The Winners and Losers," MyBroadband, October 19, 2010, http://mybroadband.co.za/news/ broadband/15951-2010-MyBroadband-Awards-The-winners-andlosers. html, accessed November 17, 2010; "About MWEB," MWEB, www.mweb.co.za/productspricing/MWEBBusiness/ AboutMWEBB usiness.aspx, accessed November 17, 2010; "Hackers Target MWEB," NewsTime, October 25, 2010, www.newstime.co.za/ ScienceandTech/Hackers_Target_M-Web/13618/, accessed November 17, 2010; "MWEB Business Tackles 'ADSL Hacking' Incident," MyBroadband, October 25, 2010, http://mybroadband.co.za/news/adsl/16077-MWEB-Businesstackles-ADSL-hacking-incident.html, accessed November 17, 2010; "MWEB Business Takes Action in 'Hacking' Incident," Moneyweb, October 25, 2010, www.moneyweb.co.za/mw/view/mw/en/ page295027?oid=5125 45&sn=2009+Detail&pid=287226, accessed November 17, 2010; "MWeb Hacked, Users' Details Exposed," TechCentral, October 26, 2010, www.techcentral.co.za/mwebhacked-users-details-exposed/18366/, accessed November 17, 2010.

Case contributed by Upasana Singh, University of KwaZulu-Natal.

CASE STUDY QUESTIONS

1. What technology issues led to the security breach at MWEB?
2. What is the possible business impact of this security breach for both MWEB and its customers?
3. If you were an MWEB customer, would you consider MWEB's response to the security breach to be acceptable? Why or why not?
4. What should MWEB do in the future to avoid similar incidents?

Early, regular, and thorough testing will contribute significantly to system quality. Many view testing as a way to prove the correctness of work they have done. In fact, we know that all sizable software is riddled with errors, and we must test to uncover these errors.

Good testing begins before a software program is even written by using a *walkthrough*—a review of a specification or design document by a small group of people carefully selected based on the skills needed for the particular objectives being tested. Once developers start writing software programs, coding walkthroughs can also be used to review program code. However, code must be tested by computer runs. When errors are discovered, the source is found and eliminated through a process called *debugging*. You can find out more about the various stages of testing required to put an information system into operation in Chapter 13. Our Learning Tracks also contain descriptions of methodologies for developing software programs that also contribute to software quality.

Review Summary

1. *Why are information systems vulnerable to destruction, error, and abuse?*

Digital data are vulnerable to destruction, misuse, error, fraud, and hardware or software failures. The Internet is designed to be an open system and makes internal corporate systems more vulnerable to actions from outsiders. Hackers can unleash denial-of-service (DoS) attacks or penetrate corporate networks, causing serious system disruptions. Wi-Fi networks can easily be penetrated by intruders using sniffer programs to obtain an address to access the resources of the network. Computer viruses and worms can disable systems and Web sites. The dispersed nature of cloud computing makes it difficult to track unauthorized activity or to apply controls from afar. Software presents problems because software bugs may be impossible to eliminate and because software vulnerabilities can be exploited by hackers and malicious software. End users often introduce errors.

2. *What is the business value of security and control?*

Lack of sound security and control can cause firms relying on computer systems for their core business functions to lose sales and productivity. Information assets, such as confidential employee records, trade secrets, or business plans, lose much of their value if they are revealed to outsiders or if they expose the firm to legal liability. New laws, such as HIPAA, the Sarbanes-Oxley Act, and the Gramm-Leach-Bliley Act, require companies to practice stringent electronic records management and adhere to strict standards for security, privacy, and control. Legal actions requiring electronic evidence and computer forensics also require firms to pay more attention to security and electronic records management.

