# chapter 4

# Managing the Information Systems Infrastructure

## preview >

As any city depends on a functioning infrastructure, companies operating in a digital world are relying on a comprehensive information systems infrastructure to support their business processes and competitive strategy. With ever-increasing speed, transactions are conducted; likewise, with ever-increasing amounts of data to be captured, analyzed, and stored, companies have to thoroughly plan and manage their infrastructure needs in order to gain the greatest returns on their information systems investments. When planning and managing their information systems architectures, organizations must answer many important and difficult questions. For example, how will we utilize information systems to enable our competitive strategy? What technologies and systems best support our core business processes? Which vendors should we partner with, which technologies do we adopt, and which do we avoid? What hardware, software, or services do we buy, build, or have managed by an outside service provider? How can the organization get the most out of the data captured from internal and external sources? How can the organization best assure that the infrastructure is reliable and secure? Clearly, effectively managing an organization's information systems infrastructure is a complex but necessary activity in today's digital world. After reading this chapter, you will be able to do the following:

1. List the essential information systems infrastructure components and describe why they are necessary for satisfying an organization's informational needs.
2. Describe solutions organizations use to design a reliable, robust, and secure infrastructure.
3. Describe how organizations can ensure a reliable and secure infrastructure, plan for potential disasters, and establish IS controls.

This chapter focuses on helping managers understand the key components of a comprehensive information systems infrastructure and why its careful management is necessary. With an increasing complexity of an organization's information needs and an increasing complexity of the systems needed to satisfy these requirements, the topic of infrastructure management is fundamental for managing in the digital world.

## Managing in a Digital World: "I Googled You!"

You're researching a paper for a physics class, and you need information on quarks. Google it (see Figure 4.1). You'd like to locate a high school classmate, but no one in your graduating class knows where she is. Google her. You're watching a movie, and a character says she "googled" a blind date. The term "google" has become so familiar to Internet users that it's often used as a verb. In fact, the term has become so common that Google is becoming concerned that its use as a verb is a copyright infringement, asking dictionaries such as *Merriam-Webster* to change their definition of Google to "to use the Google search engine to obtain information . . . on the World Wide Web."

According to the Google.com Web site, "Google is a play on the word googol, which was coined by Milton Sirotta, nephew of American mathematician Edward Kasner, and was popularized in the book *Mathematics and the Imagination* by Kasner and James Newman. It refers to the number represented by the numeral 1 followed by 100 zeros. Google's use of the term reflects the company's mission to organize the immense, seemingly infinite amount of information available on the web."

According to Google lore, company founders Larry Page and Sergey Brin argued about everything when they first met as Stanford University graduate students in computer science in 1995. Larry was a 24-year-old University of Michigan alumnus on a weekend visit; Sergey, 23, was among a group of students assigned to show him around. Both had strong opinions and divergent viewpoints, but they eventually found common ground in a unique approach to solving one of computing's biggest challenges: retrieving relevant information from a massive set of data.

By January 1996, Page and Brin had begun collaboration on a search engine called BackRub, named for its unique ability to analyze the "back links" pointing to a given Web site. Page, who had always enjoyed tinkering with machinery and had gained some notoriety for building a working printer out of Lego™ bricks, took on the task of creating a new kind of server environment that used low-end PCs instead of big expensive machines. Afflicted by the perennial shortage of cash common to graduate students everywhere, the pair took to haunting the department's loading docks in hopes of tracking down newly arrived computers that they could borrow for their network.

In 1998, Page and Brin were still operating out of a dorm room. They maxed out credit cards buying a terabyte of memory to hold their data and went looking for investors to help them further develop their search engine technology. David Filo, a friend and one of the developers of Yahoo!, told the two their technology was solid and convinced them to start up their own company.

**FIGURE 4.1**

Google products page.

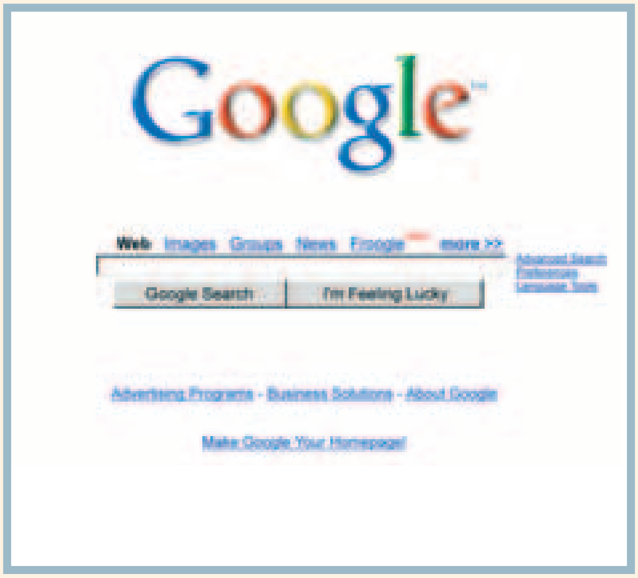


Page and Brin put out feelers for investors and found Andy Bechtolsheim, a friend of a faculty member, who wrote them a check for $100,000 after one brief meeting. Since the check was made out to "Google Inc.," Page and Brin scrambled to establish a corporation so they could deposit the check. Other investors joined, and Google Inc. began operations in September 1998 in Menlo Park, California—in a friend's garage that included a washer and dryer and a hot tub. The first employee hired was Craig Silverstein, director of technology.

From the start, Google, still in beta in 1998, handled 10,000 search queries a day. The company quickly captured the attention of the press and was extolled in *USA Today, Le Monde,* and *PC Magazine,* which named Google the best search engine of 1998.

Google quickly outgrew its garage location, and by February 1999 the company had moved into an office in Palo Alto, California, and now had eight employees and was handling more than 500,000 search queries a day.

The company continued to expand, removed the "beta" label from the search engine in 1999, and that same year moved into the Googleplex, its current headquarters in Mountain View, California.

In May 2000, Google was already the world's largest search engine answering 18 million queries a day, and was awarded a Webby Award and a People's Voice Award for technical achievement. (By the end of 2000, Google was answering 100 million search queries a day.)

On April 29, 2004, Google filed with the Securities and Exchange Commission for its initial public offering (IPO). In an unprecedented move, the IPO was sold at auction in order to make the shares more widely available. Shares were priced at $85, and Google hoped to raise $3 billion from the initial offering. Expert opinions on the success of the auction were mixed. Some said the stock price was inflated; others said the stock would eventually tank. Experts who warned of doomsday, however, were eventually proved wrong. In December 2006, Google's stock was selling for $466 a share and was expected to go to $548 in the short term.

Google has continued to innovate and move beyond the search engine market. The company offers e-mail, instant messaging, and mobile text messaging services. Other Google services include an automated news site, a Web blogging site, free imaging software, and a site for programmers interested in creating new applications. In mid-2006, Google was poised to challenge PayPal in the Internet account business and to give eBay a run for its money in the online auction business.

Google's e-mail service, like the company itself, is unique. Launched in 2004 as "Gmail," it was available to newcomers only on invitation from someone who already had the service. Gmail incorporates e-mail and instant messaging so that users can e-mail in the traditional manner and/or visit in real time.

The highest revenue generator for Google is its AdSense program. This program allows any Web site to publish advertisments on each of its pages. The Web site publisher is paid every time someone clicks on an ad originating from that page. The AdSense program also lets Web site publishers determine how many people look at the site, the cost per click, click-through rates, and so on. The AdSense program can tailor the type of ads that are placed on a Web site—that is, publishers can block ads they don't want to appear, such as competitor ads, ads concerning death or war, and ads for "adult" material.

Another Google service popular with users is froogle, which uses Google search technology to let consumers search for and compare products by product type, price, and so on. Other features include the following:

- The Google News service automatically combines news clippings from various online newspapers and provides them on one page for users' convenience.
- Google scholar helps researchers search through publications.
- Google finance searches for finance-related news and stock information.
- Other specialized search capabilities include video search, image search, mail-order catalog search, book search, blog search, and university search.

In addition, these services can also be accessed through mobile phones using the Google mobile products.

Google has clearly become a significant presence on the Internet and in users' daily lives. Look for new Google products and services at http://labs.google.com/.

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

1. To what extent was Google's initial success limited by its infrastructure?
2. If you were asked to identify the ideal location for a new Google facility, where would it be? Explain.
3. How would you rank order the various infrastructure components described in this chapter in their importance to Google's success? Explain your rationale.

Sources:

Anonymous, "Google Milestones," **http://www.google.com/corporate/history.html**

Antone Gonsalves, "Google Testing Possible eBay Competition," *Information Week* (October 25, 2005), **http://www.informationweek.com/story/showArticle.jhtml?articleID=172900366**

Eric J. Sinrod, "Google in a Patent Pickle?," *C/Net News* (January 19, 2006), **http://news.com.com/Google+in+a+patent+pickle/2010-1071_3-6027546.html**

Thomas Claburn, "Feds Seek Google Search Records in Child Porn Investigation," *Information Week* (January 19, 2006), **http://www.informationweek.com/internet/showArticle.jhtml?articleID=177101999**

## The Information Systems Infrastructure

Any area where people live or work needs a supporting **infrastructure**, which entails the interconnection of all basic facilities and services enabling the area to function properly. The infrastructure of a city, for example, includes components such as streets, power, telephone, water, and sewage lines but also schools, retail stores, and law enforcement. Both the area's inhabitants and the businesses depend on that infrastructure; cities with a good infrastructure, for example, are considered more livable than cities with poorer infrastructure and are much more likely to attract businesses and residents (see Figure 4.2). Likewise, valuable employees often choose firms with better facilities, management, and business processes.

For organizations considering where to set up a new manufacturing plant, for example, such decisions are often based on the provision of such infrastructure. Indeed, many municipalities attempt to attract new businesses and industries by setting up new commercial zones with the necessary infrastructure. In some cases, specific infrastructure components are of special importance. One such example is search engine giant Google, which has data centers located all over the world to offer the best performance to its users. Google's newest data center is nearing completion in the small town of The Dalles, Oregon, located on the banks of the

**FIGURE 4.2**

A city's infrastructure is complex and interconnected.

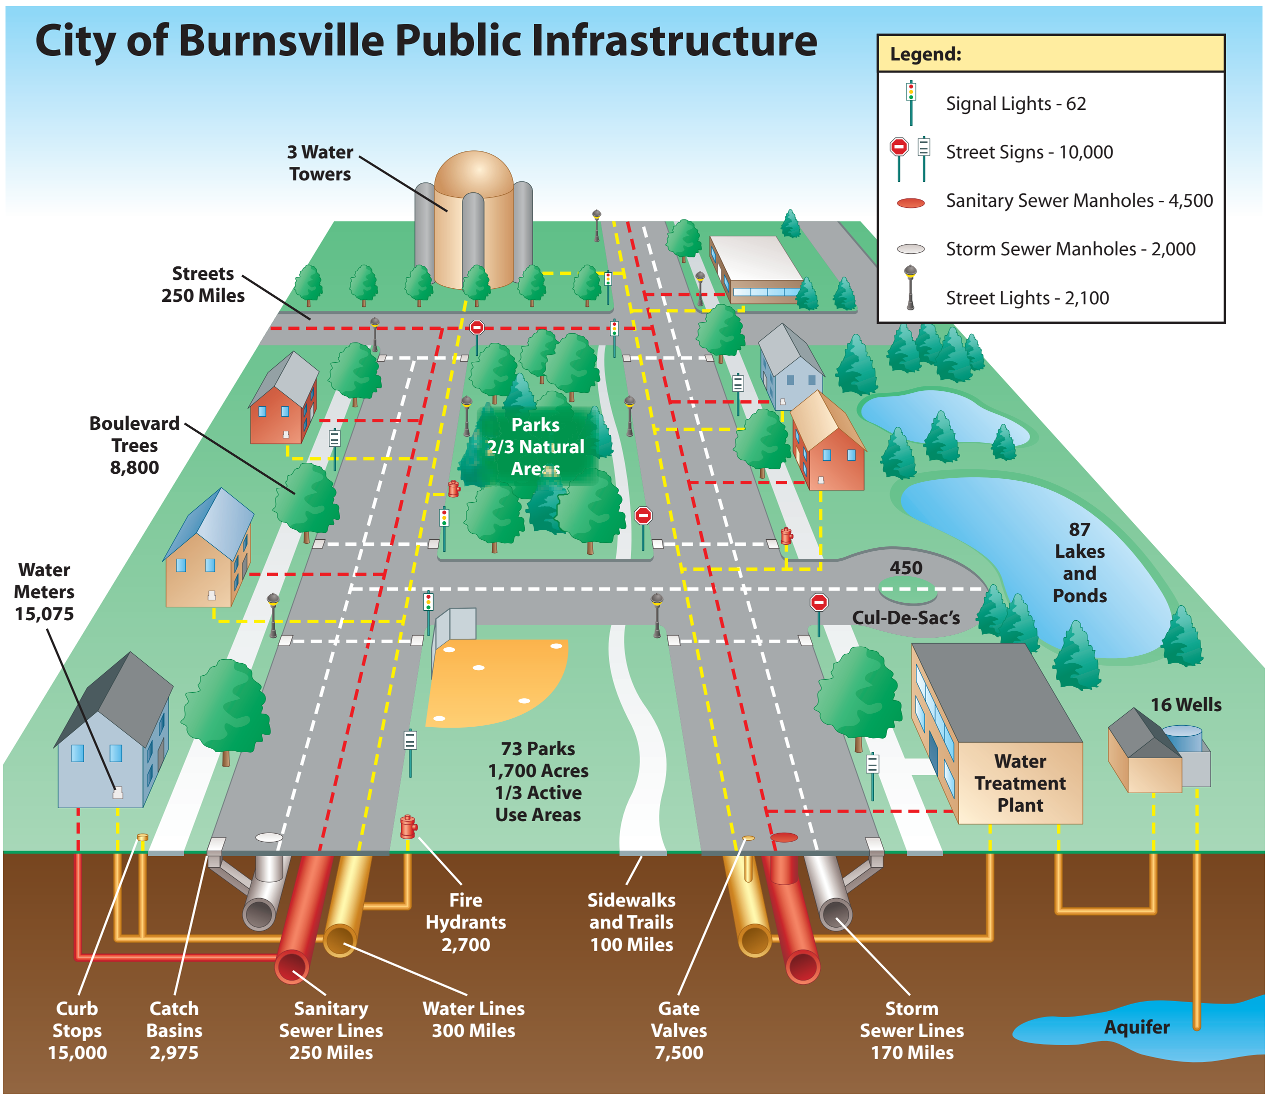


Source: http://www.burnsville.org/ftpfiles/infrabig.jpg.

**FIGURE 4.3**

Google data center in The Dalles, Oregon, is nearing completion.

Columbia River (see Figure 4.3). Why would a company such as Google choose such a rural location? First, the location offered connectivity, using a state-of-the-art fiber-optic network to provide high-speed data transfer to the Internet backbone (see Technology Briefing 4—Networking and Technology Briefing 5—The Internet and World Wide Web). Second—and maybe more important—the location on the river would give the data center access to water for its cooling needs and cheap, uninterrupted power from the nearby hydroelectric dam. As you can see from this example, companies such as Google must consider far more than just the need for increased data storage space and processing power.

For organizations operating globally, managing a comprehensive, worldwide infrastructure poses additional challenges. This is particularly acute when operating in developing nations. For example, in many parts of the world, organizations cannot count on an uninterrupted supply of water or electricity. Consequently, many of the large call centers in India that support customers around the world for companies like Dell Computers or Citibank have, for example, installed massive power generators to minimize the effects of frequent power outages or have set up their own satellite links to be independent from the local, unreliable phone networks.

## The Need for an Information Systems Infrastructure

As people and companies rely on basic infrastructures to function, businesses also rely on an **information systems infrastructure** (consisting of hardware, software, networks, data, facilities, human resources, and services) to support their decision making, business processes, and competitive strategy. **Business processes** are the activities that organizations perform in order to reach their business goals and consist of core processes and supporting processes. The **core processes** make up the primary activities in the value chain; these are all the processes that are needed to manufacture goods, sell the products, provide service, and so on (see Chapter 3—Valuing Information Systems Investments). The **supporting processes** are all the processes that are needed to perform the value chain's supporting activities, such as accounting, human resources management, and so on (see Figure 4.4).

Almost all of an organization's business processes depend on the underlying information systems infrastructure, albeit to different degrees. For example, an organization's management needs an infrastructure to support a variety of activities, including reliable communication networks to support collaboration between suppliers and customers, accurate and timely data and knowledge to gain business intelligence, and information systems to aid decision making and support business processes. In sum, organizations rely on a complex, interrelated information systems infrastructure to effectively thrive in the ever-increasing, competitive digital world.

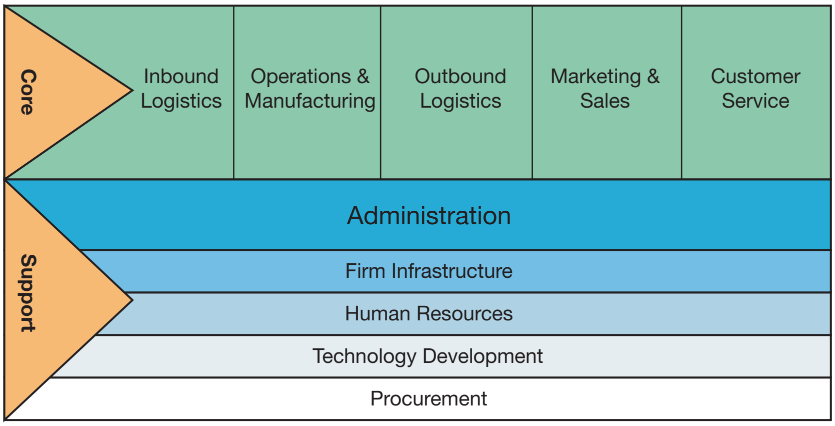


**FIGURE 4.4**

A generic value chain showing an organization's core and supporting activities.

