Part II

Information Technology Infrastructure

PART II

CANADIAN CASE STUDY

Wi-Fi in Canada: The Good, the Bad, and the Ugly

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

Bell Canada Uses Virtual Desktop IT Infrastructure to Connect 1700 Desktops

By 2004, Bell Canada employed 8000 call centre agents, one of the largest such staffs in Canada. These agents served more than 27 million customer connections throughout Canada. As the largest subsidiary of Bell Canada Enterprises (BCE), the company was facing outdated hardware and rising costs in its call centres.

By 2004, Bell Canada offered a number of services, including local, mobile, and long-distance telephone service; Internet access; and satellite television, all to both business and residential customers. It continues to offer these services. Today the company employs 44 000 people. With its current desktop systems, Bell faced hardware that was quickly becoming outdated and would need replacing along with rising costs of maintaining its current infrastructure environment.

For its call centre agents, Bell Canada needed a "lockdown" environment that could accommodate "fat clients" (a fat client is a client computer with most of its resources, e.g., processing power, software, and data on the fat client instead of the server). What this means is that the company needed to be able to extend its own local area network (LAN) to the client's building and enable a few people to work in a locked room to meet the security needs of the client. (A LAN is a network that is housed in one building or on one floor of a building; it has a limited geographic scope.) Because Bell Canada used multiple protocols for security, it was a complex environment to try to function in. Explained Martin Quigley, a senior technical consultant to Bell Canada, who worked at CGI Group, Inc., "You couldn't install programs, you couldn't install anything on the hard drive because they used file sharing. So they had to keep it in a locked room. It was difficult to manage, and it was expensive."

CGI had been providing consulting services for Bell Canada since 1993 when Bell Canada decided to outsource its IT operations, and CGI was contracted to provide IT solutions for the company. Quigley explained the current situation: "Bell Canada came to us with a project to provision, connect, and securely deploy 400 desktop environments within three weeks." Bell Canada and CGI evaluated several possible solutions before deciding to use virtual servers. They determined that using many different physical servers is difficult to manage centrally while virtual servers by definition enable multiple servers to run on one computer because updates and security must be managed on only one computer instead of several different computers. Using virtual servers is also less expensive than using the many computers required by physical servers. Users connect to a virtual server just as

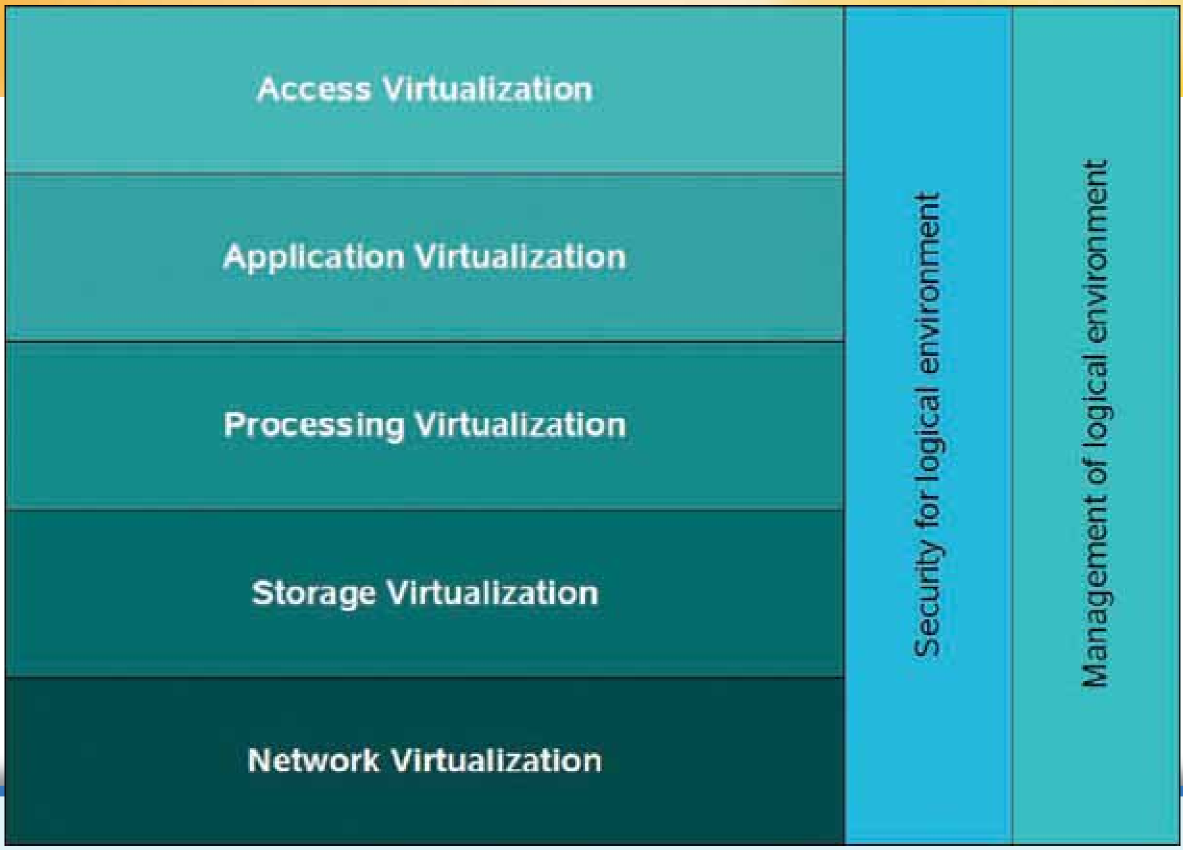

they do to any server, using the local area network.

Using this virtualization technology, Bell Canada lowered its total cost of ownership by using a single server to replace several desktops and by employing a single network connection instead of multiple, more expensive network connections. By using the servers "off-site," Bell Canada also reduced the number of site visits by staff to deal with computer problems or upgrades. By using virtual servers, Bell Canada also eliminated the need to ship hardware and applications to remote teleworkers and to have couriers ship and return older hardware from remote locations.

By using virtual servers, Bell Canada has been able to hire more people to work from remote locations who can telecommute on a daily basis. "We don't have to worry about security or the difficulty of setting up their environments," said Quigley. Neither developers nor users can tell the difference between the use of virtual servers and physical servers since, to them, they are simply connecting to "a server"; regardless of the type of server, it looks the same to everyone except the server administrator. Management of the computing environment has been greatly simplified, and it is both easy and fast to deploy additional desktops.

Since the virtual servers are attached to storage area networks (SANs) which house data physically at various locations but with networked access to the data (see Chapter 6) established by Bell Canada, backup and disaster recovery are easier and faster. Their additional backup servers can take over the load within 25 minutes of a server crashing.

When Bell Canada began this project more than 15 years ago, the company wanted to place 400 desktops in a virtual server environment. By partnering with CGI Group and its partner, VMware, Bell Canada now uses its virtual server environment for 1700 desktops, which has led to further cost savings. VMware is a company that specializes in software to consolidate physical servers into a virtual server.

It is interesting to note that a year after CGI helped Bell Canada to implement its virtual server environment in 2004, BCE decided to sell its stake in CGI, which it had acquired in 1995; the sale amounted to close to US $1 billion. BCE had purchased its 100 million shares in CGI in 1995, after Bell Canada's first consulting contract with CGI, and was considered a "friendly" buy-back. This could be said to represent an interesting use of "insider" consulting.

Sources: John D. Blau, "BCE to Sell Stake in Canadian IT Service Provider CGI," *Infoworld*, December 19, 2005, **www.infoworld.com/article/05/12/19/HNbcestake_1.html?IP%20TELEPHONY** (accessed February 13, 2009); VMware, "VMware Virtual Desktop Infrastructure Connects Bell Canada for Optimized Customer Care," 2006, **www.vmware.com/files/pdf/customers/06Q3_cs_vmw_Bell_Canada_English.pdf** (accessed February 13, 2009; CGI, "CGI Provides HR Self-Service Solution to Bell Canada's Employees, **www.cgi.com/web/en/library/case_studies/71182.htm** (accessed February 13, 2009).

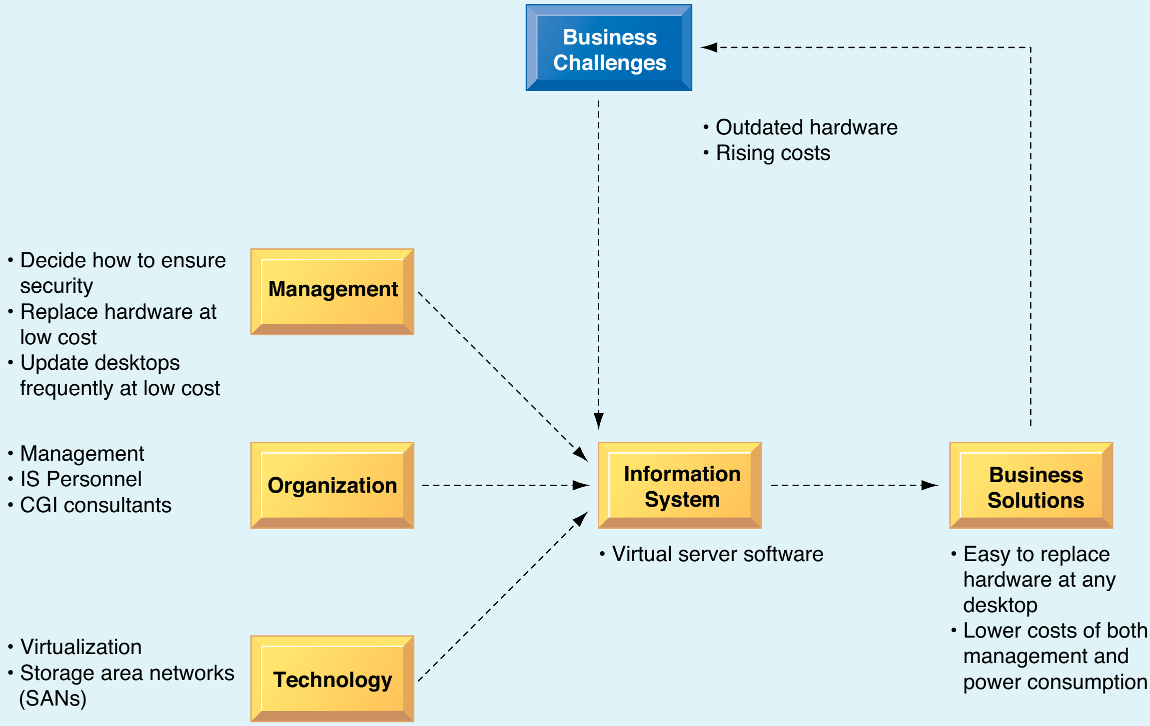

Bell Canada has an enviable track record as a successful telecommunication business. Unfortunately, its growth plans and daily operations were hampered by unmanageable and outdated server technology. Bell Canada's management felt the best solution was to replace its antiquated server infrastructure with new computer hardware and software technologies and to standardize using the technology of a single vendor—VMware. This case highlights the critical role that hardware and software investments can play in improving business performance.

The chapter-opening diagram calls attention to important points raised by this case and this chapter. Management decided that the best way to make technology promote business objectives was to overhaul and standardize IT infrastructure. Bell Canada now uses a more efficient and easily managed server environment that makes it much easier to deploy additional desktops while maintaining security and control. The entire infrastructure is easier to manage and capable of scaling up to accommodate additional desktops, growing transaction loads, and new business opportunities.

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. IT infrastructure includes investment in hardware, software, and services—such as consulting, education, and training—that are shared across the entire firm or across entire business units in the firm. A firm's IT infrastructure provides the foundation for serving customers, working with vendors, and managing internal firm business processes (see Figure 5-1).

In 2008, supplying firms around the world with IT infrastructure was a $2.09 trillion industry when telecommunications, networking equipment, and telecommunications services (Internet, telephone, and data transmission) are included. In 2009, with the global economic crisis, global IT spending is expected to decline by about 6 percent to $1.97 trillion. This is the first decline in global IT spending since 2002, immediately after the "dot-com bubble" burst. However, future predictions are that global IT spending, by 2010, will increase by between 6 percent and 15 percent (Woodie, 2009; Ortutay, 2009; Ogg, 2008). Investments in infrastructure account for between 25 and 35 percent of information technology expenditures in large firms (Weill et al., 2002).

FIGURE 5-1 Connection between the firm, IT infrastructure, and business capabilities.

