Chapter 3

Networks and Telecommunications

Chapter Outline

What You Will Learn in This Chapter

- What is the value of a single computer?
- Why are computer networks so important in today's businesses?
- What components do you need to install to create a network?
- How can multiple users share a single network?
- How is it possible that you can connect your computer to a network at the office, at home, or while on the road, even overseas?
- What is the Internet, how is it controlled, and how does it work?
- Are personal computers necessary anymore?
- What problems are you likely to encounter if you need to connect to a supplier in a different country?

Owens & Minor

How do suppliers connect to customers? Owens & Minor sells thousands of products to hospitals. It is one of three large companies in the market, competing with Cardinal Health (which owns the system built by American Hospital Supply). Like other wholesalers, Owens & Minor exists in a highly competitive environment with low margins. Whether you call it logistics or supply chain management, buyers continually search for ways to reduce their purchasing costs. Companies like Owens & Minor use information technology to reduce inventories throughout the purchasing system. By building links with customers and suppliers, Owens & Minor can immediately react to changes in demand and plan the most efficient method to deliver supplies.

Introduction

What is the value of a single computer? This question almost sounds like the old Buddhist Zen koan: What is the sound of one hand clapping? You do not need to debate the complexities of Zen to understand that the value of computers is multiplied by connecting them and sharing information. Today, many applications exist only on the network—they are built to use servers to store and process data and they are accessed across the Web from any location in the world. And communication ranging from text, to voice calls, to video conferencing has largely moved to Internet-based formats.

As shown in Figure 3.1, networks are useful on two levels. Internal company networks are used for teamwork and sharing data; the Internet connects you to external organizations, including suppliers, customers, and service providers. Actually, you use many other networks in your daily life, including cell phones, television, and even electricity. Today, almost everything you do in business relies on computer data networks. You use them for communication, sharing data, team work and collaboration, and simpler things such as backing up your data and printing.

Even though you will probably not be responsible for creating or maintaining a network, knowledge of the underlying technologies makes it easier to understand some of the problems that can arise, and helps you see how the technology is likely to change in the future so you can be prepared. Network technologies have experienced some of the most rapid changes in the past few years, and new technologies are being deployed now. You need to plan ahead so you can create new organizations and new methods of handling work.

The Internet has changed many aspects of daily life, and has the capability of fostering even more changes. To understand these opportunities and recognize the challenges that need to be overcome, you need to learn a little about how the Internet is designed and how it works. The Internet is an interesting organization because it is governed purely by committees and individuals and works because of standards. These organizations face difficult questions and the answers will dramatically influence your life and the role of businesses in the years to come. If nothing else, as an informed citizen, you need to be aware of these discussions. Many times the decisions hinge on social, economic, and political issues instead of purely technical ideas.

Trends

In many ways, the communication systems available today began with the telephone. Originally, sounds were converted into analog electrical signals, and calls were connected through physical switches. As computer technology evolved, it became possible to convert everything into digital data and the computerized switches simply transfer digital packets to a desired destination.

A key aspect of communication systems is that some connection medium is required to transmit data. Most homes have two major connections: telephone and TV cable. The phone and cable companies have been expanding their services by providing better and faster connections to houses. Thanks to the long-distance phone companies, hundreds of thousands of miles of national and international fiber optic cable connect cities and nations around the world. These backbone connections carry almost all of the Internet data, voice, and video connections.

But, the biggest trend in the past few years has been the switch to cellular phones including smart phones used as computers to access the Internet. People are increasingly giving up wired connections in favor of mobile connections that operate over radio waves. As nations move to provide more bandwidth for cellular service, data transmission capabilities for mobile phones is approaching the speeds available for fixed-line connections.

As connections and usage have grown, so have the applications and data available. Consequently, it is increasingly important to be able to connect to the Internet from anywhere at any time. Which leads to an increasing cycle: As connection capabilities improve more applications and data become available, which means connection capabilities need to improve even further, which leads to new applications, and so on.

The global economy makes life more interesting—particularly in terms of networks. Sharing data, or even using cell phones, across international borders raises many complications. Because technology continually changes, nations implement the technology at different points in time. Compatibility across borders is an even greater challenge when nations have differing goals and income levels. As managers, you need to be aware of some of the issues, both to communicate and to travel.

The objective of a network is to connect computers transparently, so that the users do not know the network exists. The network provides access to data on central computers, departmental computers, and sometimes other workers' personal computers. Networks can create shared access to fax machines, modems, printers, scanners, and other specialized hardware. To the user, it does not matter where these devices are located. The network makes them work the same as if they were on any desktop. The Internet expands these capabilities across the world. Wireless makes the services available to you wherever you travel in major cities.

Network Functions

Why are computer networks so important in today's businesses? Most companies did not seriously begin installing networks until the early 1990s. The Internet expanded into the commercial world in the mid to late 1990s.

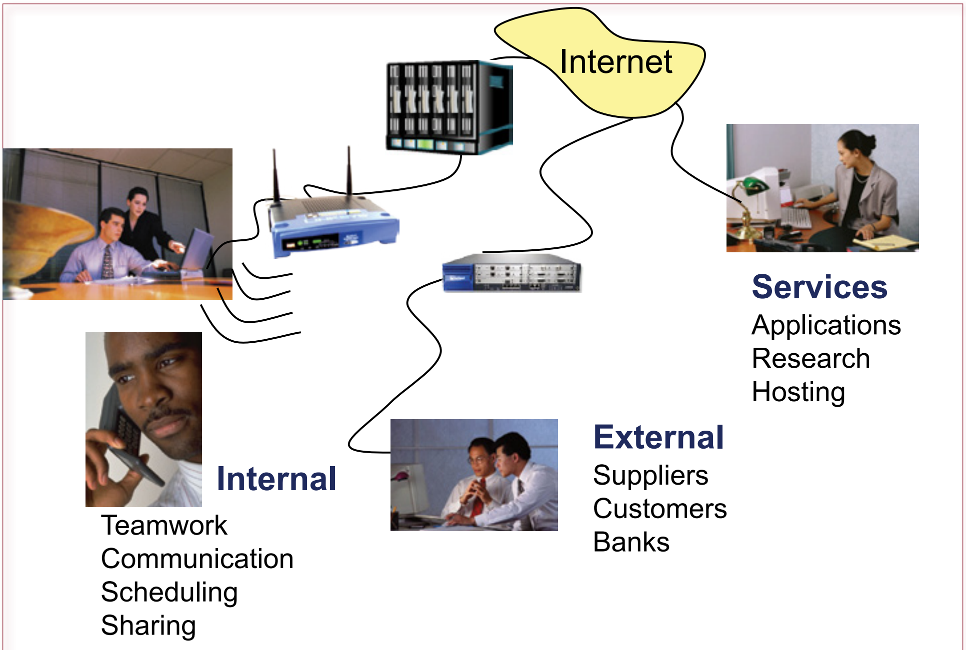

Figure 3.1

Networks. Managing an organization requires communication and teamwork. Networks facilitate teamwork and create opportunities for new organizational structures. The Internet makes it easy to connect to customers and suppliers and access new services and Web-based applications.

Perhaps the better question is: Why did it take so long? The answer: cost. Think about that time frame for a minute. The digital world as you know it is only about 15 years old, and it has already had a huge impact on business and society. But, the world is still only at the beginning of the network evolution. New opportunities in both communications and applications will continue to arise. And you will face new questions about choosing technologies and applications. To understand the value of networks, you need to look at the business applications and new procedures that have been created with networks. Then, you can start thinking about the future networks and new applications—particularly wireless and mobile applications.

Sharing Data

Sharing data is one of the most obvious uses of networks and it can make profound changes in the way an organization works. Managers can see customer and marketing data immediately as it is collected. Employees in one department can easily share data with other departments. A network facilitates the use of teams. In particular, it enables informal teams to spring up throughout the company to solve problems as they arise. Instead of waiting for a higher-level manager to appoint a team, employees can use the network to ask questions, notify others involved, and find in-house experts. A **local area network (LAN)** is commonly used to connect computers and share data within a company. A LAN uses the same technologies as any other type of network, but is recognized separately because all of the compo-

Figure 3.2

Network for transaction processing. Networks are often used to collect data in a central database. From there, the data can be queried and analyzed by managers. E-commerce sales represent transactions across the Internet.

nents of the network are owned by your company or organization (or household). Since you own everything, you get to make all of the decisions, and you do not have to pay anyone else to transmit the data.

Transactions

One of the most important reasons for connecting computers is the ability to share data. **Point of sale (POS)** systems were some of the first networks installed in businesses.

Consider a retail store with five checkout registers. Each register is actually a computer. If these computers are not connected, it is difficult to compute the daily sales for the store. At the end of the day, someone would have to manually collect the data from each computer and enter it into another computer. Also, imagine what would happen if a customer asked a clerk to determine whether a product was sold out. The clerk would have to check with each of the other clerks or perhaps call the storeroom.

As shown in Figure 3.2, with e-commerce transactions can take place across a wider network, with the Internet as the network and customer browsers as the client computers. The product data and sales transactions are stored in the central database connected to the Internet. Using a central database provides inventory data to customers. When a customer asks whether an item is in stock, the Web site can provide the answer. Managers can get daily sales figures from any location with a Web browser and an Internet connection. Payments and bills can also be handled directly online. The same diagram can be used for traditional in-house

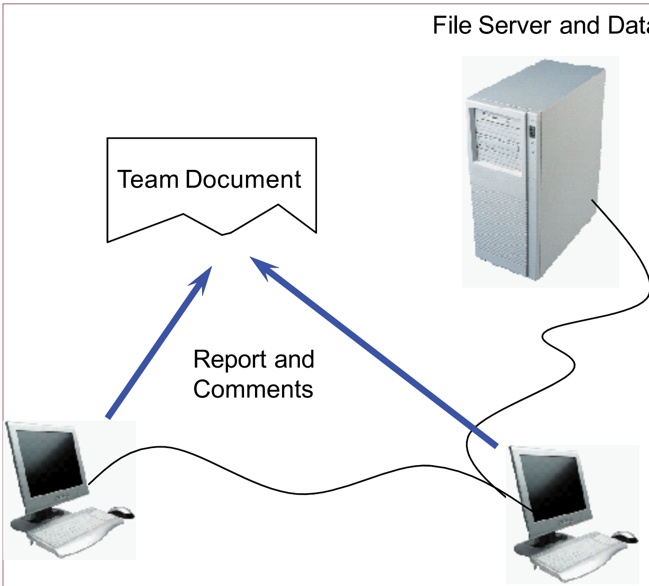

Figure 3.3

Network for decisions and collaboration. The file server holds basic data and software tools. Managers retrieve data and create reports. The reports can be shared with other managers. With collaborative software, revisions are automatically tracked and combined to form the final document.

POS systems by replacing the Internet with a simple LAN. The main point is that a central server holds all of the data and handles the actual transactions.

Decisions and Collaboration

Many types of data need to be shared in a company. But beyond simple data sharing, people often work on teams. Teamwork entails developing reports together. It requires searching for data, editing documents, and sharing comments. A portion of a network for making decisions and sharing work with team members is illustrated in Figure 3.3. Centralized data and shared documents ensure that everyone has the same version of the data and the same final documents.

In any job, it is rare for one person to work alone. Most businesses are arranged as teams. Within the teams, individual people are given specific assignments, and each team member contributes to the final product. For instance, the marketing department might have to prepare a sales forecast for the next six months. Each person could work on a specific sales region or product category. These individual pieces would then be merged into a single document. If the computers are networked, the manager's computer can pull the individual pieces from each of the individual computers. Also, each team member can be given access to the others' work to ensure that the reports use the same data and assumptions.

Groupware includes software that enables several people to work on the same document. Each individual computer has access to the master document. When one person makes a change to the document, the change is highlighted for everyone to read and approve. With existing international networks, each person might be located in a different country. E-mail systems can work as simple groupware

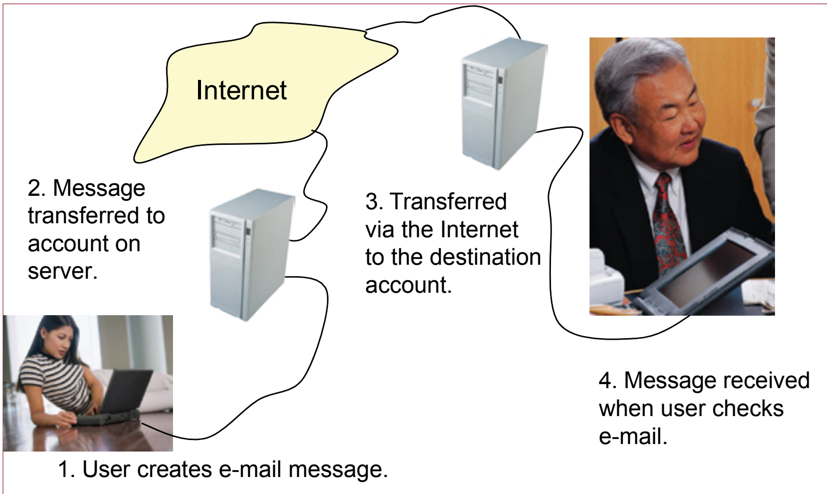

Figure 3.4

Networks for communication. With e-mail, a server holds the mail until the recipient logs on and can receive the message. Because several servers might be involved and each makes backup copies, e-mail messages are hard to delete.

tools by routing copies of files to everyone on the team. More sophisticated tools are included in Web sites built by systems like Microsoft's SharePoint. These tools are discussed in greater detail in Chapter 8.

Messages

Most people are familiar with electronic mail, or **e-mail**. With e-mail, you can send a message to any person with a mail account. Many people have come to prefer e-mail contacts to phone calls. As shown in Figure 3.4, e-mail messages are asynchronous since the sender and recipient do not have to participate at the same time. A mail server holds the message until the user logs in and retrieves the e-mail. Moreover, users can create mailing lists and send a message to many people at one time. Messages can be stored, retrieved, and searched for later usage. In most systems, the computer can automatically notify you when the recipient reads your message. You never have to worry about whether your message was received. Voice mail systems, which resemble answering machines, have some of these same advantages. However, e-mail takes less space to store. More important, e-mail can be scanned electronically to search for key topics.

E-mail has some drawbacks. For one, some people do not check the computer often enough to keep up with their mail. Another problem is that in 1991, the U.S. courts ruled that public transmission systems such as e-mail are not subject to the same legal protections for privacy as paper mail and phone calls. Unless the laws are rewritten, therefore, it is legal for employers (and almost anyone else) to read everyone's e-mail messages. Of course, the fact that it is legal does not make it ethical.

Examine the figure for a few seconds and another aspect of e-mail becomes fairly obvious. How can you delete e-mail (once it has been sent)? The answer is that it probably cannot be done. Even though some systems appear to give you

Reality Bytes: Worldwide Broadband Rankings

Akamai is an Internet company that provides services to many other companies. It specializes in distributing data around the world—putting it on servers closer to where it will be needed. This positioning gives the firm the ability to measure data flows and speeds around the world. An interesting statistic is the connection speeds of the people in each country. In 2010, the nations with the highest average connection speeds and their average speed in Mbps, were: 1. South Korea (17), 2. Hong Kong (8.6), 3. Japan (8.0), 4. Romania (6.8), 5. The Netherlands (6.5). Note that In South Korea, over 75 percent of the population connected to the Internet at speeds over 5 Mbps. The U.S. ranked 16 on the list of average speeds at 4.6 Mbps, and only 30 percent of the population had access to speeds over 5 Mbps. It is easier to build high-speed networks in nations with small physical space and large populations such as South Korea and Japan. Similarly, population areas in the U.S. such as the Northeast have more people with faster connections. But, ultimately, the speed available to the average person begins to matter—as companies rollout Web-based applications with sound and video.

Adapted from Karl Bode, Akamai: U.S. 33rd Fastest Broadband Country, www.dslreports.com, July 9, 2009. Updated data: http://www.akamai.com/stateoftheinternet

the option to delete or retract e-mail that has been sent, it generally will not work. Copies of the e-mail can exist in several locations—the original computer, the sender's e-mail server, the recipient's server, and the recipient's computer.

E-mail, and the fact that it is difficult to delete, has become an important aspect in legal cases—both private and those involving government officials. In court cases, many lawyers automatically ask for copies of all relevant e-mail in discovery motions. Consequently, many companies build **e-discovery** databases that contain copies of all e-mail messages sent through their servers. These systems have keyword search facilities so administrators can retrieve specific messages when required. The systems also have rules-based systems to handle retention policies. For example, company policies might specify that routine messages should be deleted after six months. The system can enforce retention and deletion policies automatically. Many of these systems also integrate voice mail messages with e-mail messages, so they are subject to the same retention, deletion, and discovery rules. The main point to remember is that any digital data might be stored for an unlimited time. So, although it is easy to think of electronic text and messages as ephemeral and fleeting comments, you should always remember that the potential exists for any message to be retrieved later.

You need to pay attention to two key aspects to any communication method: (1) Whether it is synchronous or asynchronous, and (2) The level of interruption it creates and the preferences of the people you work with. For example, a personal visit and phone call are synchronous, and the other person has to stop working on a task to talk with you. IM or texting is also synchronous, but users often feel more comfortable delaying these communications. E-mail is asynchronous, so each person works on his or her own schedule, and deals with mail during breaks. As the message sender, you have to choose the appropriate method of communication for each message and person. You also have to learn to allocate your own communication time, to leave sufficient time to finish your own work.

Web Site Basics

Web sites have evolved considerably, with several well-known sites providing services and applications to address many new tasks. However, many sites still exist in the early format—providing information and data. As a step up, some sites provide limited interaction, such as sales, feedback, discussion, and file sharing. At a higher level, a few sites provide greater interactivity—including games and online applications. The first decade of 2000 also saw the expansion of **social networking** sites—such as Facebook and Twitter. Social networking sites facilitate communications among groups of users. These types of sites and other capabilities are covered in more detail in Chapter 7 on e-business.

A key element of communication that is emphasized with the Internet and the Web is that communication only exists through the adoption of standards. **Standards** are simply agreements among the major users to handle tasks and data in a specific way. In terms of the Internet, standards exist for connecting networks and transmitting packets through routers. Common Web standards include **hypertext markup language (HTML)** which is the standard format for telling browsers how to display a page of text and images. These same standards can be used for other types of networks, including intranets and extranets. **Intranet** sites employ Internet technologies but use security methods to restrict access to internal users. For example, employees probably use an intranet to retrieve and update personal data in the human resources database. Similarly, **extranets** are sites specifically designed for companies that partner with your organization. For example, your suppliers can log into a special section to check on your production schedule, pick up technical specifications, or bid on jobs. Intranets and extranets use the same technologies as the common Internet, but are limited to special groups of users.

Several common Web tools can be used to support communication. Initial Web sites were designed primarily for one-way communication. They are still used for posting centralized data that is needed by many other people. However, interactive tools have also evolved to support feedback and communication within groups. Tools such as **instant messaging (IM)** can be used to send short messages to friends or coworkers. The messages are delivered instantly and can provide a way to ask questions or share ideas with team members who are online. Of course, cell phone texting services operate the same way, but the interaction requires a cell phone, or preferably a smart phone with a decent keyboard. The computerized IM tools can also be used to see when someone is online. If you are working late a night, you can quickly see which of the other team members are also online, so you know that you can ask a question without having to call everyone just to see if they are available.

A Web log or **blog** is an interesting variation of a Web site. The tools were created to make it easier to create and post content online and facilitate comments and feedback. Several sites provide public blogs enabling you to write comments on any topic. The software is also readily available for installation on internal company servers so managers can create blogs on topics that can be targeted to employees or to external partners. For instance, a blog could be used by managers to keep a log of the daily issues, problems, and solutions. Other managers in the company could skim the logs to see if problems arising in one area might cause problems in their own sections. Similarly, if a manager encounters a problem, he or she could search the blogs for similar problems, ideas on how to approach it, or even solutions that worked in the past. In essence, the blogs become a knowledge base that is accessible to other managers. Of course, the sites would be secured

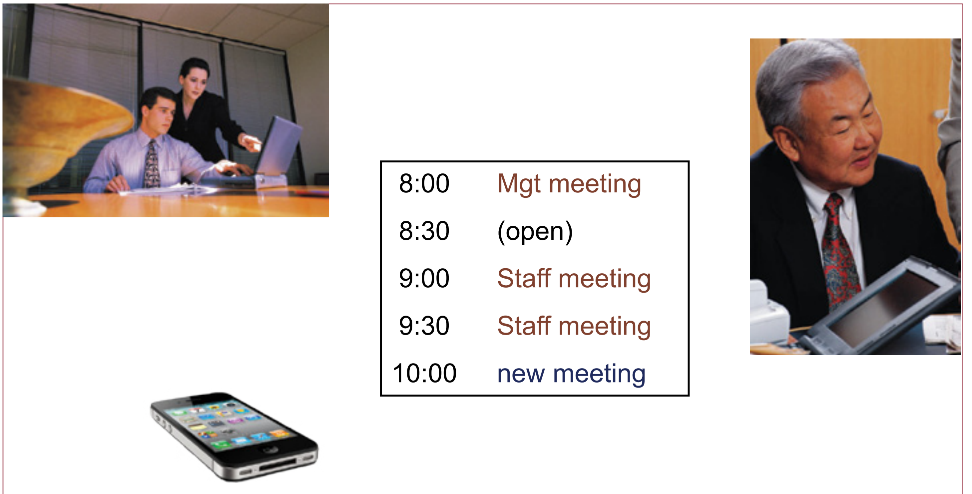

Figure 3.5

Sharing calendars. With the appropriate software, you can open your calendar to other members of your team. A team member can check the calendar and have the software find a common open time for a meeting.

so that only managers within the organization could read them. Some people are so dependent on blogs that they want to be notified whenever the information changes. **Really simple syndication (RSS)** feeds can be established so that your browser or reader automatically checks for new information and displays it automatically. It is even possible to connect your reader to your desktop so that any of your subscriptions are automatically displayed on the desktop when information changes. However, this approach is likely to be a major distraction to your work.