In order to make better decisions, managers at all levels of the organizations need to analyze information gathered from the different business processes. The processes of gathering the information as well as the information itself are commonly referred to as **business intelligence**. Whereas some of these processes obtain the information from external sources—such as marketing research or competitor analysis—other processes gather business intelligence from internal sources, such as sales figures, customer demographics, or performance indicators. While there are a variety of different systems used for gaining business intelligence (see Chapter 7—Enhancing Business Intelligence Using Information Systems), all gather, process, store, or analyze data in an effort to better manage the organization. In other words, modern organizations rely heavily on their information systems infrastructure; its components include the following (see Figure 4.5):

- Hardware
- Software
- Communications and collaboration
- Data and knowledge
- Facilities

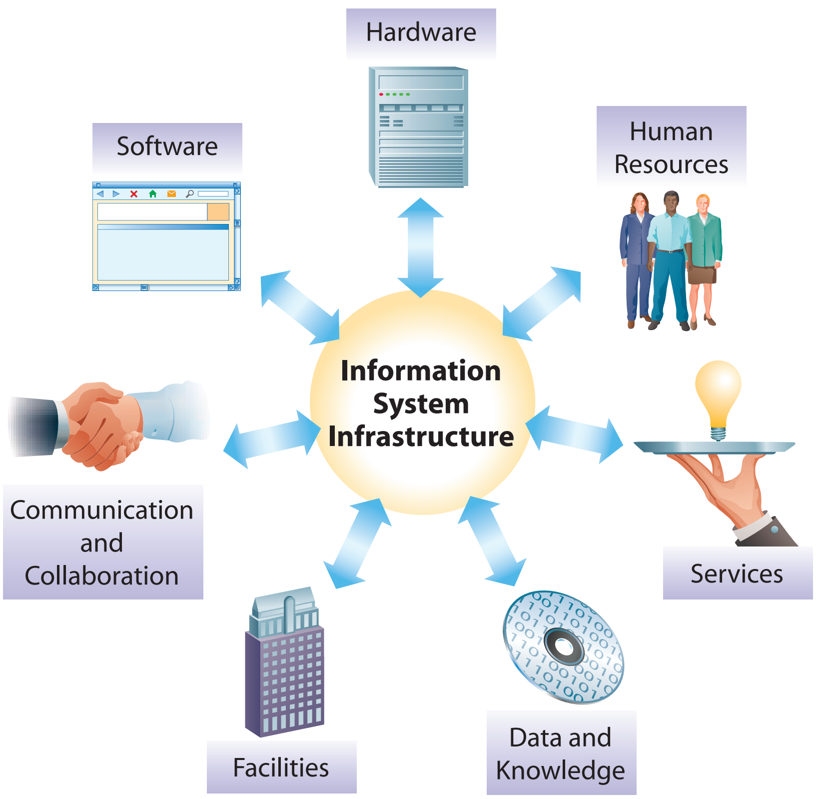


**FIGURE 4.5**

The information systems infrastructure.

- Human resources
- Services

Next, we briefly discuss each of these components and highlight their role in an organization's information systems infrastructure. To dig deeper into the technical aspects of the various infrastructure components, refer to the Technology Briefings.

**Hardware** The information systems hardware is an integral part of the IS infrastructure. This hardware consists not only of the computers used in an organization but also of networking hardware (see Technology Briefing 1—Information Systems Hardware and Technology Briefing 4—Networking) (see Figure 4.6). While the computing hardware is integral to an organization's IS infrastructure because it is needed to store and process organizational data, the networking hardware is needed to connect the different systems to allow for collaboration and information sharing.

Companies often face difficult decisions regarding their hardware. Constant innovations within the information technology sector lead to ever-increasing processor speeds and storage capacities but also to rapid obsolescence. Information systems executives therefore face countless complex questions, such as the following:

- Which hardware technologies should be chosen?
- What time interval should equipment be replaced?
- How can the information systems be secured best?
- What performance and storage is needed today? Next year?
- How can reliability be assured?

These and other questions will be addressed throughout this chapter when we discuss how different infrastructure solutions can help to support an organization's competitive strategy, decision making, and business processes.

**Software** As outlined in Technology Briefing 2—Information Systems Software, various types of software enable companies to utilize their information systems hardware and networks. This software assists organizations in executing their business processes and competitive strategy. Consequently, with increased reliance on information systems for managing the organization, effectively utilizing software resources is becoming increasingly critical and complex. For example, companies have to manage the software installed on each and every computer used, including managing updates, fixing bugs, and managing issues related to software licenses (see Figure 4.7). In addition, companies have to decide whether to upgrade their software or switch to new products and when to do so.

**FIGURE 4.6**

Hardware is an integral component of an organization's IS infrastructure.

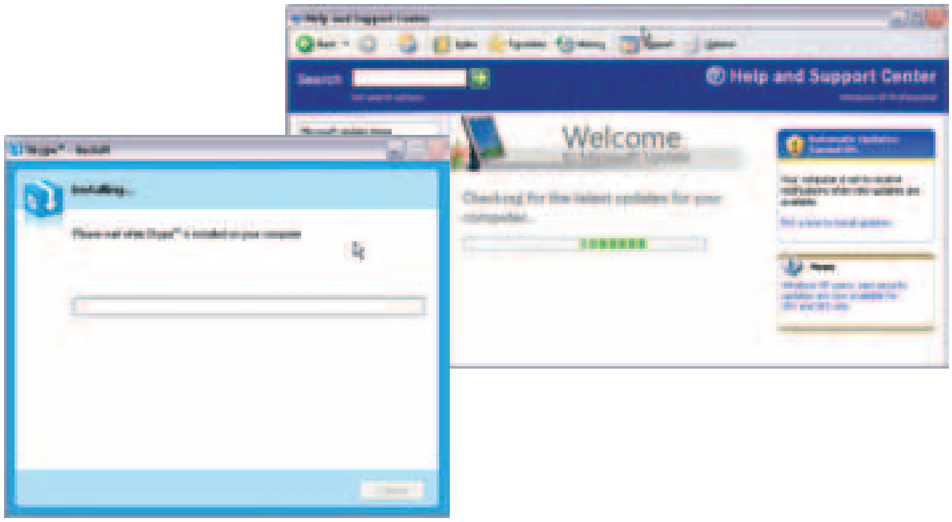


**FIGURE 4.7**

Installing and maintaining software can be a costly and time-consuming task.

Clearly, managing the software component of an IS infrastructure can be a daunting task. However, there are some developments helping organizations to better manage the software resources, and we will present these later in the chapter.

**Communications and Collaboration** As you have read in the previous chapters, one of the reasons why information systems in organizations have become so powerful and important is the ability to interconnect, allowing internal and external constituents to communicate and collaborate with each other. The infrastructure supporting this consists of a variety of components, such as the networking hardware and software (see Technology Briefings 4 and 5), that facilitate the interconnection of different computers, enabling collaboration literally around the world.

However, having a number of interconnected computers is necessary but not sufficient for enabling communication and collaboration; companies also need various other hardware and software. For example, e-mail servers, along with communication software such as Microsoft Outlook, are needed to enable a broad range of internal and external communication. Similarly, companies have to decide on whether to utilize tools such as instant messaging and which system to use for such applications (see Figure 4.8). Further, it has become increasingly important for companies to be able to utilize videoconferencing to bridge the distances between a company's offices or between a company and its business partners, saving valuable travel time and enhancing collaboration. However, as there are

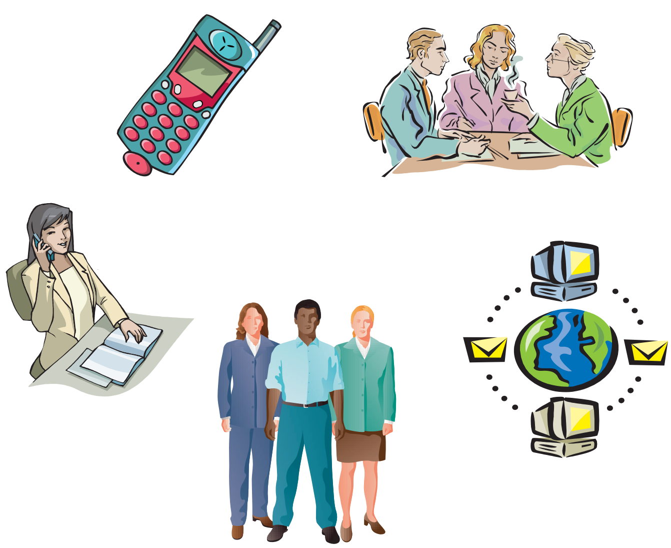

**FIGURE 4.8**

Companies have to decide how to support their communication needs.

vast differences in terms of quality, costs, and functionality of these systems, companies have to assess their communication needs and carefully decide which combination of technologies best support the goals of the organization.

**Data and Knowledge** Data and knowledge are probably among the most important assets an organization has, as data and knowledge are essential for both gaining business intelligence and executing business processes. Managing this resource thus requires an infrastructure with sufficient capacity, performance, and reliability. For example, companies such as Amazon.com need databases to store customer information, product information, inventory, transactions, and so on. Like Amazon.com, many companies operating in the digital world rely heavily on their databases not only to store information but also to analyze this information to gain business intelligence.

For example, the main data center for United Parcel Service (UPS) handles on average 10 million package tracking requests per day, with peak days approaching 20 million. To support this core business process, UPS has designed a data management architecture that includes an array of Unix-based mainframes running a massive database management system. This data management architecture has a capacity of 471 terabytes of data (approximately 471,000 gigabytes). Additionally, given that data is the lifeblood for UPS, they have replicated this infrastructure in two locations—New Jersey and Georgia—to ensure speed and reliability (see Figure 4.9).

In addition to effectively managing their data resources, organizations must also effectively manage their knowledge. In Chapter 1—Managing in the Digital World, we outlined the rise of the knowledge worker—professionals who are relatively well educated and who create, modify, and/or synthesize knowledge—and the new economy where organizations must effectively utilize their knowledge to gain a competitive advantage. Trends and options for effectively managing data and knowledge are also examined later in the chapter.

**Facilities** Although not directly needed to support business processes or business intelligence, specialized facilities are needed for the information systems infrastructure. While not every company needs facilities such as Google's data center in The Dalles, managers need to carefully consider where to house the different hardware, software, data centers, and so on. A normal desktop computer might not need much in terms of power, nor does it generate much heat; however, massive clusters of computers or **server farms** (facilities housing a vast number of servers to support the information processing needs of a large organization) have tremendous demands for reliable electricity and air-conditioning. In addition to such technical requirements, there is also the need to protect important equipment from both outside

**FIGURE 4.9**

UPS's servers handle up to 20 million requests per day.

## Ethical Dilemma

### Who Owns Company Data?

In years past, beads, land, gold, oil, animal skins, and food were all considered trade currency. Today, information is a prime bargaining chip. And, just like any of these forms of lucre, it can be stolen. For instance, a person might use information belonging to a previous employer to get a better job or to maintain old business contacts after changing jobs. The question is, Is stealing information from a company's database akin to stealing money from the bank?

The answer, of course, is "yes." The act of pilfering data from an employer or previous employer is data theft. From an organization's standpoint, stealing information is not only more harmful than stealing company stationery or a few pencils but even more harmful than stealing sophisticated computer hardware.

Unfortunately, data theft is not rare. A 2004 survey by Ibas, a data forensics firm in the United Kingdom, found that 70 percent of the respondents had stolen key information from an employer. Seventy-two percent of the people surveyed indicated that they had no ethical problems with taking proposals, presentations, contact databases, or e-mail address books when changing to a new employer. Fifty-eight percent felt that data theft ranked with exaggerating insurance claims. Thirty percent had stolen customer contact information when they left a firm. (Most thefts occurred when employees left to take another job.)

"The surprising thing is the level to which people believe this is acceptable," said Chris Watson of Ibas in an article by BBC News published online in February 2004.

Eighty percent of the employees surveyed justified their actions by using the rationale that "after all, I did the work to build that customer database and create sales leads."

Where do you stand on the issue?

Source: Anonymous, "Workplace Data Theft Runs Rampant," BBC News (February 15, 2004), http://news.bbc.co.uk/1/ hi/technology/3486397.stm

intruders and the elements, such as water or fire. The most prominent threats to an organization's IS facilities come from floods, seismic activity, rolling blackouts, hurricanes, and the potential of terrorist activities (see Figure 4.10). How can an organization reliably protect its facilities from such threats? Other issues to consider are the questions of availability; for example, can an organization afford to have its Web site unavailable for a minute, for an hour, or even for a day? Strategies for managing information systems facilities are examined later in the chapter.

**Human Resources** Another issue faced by companies is the availability of a trained workforce. Although even large facilities do not require large support staff, the staff that is needed should be well trained. This is one of the issues faced by Google's new data center in The Dalles. While the construction of the facility has created a large number of construction jobs, helping the area's unemployment situation, permanent jobs will likely require special skills so that much of the workforce will be "imported" from other regions. For this reason, many companies try to locate facilities in common areas. For example, the automobile industry has historically been centered in Detroit, while many of the technology companies have chosen areas like Austin, Boston, San Jose, and Seattle.

**Services** A broad range of services is the final infrastructure component. Over the past few years, this component has become increasingly important for many business organizations. Traditionally, an organization would perform all business processes—from acquiring raw materials to selling and servicing the final product—itself, no matter if these processes would be the organization's core competency or not. Today, with increasing pressures from the competition and the organization's shareholders, many processes that are not among a company's core competencies are being delegated to other organizations with considerably higher expertise (see Figure 4.11). For example, business

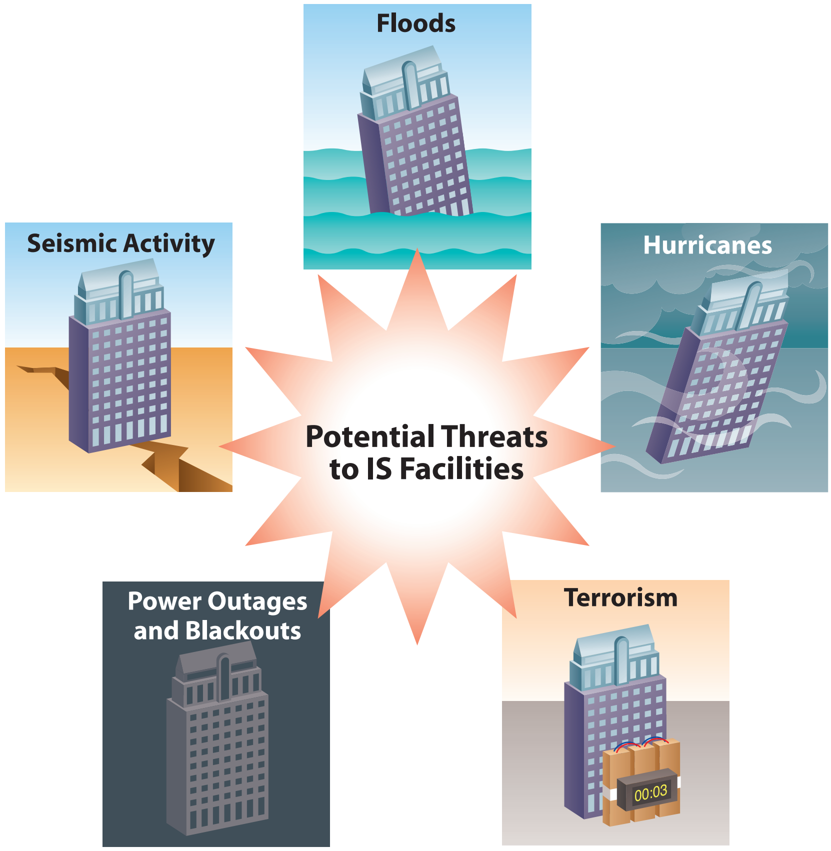


**FIGURE 4.10**

Potential threats to IS facilities vary by region include floods, hurricanes, terrorism, power outages, and seismic activity.

organizations turn over the management of their supply chains to logistics companies such as UPS, or organizations turn over the responsibility for the information systems infrastructure to organizations such as EDS. In fact, many of the solutions used to address organizations' various infrastructure needs are based on services, as you will see in the following sections.

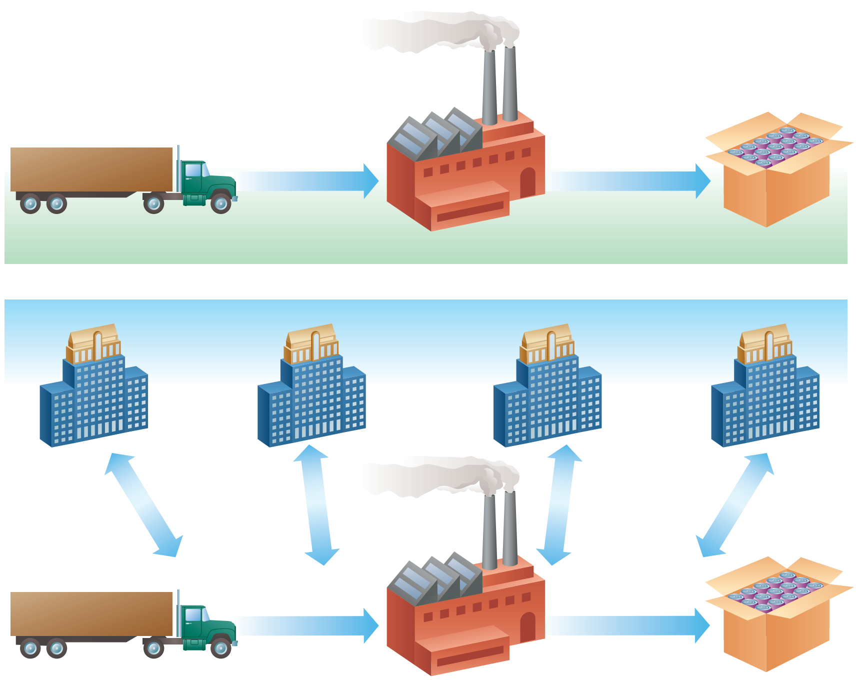

**FIGURE 4.11**

Traditionally, organizations would execute their entire value chain; today, a variety of other businesses provide services to perform noncore processes.

Clearly, there are a variety of infrastructure issues to consider for companies operating in today's digital world. In the next section, we will describe some solutions for designing and managing an organization's information systems infrastructure.

## Designing the Information Systems Infrastructure

With organizations' growing needs for a comprehensive information systems infrastructure, a number of solutions have emerged and are continuing to emerge. While some of these solutions are already common business practice, others are just now starting to be adopted. In the following sections, we will present various solutions to information systems infrastructure needs presented in the previous section.