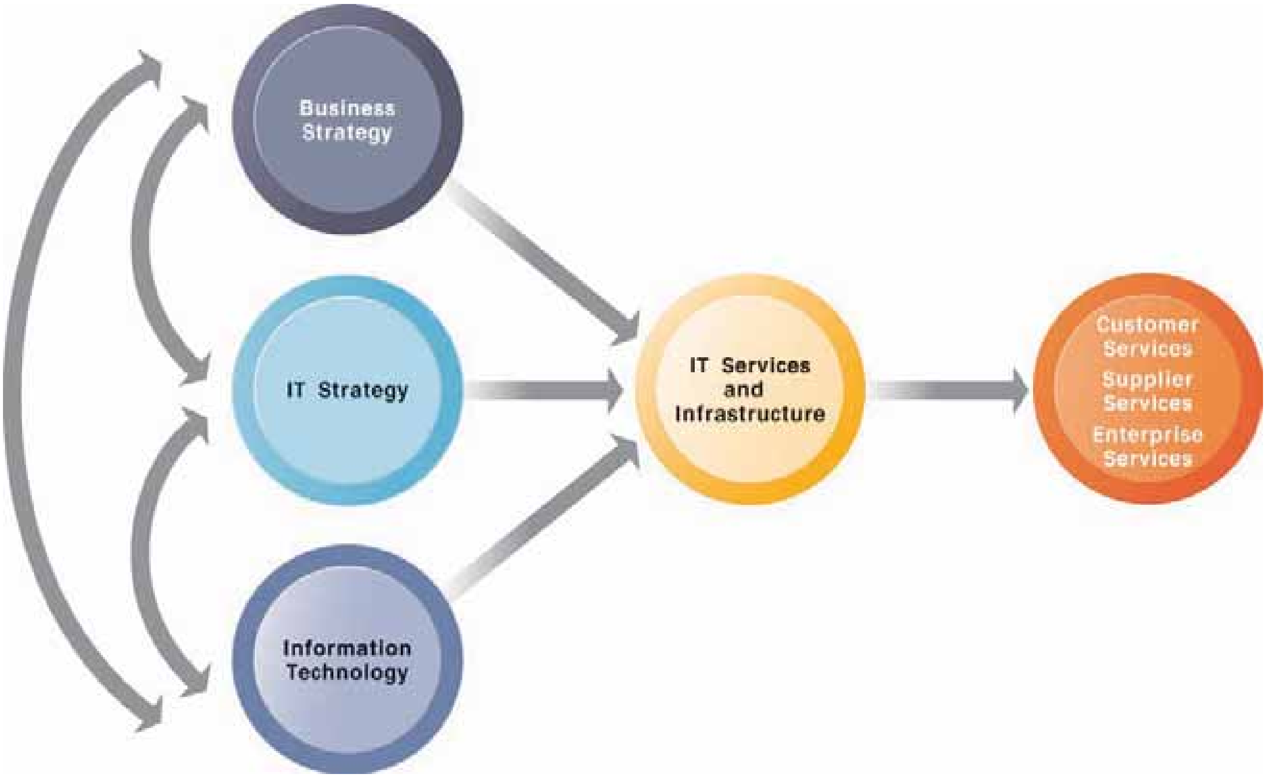

The services a firm is capable of providing to its customers, suppliers, and employees are a direct function of its IT infrastructure. Ideally, this infrastructure should support the firm's business and information systems strategy. New information technologies have a powerful impact on business and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

Defining IT Infrastructure

IT infrastructure consists of a set of physical devices and software applications that are required to operate the entire enterprise. But IT infrastructure is also a set of firm-wide 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, desktop and laptop computers, personal digital assistants (PDAs), and Internet "appliances."
- 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.
- Application software services that provide enterprise-wide capabilities such as enterprise resource planning, customer relationship management, supply chain management, and knowledge management systems that are shared by all business units.
- Physical facilities management services that develop and manage the physical installations required for computing, telecommunications, and data management services.
- IT management services that plan and develop the infrastructure, coordinate IT services with the business units, manage accounting for IT expenditures, 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, and where, when, how, and by whom they will be used.

- IT education services that provide training in system use to employees and offer managers training in how to plan for and manage IT investments.
- IT research and development services that provide the firm with research on potential IT projects and investments that could help the firm differentiate itself in the marketplace.

This "service platform" perspective makes it easier to understand the business value provided by infrastructure investments. For instance, the real business value of a fully loaded personal computer operating at 3 gigahertz that costs about $1000 or 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 Internet. The high-speed Internet service saves this employee about one hour per day in reduced wait time for Internet information. Without this PC and Internet connection, the value of this one employee to the firm might be cut in half.

Evolution of IT Infrastructure

The IT infrastructure in organizations today is an outgrowth of over 50 years of evolution in computing platforms. There have been five stages in this evolution, each representing a different configuration of computing power and infrastructure elements (see Figure 5-2). The five eras are general-purpose mainframe and minicomputer computing, personal computers, client/server networks, enterprise/Internet computing, and cloud computing.

Technologies that characterize one era may also be used in another time period for other purposes. For example, some companies, such as the Bank of Montreal, which employs an IBM z Series mainframe, still run traditional mainframe or minicomputer systems. Mainframe computers today are used 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 dominated mainframe computing from this point on.

Mainframe computers eventually became powerful enough to support thousands of online remote terminals connected to the centralized mainframe using proprietary communication protocols and proprietary data lines. The first airline reservation systems appeared in 1959 and became the prototypical online, real-time interactive computing system that could scale up to service the whole world.

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 necessitating time sharing on a single huge mainframe across the organization.

Personal Computer Era (1981 to Present) Although the first truly personal computers (PCs) appeared in the 1970s (the Xerox Alto, MITS' Altair, 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 U.S. businesses. At first using the DOS operating system, a text-based command language, and later the Microsoft Windows

Mainframe
Minicomputers

FIGURE 5-2 Eras in IT infrastructure evolution.

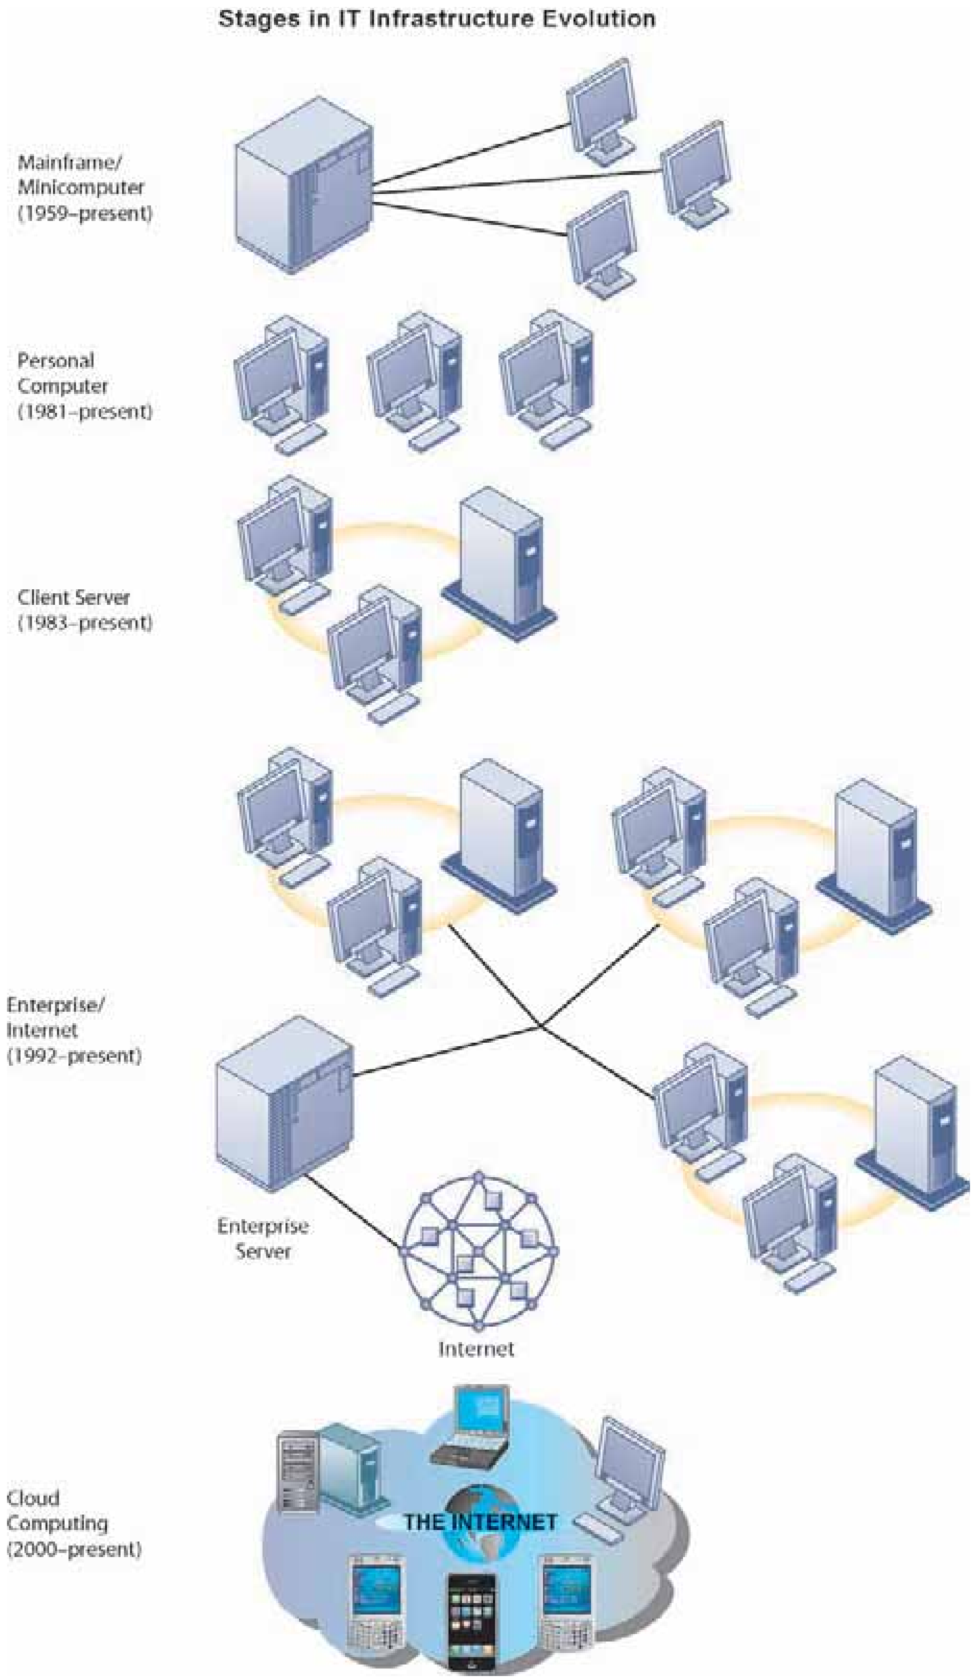

Illustrated here are the typical computing configurations characterizing each of the five eras of IT infrastructure evolution.

operating system, the **Wintel PC** computer (Windows operating system software on a computer with an Intel microprocessor) became the standard desktop personal computer. Today, 95 percent of the world's estimated 1 billion computers use the Wintel standard.

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

Client/Server Era (1983 to Present) In **client/server computing,** desktop or laptop computers called **clients** are networked to powerful **server** computers that provide the client computers with a variety of services and capabilities. Computer processing work is split between these two types of machines. The client is the user point of entry, whereas the server typically processes and stores shared data, serves up Web pages, or manages network activities. The term *server* refers to both the software application and the physical computer on which the network software runs. The server could be a mainframe, but today server computers typically are more powerful versions of personal computers, based on inexpensive Intel chips and often using multiple processors in a single computer box. 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 example, CMS, an enterprise-wide software company headquartered in Toronto, uses a multitiered client/server architecture to manage access by headquarters and global staff as well as customers and developers.

Wintel PC
Client/server computing
Clients
Server
Multitiered
N-tier
Client/server architectures
Web server
Application server

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.

FIGURE 5-3 A multitiered client/server network (N-tier).

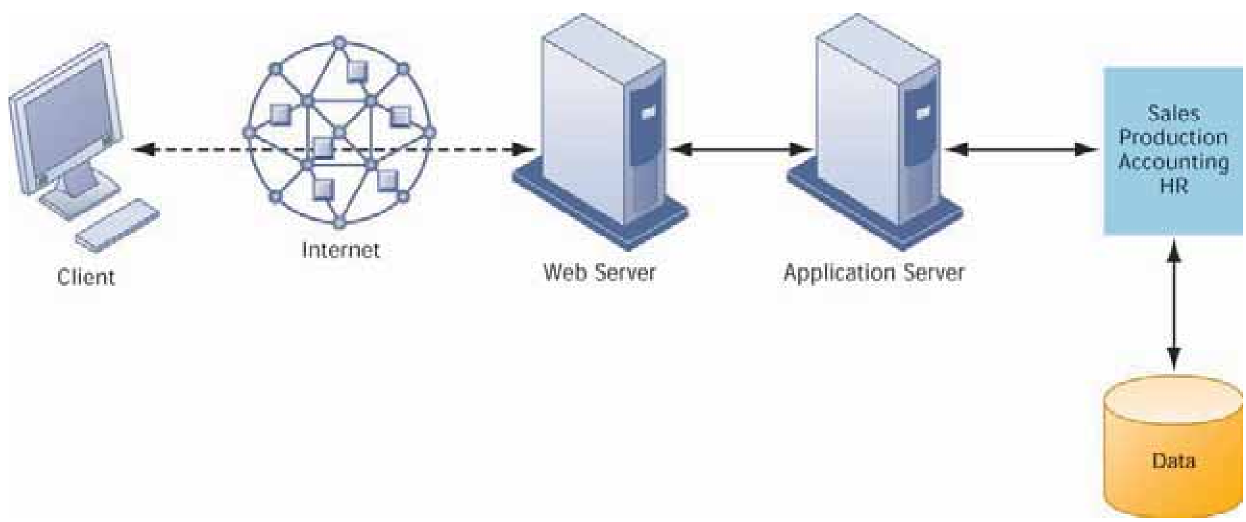

In a multitiered client/server network, client requests for service are handled by different levels of servers.

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

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

Enterprise/Internet Computing Era (1992 to Present) In the early 1990s, firms turned to networking standards and software tools that could integrate disparate networks and applications throughout the firm into an enterprise-wide infrastructure. As the Internet developed into a trusted communications environment after 1995, more and more business firms began 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, mobile phones, and other handheld devices, and it includes public infrastructures such as the telephone system, the Internet, and public network services. 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 12) and Web services (discussed in Section 5.4).

Cloud Computing Era (2000 to Present) The growing bandwidth power of the Internet has pushed the client/server model one step further, toward what is called the "Cloud Computing Model." **Cloud computing** refers to a model of computing in which firms and individuals obtain computing power and software applications over the Internet, rather than purchasing their own hardware and software. Currently, cloud computing is the fastest growing form of computing, with an estimated market size in 2009 of $9.7 billion, and a projected size of $195 billion in 2012 (Gartner, 2008; Merrill Lynch, 2008).

Hardware firms IBM, HP, and Dell are building huge, scalable cloud computing centres that provide computing power, data storage, and high-speed Internet connections to firms that rely on the Internet for business software applications. Software firms such as Google, Microsoft, SAP, Oracle, and Salesforce.com sell software applications as services delivered over the Internet. For instance, more than 500 000 firms in 2009 will use Google Apps, a suite of Internet-based desktop software applications such as word processing, spreadsheets, and calendars (Hamm, 2008; King, 2008). In 2009, more than 43 000 firms worldwide will be using Salesforce.com's customer relationship management software, some on their iPhones (see this chapter's Window on Organizations).

