PART II

Information Technology Infrastructure

CHAPTER 5
IT Infrastructure and Emerging Technologies

CHAPTER 6
Databases and Information Management

CHAPTER 7
Telecommunications, the Internet, and Wireless Technology

CHAPTER 8
Securing Information Systems

PART GOAL

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 organization? How are these technologies likely to change in the future? What technologies and procedures are required to ensure that systems are reliable and secure?

CHAPTER 5

IT Infrastructure and Emerging Technologies

LEARNING OBJECTIVES

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

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

OPENING CASE

Desktop Virtualization Pays Off

Dutch company VocaLink designs and delivers domestic and international automated payments systems and ATM switching solutions. VocaLink's switching platform connects more than 65 000 automated teller machines (ATMs), the world's busiest network, while its payments platform processes more than 3.5 billion payments per year, totalling more than $668 million daily in transactions. In addition, in Great Britain, more than 90% of salaries, 70% of bills, and almost all state benefits are processed through VocaLink. These are very sophisticated systems, and VocaLink partners with outsourcers, primarily in India, to develop the software code for its systems.

Outsourcing means contracting with others to perform some operations that could be done in-house by the company itself. As you can imagine, this creates a number of challenges, especially when the outsourcer is located half way around the world and there are cultural differences. But for VocaLink, the problem that arose while working with their outsourcer was systems integration. "Until about 18 months ago, our Indian developers would work on the code on their own servers," recalls IT director Nick Masterson-Jones. "They would then send that code back over to the UK, and we would have to spend time integrating it into our existing systems and testing those integrations." VocaLink turned to desktop virtualization to solve this issue. According to Masterson-Jones, "We wanted the developers in India to be able to write code as if it was being written here in the UK. That was the genesis of our virtualization project."

Using a combination of Citrix and VMWare desktop virtualization software, combined with blade servers and a storage area network (SAN), VocaLink created 200 virtual desktops that could be used by its Indian developers. This allowed the Indian programmers to work directly on VocaLink's core systems, eliminating the steps of integration and testing from the development cycle. This improved agility and helped their two outsourcing partners, Perot Systems and Wipro, to better collaborate.

There are unexpected consequences of most software implementation projects, such as VocaLink's desktop virtualization project. According to Masterson-Jones, "We liked the way it took the capital cost out of deploying new projects. Previously, new projects always began with buying more servers." Using desktop virtualization, projects can now be started using existing unused capacity on the company's hardware. This led VocaLink to apply virtualization to its existing corporate systems, which centralized IT support, reducing problem resolution time.

While some employees were disinclined to part with their old machines, the new thin client devices that work with the virtualization model have become the norm. Security improved, too. Masterson-Jones explained: "You cannot cut-and-paste code out of the virtual machine. Now nothing leaves our data centre." VocaLink even used its virtualization to permit employees stranded at home by a snowstorm in 2009 to work by offering remote access to the virtual desktops. Ninety-eight percent of employees reported to work from their homes. Masterson-Jones also said: "It has been exciting to measure what we've been saving in energy costs. And (making IT energy efficient) may not bring competitive advantage, but it is something you should be doing anyway." So virtualization helps with green computing, too.

Sources: Pete Swabey, "Desktop Virtualization has Helped Payments Processing Provider VocaLink Unify its Global Development Operations and has Brought Other Benefits Rather Closer to Home," *Information Age*, www.information-age.com, accessed March 1, 2011; Kenny MacIver, "The Changing Shape of the Corporate Data Centre," July 2011, http://www.i-cio.com/features/july-2011/the-changing-shape-of-the-corporate-data-center, accessed February 21, 2013; and VocaLink, www.vocalink.com, accessed February 21, 2013.

VocaLink is a world leader in automated payment and ATM systems, but its systems development efforts were hampered by the challenges of systems integration. VocaLink's IT management felt the best solution was to invest in desktop virtualization to permit its outsourced developers to develop more efficiently and effectively and test code directly on its systems. This has resulted in cost efficiencies, effectiveness, and energy savings.

The chapter-opening diagram calls attention to important points raised by this case and this chapter. Management realized that in order to improve its software development processes, it had to solve the problem of integration. The IT infrastructure investments it made in desktop virtualization had to support VocaLink's business goals and contribute to improving its performance. Other outcomes included reducing costs and the "green" goal of reducing power consumption.

By implementing desktop virtualization with blade servers, VocaLink was able to reduce wasted computer resources not used for processing, use existing resources more efficiently, and cut costs and power consumption. New software tools make it much easier to develop new applications and services. VocaLink's IT infrastructure is easier to manage, with centralized IT support, and is capable of permitting workers to telecommute when needed. This case shows that the right hardware and software investments not only improve business performance but can also contribute to important social goals, such as conservation of power and materials.

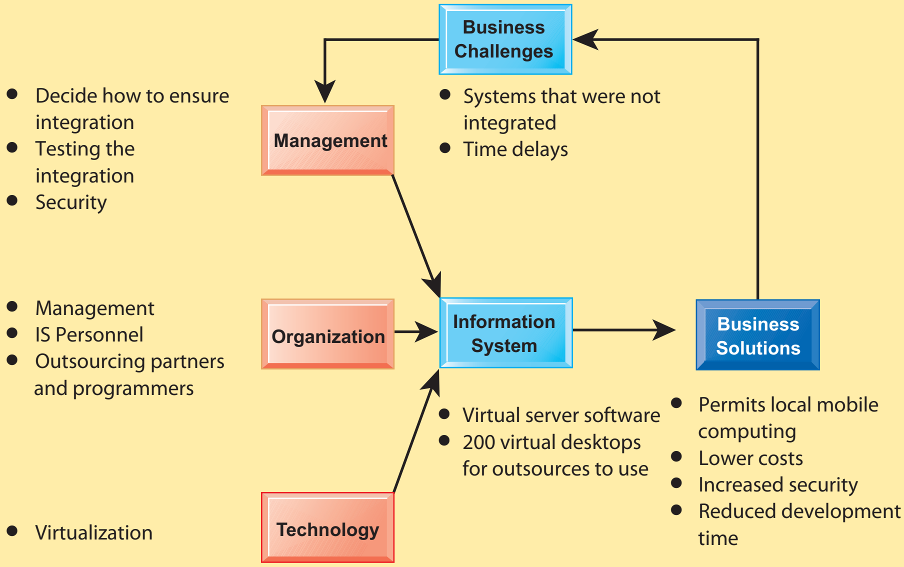

5.1 IT Infrastructure

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 investments 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).

In 2013, supplying firms worldwide with IT infrastructure (hardware and software) was estimated as a $3.78 trillion industry when telecommunications, networking equipment, and telecommunications services (Internet, telephone, and data transmission) are included (Gartner Group, January 3, 2013). This does not include IT and related business process consulting services, which add another $407 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 more than 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, mobile handheld devices, 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.

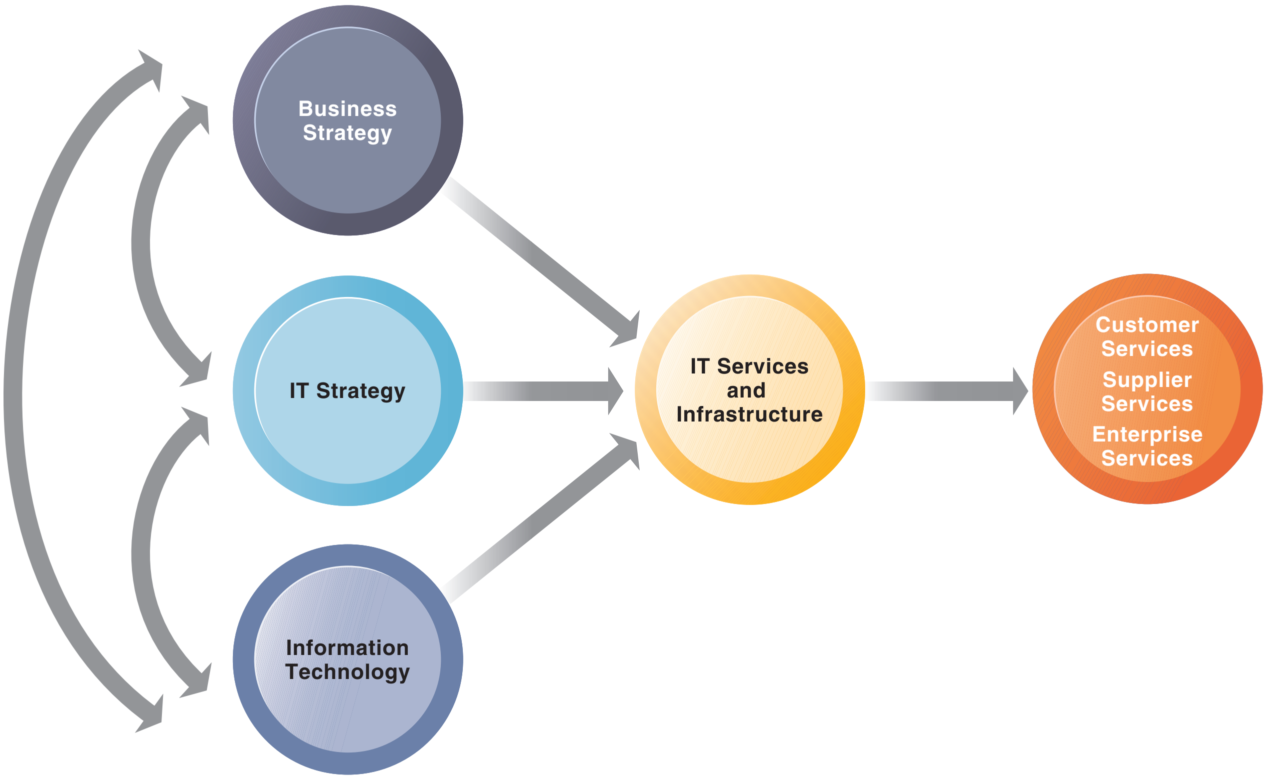

The services a firm is capable of providing to its customers, suppliers, and employees are a direct function of its IT infrastructure. Ideally, this infrastructure should support the firm's business and information systems strategy. New information technologies have a powerful impact on business and 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 IT infrastructure, coordinate IT services for business units, manage accounting for 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, how, and by whom
- IT education services that provide training in system use to employees and provide managers with 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 $1000 personal computer operating at 3 gigahertz 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 more than 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 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 centre), 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 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 2012, there were an estimated 1.2 billion PCs in the world, and 300 million new PCs are sold each year. Ninety percent are thought to run a version of Windows, and ten percent run a Macintosh OS. Nonetheless, the newer tablet market is eroding the PC market, and projections for 2013 and beyond fluctuate widely, depending on whether these projections include tablets in the PC category. The Wintel dominance as a computing platform is receding as iPhone and Android device sales increase. In 2012, Intel lost more than ten percent of its annual sales due to other chip manufacturers capitalizing on the consumer move to tablets and other mobile devices. Nearly one billion people worldwide own smartphones, and most of these users access the Internet with their mobile devices.