3. *What are the components of an organizational framework for security and control?*

What are the components of an organizational framework for security and control? Firms need to establish a good set of both general and application controls for their information systems. A risk assessment evaluates information assets, identifies control points and control weaknesses, and determines the most cost-effective set of controls. Firms must also develop a coherent corporate security policy and plans for continuing business operations in the event of disaster or disruption. The security policy includes policies for acceptable use and identity management. Comprehensive and systematic information systems auditing helps organizations determine the effectiveness of security and controls for their information systems.

4. *What are the most important tools and technologies for safeguarding information resources?*

Firewalls prevent unauthorized users from accessing a private network when it is linked to the Internet. Intrusion detection systems monitor private networks from suspicious network traffic and attempts to access corporate systems. Passwords, tokens, smart cards, and biometric authentication are used to authenticate system users. Antivirus software checks computer systems for infections by viruses and worms and often eliminates the malicious software, while antispyware software combats intrusive and harmful spyware programs. Encryption, the coding and scrambling of messages, is a widely used technology for securing electronic transmissions over unprotected networks. Digital certificates combined with public key encryption provide further protection of electronic transactions by authenticating a user's identity. Companies can use fault-tolerant computer systems to make sure that their information systems are always available. Use of software metrics and rigorous software testing help improve software quality and reliability.

Key Terms

Acceptable use policy (AUP), 323
Antivirus software, 329
Application controls, 321
Authentication, 326
Biometric authentication, 327
Botnet, 312
Bugs, 318
Business continuity planning, 324
Click fraud, 316
Computer crime, 312
Computer forensics, 320
Computer virus, 308
Controls, 306
Cybervandalism, 311

Cyberwarfare, 349
Deep packet inspection (DPI), 364
Denial-of-service (DoS) attack, 344
Digital certificates, 363
Disaster recovery planning, 356
Distributed denial-of-service (DDoS) attack, 344
Downtime, 364
Drive-by download, 341
Encryption, 362
Evil twin, 345
Fault-tolerant computer systems, 364
Firewall, 360
General controls, 353
Gramm-Leach-Bliley Act, 351
Hacker, 343
HIPAA, 351
Identity management, 355
Identity theft, 344
Information systems audit, 357
Intrusion detection systems, 361
Keyloggers, 343
Malware, 340
Managed security service providers (MSSPs), 365
Online transaction processing, 364
Password, 358
Patches, 350
Pharming, 346
Phishing, 345
Public key encryption, 363
Public key infrastructure (PKI), 364
Ransomware, 342
Risk assessment, 364
Sarbanes-Oxley Act, 351
Secure Hypertext Transfer Protocol (S-HTTP), 362
Secure Sockets Layer (SSL), 362
Security, 338
Security policy, 355
Smart card, 358
Sniffer, 343
Social engineering, 349
Spoofing, 343
Spyware, 343
SQL injection attack, 342
Token, 358
Trojan horse, 341
Two-factor authentication, 359
Unified threat management (UTM), 362
War driving, 340
Worms, 341

MyMISLab

Go to **mymislab.com** to complete the problems marked with this icon ⟲.

Review Questions

8-1 Why are information systems vulnerable to destruction, error, and abuse?

- List and describe the most common threats against contemporary information systems.
- Define malware and distinguish among a virus, a worm, and a Trojan horse.
- Explain the challenges presented in securing wireless networks.
- Define computer crime. Provide two examples of crime in which computers are targets and two examples in which computers are used as instruments of crime.
- Define identity theft and phishing and explain why identity theft is such a big problem today.
- Define denial of service (DoS) attack, explain how it differs from a distributed denial of service (DDoS) attack, and discuss how DoS and DDoS attacks are related to the use of botnets.
- Explain how software defects affect system reliability and security.

8-2 What is the business value of security and control?

- Explain how security and control provide value for businesses.
- Define computer forensics and explain what it is used for.

8-3 What are the components of an organizational framework for security and control?

- Define general controls and describe each type of general control.
- Define application controls and describe each type of application control.
- Describe the function of risk assessment and explain how it is conducted for information systems.
- Define and describe the following: security policy, acceptable use policy, and identity management.