Table 5-1 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 co-founder of Intel Corporation, the leading computer 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

TABLE 5-1 Stages in IT infrastructure evolution.

Infrastructure Dimension	Mainframe Era (1959 to Present)	PC Era (1981 to Present)	Client/Server Era (1983 to Present)	Enterprise/Internet Era (1992 to Present)	Cloud Computing Era (2000 to Present)
Signature Firm(s)	IBM	Microsoft/Intel Dell HP IBM	Novell Microsoft	SAP Oracle PeopleSoft	Google Salesforce.com IBM
Hardware Platform	Centralized mainframe	Wintel computers	Wintel computers	Multiple: • Mainframe • Server • Client	Remote servers Clients (PCs, netbooks, cell phones, smartphones)
Operating System	IBM 360 IBM 370 Unix	DOS/Windows Linux IBM 390	Windows 3.1 Windows Server Linux	Multiple: • Unix/Linux • OS 390 • Windows Server	Linux Windows Mac OS X
Application and Enterprise Software	Few enterprise-wide applications; departmental applications created by in-house programmers	No enterprise connectivity; boxed software	Few enterprise-wide applications; boxed software applications for workgroups and departments	Enterprise-wide applications linked to desktop and departmental applications: • mySAP • Oracle E-Business Suite • PeopleSoft Enterprise One	Google Apps Salesforce.com
Networking/ Telecommunications	Vendor provided: • Systems Network Windows Architecture (IBM) • DECNET (Digital) • AT&T voice	None or limited	Novell NetWare Windows 2003 Linux AT&T voice	LAN Enterprise-wide area network (WAN) TCP/IP Internet standards–enabled	Internet Wi-Fi Wireless broadband cellular networks
System Integration	Vendor-provided	None	Accounting and consulting firms Service firms	Software manufacturer Accounting and consulting firms System integration firms Service firms	SaaS (Software as a Service) firms
Data Storage and Database Management	Magnetic storage Flat files Relational databases	Dbase II and III Access	Multiple database servers with optical and magnetic storage	Enterprise database servers	Remote enterprise database servers
Internet Platforms	Poor to none (1959–1995)	None at first Later browser - enabled clients	None at first Later • Apache server • Microsoft IIS	None in the early years Later: • Intranet-and Internet-delivered enterprise services • Large server farms	Larger server farms

FIGURE 5-4 Moore's Law and microprocessor performance.

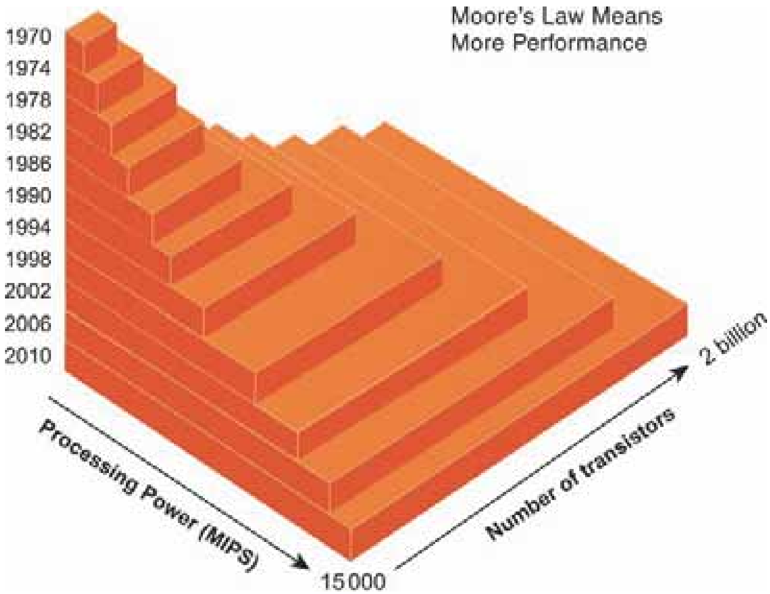

Packing more transistors into a tiny microprocessor has exponentially increased processing power.

Source: Intel, 2004, updated by the authors.

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 computing power.

FIGURE 5-5 Falling cost of chips.

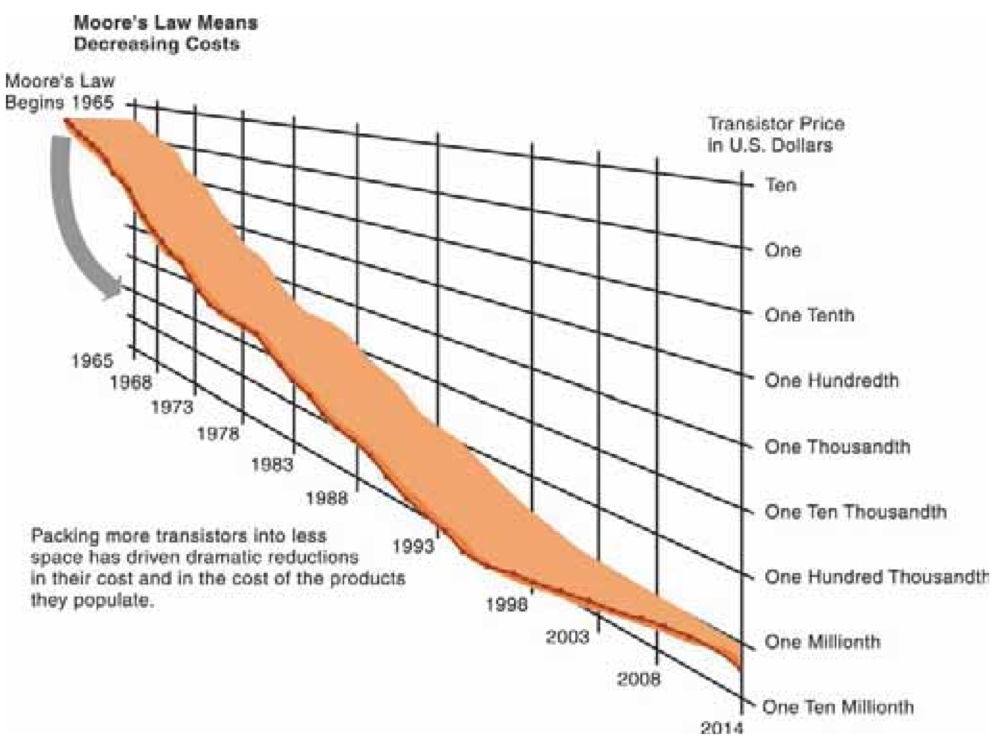

An Intel® processor today can contain as many as 1.72 billion transistors, deliver more than 10 000 MIPS, at a cost of less than 1/10 000th of a cent. That is a little less than the cost of one printed character in this book.

Source: © Intel 2004, updated by the authors.

FIGURE 5-6 Examples of nanotubes.

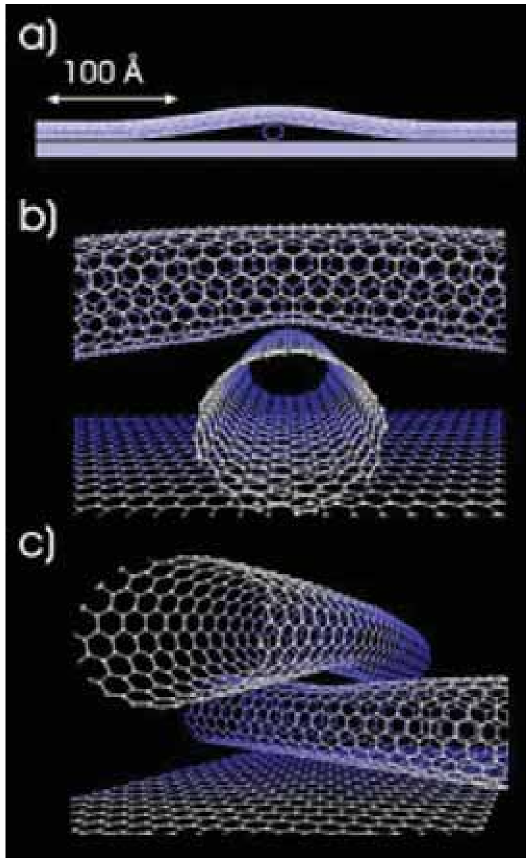

Nanotubes are tiny tubes about 10 000 times thinner than a human hair. They consist of rolled up sheets of carbon hexagons. Discovered in 1991 by researchers at NEC, they have the potential uses as minuscule wires or in ultrasmall electronic devices and are very powerful conductors of electrical current.

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 a virus, the smallest form of organic life.

By using nanotechnology, chip manufacturers can even shrink the size of transistors down to the width of several atoms. **Nanotechnology** uses individual atoms and molecules to create computer chips and other devices that are thousands of times smaller than current technologies permit. Chip manufacturers are trying to develop a manufacturing process that could produce nanotube processors economically (Figure 5-6). IBM has just started making microprocessors in a production setting using this technology.

As processor speeds increase, heat is generated that cannot be dissipated with air fans. Consumers are pressing for lower power consumption, longer battery life, and lower weight to increase laptop and handheld computer portability. For this reason, Intel and other firms are designing the next generation of chips to be less power-hungry and lower in weight. Other options include putting multiple processors on a single chip (see Section 5-3).

The Law of Mass Digital Storage A second technology driver of IT infrastructure change is the Law of Mass Digital Storage. The world produces as much as 5 exabytes of unique information per year (an exabyte is a billion gigabytes, or 1018 bytes). The amount of digital information is roughly doubling every year (Lyman and Varian, 2003). Almost all of this information growth involves magnetic storage of digital data, and printed documents account for only 0.003 percent of the annual growth.

Fortunately, the cost of storing digital information is falling at an exponential rate of 100 percent a year. Figure 5-7 shows that PC hard drive capacity has experienced a compound annual growth rate of 25 percent in the early years to more than 60 percent a year since 1990. Today's PC hard drives have storage densities approaching 1 gigabyte per square inch and total capacities of over 600 gigabytes.

Nanotechnology

FIGURE 5-7 The capacity of hard disk drives grows exponentially, 1980–2008.

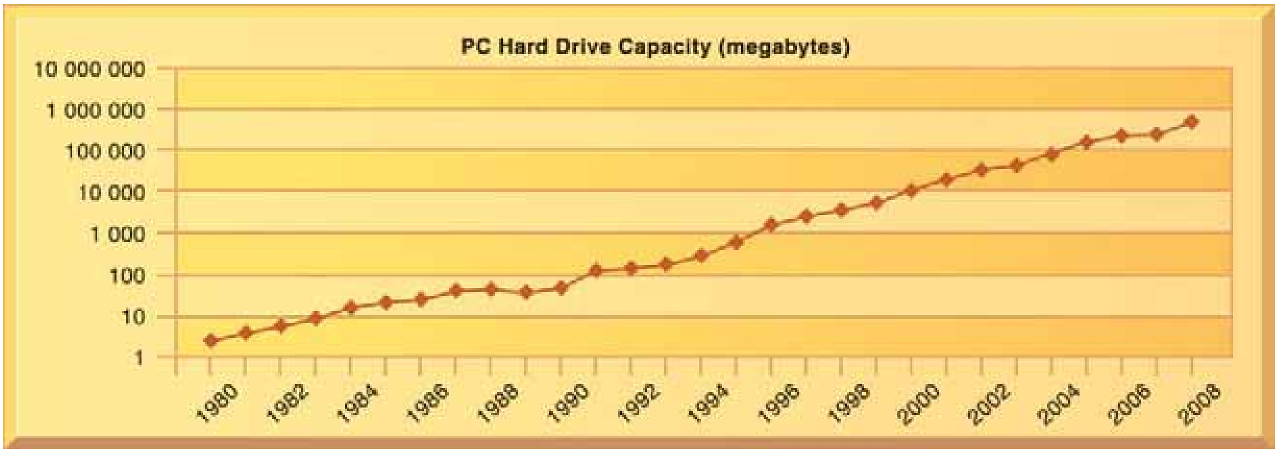

From 1980 to 1990, hard disk drive capacities for PCs experienced a rate of 25 percent annual compound growth, but after 1990, growth accelerated to more than 65 percent each year.

Source: Kurzweil (2003), updated by authors.

Figure 5-8 shows that the number of kilobytes that can be stored on magnetic disks for one dollar from 1950 to the present roughly doubled every 15 months.

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

Robert Metcalfe—inventor of Ethernet local area network technology—claimed in 1970 that the value or power of a network grows exponentially as a function of the

FIGURE 5-8 The cost of storing data declines exponentially, 1950–2010.

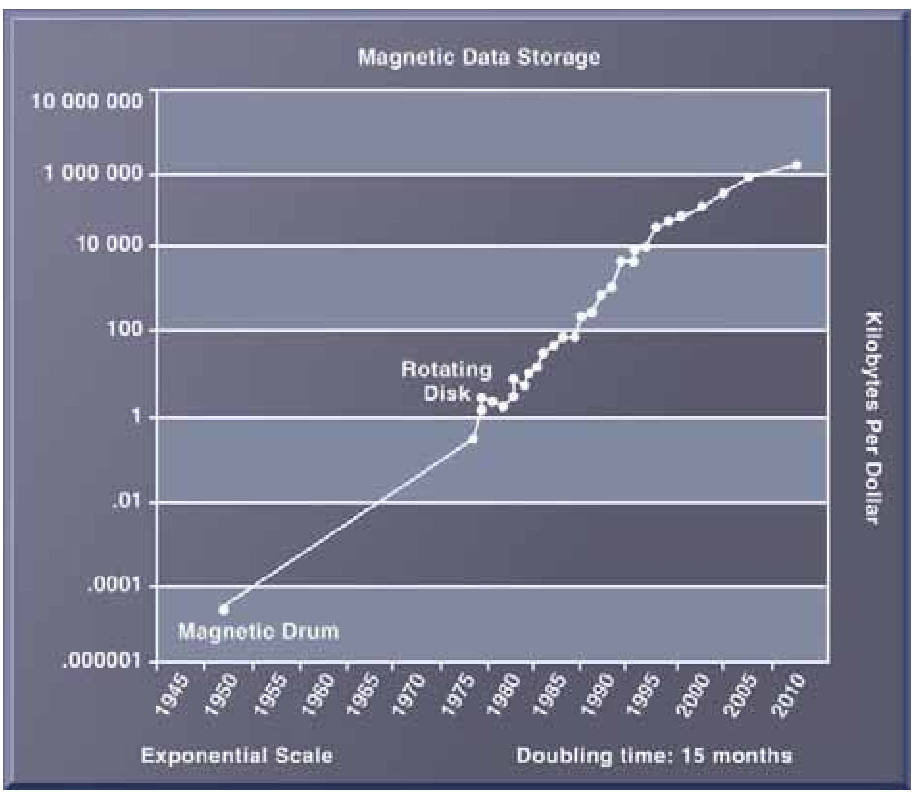

Since the first magnetic storage device was used in 1955, the cost of storing a kilobyte of data has fallen exponentially, doubling the amount of digital storage for each dollar expended every 15 months on average.