Mainframe
Minicomputers
Wintel PC

FIGURE 5.2 Eras in IT infrastructure evolution.

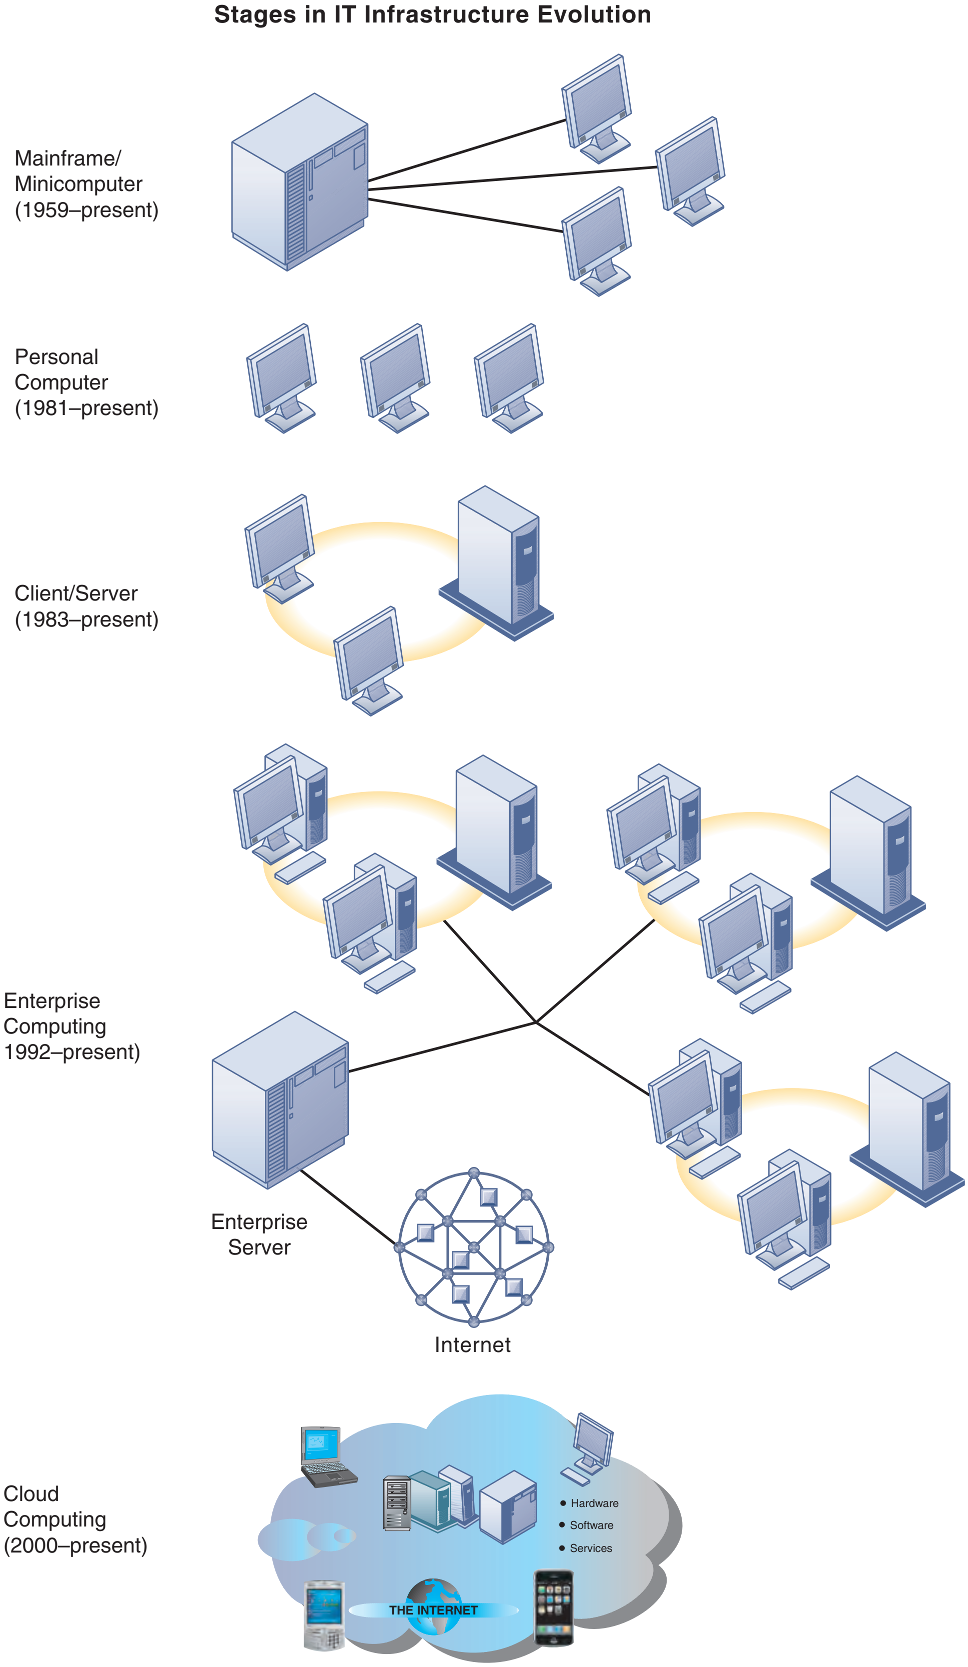

Illustrated here are the typical computing configurations characterizing each of the five eras of IT infrastructure evolution.

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, while 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 are typically 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*. While 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 available to firms.

Client/server computing
Clients
Server
Multitiered (N-tier) client/server architectures
Web server
Application server

FIGURE 5.3 A multitiered client/server network (N-tier).

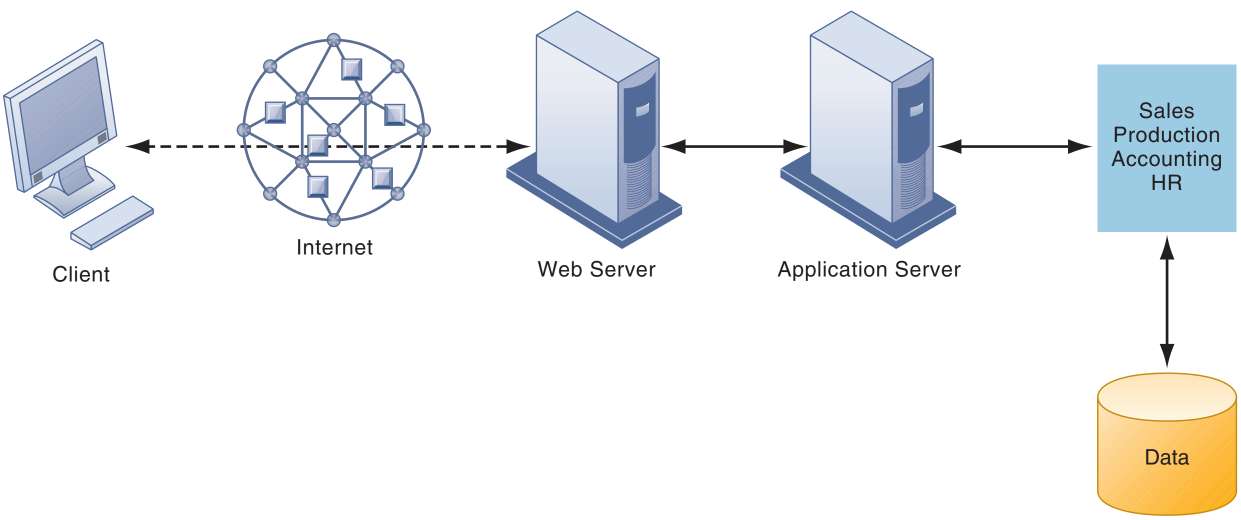

In a multitiered client/server network, client requests for service are handled by different levels of servers.

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. After 1995, as the Internet developed into a trusted communications environment, 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 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 $111 billion on public cloud services in 2012 and an estimated $211 billion by the end of 2016 (Gartner, 2012).

Hundreds of thousands of computers are located in cloud data centres, where they can be accessed by desktop computers, laptop computers, tablets, entertainment centres, smartphones, and other client machines linked to the Internet, with both personal and corporate computing increasingly moving to mobile platforms. IBM, HP, Dell, and Amazon operate huge, scalable cloud computing centres that provide computing power, data storage, and high-speed Internet connections to firms that want to maintain their IT infrastructures remotely. Software 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 modules on the Companion Website for this text 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 us look at the most important developments.

Moore's Law and Microprocessing Power In 1965, Gordon Moore, the founder of Intel, the world's largest chip manufacturer, 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.

Windows
Cloud computing
Moore's Law

FIGURE 5.4 Moore's law and microprocessor performance.

Packing over 2 billion transistors into a tiny microprocessor has exponentially increased processing power. Processing power has increased to more than 128 000 MIPS (2.6 billion instructions per second).

Source: Authors' estimate.