With the introduction of the Apple iPod, the concepts of blogs and RSS feeds was extended into audio notes that can be downloaded to the iPod and played later. These **podcasts** can be useful if you have time to listen to voice notes, but in business, it is faster to read and search written notes.

Wikipedia, the online user-supported encyclopedia has led to a new form of Web-based communication. You can use the same software to create in-house **wikis** where workers on a team can write their own descriptions, edit work by other people, and add comments. The entire system contains a search engine, so other employees can search the system for problem descriptions, work documents, and approaches or solutions used in prior cases. It is a powerful and inexpensive system when you want to record case-based data.

Calendars and Scheduling

Managers spend a great deal of time in meetings. Yet sometimes the greatest challenge with meetings is finding a time when everyone can get together. Several software packages use computer networks to solve this problem. As shown in Figure 3.5, managers enter planned meeting times and scheduled events into their personal electronic calendar file, where each event is assigned a priority number. For example, a time allotted for a haircut would be given a low priority; a meeting with a supervisor would receive a higher rating. If the CEO wants to set up a meeting, the CEO tells the computer which people are to be included, sets a pri-

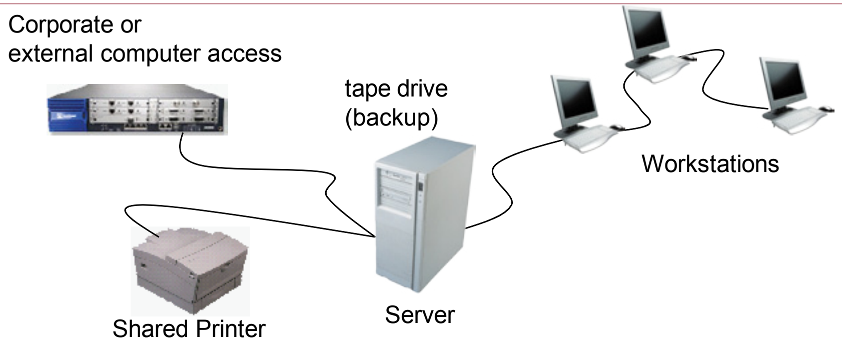

Figure 3.6

Networks for sharing hardware. The workstations use the server to perform backups. Files are picked up by the server and transferred to tape. The LAN administrator can reload a tape and restore files as needed. Networks are often used to share printers and storage devices. Networks can be used to share access to supercomputers—even if they are in a different city or different country.

ority level, and gives an approximate time. The computer then uses the network to check everyone else's schedule. When it finds an open time (overriding lower priority events if needed), it enters the meeting into each person's calendar. These systems can be useful when the calendar on your cell phone is connected to the shared calendar.

Sharing Hardware

Networks also enable workers to share expensive hardware. For example, networks are used to provide users access to special output devices, such as high-speed printers, plotters, or color printers. Networks can be used to give people access to special computers, such as when an engineer needs to use a high-speed supercomputer. Figure 3.6 shows some of the hardware devices commonly shared using a LAN: printers, tape drive backups, server processing, and external access to the Internet.

Printers

A common use of networks is to give users access to high-speed, high-quality printers. Even if each computer has a small personal printer attached, some jobs need to be handled by a high-speed laser printer. Many modern copiers now function as network printers and can collate and staple large quantities of documents. Similarly, even at $1,000 each it would be expensive to buy color laser printers for everyone who might need one, yet it might be reasonable to buy one for a department to share. With a network, users can choose from among two or three different printers by selecting them on the screen.

Another advantage is that if one printer breaks down, users can send their jobs to another printer on the network. Think about what happens if there is no network and your printer breaks down. You have to copy the file to a USB drive and interrupt someone else's work to borrow their computer to send the file to another printer. What happens if you are using a special software package that no one else has on his or her computer? You will probably have to physically move a printer from another computer desk to yours, connect the hardware, print your document,

Reality Bytes: Network Switches for Servers hit 10Gbps

As the amount of data on Internet and corporate servers continues to increase, network vendors are pushing new technologies to improve the transfer of vast amounts of data. In 2010, 10 Gigabit Ethernet is pushing into data center network design. Most data centers are built as vertical racks of servers, where each rack has a network switch at the top. In the past, these top-rack switches were connected to a middle-layer switch that then connected to the primary Internet backbones. But, hardware servers are increasingly used to host multiple virtual servers—and each virtual server pushes data across the network. Replacing the top-rack switches with 10 Gbps speeds and eliminating the middle layer speeds the data to the Internet connection. Robin Layland, an advisor to companies and vendors, notes that "Over the next few years, the old switching equipment needs to be replaced with faster and more flexible switches. This time, speed needs to be coupled with lower latency, abandoning spanning tree and support for the new storage protocols." IEEE standards are also being ratified for new switches to handle 40G and 100G Ethernet in the future. The speed increases are important in a cloud-based world where Layland notes that applications "require less than 10 microseconds of latency to function properly," compared to existing 50-100 microseconds. Jayshree Ullal, CEO of Arista Networks, a company that makes some of the new high-speed switches, notes that "most of the [existing] switches are 10:1, 50:1oversubscribed." Ultimately, companies save money by reducing the number of network switches—both in the cost of the equipment, maintenance, power, and air conditioning.

Adapted from Jim Duffy, "10G Ethernet Shakes Net Design to the Core," *Networkworld*, September 14, 2009.

and return the printer. When you are on a network, you simply select a different printer from a list displayed on your computer and go pick up the output.

Several e-commerce printing companies enable you to send your large print jobs over the Internet and have the boxes of papers shipped to you. All of the setup, pricing, and payment can be handled over the Internet.

Storage Devices

The arguments used for network printer sharing can be applied to sharing storage devices. If you have huge data files that you want to share across the organization, it is best to put them in a central location and provide network access to everyone who needs them. The central location makes it easy to upgrade the drives, provide sufficient capacity, and control the access rights and monitor security. A specialized **storage area network (SAN)** is often used to provide vast amounts of flexible storage. The disk drives are separated from the computers and connected by a high-speed network using fibre channel or high-speed LAN connections. Physically separating the secondary storage from the computer box makes it easy to expand the capacity, provide redundancy, and move the drives to safer locations. Because of the high transfer speeds across the SAN, the drives appear as simple local devices to the computer, so no software changes are needed.

Backup

Another important reason for sharing data over computer networks is that most people are not very good at maintaining backup copies of their work—especially on personal computers. If each computer is attached to a network, automatic back up can be configured two ways for individual personal computers. Individual workers can save their data files to a central file server. The network manager then makes daily (or hourly) backups of the data on the central server. A few companies even provide this service over the Internet. For a monthly fee, you can transfer your files to their server, giving you a backup copy—plus they keep backup tapes for the server. The second approach is to have backup software running on the server query each user's computer to obtain changed files and save them automatically. This approach requires that the client computers remain connected to the server and remain running, so it does not work well with laptops.

Special Processors and Grid Computing

Special computers that are relatively expensive can be attached to a network to save costs and make it easier for users to access these machines. Parallel-processing computers and other supercomputers can perform calculations thousands of times faster than ordinary computers, but they are expensive. Consider a small engineering company. For the most part, the engineers use their workstations to develop their designs. They occasionally need to run simulations or produce detailed graphics images. Both of these tasks could take many hours to perform on individual computers. With a network, each workstation can be connected to a supercomputer. When an engineer needs to perform high-speed calculations, the job is sent directly to the supercomputer. The results are returned via the network to the workstation so they can be added to the final report. More likely, a university could own the supercomputer, and the firm would lease time to run each job. If the network is designed properly, it makes no difference where the machine is located.

Grid computing extends the concept of parallel processing—instead of having multiple processors in one computer, you simply attach multiple computers together across a network. Remember that personal computers are relatively fast and inexpensive. When you can buy a new computer for $500, it becomes possible to build a huge amount of computing power by purchasing hundreds or thousands of computers and spreading the work across all of them. With special software, the job is split into multiple pieces and assigned to spare time on each computer. SETI @home is probably the best-known example of this type of computing. However, the same tools can be used within a single company—taking advantage of machines running overnight. Just keep in mind that running computers 24-hours a day substantially increases the power consumption.

Sharing Software

Networks have been used at different times to share software. When disk space was expensive, it was cheaper to put one copy of the software on a server and download it to each computer as it was needed. Today, it is possible to run software applications directly from a server, where client computers use Web browsers to connect to the server and run the software. This centralization of software and data leads to new ways to work and to manage companies. However, it raises issues of reliability and security. Some of these are covered in more detail at the end of this chapter. But the overall topics of management, centralization, and security arise in several other chapters.

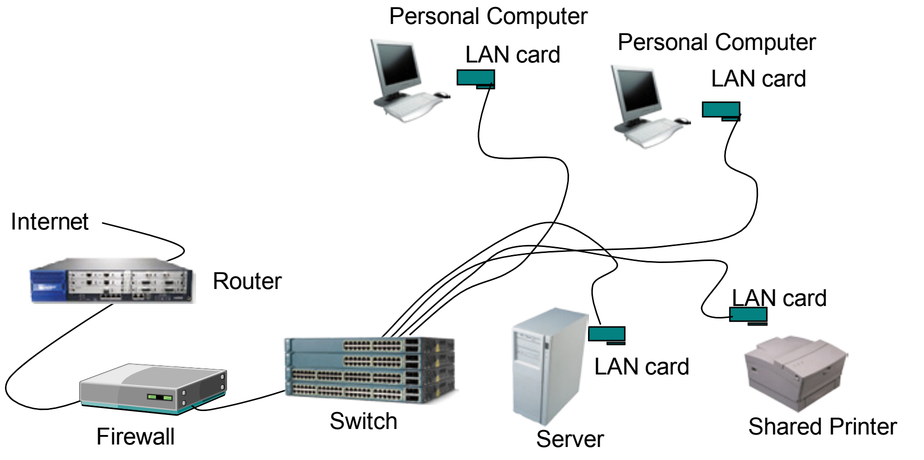

Figure 3.7

Network components. Networks require a transmission medium and each device connected must have a network card to make the connection. A switch connects the devices in the LAN, and a router connects the LAN to the Internet. A firewall is often used to filter out certain types of packets or messages so they do not transfer from the Internet to the LAN.

Components of a Network

What components do you need to install to create a network? As networks have become more important, the connection components are increasingly built into the computers. However, you still face many decisions about which technologies to use and how to solve problems, so all managers should understand the basic elements of a network shown in Figure 3.7. Each of these components (computers, transmission medium, and connection devices) is discussed in greater detail in this section.

Computers

Virtually any type of computer device can be connected to a network. The earliest computer networks consisted of one computer with several terminals and printers attached to it. These networks were fairly simple, because the one computer performed all of the work. Substantially more problems are involved in connecting several computers together. For starters, each computer needs to know that the others exist and that they have a specific location (address) on the network. The computers need to know how to send and receive information from each other. Just to connect to the network, they need LAN cards. Most desktop computers have built-in LAN cards and connectors for wired networks. Most laptops also have wireless network cards; some even have cards for cellular networks. For now, most companies run a combination of both wired and wireless networks. Most homes rely solely on wireless networks because no one wants to run cables through the house.

Computers attached to networks tend to perform one of two functions: servers or clients. Servers are computers that store data to be used by other computers attached to the network. Clients are computers used by individual people.

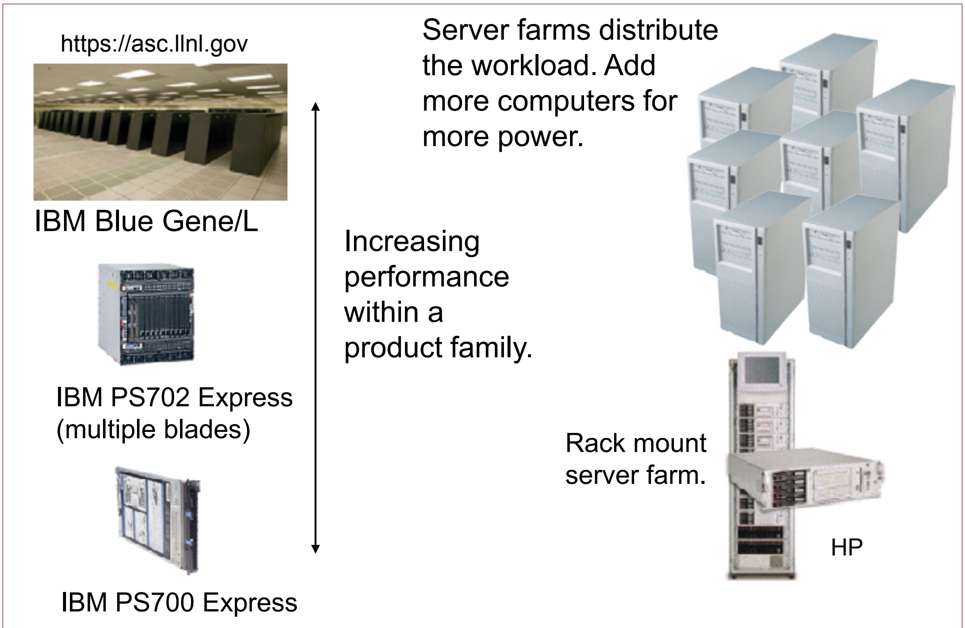

Figure 3.8

Server scalability. Two common methods of providing easy expansion of performance are (1) to purchase a faster performance server within the same product family, and (2) to build a server farm where the workload is automatically distributed across machines and new low-cost servers can be added at any time.

Sometimes a computer is used as both a client and a file server. Networks where many of the machines operate as both clients and servers are called **peer-to-peer networks.**

Servers

A wide range of servers exists today—from a simple PC to huge, expensive specialized computers. The main questions you face in choosing a server are the operating system and the issue of scalability, and the two questions are intertwined. **Scalability** is the ability to easily move up to greater performance—without having to rewrite all of your existing applications. Figure 3.8 shows two common methods used to provide scalability: (1) a vendor-provided range of servers from low-cost machines to handle small loads, to midrange, to high-capacity computers that can handle millions of users simultaneously; and (2) integration technology that enables the workload to be distributed across hundreds or thousands of small servers, known as a **server farm**.

Both approaches have their benefits. The single server is easier to configure and administer. The server farm can be expanded easily and cheaply without disrupting the existing applications. The operating system software is crucial to making a server farm work efficiently. Several vendors sell enterprise versions of software that assigns applications to the least-busy server and makes it easier to manage the server farm.

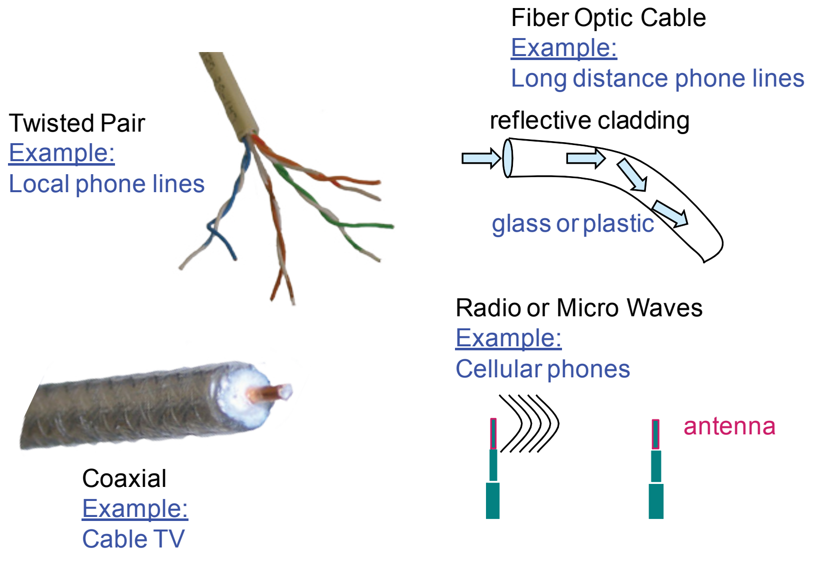

Figure 3.9

Signals can be sent via several types of media. They can be carried by electricity, light waves, or radio waves. All of these methods have advantages and disadvantages. Fiber optic cabling offers the fastest transmission rates with the least interference, but because it is relatively new, the initial cost tends to be higher.

Client Computers

The networked computers could be any type of machine. Because individual people at their own desks typically use these computers, they are often called client computer*s*. These computers need to access the network and be able to send information to at least one other computer. A **network interface card (NIC)** (or LAN card) is installed in each computer. These cards are connected together by some transmission medium (e.g., cable). In addition, these computers might have to be configured to connect to the network and set security parameters. Client computers today include laptops and PDAs that can be connected with wireless networks to enable workers to maintain connections and share data while they move around the building.

Many of the mobile devices have limited capabilities and essentially run browsers, e-mail, and calendars. Hence, more of the processing is done on the server, and the mobile device is only responsible for displaying data and basic user interface tasks. This type of environment is considerably different from the situation where all workers have desktop machines capable of handling large amounts of the processing. The thin-client cell-phone applications rely on the server, while the desktop applications simply use the server for basic data sharing. Developing applications for this new environment requires some major changes in design, programming, and security—issues that will be explained in detail in other chapters.

Transmission Media

All communication requires a **transmission medium**. As illustrated in Figure 3.9, common methods include electric wires, light waves, and radio waves. Each

Reality Bytes: Excess Capacity

Partly because of crazy predictions of Internet growth, partly because of economies of scale, and partly because of improved technologies, several long-haul data-transmission companies overbuilt when installing fiber-optic cables. In late 2002, only 2.7 percent of the fiber was being used—a whopping 97.3 percent was dark. Over time, economics prevailed and bandwidth prices dropped—at a rate of 65 percent per year. Many of the large telecommunications firms filed for bankruptcy. The bankruptcy filings generally erased the debts of the companies—enabling them to charge even lower prices when returning. Prior to 1995, telecommunications companies could send the equivalent of 25,000 one-page e-mails per second over one fiber-optic line. By 2002, advances in technology enabled them to send 25 million messages over the same line—an increase of one thousand times. The new technology called dense-wave division multiplexing (DWDM), effectively splits light into multiple colors—enabling the companies to send different messages on each color frequency. By 2011, demand had started to cut into some of the excess capacity, but not enough to begin installing significant new cables.

Adapted from Dreazen, Yochi, J., "Wildly Optimistic Data Drove Telecoms to Build Fiber Glut," *The Wall Street Journal*, September 26, 2002; and Berman, Dennis K., "Technology Races Far Ahead Of Demand and the Workplace," *The Wall Street Journal*, September 26, 2002.

method offers certain advantages and disadvantages, so they are designed for specific applications. Installation costs for all types of cables are high, and most organizations keep cabling in place at least 10 years. Consequently, you have to be careful to research the technologies when installing new cables to ensure you will be able to support future needs.

Electric Cables

The two basic types of electric cables are twisted pair and coaxial. Twisted pair is the oldest form of electrical wiring. Since electricity must travel in a closed loop, electrical connections require at least two wires. **Twisted-pair** wires are simply pairs of plain copper wires. Telephone cables are the most common example of twisted-pair wires in households. Because of the cost, most businesses have already installed twisted-pair wires—typically a specific version known as Cat 5 or Cat 6. These standards define four pairs of wires. Older, slower networks (100Base-T Ethernet) uses only two of the pairs (4 wires), but higher-speed networks (gigabit Ethernet) require all eight wires. Because of the extensive use of Cat 5 cables, the network industry has invested considerable research into maximizing the data that can be carried over that type of cable. However, newer high-speed transmissions systems might require better cables. The Cat 5e (enhanced) standard would be the minimum for installing new cables, and depending on price and location, you might consider Cat 6 or Cat 7, which is heavily shielded. But check the prices. Cat 7 cable can be five times more expensive than Cat 6

The biggest drawback to twisted-pair wires is that they are not usually shielded. Consequently, the data signal is subject to interference from other electrical appliances or cables. In particular, never run data cables next to power lines or electrical motors. Additionally, the data cable length for most LANs must stay under 90

Reality Bytes: LightSquared v. GPS

LightSquared is a startup company with a new technology to create a national wireless broadband system. The plan would entail installing 40,000 antennas around the country. In early 2011, the company signed agreements with Best Buy and Leap Wireless to enable them to sell wireless Internet access. Adding a new provider like LightSquared would ultimately provide more options for consumers and provide competition to ISPs such as the cable and phone companies. It would also provide competition to mobile phone carriers. However, the frequencies LightSquared wants to use are near the frequencies used for satellite signals—particularly low-powered GPS signals. In testing, the LightSquared devices were found to interfere with high-end GPS devices, particularly those used by aviation, police, and agricultural applications. For instance, John Deere, the tractor company, said the devices created severe interference as much as 20 miles away and "complete loss of service" between four and 22 miles. LightSquared's executive vice president for regulatory affairs and public policy, Jeff Carlisle, noted that "LightSquared and BPS can and will be able to coexist peacefully. We're committed to identify and resolving the issues through this process." Ultimately, the FCC will decide whether to move forward and what changes might have to be made. LightSquared later announced that it was moving to a different frequency farther away from the GPS signals.

Adapted from Amy Schatz, "LightSquared's Wireless Network Interferes With GPS," *The Wall Street Journal*, June 2, 2011.

meters (about 300 feet). You can extend the length by adding switches, but you will be amazed at how quickly you can run into the limit when a cable needs to turn corners and travel multiple floors. Despite these drawbacks, twisted-pair wiring is the most common LAN cabling method in use today. Largely because of the cost and the research that has led to huge performance capacities. Standard Cat-5e wiring is capable of carrying data at 1,000 megabits per second (gigabit).