Source: Kurzweil (2003) and authors.

FIGURE 5-9 Exponential declines in Internet communication costs.

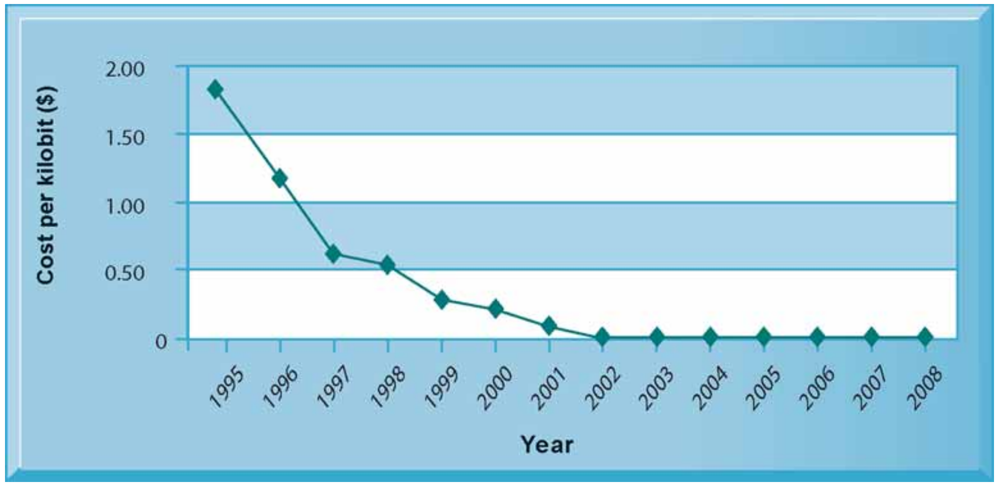

One reason for the growth in the Internet population is the rapid decline in Internet connection and overall communication costs. The cost per kilobit of Internet access has fallen exponentially since 1995. Digital Subscriber Line (DSL) and cable modems now deliver a kilobit of communication for a retail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

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 1.5 billion people worldwide now have Internet access. Figure 5-9 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 explodes.

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

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

Technology standards unleash powerful economies of scale and result in price declines as manufacturers focus on the products built to a single standard. Without these economies of scale, computing of any sort would be far more expensive than is currently the case. Table 5-2 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. Widespread adoption of Unix as the enterprise server operating system of choice made possible the replacement of proprietary and expensive mainframe infrastructure. 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 firm-wide networks, and ultimately, to the Internet.

Technology standards

TABLE 5-2 Some important standards in computing.

STANDARD	SIGNIFICANCE
American Standard Code for Information Interchange (ASCII) (1958)	A code that made it possible for computers 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 U.S. 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)	A suite of communications protocols and a common addressing scheme that enables millions of computers to connect in one giant global network (the Internet). In the last decade, it has become the default networking protocol suite for local area networks and intranets. Developed in the early 1970s for the U.S. Department of Defense by computer scientist Vinton Cerf.
Ethernet (1973)	A network standard for connecting desktop computers into local area networks that enabled the widespread adoption of client/server computing and local area networks and further stimulated the adoption of personal computers.
IBM/Microsoft/Intel Personal Computer (1981)	The standard Wintel design for personal desktop computing based on standard Intel processors and other standard devices, Microsoft DOS, and later Windows software. The emergence of this standard, low-cost product laid the foundation for a 25-year period of explosive growth in computing throughout all organizations around the globe. Today, more than 1 billion PCs power business and government activities every day.
World Wide Web (1989–1993)	Standards for storing, retrieving, formatting, and displaying information as a worldwide web of electronic pages incorporating text, graphics, audio, and video that enabled the creation of a global repository of billions of Web pages by 2004.

5.2 Infrastructure Components

IT infrastructure today is composed of seven major components. Figure 5-10 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 (often accounting firms), 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). Let us examine the size and dynamics of each of these infrastructure components and their markets, as depicted in Figure 5-10.

Computer Hardware Platforms

In 2008, 285 million PCs were shipped worldwide, with a market value of $308 billion. This component includes client machines (desktop PCs, mobile computing devices such as iPhones and BlackBerrys, and laptops) and server machines. The client machines use primarily Intel or AMD microprocessors (Gartner, 2008; Metrics 2.0, 2008).

FIGURE 5-10 The IT infrastructure ecosystem.

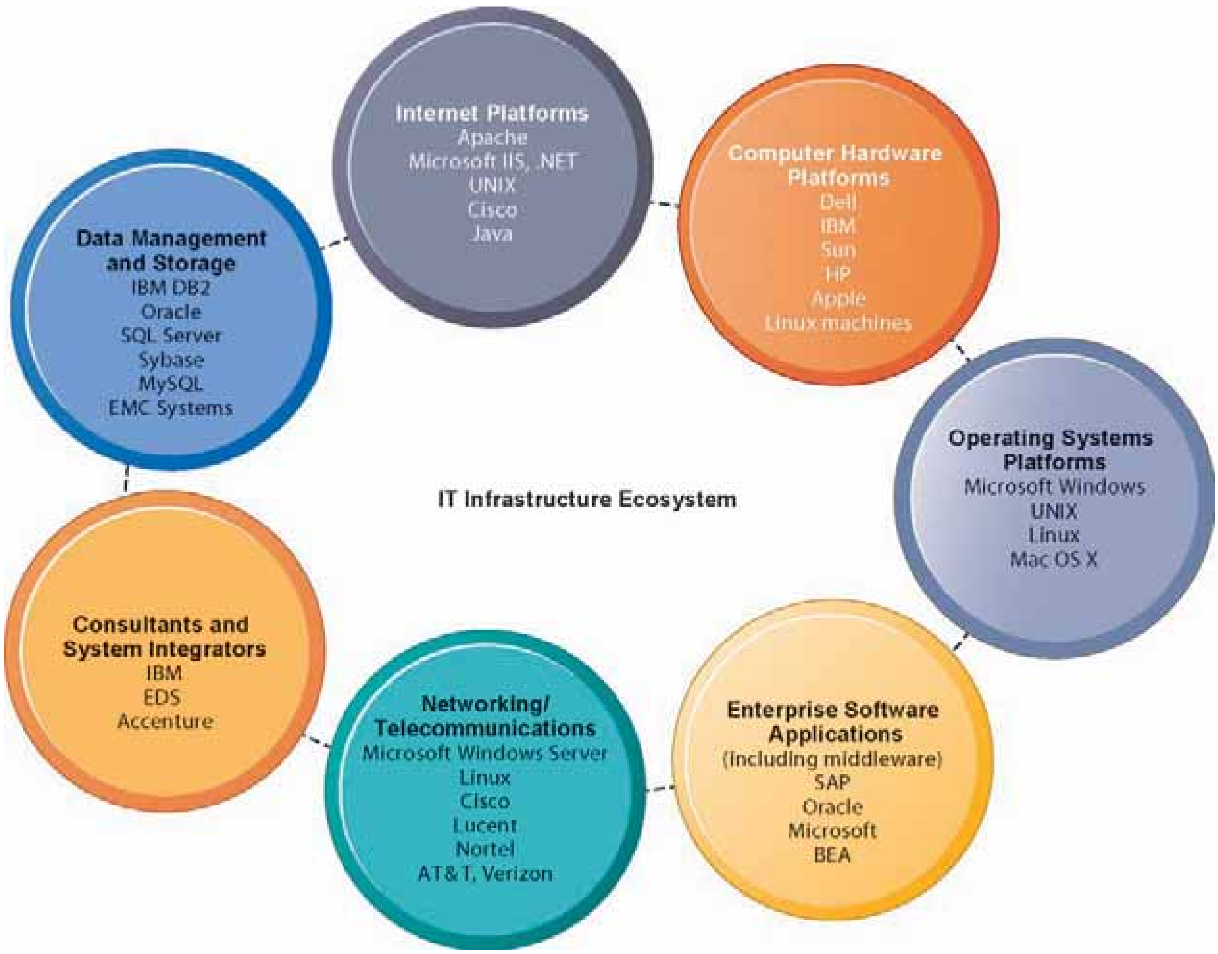

There are seven major components that must be coordinated to provide the firm with a coherent IT infrastructure. Listed here are major technologies and suppliers for each component.

Blade servers

The server market is more complex, using mostly Intel or AMD processors in the form of blade servers in racks, but also includes Sun SPARC microprocessors and IBM PowerPC chips specially designed for server use. **Blade servers** are ultra-thin computers consisting of a circuit board with processors, memory, and network connections that are stored in racks. They take up less space than traditional box-based servers. Secondary storage may be provided by a hard drive in each blade server or by external very large mass-storage drives.

A blade server is a thin, modular processing device that is intended for a single dedicated application (such as serving Web pages). It can be easily inserted into a space-saving rack with many similar servers as shown here.

The supply of computer hardware has increasingly become concentrated in top firms such as IBM, HP, Dell, and Sun Microsystems, and three chip producers, Intel, AMD, and IBM. The industry has collectively settled on Intel as the standard processor, with major exceptions in the server market for Unix and Linux machines, which might use Sun or IBM Unix processors.

Mainframes have not disappeared. The mainframe market has actually grown steadily over the last decade although 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

At the client level, 95 percent of PCs and 45 percent of handheld devices use some form of Microsoft Windows **operating system** (such as Windows Vista, Windows XP, or Windows Mobile) to manage the resources and activities of the computer. Windows comprises 70 percent of the server operating market, with 30 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 2008 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 (meaning you can increase capacity at relatively minimal expense), reliable (meaning these operating systems do not often crash), 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.

Enterprise Software Applications

In addition to software for applications used by specific groups or business units, global firms spent much more than $300 billion in 2008 on software for enterprise applications that are treated as components of IT infrastructure. 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 BEA (also acquired by Oracle in 2008) for achieving firm-wide integration by linking a 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. In general, most large firms already use enterprise applications and have developed long-term relationships with their providers. Once a firm decides to work with an enterprise vendor, switching can be difficult and costly, though not impossible.

Data Management and Storage

There are few choices for enterprise database management software, which is responsible for organizing and managing the firm's data so that it 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 together supply more than 90 percent of the North American database software marketplace. A growing new entrant is MySQL, a Linux open source relational database product available for free on the Internet and increasingly supported by HP and others. MySQL has a 25 percent share of the database software market (Evans Data Corporation, 2008).

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, Maxtor, and Western Digital. 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.

The amount of new digital information in the world is doubling every three years, driven in part by e-commerce and e-business and by statutes and regulations requiring firms to invest in extensive data storage and management facilities. Consequently, the market for digital data storage devices has been growing at more than 15 percent annually over the last five years.

Operating system
Unix
Linux
Storage area networks (SANs)

Networking/Telecommunications Platforms

Globally, firms spend more than $250 billon a year on networking and telecommunications hardware and more than a trillion dollars on networking services (consisting mainly of telecommunications and telephone company charges for voice lines and Internet access; these are not included in this discussion). Chapter 7 is devoted to an in-depth description of the enterprise networking environment, including the Internet and wireless technologies. Windows Server is the predominant local area network operating system, followed by Linux and Unix. Large enterprise-wide area networks primarily use some variant of Unix. Most local and wide area networks use the TCP/IP protocol suite as their networking standard (see Chapter 7).

The leading networking hardware providers are Cisco, Lucent, Nortel, and Juniper Networks. Telecommunications platforms are typically provided by telecommunications/telephone services companies that offer voice and data connectivity, wide area networking, and Internet access. Leading telecommunications service vendors include Bell Canada, Primus, and provincial telecommunications companies, such as SaskTel and MTS. As noted in Chapter 7, this market is exploding with new providers of cellular wireless, Wi-Fi, and Internet telephone services.

Consulting and System Integration Services

Although 20 years ago it might have been possible for a large firm to implement all its own IT infrastructure, today this is far less common. Even large firms do not have the staff, skills, budget, or necessary experience to do so. Implementing new infrastructure requires (as noted in Chapters 9 and 10) significant changes in business processes and procedures, training and education, and software integration. Leading consulting firms providing this expertise include Accenture, IBM Global Services, Electronic Data Systems, HP Technology Solutions, Infosys, and Wipro Technologies.

Software integration means ensuring the new infrastructure works with the firm's older, so-called legacy systems and ensuring the new elements of the infrastructure work with one another. **Legacy systems** are generally older transaction processing systems created for mainframe computers that continue to be used to avoid the high cost of replacing or redesigning them. Replacing these systems is cost-prohibitive and generally not necessary if these older systems can be integrated into a contemporary infrastructure.

Internet Platforms

Internet platforms overlap with, and must relate to, the firm's general networking infrastructure and hardware and software platforms. Globally, firms spend billions on Internet-related infrastructure, such as hardware, software, and management services to support a firm's Web site—including Web hosting services—and for intranets and extranets. 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 of the late 1990s led to 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. The Internet hardware server market has become increasingly concentrated in the hands of Dell, HP/Compaq, and IBM as prices have fallen dramatically.

The major Web software application development tools and suites are supplied by Microsoft (the Microsoft. NET family of development tools used to create Web sites using Active Server Pages for dynamic content), Sun (Sun's Java is the most widely used tool for developing interactive Web applications on both the server and client sides), IBM (WebSphere is IBM's suite of electronic business tools and applications), and a host of independent software developers, including Macromedia (Flash), media software (Real Media), and text tools (Adobe Acrobat). Chapter 7 describes the components of a firm's Internet platform in greater detail.