### Managing the Hardware Infrastructure

Both businesses and research facilities face an ever-increasing need for computing performance. For example, auto manufacturers, such as the GM German subsidiary Opel or Japanese Toyota, use large supercomputers to simulate automobile crashes as well as evaluate design changes for vibrations and wind noise. Research facilities such as the U.S. Department of Energy's Lawrence Livermore National Laboratory use supercomputers for simulating nuclear explosions, while others simulate earthquakes using supercomputers (see Figure 4.12); such research sites have a tremendously complex hardware infrastructure.

While not every organization faces such large-scale computing problems, the demands for computing resources are often fluctuating, leading to either having too few resources for some problems or having too many idle resources most of the time. To address this problem, many organizations now turn to *on-demand computing* for fluctuating computation needs, *grid computing* for solving large-scale problems, and *autonomic computing* for increasing reliability. In the following sections, we will discuss each of these infrastructure trends.

**On-Demand Computing** In almost every organization, demand for individual IS resources is highly fluctuating. For example, some high-bandwidth applications, such as videoconferencing, may be needed only during certain times of the day, or some resource-intensive data-mining applications may only be used in irregular intervals. **On-demand computing** is a way to address such fluctuating computing needs; here, the available resources are allocated on the basis of users' needs (usually on a pay-per-use basis). For


**FIGURE 4.12**

The Earth Simulator supercomputer creates a "virtual twin" of the earth.

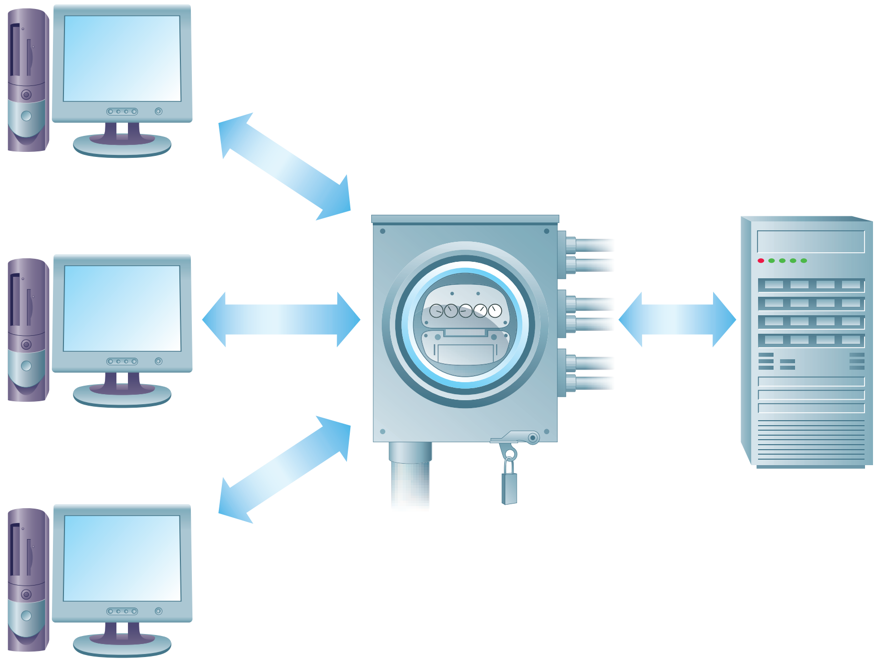

**FIGURE 4.13**

Utility computing allows companies to pay for computing services on an as-needed basis.

example, more bandwidth will be allocated to a videoconference, while other users who do not need the bandwidth at that time receive less. Similarly, a user running complex data-mining algorithms would receive more processing power than a user merely doing some word processing.

At times, organizations prefer to "rent" resources from an external provider. This form of on-demand computing is referred to as **utility computing**, where the resources in terms of processing, data storage, or networking are rented on an as-needed basis and the organization receives a bill for the services used from the provider at the end of each month (see Figure 4.13). For many companies, utility computing is an effective way for managing fluctuating demand as well as controlling costs; in essence, all tasks associated with managing, maintaining, and upgrading the infrastructure are left to the external provider and are typically bundled into the "utility" bill—if you don't use, you don't pay. Also, as with your utility bill, customers are charged not only on overall usage but also on peak usage (i.e., different rates for different times of the day).

**Grid Computing** Although today's supercomputers have tremendous computing power, some tasks are even beyond the capacity of a supercomputer. Indeed, some complex simulations can take a year or longer to calculate even on a supercomputer. Sometimes, an organization or a research facility would have the need for a supercomputer but may not be able to afford one because of the extremely high cost. For example, the fastest supercomputers can cost more than $200 million, and this does not represent the "total cost of ownership," which also includes all the other related costs for making the system operational (e.g., personnel, facilities, storage, software, and so on; see Chapter 3). Additionally, the organization may not be able to justify the cost because the supercomputer may be needed only occasionally to solve a few complex problems. In these situations, organizations have had to either rent time on a supercomputer or decided simply not to solve the problem.

However, a relatively recent infrastructure trend for overcoming cost or use limitations is to utilize **grid computing**. Grid computing refers to combining the computing power of a large number of smaller, independent, networked computers (often regular desktop PCs) into a cohesive system in order to solve problems that only supercomputers were previously capable of solving. While supercomputers are very specialized, grid computing allows organizations to solve both very large-scale problems as well as multiple (concurrent) smaller problems. To make grid computing work, large computing tasks are broken into small chunks, each of which can then be completed by the individual computers

**FIGURE 4.14**

Grid computing uses resources from various different computers located around the world.

(see Figure 4.14). However, as the individual computers are also in regular use, the individual calculations are performed during the computers' idle time so as to maximize the use of existing resources. For example, when writing this book, we used only minimal resources on our computers (i.e., we typically used only a word processor, the Internet, and e-mail); if our computers were part of a grid, the unused resources could be utilized to solve large-scale computing problems. This is especially useful for companies operating on a global scale. In each country, many of the resources are idle during the night hours, often more than 12 hours per day. Because of time zone differences, grid computing helps utilize those resources constructively. One way to put these resources into use would be to join the Berkeley Open Infrastructure for Network Computing (BOINC), which lets individuals "donate" computing time for various research projects, such as searching for extraterrestrial intelligence (SETI@home) or running climate change simulations.

However, as you can imagine, grid computing poses a number of demands in terms of the underlying network infrastructure or the software managing the distribution of the tasks. Further, many grids perform on the speed of the slowest computer, thus slowing down the entire grid. Many companies starting out with a grid computing infrastructure attempt to overcome these problems by using a **dedicated grid**. In a dedicated grid, the individual computers, or nodes, are just there to perform the grid's computing tasks; in other words, the grid consists of a number of homogeneous computers and does not use unutilized resources. A dedicated grid is easier to set up and manage and is for many companies much more cost effective than purchasing a supercomputer. As the grid evolves and new nodes are added, dedicated grids become more heterogeneous over time.

One factor that adds to the popularity of using dedicated grids is the falling cost of computing hardware. Just a few years ago, companies have attempted to utilize idle resources as much as possible and set up heterogeneous computing grids. However, the added complexity of managing heterogeneous grids poses a large cost factor so that today it is often more cost effective to set up a homogeneous, dedicated grid; in this case, the savings in terms of software and management by far offset the added costs for dedicated computing hardware in terms of both acquisition and maintenance.

**Edge Computing** Another recent trend in IS hardware infrastructure management is **edge computing**. With the decrease in cost for processing and data storage, computing tasks are now often solved at the edge of a company's network. In other words, rather than having massive, centralized computers and databases, multiple smaller servers are located closer to the individual users. This way, resources in terms of network bandwidth and access time are saved. If a computer needs several hours to compute a certain problem, it might be a good choice to send the task over a network to a more powerful computer that might be able to solve that problem faster. However, as the costs for computing power have decreased tremendously over the past years, many problems can now be computed locally within a matter of seconds, so it is not economic to send such problems over a network to a remote computer (Gray, 2004). To save resources, many businesses use edge computing for their online commerce sites. In such cases, customers interact with the servers of an edge-computing service provider (such as Akamai). These servers, in turn, communicate with the business' computers. This form of edge computing helps to reduce wait times for the consumers, as the e-commerce sites are replicated on Akamai's servers, while at the same time reducing the number of requests to the company's own infrastructure. This process not only saves valuable resources such as bandwidth but also offers superior performance that would otherwise be too expensive for organizations to offer. Akamai's services are utilized by organizations such as NBC, Fox Sports, BMW, and Victoria's Secret.

**Autonomic Computing** One major drawback of these hardware infrastructure trends and the demands for IS infrastructure in general is the increased complexity of such systems. Whereas the primary reason for having this infrastructure is the utilization of the resources, the time and money needed to manage these resources don't add value to the organization; in fact, some people believe that the costs of managing these systems undermine the benefits these systems provide, even if the organization decides to use outside services. To overcome this, academic and industry researchers (e.g. at IBM) have begun working on **autonomic computing** systems, which are self-managing, meaning they need only minimal human intervention to operate (see Figure 4.15). In other words, in a traditional computing environment, system operators often have to fine-tune the computer's configuration in order to most efficiently solve a particular type of complex problem. In an autonomic computing environment, the ultimate goal is to allow the system to do everything else on its own, completely transparent to the user. In order to achieve this, an autonomic computing system must know itself and be self-configuring, self-optimizing, self-healing, and self-protecting.

In order to optimally perform different tasks, an autonomic system must know itself; that is, it must know its configuration, capacity, and current status, but it must also know which resources it can draw on. Second, in order to be able to use different resources based on different needs, the system should be self-configuring so that the user does not have to take care of any configuration issues. Further, as any parts of a system can malfunction, an autonomic system should be self-healing so that any potential problems are detected and the system is reconfigured so as to allow the user to continue performing the tasks, even if parts of the system are not operational. Finally, as almost any computer system can be the target for an attack (see Chapter 6—Securing Information Systems), autonomic computing systems must be aware of any potential dangers and must be able to protect themselves

**FIGURE 4.15**

Autonomic computing systems have self-awareness and are self-configuring, self-optimizing, self-healing, and self-protecting.

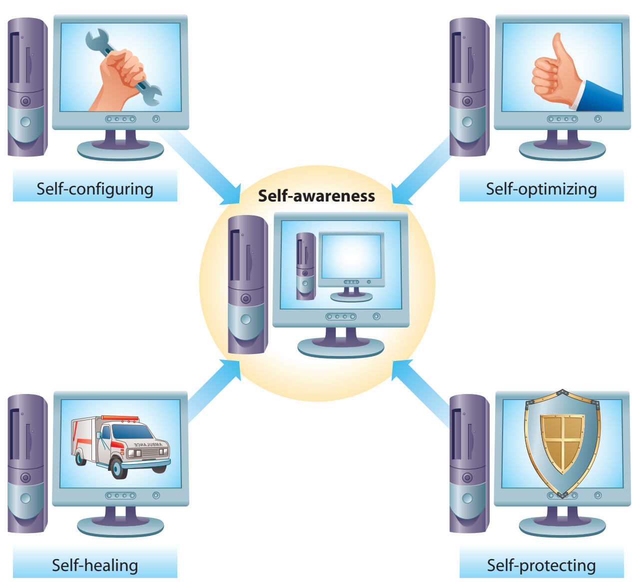

from any malicious attacks (e.g., by automatically quarantining infected parts of a system). Clearly, these are some formidable tasks researchers have to address, but considering the time and money that is currently spent on managing and maintaining IT infrastructures, autonomic computing systems are promising for the future.

## Managing the Software Infrastructure

With growing use of information systems to support organizations' business processes and the need for business intelligence, organizations have to rely on a variety of different software. However, continuously upgrading operating systems and applications software (see Technology Briefing 2) can be a huge cost factor for organizations both in terms of labor and in terms of costs for the actual products needed. To reduce such costs, many companies are now turning increasingly to using open-source software, attempting to integrate various software tools, or using application service providers for their software needs. Each of these software infrastructure management approaches is discussed next.

**Open-Source Software** Open source, seen by Friedman (2005) as one of the 10 flatteners of the world, is a philosophy that promotes developers' and users' access to the source of a product or idea (see Chapter 2—Fueling Globalization Using Information Systems). Particularly in the area of software development, the open-source movement has taken off with the advent of the Internet, and people around the world are contributing their time and expertise to develop or improve software, ranging from operating systems to applications software. As the programs' source code is freely available for use and/or modification, this software is referred to as **open-source software**.

**OPEN-SOURCE OPERATING SYSTEMS.** One of the most prevalent examples of open-source software is the operating system Linux, which was developed as a hobby by the Finnish university student Linus Torvalds in 1991. Having developed the first version himself, he made the source code of his operating system available to everyone who wanted to use it and improve on it. Because of its unrivaled stability, Linux has become the operating system of choice for Web servers, **embedded systems** (such as TiVo boxes, handheld computers, and network routers; see Figure 4.16), and supercomputers alike (as of 2006, 73 percent of the world's 500 fastest supercomputers ran Linux operating systems; Top 500, 2006).

**OPEN-SOURCE APPLICATION SOFTWARE.** In addition to the Linux operating system, other open-source software has been gaining increasing popularity because of its stability and low cost. For example, in 2006, 68 percent of all Web sites were powered by the

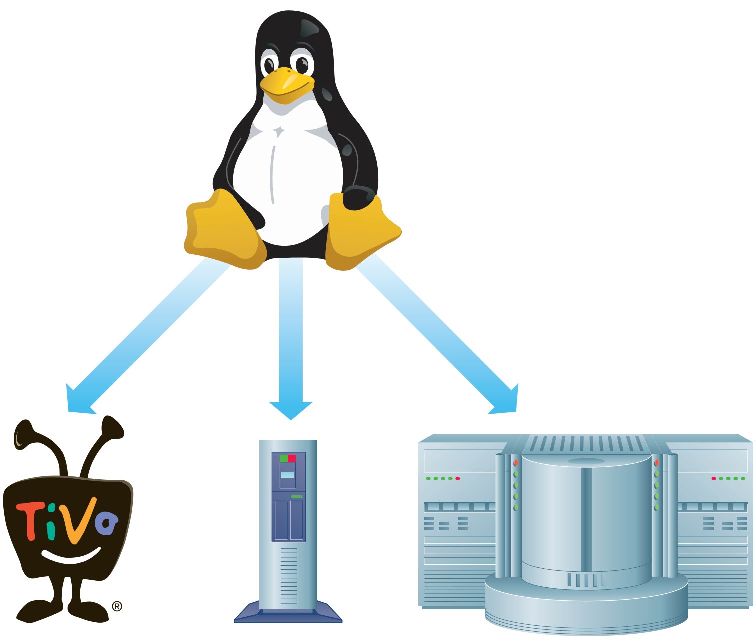


**FIGURE 4.16**

Linux is the operating system of choice for embedded systems, Web servers, and supercomputers (the penguin "Tux" is the official mascot of Linux).


**FIGURE 4.17**

The Firefox Web browser.

Apache Web server, another open-source project (Netcraft, 2006). Other popular examples of open-source application software include the Firefox Web browser (see Figure 4.17) and the office productivity suite OpenOffice. While there are many upsides to open-source software, vendors of proprietary software are still highlighting "hidden" costs of running open-source software. For example, finding an organization to provide reliable customer support can sometimes be difficult.

**Web Services** To perform business processes and for business intelligence, it is often essential to draw information from different sources or different applications. However, with the increasing complexity of an organization's software needs, it is often impossible to get all of the various applications to integrate seamlessly. In some cases, software companies (such as Microsoft and IBM) offer a wide range of products, all of which can interoperate quite well. However, business organizations sometimes shy away from being completely dependent on a single vendor for their software needs. One way to increase the independence while still being able to integrate various software applications is the use of **Web services**. Web services are Web-based software systems used to allow the interaction of different programs and databases over a network. Using Web services, companies can integrate information from different applications, running on different platforms. For example, you can use Web services offered by Google to integrate search functionality into your own Web site, or you can use Web services offered by MapQuest to provide your guests with an interactive map to your house (see Figure 4.18). You can also learn more about Web services in Technology Briefing 2.

One logical extension to using Web services is using a **service-oriented architecture**. The main goal of implementing a service-oriented architecture is the integration of different applications using Web services. In a service-oriented architecture, different repeatable business tasks, or services, are integrated to better perform various business processes. These services are typically vendor independent and can thus be used to integrate data and capabilities of different systems running on different platforms. This capability—and the reusability of different services—allows businesses to quickly react to changes in the business environment.

**FIGURE 4.18**

An example of a Web service allowing people to insert maps into their Web pages.