Coaxial cables were designed to carry more information than twisted pairs, with lower chances of interference. **Coaxial cable** (often shortened to coax) consists of a central wire surrounded by a nonconductive plastic, which is surrounded by a second wire. The second wire is actually a metallic foil or mesh that is wrapped around the entire cable. This shielding minimizes interference from outside sources. Cable television signals are transmitted on coaxial cables. Coax is capable of carrying more information for longer distances than twisted pair. But, it is more expensive and is rarely used outside the television industry.

Fiber Optics

A relatively recent invention (early 1970s) in communication uses light instead of electricity. Because light generally travels in a straight line, it could be difficult to use for communication. Fiber-optic cable allows light to travel in straight lines but still be bent around corners. A fiber-optic cable consists of a glass or plastic core that is surrounded by a reflective material. A laser light is sent down the cable. When the cable bends, the light is reflected around the corner and continues down the cable. Fiber-optic cable provides the advantages of high capacity with almost no interference. The limitation in using fiber is the higher cost of the cable and the cost of the connectors and interface cards that convert computer electrical signals into light. For example, NICs for coaxial or twisted-pair cables can be purchased for around $10, whereas NICs for fiber-optic lines that run directly to personal computers range from $50 to $400 (in 2011). The cards are harder to find and more expensive, because it is rare to use fiber lines to the desktop. A study by Partha Mitra and Jason Stark showed that even fiber-optic cables have limits. The theoretical capacity of a single fiber-optic cable is at least 100 terabits per second (*The Economist* June 28, 2001 or *Nature* June 28, 2001). In 2011, two researchers achieved this speed in demonstrations reported by *New Scientist*. Although it is possible to run fiber optics directly to your desktop, it is rarely worth the expense. Fiber-optic connections are used for long distances, in areas with electrical interference, or when connecting buildings. In the mid-2000s, Verizon, the local telephone company (not the cellular phone company) began installing thousands of miles of fiber cables directly to customer homes. The FiOS service is designed to provide high-speed Internet connections (20 mbps or more), and ultimately video on demand and other data-intensive services. Installing the fiber into homes is challenging because although fiber cable does not have to be perfectly straight, if it is bent too sharply or too many times, the light will bleed out at the corners and the signal will disappear. Corning has an answer with a new fiber that contains nano-particles to allow fiber cables to be wrapped around objects as small as a pencil and still carry the light signal. However, it might take a couple of years to create a factory that can produce the cable in large quantities (Mehta 2007).

Radio, Micro, and Infrared Waves

Radio, microwave, and infrared transmissions do not require cables. These communication methods are called **broadcasts**. Any receiver or antenna that lies in its path can pick up the signal. However, infrared transmissions and some microwaves require clear line-of-sight transmission. The major advantage of broadcast methods is portability. For example, computers can be installed in delivery vehicles to receive information from corporate headquarters. Or individuals can carry around laptops and cell phones and remain connected to the network. These computers can communicate with each other and with a central database via a broadcast network. For example, physicians in hospitals can carry small computers that automatically receive information for each patient when the physician enters the room. Any instructions can be sent directly from the physician's computer to the nursing station. In the business world, you can carry a cell phone or tablet that maintains contact on the Internet to retrieve e-mail or scan Web sites while you are in meetings, other offices, or a client's office.

Broadcast media have two primary drawbacks. First, it is more important to provide security for the transmissions. Second, broadcast transmissions are limited to a defined frequency range, which might be subject to weather, echoes, or other interference and distance problems. Consequently, because the frequency space is shared, it can become overloaded when too many people attempt to transfer large amounts of data at the same time in the same location.

The problem of limited capacity arises because only a small number of radio frequencies can be used to carry data. Most of the radio and television frequencies are already being used for other purposes. Figure 3.10 shows some of the major frequency allocations in the United States. The Federal Communications Commission (FCC) allocated the personal communication service (PCS) bands in late 1993 for use by personal communication devices such as laptop computers and personal digital assistants (PDAs). To provide these frequencies, the FCC had to

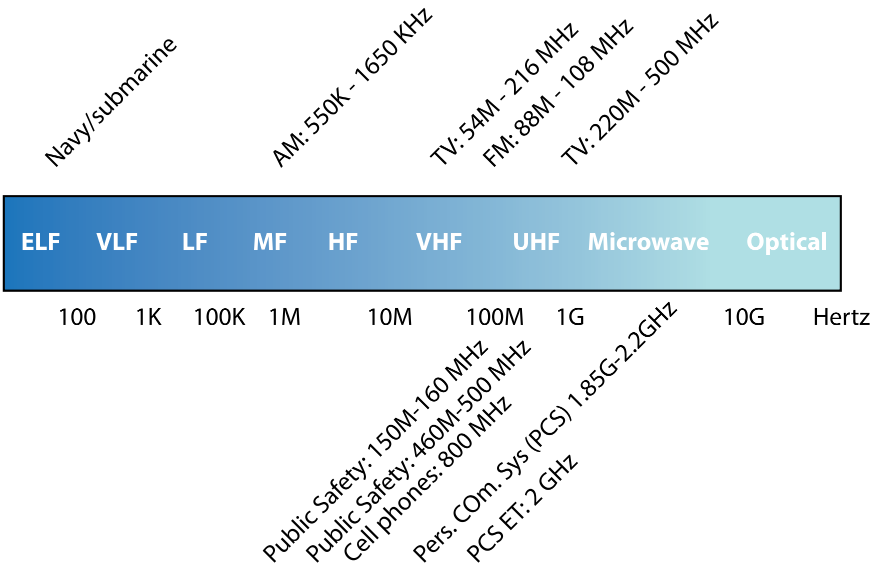

Figure 3.10

Electromagnetic frequency spectrum. Communication techniques are essentially the same for all media, because all waves physically have similar properties. However, the different frequencies affect the communication performance. Shorter wave lengths (higher frequencies) can carry more data. Some waves can travel longer distances. Others are more susceptible to interference. In any case, there are a finite number of frequencies available for communication. Hence, the frequency spectrum is allocated by governmental agencies. For complete details, see http://www.ntia.doc.gov/osmhome/allochrt.pdf.

take them away from existing users. Imagine what would happen if computers suddenly started sending information over the same radio frequency as that used by your favorite radio station. You would not be able to hear the voices on the radio, and the computers would miss most of their data because of the interference.

All governments allocate the frequency spectrum for various uses, such as radio, television, cellular phones, and garage door openers. The PCS frequencies were auctioned off to the highest bidders in 1994, raising more than $65 billion. The frequency problem is even more complicated when the signals might cross political boundaries. As a result, most broadcast telecommunications issues are established with international cooperation. Some of the overcrowding problems are being mitigated through the use of digital transmissions that cram more calls and more data into the same amount of frequency space.

Despite these problems, an increasing amount of business communication is being carried over radio networks. The International Telecommunication Union reports that in 2002, both Taiwan and Luxembourg averaged about 105 cell phone users per 100 people! In comparison with broadband Internet connections, in 2010, The Netherlands led the way where 37.8 percent of the population had broadband Internet access. The United States was ranked 14 at 27.1 percent; below Belgium, Canada, the United Kingdom, several European nations, and Korea. (source: http://www.oecd.org/document/54/0,3343,en_2649_34225_38690102_1_1_1_1,00.

Reality Bytes: Please Stop Watching Video So I Can Send e-Mail

From late 2010 to 2011, cell phone companies began rolling out fourth-generation (4G) services. Or they called them 4G service. Some, such as Verizon's LTE, have substantially faster data transmission rates. Early reports showed download speeds of 12 Mbps or more on Verizon's LTE network—although minimal competition existed so speeds might drop. Through more efficient use of the available bandwidth, LTE can provide as much as 20-times the capacity of older 3G networks. The only problem is that mobile data traffic is growing even faster. Cisco expects traffic to increase by an average of 92 percent every year from 2011 through 2015. Most of that growth will be due to video. Cisco's estimate might actually be low. Over the first four years that AT&T sold the iPhone, data traffic on its networks increased 8000 percent. Patrick Lopez, chief marketing officer for Vantrix, noted that "By the time 4G-LTE is mass market, all it will have done is allow wireless companies to keep up with the problem." The FCC notes that the wireless spectrum surplus of 255 MHz in 2011 will become a deficit of 275 MHz by 2014. The government is committed to freeing up an additional 500 MHz of spectrum for wireless communication. Wireless spectrum and efficiency is a critical issue to handle traffic between towers and users' cell phones. Backhaul, or the connections from the towers to the main Internet connections, is another problem. The big problem is that all of these changes take years to implement. In the short run, wireless providers are likely to advertise high-end video capabilities, and then charge a fortune to download the amount of data needed for high-quality video.

Adapted from David Goldman, "4G Won't Solve 3G's Problems," *CNN Online*, March 29, 2011.

html). On the other hand, advertised broadband download speeds in Japan and Portugal greatly exceed those in the other nations; running at slightly over 100 megabits per second, versus about 14 megabits in the United States (OECD data). The OECD data show monthly prices in Japan are about equal to U.S. prices. In US dollars, from $20 – 55 per month.

Although wireless Internet connections are not yet as fast as broadband connections, it is substantially less expensive to reach people with wireless connections. Already, in many nations, the number of cell phone subscribers exceeds the number of fixed telephone subscribers. According to the OECD, Korea tops the national list in terms of percentage of the population that has wireless broadband access (95 percent). In the U.S. 44 percent of the population had broadband wireless data plans.

Wireless is an increasingly important method for connecting computers and equipment. Historically, cables were the most common method of connecting machines, but people keep acquiring more devices and the tangle of cables is getting out of control. In addition, people are placing more value on portability and the need to access networks and data constantly. As a result, Figure 3.11 shows that multiple levels of wireless connections are being developed. In terms of wider network connections, you can use the cell phone system or WiMax which give you a range of 2-10 miles from an antenna. These frequencies are shared and usually regulated, so you have to purchase access through a commercial carrier (phone company). Availability is generally limited to major cities or highways, but are

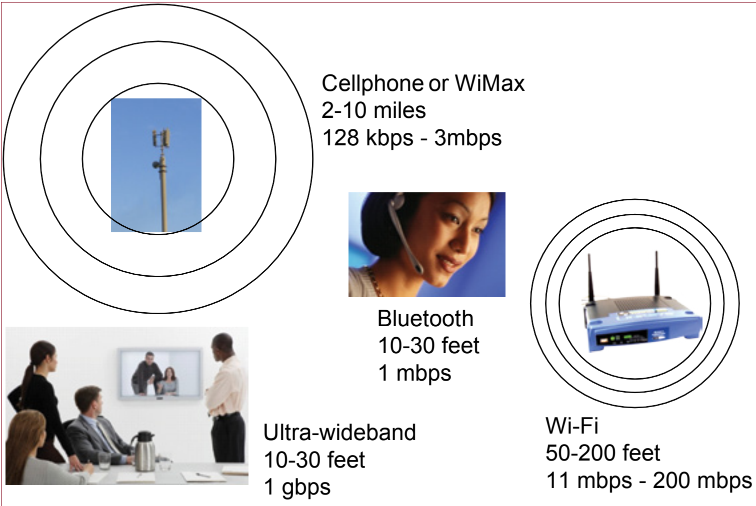

Figure 3.11

Wireless technologies. Competition for bandwidth causes tradeoffs. Wide-area technologies like cell phones and WiMax have limited bandwidth to reduce interference. Wi-Fi is a medium-range local area network technology. Bluetooth and Ultra-wideband are personal area technologies used for replacing cables for short-range connections.

acceptable for most business uses. **WiMax** is a relatively new technology that is designed to transfer wireless data over relatively long distances. Two primary frequency ranges are used—the most common one requires a license in the U.S. but not in Europe, so it is best used by commercial providers.

Most people are familiar with **WiFi** technologies because access points are relatively inexpensive and commonly installed in homes. WiFi enables multiple computers to connect to an Internet connection and is particularly useful for enabling laptop users to wander around a house or small office and remain connected to the network. Early versions were relatively slow and could support only a few users. The newest version (802.11n) can have data rates of 100-200 mbps or perhaps higher. The key defining feature of WiFi is that it has a maximum range of about 300 feet (outdoors). This range reduces interference from other devices, but means businesses have to install multiple access points to cover large buildings.

After several years of marketing by IBM, the **Bluetooth** wireless standard is becoming more popular. It has a range of 10-30 feet and version 2 has a maximum speed of about 2 mbps. The proposed version 3 increases the speed to 24 mbps. Today, Bluetooth is commonly used for connecting wireless headsets to cell phones. It could also be used to connect cell phones to each other or to other devices, but there are almost no other applications that use this technology at this time. The distance limitation and the built-in security provisions make it useful for short-distance, temporary connections.

From a technology perspective, **ultra-wideband (UWB)** is a more interesting short-range wireless technology. Unlike traditional wireless systems, UWB signals use a huge amount of frequency bandwidth (more than 500 MHz). Normally, this approach would interfere with other devices allocated to those frequencies,

LAN Name	Format	Speed (mbps)
10Base-T	Twisted pair	10
100Base-T	Twisted pair	100
Gigabit Ethernet	Twisted pair	1,000
Wireless LAN 11b, a,g	Wireless	11-54
Wireless LAN 11n	Wireless	250
fiber FDDI	Fiber optic	100
fiber ATM	Fiber optic	155
fiber high-end	Fiber optic	100,000,000 (100 terabits)

Internet Connection	Format	Speed (mbps)	Estimated Cost
Dial-up	Twisted pair	0.05	$20/month
DSL	Twisted pair	1.5+ down/0.5 up	$50/month
Cable modem	Coaxial	6+ down/1 up	$50/month
Satellite	Microwave	1.5 down/0.25 up	$50/month
Wireless/WiMax	Microwave	1.5-6 down/0.25+	$40/month
T1-lease	Twisted pair	1.544	$400-$700/month
T3-lease	Fiber optic	45	$2,500-$10,000/mo
ATM	Fiber optic	155	$15,000-$30,000/mo
OC-3	Fiber optic	155	$16,000-$20,000/mo
OC-12	Fiber optic	622	$20,000-$70,000/mo
OC-48	Fiber optic	2,488	$50,000-?/mo
OC-192	Fiber optic	9,953	
OC-768/future	Fiber optic	39,813	

Figure 3.12

Transmission capacity. Fiber-optic cables have the greatest capacity and they are immune to most traditional interference. However, they are expensive to install and repair. Most firms use twisted-pair or wireless connections within buildings and fiber-optic cables to connect buildings. You can purchase almost any Internet connection speed that you are willing to pay for. Leased line rates are negotiable and depend on distance and degree of local competition.

but UWB devices use low power levels—so that they appear as minor noise to any other device and are ignored. With the low power levels, they have a range limited to a few feet. On the other hand, UWB signals can carry huge amounts of data, in the range of 500 mbps to 1 gbps. This speed provides the ability to eliminate short cables for many devices—including video. For example, your digital camera or game machine could use UWB to connect to a television set without using cables, or your laptop could connect to the conference room projector without wires. At the moment, UWB is most commonly used for its side application—as short-range radar that can see through walls. Police and military organizations use this technology to locate people in crisis situations.

Transmission speed	Text	Image	Video (10 sec.)
Bytes	10,000	500,000	15,000,000
Bits (Bytes * 8)	80,000	4,000,000	120,000,000
	Seconds		
Dial-up 50 kbps	1.6	80	2400
DSL 1.5 mpbs	0.05	2.67	80
LAN 10 mbps	0.008	0.4	12
LAN 100 mpbs	0.0008	0.04	1.2
Gigabit 1 gbps	0.00008	0.004	0.12

Figure 3.13

Importance of transmission capacity. Text is not a problem even for slow dial-up lines, but images and video can be slow even over relatively high-speed Internet connections.

Transmission Capacity

As shown in Figure 3.12, transmission capacity is often measured in millions of bits per second. Each transmission medium has a different maximum capacity. Twisted pair can be relatively fast for short distances, but fiber optic is substantially better for longer distances. Getting a faster Internet connection is primarily an issue of cost. Fiber-optic cable might also carry a high installation cost if there are no fiber-optic connection points near your office. Today, twisted pair lines are commonly used for local area networks, but each connection run must be less than 95 meters (about 300 feet). Fiber optic lines are used between buildings and for long-distance connections.

The effect of the transmission capacity is shown in Figure 3.13. For small text and data files, the speed is not critical. Even slow dial-up lines can transfer a full page of text in a short time. The problem arises when you want to transfer more complex data like photos or even video. This figure shows why designing Web pages carefully is still so important. When over 75 percent of your clients are using dial-up lines, you need to limit pages to around 50,000 bytes, which takes at least 8 seconds to download. Most people will not wait more than 15 seconds for a page to load. Now you can see why Internet video is so difficult even with marginal quality. Even with video at broadband speeds of 1.5 mbps, site designers have to restrict the video to small sizes (one-fourth of a TV screen or smaller), use slower frame rates (as low as 15 frames per second), and employ lossy compression. These actions cut the video size by at least 1/8 (1/4 size * 1/2 frames). At that size, the 10 seconds of video could be sent in 10 seconds (80/8). Compression lets the developer specify a bit rate, and sites like YouTube tend to stay under 256 kbps, which reduces the quality.

Transmission delays are another important issue for some types of network uses. **Latency** is the time it takes for your signal to reach its destination. It is affected most by the speed at which the signal can travel across your connection medium. Fiber optic signals are fastest because they travel at the speed of light

Technology Toolbox: Creating Web Pages

Problem: You need to share information with others on the Internet.

Tools: Several tools exist to create Web pages, but at heart, Web pages are simply text files. The pages are written in the hypertext markup language (HTML). HTML consists of a few dozen tags that tell the browser how to display a page. A simple page can be written as:

```
<HTML>
  <HEAD><TITLE>Sample HTML Page</TITLE>
  <BODY>
    <H1>Section One</H1>
    <P>This is a sample paragraph on a sample page.</P>
  </BODY>
</HTML>
```

You can memorize the various tags, or you can use an editor to simply type the text and let it generate the tags. However, if you are creating pages for the Internet, avoid using document editors like Microsoft Word because the additional material it inserts makes the pages considerably larger than necessary.

Web pages display images using the <IMG SRC='myfile.jpg'> tag. You must be careful when creating image files to store them in a standard format: (1) graphics interchange format (GIF), (2) joint photographic exports group (JPEG), or (3) portable network graphics (PNG). You must be even more careful to watch the size of the file. A modern digital camera can create photographs that are 15 megabytes! Think about how long it would take a browser to download a file that large. In general, you must keep the total size of a page below 100 kilobytes.

Links are created with the anchor tag: Annual Results. The text or image between the starting and ending tags will be displayed to the user. Clicking on the link opens the file shown in quotes.

When you are creating an entire Web site, you need to use a style sheet. A style sheet contains a list of styles (fonts, sizes, colors, margins, and so on) that will be applied to the various elements on a page. For instance, you could specify that the main heading style (H1) would use an Arial typeface at 14 points in blue. The power of the style sheet is that all styles are defined in one place. By changing the style only once, every page on your site that is linked to that style sheet will automatically be displayed with the new style.

HTML is easy to use, but if you need more precise control over the page layout, you should use Adobe's portable document format (pdf). You can buy software that saves documents in this format. Browsers can download the Acrobat reader free from Adobe's Web site. This method is commonly used to distribute detailed documents such as t ax forms and posters.

Quick Quiz: Search the Web, or create a document and View the source to do the following:

1. Display a word or phrase in boldface.
2. Link a style sheet to an HTML page.
3. Display a table with three rows and four columns.
4. Display a numbered list of five items.
5. Display an icon in GIF format with a transparent background.

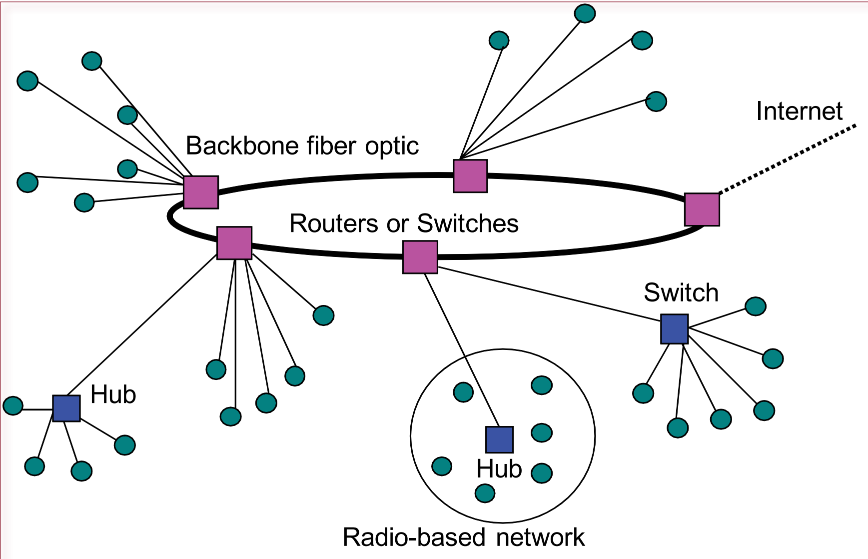

Figure 3.14

Connecting networks. Several challenges arise if you build an enterprise network. An enterprise network often connects many smaller networks established in different departments, buildings, or even nations. The hardware and software components must follow standards so they can communicate. An MIS team has to manage the overall structure to maximize efficiency, avoid duplication, and solve problems.

and require fewer repeaters. Electrical signals are also relatively fast, but the overall latency can be higher if the signal needs to travel through multiple repeaters or switches. Microwave signals tend to be the slowest and subject to interference from the weather. Satellite connections tend to have the most latency because the signal has to travel from the ground, across 23,000 miles of space to a satellite and another 23,000 miles back down to the base station. This delay can easily reach a half-second or more and makes it difficult to use satellite connections for some purposes—such as multiplayer games. Note that latency is independent from the bandwidth—so an ISP can advertise that it provides a 1 mbps satellite connection, but the delay due to latency might cause problems with how you want to use the network.