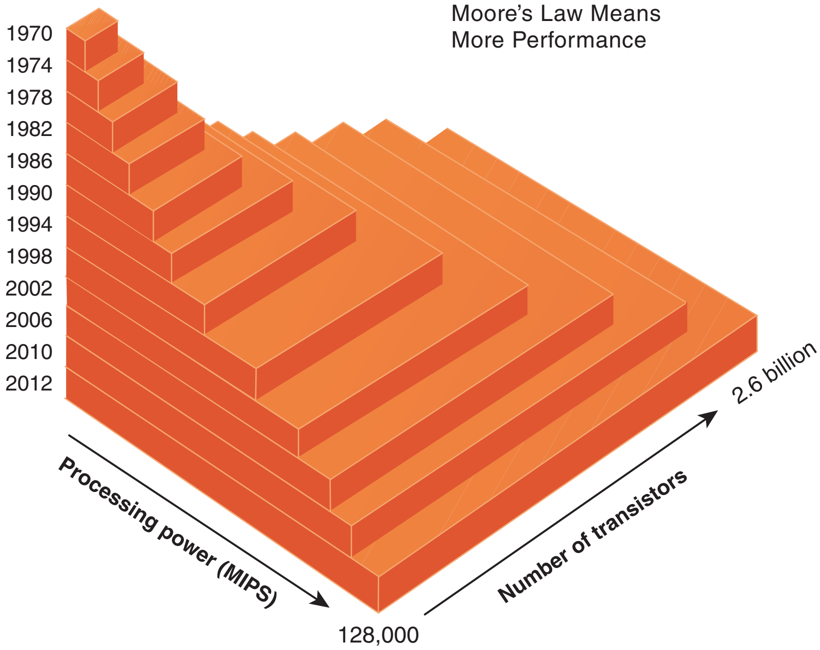

This law would later be interpreted in multiple ways. There are at least three variations of Moore's Law, none of which Moore ever stated: (1) the power of microprocessors doubles every 18 months; (2) computing power doubles every 18 months; and (3) the price of computing falls by half every 18 months.

Figure 5.4 illustrates the relationship between the 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

FIGURE 5.5 Falling cost of chips.

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

Source: Authors' estimate.

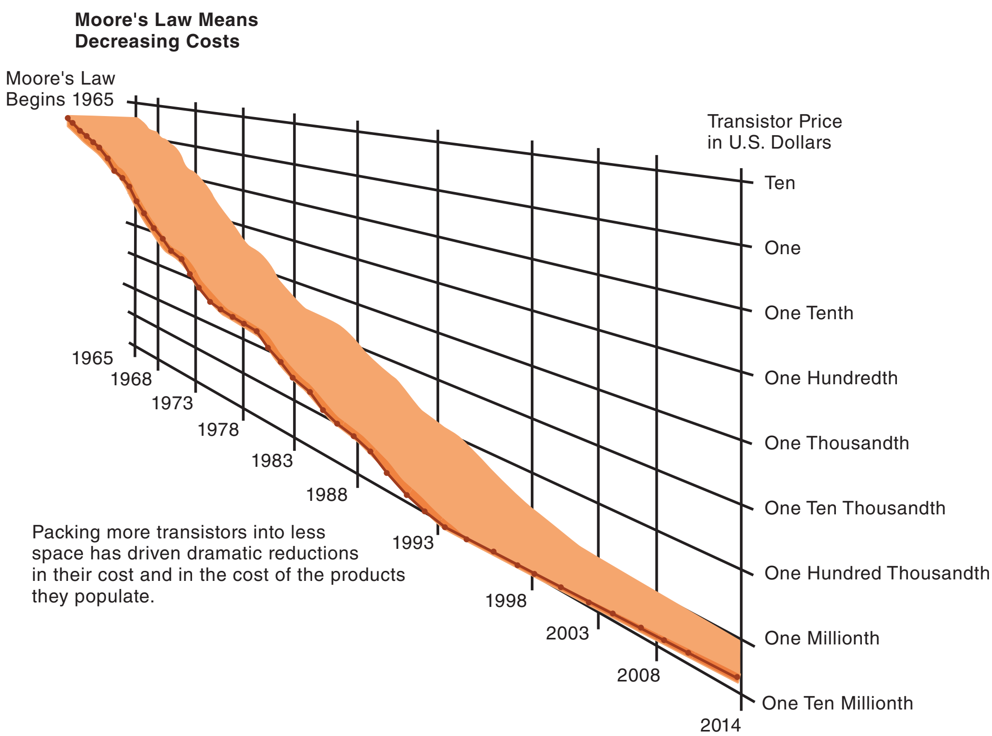

computing power. For instance, in 2012, you can buy an Intel i7 quad-core processor on Amazon for about $355, and you will be purchasing a chip with 2.5 billion transistors, which works out to about one ten-millionth of a dollar per transistor.

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. Nanotubes are tiny tubes about 10 000 times thinner than a human hair. They consist of rolled-up sheets of carbon hexagons and are very powerful conductors of electrical current; they have potential uses as minuscule wires or in ultrasmall electronic devices. Chip manufacturers are trying to develop a manufacturing process that could produce nanotube processors economically. IBM has just started making microprocessors in a production setting using this technology. In October 2012, IBM made transistors—the basic components of electronic computing—from nanometer-sized tubes of carbon and put 10 000 of them on top of a silicon chip using mainstream manufacturing processes (Takahashi, 2012).

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 5.6 shows, from 1950 to the present, the number of megabytes that could be stored on magnetic media for $1 roughly doubles every 15 months. In 2013, a 500 gigabyte hard disk drive sells at retail for about $60.

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

Nanotechnology

FIGURE 5.6 The amount of storage per dollar rises exponentially, 1950–2012.

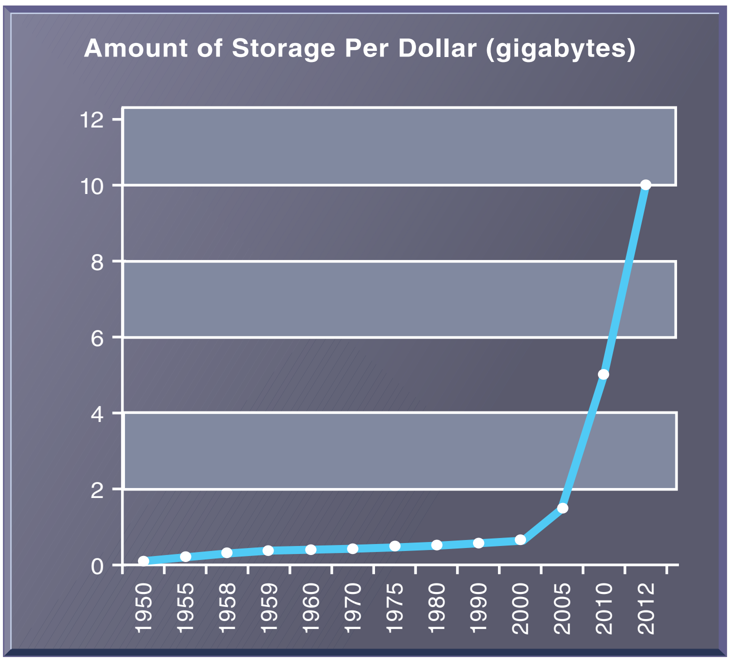

Since the first magnetic storage device was used in 1955, the amount of storage a dollar buys has risen exponentially, doubling the amount of digital storage for each dollar expended every 15 months on average.

Source: Authors' estimates.

FIGURE 5.7 Exponential declines in internet communications costs.

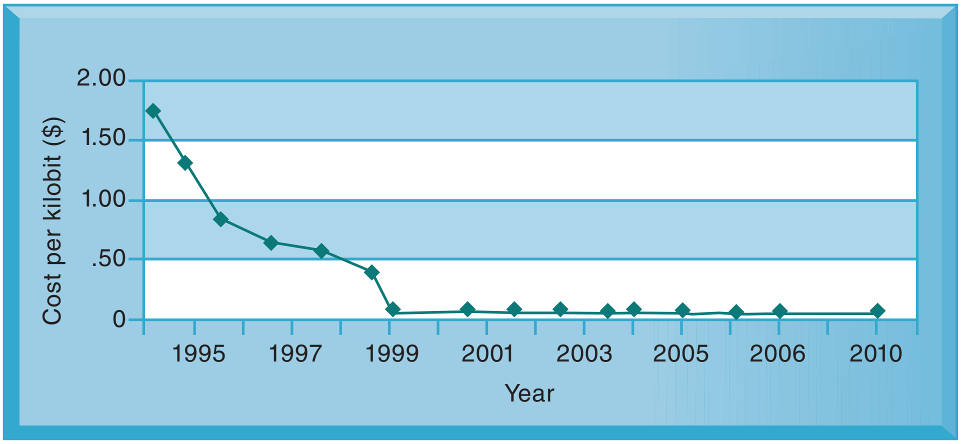

One reason for the growth in the Internet population is the rapid decline in Internet connection and overall communication costs. The cost per kilobit of Internet access has fallen exponentially since 1995. Digital subscriber line (DSL) and cable modems now deliver a kilobit of communication for a retail price of around 2 cents.

Source: Authors.

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. An estimated 2.4 billion people worldwide, or more than one-third of the earth's population, now have Internet access (Internet World Stats, 2012). Figure 5.7 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.

To take advantage of the business value associated with the Internet, firms must greatly expand their Internet connections, including wireless connectivity, and 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 Current 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 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

Technology standards

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 US 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 (1991–1993)	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.

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.

5.2 INFRASTRUCTURE COMPONENTS

IT infrastructure today is composed of seven major components. Figure 5.8 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 $456 billion on computer hardware in 2013, including servers and client devices. The server market uses mostly Intel or AMD processors in the form of blade servers in racks, but also includes Sun SPARC microprocessors and IBM chips specially designed for server use. **Blade servers** are computers consisting of only a circuit board with processors, memory, and network connections that are stored

Blade servers

FIGURE 5.8 The IT infrastructure ecosystem.