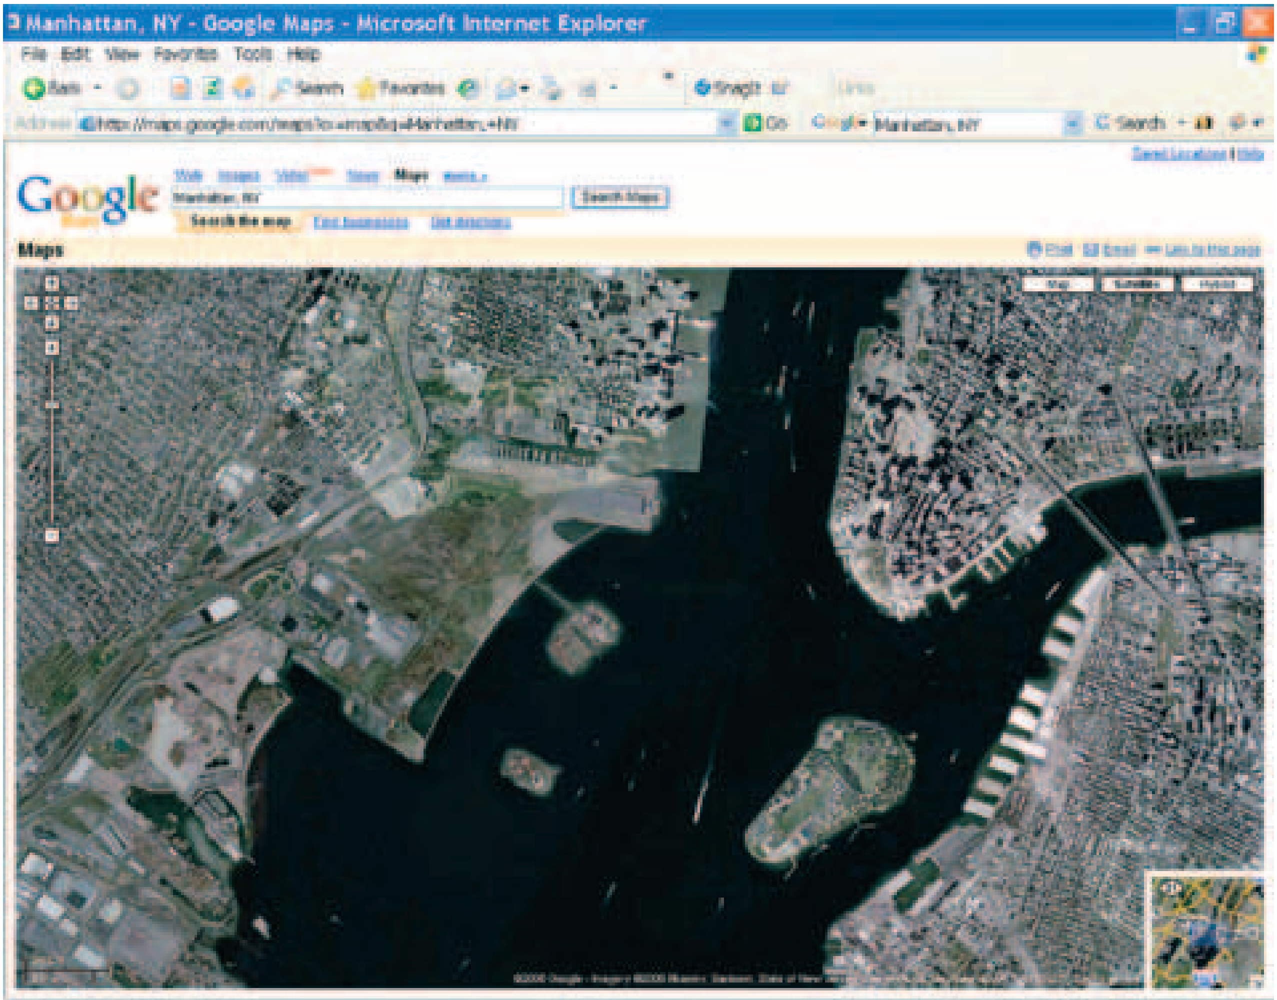


**Managing Software Assets** As organizations manage their software infrastructure, there are several issues that must be carefully managed, such as software bugs and licensing. Next, we briefly outline these issues and present some of the tools and strategies organizations are utilizing to better manage these complex tasks.

**MANAGING SOFTWARE BUGS.** With the increased complexity of software, it is almost impossible to build applications that are error free, and no matter whether such applications are operating systems, Web sites, or enterprise-wide software, there is the potential of unforeseen problems with the software. Typically, software developers will account for these unforeseen problems by incorporating a **patch management system** into the application. Patch management is typically based on an online system that checks a Web service for available patches. If the software vendor offers a new patch, the application will download and install the patch in order to fix the software bug. An example of a patch management system in wide use is the Windows Update Service. The user's operating system automatically connects to a Microsoft Web service to download critical operating system patches. While some of these patches are intended to fix bugs in the Windows operating system, many patches are built to fix security holes that could be exploited by malicious hackers.

**TABLE 4.1 Different Types of Software Licenses**

| Restrictiveness | Software Type | Rights | Restrictions | Examples |
|---|---|---|---|---|
| Full rights | Public domain software | Full rights | No restrictions; owner forsakes copyright | Different programs for outdated IBM mainframes |
| | Nonprotective open source (e.g., Berkeley software development [BSD] license) | Freedom to copy, modify, and redistribute the software; can be incorporated into commercial product | Creator retains copyright | FreeBSD operating system; BSD components in (proprietary) Mac OS X operating system |
| | Protective open source (e.g., general public license [GPL]) | Freedom to copy, modify, and redistribute the software | Modified or redistributed software must be made available under the same license; cannot be incorporated into commercial product | Linux operating system |
| | Proprietary software | Right to run the software (for licensed users) | Access to source code severely restricted; no rights to copy or modify software | Windows operating system |
| No rights | Trade secret | | Access to source code severely restricted; software is not distributed outside the organization | Google PageRank™ algorithm |

**MANAGING SOFTWARE LICENSING.** Software licensing has been a hot button topic for software companies as they lose billions in piracy and mislicensed customers (see Chapter 10—Managing Information Systems Ethics and Crime). Traditionally, software licensing is defined as the permission and rights that are imposed on applications, and the use of software without a proper license is illegal in most countries.

Most software licenses differ in terms of restrictiveness, ranging from no restrictions at all to completely restricted. Table 4.1 lists different types of software licenses, ordered in terms of restrictiveness. Note that although freeware or shareware is freely available, the copyright owners often retain their rights and do not provide access to the program's source code. For organizations using proprietary software, two types of licenses are of special importance. The first, **shrink-wrap license**, accompanies the software and is used primarily in consumer products. The shrink-wrapped contract has been named as such because the contract is activated when the shrink wrap on the packaging has been removed. The second type of licensing is **enterprise licensing**. Enterprise licensing (also known as **volume licensing**) can vary greatly and is usually negotiated. In addition to rights and permissions, enterprise licenses usually contain limitations of liability and warranty disclaimers that protect the software vendor from being sued if their software does not operate as expected.

As shown in Table 4.1, there are a variety of software licenses. For different business needs, organizations are often depending on a variety of software, each having different licenses, which can cause headaches for many organizations. Not knowing about the software an organization has can have a variety of consequences. For example, companies are not able to negotiate volume licensing options, unused licenses strain the organization's budget, or license violations can lead to fines or public embarrassment. **Software asset management** helps organizations to avoid such negative consequences. Usually, software asset management consists of a set of activities, such as performing a software inventory (either manually or using automated tools), matching the installed software with the licenses, reviewing software-related policies and procedures, and creating a software asset management plan. The results of these processes help organizations to better manage their software infrastructure by being able to consolidate and standardize their software titles, decide to retire unused software, or decide when to upgrade or replace software.

**Application Service Providers** Undoubtedly, managing the software infrastructure is a complex task, often resulting in large fluctuations in operating costs for organizations. To better control such costs, business organizations increasingly use the services of **application service providers (ASP)**. Analogous to on-demand computing, application service providers offer **on-demand software** for a variety of clients who access the applications on an as-needed basis over the Web. In other words, the software is located on the ASP's servers, and the users interact with the software using Web-enabled interfaces, such as Web browsers; while the software to perform the tasks is provided by the ASP, the organization still performs the task to be completed (such as payroll processing). For organizations, using an ASP offers a variety of benefits, such as a reduced need to maintain or upgrade software, a fixed monthly fee for services (rather than variable IT costs), and the ability to rely on a provider that has gained considerable expertise because of a large number of clients.

One example for a simple, free application service is Google calendar, which allows users to organize their schedules, share calendars, and coordinate meetings with other users. To address different business' needs, there are a variety of application service providers (see Table 4.2 for examples of different ASPs).

## Managing the Communication and Collaboration Infrastructure

The organization's communication and collaboration needs are the third major infrastructure component. As with the hardware and software infrastructure, some changes in the organizations' needs have taken place over the past years; for example, e-mail has become the communications medium of choice for many people. However, for some topics, other forms of communication are more suited, so managers turn to the telephone, instant messaging, meetings, or videoconferences. One recent trend to satisfy such diverse communication and collaboration needs is the growing convergence of computing and telecommunications.

Net Stats

### Broadband Access Increases

Recent reports show that in 2006, nearly 70 percent of active home Internet users in the United States had access to broadband connections. Web connection speed trends are shown in Figure 4.19.

**FIGURE 4.19**

Web connection speed trends in U.S. homes.

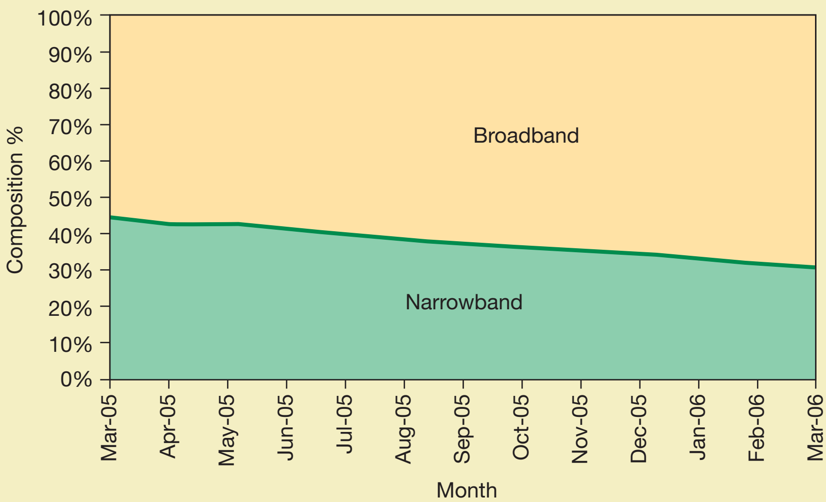


Source: Nielsen/Net Ratings. Graph found at http://www.websiteoptimization.com/bw/0604/.

**TABLE 4.2 Examples of Different Types of Application Service Providers (ASPs)**

| Type | Service Offered | Example |
|---|---|---|
| Specialist or functional ASP | Single application | ASP providing payroll processing software for companies to use |
| Vertical market ASP | Solution package for a specific industry | ASP providing property management systems for hotels |
| Enterprise ASP | Broad solutions for the needs of different organizations | ASP offering complete enterprise resource planning solutions (see Chapter 7) to different industries |
| Local ASP | Services for small businesses within a limited geographic area | ASP offering Web site design and maintenance for small-business owners in a community |

**Convergence of Computing and Telecommunications** The computing industry is experiencing an ever-increasing convergence of functionality of various devices. Whereas just a few years ago a cell phone was just a cell phone and a PDA was just a PDA (personal digital assistant; see Technology Briefing 1), such devices are now converging such that the boundaries between devices are becoming increasingly blurred. Today, an increasing number of devices offer a variety of different functionalities—formerly often available only on separate dedicated devices—to address differing needs of knowledge workers and consumers alike (e.g., phones, PDAs, cameras, music players, and so on).

In addition to a convergence of capabilities of devices, there is also increasing convergence within the underlying infrastructures. For example, in the past, the backbone networks for the telephone and Internet were distinct. Today, increasingly, most voice and data traffic shares a common network infrastructure. To facilitate this convergence, also termed **IP convergence**, the use of IP (Internet Protocol; see Technology Briefing 5) for transporting voice, video, fax, and data traffic has allowed enterprises to make use of new forms of communication and collaboration (e.g., instant messaging and online whiteboard collaboration) as well as traditional forms of communication (such as phone and fax) at much lower costs (see Figure 4.20). In the following sections, we will discuss two uses of IP for communication: voice over IP and videoconferencing over IP.

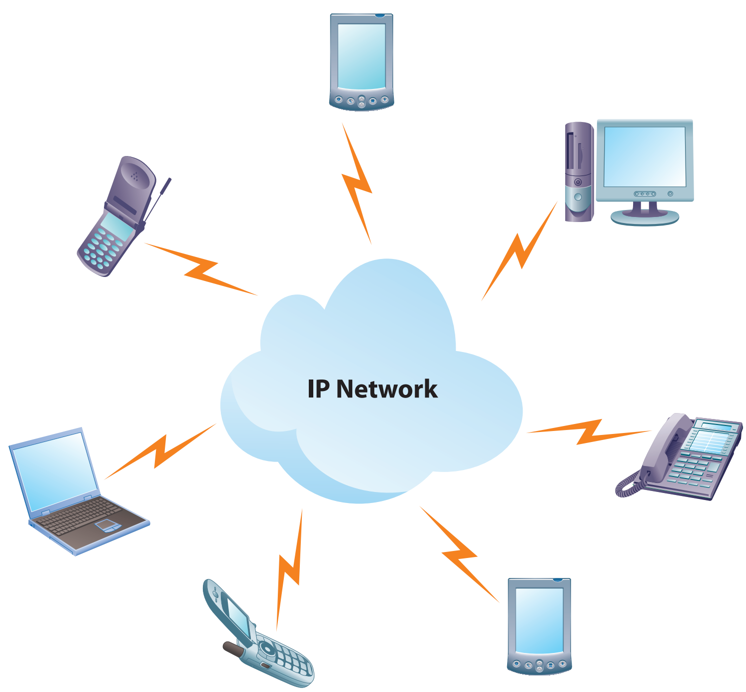


**FIGURE 4.20**

IP convergence allows various devices to communicate using IP technologies.

**VOICE OVER IP.** **Voice over IP (VoIP)** (or IP telephony) refers to the use of Internet technologies for placing telephone calls. Whereas just a few years ago the quality of VoIP calls was substandard, recent technological advances now allow the quality of calls to equal or even surpass the quality of traditional calls over (wired) telephone lines. In addition to the quality, VoIP offers a number of other benefits; for example, users can receive calls from almost anywhere they connect to the Internet. In other words, knowledge workers are not bound to their desk to receive VoIP calls; instead, using IP routing, their telephone number "follows" them to wherever they connect to the Internet. Organizations can also benefit from tremendous cost savings, as often there is almost no cost incurred over and above the costs for a broadband Internet connection (VoIP software such as Skype allows home users to make free PC-to-PC calls; see Figure 4.21).

**VIDEOCONFERENCING OVER IP.** In addition to voice communications, IP can also be used to transmit video data. Traditionally, videoconferences were held via traditional phone lines, which were not made to handle the transfer of data needed for high-quality videoconferencing. Some companies also used dedicated digital lines for videoconferencing; however,

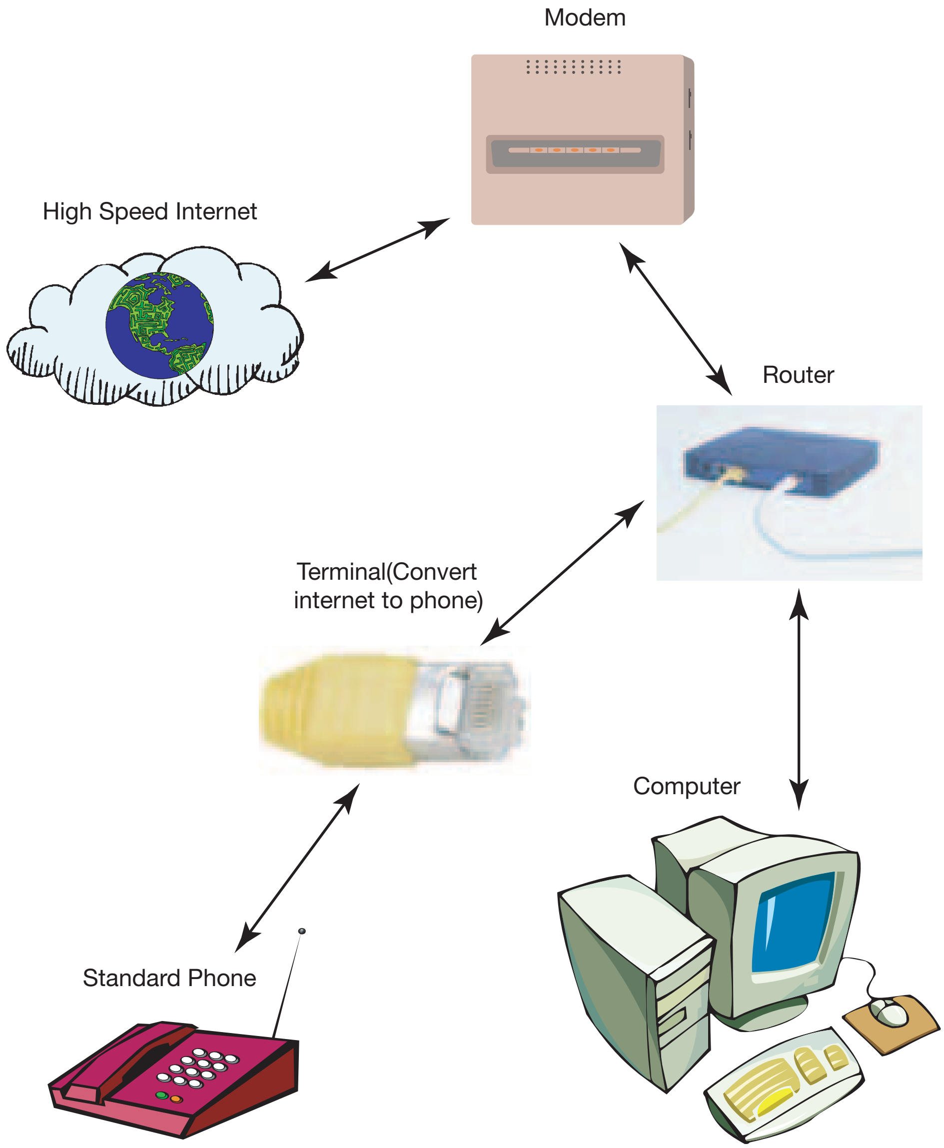


**FIGURE 4.21**

VoIP technology enables organizations and individuals to reduce their telecommunications costs.

**FIGURE 4.22**

The HP HALO meeting room features life-size images.

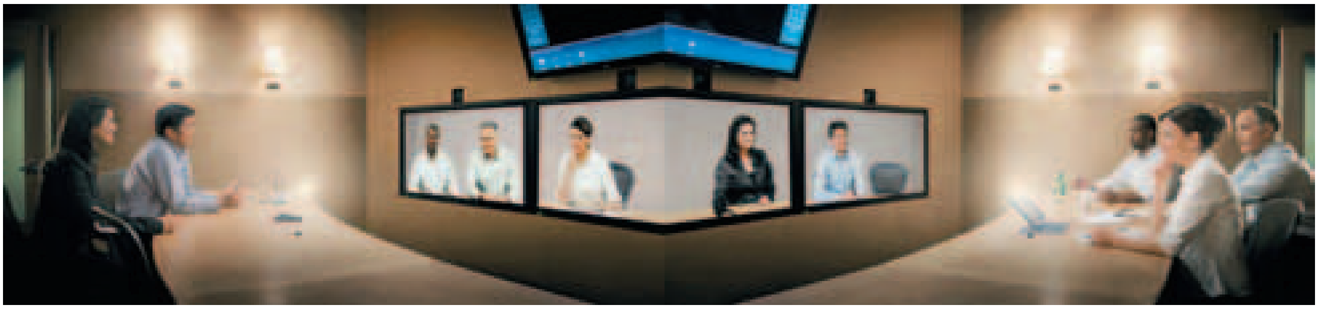

this was a very costly option. Similar to VoIP, the Internet also helped to significantly reduce costs and enhance the versatility of videoconferences by enabling **videoconferencing over IP**.