Connection Devices

To reduce overall traffic, larger organizations often find it beneficial to build the corporate network from a set of smaller networks. Both large and small companies use similar techniques to connect their networks to the Internet. Figure 3.14 shows a common configuration. Computers within a building or smaller area are linked into a hub, switch, or router. This interconnection device is then linked to the backbone, which is typically a fiber-optic line. A specially-configured router then connects to the Internet.

If you look at the physical box, it can be difficult to tell the difference between a hub, router, and switch. However, internally, they function quite differently, and you need to understand those differences to decide how a network should be configured. Hubs are the simplest connection devices. They essentially act like a gi-

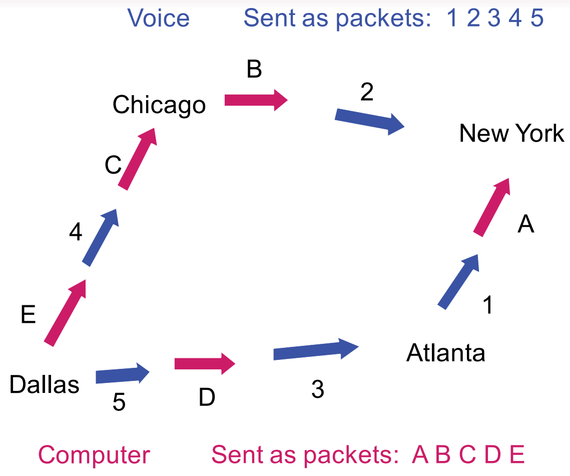

Figure 3.15

Network packets. All data is converted to digital form and broken into small packets. Each packet contains a destination and source address, along with sequencing instructions. The packets can be sent over different routes or mixed with other packets. At their destination, the original message is restored by the network.

ant junction box. Any device connected to a hub shares all of the lines with the other devices. It is just a mechanism to plug several computers together. Switches and routers actually examine every packet that passes through them and decide where to send each packet. They are actually specialized computers that can be programmed to identify network problems and intelligently route traffic to the fastest route.

Routers and switches are crucial to improving efficiency in large networks. In many respects, a router works like a post office. It examines the destination address of each packet of information and selects the best way to send that packet. Routers and switches improve performance by choosing the path of the message and segmenting large networks into smaller pieces. For example, router segments might be assigned to each department of a large company, where each department has its own server. Most of the messages stay within the specific department, so they do not take up bandwidth across the entire company. Only messages that truly need to be sent to other areas will be transmitted outside the departmental segment.

The main differences between switches and routers consist of how each one identifies other objects on the network. Switches use physical addresses that are assigned to each network interface card. Routers use a logical address that is easier to change. Routers also contain more intelligent routing algorithms. They can analyze traffic patterns and communicate more data with other routers to find the best transmission path. Switches are commonly used to isolate traffic within a local area network, and routers transfer data across longer distances and more complex networks such as the Internet. It is possible to build a network entirely based on routers, but it is usually cheaper to use switches where possible. Just keep in mind that most switches cannot transfer data beyond seven hops, so if you need lots of connections, you will have to add routers.

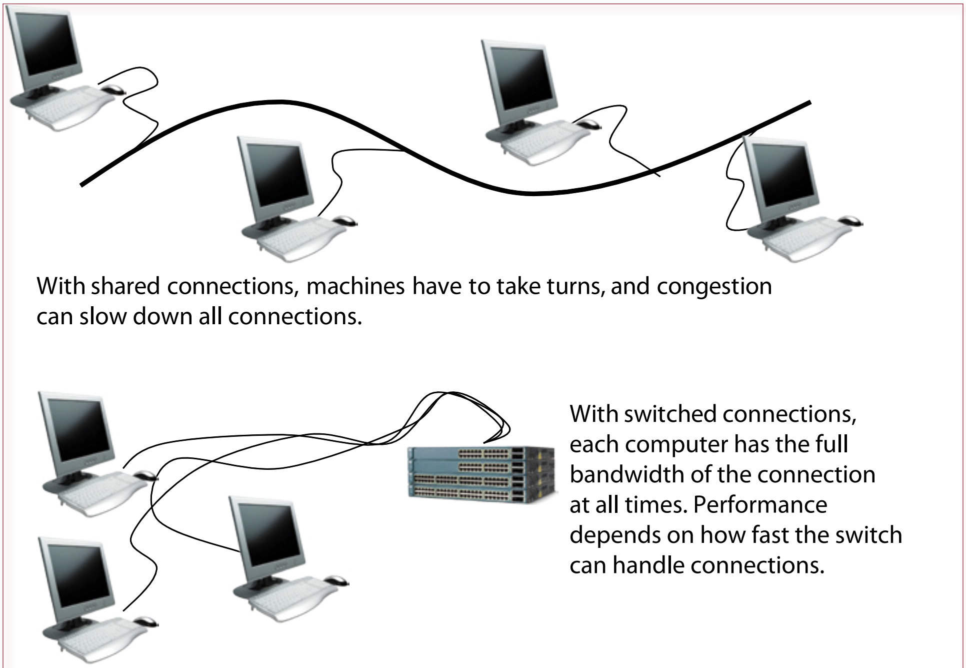

Figure 3.16

Shared connections. Some networks rely on sharing the transmission medium with many users. Sharing means computers must take turns using the network. Sometimes one or two highly active users can slow down the network for everyone, so you will not really get the listed transmission capacity.

Network Structure

How can multiple users share a single network? Transferring data between two devices is relatively easy—you just find a way to physically connect the two computers and install some software to handle the communication. You might be able to build a network that works the same way—by connecting each computer to every other computer. But eventually the cabling costs would be huge. Instead, at some point communications travel over a shared link. Think about a network between two cities (say Los Angeles and Boston). Computers within each city might be connected to each other, but they all share one long-distance line to transfer data between the two cities. Almost all networks are built using shared communication links. Several methods exist to share data connections, and the method can affect the overall performance of the network. But, it is also important to remember that the number of people and their usage of the network will affect your performance. Shared networks require careful management—particularly when a few "bandwidth hogs" exist.

Network Packets

A key feature in sharing networks is that data is transmitted in small batches. All data, including voice and video, is converted to digital form. The data is split into smaller **packets** that contain a source address, destination address, and a sequencing method. The packets tend to be a uniform size (a couple thousand bytes) and

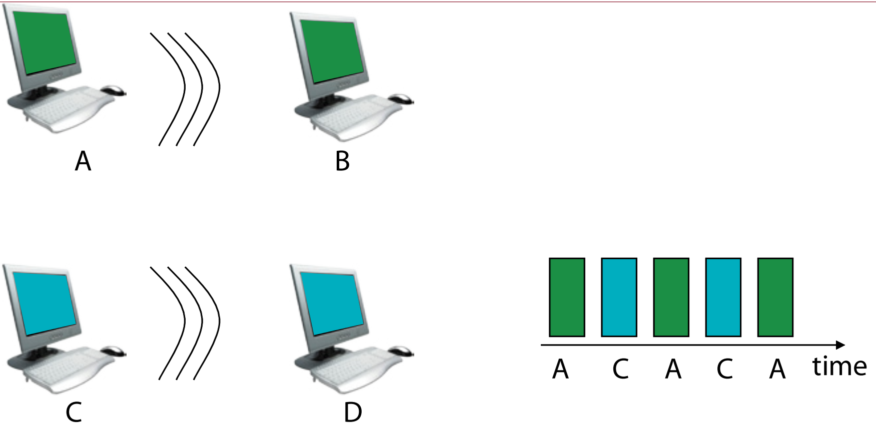

Figure 3.17

Time division. Only one machine or person can send data at a time. Data is limited to short packets to prevent one communicator from monopolizing the time.

prevent one user from monopolizing the connection. They also make it easy for a single network to handle data, voice, images, and video. The network simply focuses on routing the packets to the appropriate destination. Figure 3.15 shows an example of voice and computer data being sent over the same network. The switches and routers choose the best path to send the data when each packet arrives. The network delivers the packets to the appropriate receivers in the proper sequence.

Today, most networks use the protocols defined in the Internet standards. The Internet is capable of carrying any type of data packet; however, one catch with the standard Internet is that there is no service guarantee. Your packets could be delayed or lost along the way. Most applications have a method for recognizing lost packets but it takes time to retransmit the packet. This issue is minor for Web sites and e-mail, but can become annoying if you want to use the connection for phone calls. Newer versions of the Internet protocols provide support for guaranteed levels of service. But, so far network companies have not devised methods to implement and charge for **quality of service (QOS)**. Part of the problem lies with the need to communicate across networks owned by multiple companies. There is also a question of whether people are willing to pay extra for a defined level of quality. However, it is possible that someday you will be able to conduct video meetings by reserving a certain level of speed at a set time.

Shared Connections

Figure 3.16 shows that individual machines can be connected to a shared medium or they can be connected individually to a switch. However, the switch uses a shared medium to connect to other switches and routers; so ultimately, connections are shared at some point on almost any network. Wired local area networks commonly use switches to separate traffic within departments or physical areas. Wireless (and cell phone) networks require all computers to share the same radio frequencies. Some of the most advanced sharing technologies have been created to handle the huge increases in traffic experienced with cell phone networks.

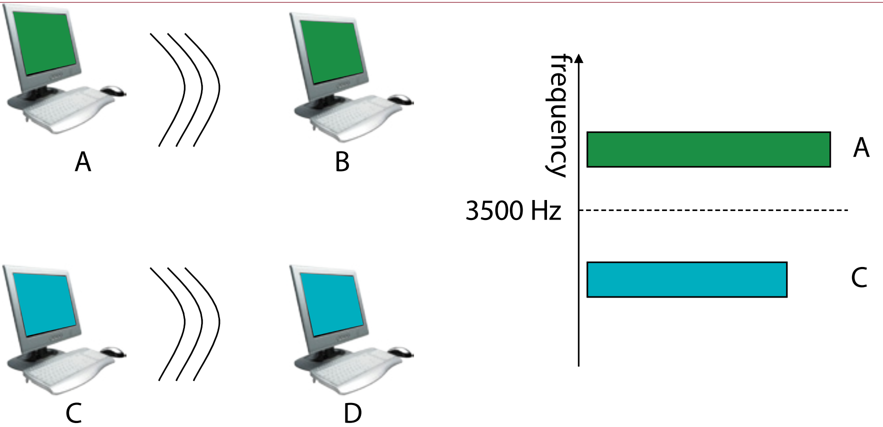

Figure 3.18

Frequency division. Each pair communicates at a different frequency (pitch). Radio and television stations work this way by using assigned frequency ranges (channels).

In theory, all communication systems that share a communication medium share the same problems. The solutions to the problems are similar regardless of the communication technology. In practice, engineers point out that the techniques require considerable modification depending on the physical connection method, but you will have to study the engineering books to understand those details. For now, think about the protocols that people have created to handle a common shared-communication environment: the classroom.

What happens if everyone tries to talk at the same time in a classroom? Is it possible for more than one person to talk at one time? The common answer is that only one person is allowed to talk at a time. A signaling mechanism (raising your hand or being recognized by the leader or instructor) is used to identify which person gets to speak. Typically, the leader also limits the time any person is allowed to speak. Figure 3.17 illustrates the basic process of **time division multiplexing (TDM)**. Only one machine or person can communicate at a given time. Rules or protocols determine who gets to send a packet, what to do if two machines try to communicate at the same time, how to address the packets, and the length of the packets.

If you think about the problem for a while—and engineers have been thinking about the problems for many years—you might find a second solution. If everyone's hearing is good enough, it would be possible for two people to carry on a conversation at the same time in the same room. Simply ask one pair to communicate at one frequency (sing soprano), and the other pair to communicate at a different frequency (sing bass). In electronics terms, it is relatively easy to share a communication line by using this **frequency division multiplexing (FDM)**. Figure 3.18 shows the basic process—it has been in common use for decades. Radio and television stations are assigned separate channels (frequency ranges) so they can broadcast at the same time without interfering with each other. The telephone network was an early adopter of this approach—limiting the frequency range of conversations to about 3000 Hertz and shifting connections to different frequen-

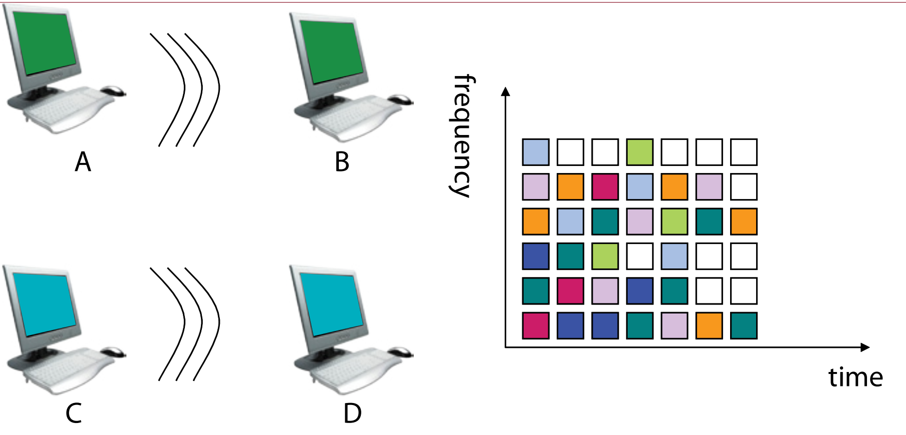

Figure 3.19

Spread spectrum. Radio networks share the scarce airwaves by breaking communications into small packets that are sent at different frequencies and time slots. In a given time slot, if one frequency is not available, the transmitter shifts to the next frequency. If all frequencies are being used, the transmitter must wait for the next time slot.

cies so they could be simultaneously transmitted across expensive long-distance lines.

With a little more thought—and a focus on packets—you will see that it is possible to combine both frequency and time division into **direct spread spectrum (DSS)**. As an interesting historical footnote, the original patent for the technique was granted to actress Hedy Lamar and composer George Antheil in 1942, during World War II. Hedy recognized that a piano-roll mechanism could be used to shift radio frequencies at random to prevent opponents from intercepting or jamming radio transmissions.

Figure 3.19 illustrates the process. When a computer or cell phone wants to transmit a packet, it checks for an open frequency slot. At each time interval, that frequency might shift, so the receiving computers listen to all available frequency slots and check the destination addresses of each packet. Today, the key benefit lies in its ability to fully utilize the available bandwidth. Allocating fixed frequencies to everyone is inefficient because many of the frequencies will be unused at any point in time. Splitting the communication by time and frequency enables each transmitter to choose any available open frequency. Remember that communication frequencies are tight in every nation. Yet, researchers have shown there is plenty of bandwidth available—if all transmissions were converted to spread spectrum. Because of the technical and political issues, this change is unlikely to happen soon. However, almost all new applications use spread spectrum technologies within a defined frequency range.

One cellular phone manufacturer (Qualcomm) invented a radically new method of sharing a communication medium. Today, the Verizon network uses this technique to provide high-speed voice and data communications to its cell phone customers. Newer proposals for cell phone networks have discussed using similar technologies. The method begins with DSS, but every transmission also incorpo-

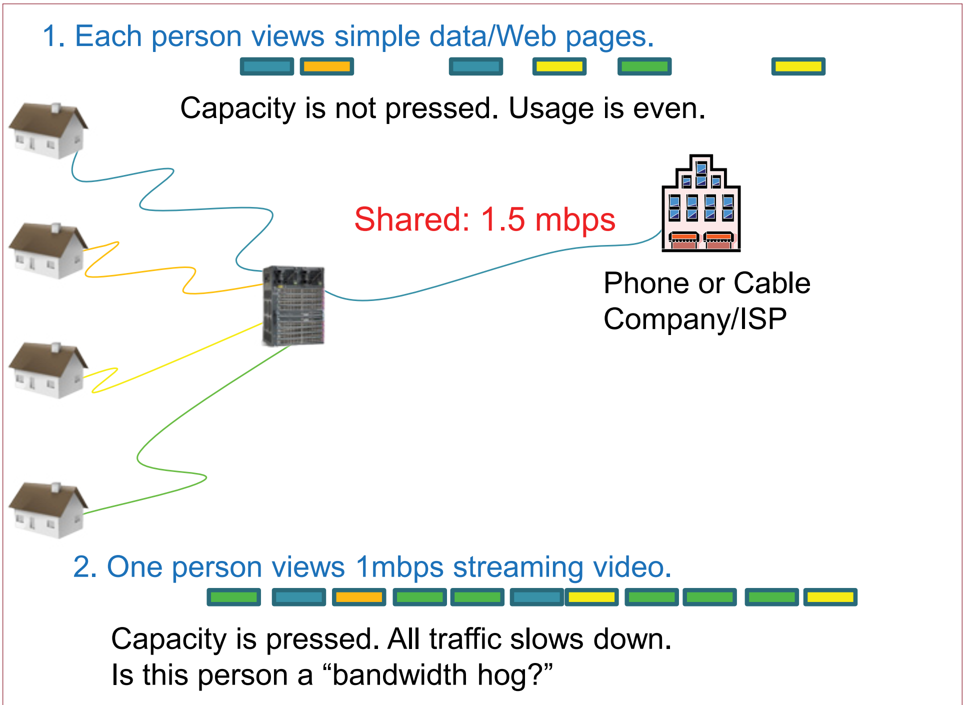

Figure 3.20

Managing shared networks. Some network applications use significantly more bandwidth or capacity. If a small number of advanced users eat up the shared capacity, does that make them bandwidth hogs?

rates a unique ID code. By adding some mathematical calculations to every device (including the cell phone), multiple devices can transmit at the same time on the same frequency—using the math to recover separate data packets. This code-division multiple access (CDMA) method is efficient, but patent and licensing battles have made other providers reluctant to adopt it.

Managing Shared Networks

As the number of users increase on a network, it is likely that the network will encounter congestion problems. For example, the AT&T cell phone network was widely criticized when it was the sole provider of the Apple iPhone. In some major cities including New York and San Francisco, many users reported problems with dropped phone calls. Similar problems can happen with any shared connection when traffic reaches the bandwidth capacity. For example, most cell phone networks face problems if a major crisis hits a particular location and thousands or millions of people all try to use their phones at the same time. Related problems happen with data networks. Most Internet providers provide data capacity based on the "typical" usage rates. For example, as shown in Figure 3.20, an Internet connection might have a data capacity of 1.5 mbps, and that line might be connected to 10 households. As long as each household uses typical Web browsing at random times, the data will transmit quickly enough that interference will be low and users will mostly see the full transfer rate. But, what happens if three

Reality Bytes: Video Traffic from NetFlix Dominates the Internet

For many years, peer-to-peer file sharing dominated the traffic flow on the Internet—particularly in the United States. Many people used BitTorrent to exchange large data files—particularly video, and often illegally. By May 2011, NetFlix, with its online subscriptions, had taken over the number one position in terms of total data transferred on the Internet. The traffic account for 22.2 percent of all U.S. broadband traffic. BitTorrent was still a close second at 21.6 percent. The percentages are based on the actual usage or volume of traffic, not on the bandwidth capacity. The main reason for the increase is the relatively low price of $9 per month for a subscription. Experts estimate that it costs NetFlix between 2-10 cents to deliver a movie online, based on server and bandwidth costs. The licensing costs are higher but the exact cost is not available.In some ways, today's Internet business world is like the early days of business—you simply did not trust anyone you did not know. As economies grew, states learned to control fraudulent businesses by requiring them to register. Today, legitimate businesses are registered with at least one state government and must file annual forms and provide legal contact information. Most of this data is accessible online. Is it really any different to require people to file legitimate contact information for online registrations?

Adapted from Ryan Singel, "Most Content Online is Now Paid For, Thanks to Netflix," *CNN Online* (Wired), May 18, 2011.

users start watching hours of streaming video that requires 500 kbps per person? These three people will effectively use up the entire capacity of the shared line. Streaming video can be a problem on many networks. The BitTorrent data transfer program can also be a problem because it tries to utilize as much bandwidth as possible and is often used to transfer huge files.

Internet and cell phone providers (along with your university network managers) have faced congestion issues for years. As new network applications have emerged—particularly video—the capacity demand has continued to increase. In any shared network, some early adopters are likely to use the new applications more than other people. Consequently, their bandwidth or capacity usage will be higher than other users. Is that bad? Some commercial networks have approached the problem by labeling them **bandwidth hogs** and imposing penalties for anyone who uses the network more than others.

Managing a commercial network is more complicated than it appears—and it already seems hard enough. In addition to installing lines and keeping the network switches and routers running, the commercial providers have to determine pricing and find ways to encourage people to choose their network, yet keep costs low enough to earn a profit. And that does not begin to deal with any of the security issues covered in Chapter 5. To obtain customers, many companies offered "all-you-can-eat" data plans for a fixed monthly fee. But, new applications were developed including video, large-file sharing (BitTorrent), and Web video conferencing; and these new applications require considerably more bandwidth than simple e-mail and Web browsing. Some people adopted the new applications quickly, with the potential to cause congestion for everyone. Ideally, network companies would expand their capacity and provide higher speeds for everyone—but that increases costs, and some people are likely to consume bandwidth no matter how much is available.

- Prioritize Traffic
    - Slow down some users—perceived hogs.
    - Slow down based on type of traffic.
        - . Packeteer—examine packets to identify.
        - . Connection port (rare).
    - Sell quality of service (rare yet).
- Pricing Mechanisms with Data Caps
    - Overage fees.
    - Differential pricing.
    - Time-of-day pricing (rare yet).
    - Potential problems as speed increases (4G cell).
        - . Medium quality is at least 1 GB/movie.
        - . Common data cap of 5 GB/month.
        - . At 5 mbps, 5 GB cap reached in 133 minutes.
- “Net Neutrality” Proposal
    - What if a commercial network tries to slow down traffic from a competitor?
    - Example: Comcast owns NBC (2011/01)
    - But can “neutrality” be defined to still allow networks to manage traffic?