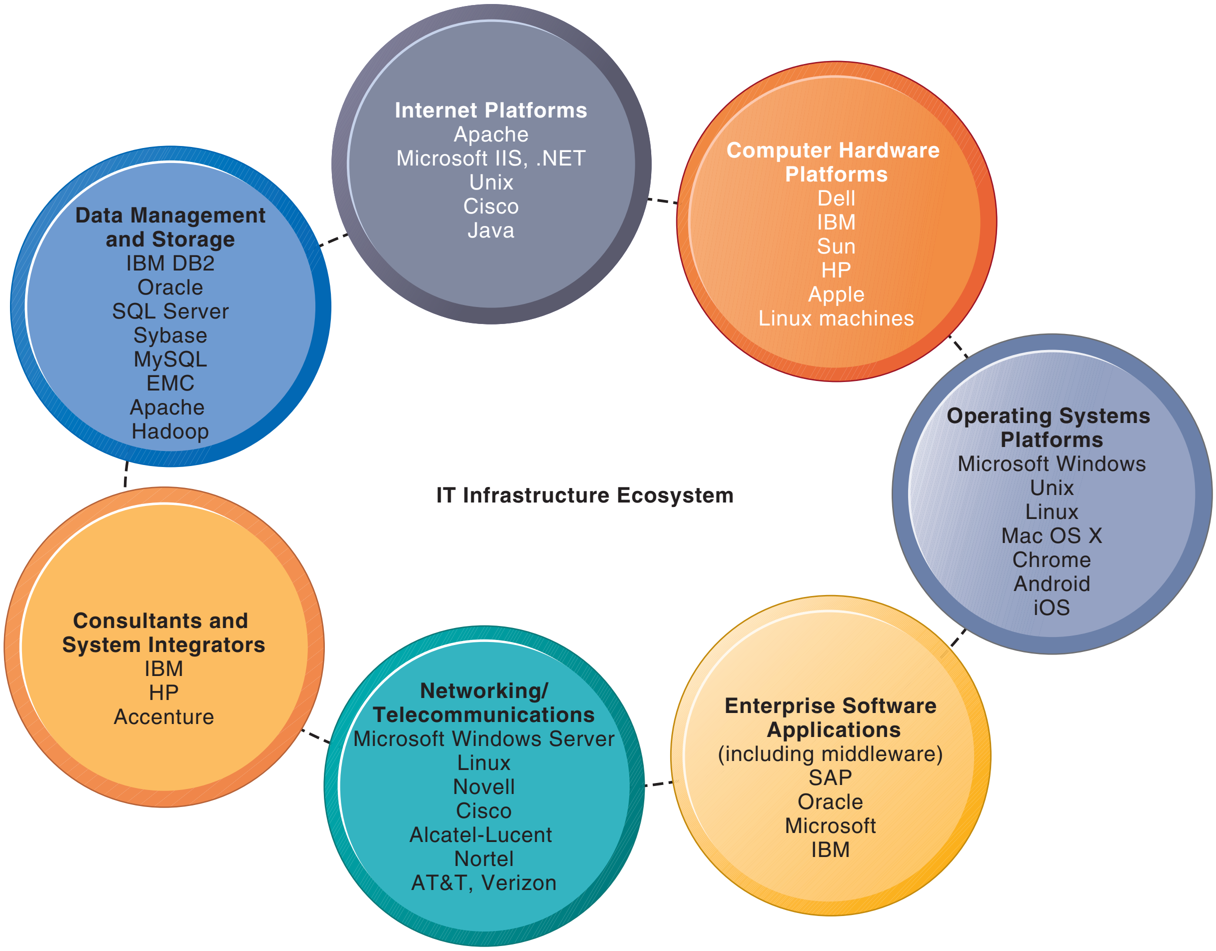

There are seven major components that must be coordinated to provide the firm with a coherent IT infrastructure. Listed here are the major technologies and suppliers for each component.

in racks. They take up less space than traditional box-based servers. Secondary storage may be provided by a hard drive in each blade server or by external mass-storage drives.

The marketplace for computer hardware has increasingly become concentrated in a few top firms, including IBM, HP, Dell, and Sun Microsystems (acquired by Oracle), and three chip producers: Intel, AMD, and IBM. The industry has collectively settled on Intel as the standard processor for business computing, with major exceptions in the server market for Unix and Linux machines, which might use Sun or IBM processors.

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 centres. The mainframe is still the digital workhorse for banking and telecommunications networks. 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

Unix

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 and cloud-connected computers.

Google's **Chrome OS** provides a lightweight operating system for cloud computing using netbooks that are manufactured by Google partners such as Samsung and Acer. Programs are not stored on the user's PC 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 is 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 permits use of a mouse and keyboard.

Enterprise Software Applications

In 2013, firms worldwide are expected to spend about $306 billion 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 firm-wide 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, with a staggering 1.3 zettabytes to be created in 2016 alone (Moscaritolo, 2012). All the tweets, blogs, videos, e-mails, and Facebook postings, as well as traditional corporate data, add up to about 38 million DVDs

Linux
Operating system
Chrome OS
Android
iOS
Multitouch
Windows 8

per hour, or one billion terabyte drives (each of which holds one billion bytes) (Moscaritolo, 2012).

With the amount of new digital information in the world growing so rapidly, the market for digital data storage devices has been growing 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 $416 billion for telecommunications equipment in 2013 and another $1.73 billion on telecommunications services (Gartner, 2012). 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 typically 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, Canada's Nortel, 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 Rogers Communications and Bell Canada. 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 Expression 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 a firm's Internet platform in greater detail.

Consulting and Systems 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,

Storage area networks (SANs)
Web hosting service

Infosys, and Wipro Technologies, as well as major accounting firms, such as Deloitte Canada.

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 today's infrastructure.

5.3 CONTEMPORARY HARDWARE PLATFORM TRENDS

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 will look at eight hardware trends: the mobile digital platform, consumerization of IT, grid computing, virtualization, cloud computing, green computing, high-performance/power-saving processors, and autonomic computing.

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 ebook readers such as Amazon's Kindle or Chapters' Kobo 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 use smartphone applications that drill down into vehicle sales information, financial performance, manufacturing metrics, and project management status.

Consumerization of IT and BYOD

The popularity, ease of use, and variety 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, such as Google and Yahoo search, Gmail, Google Apps, Dropbox (see Chapter 2), and even Facebook and Twitter, as well.

Consumerization of IT is forcing businesses, especially large enterprises, to rethink the way they obtain and manage information technology equipment and services. Historically, at least in large firms, the central 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

Legacy systems
Tablet computers
Consumerization of IT

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 to make sure they serve the needs of the business. The Window on Management box explores some of the management challenges created by BYOD and IT consumerization.

Grid Computing

Grid computing involves connecting geographically remote computers into a single network to create a virtual supercomputer by combining the computational power of all computers on the grid. Grid computing takes advantage of the fact that most computers use their central processing units on average only 25 percent of the time for the work they have been assigned, leaving these idle resources available for other processing tasks. Grid computing was impossible until high-speed Internet connections enabled firms to connect remote machines economically and move enormous quantities of data. Grid computing requires software programs to control and allocate resources on the grid.

The business case for using grid computing involves cost savings, speed of computation, and agility. For example, Royal Dutch/Shell Group is using a scalable grid computing platform that improves the accuracy and speed of its scientific modelling applications to find the best oil reservoirs. This platform, which links 1024 IBM servers running Linux, in effect creates one of the largest commercial Linux supercomputers in the world. The grid adjusts to accommodate the fluctuating data volumes that are typical in this seasonal business. Royal Dutch/Shell Group claims the grid has enabled the company to cut processing time for seismic data while improving output quality and helping its scientists pinpoint problems in finding new oil supplies.

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 so that it acts like many different machines. 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 or grid computing. Virtualization makes it possible for a company to handle its computer processing and storage using computing resources housed in remote locations. VMware is the leading virtualization software vendor for Windows and Linux servers.

By providing the ability to host multiple systems on a single physical machine, virtualization helps organizations increase equipment utilization rates, conserving data centre space and energy. Most servers run at just 15–20 percent of capacity, and virtualization can boost server utilization rates to 70 percent or higher. Higher utilization rates translate into fewer computers required to process the same amount of work, as illustrated by VocaLink's efforts to use virtualization discussed in the chapter-opening case. 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, as is the case with cloud computing.

Cloud Computing

Grid computing
Virtualization

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-

WINDOW ON MANAGEMENT

Should you Use your Iphone for Work?

Look around. On the street, at restaurants, sports events, and stores, you will find many people using their smartphones. And many people are starting to use these devices on the job as well. According to a Juniper Research report, the number of employees who use personal devices at work will increase to 350 million by 2014. About 150 million people use their own mobile phones and tablets at the workplace today. Forty-one percent of Canadians used smartphones in 2011, compared to twenty-eight percent worldwide. Interestingly, the largest trend in the Canadian market was in young men between the ages of 22 and 30.

If almost everyone has a personal smartphone, why not use it for work? Employees using their own smartphones permit companies to enjoy all of the same benefits of a mobile workforce without spending their own money on these devices, but IT departments need to overcome several logistical hurdles to make that vision a reality. Using personal devices for business poses difficult problems for companies, including security, inventory management, support, integrating mobile devices into pre-existing IT functions and systems, and measuring return on investment. In other words, it's not that simple.

A significant amount of corporate IT resources are dedicated to managing and maintaining a large number of devices within an 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 the 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.

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 phones, such as switching cellular carriers, changing their phone numbers, or buying new mobile devices, 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 can be a nightmare. In addition, firms need to specify what corporate uses the individual may use on the phone and when the company will or will not reimburse employees for corporate use. In addition, employee privacy is an issue that must be considered (see Chapter 4). Establishing corporate policies governing smartphone use is a necessity for companies today.

With the variety of phones and operating systems available, providing adequate technical support for every employee can be difficult. When employees are not able to access critical data or encounter other problems with their mobile devices, they 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 technical support much easier. Mobility introduces a new layer of variety and complexity to technical support that companies need to be prepared to handle.

A firm's software development teams can benefit from having one person specifically focused on ensuring that new applications will be easily usable and useful on smartphones. Many companies are integrating these "mobility experts" into core IT functions and software development. Unless applications and software can be used on mobile devices to connect to the firm's existing IT platform and company-wide customer relationship management (CRM), supply chain management (SCM), and enterprise resource planning (ERP) systems, a business smartphone is just a phone; mobility experts can help a company leverage mobility more effectively.

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 is not 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 will find a detailed discussion of mobile security issues in Chapter 8.

A number of software products have emerged to help companies manage diverse mobile platforms. Sybase Afaria, Trellia, Microsoft Systems Center Device Manager, and Odyssey Software Athena have capabilities for configuring devices remotely, enforcing different sets of policies for different users and devices, and managing applications running on all of them.