While for some videoconferences desktop videoconferencing equipment (consisting of a webcam, a microphone, speakers, and software such as Microsoft Office Live Meeting or Skype) may be sufficient, for others higher-end equipment may be needed. Such infrastructure can include specific videoconferencing hardware, or it can even be a $400,000 HP HALO meeting room featuring life-size images allowing people from across the globe to meet as if they were sitting in the same room (see Figure 4.22). In contrast to other applications, with the HALO room, HP provides a videoconferencing service to its customers, offering features such as access to a dedicated network infrastructure or support services for a fixed monthly fee. We will discuss videoconferencing in more detail in Chapter 7.

**Increasing Mobility** Changes in communication media—such as the growth of e-mail or instant messaging—has led to changes in the way we communicate. In today's digital world, knowledge workers desire being connected, whenever and wherever they are, so that they are able to quickly respond to any communication or use any spare minute to clean up their e-mail in-box. One infrastructure component supporting this need for connectivity is the provision of a wireless infrastructure.

**WIRELESS INFRASTRUCTURES.** Today's knowledge workers use two primary categories of wireless devices for their communication needs: (1) communication devices (such as cell phones) that use the public telephone infrastructure and (2) wireless devices capable of connecting to an organization's internal network. The convergence of devices and infrastructures allows the sending and receiving of e-mails using a cell phone. Thus, knowledge workers no longer need a laptop computer, nor do they need to be connected to an organization's network to get useful tasks completed. Similarly, using Web-enabled cell phones or PDAs, knowledge workers can access their company's networks and informational resources (see Chapter 5 for a discussion of corporate intranets). However, for many applications, having access to an organization's network can offer many advantages in terms of speed, ease of use, and so on. Thus, organizations are increasingly using wireless infrastructures.

When deciding on installing wireless infrastructures, organizations have to take great care to address issues such as the standards being used as well as the security of the network. For example, wireless networks based on the 802.11 family of standards are the norm today (see Technology Briefing 4). However, with the widespread use, misuse also abounds; as most of today's new laptops are equipped with the necessary hardware to connect to 802.11 networks, people are often trying to find a way to surf the Internet for free while on the road or even try to (illegally) access organization's networks (also called drive-by hacking or "war driving"; see Figure 4.23 and Chapter 6—Securing

**FIGURE 4.23**

Accessories for "war driving" are easily obtainable, causing security concerns for organizations.

Source: http://shop.netstumbler.com/SearchResult.aspx?CategoryID=26.

## Key Enabler

### Cognitive Radio

Wireless transmission is now the name of the game. Increasingly, cell phones, Wi-Fi, and satellite communication are saturating the airwaves with our personal data. Privacy is one concern, but also of concern is overcrowding of the airwaves. (Try making a cell phone call in Arizona during the peak "snowbird" season, and you will understand the problem.) What happens when calls are dropped? Usually one of three conditions: (1) wireless frequencies are full, (2) wireless towers are too few and far between, or (3) the environment (weather, high mountains, and so on) interferes.

Fortunately, researchers at Virginia Polytechnic Institute and State University in Blacksburg, Virginia, have come up with a plan for increasing the carrying capacity of airwaves. The technology—called "cognitive radio"—was originally designed for use at disaster sites since it gives a wireless signal a certain degree of intelligence. That is, it allows the transmitter to detect whether certain segments of the radio spectrum are in use. The signal can then switch to unused portions of the spectrum.

The radio spectrum in the United States is fixed by the Federal Communications Commission (FCC), with much of the spectrum being allotted for signals that are rarely used. These open areas are used intelligently by the cognitive radio signal to allow maximum carrying capacity.

Cognitive radio's current capabilities include location sensing, detection of other transmission devices, changing frequency, and even adjusting power output of the signal. These capabilities and others provide cognitive radio with the ability to adapt to conditions in real time, thereby maximizing the radio spectrum.

Intel has been a leader in the commercialization of this technology, and the company is now building reconfigurable chips that will analyze the environment and select the optimal protocol and frequency for sending data. The FCC has made special allowances so that these new devices can be tested on a large scale.

Sources: Fette, B. 2004. Cognitive Radio Shows Great Promise. *COTS Journal*. http://www.cotsjournalonline.com/ home/article.php?id=100206

Savage, J. 2006. Cognitive Radio. *Technology Review*. http://www.technologyreview.com/read_article.aspx?ch=specialsections&sc=emergingtech&id=16471

Niknejad, K. 2005. Cognitive radio: a smarter way to use radio frequencies. *Columbia News Service*. http://jscms.jrn.columbia.edu/cns/2005-04-19/niknejad-smartradio

Information Systems). In many cases, having an (unsecured) wireless network is equal to having a live network cable lying in the company's parking lot. Clearly, securing wireless infrastructures still poses challenges for organizations, which have to strike a balance between providing relatively easy access for authorized individuals and limiting access for outsiders.

## Managing the Data and Knowledge Infrastructure

To support more effective business processes and gather business intelligence, organizations have to find ways to effectively manage data from different sources as well as manage their internal knowledge. Thus, companies have turned to data-mining and knowledge management tools, which we will discuss in the following sections.

### Data Mining

**Data mining** is a method companies use to sort and analyze information to better understand their customers, products, markets, or any other phase of their business for which data has been captured. With data-mining tools, you can graphically drill down from summary data to more detailed data, sort or extract data based on certain conditions, and perform a variety of statistical analyses, such as trend analysis, correlation analysis, forecasting, and analysis of variance. The next sections describe how data is being collected when organizations interact with their clients, and how these data are analyzed using data mining techniques.

**ONLINE TRANSACTION PROCESSING.** Fast customer response is fundamental to having a successful Internet-based business. **Online transaction processing (OLTP)** refers to immediate automated responses to the requests of users. OLTP systems are designed specifically to handle multiple concurrent transactions from customers. Typically, these transactions have a fixed number of inputs, such as order items, payment data, and customer name and address, and a specified output, such as total order price or order tracking number. In other words, the primary use of OLTP systems is gathering new information, transforming that information, and updating information in the system. Common transactions include receiving user information, processing orders, and generating sales receipts. Consequently, OLTP is a big part of interactive electronic commerce applications on the Internet. Since customers can be located virtually anywhere in the world, it is critical that transactions be processed efficiently (see Figure 4.24). The speed with which database management systems can process transactions is, therefore, an important design decision when building Internet systems. In addition to which technology is chosen to process the transactions, how the data is organized is also a major factor in determining system performance. Although the database operations behind most transactions are relatively simple, designers often spend considerable time making adjustments to the database design in order to "tune" processing for optimal system performance. Once an organization has all this data, it must design ways to gain the greatest value from its collection; although each individual OLTP system could be queried individually, the real power for an organization comes from analyzing the aggregation of data from different systems. Online analytical processing is one method being used to analyze these vast amounts of data.

**ONLINE ANALYTICAL PROCESSING.** **Online analytical processing (OLAP)** refers to the process of quickly conducting complex analysis of data stored in a database, typically using graphical software tools. The chief component of an OLAP system is the **OLAP server**, which understands how data is organized in the database and has special functions for analyzing the data. OLAP tools enable users to analyze different dimensions of data beyond data summary and data aggregations of normal database queries (see Technology Briefing 3—Database Management). For example, OLAP can provide time-series and trend analysis views of data, data drill-downs to deeper levels of consolidation, and the ability to answer "what if" and "why" questions. An OLAP query for Amazon.com might be, "What would be the effect on profits if wholesale book prices increased by 10 percent and transportation costs decreased by 5 percent?" Managers use the complex query capabilities of an OLAP system to answer questions within executive information systems, decision support systems, and enterprise resource planning (ERP) systems (each of these systems is described in later

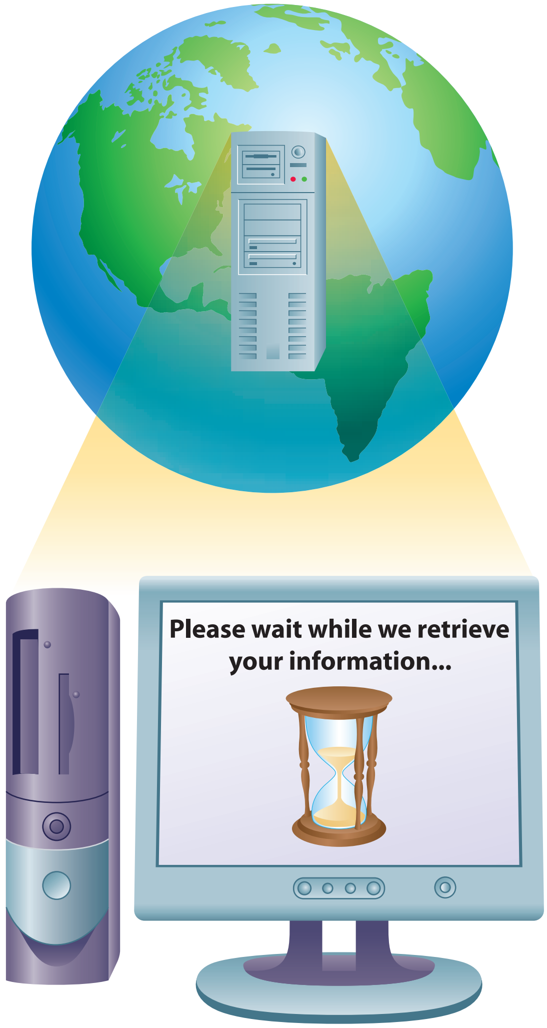


**FIGURE 4.24**

Global customers require that online transactions be processed efficiently.

chapters). Given the high volume of transactions within Internet-based systems, analysts must provide extensive OLAP capabilities to managers in order to gain the greatest business value.

**MERGING TRANSACTION AND ANALYTICAL PROCESSING.** The requirements for designing and supporting transactional and analytical systems are quite different. In a distributed online environment, performing real-time analytical processing diminishes the performance of transaction processing. For example, complex analytical queries from an OLAP system require the locking of data resources for extended periods of execution time, whereas transactional events—data insertions and simple queries from customers—are fast and can often occur simultaneously. Thus, a well-tuned and responsive transaction system may have uneven performance for customers while analytical processing occurs. As a result, many organizations replicate all transactions on a second database server so that analytical processing does not slow customer transaction processing performance. This replication typically occurs in batches during off-peak hours, when site traffic volumes are at a minimum.

The systems that are used to interact with customers and run a business in real time are called **operational systems**. Examples of operational systems are sales order processing and reservation systems. The systems designed to support decision making based on stable point-in-time or historical data are called **informational systems**. The key differences

**TABLE 4.3 Comparison of Operational and Informational Systems**

| Characteristic | Operational System | Informational System |
|---|---|---|
| Primary purpose | Run the business on a current basis | Support managerial decision making |
| Type of data | Current representation of state of the business | Historical or point-in-time (snapshot) |
| Primary users | Online customers, clerks, salespersons, administrators | Managers, business analysts, customers (checking status, history) |
| Scope of usage | Narrow and simple updates and queries | Broad and complex queries and analyses |
| Design goal | Performance | Ease of access and use |

between operational and informational systems are shown in Table 4.3. Increasingly, data from informational systems are being consolidated with other organizational data into a comprehensive data warehouse, where OLAP tools can be used to extract the greatest and broadest understanding from the data (see Figure 4.25).

**Data Warehousing** Large organizations such as Wal-Mart, UPS, and Alaska Airlines have built **data warehouses**, which integrate multiple large databases and other information sources into a single repository. This repository is suitable for direct querying, analysis, or processing. Much like a physical warehouse for products and components, a data warehouse stores and distributes data on computer-based information systems. A data warehouse is a company's virtual storehouse of valuable data from the organization's disparate information systems and external sources. It supports the online analysis of sales, inventory, and other vital business data that have been culled from operational systems. The purpose of a data warehouse is to put key

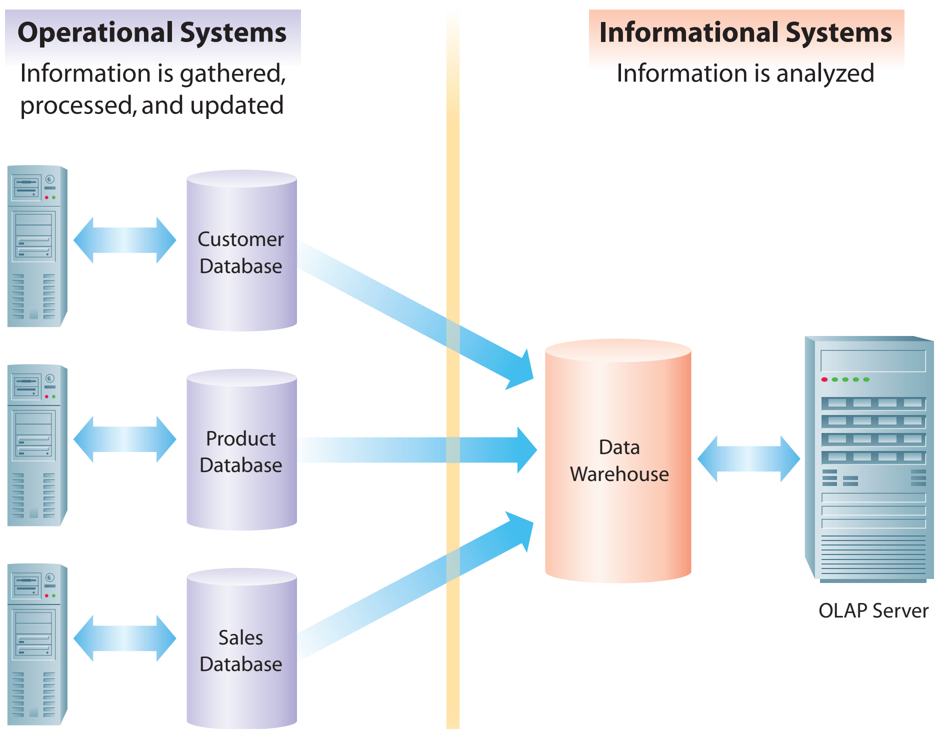


**FIGURE 4.25**

Enhancing business intelligence by combining data from different sources.

**TABLE 4.4 Sample Industry Uses of Data Warehousing**

| Use of Data Warehousing | Representative Companies |
|---|---|
| **Retail** | |
| Analysis of scanner checkout data | Amazon.com |
| Tracking, analysis, and turning of sales promotions and coupons | Costco |
| | CVS Corporation |
| Inventory analysis and redeployment | Home Depot |
| Price reduction modeling to "move" the product | Office Depot |
| Negotiating leverage with supplies | Sears |
| Frequent buyer program management | Target |
| Profitability analysis | Walgreen |
| Product selections of granular market segmentation | Wal-Mart |
| | Williams-Sonoma |
| **Telecommunications** | |
| Analysis of call volume, equipment sales, customer profitability, costs, inventory | AT&T |
| | Cingular Wireless |
| Inventory analysis and redeployment | Comcast Cable |
| Purchasing leverage with suppliers | Hong Kong CSL |
| Frequent buyer program management | Telefonica SA |
| Resource and network utilization | T-Mobile |
| Problem tracking and customer service | Verizon |
| **Banking and Financing** | |
| Relationship banking | Bank of America |
| Cross-segment marketing | Citigroup |
| Risk and credit analysis | Goldman Sachs |
| Merger and acquisition analysis | Merrill Lynch |
| Customer profiling | Morgan Stanley |
| Branch performance | UBS AG |
| Portfolio management | Wells Fargo |
| **Automotive** | |
| Inventory and supply chain management | DaimlerChrysler AG |
| Resource utilization | Ford |
| Negotiating leverage with suppliers | General Motors |
| Warranty tracking and analysis | Honda |
| Profitability analysis and market segmentation | Toyota |

business information into the hands of more decision makers. Table 4.4 lists sample industry uses of data warehouses. Data warehouses can take up hundreds of gigabytes (even terabytes) of data. They usually run on fairly powerful mainframe computers and can cost millions of dollars.

Data warehouses represent more than just big databases. An organization that successfully deploys a data warehouse has committed to pulling together, integrating, and sharing critical corporate data throughout the firm.

**Data Marts** Rather than storing all enterprise data in one data warehouse, many organizations have created multiple data marts, each containing a subset of the data for a single aspect of a company's business—for example, finance, inventory, or personnel. A **data mart** is a data warehouse that is limited in scope. It contains selected information from the data warehouse such that each separate data mart is customized for the decision support applications of a particular end-user group. For example, an organization may have several data marts, such as a marketing data mart or a finance data mart, that are customized for a particular type of user. Data marts have been popular among small and medium-sized businesses and among departments within larger organizations, all of which were previously prohibited from developing their own data warehouses because of the high costs involved.

Williams-Sonoma, for example, known for its high-class home furnishing stores, is constantly looking to find new ways to increase sales and reach new target markets. Some of the most important data are coming from their catalog mailings, a database that contains 33 million active U.S. households. Using SAS data-mining tools and different models, Williams-Sonoma can segment customers into groups of 30,000 to 50,000 households and can predict the profitability of those segments based on the prior year's sales. These models resulted in the creation of a new catalog for a market segment that had up to then not been served by Williams-Sonoma. Now, for example, Williams-Sonoma markets a variety of new products, such as fringed lamps, chic furniture, and cool accessories, to an identified market segment using its Pottery Barn Teen catalog.

Data marts typically contain tens of gigabytes of data as opposed to the hundreds of gigabytes in data warehouses. Therefore, data marts can be deployed on less powerful hardware. The differences in costs between different types of data marts and data warehouses can be significant. The cost to develop a data mart is typically less than $1 million, while the cost for a data warehouse can exceed $10 million. Clearly, organizations committed to getting the most out of their data must make a large investment in database technology.

While hard data from inside and outside sources are very important to an organization's success, another key ingredient is the employee's knowledge. However, capturing this knowledge and using it for business intelligence processes can be a formidable task. To better capture and use employee's knowledge, companies are turning to knowledge management.