- Explain how information systems auditing promotes security and control.

8-4 What are the most important tools and technologies for safeguarding information resources?

- Name and describe three authentication methods.
- Describe the roles of firewalls, intrusion detection systems, and antivirus software in promoting security.
- Explain how encryption protects information.
- Describe the role of encryption and digital certificates in a public key infrastructure.
- Distinguish between disaster recovery planning and business continuity planning.
- Identify and describe the security problems posed by cloud computing.
- Explain the actions companies should take to secure mobile platforms.

Discussion Questions

8-5 Security isn't simply a technology issue, it's a business issue. Discuss.

8-6 Who poses the biggest security threat: insiders or outsiders?

8-7 Suppose your business had an e-commerce Web site where it sold goods and accepted credit card payments. Discuss the major security threats to this Web site and their potential impact. What can be done to minimize these threats?

Hands-On MIS Projects

The projects in this section give you hands-on experience analyzing security vulnerabilities, using spreadsheet software for risk analysis, and using Web tools to research security outsourcing services.

Management Decision Problems

8-8 Gifty is an online e-tailer for handmade gifts. Customers can purchase either via its Web site or via a mobile app. Prepare a security analysis for this Internet-based business. What kinds of threats should it anticipate? What would be their impact on the business? What steps can it take to prevent damage to its Web sites and continuing operations?

8-9 A survey of your firm's IT infastructure has identified a number of security vulnerabilities. Review the data on these vulnerabilities, which can be found in a table in MyMISLab™. Use the table to answer the following questions:

- Calculate the total number of vulnerabilities for each platform. What is the potential impact of the security problems for each computing platform on the organization?
- If you only have one information systems specialist in charge of security, which platforms should you address first in trying to eliminate these vulnerabilities? Second? Third? Last? Why?
- Identify the types of control problems illustrated by these vulnerabilities and explain the measures that should be taken to solve them.
- What does your firm risk by ignoring the security vulnerabilities identified?

Improving Decision Making: Using Spreadsheet Software to Perform a Security Risk Assessment

Software skills: Spreadsheet formulas and charts
Business skills: Risk assessment

8-10 This project uses spreadsheet software to calculate anticipated annual losses from various security threats identified for a small company.

Mercer Paints is a paint manufacturing company located in Alabama that uses a network to link its business operations. A security risk assessment requested by management identified a number of potential exposures. + exposures, their associated probabilities, and average losses are summarized in a table, which can be found in MyMISLab. Use the table to answer the following questions:

- In addition to the potential exposures listed, identify at least three other potential threats to Mercer Paints, assign probabilities, and estimate a loss range.
- Use spreadsheet software and the risk assessment data to calculate the expected annual loss for each exposure.
- Present your findings in the form of a chart. Which control points have the greatest vulnerability? What recommendations would you make to Mercer Paints? Prepare a written report that summarizes your findings and recommendations.

Improving Decision Making: Evaluating Security Outsourcing Services

Software skills: Web browser and presentation software
Business skills: Evaluating business outsourcing services

8-11 This project will help develop your Internet skills in using the Web to research and evaluate security outsourcing services.

You have been asked to help your company's management decide whether to outsource security or keep the security function within the firm. Search the Web to find information to help you decide whether to outsource security and to locate security outsourcing services.

- Present a brief summary of the arguments for and against outsourcing computer security for your company.
- Select two firms that offer computer security outsourcing services, and compare them and their services.
- Prepare an electronic presentation for management summarizing your findings. Your presentation should make the case on whether or not your company should outsource computer security. If you believe your company should outsource, the presentation should identify which security outsourcing service you selected and justify your decision.

Collaboration and Teamwork Project

8-12 In MyMISLab, you will find a Collaboration and Teamwork Project dealing with the concepts in this chapter. You will be able to use Google Drive, Google Docs, Google Sites, Google+, or other open source collaboration tools to complete the assignment.