Novo Nordisk, headquartered in Denmark, manufactures and markets pharmaceutical products and services throughout the world. Its 2000-member sales force operates in 25 different countries and uses a diverse assortment of mobile phones, smartphones, and mobile handhelds. To manage all of these devices centrally, Novo Nordisk implemented Sybase Afaria. Using Afaria, the company's internal IT department can deploy new applications to mobile devices quickly and without extensive end-user interaction. A new mobile phone user just needs to answer "yes" to Novo Nordisk's configuration process and the installation happens automatically. Afaria also has features for enabling individual countries or regions to provide their own local support, a necessity since each Novo Nordisk market has it own data connections, policies, and requirements.

Another approach to mobile device management is virtualization. Companies can install software such as Citrix Systems XenDesktop that runs Windows desktops and individual applications on any device, regardless of the operating system. Employees then use that software to access their entire desktop on their smartphones and mobile handhelds and are thus able to use the same programs on the road that they use in the office. The virtualization software has built-in security features that allow corporations to prohibit saving data on local devices, to encrypt all corporate data without touching employees' personal applications and data, and to remotely erase data in the event of a security breach. India's Anthem Group, a leading provider of pharmaceuticals and biotechnology services, implemented Citrix XenDesktop to enable employees to remotely access data because this virtualization solution runs on all devices with minimal bandwidth consumption.

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 firm's long-term plans, companies will end up wasting their money on unnecessary devices and programs. One of the biggest worries that managers have about mobility is the difficulty of measuring return on investment. Many workers swear by their mobile devices, and the benefits are too significant to ignore, but quantifying how much money is earned or saved by going mobile can be difficult.

Sources: Boonsri Dickinson, "Security Headaches: BYOD Users Expected to Double by 2014," *Information Week*, August 8, 2012; "Anthem Group Enables Secure Remove Access with Citrix XenDesktop and XenServer," expresscomputeronline.com, accessed July 20, 2012; "So You Want to Use Your IPhone for Work: Uh-Oh," *The Wall Street Journal*, April 25, 2011; Samuel Greengard, "Managing Mobility in the Enterprise," *Baseline*, January 28, 2011; QMI Agency, "Canadians Leading the Way on Smartphone Use, Survey Says," *Toronto Sun*, June 24, 2011; Dell Computer, "Management Madness: How to Manage Mobile Devices," mkting.cio.com, accessed July 18, 2011; and Dell Computer, "Is Your Infrastructure Mobile-Ready?" www.cio.com, accessed May 21, 2011.

CASE STUDY QUESTIONS

1. What are the advantages and disadvantages of allowing employees to use their personal smartphones for work?
2. What management, organization, and technology factors should be addressed when deciding whether to allow employees to use their personal smartphones for work?
3. Allowing employees to use their own smartphones for work will save the company money. Do you agree? Why or why not?

needed basis from any connected device and location. Figure 5.9 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 the amount of resources actually used.

Cloud computing consists of three different types of services:

- **Cloud infrastructure as a service:** 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

FIGURE 5.9 Cloud computing platform.

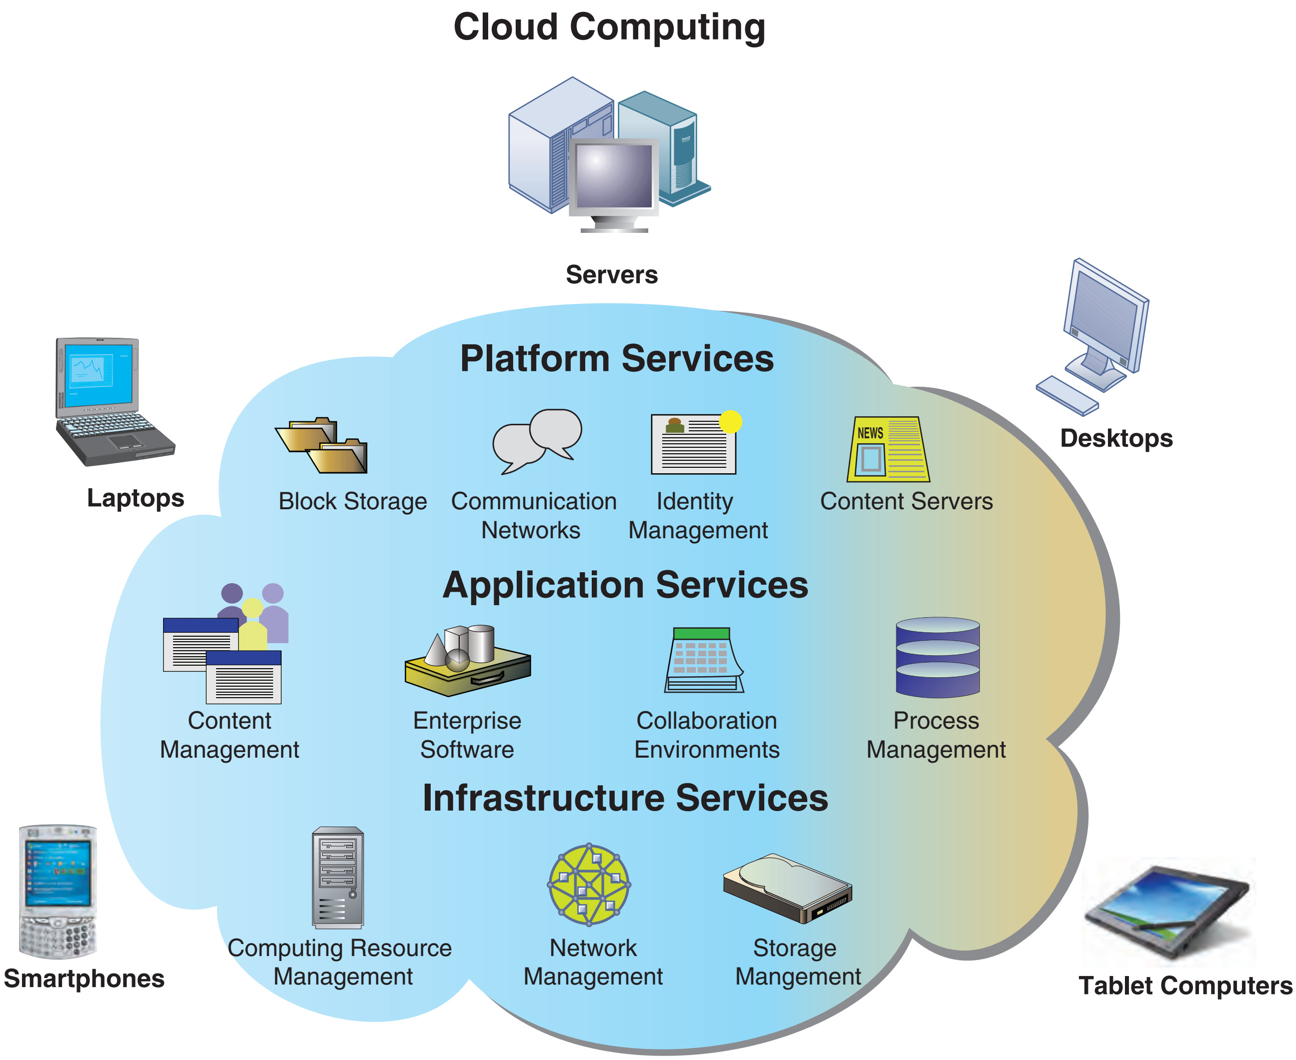

In cloud computing, hardware and software capabilities are a pool of virtualized resources provided over a network, often the Internet. Businesses and employees have access to applications and IT infrastructure anywhere, at any time, and on any device.

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 chapter-ending case study.)

- **Cloud platform as a service:** 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.
- **Cloud software as a service:** 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 Apps also has a pared-down free version. 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 groups. 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 chapter-ending case study.)

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 these services.

In February 2011, IBM opened a $42 million cloud computing centre in Toronto to allow businesses and organizations to develop, host, and test applications on a pay-as-you-go model. IBM said confidential information is protected and kept securely resident in Canada, in accordance with Canadian privacy laws. According to Canada-based IDC analyst Mark Schrutt, the market for cloud delivery capabilities in Canada was expected to grow by 40% over the next three years and would break the $1 billion mark by 2015. The cloud computing centre added to a network of 16 data centres IBM operates across Canada. Among the first clients to take advantage of the IBM Compute Cloud Centre's offerings in Canada was VisionMax, a Canadian-based custom software development firm offering IT solutions.

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 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. Another limitation of cloud computing is that users become dependent on the cloud computing provider, and this may not necessarily be desirable, as discussed in the chapter-ending case. 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 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 centres 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 centres owned by computer hardware vendors. You can find out more about cloud computing online in the Learning Tracks for this chapter.

Public cloud
Private cloud
Utility computing
On-demand computing
Hybrid cloud

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, mobile devices, and networking and communications systems to minimize their 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. In Canada, it is estimated that the energy used for PC and notebook use, not including smartphones or tablets, could power 7 500 SUVS driving around the world (University of Alberta, 2013). It is easy, therefore, to see that cutting power consumption in data centres has become both a serious business and environmental challenge. The Window on Technology box examines this problem.

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 pro-
r cores (which perform the reading and execution of computer instructions) on a
chip. A **multicore processor** is an integrated circuit to which two or more proces-
ores have been attached for enhanced performance, reduced power consumption,
more efficient simultaneous processing of multiple tasks. This technology enables
r more processing engines with reduced power requirements and heat dissipation
rform tasks faster than a resource-hungry chip with a single processing core. Today
vill find PCs with dual-core, quad-core, six-core, and eight-core processors.
tel and other chip manufacturers have developed microprocessors that minimize
r consumption, which is essential for prolonging battery life in small mobile digital

lthough we made the e as per your ctions, but as e design, the ntal line for ox appears the white See page 145 design PDF.

Green computing or green IT
Multicore processor

WINDOW ON TECHNOLOGY

Green Data Centres: Good for Business?

What is too hot to handle? It might very well be your company's data centre, 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 centres 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 and on corporate operating costs.