**Increasing Business Intelligence with Knowledge Management** There is no universal agreement on what exactly is meant by the term "knowledge management." In general, however, **knowledge management** refers to the processes an organization uses to gain the greatest value from its knowledge assets. In Chapter 1, we contrasted data and information as well as knowledge and wisdom. Recall that data are raw material—recorded, unformatted information, such as words or numbers. Information is data that have been formatted and organized in some way so that the result is useful to people. We need knowledge to understand relationships between different pieces of information; wisdom is accumulated knowledge. Consequently, what constitutes **knowledge assets** are all the underlying skills, routines, practices, principles, formulas, methods, heuristics, and intuitions, whether explicit or tacit. All databases, manuals, reference works, textbooks, diagrams, displays, computer files, proposals, plans, and any other artifacts in which both facts and procedures are recorded and stored are considered knowledge assets (Winter, 2001). From an organizational point of view, properly used knowledge assets enable an organization to improve its efficiency, effectiveness, and, of course, profitability.

Knowledge assets can be distinguished as being either explicit or tacit (Santosus and Surmacz, 2001). **Explicit knowledge assets** reflect knowledge that can be documented, archived, and codified, often with the help of information systems. Explicit knowledge assets reflect much of what is typically stored in a database management system. In contrast, **tacit knowledge assets** reflect the processes and procedures that are located in a person's mind on how to effectively perform a particular task (see Figure 4.26). Identifying key tacit knowledge assets and managing these assets so that they are accurate and available to people throughout the organization remains a significant challenge.

Tacit knowledge assets often reflect an organization's **best practices**—procedures and processes that are widely accepted as being among the most effective and/or efficient. Identifying how to recognize, generate, store, share, and manage this tacit knowledge is the primary objective for deploying a knowledge management system. Consequently, a **knowledge management system** is typically not a single technology but rather a collection of technology-based tools that include communication technologies—such as e-mail, groupware, instant messaging, and the like—as well as information storage and retrieval systems—such as a database management system—to enable the generation, storage, sharing, and management of tacit and explicit knowledge assets (Malhotra, 2005).

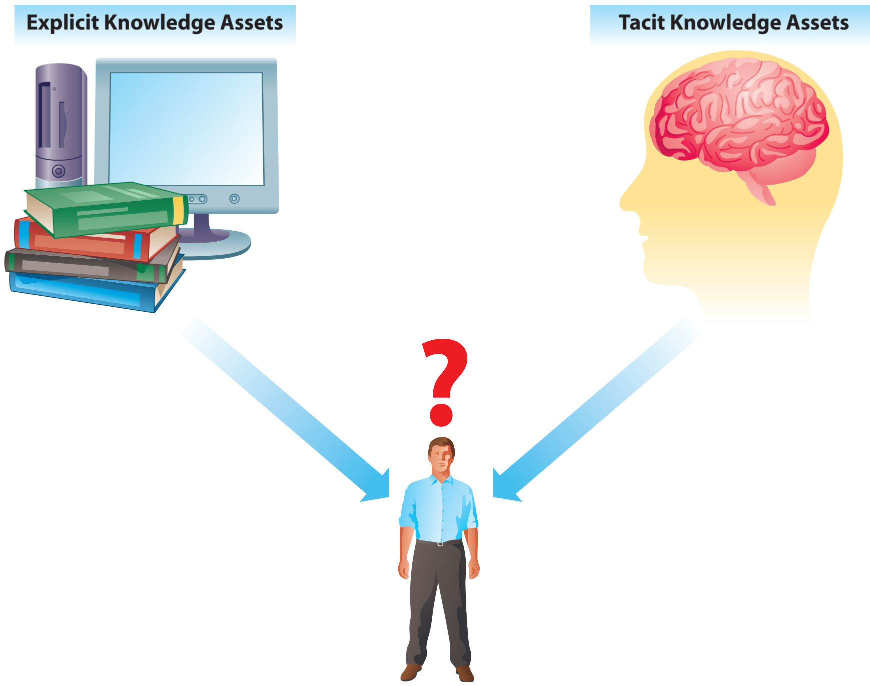


**FIGURE 4.26**

Explicit knowledge assets can easily be documented, archived, and codified, whereas tacit knowledge assets are located in a person's mind.

## Managing the Facilities Infrastructure

As described previously, data and knowledge are key assets for most business organizations. Thus, it is vital to ensure that the different IS infrastructure components (such as hardware, software, communication, and data and knowledge) are available when needed and are unaffected by potential outside influences. Given various threats to IS infrastructure components (such as storm, power outages, earthquakes, and so on), organizations have to take great care on where and how to house the infrastructure. In the following sections, we will discuss different aspects of facilities and how they are secured from outside influences.

**Ensuring Availability** As many potential causes of disasters cannot be avoided (there's no way to stop a hurricane), organizations should attempt to plan for the worst and protect their infrastructure accordingly. For companies operating in the digital world, the information systems infrastructure is often mission critical, so special care has to be taken to secure it. Whereas some applications can tolerate some downtime in case something malfunctions or disaster strikes, other applications (such as UPS's package tracking databases) can't tolerate any downtime—these companies need 24/7/365 reliability.

In order to provide for uninterrupted service, the infrastructure is usually housed in high-availability facilities; such facilities are equipped with different features to assure availability and reliability (see Figure 4.28). The facilities for UPS in Atlanta, Georgia, and Mahwah, New Jersey, are prime examples for such high-availability facilities. To ensure uninterrupted service, the data centers are self-sufficient, and each can operate for up to two days on self-generated power. The power is needed not only for the computers but also for air conditioning, as each facility needs air-conditioning capacity equaling that of or more than 2,000 homes. In case power fails, the cooling is provided using more than 600,000 gallons of chilled water, and the UPS facilities even have backup wells in case the municipal water supply should fail. Other protective measures include raised floors (to protect from floods) and buildings designed to withstand winds of 200 miles per hour. As you can imagine, such facilities are highly complex, and monitoring the operations can be a difficult task, as there are more than 10,000 data points to observe (such as temperature readings, power surges, and so on). To help manage this infrastructure component, these facilities have been designed so that they can be monitored from a single laptop computer.

Many (especially smaller) organizations do not need facilities the size of one of the UPS data centers; instead, they may just need space for a few servers. For such needs,

## Change Agents

### Larry Page and Sergey Brin, Cofounders of Google (Figure 4.27)

Thanks to Larry Page's and Sergey Brin's creation, everyone who uses the Internet knows that to "google" means to search for information. The coined verb has come into popular use since 1998, when Google, an Internet search engine, first went online. In August 2004, Brin and Page took the company public on NASDAQ and that same year announced first-quarter results as a public company with record revenues of $805.9 million.

Google is unique among search engines in that it indexes more content than most other search engines and finds relevant Web pages in less time. It can also search through a database of almost a million pictures, can personalize searches using a personal profile users create, and produce maps in conjunction with local searches such as "restaurants" in addition to offering free e-mail accounts. Additionally, Google is constantly releasing new applications and fighting for much more turf in the search engine business. Today, Brin and Page are reportedly worth approximately $12.8 billion each. (They recently bought a Boeing 767 for their private use.) Avid environmentalists, the two Google founders drive hybrid cars and encourage their employees to do the same. Brin and Page also came up with several innovative ideas to keep their employees' morale high; among the perks for Google employees are roller-hockey games in the company's parking lot twice a week, on-site workout and massage rooms, or one day per week to spend on their favorite innovations and projects.

In accordance with the company's informal motto "don't be evil," Brin and Page have established Google.org as Google's philanthropic arm (which also includes the Google Foundation). Drawing on Google's resources such as talent or technology, Google.org tries to address some of the world's most pressing problems—including poverty, energy, and the environment—by funding studies to improve water supplies in Western Kenya or programs to increase literacy in India.

Sources: http://en.wikipedia.org/wiki/Lawrence_E._Page
http://en.wikipedia.org/wiki/Sergey_Brin
http://news.bbc.co.uk/2/hi/business/3666241.stm
http://www.google.org/

**FIGURE 4.27**

Sergey Brin and Larry Page, cofounders of Google.


Source: http://www.google.com/press/images.html.

**FIGURE 4.28**

High-availability facilities feature sturdy construction, backup generators, air conditioning, access control, intrusion detection systems, and fire suppression systems.


**FIGURE 4.29**

Collocation facilities allow organizations to rent secure space for their infrastructure.

Source: http://www.sungard.com/corporate/general_pictures.htm.

companies can turn to **collocation facilities**. Organizations can rent space (usually in the form of cabinets or shares of a cabinet; see Figure 4.29) for their servers in such collocation facilities. Organizations managing collocation facilities provide the necessary infrastructure in terms of power, backups, connectivity, and security.

**Securing the Facilities Infrastructure** An organization's information systems infrastructure always needs to be secured to prevent it from outside intruders. Thus, no matter whether your server is located in a cabinet within your organization or you have rented space in a collocation facility, you should have physical safeguards in place to secure the equipment. Common safeguards include access control, closed-circuit television monitoring (see Chapter 6), and intrusion detection systems. We will discuss these safeguards and other security issues in later chapters.

## Managing the Human Resource Infrastructure

With the increased sophistication of the information systems infrastructure, organizations trying to manage their own infrastructure are facing the need for a highly trained workforce. However, access to the necessary human resource infrastructure is not a given in many rural areas. Over time, certain areas have become known for the availability of talented staff in a certain sector, and thus organizations operating in that sector tend to set up shop in such areas. Such areas are often characterized by a high quality of life for the people living there, and it is no surprise that many companies in the information technology sector are headquartered in Silicon Valley, California, or Seattle, Washington. In other areas, organizations can often not depend on an existing human resource infrastructure and have to attract people from other areas. In such cases, organizations have to find ways to both attract and retain employees.

Human resource policies provide another approach for assuring an adequate supply of skilled personnel. For example, many organizations provide educational grants or expense-matching programs to encourage employees to improve their education and skills. Typically, after receiving continuing education benefits, employees must agree to remain with the organization for some specified period of time or be forced to repay the employer. Other human resource policies, such as telecommuting, flextime, and creative benefit packages, can help to attract and retain the best employees.

With increasing globalization, other regions throughout the world are boasting about their existing human resource infrastructure. One such example is the Indian city of Bangalore, where, over a century ago, Maharajas started to lure talented people to the region to build a world-class human resource infrastructure. Although this has certainly helped to attract top Indian companies and multinational corporations alike, many companies have recently started complaining about other infrastructure issues, such as bad roads, power outages, housing conditions, traffic jams, and heavy rains. Clearly, for an area, just having a good human resource infrastructure is not sufficient, as organizations have to

**TABLE 4.5 Organizations Use Different Services to Support Their Infrastructure Needs**

| IS Infrastructure Component | Service | Example |
|---|---|---|
| Hardware | Utility computing | Organizations pay for processing or data storage on an as-needed basis |
| Software | Application service provider (ASP) | Organizations use a payroll system hosted on an ASP's server |
| Communication and collaboration | Videoconferencing | Organizations install HP HALO rooms and pay a monthly fee for usage and support |
| Data and knowledge | ASP | Data from applications hosted on an ASP's server is stored by the provider |
| Facilities | Collocation facility | Companies rent space for their servers in a collocation facility |

balance all their infrastructure needs when deciding where to move their headquarters or where to set up a new local subsidiary.

### Managing the Services Infrastructure

When operating in today's digital world, organizations have to rely on a complex information systems infrastructure. For many (especially smaller) organizations, maintaining such infrastructure is beyond their means because of the costs for maintaining and upgrading hardware and software, employing in-house experts for support, and so on. Thus, organizations big and small are turning to outside service providers for their infrastructure needs. As you read throughout this chapter, organizations can use different services to support their infrastructure needs. Table 4.5 provides examples for different types of services offered. Next, we will discuss an additional form of services, namely, outsourcing.

## Brief Case

### Toyota's Savvy Chief Information Officer

Some jingoistic Americans who refuse to buy "foreign" cars have also boycotted Toyota. They might be surprised to learn that, in fact, Toyota operates 14 manufacturing facilities in North America. Toyota employs over 37,000 Americans, and in 2006 the company produced 15 million cars in North America—a 20-year high—investing nearly $16.2 billion in local manufacturing and sales.

As with any large corporation, Toyota's IT investments are also large. Experts report, however, that until 2005, Toyota's IT departments had failed to keep up with increases in production and sales—especially in sales and customer support services.

After closely examining Toyota's IT services in 2003, Barbra Cooper, Toyota's chief information officer, found that these failures were due not to poor software choices, poor implementation, lack of appropriate IT personnel, or lack of expertise but rather to lack of trust in and respect for Toyota's American IT departments. Cooper had been hired in 1996 to oversee IT projects and realized early on that the IT and business departments seldom communicated closely or worked together to solve IS problems. During her retrospective examination of Toyota's IT services, Cooper realized that a major problem was that the IT staff was buried under the weight of six enterprise-wide projects and that staff was spread so thin that they could not adequately address any of the projects. Called the "big six," the projects included a new extranet for dealers, a PeopleSoft ERP implementation, a sales management system, a forecasting system, an advanced warranty system, and a documents management system.

Toyota's head office (business vs. IT) believed all six projects were vital and failed to realize that the IT department could not possibly successfully complete them. Cooper sought to bring the business and IT departments together so that priorities, time lines, and practical delivery dates could be discussed by both sides. Cooper also sought corporate headquarters' help in defining a new IT projects approval process in which the IT department would be involved in all planning phases.

In the beginning, some business managers and other employees did not like the new cooperative planning process, probably because they were reluctant to concede the control they had previously enjoyed over IT projects. They would no longer be able to develop their own IT initiatives and simply deliver their ideas to the IT department for implementation.

The new IT project approval process that Cooper recommended would also result in extensive changes in her IT department. Thus, already-stretched IT employees feared that their responsibilities would increase even more. In the past, Cooper had been protective of her employees' time and of projects undertaken by the IT department, and now she resolved to ensure that IT capabilities be considered before projects were assigned.

As Toyota's IT department met deadlines and delivered high-quality information systems, Cooper and her department gained respect. Toyota's business executives were further impressed when, in 2005, IT services reported a reduction in project costs of 16 percent—translating into several million dollars in savings. As a result, Cooper was given the go-ahead to revamp Toyota's IS model worldwide.

#### Questions

1. What is more important, leadership or infrastructure?
2. Can you have great information systems infrastructure without a great leader?

Sources: Thomas Wailgum, "The Big Fix," *CIO* (April 15, 2005), http://www.cio.com/archive/041505/toyota.html

http://www.toyota.com/about/operations/manufacturing/index.html

**Outsourcing** As we have defined in Chapter 1, outsourcing is the turning over of partial or entire responsibility for information systems development and/or management to an outside organization. In a prior section, we talked about application service providers, where an organization uses software services from another organization. In contrast, outsourcers such as Accenture provide services to organizations. Such services can be conducting business processes, such as finance, accounting, or human resources, or the services provided can be the development and maintenance of software or the management of an organization's technology infrastructure. As with other infrastructure solutions such as on-demand computing or on-demand software, outsourcing can help a company focus on its core processes without having to worry about supporting processes. While outsourcing has seen a tremendous increase over past years, outsourcing is typically limited to noncore business functions. However, there are some noncore business functions that organizations tend to keep within their own realm. For example, although more and more companies outsource the management of their information systems infrastructure, only very few outsource information systems security, as it is regarded as being critical for an organization's survival (CSI, 2006).

## Ensuring a Reliable and Secure Infrastructure

In the previous sections, you have read about how organizations get most out of their data, how to maximize available resources, and how to use outside providers for the infrastructure needs. While these are all very important issues for an organization to manage, there is one issue that is even more critical for effective infrastructure management. Specifically, what happens when something goes seriously wrong? For organizations such as Amazon.com, a network outage can quickly lead to millions of dollars in terms of lost revenue. Unfortunately, there are many events that can lead to catastrophic system failures, such as natural disasters, criminal activities, or just plain accidents. The most common causes of disaster are power outages, hardware failures, and floods (see Figure 4.30). How can companies managing in a digital world avoid such disasters? In a prior section, you learned how companies proactively try to avoid a disaster by building and maintaining high-availability facilities or by renting space in a collocation facility. In the following


**FIGURE 4.30**

Power outages, hardware failures, and floods are the most common causes of disaster.

section, you will learn how organizations can attempt to limit the impacts of a disaster or plan for and recover from a disaster should one happen.

## Disaster Planning

In some cases, all attempts to provide a reliable and secure information systems infrastructure are in vain, and disasters cannot be avoided. Thus, organizations need to be prepared for when something catastrophic occurs. The most important aspect of preparing for disaster is creating a **disaster recovery plan**, which spells out detailed procedures for recovering from systems-related disasters, such as virus infections and other disasters that might cripple the information systems infrastructure. This way, even under the worst-case scenario, people will be able to replace or reconstruct critical files or data, or they will at least have a plan readily available to begin the recovery process, including a list of service providers. A typical disaster recovery plan includes information that answers the following questions:

- What events are considered a disaster?
- What should be done to prepare the backup site?
- What is the chain of command, and who can declare a disaster?
- What hardware and software is needed to recover from a disaster?
- Which personnel are needed for staffing the backup sites?
- What is the sequence for moving back to the original location after recovery?

**Backup Sites** **Backup sites** are critical components in a disaster recovery plan, as they allow businesses to continue functioning in the event a disaster strikes; in other words, backup sites can be thought of as a company's office in a temporary location. Commonly, a distinction is made between cold and hot backup sites. These are discussed next.

**COLD BACKUP SITES.** A **cold backup site** is nothing more than an empty warehouse with all necessary connections for power and communication but nothing else. In the case of a disaster, a company has to first set up all necessary equipment, ranging from office furniture to Web servers. While this is the least expensive option, it also takes a relatively longer time before a company can resume working after a disaster.

**HOT BACKUP SITE.** A **hot backup site**, in contrast, is a fully equipped backup facility, having everything from office chairs to a one-to-one replication of the most current data. In the event of a disaster, all that has to be done is to relocate the employees to the backup site to continue working. Obviously, this is a very expensive option, as the backup site has to be kept fully equipped and all the information systems infrastructure duplicated. Further, hot backup sites also have a redundant backup of the data so that the business processes are interrupted as little as possible. To achieve this redundancy, all data are **mirrored** on a separate server (i.e., everything is stored synchronously on two independent systems). This might seem expensive, but for a critical business application involving customers, it may be less expensive to run a redundant backup system in parallel than it would be to disrupt business or lose customers in the event of catastrophic system failure.