Legacy systems
Web hosting service

5.3 Contemporary Hardware Platform Trends

Although the cost of computing has fallen exponentially, the cost of the IT infrastructure has actually expanded as a percentage of corporate budgets. Why? The costs of computing services (consulting, systems integration) and software are high, and the intensity of computing and communicating has increased as other costs have declined. For instance, employees now use much more sophisticated applications, requiring more powerful and expensive hardware of many different types (laptop, desktop, mobile handheld computers).

Firms face a number of other challenges. They need to integrate information stored in different applications and even on different platforms (telephone, legacy systems, intranet, Internet sites, desktop, and mobile devices). Firms also need to build resilient infrastructures that can withstand huge increases in peak loads and routine assaults from hackers and viruses while conserving electrical power. Firms need to increase their service levels to respond to growing customer and employee expectations for service. The trends in hardware and software platforms we now describe address some or all of these challenges.

The Emerging Mobile Digital Platform

As computing increasingly takes place over the network, new mobile digital computing platforms have emerged. Communication devices such as cell phones and smartphones such as the BlackBerry and iPhone have taken on many functions of handheld computers, 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, low-cost lightweight subnotebooks called **netbooks,** optimized for wireless communication and Internet access, with core computing functions such as word processing, and digital e-book readers such as Amazon's Kindle with some Web access capabilities. More and more business computing is moving from PCs and desktop machines to these mobile devices; managers increasingly use these devices to coordinate work and communicate with employees.

Netbooks

The Eee PC netbook designed by ASUS weighs only two pounds and uses a Linux-based operating system to provide wireless Internet access and a series of open source desktop productivity tools. Portability, ease of use, and low cost have made netbooks increasingly popular computing platforms.

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 in the United States 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 across the connection.

Grid computing requires software programs to control and allocate resources on the grid. Client software communicates with a server software application. The server software breaks data and application code into chunks that are then parcelled out to the grid's machines. The client machines can perform their traditional tasks while running grid applications in the background.

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.

Cloud Computing and the Computing Utility

Earlier in this chapter, we introduced cloud computing, in which hardware and software capabilities are provided as services over the Internet (also referred to as "the cloud"). Data are permanently stored in remote servers in massive data centres and accessed and updated over the Internet using clients that include desktops, notebooks, netbooks, entertainment centres, and mobile devices. For example, Google Apps provides common business applications online that are accessed from a Web browser, while the software and user data are stored on the servers.

Since organizations using cloud computing generally 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 (or are billed on a subscription basis). You will hear the terms **on-demand computing** or **utility computing** used to describe these services. The Window on Organizations on Salesforce.com on page XXX and the chapter-ending case on Amazon.com describe examples of cloud computing hardware and software services providers.

Some analysts believe that cloud computing represents a sea change in the way computing will be performed by corporations as business computing shifts out of private data centres into "the cloud" (Carr, 2008). This remains a matter of debate. Cloud computing is more immediately appealing to small and medium-size 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 most likely scenario is a hybrid computing model in which firms will use their own infrastructure for their most essential core activities and adopt cloud computing for less critical systems. 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.

Grid computing
On-demand computing
Utility computing

Autonomic Computing

Computer systems have become so complex that some experts believe they may not be manageable in the future. With operating systems, enterprise, and database software weighing in at millions of lines of code, and large systems encompassing many thousands of networked devices, the problem of managing these systems is both complex and expensive.

It is estimated that one-third to one-half of a company's total IT budget is spent preventing or recovering from system crashes. About 40 percent of these crashes are caused by operator error. The reason is not that operators are not well trained or do not have the right capabilities; rather, it happens because the complexities of today's computer systems are too difficult to understand, and IT operators and managers are under pressure to make decisions about problems in seconds.

One approach to dealing with this problem from a computer hardware perspective is to employ 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. Imagine, for instance, a desktop PC that could know it was invaded by a computer virus. Instead of blindly allowing the virus to invade, the PC would identify and eradicate the virus or, alternatively, turn its workload over to another processor and shut itself down before the virus destroyed any files.

A few of these capabilities are present in desktop operating systems. For instance, virus and firewall protection software can 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.

Autonomic computing

Technology

WINDOW ON TECHNOLOGY

Computing Goes Green

Computer rooms are becoming too hot to handle. Data-hungry tasks such as video on demand, music downloads, exchanging photos, and maintaining Web sites require more and more power-hungry machines. The number of servers in corporate data centre servers increased to 29 million worldwide. During the same period, the total annual cost of electricity for data centres jumped from $3.9 billion to $8.8 billion across the world.

What's more, the heat generated from all of these servers is causing equipment to fail. Firms are forced to spend even more on cooling their data centres or to find other solutions. Some organizations spend more money to keep their data centres cool than they spend to lease the property itself. Cooling costs have helped raise the average annual utility bill of a 9000 square metre data centre to $7.2 million. It is a vicious cycle, as companies must pay to power their servers and then pay again to keep them cool and operational. Cooling a server requires roughly the same number of kilowatts of energy as running one. All of this additional power consumption has a negative impact on the environment as well as on corporate operating costs. A major data centre can have power consumption costs that run into the millions of dollars. Saving a large proportion of these costs can have a significant effect on the bottom line for these companies.

BT Global Services, an international company specializing in a variety of IT consulting and other professional services that has offices in Canada, has managed over the last several years to cut its data centre power consumption costs by between 60 and 70 percent and managed to reduce costs as well. According to Steve O'Donnell, global head of Data Centre and Customer Experience Management at BT, "We're really very far ahead in thought leadership on this. We want the market to follow and reduce the global carbon footprint for IT services." BT even wants to increase its savings on power consumption to 80 percent by 2016.

In Kingston, Ontario, Queen's University sent out a bulletin to bring peoples' attention to computer usage and practices, including turning computers off at night and using the energy savings settings on terminals to save money. "There are about 8500 computers on the Queen's campus, and there's significant energy savings to be had if they were turned off in the evenings and if staff and students used the energy [saving] settings on their computers," said John Witjes, engineering director of physical plant services at Queen's.

Another cooling solution comes from Degree Controls Inc., based in Milford, New Hampshire. Degree Controls installs floor tiles equipped with powerful fans that blow cool air directly onto servers. The tiles cost $2200 each.

Some of the world's most prominent firms are tackling their power consumption issues with one eye toward saving the environment and the other toward saving dollars. Google, Microsoft, and HSBC are all building data centres that will take advantage of hydroelectric power. Salesforce.com plans to offset its carbon footprint by investing in renewable energy projects and alternative energy sources.

Hewlett Packard is working on a series of technologies to reduce the carbon footprint of data centres by 75 percent, replace the copper wiring on microprocessors with light pulses, and develop new software and services to measure energy use and carbon emissions. It reduced its power costs by 20 to 25 percent through a consolidation of servers and data centres. None of these companies claim that their efforts will save the world, but they do demonstrate recognition of a growing problem and the commencement of the **green computing** era.

IT managers also have hardware and software options that conserve power. Some organizations are choosing to use thin client computers, which are very basic terminal machines that connect directly to servers and consume significantly less power than normal desktop clients. A call centre operated by Verizon Wireless in Chandler, Arizona, replaced 1700 PCs with thin clients from Sun Microsystems and saw its power consumption go down by one-third. Sun states that on average its thin clients use less than half of the electricity that PCs require. Enterprise Rent-A-Car expects a move from PCs to HP thin clients will cut energy consumption by 5 million kilowatt-hours, save about $500 000 annually, and reduce CO_2 emissions by 6.5 million pounds per year, *Computerworld* reports. The network of more than 45 000 thin client terminals—operated through 743 terminal servers—connects Enterprise rental offices throughout Canada, the United States, the United Kingdom, Ireland, Germany, and Puerto Rico. Thin clients use only 13.6 watt-hours and 2.4 watt-hours of electricity in active and passive states, respectively, compared with the 77.1 watt-hours and 1.8 watt-hours consumed by PCs.

Microsoft's Windows Vista operating system has enhanced sleep features that reduce power consumption by much greater margins than the standby modes in previous versions of Windows. In sleep mode, computers may draw as little as 3 to 4 watts of power versus 100 watts for an idle computer that is not asleep. Businesses also have the option of using more efficient chips in their servers. Virtualization is a highly effective tool for more cost-effective greener computing because it reduces the number of servers required to run a firm's applications. As we saw in the chapter-opening case about Bell Canada, virtualization yields cost reductions and better management as well as the power savings associated with the virtual server environment.

Going green is more than saving power though. It also refers to the need to reduce the use of hazardous materials in computing devices, maximize energy efficiency during the product's lifetime, and promote recyclability or biodegradability of defunct products and factory waste. Many of the components that go into computers are either harmful, such as mercury and lead, or hard to recycle. The Silicon Valley Technology Coalition believes that there are more than 500 million obsolete computers in North America with an additional 130 million cell phones being disposed of annually.

What happens to these devices? Some are, of course, given to organizations and individuals who may or may not be able to use the outdated technology. Others are sent away to be decomposed and recycled. The Silicon Valley Toxics Coalition estimates that 80 percent of the post-consumer e-waste collected for recycling is shipped abroad to countries such as China, India, and Pakistan. Still others are simply left as garbage. Movements are under way to reduce, reuse, and recycle computers, their components, and their packaging in a safe environment.

A variety of government and industry initiatives have been implemented to help move manufacturers and others toward green computing. These include the Energy Star Program adopted by Canada, the European Union's directives 2002/95/EC (RoHS) on the reduction of hazardous substances, the Green Grid consortium of major industry manufacturers, and the International Professional Practice Partnership promulgated by the International Federation for Information Processing (IFIP) which includes certification processes. It is hoped that these and additional initiatives will progress toward global green computing.

To Think About

1. What business and social problems does data centre power consumption cause?
2. What solutions are available for these problems? Which are the most environment-friendly?
3. What are the business benefits and costs of these solutions?
4. Should all firms move toward green computing? Why or why not?

MIS in Action

Perform an Internet search on the phrase *green computing* and then answer the following questions:

1. How would you define green computing?
2. Who are some of the leaders of the green computing movement? Which corporations are leading the way? Which environmental organizations are playing an important role?
3. What are the latest trends in green computing? What kind of impact are they having?
4. What can individuals do to contribute to the green computing movement? Is the movement worthwhile?

Sources: Jennifer Pritchett, "Queen's University Saving Green by Going Green," *The Kingston Whig Standard*, November 21, 2007; "Green IT Comes of Age: Data Centres That Use 60% percent Less Power," **www.biggerthinking.com/en/efficient_data_centres.aspx** (accessed February 13, 2009); Scott Ferguson, "Cooling the Data Center," *eWeek*, June 9, 2008; Rob Bernard, "Microsoft's Green Growth," *eWeek*, April 7, 2008; Eric Chabrow, "The Wild, Wild Cost of Data Centers," *CIO Insight*, May 2008; Jim Carlton, "IT Managers Make a Power Play," *The Wall Street Journal*, March 27, 2007; "IT Managers Find Novel Ways to Cool Powerful Servers," *The Wall Street Journal*, April 10, 2007; and Marianne Kolbasuk McGee, "Data Center Electricity Bills Double," *Information Week*, February 17, 2007.

Virtualization and Multicore Processors

As companies deploy hundreds or thousands of servers, many have discovered that they are spending almost as much on electricity to power and cool their systems as they did on purchasing the hardware. Energy consumed by data centres more than doubled between 2000 and 2008. In a 2008 survey of North American executives and IT managers regarding their data centres, energy conservation was the most difficult issue to resolve with their current tools. Respondents felt that managing the total cost of power was the second most difficult task, and many of the respondents noted that their interest and work with virtualization technology were influenced by the hope of ultimate energy savings. Some of the survey results regarding energy usage:

- 33 percent of companies have implemented server virtualization for energy-saving goals.
- 55 percent of companies can measure power usage in their data centre, primarily at the UPS level.
- 83 percent consider the ability to measure power consumption at the entire data centre level as "valuable" or "extremely valuable" ("Green IT," 2009; "White Paper," 2006).

Cutting power consumption in data centres is now a major business challenge. The Window on Technology: Computing Goes Green, examines this problem. As you read this Window on Technology, try to identify the alternative solutions for this problem and the advantages and disadvantages of each.

The Window on Technology described organizations curbing hardware proliferation and power consumption by using virtualization to reduce the number of computers required for processing. A virtual server uses the hardware of one physical computer to provide the appearance, services, and capabilities of multiple servers, each of which appear to be running their own operating system and providing a unique operating environment. Server virtualization enables companies to run more than one operating system at the same time on a single machine. **Virtualization** means that servers can be accessed in ways that are not restricted by physical configuration or geographic location. Most servers run at just 10 to 15 percent of capacity, and virtualization can boost utilization server utilization rates to 70 percent or higher. Higher utilization rates translate into fewer computers required to process the same amount of work. As we saw in the opening case, it is possible to solve a number of management problems and save costs by using virtualization.

For example, Liverpool Women's NHS Foundation Trust, England's largest specialist in providing health care services for women, saved more than 70 percent of its power consumption costs by combining 30 physical servers into only 4 physical servers through virtualization. Administrative costs of server management also fell by 70 percent. Server virtualization software runs between the operating system and the hardware, masking server resources, including the number and identity of physical servers, processors, and operating systems, from server users. VMware is the leading server virtualization software vendor for Windows and Linux systems. Microsoft offers its own Virtual Server product and has built virtualization capabilities into the newest version of Windows Server.

In addition to reducing hardware and power expenditures, virtualization allows businesses to run their legacy applications on older versions of an operating system on the same server as newer applications. Virtualization also facilitates centralization of hardware administration.

Multicore Processors Another way to reduce power requirements and hardware sprawl is to use multicore processors. A **multicore processor** is an integrated circuit to which two or more processors have been attached for enhanced performance, reduced power consumption, and more efficient simultaneous processing of multiple tasks. This technology enables two processing engines with reduced power requirements and heat dissipation to perform tasks faster than a resource-hungry chip with a single processing core. Today you will find dual-core processors in PCs and quad-core processors in servers. Sun Microsystems' UltraSPARC T2 chip for managing Web applications has eight processors and will soon be followed by a 16-core processor.