Figure 3.21

Options for managing network traffic. Commercial networks are trying several methods to reduce the congestion caused by advanced users. Prioritizing certain types of traffic can be a good solution, but it could also be abused. Pricing based on usage makes the most sense economically, but it can be harder to implement and harder to sell to customers.

Figure 3.21 outlines some of the options available to manage traffic on networks to reduce congestion. Most commercial and university networks use a combination of the choices. In some ways, prioritizing traffic can be the most neutral. Network devices can monitor the type of data carried by each packet and determine how quickly to send the packet. High-bandwidth, low-value packets are delayed when network usage is high. For example, e-mail is often given high priority, while online games and file sharing are given low priorities. The packets eventually get delivered, but users of those applications see substantially slower data rates. A few providers have implemented plans to reduce the data speeds for specific users. Once a person exceeds a certain amount of data transfer for the month, their future usage will be slowed down for a month.

Cell phone providers have become the most creative in terms of pricing. An economist would observe that congestion pricing mechanisms are more efficient than arbitrary limits. The catch is that marketers want to sell phones and services by offering more open-ended plans, and customers get nervous about data caps. Most vendors have implemented simple tiered data plans, where customers pick a plan by estimating monthly usage levels.

However, data caps are often set relatively low and they can be easy to exceed. For example, consider a monthly cap of 5 gigabytes (GB). Then consider a 4G cell phone that can download data at 5 megabits per second. That phone would hit the data cap after only 133 minutes—or about one high-quality movie. (Calculation: 5 billion Bytes * 8 bits/Byte / 5 million bits/second / 60 seconds/minute).

Before you make any hasty decisions about the managers of networks, think about the problem from a business perspective for a minute. What happens to the network if ten people want to stream a movie at 5 mbps at the same time? Effectively, that generates 50 mbps demand. How much capacity can be built into the network? What if 100 people do the same thing? That amounts to 500 mbps simultaneous demand. How are you going to build a network that can handle that level of traffic? And you are still looking at only 100 people on the network. Digitized voice is only 64 kbps, so the high-definition movie takes about 80 times the bandwidth as a phone call. Either you need a massive increase in network capacity or some way to discourage people from using high-bandwidth applications. Although, you might want to talk to the marketing people about toning down the advertising telling people they can watch movies on their cell phones.

Standards

How is it possible that you can connect your computer to a network at the office, at home, or while on the road, even overseas? Standards are agreements among vendors, customers, and nations. If everyone follows the standards, equipment can be connected, data can be shared, and you can connect your computer to the network anywhere in the world. But evolving technologies and competition significantly complicate these tasks. Setting standards and moving to new ones is always a challenging task. However, the Internet exists because of standards.

The Need for Standards

Standards are important with networks. Many different types of computers and connection devices exist. Each computer and network company has its own idea of which methods are best. Without standards, there is no way to connect computers or networks produced by different vendors. Standards are also supposed to protect the buyers from obsolescence. If you install network equipment that meets existing standards, you should be able to buy products in the future that will work with it.

Unfortunately, the process is complicated because several standard-setting organizations exist. Each major country has its own standards organization such as ANSI, the American National Standards Institute, plus several international organizations exist, such as ISO and the ITU (International Telecommunications Union, renamed from CCITT), charged with defining computer and communication standards. In addition, manufacturers of computers and telecommunications equipment try to define their own standards. If one company's products are chosen as a standard, it gains a slight advantage in design and production.

It is not likely that typical managers will be involved in the issues of setting and choosing communication standards. Yet, as a consumer, you need to be aware that there are many standards and sometimes variations on the standards. (In this industry, a standard does not mean there is only one way to do something.) When you are buying telecommunications equipment, the goal is to purchase equipment that meets popular standards. It is not always easy to decide which standards will become popular and which ones will be abandoned.

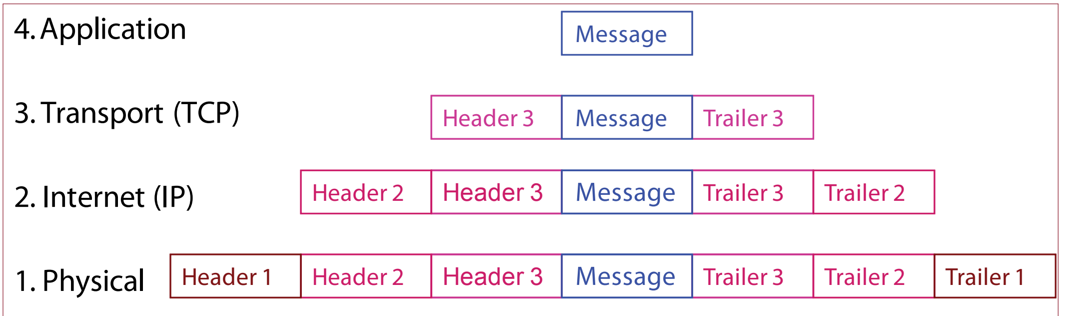

Figure 3.22

TCP/IP reference model. The model illustrates how data from an application like e-mail is turned into packets at the transport layer, routed to the destination at the IP layer, and physically transferred as bits at the physical layer.

A Changing Environment

Why are there so many standards? It would be far simpler if everyone could agree to use one standard and build products that are compatible. The problem with this concept is that technology is continually changing. Thirty years ago, phone companies never considered using digital transmission over fiber-optic cables, which is the dominant form of long-distance transmission used today. When the Internet was first created, the technology of the time would not handle the billions of addresses we need today.

As each technology is introduced, new standards are created. Yet we cannot discard existing standards because it takes time for firms to convert to the new technology. Besides, as manufacturers gain experience with a technology, they add features and find better ways to use the products. These alterations usually result in changes to the standards. An additional complication is that many companies are modifying their products at the same time. It is hard to determine in advance which changes are worthwhile and should be made standards.

The net result is that standards can be useful, but managers have to be careful not to rely too much on a standard. First, remember that even if two products support a standard, they still might not work well together. Second, if you choose a standard for your department or company, remember that technology changes. Corporate standards should be reevaluated every year or so to make sure they are still the best solution. Global standards are required to ensure compatibility of networks and efficient routing of data.

Internet TCP/IP Reference Model

The Internet transmission protocol standard is an important standard. It also helps you understand how networks operate and how the Internet works. As shown in Figure 3.22, the TCP/IP (transmission control protocol/Internet Protocol) reference model breaks the process into four layers: application, transport, Internet, and physical. Breaking the process into separate layers is critical to building large networks. Each layer can be handled independently of the others. For example, at the physical layer, replacing a wired connection with a wireless one should not affect any of the higher layers. The physical layer devices simply have to provide the same functionality to the Internet layer.

Technology Toolbox: Transferring Files on the Internet

Problem: You need to transfer data files to a Web site.

Tools: It is relatively easy to transfer files from your computer to another computer connected on a local area network. Assuming that you have the proper security permissions, you can see the other computer as if it were another disk drive and drag-and-drop the file. But no company is going to open its network to allow you to do the same thing across the Internet. So, how do you transfer files to a Web server if it is not located on your LAN?

Connection	Strengths	Weaknesses
FTP	Inexpensive and easy to use.	No security.
FrontPage	Very easy to use with secure login.	Limited availability and users have to purchase client software.
WebDAV	Internet standard and can be secure.	Web server security is weaker by allowing directory browsing.
VPN	Very secure.	Need to purchase server and client software and requires extra setup steps.

Surprisingly, given how frequently people need to transfer files, this question does not yet have a good answer. An older method still in use is the file transfer protocol (FTP). Your Web browser probably supports drag-and-drop with ftp. You can enter an address like ftp://myserver.com/www, and the browser will connect and display a list of files in the folder. Then you can drag-and-drop your files or an entire folder onto the folder window and it will be transferred to the server. The drawback to FTP is that passwords are sent in the clear and might be intercepted by hackers. Most companies today require that you transfer files using secure FTP tools that automatically encrypt login and file data. These systems require your client computer to have special software that matches the server's encryption.

The Internet community has developed the Web distributed authoring and versioning (WebDAV) service to help solve the problems of sharing files. As an open standard, it is available on many servers. If you have an account on a server that has WebDAV enabled, you can use the Internet Explorer to connect and transfer files. Just be sure to use the File/Open command, enter the name of the site, and check the box to "Open as a Web Folder." WebDAV is an extension of the HTTP protocols, so it supports secure login and can take advantage of common encryption tools.

The other approach to transferring files to a corporate server is to establish a virtual private network (VPN) connection first so that all communications are encrypted. Then you are free to use FTP or any other common protocol inside the VPN tunnel because no one can intercept and decrypt the messages. But you need control over the server or a special server to configure VPN.

Most hosting companies use Web file uploads, which can be secure if Web security is installed on the server. But, transferring files has been cumbersome and generally only one file can be uploaded at a time. The new HTML5 definition might improve this process because it supports file upload with a drag-and-drop feature built into the browsers.

Quick Quiz:

1. Which methods can you use to transfer files to a university server?
2. Why is FTP considered a security threat?
3. What other objections exist to FrontPage?

Notice that moving down, each layer takes the data from the prior layer and adds header and trailer information. This additional data is necessary for each layer to perform its function, but it means that more data must be transferred. For example, even if your physical connection can transmit data at 10 mbps, a 10-megabit file cannot be transferred in one second. Depending on the application and the network details, the overhead from the layers can be 20 percent or higher.

Note that there is a competing seven layer network model that defines the ISO-OSI view of networks. The model is similar in concept to the TCP/IP model, but it breaks level 1 into two layers, and level 4 into three layers. Since this is not a networking book and the Internet protocol dominates, there is no reason to cover it here. You can find details on the Internet if you ever need them.

Subnet/Physical Layer

The purpose of the subnet or physical layer is to make the connection between two machines and physically transfer bits of data. It is directly related to hardware. Standards exist to specify constraints on voltage, type of wire, frequency of signals, and sizes of physical connectors. Raw data bits are transferred at this stage. Many different technologies exist, including wireless, wired, and fiber-optic lines.

The network interface card is a critical component of the physical layer. In addition to handling the data transfer, each card has a unique **media access control (MAC)** address. This address is a globally unique number that is assigned to the card by the manufacturer. Switches use this number to identify every device on the network and send packets to the correct device. In fact, the technical difference between switches and routers is that switches work at Level 1 and routers at Level 2.

Internet/Network Layer

The Internet layer is concerned with routing messages to the appropriate location. In particular, it selects the appropriate path for a message in networks where there is a choice. Its main advantage is that it separates the higher layers from the physical transmission of data. The network layer handles connections across different machines and multiple networks. The Internet Protocol (IP) is the standard used in routing packets on the Internet network. With IP, each packet is treated independently of the others and each can follow a different route to the destination. Each machine must have a globally unique address, so a mechanism is established to assign numbers to machines. It would be difficult to use the physical MAC addresses because they are randomly assigned. A logically assigned IP address is better because the number can be used to help route each packet by using portions of the number to segment the network. Plus, if you need to change a physical network card, you can keep the same IP address and the Internet will continue to function correctly. The current version of IP (IPv4) uses a 32-bit address, which is beginning to cause some problems because IPv4 supports a maximum of about 2 billion addresses, and the world is beginning to run out of numbers.

The newer IPv6 standard supports 128-bit addresses, but it will take time to phase in the new system. Newer servers and operating systems support the IPv6 protocol, but it will take time to update all of the routers on the Internet. In the meantime, most companies are using intermediate steps to allocate the IPv4 numbers. Windows 7 automatically activates IPv6 alongside IPv4 so it can support both types of connections.

Reality Bytes: NEC Lab Sets Fibre Speed Record

In March 2011, NEC Laboratories in Princeton, New Jersey announced they had pushed more than 100 terabits of information per second through a single optical fiber. That amounts to transferring roughly 250 dual-layer Blu-ray discs in one second. Tim Wang from the Lab noted the feat marks "a critical milestone in fibre capacity." By comparison, Tim Strong of Telegeography noted that total capacity between New York and Washington D.C., one of the world's busiest Internet routes, is only a few terabits per second. But, he also noted that traffic has been growing by about 50 percent per year. Typically, optical fiber capacity is increased by splitting transmissions into multiple frequencies—or colors. Additional capacity can be gained by using different polarities, amplitudes, and phases of light. Dayou Qian of NEC reported sending 101.7 terabits per second through 165 kilometers of fiber. His team used 370 separate lasers and several kilowatts of power to accomplish this feat. At the same conference, Jun Sakaguchi of Japan's National Institute of Information and Communications Technology in Tokyo also reported reaching 100 terabits per second. His team used a special fiber with seven internal light-carrying cores. Using a completely different method, Wolfgang Freude and his team at the Karlsruhe Institute of Technology in Germany achieved transfer rates of 26 terabits per second—using a single laser. Although somewhat slower than the 100 terabits record, the cost of implementing Freude's method is substantially lower. His team used orthogonal frequency division multiplexing (multiple colors) and a fast Fourier transform to separate the colors on the receiving end. Each color arrives at slightly different times and the process could be embedded onto a silicon chip.

Adapted from Jeff Hecht, "Ultrafast Fibre Optics Set New Speed Record," New Scientist, April 29, 2011; and Jason Palmer, "Laser Puts Record Data Rate Through Fibre," BBC Online, May 22, 2011.

Transport Layer

The transport layer is responsible for dividing the application data into packets and providing logical connections to the network. The transport control protocol (TCP) is commonly used on the Internet to handle these connections. TCP supports multiple applications at the same time by creating numbered ports. For example, e-mail is usually transferred through port 25, and Web data through port 80. TCP sends the data packet to the specified port on the desired machine. TCP on the host machine listens to these ports and sends the incoming data to the appropriate application server. TCP also monitors the packets to see if any are lost in transmission. If so, the recipient machine can request that the missing packet be re-sent, providing a highly reliable connection between two machines.

The Internet also supports the user datagram protocol (UDP), which is a highly simplified transport method. Most important, it does not guarantee that a packet will be transferred. Generally, users do not get to choose between TCP and UDP. This choice is made by the software developer at the network level. But why would anyone choose UDP when there is no assurance that the packets will be delivered? The main reason is speed. Because UDP is so simple, it adds only a tiny overhead to each packet, which makes it useful for large transfers of data, such as large files and streaming multimedia. If necessary, the application can check at the end to ensure that all data was transferred.

Application Layer

The application layer consists of tools and services for the user. Typical Internet applications include e-mail, file transfer (FTP), and Web browsing with the hypertext transfer protocol (HTTP). These systems work because developers have agreed to follow basic standards.

With the TCP/IP reference model, applications are responsible for incorporating authentication and compression. Not having a standard underlying method for handling security has caused some problems with TCP/IP. Few applications actually have any security, several incompatible variations of security systems have been created, and hackers have been able to write programs that attack the underlying, unprotected layers. Security is one of the main problems being addressed in IPv6 and Internet2.

The Internet

What is the Internet, how is it controlled, and how does it work?

The Internet is a loose collection of computer networks throughout the world. It began as a means to exchange data among major U.S. universities (NSFnet of the National Science Foundation) and connections to various military organizations and U.S. defense suppliers (Arpanet of the Advanced Research Projects Agency). No one knows how many computers or networks are currently connected by the Internet. The numbers have been increasing exponentially since the early 1990s, so any number is likely to be wrong within a few days. To give you some idea of the Internet's astounding growth, in January 1993, there were 1.313 million host computers. In January 1994, there were 2.217 million hosts located in more than 70 countries. In 2000, Telecordia estimated that the Internet had exceeded 100 million hosts. By 2007, Internet Systems Consortium counted over 400 million host computers. Billions more people connect to these hosts using computers, tablets, cell phones, and so on. In 1994, over 20 million people had access to at least e-mail services. As of mid-1994, commercial use (as opposed to university and government use) accounted for 50 percent of the Internet usage. By 2004, over 500 million people worldwide had access to the Internet. Measuring the Internet is difficult, since machines are not always connected. Most studies use some type of survey to estimate the size. However it is measured, usage of the Internet continues to grow.

What exactly is the Internet? At heart, the Internet is just a communication system for computers. It is defined by a set of standards that allow computers to exchange messages. The most amazing aspect of the Internet is that there really is no single person or group in charge. Anyone who wishes to connect a computer to the Internet simply agrees to pay for a communication link—via an Internet service provider (ISP)—and to install communications hardware and software that supports the current Internet standard protocols. The person or company is given a base address that allows other computers to identify users on the new computer. Standards are defined by a loose committee, and addresses are controlled by another small committee. The committees are convened purely for the purpose of speeding the process; all decisions are up to the organizations connected to the network. Participation in the Internet is voluntary, and there are few rules, just standard practices and agreements. From a business or consumer viewpoint, there are two primary aspects to the Internet: establishing a connection and using the Internet.

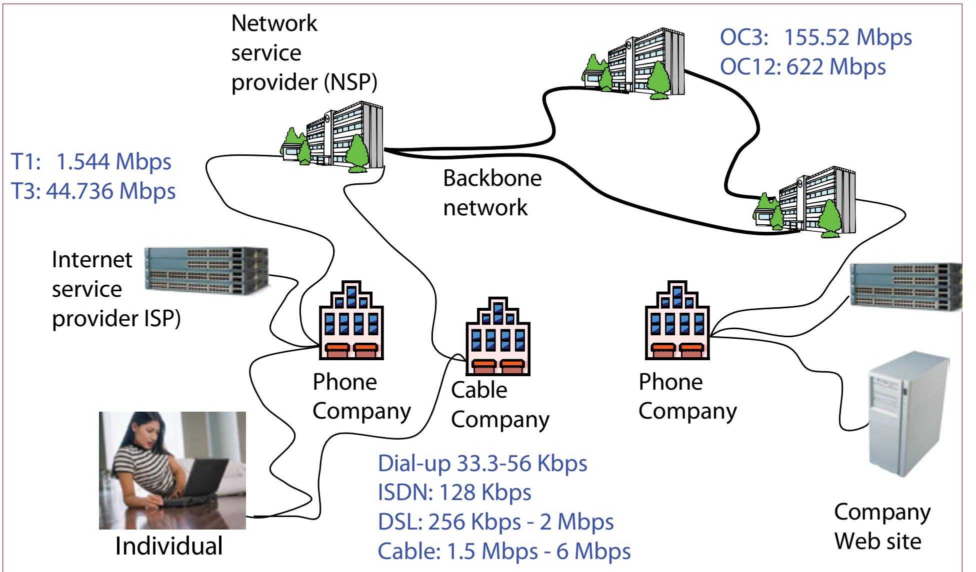

Figure 3.23

Internet connections. Each computer must be connected to others. The Internet has a connection hierarchy. Companies and individuals typically use phone company lines to connect to an ISP. The ISP connects to an NSP, which routes data over the high-speed backbone network to the destination NSP, down to the other ISP, and to the final computer. Each step may involve several computers.

How the Internet Works

The Internet is a communication system; it is a method of connecting computers together. So the first step in determining how the Internet works is to understand how your computer connects to others. As shown in Figure 3.23, the Internet has a hierarchy of service providers. Individuals pay a local **Internet service provider (ISP)** for access to the Internet. In turn, local ISPs pay an upstream network service provider (NSP) for access to their systems and features. Each connection must be made over a communication link. Local links are typically made over telephone wires, but cable companies also provide service over their coaxial lines. A few companies provide satellite connections. Some wireless providers exist, but availability, features, and pricing vary by location. Most ISPs also utilize phone company lines to connect to their NSP, but they lease dedicated, full-time lines that provide faster service. The largest NSPs (Tier 1) also provide backbone service. That is, they route communications over their own fiber-optic lines that are installed across the United States. Increasingly, NSPs are also phone companies. Some started as phone companies and expanded into the Internet; others started with the Internet and gradually offered voice services.

If you search for an ISP or look at the market for Internet connections, you will find several providers using different technologies. The Internet backbone has been relatively stable for several years, but the industry battles have been taking place in the **last mile** of the Internet—building connections directly to customers. The primary providers have been the local phone company and cable TV providers. They are important because they have existing lines into customer houses

Backbone Providers		**Network Service Provides**	
AT&T	Sprint	AT&T	Qwest
Level 3	MCI (UUNet)	Cable & Wireless	Sprint
Qwest		IBM	MCI
Phone Companies		**Cable Companies**	
Regional Bell operating companies (3 RBOCs)		Comcast	
Competitive local exchange carriers (CLECs)		Cox Communications	
		Regional	
Leading ISPs		**Satellite**	
AT&T		Direct Satellite	
Comcast		Wild Blue	
Verizon		Starband	
America Online			

Figure 3.24

Leading Internet providers. There are thousands of ISPs and cable companies. This list provides only some of the large companies in each category.

and businesses. Some technical differences exist between a digital subscriber line (DSL) and cable modem broadband connection, but ultimately customers make a choice based on availability and secondary considerations. For example, both tend to require that you purchase additional services, such as local phone service or cable-TV. The electricity industry has also been trying for several years to enter the ISP market—because they provide the third set of wires into households. However, transmitting data over noisy (and old) power lines has presented many technical obstacles. In 2006, the industry claimed to have found reasonably low-cost solutions, but it is not clear that the industry knows how to compete with the other providers. A few cities (and companies such as Google) have attempted to cover entire cities with WiFi Internet access—effectively offering free Internet connections. However, WiFi access points have a limited range of a few hundred feet. Although each access point is relatively inexpensive, it costs money to run cables to all of the points needed to cover a wide area.