Companies are now looking to green computing for solutions. The standard for measuring data centre energy efficiency is Power Usage Effectiveness (PUE). This metric is a ratio of the total annual power consumed by a data centre 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 centres has hovered around 2.0. That means the data centre is using twice the amount of electricity that is 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 centre size, the types of servers and their uses, the proficiency of monitoring software, building architecture, and the climate outside the facility. New data centre 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 a firm's IT infrastructure. Toronto's Data Centers Canada (DCC) uses 'free cooling' — a process in which cold air is inducted from outside during the winter months to cool server cabinets. With free cooling, DCC can turn off the air conditioning and chillers altogether, allowing the company to save between $3000 and $4000 each month on electricity costs. They also use 'cold air containment' so that cool air flowing into the cabinets stays contained and does not get absorbed by warmer air. Other energy-saving measures that DCC has been considering include LCD lighting to help minimize cooling costs, UPS battery systems which plug directly into an on-site generator, and 'fly wheel' UPS systems which do not end up in a landfill. Today, DCC's electricity costs are between 30 and 40 percent less than what its customers would pay at a competitor's data centre. Microsoft is also trying to encourage energy-saving software practices by charging business units by the amount of power they consume in the data centre rather than the space they take up on the floor.

Other tools and techniques are also available to make data centres more energy efficient. Some of the world's most prominent firms are leading the way with innovative programs and technologies for reducing power consumption. Google and Microsoft are building data centres that take advantage of hydroelectric power. Hewlett-Packard is working on a series of technologies to reduce the carbon footprint of data centres by 75 percent, along with new software and services to measure energy use and carbon emissions. It reduced its power costs by 20 to 25 percent through consolidation of servers and data centres. Since these companies' technologies and processes are more efficient than most other companies, using their online software services in place of in-house software may also count as a green investment. In April 2011, Facebook publicly posted the specifications for the design of its data centres, including motherboards, power supplies, server chassis, server racks, and battery cabinets, as well as data centre 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 centre'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 centre design was implemented at Facebook's Prineville, Oregon data centre. 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 centre reports its PUE is 1.07, one of the lowest in the world. By using fans and running warmer than average processor temperatures, Google's newest data centres deliver a PUE rating of 1.16. Yahoo's new Lockport, N.Y., data centre 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.

In addition to lowering IT costs, using cloud computing services may save energy as well. Cloud computing centres pack in servers that have been optimized for virtualization and 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 handling the computing for large numbers of companies. Experts note that it is 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 are not deploying their existing technology resources efficiently or using green measurement tools.

Sources: "How Facebook's Data Centre Leads by Example," *CIO Insight*, August 20, 2012; Sam Greengard, "IT Gets Greener," *Baseline*, April 11, 2012; Mark Lafferty, "Maybe It's Easy Being Green," *PC World*, July 24, 2012; Charles Babcock, "Cloud's Thorniest Question: Does It Pay Off?" *Information Week*, June 4, 2012; "New Study: Cloud Computing Can Dramatically Reduce Energy Costs and Carbon Emissions," *AT&T*, July 21, 2011; Joseph F. Kovar, "Data Centre Power Consumption Grows Less Than Expected: Report," *CRN*, August 10, 2011; Kenneth Miller, "The Data Centre Balancing Act," *Information Week*, May 16, 2011; Vivian Wagner, "Green Data Centres Are Where It's At," *TechNewsWorld*, April 20, 2011; Bill Kenealy, "Facebook Reveals Its Server Secrets," *Information Management*, April, 11, 2011; and BlackIron Data, "Toronto Data Centre is Early Adopter of Most Efficient Cooling Technology from Emerson," February 20, 2013, accessed February 21, 2013.

CASE STUDY QUESTIONS

1. What business and social problems are caused by data centre power consumption?
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?

devices. Highly power-efficient microprocessors, such as ARM, Apple's A4, A5, and A6 processors, and Intel's Atom are in netbooks, digital media players, and smartphones. The dual-core A6 processor used in the iPhone 5 is 22% smaller than the A5 and therefore saves battery life over its predecessors (Sakr, 2012).

Autonomic Computing

With large systems encompassing many thousands of networked devices, computer systems have become so complex today that some experts believe they may not be manageable in the future. One approach to this problem is autonomic computing. **Autonomic computing** is an industry-wide effort to develop systems that can configure themselves, optimize and tune themselves, heal themselves when broken, and protect themselves from outside intruders and self-destruction.

You can glimpse a few of these capabilities in desktop systems. For instance, virus and firewall protection software are able to detect viruses on PCs, automatically defeat the viruses, and alert operators. These programs can be updated automatically as the need arises by connecting to an online virus protection service such as McAfee. IBM and other vendors are starting to build autonomic features into products for large systems.

5.4 CONTEMPORARY SOFTWARE PLATFORM TRENDS

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 as many as 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. Open-source tools are being used on netbooks as inexpensive alternatives to Microsoft Office. Major hardware and software vendors, including IBM, HP, Dell, Oracle, and SAP, now offer Linux-compatible versions of their products. You can find out more about the Open Source Definition from the Open Source Initiative and the history of open-source software at the online 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, netbooks, 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 major force in local area networks, Web servers, and high-performance computing work. IBM, HP, Intel, Dell, and Oracle have made Linux a central part of their offerings to corporations.

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.

Autonomic computing
Open-source software

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.

The Java platform has migrated into cell phones, smartphones, automobiles, music players, game machines, and finally, 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. Oracle Corporation estimates that 3 billion devices are running Java, and it is the most popular development platform for mobile devices running the Android operating system (Taft, 2012). 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 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 (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. This is one reason Apple dropped support for Flash on its mobile devices. 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 will also make it easier for Web pages to function across different display devices, including mobile devices and desktops, and it will support the storage of data offline for apps that run over the Web. Web pages will execute more quickly and look like smartphone apps. Although HTML5 is still under development, elements are already being used in a number of Internet tools, including Apple's Safari browser, Google Chrome, and recent versions of the Firefox Web browser. Google's Gmail and Google Reader have adopted parts of the HTML5 standard as well. Web sites listed as "iPad ready" are making extensive use of HTML5 including the Web sites of Saskatoon's FreshnLocal and BC Business Online.

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 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.

Java
Web browser
HTML (hypertext markup language)
HTML5
Web services

TABLE 5.2 Examples of xml.

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

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. While 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 postal 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 a 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.

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 5.10).

XML
eXtensible markup language
Service-oriented architecture (SOA)

FIGURE 5.10 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.

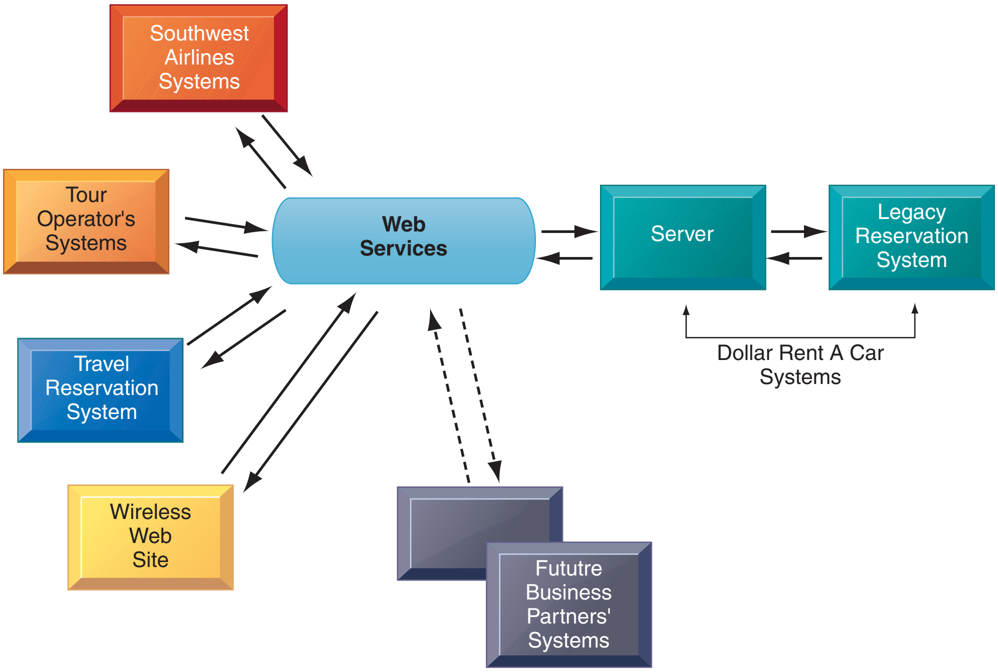

Software Outsourcing and Cloud Services

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 most will purchase or rent at least some of their new software applications from external sources. Figure 5.11 illustrates the rapid growth in external sources of software for U.S. firms. In Canada, as far back as 2010, almost $15 billion of software development was projected to be outsourced to other countries (Kavur, 2010).

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 lower-wage areas of the world. According to industry analysts, spending on offshore IT outsourcing services was approximately $255 billion in 2012 (Gartner, 2012). The largest outsourcing expenditures are to domestic U.S. firms providing middleware, integration services, and other software support often required to operate larger enterprise systems.

Software package
Outsourcing

FIGURE 5.11 Changing sources of firm software.

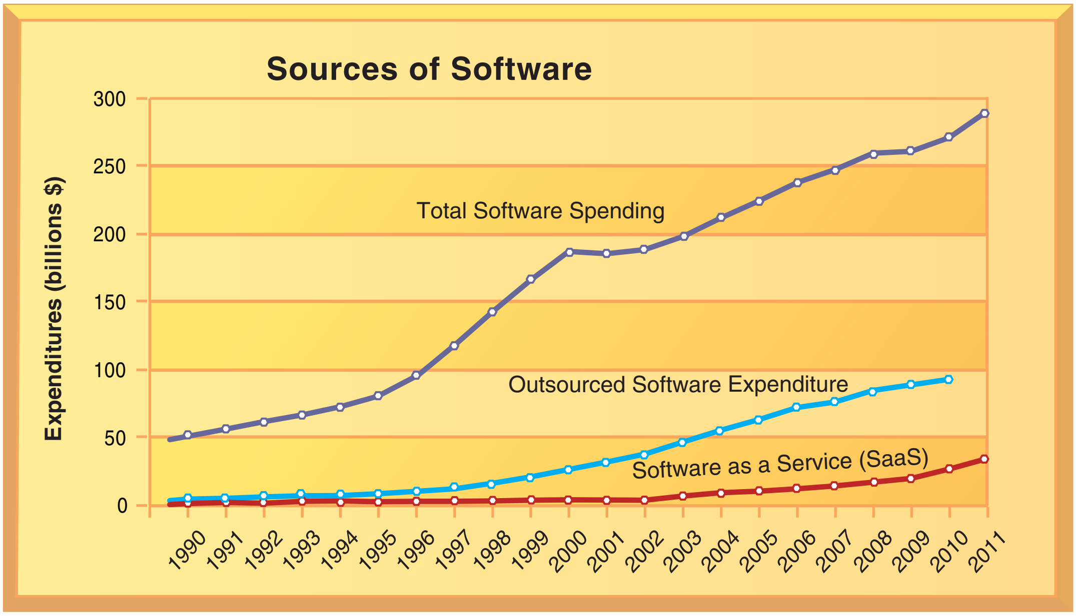

In 2012, U.S. firms were projected to spend more than US $279 billion on software. About 35 percent of that (US $98 billion) would originate outside the firm, either from enterprise software vendors selling firmwide applications or individual application service providers leasing or selling software modules. Another 4 percent (US $11 billion) would be provided by SaaS vendors as an online cloud-based service.