Virtualization
Multicore processor

5.4 Contemporary Software Platform Trends

There are five major themes in contemporary software platform evolution:

- Linux and open source software
- Java and Ajax
- Web services and service-oriented architecture
- Software mashups and Web 2.0 applications
- Software outsourcing

Linux and Open Source Software

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

Open source software is based on the premise that it is superior to commercially produced proprietary software because thousands of programmers around the world working for no pay can read, perfect, distribute, and modify the source code much faster, and with more reliable results, than small teams of programmers working for a single software company. The open source movement has been evolving for more than 30 years and has demonstrated after many years of effort that it can produce commercially acceptable, high-quality software.

Now many thousands of open source programs are available from hundreds of Web sites. Popular open source software tools include the Linux operating system, the Apache HTTP Web server, the Mozilla Firefox Web browser, and the 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, Hewlett-Packard, Dell, Oracle, and SAP, now offer Linux-compatible versions of their products. You can find out more out more about the Open Source Definition from the Open Source Initiative, a group that was founded 20 years ago to promote the use of open source software. You can also learn about the history of open source software at the Learning Tracks for this chapter.

Linux Perhaps the most well-known open source software is Linux, an operating system related to Unix. Linux was created by the Finnish programmer Linus Torvalds and first posted on the Internet in August 1991. Linux applications are embedded in cell phones, smartphones, netbooks, and other handheld devices. 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.

Linux is currently a small but rapidly growing presence on the desktop, especially as an operating system for Internet-enabled netbooks. It plays a major role in the back office running local area networks, Web servers, and high-performance computing work, with 20 percent of the server operating system market. IBM, HP, Intel, Dell, and Sun have made Linux a central part of their offerings to corporations. More than two dozen countries in Asia, Europe, and Latin America have adopted open source software and Linux.

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

Software for the Web: Java and Ajax

Open source software
Java

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

Web. Java was created by James Gosling and the Green Team at Sun Microsystems in 1992. Nearly all Web browsers come with a Java platform built in. More recently, the Java platform has migrated into cellular 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. For each of the computing environments in which Java is used, Sun has 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 is a very robust language that can handle text, data, graphics, sound, and video, all within one program if needed. Java enables PC users to manipulate data on networked systems using Web browsers, reducing the need to write specialized software. A **Web browser** is an easy-to-use software tool with a graphical user interface for displaying Web pages and for accessing the Web and other Internet resources. Microsoft's Internet Explorer, Mozilla Firefox, and Netscape 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.

The rapid deployment of Java was hindered in the past because of disagreements between Sun Microsystems and Microsoft over Java standards. In April 2004, under pressure from major customers such as General Motors, Microsoft agreed to stop distributing the Microsoft Java Virtual Machine (MSJVM) it had developed for its proprietary version of Java and to cooperate with Sun in the development of new technologies, including Java.

Ajax Have you ever filled out a Web order form, made a mistake, and then had to start all over again after a long wait for a new order form page to appear on your computer screen? Or visited a map site, clicked the North arrow once, and waited for an entire new page to load? **Ajax** (Asynchronous Java Script and XML) is another Web development technique for creating interactive Web applications that prevents all of this inconvenience.

Ajax allows a client and server to exchange small pieces of data behind the scenes so that an entire Web page does not have to be reloaded each time the user requests a change. So if you click North on a map site, such as Google Maps, the server downloads just that part of the application that changes with no wait for an entirely new map. You can also grab maps in map applications and move the map in any direction without forcing a reload of the entire page. Ajax uses JavaScript programs downloaded to your client to maintain a near continuous conversation with the server you are using, making the user experience more seamless.

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

The foundation technology for Web services is **XML,** which stands for **Extensible Markup Language.** This language was developed in 1996 by the World Wide Web Consortium (W3C, the international body that oversees the development of the Web) as a more powerful and flexible markup language than hypertext markup language (HTML) for Web pages. **Hypertext markup language (HTML)** is a page description language for specifying how text, graphics, video, and sound are placed on a Web page document. 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 5-3 illustrates some sample XML statements.

Web browser
Ajax
Web services
XML
Extensible Markup Language
Hypertext markup language (HTML)

TABLE 5-3 Examples of XML.

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

By tagging selected elements of the content of documents for their meanings, XML makes it possible for computers to manipulate and interpret 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. *SOAP*, which stands for *Simple Object Access Protocol*, is a set of rules for structuring messages that enables applications to pass data and instructions to one another. *WSDL* stands for *Web Services Description Language*; it is a common framework for describing the tasks performed by a Web service and the commands and data it will accept so that the Web service can be used by other applications. *UDDI*, which stands for *Universal Description, Discovery, and Integration*, enables a Web service to be listed in a directory of Web services so that the Web service can be easily located. Companies discover and locate Web services through this directory much as they would locate services in the Yellow Pages. Using these protocols, a software application can connect freely to other applications without custom programming for each different application with which it wants to communicate. All computers share the same standards.

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

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

Here is a real-world example of an interorganizational SOA at work. 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 SouthwestAir.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 uses 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 mobile phones and PDAs. It does not have to write new software code for each new partner's information systems or each new wireless device (see Figure 5-11).

Service-oriented architecture (SOA)

FIGURE 5-11 How Dollar Rent A Car uses Web services.

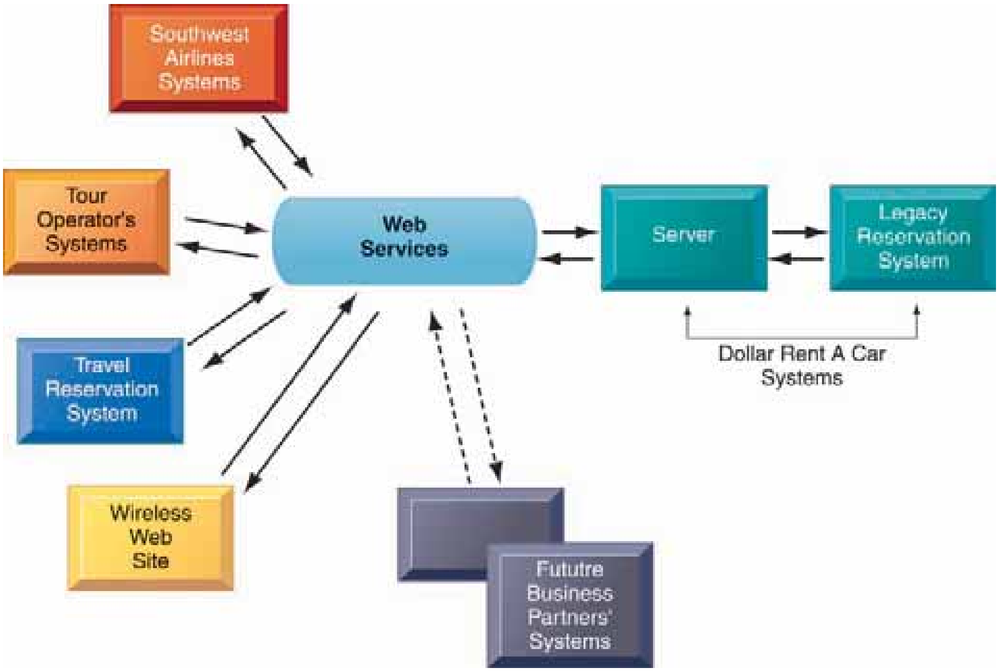

Dollar Rent A Car uses Web services to provide a standard intermediate layer of software to "talk" to other companies' information systems. Dollar Rent A Car can use this set of Web services to link to other companies' information systems without having to build a separate link to each firm's systems.

Mashups and Widgets

In the past, software such as Microsoft Word or Adobe Illustrator came in a box and was designed to operate on a single machine. Increasingly, software is downloadable from the Internet and 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 software from different sources and combine it in order to produce a application that is "greater than" the sum of its parts.

Part of the movement called Web 2.0 (see Chapter 7), and in the spirit of musical mashups, 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. One area of great innovation is the mashup of mapping and satellite image software with local content. For instance, Zoocasa is a new real estate search engine in Canada that is using Google Maps to display real estate listings. The Zoocasa home page is a simple search box in which the user enters a city or neighbourhood and can define search criteria by price, number of bedrooms, and number of bathrooms. The Zoocasa search results are then presented on a Google Map and listed in a sidebar beside the map. One innovative feature of the map is that if you click on a property marker, the sidebar automatically scrolls to display the property details. Google, Yahoo!, and Microsoft now offer tools to allow other applications to pull in information from their map and satellite images with relatively little programming.

You have performed a mashup if you have ever personalized your Facebook profile or your blog with a capability to display videos or slide shows. The small pieces of software code that enable users to embed content from one site into a Web page or another Web site are called widgets. **Widgets** are small software programs that can be added to Web pages or placed on the desktop to provide additional functionality. For example, the Flixter widget on Facebook profiles transports users to a place where they can list the films they have seen along with their ratings and reviews, view their friends' ratings and reviews, and what is playing in theatres.

Mashups
Widgets

Web widgets run inside a Web page or blog. Desktop widgets integrate content from an external source into the user's desktop to provide services such as a calculator, dictionary, or display of current weather conditions. The Apple Dashboard, Microsoft Windows Vista Sidebar, and Google Desktop Gadgets are examples of desktop widgets.

Widgets also provide storefront windows for advertising and selling products and services. Random House Inc. has a widget that enables visitors to its Web site to click through to purchase new book releases from its online store. Amazon.com and Walmart have toolbar widgets that enable surfers to search their Web stores while staying on their social network or another personal page. Widgets have become so powerful and useful that Facebook and Google launched programs to attract developers of widgets for their Web sites.

Software Outsourcing

Today most business firms continue to operate legacy systems that continue to meet a business need and that would be extremely costly to replace. But they will purchase most of their new software applications from external sources. There are three external sources for software: software packages from a commercial software vendor, software services from an online service provider, and outsourcing custom application development to an outside software firm, often offshore firms in low-wage areas of the world.

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, supply chain management, and finance to 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 12 discusses enterprise systems in detail.

Software as a Service (SaaS) It is clear that software will increasingly be delivered and used over networks as a service. Earlier in this chapter, we described cloud computing, in which software is delivered as a service over the Internet. 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 on either 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, including salesforce automation, partner relationship management, marketing, and customer service. It includes tools for customization, integrating its software with other corporate applications, and integrating new applications to run other parts of the business. The Window on Organizations provides more detail on these capabilities.

Companies considering the SaaS model need to carefully assess the costs and benefits of the service, weighing all people, organizational, and technology issues, including the ability to integrate with existing systems and deliver a level of service and performance that is acceptable for the business. In some cases, the cost of renting software will add up to more than purchasing and maintaining the application in-house. Yet there may be benefits to paying more for software as a service if this decision allows the company to focus on core business issues instead of technology challenges.

Software package
Software as a Service (SaaS)
Outsourcing

Software Outsourcing A third external source of software is **outsourcing,** in which a firm contracts custom software development or maintenance of existing legacy programs to outside firms, frequently firms that operate offshore in low-wage areas of the world.

According to the Gartner Group, worldwide outsourcing totalled over $539 billion in 2008, and it is growing at about 8 percent a year (Gartner, 2008). The largest expenditure here is paid to North American firms providing middleware, integration services, and other software support that are often required to operate larger enterprise systems. IBM has the largest share of this global market (8 percent) followed by EDS (5 percent) and ADP (3 percent). About 50 percent of all global outsourcing originates in the financial services sector.

For example, in March 2008, Royal Dutch Shell PLC, the world's third-largest oil producer, signed a five-year, $4.88 billion outsourcing deal with T-Systems International GmbH, AT&T, and Electronic Data Systems (EDS). The agreement assigned AT&T responsibility for networking and telecommunications, T-Systems for hosting and storage, and EDS for end-user computing services and integration of the infrastructure services. Outsourcing this work will help Shell cut costs and focus on systems that improve its competitive position in the oil and gas market.

Offshore outsourcing firms have primarily provided lower-level maintenance, data entry, and call centre operations. However, with the growing sophistication and experience of offshore firms, particularly in India, more and more new-program development is taking place offshore. Chapter 9 discusses offshore software outsourcing in greater detail.

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

Service-level agreement (SLA)

Organizations

WINDOW ON ORGANIZATIONS

Salesforce.com: Software as a Service Goes Mainstream

Salesforce.com has been considered one of the most disruptive technology companies of the past few years and is credited with single-handedly shaking up the software industry with its innovative business model and resounding success. This company provides CRM solutions in the form of software as a service leased over the Internet, as opposed to software bought and installed on machines locally. It was founded in 1999 by former Oracle executive Marc Benioff and has since grown to 2600 employees and earned $913 million in revenue in 2007. Salesforce.com has more than 43 000 corporate customers and well over 1 million subscribers.

The company attributes its success to the many benefits of its on-demand model of software distribution. The on-demand model eliminates the need for large upfront capital investments in systems and lengthy implementations on corporate computers. Subscriptions start as low as $11 per month per user for the pared-down Group version for small sales and marketing teams, with monthly subscriptions for more advanced versions for large enterprises starting around $80 per user. Salesforce.com implementations take up to three months. There is no hardware for subscribers to purchase, scale, and maintain; no operating systems, database servers, or application servers to install; no consultants and staff; and no expensive licensing and maintenance fees. The system is accessible via a standard Web browser, and Salesforce.com continually updates its software behind the scenes. There are tools for customizing some features of the software to support a company's unique business processes. Salesforce.com's solutions offer better scalability than those provided by large enterprise software vendors because they eliminate the cost and complexity of managing multiple layers of hardware and software.

Benioff believes that all of these advantages will inevitably lead to "the end of software," or, more appropriately, a new future of software in which the software-as-a-service model will supplant the current model and become the new paradigm. However, it is still too soon to tell whether this prediction will turn out to be true. Salesforce faces significant challenges as it continues to grow and refine its business.