A greater challenge to the dominance of the phone and cable companies is wireless access via fourth generation (4G) technologies. Sprint was the first to market with WiMax service provided by Clear. In late 2010, Verizon began offering their 4G cell phone service long-term-evolution (LTE). With few users and minimal congestion, early adopters reported speeds of 6-10 mbps for downloaded data. Upload data rates are substantially slower—even less than 100 kbps in many cases. But with these numbers, cell phone service might be usable as a full-time Internet connection. In fact, most companies sell WiFi add-on services that enable cell phones to operate as WiFi access points—providing Internet access to nearby computers and tablets.

You should understand the foundations of the Internet, because someone has to pay for each connection. Current pricing policies are to charge for the initial communication link and for the point-of-contact Internet service. For example,

General Data Centers	Specialty Hosting Companies
Equinix Cybercon Savis Telecity (Europe)	IBM AT&T EDS Thousands of small, regional providers

Figure 3.25

Leading data center providers. The general data centers simply provide space, electricity, cooling, and high-speed Internet connections. The specialty providers usually sign contracts to provide individual services.

an individual pays the phone company for the local phone line and pays the ISP for basic services. The ISP pays the phone company for the next link and pays the NSP for access services. Figure 3.24 lists some of the largest providers in each category. You can check with them for current prices and services.

The charging mechanism is similar for companies that wish to establish Web sites. The catch is that the costs are higher because the company needs faster communication services. The phone company charges more money for a faster link (e.g., $500 to $1,000) per month for a T1 line). The ISP also charges more money for the increased traffic because it needs faster equipment and faster connections to the NSP.

The Internet service connection business is completely based on economies of scale. The high-speed fiber networks (OC3 and OC12) can handle a vast number of transmissions, but they carry a high fixed cost. The backbone providers make money by selling smaller increments of bandwidth to the ISPs, which incorporate a sufficient profit. Many of the NSPs are backbone providers and increasingly they also offer ISP services.

Video and Heavy Use Complications

The problem of video is relatively complex. Perhaps you have a high-speed Internet connection and could receive high-quality video, say 1 mbps. But, think about what will happen if as few as 100 people try to view that same high-resolution video at the same time. The provider's server has to stream 100 copies of that video simultaneously—eating up 100 mbps. Hence the server (such as YouTube) needs to have at least an OC-3 Internet connection at $20,000 per month just to handle 100 simultaneous viewers. The same problem can arise if millions of visitors suddenly find a new Web site—as happens when a national radio or television program highlights a site. How is it possible to handle that much traffic?

One answer is that several specialty hosting companies have evolved with massive Internet connections. They have built data centers with multiple OC-48 (or better) Internet connections. You can contract with them to **co-locate (colo)** your servers in their buildings. For a monthly fee, they provide the space, power—with UPS, and air conditioning, along with the high-speed Internet connection. Figure 3.25 lists some of the biggest providers with multiple data centers. Larger companies might contract directly with these firms; smaller companies will work through intermediaries who already have co-location arrangements with these companies. The specialty providers, such as IBM, tend to be more expensive and negotiate individual contracts to perform specific jobs.

Even if you have a high-speed connection to the Internet, it is still difficult and expensive to handle thousands or millions of Web site visitors. A second answer is provided by Akamai, a company that was created specifically to reduce the load on Internet servers. The company has installed thousands of servers in over 70 countries that hold duplicate copies of your content. As shown in Figure 3.26, individual users automatically retrieve data from the nearest server. Users see only your site but get a faster response. More importantly, the load on your company's server and Internet connection are significantly reduced because the individual requests are handled at multiple points. Distributed network servers are the only way to handle the distribution of video content, because even a few hundred users could overwhelm the network capacity of a single server. Amazon provides a similar service that smaller companies can use for smaller tasks. This service (S3) is described in more detail in Chapter 7.

Broadcast streams are slightly different from simple video files. You might have an event that you want to broadcast to multiple users at the same time—such as a ball game. Your server could generate a single feed, and this feed could be duplicated and split to thousands of users. If handled correctly, your server and Internet connection generate only a single data stream. This stream is then split and duplicated at routers along the way to multiple clients. Some companies, notably MLB, are working on these technologies, but problems still arise when too many people access the streams.

Telephones and The Internet

All phone calls today are converted to digital signals and sent over a network. So, why should you pay for long distance phone calls? Of course, most people have cell phones which rarely charge extra for long-distance calls within a country. But, they still use plan minutes and generally require higher payments for international calls. You already pay a fixed monthly fee for Internet access, and the Internet is

Figure 3.26

Distributing content through Akamai. By distributing your content to servers at the "edge" of the Internet, customers retrieve data from multiple points, reducing the load on your server and Internet connection.

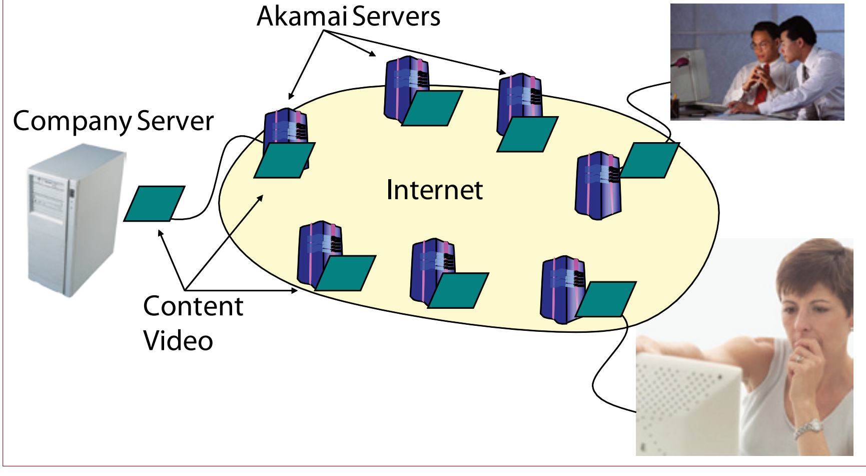

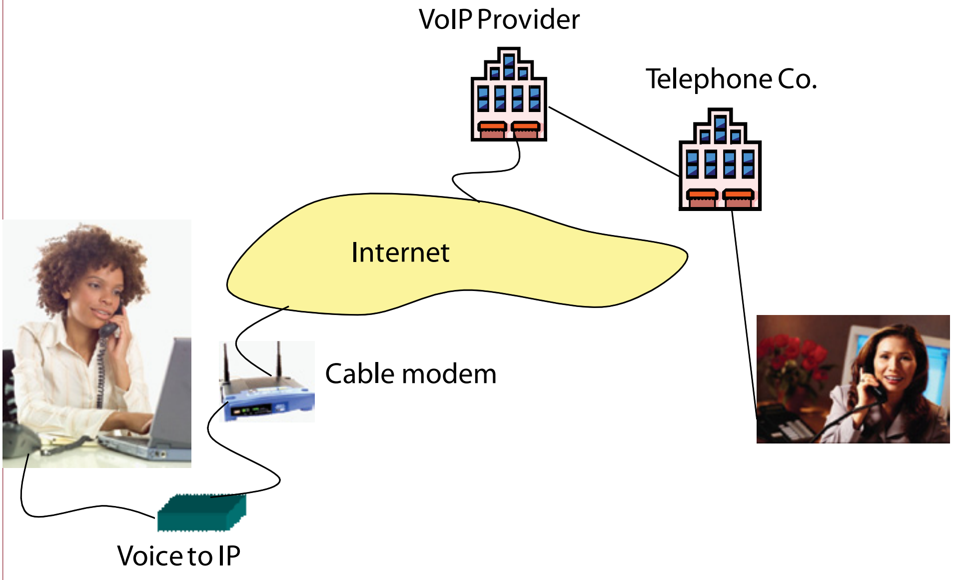

Figure 3.27

Voice over IP. The phone, or a separate device, digitizes the voice and sends it as IP packets on the Internet. A third-party company converts the Internet packets back to telephone packets and sends them across the phone network.

worldwide, so you could just use the Internet connection to handle voice calls. **Voice over IP (VoIP)** is a technology that converts your phone call directly to Internet packets and ships them to the destination. The technology is straightforward and relatively inexpensive. However, Figure 3.27 shows an important complication. To use your VoIP phone to connect to people with regular or cellular phone service, you have to contract with a company that can convert the Internet packets back into the plain old telephone system (POTS). The process is particularly effective for international calls because the Internet packets can be converted as close as possible to the destination—avoiding international and long distance charges. In the U.S., Vonage played that role until 2007 when Verizon, one of the phone companies, won a major patent-infringement case against Vonage. The phone companies were already pressured by the consumer switch to cell phones; losing long-distance revenue to third-parties like Vonage presents a major competitive threat. On the other hand, if the person you wish to call has Internet access, you can simply use a direct PC-to-PC connection and skip the phone company and third-party completely. If you can directly connect to the other computer, you can use voice or even video connections for no additional charge—anywhere in the world. Skype is a newer company that now dominates VoIP calls—partly because they offer free calls for computer-to-computer access, including video chat if you have Web cameras. Skype also offers relatively inexpensive monthly fees even for international calls to land lines and cell phones. You can even get your own phone number for your Skype/computer account.

Reality Bytes: More Data More Speed

Check any wireless indicator and it is obvious that many households have installed wireless routers. Wireless access is popular for almost all computers, both for homes and businesses. The initial wireless transfer rates were only 11 mbps on the 802.11b standard. Then the 802.11g standard was created to support 54 mbps transfer rates. After a few years and some arguments over technology, the 802.11n standard supported a top rate of 450 mbps. Although 802.11n routers are widely available in 2011, most operate at about a third of that peak speed or about 150 mbps. All transmission methods have overhead and do not actually transfer data at the peak rates. Typically, 10 percent or so is used for overhead such as packet identifiers and addresses. But, wireless is relatively inefficient, more often operating at 70 percent of the peak transfer rates. Rates are even lower when multiple devices are heavily using the same router. But, technology development continues. The 802.11ac standard defines transfer rates of 1 gigabit per second (gbps) and the 802.11ad standard will support a peak rate of 7 gbps. Both are expected to reach the market in 2012 or 2013. The new technologies will probably operate in the 60 GHz frequency band, compared to the 2.4 GHz band used by older devices. This range provides more data capacity, but the radio waves do not travel very far and are absorbed by walls, trees, and water. Why would anyone need such high transfer rates? The simple answer is: video. The 7 gbps speed is fast enough to transfer uncompressed high-definition video. Which means almost all devices could operate without wires. Wireless does have a problem with latency—the delay between sending a signal and receiving a response. Latency is not an issue with movies or basic displays, but it does cause problems for games or similar interactions, where responses need to be almost instantaneous.

Adapted from Stephen Shankland, "Coming to a Network Near You: Faster Wi-Fi," CNet News, May 31, 2011.

Internet Addresses

In any network, each computer must have a unique address. At the local level, network cards contain a **media access control (MAC)** address. This number is used by switches to route local traffic. On the Internet, every computer is assigned a logical IP number.

Currently, these addresses are 32-bit numbers, typically written as 4 bytes separated by dots. For example, your machine might be assigned 138.9.15.2 as an address. However, 32-bit numbers will identify a maximum of 4 billion machines. In practice, perhaps 80 percent of those numbers are usable because some values cannot be assigned and others are assigned in bulk to companies but not in use. Hence, an Internet committee designed a new numbering system consisting of 128 bits, which allows for several millions of numbers to be assigned to every person likely to live on the planet. The new system is known as **IPv6** and is being phased in gradually.

Servers need to be assigned IP addresses that do not change, or seldom change. Client computers can be assigned numbers from a local router or computer using the **dynamic host control protocol (DHCP)**. The DHCP server has a pool of address numbers and gives one to each computer as it connects to the network.

In most organizations, the process of assigning addresses is more complex. Large companies and universities were the first to encounter problems with the

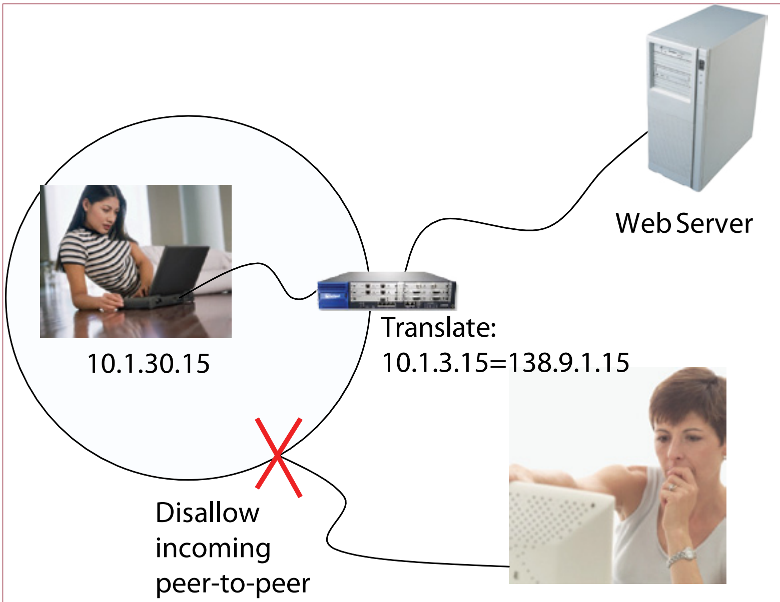

Figure 3.28

Network address translation. Internally, machines are assigned numbers that cannot be used outside the local network. Requests sent to external servers are temporarily assigned a real IP address by the router. Any external attempt to reach an internal computer is automatically blocked.

limited number of IPv4 addresses. Consequently, they purchased Internet routers that implemented **network address translation (NAT),** which is commonly used by businesses and even home users today. As shown in Figure 3.28, with NAT, your house (or office) might be assigned one Internet address that connects to a router. The router then assigns completely different numbers to all of the machines in the home or office. When a local computer requests a page or service from the Internet, the router tags the address so it can identify which machine made the request. All replies are returned to the same port, and the router directs the connection to the original machine. Essentially, the router translates or changes the internal addresses into the single external number. Client computers are assigned internal IP numbers that cannot be used outside the NAT. Typically they start with 10 or 192.168. Essentially, one real IP address is multiplied to work with several local devices.

NAT also creates an additional security level because no one outside this local network can directly access any of the computers. However, sometimes you actually want to connect to other computers. In particular, it causes complications if you want to use groupware or peer-to-peer systems across the NAT router. If you and the person you wish to contact are both running computers behind different NAT connections, you will not be able to make a direct connection. In this situation, you will need the use of a third-party server. Cisco WebEx and Microsoft LiveMeeting perform this role for online conferences. Skype performs this function for voice phone calls that use the Internet, and SlingPlayer does the same for watching television remotely. Both parties first establish a connection to the serv-

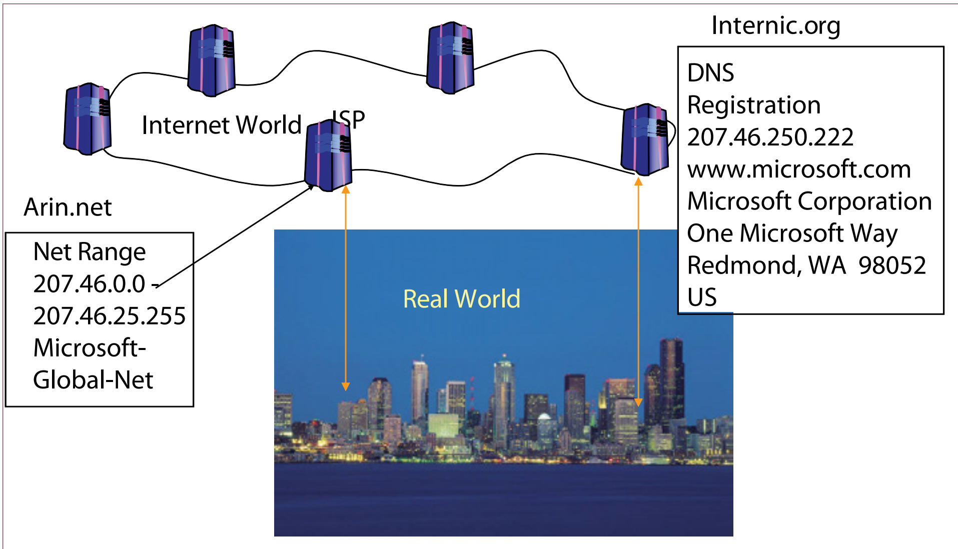

Figure 3.29

Domain Name Registration. Internet entities are anchored to the real world through the DNS registration and their ISP. But, some firms lie in their registration, and some ISPs are controlled by unethical people.

er. The server has a fixed IP address, grabs the currently-assigned IP addresses of the two parties and connects them to each other. Companies can also run in-house servers with fixed IP addresses to handle the initial connection. Using in-house computers is more secure because the communications do not have to go through a third party.

NAT also makes it considerably more difficult to track down individual users on the Internet. If you receive a message (spam) from someone, or want to find the true owner of a Web site, you begin with only the shared IP address. This address will point you to the company that hosts the original user. You can find the hosting company by entering the IP address at one of the registry sites, such as www.arin.net. But to find the specific person, the hosting company (or university) will have to provide you with detailed records of who was using the computer at the specific time. Most of these organizations will reveal the information only when presented with a court order (if they kept historical records). This process makes it more difficult and expensive to track down fraud and illegal activities on the Internet.

Domain Names

Numbers are difficult for most people to remember. So the Internet utilizes a system where a **domain name system (DNS)** server converts names to numbers. Anyone can apply for a name and pay a nominal fee (perhaps $20 per year) to use that name. Of course, names must be unique, so sometimes disagreements arise over popular names. Names are followed by a suffix such as .com, .edu, or .gov. Several newer suffixes were added in 2002 and 2003, but the .com and .net suffixes are still the most coveted. One interesting feature of the names is that each country is given a specific suffix. Some of the countries sell names using their unique suffix. For example, the small nation of Tuvalu allows companies to use its suffix (.tv) for a fee.

As a business, you will need to register your domain name to make it accessible to other users. Several companies act as registrars and you pay them to officially link your name to your IP address. In terms of data, the connection is ultimately performed with the DNS system, and you can either run your own DNS computer or pay another company to handle the updates. Figure 3.29 shows that DNS registration is a critical element to electronic commerce. Registration provides a tie between the Internet world and the real world. A second tie is provided by the connection of the firm to its ISP. In a legal situation, the ISP can be forced to reveal the true identity and address of its client. A third anchor is created when a company installs a digital security certificate. The certificate authority verifies the real-world identity of the client and binds it to the certificate.

You can use the Whois facilities of the domain registrars (e.g., NetworkSolutions, or Internic.org) to look up the registration data. You can use the Internet registries to identify ISPs based on the IP address. These registries are organized by geographical location and include ARIN (North America), RIPE_NCC (Europe, the Middle East and Central Asia), APNIC (Asia/Pacific), LACNIC (Latin America), and AfriNIC (Africa). However, some ethically-challenged people lie on their DNS registrations. A few even more ethically-challenged people operate ISP services specifically to support spammers and hackers. As a consumer, the answer is relatively easy: never, ever send money to a company that lies on its DNS registration!

If you want to run your own Web site, you will want to create and register your own domain name. The exact IP address will be given to you by your hosting ISP, but you have to register for the domain name that will be linked to the address. This registration information is kept in global **whois** databases. For example, you can use the networksolutions.com database to enter a Web site name and retrieve the registered owner of the site. When you register a site name, you are supposed to provide accurate information on your name, address, and phone number. However, at this point in time, the registrars have been weak at verifying the data. The resulting misleading information makes it difficult to track down spammers and other owners of fraudulent sites. In a sense, the Internet today is like the business world in the 1800s—any fly-by-night scammer can set up a business Web site. In the business world, this situation was partly solved by requiring businesses to obtain licenses and register with the government. If you have any doubts about the validity of a company, you should check to ensure that the business is legally registered with the state government. Most states have searchable online databases and provide free access through the office of the secretary of state.

Domain names can consist of several parts separated by dots. The names are read from right to left. The text on the far right is the top-level domain (TLD). The most familiar TLDs are the six that were created first: com, net, org, edu, gov, and mil. Several others have been added, such as info and biz, and new ones are proposed and debated each year. There are also hundreds of country codes (such as us, ca, uk, and tv). Around 2010, the Internet committees introduced the concept of generic TLDs (gTLD). In this case, companies or individuals will be able to purchase their own unique TLD (such as .ibm or .houses). The purchaser of a gTLD will act as its own registrar and be able to assign or sell domain names within that TLD. However, the initial cost (evaluation fee) for a gTLD is $185,000 and costs increase from that point; including the ability to run your own registry.

The second name is the name you must pay for and register with a DNS registrar. It usually is the name of a company or service, but could be almost anything.

Common examples include microsoft.com, google.com, pacific.edu, and sec.gov. Anything to the left of the registered domain can be created within your organization's network. Many times it is used to represent a department or specific server within the organization. However, there are no fixed rules so the names could represent almost anything. For example, you might create marketing.mycompany.com for the marketing department server. Why do you care? If you need to setup networks, you have to learn more of these details, and you have to design a naming convention that is easy to use. But, even as a typical user you should be familiar with the basic structure so you can recognize bad or dangerous domains. For example, accounts.citibank.com.xi92lai.293aafa.com is most definitely *not* an addressed owned by Citicorp. When you see that address in an e-mail message, you immediately know it is a fake—most likely created to steal your bank password.

For years, the Internet committees have realized that the domain name system is based on the English language and that many companies would like to create domain names in local languages. However, a name system based on multiple languages presents several technical challenges—notably, ensuring that names in different languages or alphabets are never reduced to an existing name in a different language. Another problem is that the DNS database was designed to hold only Latin-based characters. Most of these issues have been resolved and by 2011, several TLDs became available in non-English languages including Arabic and Chinese.