**CHOOSING A BACKUP SITE LOCATION.** Thinking about the location of redundant systems is an important aspect of disaster planning. If a company relies on redundant systems, all of which are located within the same building, a single event can incapacitate both systems. Similarly, events such as a hurricane can damage systems that are located across town from each other. Thus, even if the primary infrastructure is located in-house, when having redundant systems, it pays to have the backup located in a different geographic area to minimize the risk of a disaster happening to both systems.

## Designing the Recovery Plan

When planning for disaster, two objectives should be considered by an organization: recovery time and recovery point objectives. **Recovery time objectives** specify the maximum time allowed to recover from a catastrophic event. For example, should the organization be able to resume operations in minutes, hours, or days after the disaster? Having completely redundant systems helps to minimize the recovery time and might be best suited for mission-critical applications, such as e-commerce transaction servers. For other applications, such as data mining, while important, the recovery time can be longer without disrupting primary business processes.

Additionally, **recovery point objectives** specify how current the backup data should be. Imagine that your computer's hard drive crashes while you are working on a term paper. Luckily, you recently backed up your data. Would you prefer the last backup to be a few days old, or would you rather have the last backup include your most recent version of the term paper? Having completely redundant systems that mirror the data helps to minimize (or even avoid) data loss in the event of a catastrophic failure.

The disaster recovery plan is just one part of an overall plan for effectively managing the information systems infrastructure. In Chapter 6, we outline a comprehensive information systems security plan that includes disaster recovery.

### Information Systems Controls, Auditing, and the Sarbanes-Oxley Act

As you have seen, there are a variety of components to consider when managing the information systems infrastructure. No matter how organizations choose to manage their infrastructure, **information systems controls** have to be put into place to control costs,

## BlackBerry

What would physicians, emergency medical technicians, trial attorneys, government decision makers, military officers, and countless other professionals do without their pagers, cell phones—and BlackBerries?

In 1999, Research in Motion (RIM), based in Waterloo, Ontario, introduced a wireless device it named BlackBerry, after the berry-like buttons on the product. A BlackBerry fits in a user's palm and is operated using a trackwheel and buttons. When first introduced, BlackBerries concentrated on e-mail, but they now support push e-mail (received in real time), mobile telephoning, Internet faxing, text messaging, Web browsing, and other wireless information services. BlackBerries are noted for their reliability, offering connectivity when other telecommunication devices fail.

In the early 2000s, NTP, Incorporated, a Virginia-based patent-holding company, sent notice of their wireless telecommunications patents to several fledgling wireless companies, offering to license their patents to them. None of the companies bought patent licenses from NTP. NTP then sued one of the companies, RIM, claiming patent infringement. RIM claimed in court that a functional wireless e-mail system was already in the public domain before NTP's inventions. The jury, however, found for NTP, and RIM was fined several million dollars. The case continued through several appeals but was finally settled in 2006. RIM agreed to pay NTP $612.5 million "in full and final settlement of all claims." RIM also announced the use of newly developed technology that would remove all question of whether the BlackBerry used NTP-patented technology. BlackBerry users, who totaled more than 3 million in March 2006, were relieved when the case was settled since many had feared that the beloved devices would be shut down by court order. In fact, the U.S. Department of Defense had testified during the patent infringement lawsuit that loss of the BlackBerry network would be a threat to national security since so many government employees used the device.

Although NTP's patent infringement case was settled, the situation raised several questions. When similar innovations are developed independently but simultaneously, who actually "owns" the invention? If a company pays to license a patent, how long should the company pay royalties on the product using the patent? And, perhaps most important, how secure is intellectual property in a flat world, where technological advances emerge virtually overnight?

Sources: Wikipedia.org/BlackBerry: http://en.wikipedia.org/wiki/BlackBerry

Mark Heinzl and Amol Sharma, "RIM to Pay NTP $612.5 Million to Settle BlackBerry Patent Suit," *Wall Street Journal* (March 4, 2006), http://online.wsj.com/article_email/SB114142276287788965-lMyQjAxMDE2NDAxMzQwMjMyWj.html gain and protect trust, remain competitive, or comply with internal or external governance (e.g., the Sarbanes-Oxley Act, discussed later in this section). Such controls, which help ensure the reliability of information, can consist of a variety of different measures, such as policies and their physical implementation, access restrictions, or record keeping, to be able to trace actions and transactions and who is responsible for these. IS controls thus need to be applied throughout the entire IS infrastructure. To be most effective, controls should be a combination of three types of controls:

- Preventive controls (to prevent any potentially negative event from occurring, such as by preventing outside intruders from accessing a facility)
- Detective controls (to assess whether anything went wrong, such as unauthorized access attempts)
- Corrective controls (to mitigate the impacts of any problem after it has arisen, such as restoring compromised data)

One way to conceptualize the different forms of controls is by a hierarchy ranging from high-level policies to the implementation at the application level (see Figure 4.31 for the hierarchy of controls; note that the categories are not necessarily mutually exclusive); Table 4.6 gives a brief explanation of the different types of controls and presents examples for each. You have learned about a variety of IS controls in prior sections, and while reading this book, you will continue to come across the different elements of control. In the following sections, we will describe how companies use IS auditing to assess the IS controls in place and whether further IS controls need to be implemented or changed.

**Information Systems Auditing** Analyzing the IS controls should be an ongoing process for organizations. However, often it can be beneficial for organizations to periodically have an external entity review the controls so as to uncover any potential problems. An **information systems audit**, often performed by external auditors, can help organizations assess the state of their information systems controls to determine necessary changes and to help ensure the information systems' availability, confidentiality, and integrity. The response to the strengths and weaknesses identified in the IS audit is often determined by the potential risks an organization faces. In other words, the IS audit has to assess whether the IS controls in place are sufficient to address the potential risks. Thus, a major component of the IS audit is a **risk assessment**, which aims at determining what

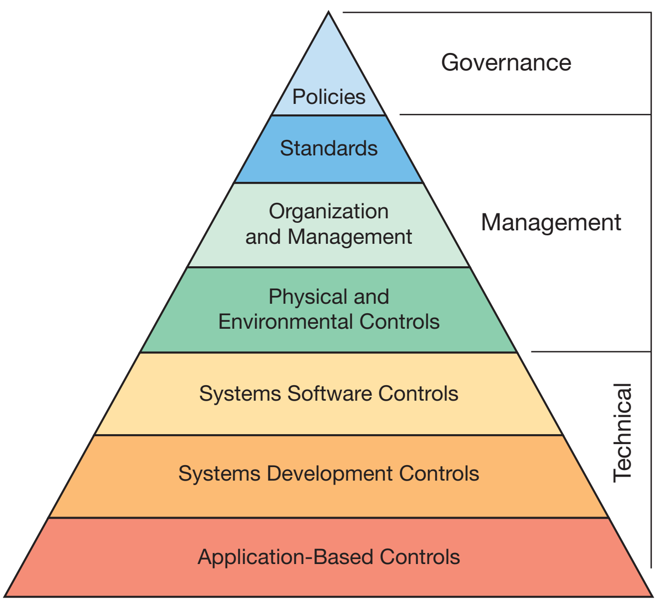


**FIGURE 4.31**

Hierarchy of IS controls.

Source: http://infotech.aicpa.org/

**TABLE 4.6 Different Types of Information Systems Controls**

| Type of Control | What Is It For? | Examples |
|---|---|---|
| Policies | Define aims and objectives of the organization | General policies about: Security and privacy, Rights of access, Data and systems ownership, End-user development, Access to sensitive areas (e.g., high-availability facilities), Disaster planning |
| Standards | Support the requirements of policies | Standards about:<br>Systems development process<br>Systems software configuration<br>Application controls<br>Data structures<br>Documentation |
| Organization and management | Define lines of reporting to implement effective control and policy development | Policies about:<br>Security and use<br>Account authorization<br>Backup and recovery<br>Incident reporting |
| Physical and environmental controls | Protect the organization's IS assets | High-availability facilities<br>Collocation facilities |
| Systems software controls | Enable applications and users to utilize the systems | Control access to applications<br>Generate activity logs<br>Prevent outside intrusion (e.g., by hackers) |
| Systems development and acquisition controls | Ensure that systems meet the organization's needs | Document user requirements<br>Use formal processes for systems design, development, testing, and maintenance |
| Application-based controls | Ensure correct input, processing, storage, and output of data; maintain record of data as it moves through the system | Input controls (such as checking the input into a Web form)<br>Processing controls<br>Output controls (comparing the outputs against intended results)<br>Integrity controls (ensure that data remains correct)<br>Management trail (keep record of transactions to be able to locate sources of potential errors) |

type of risks the organization's IS infrastructure faces, the criticality of those risks to the infrastructure, and the level of risks the organization is willing to tolerate. To determine the potential risks, an organization needs to first identify the infrastructure components that are at risk, identify potential vulnerabilities, and map the vulnerabilities to the potential threats. Then the probability of each event's occurring and its potential impact should be estimated. This process will lead to several important questions that must be answered, including the following:

- What are the costs of restoring the data in case of a threat event?
- What are potential legal costs in case confidential data are lost?
- What are the costs to the business if core systems are unavailable for a certain period of time?

Depending on the nature of the risks, the level of risk tolerance, and the severity of the risks identified, an organization can follow various steps. These steps are reducing or eliminating the risk (by implementing stricter IS controls), sharing or transferring the risk (e.g., by outsourcing certain functions to a highly skilled service provider), or just accepting the risk (in case the risk is deemed not critical to an organization's success) (see also Chapter 6).

Once the risk has been assessed, auditors have to evaluate the organization's internal controls. During such audits, the auditor tries to gather evidence regarding the effectiveness of the controls. However, testing all controls under all possible conditions is very inefficient and often infeasible. Thus, auditors frequently rely on **computer-assisted auditing tools (CAAT)**, which is specific software to test applications and data, using test data or simulations. In addition to using specific auditing tools, auditors use audit sampling procedures to assess the controls, enabling the audit to be conducted in the most cost-effective manner. Once the audit has been performed and sufficient evidence has been gathered, reports are issued to the organization. Usually, such reports are followed up with a discussion of the results and potential courses of action.

**The Sarbanes Oxley Act** Performing an IS audit can help an organization reduce costs or remain competitive by identifying areas where IS controls are lacking and need improvement. Another major factor that has contributed to a high demand for IS auditors is the need to comply with government regulations, most notably the **Sarbanes-Oxley Act** (hereafter S-OX) of 2002. Formed as a reaction to large-scale accounting scandals that led to the downfall of corporations such as WorldCom and Enron, the act primarily addresses the accounting side of organizations. However, given the importance of an IS infrastructure and IS controls for an organization's financial applications, it is of major importance to include IS controls in compliance reviews.

According to S-OX, companies have to demonstrate that there are controls in place to prevent misuse or fraud, controls to detect any potential problems, and effective measures in place to correct any problems; S-OX goes so far that corporate executives face jail time and heavy fines if the appropriate controls are not in place or are ineffective. The information systems architecture plays a key role in S-OX compliance, given that many controls are information-systems based, providing capabilities to detect information exceptions and to provide a management trail for tracing exceptions. However, S-OX itself barely addresses IS controls specifically; rather, it addresses general processes and practices, leaving companies wondering how to comply with the guidelines put forth in the act. Further, it is often cumbersome and time consuming for organizations to identify the relevant systems to be audited for S-OX compliance. Thus, many organizations find it easier to review their entire IS infrastructure, following objectives set forth in guidelines such as the **control objectives for information and related technology (COBIT)**—a set of best practices that help organizations both maximize the benefits from their IS infrastructure and establish appropriate controls.

Another issue faced by organizations because of S-OX is the requirement to preserve evidence to document compliance and for potential lawsuits. Since the inception of S-OX, e-mails and even instant messages have achieved the same status as regular business documents and thus need to be preserved for a period of time, typically up to seven years. Failure to present such documents in case of litigious activity can lead to severe fines being imposed on companies and their executives, and courts usually will not accept the argument that a message could not be located. For example, the investment bank Morgan Stanley faced fines up to $15 million for failing to retain e-mail messages. On the surface, it seems easiest for an organization to simply archive all the e-mail messages sent and received. However, such a "digital landfill" where everything is stored can quickly grow to an unmanageable size, and companies cannot comply with the mandate to present any evidence in a timely manner. Thus, many organizations turn to e-mail management software that archives and categorizes all incoming and outgoing e-mails based on key words. Even using such specialized software, finding e-mails related to a certain topic within the archive can pose a tremendous task: some analysts estimate that a business with 25,000 employees generates over 4 billion e-mail messages over the course of seven years (not counting any increase in e-mail activity), which will be hard to handle for even the most sophisticated programs.

## Industry Analysis

### Radio

The new satellite and high-definition radio stations have been called "jukeboxes on steroids." Generally, there are no ads and no DJs—just uninterrupted tunes.

Satellite radio, also known as subscription radio, operates via digital signals received from low-orbiting satellites. In the absence of paid advertising, subscribers pay monthly fees for a bundle of music "channels," each one devoted to talk shows or to various types of music, such as country, hip-hop, classical, and so on.

Currently, the prominent players in the satellite radio market are XM Radio and Sirius in North America and WorldSpace in Europe, Asia, and Africa. All of the signals are proprietary—that is, services do not share equipment. Instead, each service owns its own licensed equipment for receiving and relaying signals. The popularity of satellite radio is beginning to affect traditional FM market share, as evidenced by the gravitation of some big names in radio, such as Howard Stern, to satellite stations.

In addition to satellite radio, high-definition (HD) radio has recently become popular with listeners. HD radio is similar to the satellite model in that a digital signal provides CD-quality listening. This technology claims to eliminate the static and hiss associated with the FM analog signal.

The digital medium also offers a promising multicasting feature using current radio equipment. "Multicasting" is the term used for subdividing a radio station's signal into multiple channels. Multicasting means that a station's listeners have a choice of stations that carry talk shows, various genres of music, sports, and so on, much like TV signals are divided into several channels. The potential for current radio stations is limitless since they can now offer several types of high-quality niche programming without having to buy additional equipment.

Clearly, the radio industry is joining other major industries in becoming globalized and digitized to become successful in today's world.

#### Questions

1. Contrast the infrastructure required to launch a satellite versus traditional radio station.
2. Today there are thousands of AM/FM stations. Forecast their future and provide a strategy for retaining and gaining market share.

Sources: Anonymous, "A New Radio Format, or Just Another Name?," *Audio Graphics* (April 24, 2006), http://www. audiographics.com/agd/042406-3.htm

Anonymous, "Business Briefs" (May 10, 2006), http://ws.gmnews.com/news/2006/0510/Business/013.html

## Key Points Review

1. ***List the essential information systems infrastructure components and describe why they are necessary for satisfying an organization's informational needs.*** Modern organizations heavily rely on information systems infrastructure; its components include hardware, software, communications and collaboration, data and knowledge, facilities, human resources, and services. While the computing hardware is integral to an organization's IS infrastructure, as it is also needed to store and process organizational data, networking hardware is needed to connect the different systems to allow for collaboration and information sharing. Software assists organizations in executing their business processes and competitive strategy. Consequently, with increased reliance on information systems for managing organizations, effectively utilizing software resources is becoming increasingly critical and complex. Communication and collaboration is one of the reasons why information systems have become so powerful and important to modern organizations. The ability to interconnect computers, information systems, and networks ultimately allows the interconnection of both internal and external business processes, facilitating improved communication and collaboration. Data and knowledge are probably among the most important assets an organization has, as data and knowledge are

essential for both gaining business intelligence and executing business processes. Facilities, though not directly needed to support business processes or business intelligence, are necessary for the information systems infrastructure. Human resources are also an important component in information system infrastructure. Although even large facilities do not require large support staff, the staff that is needed should be well trained. Finally, a broad range of services are needed to support the information system infrastructure.

2. ***Describe solutions organizations use to design a reliable, robust, and secure infrastructure.*** Many organizations now turn to on-demand computing for fluctuating computation needs, utility computing for "renting" of resources, grid computing for solving large-scale problems, edge computing for providing a more decentralized use of resources, and autonomic computing for increasing reliability. To manage the ever-increasing complexity of software needs, organizations turn to open-source software to increase their independence, use Web services to integrate different applications housed on different systems, and use patch management and software asset management systems to keep their systems current. In other cases, organizations want to free themselves from having to address such issues and use on-demand software provided by application service providers. The convergence of computing and telecommunications has helped organizations address their diverse communication needs, such as by enabling voice over IP or videoconferencing over IP. Often, companies implement wireless infrastructures to increase their employees' mobility. To support more effective businesses and to gather business intelligence, organizations have to find ways to manage vast amounts of data, usually using online transaction processing and online analytical processing. Data warehouses and data marts support the integration and analysis of large data sets. Knowledge management systems are a family of tools helping to organize, store, and retrieve a company's tacit and explicit knowledge assets. Finally, organizations have to manage their facilities to ensure security and availability, have to manage the human resource infrastructure to attract and retain qualified personnel, and have to manage the use of different services, often using outsourcing.

3. ***Describe how organizations can ensure a reliable and secure infrastructure, plan for potential disasters, and establish IS controls.*** In some cases, all attempts to provide a reliable and secure information systems infrastructure are in vain, and disasters cannot be avoided. Thus, organizations need to be prepared when something catastrophic occurs. The most important aspect of preparing for disaster is creating a disaster recovery plan, which spells out detailed procedures for recovering from systems-related disasters, such as virus infections and other disasters that might cripple the information systems infrastructure. This disaster plan should include decisions about where to back up and whether this backup site be hot or cold. A hot backup site completely replicates data and facilities; a cold site is an empty facility with only power and network connectivity. Reliability is also enhanced through the design of a comprehensive disaster recovery plan that outlines recovery goals and tactics. IS controls can help ensure a secure and reliable infrastructure; such controls should be a mix of preventive, detective, and corrective controls. To assess the efficacy of these controls, organizations frequently conduct information systems audits to determine the risks an organization faces and how far the IS controls can limit any potentially negative effects. Further, organizations perform IS audits to comply with government regulations, most notably, the Sarbanes-Oxley Act of 2002. According to S-OX, companies have to demonstrate that there are controls in place to prevent misuse or fraud, controls to detect any potential problems, and effective measures to correct any problems; S-OX goes so far that a business executive could face heavy fines or substantial jail time if appropriate controls are not in place or are ineffective. Performing thorough IS audits on a regular basis can help assess compliance to these regulations.