The first challenge comes from increased competition, both from traditional industry leaders and new challengers hoping to replicate Salesforce's success. Microsoft, SAP, and Oracle have rolled out subscription-based versions of their CRM products in response to Salesforce. Smaller competitors like NetSuite also have made some inroads against Salesforce's market share. In addition, Siebel Systems, the world's leading CRM supplier bought in 2005 by Oracle, has been a major competitor since Salesforce.com was founded.

Analysts predict that Microsoft has a chance to compete with Salesforce by developing merely an acceptable on-demand CRM product because of the average customer's already-established familiarity with Microsoft applications. Also Microsoft plans to offer its product at half the price of Salesforce.com, using a tactic it has employed with great effect in other marketplaces to pressure competitors. Salesforce.com still has plenty of catching up to do to reach the size and market share of its larger competitors. As of 2007, SAP's CRM market share was 25.7 percent, compared to only 7 percent for Salesforce.com. IBM's customer base includes 9000 software companies that run their applications on their software and that are likelier to choose a solution offered by IBM over Salesforce.com.

Another challenge for Salesforce.com is to expand its business model into other areas. Salesforce is currently used mostly by sales staff needing to keep track of leads and customer lists. One way the company is trying to provide additional functionality is through a partnership with Google and more specifically Google Apps. Salesforce.com is combining its services with Gmail, Google Docs, Google Talk, and Google Calendar to allow its customers to accomplish more tasks via the Web.

The partnership between Salesforce.com and Google represents a united front against Microsoft intended to cut into the popularity of Microsoft Office. Currently, Salesforce.com describes the partnership as "primarily a distribution deal," but it could grow stronger based on the idea that businesses prefer to manage CRM in one place. Salesforce.com and Google both hope that their Salesforce.com for Google Apps initiative will galvanize further growth in on-demand software.

Salesforce opened up its Force.com application development platform to other independent software developers and listed their programs on its AppExchange. Using AppExchange, small businesses can go online and easily download more than 800 software applications, some add-ons to Salesforce.com and others that are unrelated. 24 Hour Fitness, the world's largest privately owned and operated fitness chain, uses App Exchange along with Salesforce.com's Enterprise edition for companywide salesforce automation and customer service. One of its App Exchange applications integrates Hoover's database of up to 21 million companies and 28 million executives with Salesforce, and another allows users to easily create and distribute on-demand surveys and response forms.

The question is whether the audience for the App Exchange application platform will prove large enough to deliver the level of growth Salesforce wants. Some analysts believe the platform may not be attractive to larger companies for their application needs.

A third challenge is availability. Salesforce.com subscribers depend on the service being available 24/7. But occasional outages have occurred (see the Chapter 8 Window on Technology), making some companies rethink their dependency on software as a service. Salesforce.com provides tools to assure customers about its system reliability and also offers PC applications that tie into its services so users can work offline.

To Think About

1. What are the advantages and disadvantages of the software-as-a-service model?
2. What are some of the challenges facing Salesforce as it continues its growth? How well will it be able to meet those challenges?
3. What kinds of businesses could benefit from switching to Salesforce and why?
4. What factors would you take into account in deciding whether to use Saleforce.com for your business?

MIS in Action

Explore the Salesforce.com Web site. Go to the App Exchange portion of the site, and examine the applications available for each of the categories listed. Then answer the following questions:

1. What are the most popular applications on App Exchange? What kinds of processes do they support?
2. Could a company run its entire business using Salesforce.com and App Exchange? Explain your answer.
3. What kinds of companies are most likely to use App Exchange? What does this tell you about how Salesforce.com is being used?

Sources: J. Nicholas Hoover, "Service Outages Force Cloud Adopters to Rethink Tactics," *InformationWeek*, August 16, 2008, **www.informationweek.com/news/services/saas/showArticle.jhtml?articleID=210004236** (accessed November 19, 2009); Jay Greene, "Google and Salesforce: A Tighter Bond," *BusinessWeek*, April 15, 2008; Mary Hayes Weier, "Salesforce, Google Show Fruits of Their Collaboration,"*InformationWeek*, April 21, 2008; John Pallatto and Clint Boulton, "An On-Demand Partnership," *eWeek*, April 21, 2008; Gary Rivlin, "Software for Rent," *The New York Times*, November 13, 2007; Steve Hamm, "A Big Sales Job for Salesforce.com," *BusinessWeek*, September 24, 2007; Mary Hayes Weier, "Salesforce.com" and Marianne Kolbasuk McGee, "Salesforce as B-to-B Broker," *InformationWeek*, December 10, 2007; and Salesforce.com, report on Form 10-K for the fiscal year ended January 31, 2008, filed with the SEC on 2/29/08.

5.5 Management Issues

Creating and managing a coherent IT infrastructure raises multiple challenges: dealing with platform and technology change, management and governance, and making wise infrastructure investments.

Dealing with Platform and Infrastructure Change

As firms grow, they can 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 licences? How well does the infrastructure scale? **Scalability** refers to the ability of a computer, product, or system to expand to serve a large number of users without breaking down. New applications, mergers and acquisitions, and changes in business volume all impact computer workload and must be considered when planning hardware capacity.

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

Management and Governance

A long-standing issue among information system managers and CEOs has been the question of who will control and manage the firm's IT infrastructure. Chapter 2 introduced the concept of IT governance and described some issues it addresses. Other important questions about IT governance include: 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 firm 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 underinvesting 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 its own IT infrastructure components or rent them from external suppliers. As we discussed earlier, a major trend in computing platforms—both hardware and software—is to outsource to external providers. The decision either to purchase your own IT assets or rent them from external providers is typically called the *rent versusbuy* decision.

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. Although similar to Porter's competitive forces model, which we discussed in Chapter 3, this model has six components as compared to Porter's original five forces. The following section describes these six forces or factors and their impact on the firm's IT services and infrastructure.

Scalability

FIGURE 5-12 Competitive forces model for IT infrastructure.

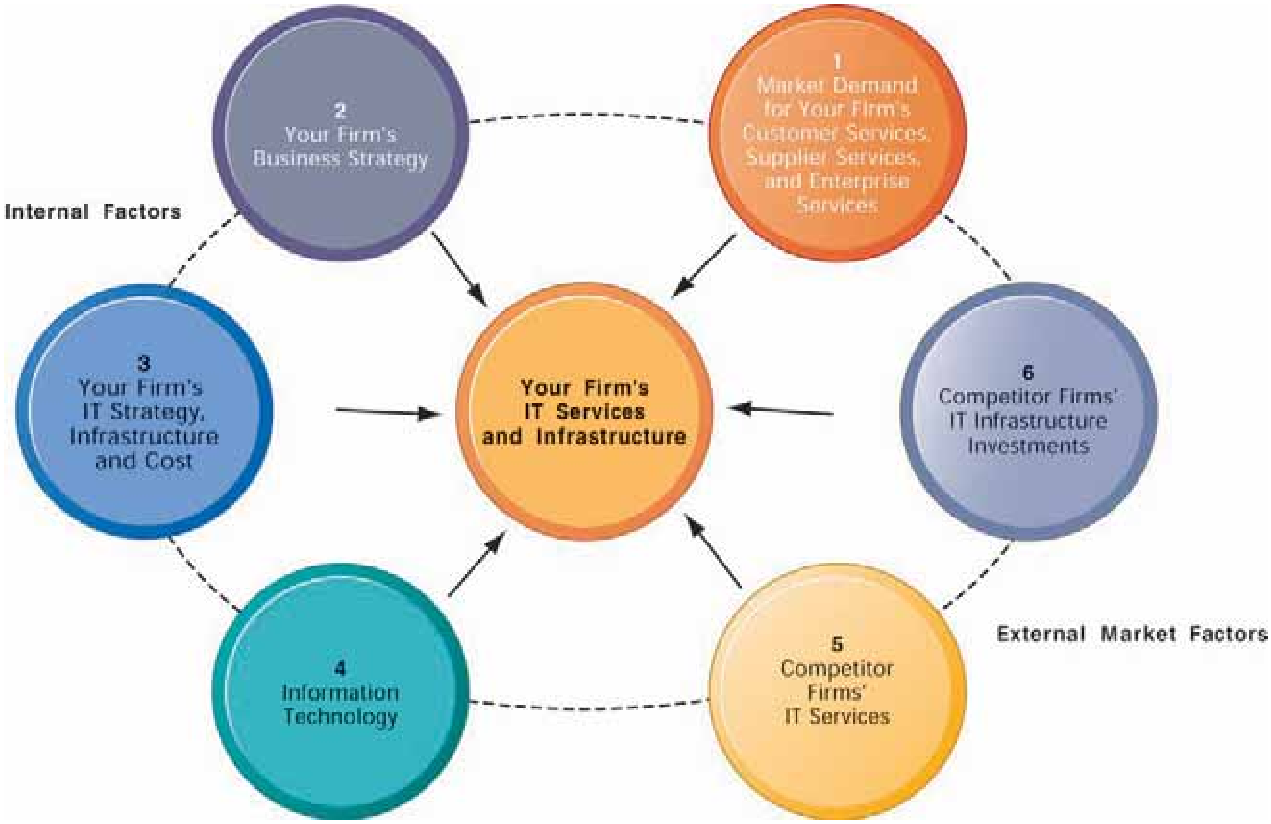

There are six factors you can use to answer the question, "How much should our firm spend on IT infrastructure?"

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

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

Your Firm's Information Technology (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 total IT infrastructure costs. You will want to perform a total cost of ownership analysis (see the discussion later). If your firm has no IT strategy, you will need to devise one that takes into account the firm's five-year strategic plan.

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

TABLE 5-4 Total cost of ownership (TCO) cost components.

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

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

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

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

Total Cost of Ownership of Technology Assets

In benchmarking your firm's expenditures on IT infrastructure with that of your competitors, you will need to consider a wide range of costs. The actual cost of owning technology resources includes the original cost of acquiring and installing hardware and software, as well as ongoing administration costs for hardware and software upgrades, maintenance, technical support, training, and even utility and real estate costs for running and housing the technology. The **total cost of ownership (TCO)** model can be used to analyze these direct and indirect costs to help firms determine the actual cost of specific technology implementations. Table 5-4 describes the most important TCO components to consider in a TCO analysis. Figure 5-13 illustrates the concept of TCO.

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. Hidden costs for support staff, downtime, and additional network management can make distributed client/server architectures—especially those incorporating mobile wireless devices—more expensive than centralized mainframe architectures. Hardware and software acquisition costs account for only about 20 to 40 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.

These firms could reduce their TCO through greater centralization and standardization of their hardware and software resources, as did Bell Canada, described in the

Total cost of ownership (TCO)

FIGURE 5-13 The Iceberg Model for total cost of ownership

An iceberg illustrates the concept of TCO, where most of the content is hidden.

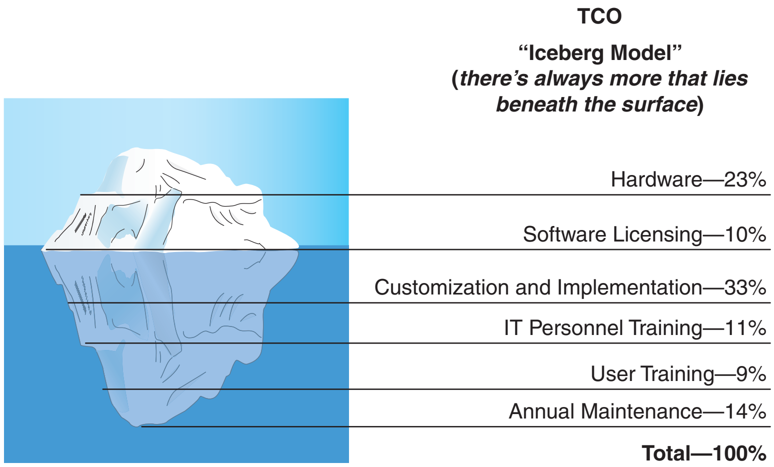

chapter-opening case. Companies could reduce the size of the information systems staff required to support their infrastructure if the firm minimizes the number of different computer models and pieces of software that employees are allowed to use. In a centralized infrastructure, systems can be administered from a central location and troubleshooting can be performed from that location.

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/Internet computing era, and the cloud computing era. Drivers include Moore's Law, the law of mass data storage, and Metcalfe's Law. 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, implying that the price of computing falls in half. The Law of Mass Digital Storage deals with the exponential decrease in the cost of storing data, stating that the number of kilobytes of data that can be stored on magnetic media for $1 roughly doubles every 15 months. Metcalfe's Law helps shows that a network's value to participants grows exponentially as the network takes on more members. Also driving exploding computer use is the rapid decline in costs of communication and growing agreement in the technology industry to use computing and communications standards.

3. ***What are the current trends in computer hardware platforms?***
   The emerging trends in mobile digital computing platform, grid computing, mobile platforms, on-demand cloud computing, and autonomic computing demonstrate that, increasingly, computing is taking place over a network. 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 with which to attack large computing problems. Cloud computing is a model of computing in which firms and individuals obtain computing power and software applications over the Internet, rather than purchasing and installing the hardware and software on their own computers. In autonomic computing, computer systems have capabilities for automatically configuring and repairing themselves.

   Virtualization organizes computing resources so that their use is not restricted by physical configuration or geographic location. Server virtualization enables companies to run more than one operating system at the same time on one computer. A multicore processor is a microprocessor to which two or more processors have been attached for enhanced performance, reduced power consumption, and more efficient simultaneous processing of multiple tasks.

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

Contemporary software platform trends include the growing use of Linux, open source software, Java and Ajax, Web services, mashups and widgets, and software outsourcing. Open source software is produced and maintained by a global community of programmers and is 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. Ajax allows a client and server to exchange small pieces of data behind the scenes so that an entire Web page does not have to be reloaded each time the user requests a change.

Web services are loosely coupled software components based on open Web standards that are not product-specific and can 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. Mashups and widgets are the building blocks of new software applications and services based on the cloud computing model. 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 software services (SaaS).