Information Security Threats and Policies in Europe

CASE STUDY

The IT sector is one of the key drivers of the European economy. It has been estimated that 60 percent of Europeans use the Internet regularly. Additionally, 87 percent own or have access to mobile phones. In 2009, the European broadband market was the largest in the world. These facts demonstrate the importance of ensuring the security and safe operation of the Internet for the well-being of the European economy. The safety and security of the Internet have been threatened in recent years as Internet-based cyber attacks have become increasingly sophisticated.

In 2007, Estonia suffered a massive cyber attack that affected the government, the banking system, media, and other services. The attack was performed using a variety of techniques, ranging from simple individual ping commands and message flooding to more sophisticated distributed denial-of-service (DDoS) attacks. Hackers coordinated the attack by using a large number of compromised servers organized in a botnet distributed around the world. A botnet is a network of autonomous malicious software agents that are under the control of a bot commander. The network is created by installing malware that exploits the vulnerabilities of Web servers, operating systems, or applications to take control of the infected computers. Once a computer is infected it becomes part of a network of thousands of "zombies," machines that are commanded to carry out the attack.

The cyber attack on Estonia started in late April 2007 and lasted for almost 3 weeks. During this period, vital parts of the Estonian Internet network had to be closed from access from outside the country, causing millions of dollars in economic losses.

At around the same time, Arsys, an important Spanish domain registration company, was also targeted by international hackers. Arsys reported that hackers had stolen codes that were then used to insert links to external servers containing malicious codes in the Web pages of some of its clients.

In 2009, an estimated 10 million computers were infected with the Conflicker worm worldwide. France, the UK, and Germany were among the European countries that suffered the most infections. The French navy had to ground all military planes when it was discovered that its computer network was infected. In the UK, the worm infected computers in the Ministry of Defense, the city of Manchester's city council and police IT network, some hospitals in the city of Sheffield, and other government offices across the country. Computers in the network of the German army were also reported as infected. Once installed on a computer, Conflicker is able to download and install other malware from controlled Web sites, and thus infected computers could be under full control of the hackers.

More recently, a sophisticated malware threat targeting industrial systems was detected in Germany, Norway, China, Iran, India, Indonesia, and other countries. The malware, known as Stuxnet, infected Windows PCs running the Supervisory Control and Data Acquisition (SCADA) control system from the German company Siemens. Stuxnet was propagated via USB devices. Experts estimated that up to 1,000 machines were infected on a daily basis at the peak of the infection. The malware, hidden in shortcuts to executable programs (files with extension .lnk), was executed automatically when the content of an infected USB drive was displayed. Employing this same technique, the worm was capable of installing other malware. Initially, security experts disclosed that Stuxnet was designed to steal industrial secrets from SIMATIC WinCC, a visualization and control software system from Siemens. However, data gathered later by other experts indicates that the worm was actually looking for some specific programmable logic controller (PLC) devices used in a specific industrial plant, a fact that points to the possibility that the malware was part of a well-planned act of sabotage. Even though none of the sites infected with Stuxnet suffered physical damage, the significance that such a sophisticated threat represents to the industrial resources in Europe and other parts of the world cannot be underestimated.

As of 2001, EU member states had independent groups of experts that were responsible for responding to incidents in information security. These groups lacked coordination and did not exchange much information. To overcome this, in 2004 the European Commission established the European Network and Information Security Agency (ENISA) with the goal of coordinating efforts to prevent and respond

more effectively to potentially more harmful security threats. ENISA's main objectives are to secure Europe's information infrastructure, promote security standards, and educate the general public about security issues.

ENISA organized the first pan-European Critical Information Infrastructure Protection (CIIP) exercise, which took place in November 2010. This exercise tested the efficiency of procedures and communication links between member states in case an incident were to occur that would affect the normal operation of the Internet. ENISA acts as a facilitator and information broker for the Computer Emergency Response Teams (CERT), working with the public and private sectors of most EU member states.