## Key Terms

embedded systems 141
enterprise licensing 144
explicit knowledge assets 154
grid computing 138
hot backup site 161
information systems audit 163
information systems controls 162
information systems infrastructure 130
informational systems 151
infrastructure 129
IP convergence 146
knowledge assets 154
knowledge management 154
knowledge management system 156
mirror 161
OLAP server 150
online analytical processing (OLAP) 150
online transaction processing (OLTP) 150
on-demand computing 137
on-demand software 144
open-source software 141
operational systems 151
patch management system 143
recovery point objectives 161
recovery time objectives 161
risk assessment 163
Sarbanes-Oxley Act 165
server farms 134
service-oriented architecture 142
shrink-wrap license 144
software asset management 144
supporting processes 130
tacit knowledge assets 154
utility computing 138
Web services 142
videoconferencing over IP 147
voice over IP (VoIP) 146
volume licensing 144

## Review Questions

1. List three reasons why Google chose The Dalles, Oregon, for its newest data center.
2. Describe what business intelligence (BI) is and how organizations use BI to gain a competitive advantage.
3. What is on-demand computing, and how can organizations use this technology to cut costs?
4. Define grid computing and describe its advantages and disadvantages.
5. List the five elements that are essential for autonomic computing.
6. Describe why companies would choose to implement open-source software over fully licensed software.
7. List and describe the two main types of software licenses.
8. Describe what is meant by the term IP convergence.
9. Describe the two main categories of wireless device for communication needs.
10. Compare and contrast data warehouses and data marts.
11. Define outsourcing and how it is used in organizations today.
12. Describe how the Sarbanes-Oxley Act impacts the management of the information systems infrastructure.

## Self-Study Questions

Visit the Interactive Study Guide on the text Web site for additional Self-Study Questions: **www.prenhall.com/ Jessup**.

1. __________ processes are the activities organizations perform in order to reach their business goals.
   A. core
   B. support
   C. business
   D. functional
2. In modern organizations, information system infrastructure heavily relies on all of the following except for __________.
   A. hardware
   B. communication
   C. human resources
   D. services
   E. finance
3. __________ is a special form of on-demand computing which is typically used to solve large-scale computing problems.
   A. grid
   B. utility
   C. access
   D. edge
4. Rather than having massive centralized computers and databases, smaller servers are now located closer to individual users. This is called __________ computing.
   A. edge
   B. grid
   C. utility
   D. access
5. Which of the following is *not* an example of open-source software?
   A. Linux
   B. OpenOffice
   C. Apache
   D. Windows Vista
6. This management system allows developers to account for unforeseen problems after the software was shipped to the customer.
   A. account management system
   B. patch management system

C. software bug system
D. software inventory system

7. This system helps organizations avoid the negative impacts of installing unlicensed or private software.
A. software asset management
B. patch management system
C. software bug system
D. software inventory system

8. What method is used to sort and analyze information to better understand an organization's customers, products, market, and so on?
A. OLTP
B. Web services
C. OLAP
D. data mining

9. __________ knowledge assets reflect the process and procedures that are located in a person's mind.
A. tacit
B. explicit
C. implicit
D. real

10. Which is a specific software tool used by auditors to test applications and data or run simulations of business transactions and processes?
A. Sarbanes-Oxley
B. Web services
C. computer-assisted auditing tools (CAAT)
D. risk assessment system

Answers are on page 171.

## Problems and Exercises

1. Match the following terms with the appropriate definitions:
   i. Utility computing
   ii. Web services
   iii. Shrink-wrap license
   iv. Application service provider
   v. Voice over IP
   vi. OLTP
   vii. Knowledge management
   viii. Operational system
   ix. Data warehouse
   x. OLAP
   a. The integration of large databases into a single repository
   b. The use of the Internet technology for placing telephone calls
   c. A type of license that accompanies software primarily used by consumers
   d. Technology used for immediate automated response to requests of the users
   e. Technology used in transaction and analytical processing to interact with customers and run a business in real time
   f. The process an organization uses to gain the greatest value from its knowledge assets
   g. A form of on-demand computing where resources are rented on an as-needed basis
   h. Web-based software systems that allow the interaction of different programs over a network
   i. Internet technology that provides access to application software via a Web browser
   j. A graphical software tool that provides complex analysis of data
2. Akamai (www.akamai.com) distributes on average 10 to 20 percent of the Web's content on any given day. What function do they provide, and how do they accomplish this?
3. Do you feel that there should be human involvement in the resolution of network, hardware, and software problems? Autonomic computing allows for the computers to self-configure; describe some societal, ethical, and technical problems that might arise from this.
4. Several Fortune 500 companies now use voice over IP (VoIP) for virtually all of their corporate communication. What are the advantages to this technology for business infrastructure? What are some of the consumer VoIP technologies that are currently in place? Which one of these technologies do you see lasting?
5. The HP Halo system allows for videoconferences to mimic face-to-face conferences. Imagine you were in charge of designing a new videoconference system. What are some of the features you would include? Which features do you feel are less important?
6. Based on your experience with online transaction processing systems (in everyday life or in the workplace), which ones seem to work best? What characteristics did you judge the success of these systems by? Would you make any adjustments to the system?
7. Interview an IS employee within an organization at a university or a workplace. What are the important issues for the infrastructure of the organization? Have they had any experiences with data loss due to infrastructure problems? What backup plans do they have in place in case of disaster?
8. Describe how various systems in this chapter might enable organizations to become prepared for a disaster

such as a fire or flood. What technologies in particular are of critical importance?

9. Choose an organization with which you are familiar that uses databases. Then brainstorm how this organization could use data mining to gain an understanding of their customers, products, or marketplace. What are some key pieces of information that should be data mined for this organization?
10. Using a search engine, enter the key word "data warehousing." Who are the large vendors in this industry? What type of solutions do they offer to their clients? Do you see any common trends in the data warehousing businesses?
11. Using a Web browser, go to http://www.kmworld.com/. Looking at this site's current stories on knowledge management, find some current trends in knowledge management. Under the "Solutions" link, choose an industry of interest to you. What types of solutions are offered for this industry?
12. Find any Web site that could be considered a "knowledge management" site. What knowledge does the site capture? What are the important attributes included in this site? What suggestions would you have to make the site a better place to find or store knowledge?
13. Interview an IS professional and ask him or her about open-source software. Does he or she see all types of information systems to be candidates for open-source software? Additionally, find out what systems are most likely and least likely to be open source.
14. Browse BOINC's different grid computing projects at http://boinc.berkeley.edu/. Which projects seem most interesting? For which projects would you "donate" your computing resources? If you would not donate any resources, why not?
15. Interview an IS professional about his or her company's use of software asset management processes. How does he or she keep track of the different software installed? If anyone asked you about the software installed on your computer, would you know what you have installed? Would you be able to produce the licenses for each software installed?

## Application Exercises

The existing data files referenced in these exercises are available on the Student Companion Web site: **www.prenhall.com/Jessup**.

### Spreadsheet Application: Tracking Frequent Flier Mileage

You have recently landed a part-time job as a business analyst for Campus Travel. In your first meeting, the operations manager learned that you are taking an introductory MIS class. As the manager is not very proficient in using office software tools, he is doing all frequent flier mileage in two separate Excel spreadsheets. One is the customer's contact information, and the second is the miles flown. Being familiar with the possibilities of spreadsheet applications, you suggest setting up one spreadsheet to handle both functions. To complete this, you must do the following:

1. Open the spreadsheet frequentflier2.csv. You will see a tab for customers and a tab labeled "miles flown."
2. Use the vlookup function to enter the miles flown column by looking up the frequent flier number (Hint: If done correctly with absolute references, you should be able enter the vlookup formula in the first cell in the "miles flown" column and copy it down for all the cells).
3. Use conditional formatting to highlight all frequent fliers who have less than 4,000 total miles.
4. Finally, sort the frequent fliers by total miles in descending order and print out the spreadsheet.

### Database Application: Building a Knowledge Database

Campus Travel seems to be growing quite rapidly. Now they have franchises in three different states, totaling 16 locations. As the company has grown tremendously over the past few years, it has become increasingly difficult to keep track of the areas of expertise of each travel consultant; often, consultants waste valuable time trying to find out who in the company possesses the knowledge about a particular region. Impressed with your skills, the general manager of Campus Travel has asked you to add, modify, and delete the following records from its employee database:

1. Open employeedata.mdb.
2. Select the "employee" tab.
3. Add the following records:
   a. Eric Tang, Spokane Office, Expert in Southwest, Phone (509)555-2311
   b. Janna Connell, Spokane Office, Expert in Delta, Phone (509)555-1144
4. Delete the following record:
   a. Carl Looney from the Pullman office
5. Modify the following:
   a. Change Frank Herman from the Pullman office to the Spokane office
   b. Switch Ramon Sanchez's home number to (208)549-2544

## Team Work Exercise: Your Personal Communication Infrastructure Assessment

Work in a team of four or five students and have each person list the number of wired telephone calls, cellular telephone calls, instant messages, e-mail messages, and so on. For each instance, also inventory the total time spent, who the call was to, and the purpose of the call. Have each person also document his or her demographic and relevant personal information (e.g., age, gender, relationship status, children, and so on). Combine the results of all team members and discuss patterns and anomalies.

## Answers to the Self-Study Questions

| | | | | |
|---|---|---|---|---|
| **1.** C, p. 130 | **2.** E, p. 130 | **3.** A, p. 138 | **4.** A, p. 139 | **5.** D, p. 141 |
| **6.** B, p. 143 | **7.** A, p. 144 | **8.** D, p. 150 | **9.** A, p. 154 | **10.** C, p. 165 |

case 1

## Films Go Digital at the Sundance Film Festival

For independent films, the Sundance Film Festival is what film festivals in Cannes or Berlin are for more mainstream movies. This annual festival is sponsored by the (nonprofit) Sundance Institute based in Sundance, Utah, and was founded by actor Robert Redford more than 20 years ago. As other film festivals, the Sundance Film Festival chooses the best feature and short films in a variety of categories.

For the first time in Sundance history, in 2006 all entries in the short films category were available on the festival Web site. In fact, to date, Sundance is the only film festival to premiere films online. This new feature was so popular that after the films first aired online, there were over 40,000 e-mails expressing gratitude for the free screening.

Airing films online is not the only way the Sundance Film Festival uses information systems in innovative ways. The festival offers many filmmakers workshops and training sessions, and in 2006, a majority of the training was on digital technology since 30 percent of all films produced that year were in digital format. For many independent filmmakers, digital technology has opened vast opportunities, as it allows producing studio-quality films without having to rely on expensive lighting, film development, and postproduction facilities. Thus, people who could never afford all the necessary equipment can now produce movies digitally. Further, digital cameras and projectors and advances in software have made the transition from celluloid to digital more attainable for filmmakers, who until recently used traditional technology. It is no surprise that sponsors for the 2006 Sundance Film Festival included tech giants Hewlett Packard, Sprint, Adobe Systems, and Sony.

Increasingly, independent movie producers see the Web as a powerful distribution tool. In 2006, for example, entries for the Sundance Festival's online film category totaled nearly 300, creating the need to further divide this category into the subcategories of animation, live action, and interactive.

The success of the online presentations of the 2006 Sundance Film Festival proves that movie aficionados appreciate the opportunity to watch films on Web sites as well as in theaters or on home TV sets, and that digital creation and distribution of movies is here to stay.

### Questions

1. What other industries/events that currently are not posting creative content on the Web could benefit by following Sundance's lead?
2. Choose the perspective of either the artist or the general public and argue who benefits most from posting content onto the Web.
3. List the pros and cons for posting creative content on the Web.

Sources: Michelle Meyers, "Tech Plays Supporting Role at Sundance Festival," *CNet News* (January 18, 2006), http://news.com.com/Tech+plays+supporting+role+at+Sundance+festival/2100-1025_3-6028354.html?part=rss&tag=6028354&subj=news

Gwendolyn Mariano, "Sundance to Roll Film on Web Festival," *CNet News* (December 15, 2000), http://news.com.com/Sundance+to+roll+film+on+Web+festival/2100-1023_3-249997.html?tag=nl

http://festival.sundance.org/2006/

case 2

## e-Enabling the Air Transport Industry: Managing the Information Systems Infrastructure

The air transport industry has always been about reaction time. Acquiring and utilizing the most accurate and timely information is a key to success. To be most effective, large amounts of data from a variety of sources must be collected, analyzed, and shared in real time, using a variety of different infrastructure components. To get the right information to the right person at the right time is a huge challenge given the variety of people who need timely information to be most effective at their jobs.

For example, pilots must have the most up-to-date information about their aircraft, the weather, and the status of air traffic control in order to fly as efficiently and safely as possible. Mechanics need to know what condition the airplanes coming to their sites are in and how best to address any problems those airplanes might have. Airline flight operations centers need to know that they will have the aircraft and crew necessary in the right places at the right times for upcoming flights. Ground operations need to know where to bring fuel and catering items to service airplanes as they arrive at their gates. And passengers need to know what they are supposed to do if their flights are delayed or rerouted. The information technology revolution has gone a long way in helping bring all the information generated in the air transport enterprise to bear in order to increase safety, security, and efficiency. In fact, over the past several years, there have been vast improvements in the collection and dissemination of accurate and timely information throughout the various processes for supporting and operating an airplane. Nevertheless, the current systems and capabilities still have major shortcomings.

Specifically, the airplane itself, the prime generator of revenue and the prime consumer of operating funds, regularly unplugs from the rest of the information enterprise for as much as 14 hours at a time while in flight. Pilots have to navigate on the basis of weather reports that are hours old. Mechanics can't diagnose problems and implement solutions until the airplane is parked at the gate. Airline operations centers don't know whether an aircraft might have a mechanical problem or lack a crew for the next flight until it's time for that next flight. And passengers on delayed or rerouted flights have to scramble to locate new gates or rush to make connections.

Clearly, there are major gains to be made with improved data collection, analysis, and sharing. Because of this opportunity, the Boeing Company recently unveiled its e-Enabled Advantage, an effort to tie the entire air transport system into a seamless IS infrastructure that shares applications and data. As discussed in earlier chapters, e-Enabling creates a common onboard information and communication infrastructure for the benefit of passengers, flight and cabin crews, airline operations, system performance, and the overall industry. Indeed, Boeing's newest airliner, the 787, was code-named 7E718, and "e-Enabled" is one of the concepts addressed in the "E." But customers don't have to wait for the 787 to take advantage of the e-Enabled air transport system. Key infrastructure components of the e-Enabled Advantage, such as the Jeppesen Electronic Flight Bag or Airplane Health Management, are in production today.

In order to make e-Enablement possible, many different data needs have to be satisfied (see Figure on page 173). On the flight deck, for example, the Jeppesen Electronic Flight Bag (EFB) gives flight crews a sharper strategic picture of where they are, where they are going, and what waits over the horizon. The EFB offers the most up-to-date navigational information, live weather reports, instant access to flight

The Boeing 787 Dreamliner.

Source: http://www.boeing.com/companyoffices/gallery/images/commercial/787/k63304-2.html.


Data needs of operating a modern aircraft.

and aircraft data, airport surface positional awareness, cabin–to–flight deck surveillance, and more. In the airline operations center, various systems and data can be integrated to give managers and planners advance knowledge of possible schedule disruptions and a wide range of options to mitigate any problems. For example, each airplane contains an onboard data server that can be coupled with Boeing's Communication Navigation Surveillance/Air Traffic Management applications, various simulation and analysis products, crew management applications, and Boeing's Integrated Airline Operations Center to produce forward-looking plans to maximize operational performance. Likewise, various systems provide flight attendants with access to detailed information on their customers' needs, helping to give passengers a more enjoyable flight. In addition, the airline cabin can be connected to the airline's credit card verification systems on the ground, freeing flight attendants from having to carry thick wads of cash for beverage service, duty-free shopping, or other transactions. Catering and duty-free inventories will be updated automatically, ensuring that airlines can keep control of their stocks, improve oversight, and make sure every flight has the cabin items it needs.

Airlines such as Hong Kong's Cathay Pacific and Etihad Airways of the United Arab Emirates use Boeing's Airplane Health Management (AHM) program to monitor airframe systems information across the entire fleet in real time using high-bandwidth information flows. With AHM, engineers and maintenance personnel are able to examine system behavior while the airplane is in flight, quickly determine whether repairs can be deferred until the next scheduled maintenance, and inform the airline operations center whether necessary repairs can be completed at the destination without disrupting the flight schedule. In fact, AHM helps determine if a problem is developing over time and allows the airline to fix critical systems before they break, avoiding costly delays and potential catastrophes.

The e-Enabled air transport system is Boeing's vision of a day when the airplane is just another node on an enterprise-wide information network, ensuring that everyone in the system has all the information they need to react at the very best moment. The variety of different interrelated components of the e-Enabled Advantage require a solid underlying communication infrastructure to link the airplane to the airline enterprise. However, connecting a moving object to a communications infrastructure has never been an easy feat, and especially since Boeing has decided to discontinue the Connexion project (see Chapter 1), other ways have to be found to connect the airplanes to the ground systems. One way to enable this communication is through the use of ground-based networks, such as existing cell phone networks, or dedicated ground-based systems. However, for communications taking place while traveling across the oceans, this solution is not feasible. Thus, satellite-based systems have to be used. For example, telecommunications provider OnAir uses capacity on Inmarsat satellites to establish connections between the airplanes and their ground networks. As most aircraft are already equipped with Inmarsat communication systems, OnAir's systems can piggyback on these existing infrastructures.

## Questions

1. What other infrastructure components could be important for e-Enabling an airline? Which infrastructure components would be most important?
2. Identify, from your own experience, ways that e-Enabling will improve your flight experience.
3. Identify several ways that an airline could market its adoption of the e-Enabled Advantage to current and potential customers.

Sources: http://www.boeing.com/news/frontiers/archive/2003/august/i_ca1.html

Boeing's Airplane Health Management to Monitor Cathay Pacific's 777 and 747 Fleets (2006), http://www.boeing.com/news/releases/2006/q2/060427b_nr.html