Internet Committees

The Internet works because of the adoption of standards. Everyone connecting a device to the Internet agrees to follow the standards. But, these standards need to evolve to handle new technologies. The standards are maintained and discussed by several Internet committees. The **Internet Corporation for Assigned Names and Numbers (ICANN),** the **Internet Engineering Task Force (IETF),** and the **Internet Assigned Numbers Authority (IANA)** are publicly run organizations in charge of establishing many of these standards. These organizations are heavily dependent on volunteers and rely on public comments to design new standards. The IETF is the most technical group and it deals with most of the routing and HTML standards. IANA controls the allocation of IP address segments, TLDs, and other assigned numbers. It is largely concerned with technical issues.

In some ways, ICANN is a more interesting group. It was not founded until 1998. It coordinates the Internet naming system. In actuality, it is the political committee where everyone gets to argue over various directions for the Internet. For example, it is the focus of discussions for whether new TLDs should be introduced, or when internationalized domain names should be allowed. Originally, these tasks were handled by U.S. organizations. ICANN is a not-for-profit corporation that is not directly controlled by the government, but it is founded in California. At various times, several nations have suggested that ICANN should be run as an international body—perhaps even within the United Nations. U.S. politicians have resisted these attempts to change control of ICANN.

Internet 2

Originally, the U.S. government funded much of the Internet design and development. By 1995, the U.S. government had discontinued almost all funding, and the Internet was largely financed and controlled by private organizations. From 1994, the commercial use of the Internet increased exponentially. In 1996, 34 university

Reality Bytes: Faster Speeds but Data Caps

In April 2011, AT&T announced that it would limit broadband customers to no more than 250 gigabytes of Internet data transfer per month. DSL users would be limited to 150 GB. The 250 GB value matches that imposed by Comcast. AT&T claimed that only 2 percent of its users would be affected by the caps, but that people in that group use up 20 percent of the network bandwidth. Data transfers above the caps will be priced at $10 for each 50 GB. Mark Siegel, spokesman for AT&T noted that the company's customers averaged 18 GB per month. He also claimed that the reason for the change was because customers said "that the people who use the most should pay more." The 250 GB cap could be reached by watching 109 hours of high-definition video from Netflix—or about 3.6 hours every day for a month. The comparison to video is important because customers are moving to online video and dropping TV services from AT&T and Comcast. Vince Vittore, a broadband analyst with the Yankee Group, noted that the data cap "isn't absolutely necessary. It's mostly a move to prevent customers from cutting off video services." Note that Cisco forecasted that video on-demand usage will double every 2.5 years. Verizon Wireless followed with data caps and substantially higher prices for heavy users in July 2011 (Web site).

Adapted from David Goldman, "AT&T Starts Capping Broadband," *CNN Online*, May 3, 2011.

participants decided that they needed faster connections (the number of participants expanded to 100 in 1999, 205 in 2003, and 381 in 2007 with 58 international participants). In early 2011, the list of paying participants included 260 organizations and 147 sponsored participants according to the lists on www.internet2.edu. With the support of the government and industry, they began creating Internet 2 (http://www.internet2.edu).

The two most important proposed features of Internet 2 are high-speed connections and quality-of-service provisions. The overall objective is to provide a transmission network that can support full-speed video and other high-bandwidth applications. To understand the change, consider that most existing "high-speed" Internet connections are in the range of 1 mbps to 50 mbps. The Internet 2 calls for gigabit connection points and a minimum connection of 155 mbps. The cost of connection depends on the connection speed, where a 1 gbps connection costs $250,000 a year.

A related, but more fundamental, change is the ability to specify a desired level-of-service quality. Currently, if traffic increases on the Internet, all communications slow down. This situation is annoying but not troublesome for simple tasks like sending e-mail or browsing a Web site. On the other hand, full video transfer requires a constant minimum level of transmission capacity. So participants need a mechanism to tell all components that a specific set of messages should take priority to receive a certain level of service. Some people have suggested that the system should enable participants to pay a fee to gain their desired levels of service, for each type of message. For example, basic e-mail messages would be free if there is no rush in delivering them. But to reserve a time slot for videoconferencing, participants would pay an additional fee. Then all of the Internet 2 components would give the video packets a higher priority. So far, there has been

Figure 3.30

Mobile commerce. Wireless connections offer new capabilities, such as e-mail and Internet browsers on cell phones.

no agreement on whether additional fees should be charged, or on how the quality-of-service issues can be resolved. Although the system is being designed for academic and government users, the industry participants (e.g., Cisco) ultimately intend to transfer any useful technologies to the commercial Internet. Businesses could find many uses for high-speed connections and service-quality guarantees. For starters, better video transfer may finally open the way for desktop videoconferencing to replace travel to meetings.

Wireless Networks and Mobile Commerce

Beginning with cell phones in the 1990s people and businesses have become fascinated with wireless communication. Wireless Internet and mobile commerce have the potential to revolutionize the Internet, business, and society. Technologically, wireless is different from the traditional Internet in only two ways: (1) the transmission medium is microwave radio, and (2) the client devices are smaller with smaller screens and less computing power. Yet wireless connections open thousands of new possibilities. The client devices can consist of anything from enhanced cell phones, such as the device shown in Figure 3.30, to laptops and digital tablets. Many of these devices are more powerful than some older computers. Most cell phones and tablets have built-in cameras, Web browsers, and their own programming or application capabilities. Touch screens have made it possible to enter basic text and notes on these portable devices.

To date, the wireless communication presents the greatest challenges. As the portable devices become more sophisticated, people want to use them for everything from simple text to voice to video. And video includes video calls as well as downloaded high-definition movies. The catch is that 3G cellular networks were designed to handle perhaps 2 mbps of data transmission. Even the early 4G sys-

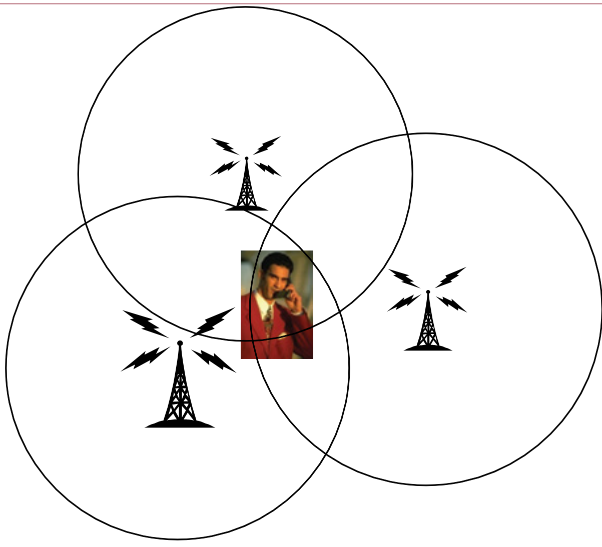

Figure 3.31

Location knowledge in m-commerce. Businesses could contact potential customers on learning their location through the wireless system.

tems provide download speeds of *only* 3 – 10 mbps. Upload speeds are substantially lower, often less than 1 mbps. Transmitting video on these wireless connections can be slow for individual users. Worse, the cell-phone space is shared by everyone in the area. Throw a few thousand people into the same location all trying to send and receive large images or video and the network slows to a crawl for everyone—as AT&T learned when it was the first to sell the Apple iPhone. Despite the increased money pumped into expanding the network, the growth of cell phones and data applications has expanded even faster. At some point, the cell-phone companies might be able to install enough capacity to handle the increasing traffic, but it is likely that data caps and usage charges will remain for heavy users.

Some of the more interesting possibilities of m-commerce come from the ability to use the phone to find and pay for virtually anything you might need. Rather than rely on credit cards or cash, your cell phone could record all of your transactions and give you complete instantaneous access to your data. But an incredibly reliable and secure system must exist before consumers will trust it. In 2011, several phone vendors began rolling out **near-field communication (NFC)** payment systems. NFC is a very-short range wireless communication method where either the cell phone itself or a tag placed within the cell phone can carry basic payment data. When waved next to a payment terminal, the ID is used to transfer money from your account to a vendor. Some credit cards—particularly those used in Europe use similar technology. The actual source of the funds (your bank account, credit card, cell phone bill, or PayPal) depends on the system being used. Several companies have created competing versions and it is not clear which one might win the largest market share.

Another interesting possibility for m-commerce comes from the ability to use location to identify potential customers. Figure 3.31 shows that it is relatively easy to identify a person's location based on the way the cellular system works using signal strength and triangulation. Most cell phones also have built-in GPS receivers and can locate your position down to a few feet. Currently, people can turn off most of the location broadcast data on their cell phones, but it is available for 911 emergency calls.

Think about the business opportunities that this system could provide, even from the perspective of a consumer. As you walk into a mall, you could enter your shopping list into a browser on your phone, which would contact all of the local stores for prices and availability, and then provide a local map and directions to each item. You could enter notes and make your selection through the browser, then press a couple of keys and purchase the items. Most cell phones and search engines provide a method to find companies based on your current location—such as the closest gas stations or restaurants that carry specific types of food. Similarly, some companies use a system that automatically tracks your location and notifies your friends (www.loopt.com). A few companies such as www.foursquare.com have built a system to encourage people to *check-in* to report their specific location in commercial stores.

Businesses have even more opportunities to collect data on customers, such as sending advertising to your phone and tracking the number of browsers versus buyers. But marketing systems have greater knowledge of exactly what products and features each customer wants. For instance, the systems could offer a mini-survey to each customer who chooses to buy a competitor's product to find ways to improve products. Of course, the privacy aspects are interesting. They are examined in Chapter 14.

Cloud Computing

Are personal computers necessary anymore? As network speeds increase, new computing methods are being created and rolled out to businesses and individuals. The generic term cloud computing is often used to describe applications where the server, and most of the software, are run on computers stored on Internet servers—usually leased from a third party. Figure 3.32 shows the basic concept. The main data and applications are running on servers located on the Internet. The results, whether data, images, or even video, are transmitted and displayed on client computers, laptops, or even smart phones and tablets. The key is that the main applications and data are centralized. Ideally, the user can access the data and applications using any device from almost any location. For these applications, the user needs only a simple client display device, often only running a Web browser. The data and applications remain on the servers.

Scalability and reliability are important aspects to cloud computing. The applications and data are run on distributed servers connected by high-speed networks. If one server fails, the load is automatically picked up by the other servers. Similarly, new servers can be added at any time to extend the performance of the collection. Data is stored on storage area networks—also connected with high-speed networks. Most service providers also have multiple server clusters running in different physical locations around the world. Data is replicated across the systems so it is closer to any user on the Internet. Each local cluster has its own high-capacity Internet connections. Most sites have dual connections from different providers. If one network provider goes down or has a line cut, the system automatically routes data through the backup connection.

Cloud service providers also use virtual machines to lease space on their hardware to multiple customers. Multiple VMs can run on one physical server. Between the hardware, networks, and personnel, the cloud providers make money through economies of scale. Consider the network costs alone. A T1 line providing 1.544 mbps might cost $400 a month to a single company. But, a giant provider can get an OC3 line for perhaps $20,000 a month that provides 155 mbps; or about half the cost per megabit per second. Plus, the higher-capacity line can be allocated across even more customers because most of them will have minimal traffic. In a large data center, the monthly network and power costs constitute a large portion of the expenses. A handful of employees with the appropriate software can monitor hundreds or thousands of servers, so personnel costs are relatively low. Selling capacity to multiple businesses helps cover the monthly costs. The more services they can offer, the more money they can attract.

But, it is difficult to answer the question of whether everything can be shifted to the cloud. Even if you completely trust the servers and the network, is there some data that you might wish to keep on your own computers? Or, is it always going to be easier to pay someone else to handle the software applications and data? These are the main questions that need to be answered.

Global Telecommunications

What problems are you likely to encounter if you need to connect to a supplier in a different country? Business firms are becoming more dependent on international markets. This internationalization increases the demands on the telecommunications system. The international transmission of data is becoming part of the daily business routine. A manufacturing company may have factories in several different countries, with the headquarters located in yet another country. Supplies have to be sent to factories. Finished and intermediate products have to be tracked. Customer demands have to be followed in each country. Quality control and warranty repair information have to be followed from the supplier through the factory and out to the customers. Financial investments have to be tracked on stock markets in many countries. Different accounting and payroll rules have to be provided for each country. Basic accounting and financial data have to be available to management at any time of day, in any part of the organization.

Creating networks across international boundaries creates many problems. Some of the complications are technical, some are political or legal, and others are cultural.

Technical Problems

The biggest technical complication is that each country may have its own telecommunications standards. Today, most people rely on cell phones—particularly in smaller, developing countries. Just as in the U.S., several different standards exist for cell phones. Consequently, a cell phone from one country might or might not work in a different area. Many of the European nations use a version of the GSM standard, so phones on those systems will generally work in multiple areas. However, some American travelers have learned the hard way that between roaming and international charges, using a phone in a different nation can be extremely expensive. Many European travelers purchase new SIM cards (subscriber identity module) to work in each country. A new SIM card essentially enables you to make local phone calls without roaming charges, but international calls can still be expensive. To handle those, frequent travelers can use VoIP systems such as Skype and find a WiFi connection to make low-cost international calls over the Internet.

One possible way to avoid the public telecommunications hassles is to use satellite phones. Some of these phones (notably Iridium) will work almost anywhere on the planet. Unfortunately, they do not work well inside large buildings because the signal gets blocked. Still, they can provide international reach even in remote areas and support small data transfers as well as voice. The traveler needs to purchase a special phone and sign up for a package of minutes. The per-minute cost (in 2010) averages about $0.90 – 1.30, which might seem high to people used to unlimited calls. However, even in the U.S., some cell phone plans charge $0.45 a minute; and international roaming charges in Europe can often exceed $1 a minute and some people have been charged over $3 a minute.

Fortunately, the Internet has made it easier to share data, images, and video regardless of location. So, although the U.S., Europe, and Asia use different television standards, the computers all handle the standard types of images and video used on the Internet. You might have to ensure that the recipient has an up-to-date computer to handle some newer types of video, but as time goes by, most computers should be acceptable. However, keep in mind that it might not be possible to plug your cell phone or laptop into a monitor in a different nation if you want a larger display. Still, you can usually transfer the data and use a local system.

Legal and Political Complications

Some important problems can be created when a firm wants to transmit information across national boundaries. These transfers are called **transborder data flows (TBDFs)**. The problem arises because the information has value to the sender. Some of the biggest complications today are political. Some nations exert strong control over the Internet. For example, in China the government physically owns and runs all of the Internet routers. That gives them complete control over traffic. The government commonly filters traffic for certain keywords and blocks sites that it deems to be against public policy. As another example, during the Egyptian turnover in January 2011, the existing president (Mubarak) appears to have ordered the shutdown all Internet routers in an attempt to prevent protestors from communicating. He also shut down all cell phone service for a couple of days. [Vijayan 2011] Ultimately, all he managed to do was demonstrate to all citizens that he had no respect for them and he was forced to step down.

Several other nations have expressed concerns over cell phones—particularly the Research-in-Motion Blackberry which promises strong encryption. In 2010, some Middle Eastern countries along with India threatened to ban Blackberry unless the company provided a way for government authorities to intercept civilian communications. They are not alone. The Obama administration also worked on a plan to obtain access to decrypted messages using wiretap authority [Rashid 2010]. The issues of tradeoffs between privacy and government monitoring are explored in more detail in Chapter 14. For now, it is important realize that technologies commonly used in one country might be banned in another country. It is important for travelers and businesspeople to keep up with the laws of the nations in which they operate.

Another important issue revolves around typical marketing data about customers. It is common for marketing departments to maintain huge databases. These databases contain customer names, addresses, phone numbers, estimated income levels, purchases, and other marketing information. Problems have arisen because the western European nations have much stricter laws concerning privacy than the United States. In most European nations, it is illegal to sell or trade customer data to other companies. It must also be stored in protected files that cannot be seen by unauthorized employees or outsiders. In most cases, it is the responsibility of the company to prove it is meeting the requirements of the law. In many cases, this requirement means that customer data must be maintained on computers within the original nation. Also, this data cannot then be transmitted to computers in other countries. As a result, the multinational company may be forced to maintain computer facilities in each of the nations in which it does business. It also needs to impose security conditions that prevent the raw data from being transmitted from these computers.

There is one more important political issue involving international computer centers. Many nations, especially the developing nations, change governments quite often, as well as abruptly. There are many nations where terrorist activities are prevalent. Oftentimes, large multinational companies present tempting targets. Because computer centers tend to be expensive, special security precautions need to be established in these countries. Probably the most important step is to keep the computer center away from public access. Several U.S. security specialists publish risk factors and suggested precautions for each country. They also provide security analysis and protection—for a fee.

A host of other political complications affect any multinational operation. For example, each nation has different employment laws, accounting rules, investment constraints, and local partnership requirements. Most of these can be surmounted, but they usually require the services of a local attorney.

Cultural Issues

All of the typical cultural issues can play a role in running multinational computer networks. The work habits of employees can vary in different nations. It may be difficult to obtain qualified service personnel at some times of day or night. These issues can be critical for computer networks that need to remain in operation 24 hours a day. In many nations, it is still considered inappropriate to have female managers when there are male subordinates. Collecting information may be exceedingly difficult or even culturally forbidden. In some countries, you will lose a customer if you try to obtain marketing data such as age and income.

In some nations, the connections between suppliers and customers are established by culture. For instance, in Japan, distribution of products is handled by only a few large firms. These companies have established relationships with the suppliers and retail outlets. In any country, it can be difficult for an outside firm to create a relationship between suppliers and customers. Trying to build computer networks with the various companies could cause severe repercussions. The established firms may think you are trying to steal their knowledge or information.

Summary

One of the most important concepts in MIS is the necessity of sharing data. But networks have gone beyond simple data to becoming a major part of the use of computers and accessing information. Many applications and business tools rely on almost-constant connection to the Internet and other devices. As a manager, you need to work in this world where you can search and connect to data and people from many different parts of a company. It also means that you need to be accessible and make your knowledge and skills available to other workers on the team.

Communication and networks require standards, and they need to evolve over time. One of the greatest shifts taking place now is the move to high-speed wireless communication. Physically, computer components have shrunk far enough so that even relatively small cell phones carry huge amounts of computer power. As people carry these devices everywhere, they want access to everything from anywhere in the world. The current challenge lies in building and paying for a wireless infrastructure that can handle the increased load created by everyone demanding high-speed access at the same time. The second challenge for business lies in determining how to use these connections and applications to improve management and increase profits.

Building networks requires specialized workers and skills, but managers need to keep up with some of the changing technology to make beneficial choices. As technologies and prices change, managers often need to determine when to implement new technologies and when to wait. Making these decisions often requires understanding the underlying technologies and trends. Managers also need to be aware of the fundamental economics, particularly the economies of scale involved in networks. Decisions need to be made about leasing space and network capacity versus building and managing your own connections. And these decisions need to be reviewed as prices change and new firms enter the industry.

Globally, telecommunications reach around the world. But capabilities and technologies vary by nation. And the U.S. is not always at the top of the world. Still, if a company wants to connect to customers and suppliers in other nations, managers need to be aware of potential conflicts. Technology differences remain, pricing issues can be huge barriers, and political restrictions can cause complications with many types of communications and applications. Most of the problems can be avoided or bypassed but it helps if you search for them and make plans ahead of time.

A Manager's View

The goal of telecommunications is to remove location as a factor in management and decisions. Technologies exist to enable you to connect to almost anyone, anywhere in the world to share data, voice, and video. However, bandwidth or the speed of transferring data varies considerably, and transferring large amounts of data can be expensive. Still, as speed and availability increase, new businesses and new ways of operating businesses are being developed. Opportunities exist to create expand and profit as the world becomes more interconnected.

Key Words

bandwidth hogs
blog
Bluetooth
Broadcasts
coaxial cable
co-locate (colo)
direct sequence spread spectrum (DSSS)
domain name system (DNS)
dynamic host control protocol (DHCP)
e-discovery
e-mail
extranets
file transfer protocol (FTP)
frequency division multiplexing (FDM)
grid computing
Groupware
hypertext markup language (HTML)
instant messaging (IM)
Internet service provider (ISP)
Internet Assigned Numbers Authority (IANA)
Internet Corporation for Assigned Names and Numbers (ICANN)
Internet Engineering Task Force (IETF)
Intranet
IPv6
last mile
latency
local area network (LAN)
media access control (MAC)
near-field communication (NFC)
network address translation (NAT)
network interface card (NIC)
packets
peer-to-peer networks.
podcasts
point of sale (POS)
portable document format (pdf)
really simple syndication (RSS)
scalability
server farm
social networking
standards
storage area network (SAN)
time division multiplexing (TDM)
transborder data flows (TBDFs)
transmission medium
twisted-pair
ultra-wideband (UWB)
voice over IP (VoIP)
whois
WiFi
Wikis
WiMax

Web Site References

Financial News and Quotes	
Big Charts	bigcharts.marketwatch.com
Bloomberg	www.bloomberg.com
Dun & Bradstreet	www.dnb.com
Global Financial Data	www.findata.com
Reuters	www.reuters.com
SEC Edgar	www.sec.gov
Motley Fool	www.fool.com
Wall Street Journal	wsj.com
Yahoo Finance	finance.yahoo.com
Discount Online Trading	
TDAmeritrade	www.tdameritrade.com
E*Trade	www.etrade.com
Fidelity	www.fidelity.com
Charles Schwab	www.schwab.com
ScotTrade	www.scottrade.com

Review Questions

1. Does it matter where data is stored anymore? Are networks reliable and available enough to assume data and processing can be handled anywhere?
2. Why do businesses use networks to share hardware?
3. How is cloud computing like renting software?
4. List the main components of a network.
5. List the types of transmission media that are available. How do they compare in transmission rates and cost?
6. What are the main advantages and drawbacks to wireless networks?
7. How do networks handle multiple users at the same time?
8. Why are standards so important in networks?
9. What organization is responsible for creating (or not) new top-level domains on the Internet? Should this group be adopted and controlled by the United Nations?
10. What type of network service do you need to run a Web server? Newer (LTE) phone speeds are reported to be over 10 mbps; would it be possible to use those to run a Web server?
11. Why are some cell phone networks so slow?
12. Why is the domain name system so important and what is the role of domain registrars?
13. What is the Internet2 and how will it affect businesses?
14. What problems arise with global telecommunications?
15. Describe the process by which a computer is assigned an IP address.
16. Some people and a few members of the Internet committees have suggested a substantial increase in the number of top-level Internet domains. Why is this a bad idea for businesses?