The European Commission has recently launched the Digital Agenda for Europe. The goal of this initiative is to define the key role that information and communication technologies will play in 2020. The initiative calls for a single, open European digital market. Another goal is that broadband speeds of 30Mbps be available to all European citizens by 2020. In terms of security, the initiative is considering the implementation of measures to protect privacy and the establishment of a well-functioning network of CERT to prevent cybercrime and respond effectively to cyber attacks.

Sources: "Digital Agenda for Europe," European Commission, August 2010, http://ec.europa.eu/information_society/ digitalagenda/index_en.htm, accessed October 20, 2010; "The Cyber Raiders Hitting Estonia," BBC News, May 17, 2007, http://news.bbc.co.uk/2/hi/europe/6665195.stm, accessed November 17, 2010; Robert McMillan, "Estonia Ready for the Next Cyberattack," Computerworld, April 7, 2010, www.computerworld.com/s/article/9174923/Estonia_readies_for_the_next_cyber attack, accessed November 17, 2010; "Another Cyber Attack Hits Europe," Internet Business Law Services, June 18, 2007, www.ibls.com/internet_law_news_portal_view.aspx?id=1782&s=latestnews, accessed November 17, 2010; "New Cyber Attack Hits Norway," Views and News from Norway, August 30, 2010, www.newsinenglish.no/2010/08/30/new-cyber-attacks-hit-norway, accessed November 17, 2010; Gregg Keiser, "Is Stuxnet the 'Best' Malware Ever?" Computerworld, September 16, 2010; Robert McMillan, "Was Stuxnet Built to Attack Iran's Nuclear Program," Computerworld, September 21, 2010, www.computerworld.com/s/article/9186920/Was_Stuxnet_built_to_attack_Iran_s_nuclear_program_, accessed November 17, 2010; Ellen Messmer, "Downadup/Conflicker Worm. When Will the Next Shoe Fall?" Network World, January 23 2009, www.networkworld.com/news/2009/012309-downadup-conflicker-worm.html?hpg1=bn, accessed November 17, 2010; Erik Larkin, "Protecting Against the Rampant Conflicker Worm," *PCWorld*, January 16, 2009; "War in the Fifth Domain," *Economist*, July 1, 2010, www.economist.com/node/16478792, accessed November 17, 2010.

CASE STUDY QUESTIONS

8-13 What is a botnet?

8-14 Describe some of the main points of the Digital Agenda for Europe.

8-15 Explain how a cyber attack can be carried out.

8-16 Describe some of the weaknesses exploited by malware.

Case contributed by Daniel Ortiz-Arroyo, Aalborg University.

MyMISLab

Go to **mymislab.com** for the following Assisted-graded writing questions:

8-17 Describe three spoofing tactics employed in identity theft using information systems.

8-18 Describe four reasons why mobile devices used in business are difficult to secure.

Chapter 8 References

"Devastating Downtime: The Surprising Cost of Human Error and Unforeseen Events." Focus Research (October 2010).

Boulton, Clint."American Airlines Outage Likely Caused by Software Quality Issues." *Wall Street Journal* (April 17, 2013).

Boyle, Randall J. and Raymond R. Panko, Raymond R. Corporate *Computer Security 4/e*. Upper Saddle River NJ: Prentice-Hall (2015).

Breedon, John II. "Trojans Horses Gain Inside Track as Top Form of Malware." GCN (May 6, 2013).

Cavusoglu, Huseyin, Birendra Mishra, and Srinivasan Raghunathan. "A Model for Evaluating IT Security Investments." *Communications of the ACM* 47, No. 7 (July 2004).

Crossman, Penny. "DDoS Attacks Are Still Happening-and Getting Bigger." *Information Management* (July 29, 2014).

Dey, Debabrata, Atanu Lahiri, and Guoying Zhang. "Quality Competition and Market Segmentation in the Security Software Market. *MIS Quarterly* 38, No. 2 (June 2014).