Sources: BEA National Income and Product Accounts, 2012; authors' estimates.

For example, Cemex, Mexico's largest cement manufacturer, signed a 10-year $1 billion outsourcing contract with IBM in July 2012. Under the terms of the contract, 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 centre 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 are 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 centres 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.

Software as a service (SaaS)

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 defining 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 an online 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 slideshows.

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, a Canadian has created a mashup that links Environment Canada's weather radar maps with Google maps. Amazon uses mashup technologies to aggregate product descriptions with partner sites and user profiles.

Apps are small pieces of software 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.

An estimated 1 billion people used apps in 2012 (eMarketer, 2012). A 2012 report by ICTC estimates that there were 22 800 app developers in Canada in 2012, who generated $775 million in revenue in that year. By 2016, app-generated revenue is expected to be $2.2 billion (Information and Communications Technology Council, 2012). By 2012, more than 32 billion apps had been downloaded. Many are free or purchased for a small charge, much less than for conventional software. There are already more than 700 000 apps for the Apple iPhone and iPad platform and a similar number that run on devices using Google's Android operating system; in addition, Microsoft and others offer hundreds of thousands of apps. The success of these mobile platforms depends in large part on the quantity and 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.

Some downloaded apps do not access the Web, but many do, providing faster access to Web content than traditional Web browsers. 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, conduct inquiries on an enterprise-wide system, and dictate voice messages (see the Chapter 1 Window on Management). There are also a huge number of e-commerce apps for researching and buying goods and services online.

5.5 MANAGEMENT ISSUES

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.

Service level agreement (SLA)
Mashups
Apps

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

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 its 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 they allow 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

Scalability
Total cost of ownership (TCO)

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, etc.
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

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 $1000 to $3000, according to various consultant estimates. Gains in productivity and efficiency from equipping employees with mobile computing devices must be balanced against increased costs to integrate these devices into the firm's IT infrastructure and to provide 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 20-35 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.

Competitive Forces Model for IT Infrastructure Investment Figure 5.12 illustrates a competitive forces model you can use to address the question of how much your firm should spend on IT infrastructure.

1. **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 about slow responses to their queries about price

FIGURE 5.12 Competitive forces model for IT infrastructure.

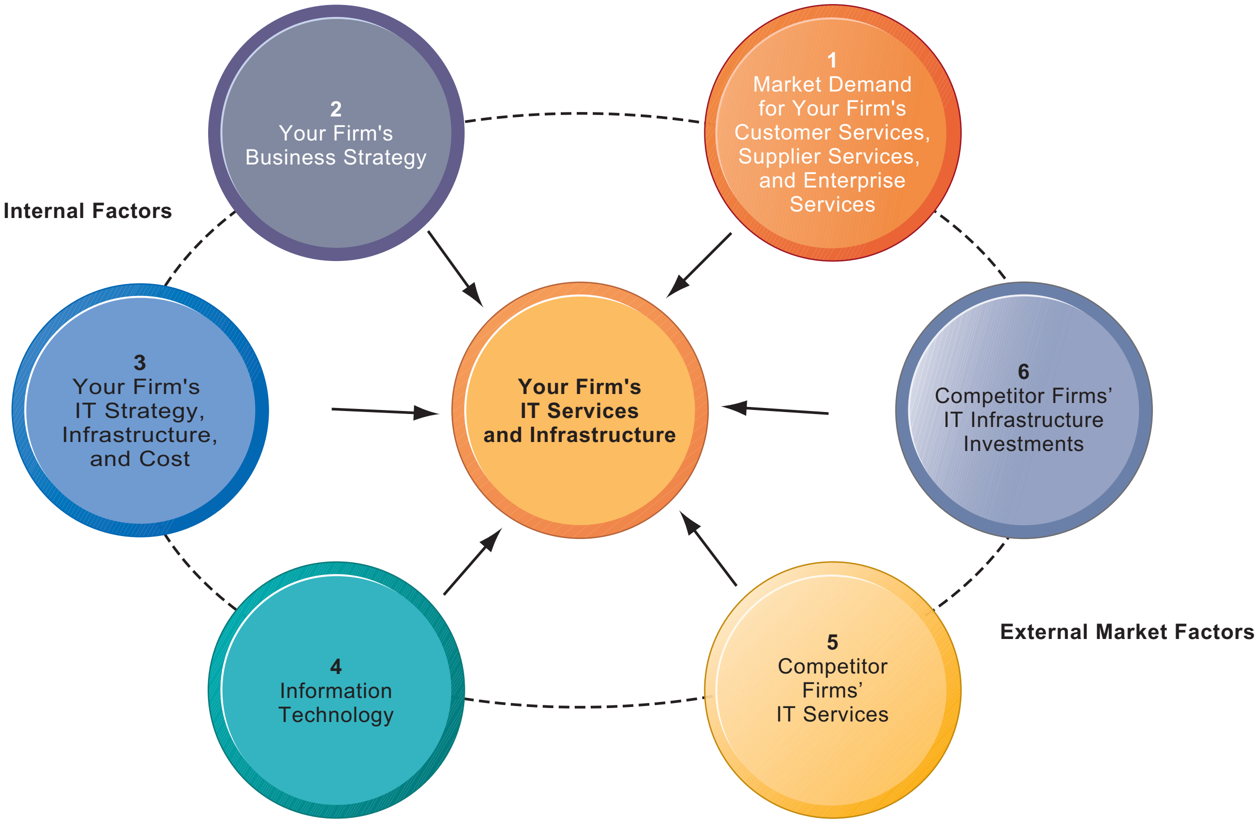

There are six factors you can use to answer the question, "How much should our firm spend on IT infrastructure?"

and availability? Are employees complaining about difficulty finding the right information for their jobs? Are suppliers complaining about difficulties discovering your production requirements?

2. **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.
3. **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.
4. **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 technologies 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 newer technologies.
5. **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.

6. **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' 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 less expensive ways of providing services, which 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.

Learning Track Modules

The following Learning Tracks provide content relevant to topics covered in this chapter:

1. How Computer Hardware and Software Work
2. Service Level Agreements
3. The Open Source Software Initiative
4. Comparing Stages in IT Infrastructure Evolution
5. Cloud Computing

REVIEW SUMMARY

1. *What is IT infrastructure and what are its components?*

   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. 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.

2. *What are the stages and technology drivers of IT infrastructure evolution?*

   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 halves. 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.

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

   Increasingly, computing is taking place on a mobile digital platform. Grid computing involves connecting geographically remote computers into a single network to create a computational grid that combines the computing power of all the computers on the network. 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 impacts on the environment and computing costs. In autonomic computing, computer systems have capabilities for automatically configuring and repairing themselves. Power-saving processors dramatically reduce power consumption in mobile digital devices.

4. *What are the current trends in 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, as well as 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

Review Questions

1. *What is IT infrastructure and what are its components?*
   - *Define IT infrastructure from both a technology and a services perspective.*
   - *List and describe the components of IT infrastructure that firms need to manage.*
2. *What are the stages and technology drivers of IT infrastructure evolution?*
   - *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.*
3. *What are the current trends in computer hardware platforms?*

- *Describe the evolving mobile platform, grid computing, and cloud computing.*
- *Explain how businesses can benefit from autonomic computing, virtualization, green computing, and multicore processors.*

4. *What are the current trends in software platforms?*
   - *Define and describe open source software and Linux, and explain their business benefits.*
   - *Define Java and HTML5, and explain why they are important.*
   - *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. *What are the challenges of managing IT infrastructure and management solutions?*
   - *Name and describe the management challenges posed by IT infrastructure.*
   - *Explain how using a competitive forces model and calculating the TCO of technology assets help firms make good infrastructure investments.*

Discussion Questions

1. *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?*
2. *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?*
3. *What are the advantages and disadvantages of cloud computing?*

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

1. The University of Guelph Medical Centre (UGMC), a fictitious organization, relies on information systems to operate 19 hospitals, a network of other care sites, and international and commercial ventures. Demand for additional servers and storage technology is growing by 20 percent each year. UGMC sets up a separate server for every application, and its servers and other computers are running a number of different operating systems, including several versions of Unix and Windows. UGMC is managing 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?
2. 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 centre, and 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

In this exercise, you will use spreadsheet software to calculate the cost of desktop systems, printers, and software.

Use the Internet to obtain pricing information on hardware and software for an office of 30 people. You will need to price 30 PC desktop systems (monitors, computers, and keyboards) manufactured by Lenovo, Dell, and HP. (For the purposes of this exercise, ignore the fact that desktop systems usually come with preloaded software packages.) Also obtain pricing on 15 desktop printers manufactured by HP, Canon, and Dell. Each desktop system must satisfy the minimum specifications shown in tables which you can find on the Companion Website.

Also obtain pricing on 30 copies of the most recent versions of Microsoft Office, Lotus SmartSuite, and Apache OpenOffice (formerly Oracle Open Office), and on 30 copies of Microsoft Windows 7 Professional or Windows 8 Pro. Each desktop productivity package should contain programs for word processing, spreadsheets, databases, and presentations. Prepare a spreadsheet showing your research results for the software and the desktop system, printer, and software combination offering the best performance and pricing per worker. Because every two workers share one printer (15 printers/30 systems), your calculations should assume only half a printer cost per worker.