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

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

Key Terms

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. What is Metcalfe's Law?*
- *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, 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 Ajax, and explain why they are important.*
- *Define and describe Web services and the role played by XML.*
- *Define and describe software mashups and widgets.*
- *Name and describe the three external sources for software.*

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

- *Name and describe the management challenges posed by IT infrastructure.*
- *Explain how using a competitive forces model and calculating the total cost of ownership (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?*

Collaboration and Teamwork: Evaluating Server Operating Systems

Form a group with three or four of your classmates. One group should research and compare the capabilities and costs of Linux versus the most recent version of the Windows operating system for servers. Another group should research and compare the capabilities and costs of Linux versus Unix. If possible, use Google Sites to post links to Web pages, team communication announcements, and work assignments; to brainstorm; and to work collaboratively on project documents. Try to use Google Docs to develop a presentation of your findings for the class.

Visit the MyMISLab Web site at **www.pearsoned.ca/mymislab.** This online homework and tutorial system puts you in control of your own learning with study and practice tools directly correlated to this chapter's content.

Learning Track Module

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

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 Vancouver Medical Centre (UVMC)—a fictional medical centre—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 has been growing by 20 percent each year. UVMC was setting up a separate server for every application, and its servers and other computers were running a number of different operating systems, including several versions of Unix and Windows. UVMC had to manage technologies from many different vendors, including Hewlett-Packard (HP), Sun Microsystems, Microsoft, and IBM. Assess the impact of this situation on business performance. What factors and management decisions must be considered when developing a solution to this problem?
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. 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 alternative desktop systems.

You have been asked to obtain pricing information on hardware and software for an office of 30 people. Using the Internet, get pricing for 30 PC desktop systems (monitors, computers, and keyboards) manufactured by Lenovo, Dell, and HP/Compaq as listed at their respective corporate Web sites. (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 Hewlett-Packard, Canon, and Dell. Each desktop system must satisfy the minimum specifications shown in the following table:

Minimum Desktop Specifications

Processor speed	Dual-core 2 GHz
Hard drive	250 GB
RAM	3 GB
DVD-ROM drive	16 x
Monitor (diagonal measurement)	17 inches

Each desktop printer must satisfy the minimum specifications shown in the following table:

Minimum Monochrome Printer Specifications

Print speed (black and white)	20 pages per minute
Print resolution	600 × 600
Network-ready?	Yes
Maximum price/unit	$700

After pricing the desktop systems and printers, obtain pricing on 30 copies of the most recent versions of Microsoft Office, Lotus SmartSuite, and Sun Star Office desktop productivity packages and on 30 copies of Microsoft Windows Vista Business. The application software suite packages come in various versions, so be sure that each package contains programs for word processing, spreadsheet analysis, database analysis, graphics preparation, and e-mail.

Prepare a spreadsheet showing your research results for the desktop systems, printers, and software. Use your spreadsheet software to determine the desktop system, printer, and software combination that will offer both the best performance and pricing per worker. Because every two workers will share one printer (15 printers/30 systems), assume only half a printer cost per worker in the spreadsheet. Assume that your company will take the standard warranty and service contract offered by each product's manufacturer.

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

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

In this exercise, you will use software at various online travel sites to arrange transportation and lodging for a large sales force to attend a sales conference at two alternative locations. You will use that information to calculate total travel and lodging costs and decide where to hold the conference.

The fictional Foremost Composite Materials Company is planning a two-day sales conference for October 15–16, starting with a reception on the evening of October 14. 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 $105 000 for the representatives' room rentals. The hotel must also have such services as overhead and computer projectors as well as a business centre and banquet facilities. It also should have facilities for the company representatives to be able to work in their rooms and to enjoy themselves in a swimming pool or gym facility. The company would like to hold the conference in either Whistler, BC or Montreal, QC.

Foremost usually likes to hold such meetings in Hilton- or Marriott-owned hotels. 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.

Link to the two sites' home pages, and search them to find a hotel that meets Foremost's sales conference requirements. Once you have selected the hotel, locate flights arriving the afternoon prior to the conference because the attendees will need to check in the day before and attend your reception the evening 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 $60 per attendee per day for food.

1. What was your final budget?
2. Which did you select as the best hotel for the sales conference and why?

CASE STUDY Amazon's New Store: Utility Computing

Looking for a good deal on that DVD box set of *The West Wing* or the last *Harry Potter* book? Since opening as an online bookstore in 1995, Amazon.com has morphed into a virtual superstore with product offerings in 36 categories, including furniture, jewellery, clothing, and groceries. But what if what you really need is a place to store several terabytes of data? Or the computing power of 100 Linux servers? Now you can get those from Amazon too.

Over its first 12 years, Amazon.com committed $2 billion to refine the information technology infrastructure that was largely responsible for making it the top online retailer in the world. Following the burst of the dot-com bubble in 2001, Amazon focused heavily on modernizing its data centres and software so that it could add new features to its product pages such as discussion forums and software for audio and video.

In March 2006, Amazon introduced the first of several new services that founder Jeff Bezos hoped would transform its future business. With Simple Storage Service (S3) and later Elastic Compute Cloud (EC2), Amazon entered the cloud utility computing market. The company had realized that the benefits of its $2 billion investment in technology could also be valuable to other companies.

Amazon had tremendous computing capacity, but like most companies, used only a small portion of it at any one time. Moreover, the Amazon infrastructure was considered by many to be among the most robust in the world. So, the one-time bookseller exposed the guts of its entire system over the Internet to any developer who could make use of it. Amazon began to sell its computing power on a per-usage basis, just as a power company sells electricity.

S3 is a data storage service that is designed to make Web-scale computing easier and more affordable for developers. Customers pay 15 cents per gigabyte of data stored per month on Amazon's network of disk drives. There is also a charge of 20 cents per gigabyte of data transferred. The service has neither a minimum fee nor a startup charge. Customers pay for exactly what they use and no more. Data may be stored as objects ranging in size from 1 byte to 5 gigabytes, with an unlimited number of objects permitted. Using S3 does not require any client software, nor does it require the user to set up any hardware. Amazon designed S3 to provide a fast, simple, and inexpensive method for businesses to store data on a system that is scalable and reliable. S3 promises 99.99 percent availability through a mechanism of fault tolerance that fixes failures without any downtime.

Working in conjunction with S3, EC2 enables businesses to utilize Amazon's servers for computing tasks, such as testing software. Using EC2 incurs charges of 10 cents per instance-hour consumed. An instance supplies the user with the equivalent of a 1.7 GHz x86 processor with 1.75 GB of RAM, a 160 GB hard drive, and 250 megabits per second of bandwidth on the network. The service also includes 20 cents per GB of data traffic inbound and outbound per month, as well as the standard S3 pricing for storing an Amazon Machine Image (AMI), which contains the applications, libraries, data, and configuration settings that a business uses to run its processes.

According to Adam Selipsky, vice president of product management and developer relations for Amazon Web Services (AWS), Amazon is really a technology company that can bring a wealth of engineering prowess and experience to independent developers and corporations by allowing them to run their processes on Amazon's computer systems. Selipsky also emphasizes that AWS is not simply about providing great amounts of storage capacity and server time. AWS creates the opportunity for others to work at Web scale without making the mistakes that Amazon has already made and learned from. *Simplicity* and *ease of use* are not generally terms that go along with building a Web-scale application, but they are major selling points for AWS. Users build on the services through Application Programming Interfaces (APIs) made available by Amazon.

From the very beginning, customers have responded strongly to S3 and EC2. Bezos targeted micro-sized businesses and Web startups as customers for AWS, but the services have also attracted some mid-size businesses and potential big players in e-business.

MileMeter Inc. is a Dallas-based startup that plans to sell auto insurance by the mile. It initially ran its own server in a data centre but moved most of its applications onto "virtual" computers in Amazon's EC2. CEO Chris Gay said, "I don't need to have a systems administrator or a network administrator. I don't have to worry about hardware becoming irrelevant."

Webmail.us provides e-mail management services for thousands of companies around the world from its Blacksburg, Virginia, headquarters. When the company needed to increase its short-term storage capacity and the redundancy of its primary data backups, it selected S3 as its storage provider. Webmail.us sends more than a terabyte of data to Amazon to store with S3 every week. Bill Boebel, cofounder and chief technology officer of Webmail.us, was very pleased that his company was able to create a simple interface with which Amazon can accept the abundant small files that his company manages. Other backup systems have had difficulty handling the typical Webmail.us backup load, and most hosting companies would require a custom application to handle such data. Webmail.us even used EC2 to develop its storage interface. According to CEO Pat Matthews, Amazon immediately reduced his company's data backup costs by 75 percent.

Powerset is an up-and-coming search engine company based in San Francisco that wants to focus its time and the $15.25 million it has raised on its core business, natural language search technology. By using S3 and EC2, Powerset saves upfront cash expenditures and eliminates the risk that building an infrastructure will take longer than expected. Many of the traditional utility computing suppliers charge around one dollar per CPU hour, or ten times what Amazon charges.

Powerset's CEO Barney Pell says that the pay-as-you go model is very important because his company does not know how fast it will grow. What he does know is that the demand for Powerset's service will come in bursts, and trying to predict hardware needs is a dangerous game. If Powerset overestimates its peak usage capacity, the company will waste money on unnecessary hardware. If the company underestimates peak usage, it could fail to meet its users' expectations and damage its business. With AWS in place, Powerset never has to worry about being unable to add computing power when a spike in usage occurs.

SmugMug Inc., an online photo-sharing startup, was immediately drawn to the ease with which it could back up photos on Amazon's S3. Storing its users' photos on Amazon's devices prevents SmugMug from having to purchase its own additional storage and saved the company $1 million during the first year it used Amazon's services.

As with any large business initiative, there are issues for Amazon to confront before anyone can declare AWS to be a successful venture. Larger businesses may be more inclined to use a more established company, especially one with more experience hosting core applications and data. Currently, Amazon's flexible, pay-as-you-go model gives the company a competitive advantage over companies that require service contracts.

However, according to Daniel Golding, vice president of Tier 1 Research, the established companies, such as IBM, Hewlett-Packard, and Sun Microsystems, may follow Amazon's lead and offer utility computing without service-level agreements (SLAs). Complicating the matter is that some companies are wary of using a supplier that does not offer SLAs, which guarantee the availability of services in terms of time. Golding suggests that Amazon may have launched a major shift in the industry, but others may reap the rewards while Amazon may suffer for it.

One more challenge for Amazon is the viability of AWS itself. Will the services actually function as planned? The company's track record with new technology projects is mixed. Amazon launched its A9.com search site with much fanfare, but the site never really caught on with users. Moreover, the growth of AWS could be harmful to Amazon's Web services line as well as to its retail line if Amazon does not position itself to handle a dramatic increase in demand on its infrastructure. AWS customers could drop the service, and Amazon.com could falter.

In January 2007, February 2008, and July 2008, Amazon's S3 servers experienced significant outages, with service lost for eight hours in July. The January 2007 problem was caused by faulty hardware installed during an upgrade, and it was resolved quickly. The July 2008 outage was more problematic. Amazon reported that components were unable to interact properly due to a problem with "internal system communications." Amazon promised to provide a fuller explanation once it determined the root cause. Some users were critical and questioned whether Amazon was capable of being their solution for hosted storage going forward.

AWS has charmed some high-profile clients. Microsoft uses S3 to increase software download speeds for its users. Linden Lab, creator of the online virtual world Second Life, uses the service to alleviate the pounding its servers take when the company releases its frequent software upgrades. The Nasdaq stock exchange uses S3 to host data for Nasdaq Market Replay, an application that lets companies play back historical market data in real time. However, Nasdaq is reluctant to use an online service for transactional or highly secure data.

For Canadians wishing to use Amazon's Web services and S3 architecture and applications, Amazon presents a double-edged sword. The services appear to have great potential, but issues such as currency and tax conversion and transborder data laws (see Chapter 4) and regulations may prevent Canadian companies from using Amazon's service because its servers and services are all resident in the United States, not in Canada.

To better support large accounts, Amazon started round-the-clock phone support and credits if S3 availability falls below 99.99 percent in a single month. For now, the potential of AWS is being converted into performance mostly by tech-savvy developers with financial backing. More than 370 000 developers, ranging from individuals to large companies, have signed up. As more developers contribute and the services evolve, Amazon hopes one day to make it possible for anyone with an idea and an Internet connection to begin to put together the next Amazon.com.

CASE STUDY QUESTIONS:

1. What technology services does Amazon provide? What are the business advantages to Amazon and to subscribers of these services? What are the disadvantages to each? What kinds of businesses are likely to benefit from these services?
2. 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.
3. Search the Internet for companies that supply utility computing. Select two or three such companies, and compare them to Amazon. What services do these companies provide? What promises do they make about availability? What is their payment model? Who is their target client? If you were launching a Web startup business, would you choose one of these companies over Amazon for Web services? Why or why not? Would your answer change if you were working for a larger company and had to make a recommendation to the CTO?
4. Name three examples of IT infrastructure hardware components and software components that are relevant to this case. Describe how these components fit into or are used by Amazon's Web services and/or the customers that subscribe to these services.
5. Think of an idea for a Web-based startup business. Explain how this business could utilize Amazon's S3 and EC2 services.
6. What are the legal, ethical, and privacy issues for Canadian businesses that might want to use the AWS and S3 services offered by Amazon? How would that affect the potential for these services in Canada? Is there a solution to this problem?

Sources: Thomas Claburn, "Amazon's S3 Cloud Service Turns into a Puff of Smoke," *InformationWeek*, July 28, 2008; Chris Preimesberger, "Perils in the Cloud," *eWeek*, August 4, 2008; Jessica Mintz, "Amazon's Hot New Item: Its Data Center," Associated Press, February 1, 2008; J. Nicholas Hoover, "Ahead in the Cloud: Google, Others Expand Online Services," *InformationWeek*, April 14, 2008; J. Nicholas Hoover and Richard Martin, "Demystifying the Cloud," *InformationWeek*, June 23, 2008; Edward Cone, "Amazon at Your Service," *CIO Insight*, January 7, 2007; Robert D. Hof, "So You Wanna Be a Web Tycoon? Amazon Can Help," **www.webworkerdaily.com**, January 24, 2007; and Thomas Claburn, "Open Source Developers Build on Amazon Web Services," TechWeb.com, January 12, 2007.