Donohue, Brian. "Malware C&C Servers Found in 184 Countries." ThreatPost.com (August 2, 2013).

Galbreth, Michael R. and Mikhael Shor. "The Impact of Malicious Agents on the Enterprise Software Industry." *MIS Quarterly* 34, No. 3 (September 2010).

Grossman, Lev. "The Code." Time(July 21, 2014).

Hui , Kai Lung, Wendy Hui and Wei T. Yue. "Information Security Outsourcing with System Interdependency and Mandatory Security Requirement." *Journal of Management Information Systems* 29, No. 3 (Winter 2013).

Javelin Strategy & Research."2014 Identity Fraud Study." (2014).

Kaplan, James, Chris Rezek, and Kara Sprague."Protecting Information in the Cloud." *McKinsey Quarterly* (January 2013).

Karlovsky, Brian. "FireEye Names Malware's Favorite Targets, Sources." *Australian Reseller News* (March 2, 2014).

McAfee." Mobile Malware Growth Continuing in 2013." (February 21, 2013).

Osterman Research. "The Risks of Social Media and What Can Be Done to Manage Them. Commvault (June 2011).

Panda Security. "Annual Report Pandalabs 2013." (2014).

Perez, Sarah, "AY Media-Owned Blogging Piatform Typepad Enters Day 5 of On-And-Off DDoS Attacks." Techcrunch.com (April 21, 2014).

Ponemon Institute. "2013 Cost of Cybercrime Study: United States (October 2013).

Ponemon Institute. "2014 Cost of Data Breach Study: United States (2014).

Reisinger, Don. "Android Security Remains a Glaring Problem: 10 Reasons Why." *eWeek* (March 2, 2014).

Ribeiro, John. "Hacker group targets Skype social media accounts," *Computer world* (January 2, 2014).

Sadeh, Norman M. "Phish Isn't Spam." *Information Week* (June 25, 2012).

Scharr, Jill. "Fake Instagram 'Image Viewers' Are Latest Malware Fad." Tom's Guide (May 8, 2014).

Schwartz, Matthew J. "Android Trojan Looks, Acts Like Windows Malware." *Information Week* (June 7, 2013).

Sengupta, Somini. "Machines that Know You without Using a Password." New York Times (September 10, 2013).

Solutionary. "Solutionary Security Engineering Research Team Unveils Annual Global Threat Intelligence Report ." (March 12, 2013).

Spears. Janine L. and Henri Barki. "User Participation in Information Systems Security Risk Management." *MIS Quarterly* 34, No. 3 (September 2010).

Stallings, William H. and Lawrie Brown. *Computer Security: Principles and Practice 3/e*. Upper Saddle River, NJ: Prentice Hall (2015).

Symantec. "State of Mobility Global Results 2013." (2013).

_____. "Symantec Internet Security Threat Report." (2014).

Temizkan , Orcun, Ram L. Kumar , Sungjune Park and Chandrasekar Subramaniam. "Patch Release Behaviors of Software Vendors in Response to Vulnerabilities: An Empirical Analysis . " *Journal of Management Information Systems* 28, No. 4 (Spring 2012).

Vance , Anthony, Paul Benjamin Lowry and Dennis Eggett. "Using Accountability to Reduce Access Policy Violations in Information Systems ." *Journal of Management Information Systems* Volume 29 Number 4 Spring 2013.

Yadron, Danny. "Companies Wrestle with the Cost of Cybersecurity." *Wall Street Journal* (February 25, 2014).

Yan Chen ; K. Ram Ramamurthy and Kuang-Wei Wen." Organizations' Information Security Policy Compliance: Stick or Carrot Approach?" *Journal of Management Information Systems* 29, No. 3 (Winter 2013).

Zhao, Xia, Ling Xue and Andrew B. Whinston." Managing Interdependent Information Security Risks: Cyberinsurance, Managed Security Services, and Risk Pooling Arrangements ." *Journal of Management Information Systems* 30, No. 1 (Summer 2013).

376