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

Software skills: Internet-based software

Business skills: Researching transportation and lodging costs

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 125 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 $150 000 for the representatives' room rentals. The company would like to hold the conference in either Whistler, BC, or Montreal, QC, at a Hilton- or Marriott-owned hotel.

Use the Hilton and Marriott 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 Toronto (54), Vancouver (32), Quebec City (22), Edmonton (19), and Winnipeg (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

On the Companion Website, you will find a Collaboration and Teamwork Project dealing with the concepts in this chapter. You will be able to use Google Sites, Google Docs, and other open source collaboration tools to complete the assignment.

Visit the Companion Website at www.pearsoncanada.ca/laudon to access valuable study tools including

- *Self-Study Quizzes*
- *Glossary Flashcards*
- *Data Files for the Hands-On MIS activities*
- *Learning Track Modules*
- *Collaboration and Teamwork Projects*
- *videos, and more!*

CASE STUDY

Should Businesses Move to the Cloud?

Cloud computing has just begun to take off in the business world. Larger Canadian companies are opening their own cloud facilities. In May 2013, Canadian Tire opened a cloud computing facility in Winnipeg, MB. Yet the biggest player in the cloud computing marketplace is one you might not expect: Amazon. Under its Web Services division (AWS), Amazon has streamlined cloud computing and made it an affordable and sensible option for companies ranging from tiny Internet start-ups to established companies like FedEx.

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 $49 billion a year retail business. To make harnessing the cloud simpler, Amazon added an automated service called CloudFormation that helps customers get the right amount of computing resources for their needs. Customers provide the amount of server space, bandwidth, storage, and any other services they require, and AWS can automatically allocate those resources.

Since its launch in March 2006, AWS has continued to grow in popularity, with $1 billion in business in 2011 and hundreds of thousands of customers across the globe. In fact, Amazon believes that AWS will someday become more valuable than its retail operation. Amazon's sales pitch is that you do not pay a monthly or yearly fee to use their computing resources—instead, you pay for exactly what you use. For many businesses, this is an appealing proposition because it allows Amazon to handle all of the maintenance and upkeep of IT infrastructures, leaving businesses free to spend more time on higher-value work.

The difference between cloud computing today and the cloud computing of the past is the scale of today's clouds and the amount of digital data requiring storage, which has increased exponentially over the past few years. Web companies used to build dozens of data centres, at often up to $500 000 000 per centre. Now, leading cloud companies such as Amazon, Google, and Microsoft have built software that uses automated methods to spread data across the globe and control thousands of servers, and they have refined data centre designs with the goal of increasing efficiency. More than ever, companies are turning to cloud computing providers like these for their computing resources.

Zynga is a good example of a company using cloud computing to improve its business in a new way. Zynga is the developer of wildly popular Facebook applications like *FarmVille*, *Mafia Wars*, and many others. With more than 290 million monthly active users, Zynga's computing demands are already significant. When Zynga releases a new game, however, it has no way of knowing what amount of computing resources to dedicate to the game. The game might be a mild success or a smash hit that adds millions of new users. The ability to design applications that can scale up in the number of users quickly is one of Zynga's competitive advantages.

Because of the uncertainty surrounding resource usage for new game launches, Zynga uses Amazon's cloud computing platform to launch new offerings. That way, it can pay only for the resources it ends up using, and, once game traffic stabilizes and reaches a steady number of users, Zynga moves the game onto its private zCloud, which is structurally similar to Amazon's cloud but operates under Zynga's control in data centres on the east and west coasts. Zynga's own servers handle 80 percent of its games. (Zynga recently started selling extra capacity on zCloud to other game-makers.) To streamline the process of moving application data from Amazon to the zCloud, Zynga has automated many computing tasks, selected hardware and chip configurations that are very similar to Amazon's, and makes significant use of virtualization.

There are a few reasons why Zynga is well-suited to use this combination of public and private clouds. The first is its business model, which involves games that have a tendency to be boom or bust. Rather than spending on computing resources of its own before the launch of each game, it is much more cost-effective to use Amazon's cloud services until Zynga can more accurately predict the computing power it needs. As a recent start-up, Zynga lacks the accumulated legacy systems and infrastructure typically found in older companies. The more systems a company has, the tougher it is to integrate its applications and data with cloud systems.

Although the consequences for server downtime are not as catastrophic for Zynga as they would be for a financial services firm, Zynga still needs 99.9 percent uptime. On its own financial reports, Zynga recognized that a significant majority of its game traffic had been hosted by a single vendor and any failure or significant interruption in its network could negatively impact operations. Amazon Web Services had an outage for several hours in April 2011 that made it impossible for users to log into some of Zynga's games.

However, owning data centres also comes with risks. If the demand for Zynga's games were to drop dramatically, Zynga would have too much IT infrastructure on its hands, and losses could result. The most likely scenario has Zynga owning part of its data centres and relying on external services such as Amazon for the rest.

Not all companies use cloud computing in the same way that Zynga does, but many do. Outback Steakhouse was not sure how popular an upcoming coupon promotion would be, so the company used Microsoft's Azure cloud to launch the promotion. Outback ended up selling an unexpectedly large 670 000 coupons. Using the cloud, Outback was able to avoid taxing in-house systems unnecessarily. Redwood Global, based in Toronto, is Canada's fastest-growing IT and financial services staffing firm. It has 30 full-time employees, and an additional 180 contract staff who work remotely at client sites located across North America—from Toronto, ON to Dallas, TX. Redwood routinely updated its systems and software, but found that on-premises solutions were not scalable or reliable enough to provide needed support to its staff and, ultimately, its clients. Redwood moved to Microsoft Office 365, a cloud-based system that supports messaging, file sharing, email, calendars, and the usual suite of office applications.

Start-ups and smaller companies are finding that they no longer need to build a data centre. With cloud infrastructures like Amazon's readily available, they have access to technical capability that was formerly only available to much larger businesses. For example, online crafts marketplace Etsy uses Amazon computers to analyze data from the 1 billion monthly views of its Web site. Etsy can then use its findings to create product recommendation systems that allow customers to rank which products they like best and to generate a list of 100 products they might enjoy. Etsy's engineers and managers are excited about their ability to handle these types of issues on someone else's computer systems.

IBM, Cisco, and other traditional data centre giants realize that cloud computing is a threat to their technology infrastructure business. As a solution to rising computing costs, they have been steering their customers toward virtualization software, which allows them to run many more applications on each individual server that they buy. There are also many companies that simply have too much legacy technology to use the cloud effectively. For example, Credit Suisse has 7000 software applications running on its systems that have been developed over the past 20 years. Ensuring that all of these applications would work the same way in the cloud would be more trouble than it is worth.

Many other companies share Zynga's concern about cloud reliability and security, and this remains a major barrier to widespread cloud adoption. Amazon's cloud experienced significant outages in April and August 2011 and again on June 14 and 29, 2012. Normally, cloud networks are very reliable, 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.

According to Amazon, a simple network configuration error caused a major multiday service outage in Amazon's east coast region from April 21–24, 2011. Amazingly, the error was most likely a simple human error made during a routine network adjustment. Sites affected included Reddit, Foursquare, Engine Yard, HootSuite, Quora, Zynga, and many more. On June 14 and June 29, 2012, AWS suffered outages due to power failures in its primary east coast data centre in North Virginia. Many popular Web sites, including Netflix, Heroku, Quora, and Pinterest, as well as Web sites of smaller companies, were knocked offline for hours.

The outages were proof that the vision of a cloud with 100 percent uptime is still far from reality. Experts have conflicting opinions of how serious this issue is. A June 2012 report issued by the Paris-based International Working Group on Cloud Computing Resiliency estimated that the major cloud computing services were down about 10 or more hours per year, with average availability at 99.9 percent or less. Even this small amount of downtime can lead to large revenue losses for firms that need 24/7 availability. Nevertheless, some large cloud users such as Netflix believe that overall cloud service availability has steadily improved. Neil Hunt, Netflix's chief product officer, believes the cloud is becoming more reliable and that AWS gives Netflix much larger scale and technical expertise than it would otherwise have. A number of experts recommend that companies for whom an outage would be a major risk consider using another computing service as a backup.

Still, cloud computing has gone mainstream, and the major cloud providers have the sales numbers to prove it. Amazon, Microsoft, Google, Apple, and other cloud providers have to continue to work to avoid outages, while other companies must decide whether the cloud is right for them and, if so, how to most effectively use the cloud to enhance their businesses.

Sources: Charles Babcock, "How Game-Maker Zynga Became a Cloud Vendor," *Information Week*, May 14, 2012; Charles Babcock, "Cloud's Thorniest Question: Does It Pay Off?" *Information Week*, June 4, 2012; Zack Whittaker, "Amazon Explains Latest Cloud Outage: Blame the Power," ZDNet, June 18, 2012; Stuart J. Johnston, "Cloud Outage of 13 Providers Reveals Downtime Costs," searchcloud-computing.com, June 22, 2012; Charles Babcock, "4 Companies Getting Real Results from Cloud Computing," *Information Week*, January 15, 2011; Charles Babcock, "Amazon Launches CloudFormation to Simplify App Development," *Information Week*, February 28, 2011; Ashlee Vance, "The Cloud: Battle of the Tech Titans." *Bloomberg Businessweek* (March 3, 2011); Peter Svensson and Steve Lohr, "Amazon's Trouble Raises Cloud Computing Doubts," *The New York Times*, April 22, 2011; Charles Babcock, "Post Mortem: When Amazon's Cloud Turned on Itself," *Information Week*, April 29, 2011; Patrick Thibodeau, "Amazon Cloud Outage Was Triggered by Configuration Error, "*Computerworld*, April 29, 2011; "Canadian Staffing Firm Gains Scalability, Reclaims $23,000 Annually with the Cloud," Microsoft, March 28, 2012, available at http://www.microsoft.com/Canada/casestudies/Case_Study_Detail.aspx?casestudyid=710000000339, accessed February 21, 2013; and Charles Babcock, "Zynga's Unusual Cloud Strategy is Key To Success," *Information Week*, July 1, 2011; "Canadian Tire to Open Advanced Cloud Computing Centre in Downtown Winnipeg" TechVibes, http://www.techvibes.com/blog/canadian-tire-advanced-cloud-computing-centre-2013-05-03, accessed June 18, 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?