Exercises

1. Some cell phones have the ability to function as a Wi-Fi to connect other devices to the Internet using the cell phone network. Most tablets have an optional cell phone connection. Choose a tablet and a cell phone from the same service. Compare the costs of paying to connect the tablet directly to the cell network versus paying for the Wi-Fi connector and connecting through a cell phone.
2. Find at least one method to connect a printer to your home (or school) network so that it can be shared by all users on the network. How would you limit access to specific people, or track the number of pages printed by each person?

3. Using the Internet, find at least two software packages that will back up data across a LAN. Briefly explain how the software functions and what components need to be installed. Estimate the price of the software for a network of four servers and 100 clients.
4. Research to find the service that generates the most amount of traffic on the Internet today.
5. Design your "perfect" wireless communication device. Possibly draw it, but identify the features that you would want. Rank the features by importance (descending order) and write a brief reason for your top three selections.
6. Assume you want to start an e-business. What steps do you have to take to obtain and establish a domain name for the business? How much will it cost?
7. Choose two existing Web site names. To make it more interesting, one of them should be from a spam site. Use the Internet resources to obtain as much information about the company and owner as you can. Note, you should also check the office of the appropriate secretary of state. [Answer depends on companies, but at least provide the output from www.arin.net (or www.apnic.net) and www.networksolutions.com (or www.internic.org)]
8. Check vendor advertising and identify the costs and speed of the following services:

   a. DSL from the phone company.
   b. Cable modem.
   c. Cell phone Internet access.
   d. Dial-up Internet access.
   e. ClearWire.
   f. Wireless access at a local hot spot.

9. Identify the cost of running a server on Amazon's cloud system using an entry-level Web server running Microsoft Windows Server. Assume the site you create will contain about 5 GB of data and handle about 500 gigabytes of data transfer (outbound) in a month.
10. Estimate how long it would take to transmit the following information.

    a. A 5 megabyte image file over a 3G cell phone (70 kbps).
    b. An e-mail message with a 120 megabyte video clip attachment over a 5 mbps cable-modem connection.
    c. A 1 gigabyte data folder over a 100 mbps local area network connection.
    d. A 1 gigabyte data folder onto an SSD drive that can write 270 MB/second.
    e. If possible, test the above examples with real data.
    f. Why would the actual transfer times be slower than the estimated numbers?

11. Identify at least three major features provided by HTML 5.

Technology Toolbox

12. Create three Web pages that are linked. Include at least one image.
13. Create a style sheet for at least two Web pages and demonstrate how the look of the site can be changed by altering the style sheet.
14. Create a simple Web page using Microsoft Word and save it. Create a similar page using either straight HTML or an HTML editor such as FrontPage. Compare the two files and comment on the differences.
15. Find three Web hosting companies and identify the file transfer methods supported by each company. Does the company charge extra for some methods?
16. Briefly describe the purpose of server programming for Web pages and how it differs from client-side scripting.

Technology Toolbox

17. How are files uploaded to YouTube?
18. How are files uploaded to the Amazon S3 service?
19. Find a commercial Web hosting company and list the methods that it supports to upload files.

Teamwork

20. For each member in the team identify (a) how many e-mail messages they send a day, (b) the number of text messages sent per day, (c) the number of phone calls made per day. Combine the results from all team members and comment on any patterns. Why would people send more cell phone text messages than e-mails?
21. As a group answer the following two questions. Is it possible to live with only one wireless device? Can it be standardized within the group?
22. You have an American cell phone and will be traveling to Europe for two months. You want to have cell phone service so that your friends and coworkers can reach you in an emergency. You also want to be able to call people and places in Europe while you are there. Each team member should find a different way to provide this service. Share your results, compare the costs and benefits and choose a plan.
23. Each person should find a software package that is used to create Web pages. Summarize the basic features, ease of use, target market, and its price. Combine the data and recommend a package that could be used by a marketing department to design basic Web pages.
24. If a network is available with Microsoft Outlook (or similar package) on each client, add each team member to the Contacts list. Each person should enter a few items in a personal calendar to block out some times for one week. Then share the calendars and use the system to schedule a meeting.

25. Interview managers, friends, or family members who have jobs and identify where their primary files are stored. Are they stored on local computers, on a department file server, or at a more distant location? Who is responsible for backup? What is the process for recovering lost files? Combine each result into a document. Create a table and a chart to summarize the results. Write a short paragraph suggesting improvements.

26. If you have access to computers running Windows 7 or above, have each team member determine the IPv6 address of a machine and compare the results. Hint: You can always use the command line program ipconfig.

Rolling Thunder Database

27. Design a network for the Rolling Thunder Bicycle Company. Identify who will need access to the network; how many workstations you will need (and where to place them); the data, input forms, and reports users will need. Using the existing data, estimate the storage requirements and transmission needs. Specify how changes and growth will affect the type of network needed.

28. Describe how the Internet could be used to increase sales at Rolling Thunder Bicycles.

29. Rolling Thunder Bicycles wants to expand international sales. What changes would need to be made to the application? What problems would you expect to encounter, and how would you overcome these potential problems?

30. The manager of Rolling Thunder Bicycles wants to access the database from home. Describe some options of how this could be done. What potential problems might arise?

31. If you have access to laptops and a wireless network, test the database running over a wireless connection. If five or six people want to use the system with wireless devices, what problems might arise? Do some research: Is there a system that can help?

Additional Reading

"The capacity of optical fibre," *The Economist*, June 28, 2001. [Theoretical capacity of fiber-optic cables.]

Federal Communications Commission, *Broadband Performance*, OBI Technical Paper No. 4, 2010. http://www.fcc.gov/Daily_Releases/Daily_Business/2010/db0813/DOC-300902A1.pdf. [FCC evaluation of U.S. broadband usage and speeds, suggesting that actual consumer speeds are about half the advertised rates.]

Jeff Hecht, "Ultrafast Fibre Optics Set New Speed Record," *New Scientist*, April 29, 2011. [Two researchers show demonstration speeds of 100 terabits per second in fiber.]

"The Internet, untethered," *The Economist*, October 13, 2001. [Special section on wireless issues with examples from Japan.]

Kahn, J. M. and K.P. Ho. "Communications technology: A bottleneck for optical fibres." *Nature*, June 28, 2001, pp. 1007-1009. [Theoretical capacity of fiber-optic cables.]

Kahn, Robert E. "The Role of Government in the Evolution of the Internet, *Communications of the ACM*, August 1994, p. 15-19. [Early days of the Internet.]

Kurose, James F and Keith W. Ross, *Computer Networking: A top-down approach featuring the Internet*, 2001, Addison Wesley Longman: Boston. [Technical details on networking and TCP/IP.]

Mehta, Stephanie, "Bend it Like Corning," *Fortune*, August 6, 2007. [New fiber optic cable from Corning that can handle tight bends.]

Fahmida Y. Rashid, "Obama Wants to Wiretap the Internet," *eWeek*, September 27, 2010. [Several countries and the U.S. lean on Blackberry to provide decryption keys so police can listen in on calls and messages.]

Jaikumar Vijayan, "Egypt's 'Net Blockage an 'Armageddon Approach'," *Computerworld*, January 28, 2011. [Initial report on Egypt's shutdown of the Internet and cell phones.]

Wiggins, R. "How the Internet Works," *Internet World*, 8(10), 1996. [Basic explanation of terms and connection points.]

Waltner, C. Meet your connection, *Fortune Technology Review*, Summer 2001, 59-66. [Discussion of business Internet connection costs and options.]

Cases: Wholesale Suppliers

The Industry

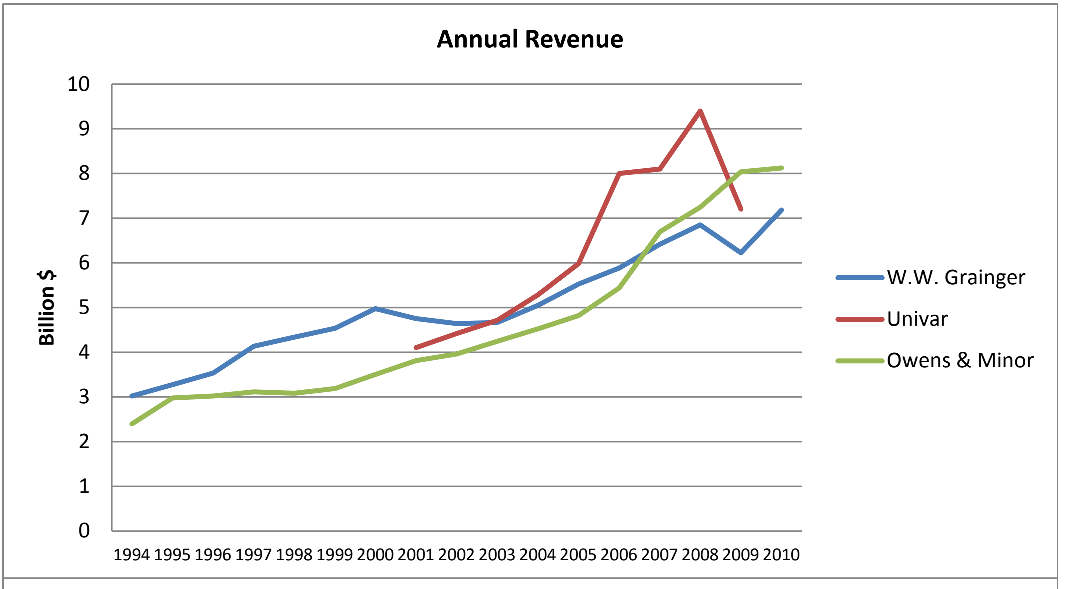

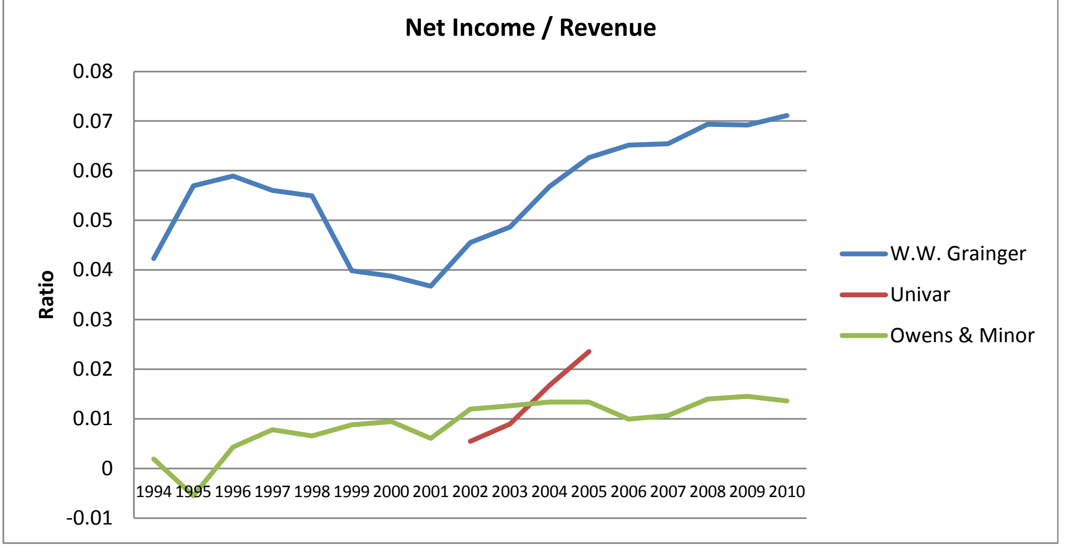

Wholesale suppliers are the original business-to-business providers. In many ways, they are like any other manufacturer or retailer. The difference is that their primary customers are other businesses. Selling to other businesses often means that the company supplies bulk items, but some wholesalers specialize in small, hard-to-find items. Some wholesalers provide commodities, such as raw chemicals. That raises one of the big challenges: differentiating your company from the competitors. What makes one company's products different from the others? In most cases, it is not the product that matters, but the customer service. Sometimes location is a key aspect to providing service, so small firms can survive by carving a local niche. But the big firms often have an advantage in economies of scale, driving costs down by serving large geographical regions through distribution centers. Some of the large commodity providers are global in scope. But size raises a new set of problems: managing a global company that has grown through multiple mergers. Communicating within the company is difficult enough. Communicating a consistent message to customers can be exceedingly hard.

Payments are another challenging issue for wholesalers. Manufacturers often squeeze vendors by stretching payments to the future. Small firms often pay the price in these cases, because it puts pressure on their cash flow when they do not receive payment for 90 days. Technology cannot always solve this problem because it is a power balance between vendor and customer. However, it can at least track the payment history and make the entire process transparent to both sides.

B2B Ecommerce

Communication is the entire purpose of electronic networks. Internal networks are useful for collaboration and coordinating employees. For customers, the big leap in the last decades has been e-commerce—particularly business-to-business (B2B) sites. B2B sites can be set up by one company, focused on providing a single interface for its customers. Alternatively, it might be a site shared across the industry by multiple suppliers and many customers. Several auction-based Web sites were created around 2000 to increase competition and make it easier for customers to find and purchase various items. Many of these sites later failed; however, a few remain in some key industries. Both sellers and purchasers have to decide what type of B2B site works best for each type of product.

A shared site can be run by a neutral party, with costs shared by everyone. Hence, it can be easier for smaller firms to participate. But auction-based sites tend to drive down the prices of commodities, so sellers might not want to use them. From the perspective of the buyers, they might prefer a tighter relationship with one or two suppliers. If you compete only on the basis of price, suppliers might not be responsive when crises arise. Ultimately, one of the big questions for buyers is how much of the logistics they want to handle themselves, versus how much they want to outsource to the vendors.

Logistics

Distribution is a key factor in the wholesale industry. Suppliers need to get products to the customers. Sometimes the products are bulky. In the case of chemicals, they might be volatile and difficult to transport. But that is the point of customer service by the supplier. On the other hand, logistics involves more than just transportation. It includes tracking orders and payments, timing the deliveries, finding rare products, tracking shipments, and recording everything in a format that can be shared. Elemica is a multi-vendor Web site for the chemical industry. Its Web site notes that an average international shipment requires 6 participants, 20 documents, and between 4 and 12 weeks of time. One of Elemica's selling points is a software tool to track the entire process and integrate the data between the supplier and purchaser enterprise resource planning (ERP) systems.

Wholesalers also have to make difficult decisions about inventory levels. Increasingly, through just-in-time ordering systems, manufacturers have pushed inventories back to the suppliers. Consequently, suppliers have to continually evaluate and forecast customer demands, compare them to world supply levels, and determine how much to stock of each item in every location. Linking to customer databases can be a useful step, because it enables suppliers to quickly spot trends and make more accurate forecasts (Babcock 2003).

Additional Reading

Babcock, Charles, "Distribution: Distributors Look To 'Crystal Balls'," *Information Week*, September 22, 2003.

Elemica Web site: http://www.elemica.com

Case: Owens & Minor

Founded in 1882, Owens & Minor (ticker: OMI) is a leading distributor of medical and surgical supplies to hospitals in the United States. It is also a health care supply chain management company that can run the entire logistics side for hospitals. In 2006, the company distributed almost 120,000 products from over 1,000 suppliers. It had 41 distribution centers and a total of 3,200 employees distributing products to about 4,000 customers (annual report 2006). By 2010, the number of customers had increased to 4,400 supported by 52 distribution centers and 4,800 employees (annual report 2010).

The company purchases medical items in bulk from suppliers and stores them in its warehouses, which are close to customers. Most warehouses deliver within a 200-mile radius. Customers can order through a variety of systems, including a stockless, automated system where Owens & Minor employees deliver individual items as needed. Although the company purchases items from almost all of the medical manufacturers, about 16 percent of the net sales in 2003 consisted of products from Johnson & Johnson and 14 percent from subsidiaries of Tyco International, which includes the Kendall Company (annual report). Its major nationwide competitors are Cardinal Health (whose roots lie with American Hospital Supply Corporation), and McKesson. In 2006, sales revenue reached $5.5 billion with a net income of $49 million, which represent a 16 percent increase in revenue but a 25 percent decline in profit from 2005. But the move into new offices accounts for some of the increased costs (2006 Annual Report).

Delivering Supplies

Until 1998, O&M was simply delivering supplies to hospitals. In that year, the company was hit with the loss of substantial revenue when HCA Inc., the biggest health care provider in the United States, decided to take its business to another company—immediately reducing O&M's revenue by 11 percent. The company decided that the answer was to go beyond simply selling and delivering products. It needed to provide entire supply chain management tools to hospitals. The company estimates that hospitals spend an additional 40 percent of their supply costs just managing the logistics, tracking orders, and restocking. David Guzman, hired as CIO in 2000, notes that "there's only so much blood you can squeeze from a turnip in terms of the product price. But there's a lot you can do with the logistics costs" (Kontzer 2003). Consequently, O&M decided to extend itself deeper into the supply chain on both the hospital and manufacturer ends. CEO Bilmer Minor III notes that "our business has been built on supply-chain services and customer service. We realized our difference would be based on providing information" (Kolbasuk 2001).

Logistics Services

As a distributor, O&M understands the importance of networks for connecting to customers and suppliers. The company survives on price differentials—buying products at discounts from manufacturers and distributing them to hospitals. With fierce competition, it has to use technology to hold down costs to make a profit—about 1 percent on sales (Stahl 2004).

One challenge that O&M faces is that its customers (hospitals, nurses, and physicians) are several years behind in technology. Craig Smith, president of O&M, estimates they are as much as 20 years behind. Smith believes that radio frequency identification (RFID) chips can make a substantial difference by helping hospitals track everything they own from beds to drugs. "We're headstrong into RFID. It will have a significant impact on our business" (Stahl 2004).

Beyond the simple warehouse and transportation services, O&M has created several tools to help hospitals reduce their costs. The OMSolutions division is a professional services unit that handles consulting and outsourcing to help hospitals reduce costs. The team handles everything from stockroom and process redesign within a hospital to outsourced management of supplies (annual report). The CostTrack system is another tool to help customers analyze their costs. Using activity-based costing, the system shows customers exactly how much each step of the purchasing process is costing. It enables hospitals to then choose exactly which parts of the process they want to outsource to O&M. In 2003, 32 percent of O&M sales were generated through the CostTrack system (annual report).

One big problem that hospitals face is that they buy supplies from many sources. As much as 50 percent of the purchases are direct from manufacturers. Other items might be purchased from multiple suppliers and distributors. So O&M developed the Wisdom2 analytical tool. It collects all purchasing data from a hospital and makes it available as a single dataset. It tracks purchases from competitor suppliers as well as directly from the manufacturers. Judy Springfield, director of corporate standardization and contracting at Baylor Health Care System (a nine-hospital network), says that "we've just been dying for data" (Kontzer 2003). The Wisdom2 system provides the detailed data she wants, across all of the purchases, and even puts it into a spreadsheet. Wisdom and Wisdom2 are also useful to O&M directly.

Ultimately, Owens & Minor would like to simplify the ordering process even more—to the point of automated replenishment. If the system can forecast usage rates, O&M employees could deliver new items to the hospital and track the entire process, without needing a separate order from the hospital. Guzman thinks the process is feasible, "You'd be surprised how remarkably predictable demand is in the healthcare system. And I don't mean in general. I mean system by system. Massachusetts General is different from Stanford Hospital, but Stanford in its own right is predictable, as is Massachusetts General" (Kontzer 2003).

O&M also uses technology and networks to improve its operations. Remember that margins are extremely tight, and saving costs is critical. A new warehouse-management system in 2001 that uses wireless technology to guide workers increased productivity by 20 percent (Kolbasuk 2001). Notice that O&M manages to record over $4 billion of sales a year with only about 3,200 total employees.

New Systems

To provide new services, Owens & Minor needed new systems. But before building new systems, the company first needed to consolidate its information technology team. In 2002, the company canceled an outsourcing contract with IBM and expanded an arrangement it had with Perot Systems. The goal was to consolidate data centers to save costs, but also to begin creating new technology systems (Vijayan 2002). The company's main goal is to redesign the legacy systems so they use Web services to integrate all of the components. Ultimately, every information system the company has will need to connect to the OMDirect Web portal. Integrating the systems is still difficult and requires a new system architecture that uses XML to transfer the data. Guzman notes that "you can't go out and buy Web services. It's clear you have to be the one to build [them]" (Murphy and Bacheldor 2003), so the process is scheduled to take three years.

In 2006, Owens & Minor extended its outsourcing agreement with Perot Systems, committing to migrating their mainframe systems through 2014. Perot Systems also helped transfer and integrate data from the purchase of McKesson. The most important step was linking the O&M systems with those of the new customers (2006 Annual Report). The purchase of McKesson reduces the number of national providers to two: O&M and Cardinal Health. With the purchase of Access Diabetic Supply, LLC in 2005, O&M has become the third-largest mail-order supplier of diabetes supplies; but faced problems with not getting paid by many end-customers. The company prides itself on being an early adopter of computer technology—largely used to reduce costs and improve productivity. The company distributes 180,000 products from 1,200 suppliers, so tracking inventory and purchases are important. They also provide software management tools for operating rooms across the nation, including QSight, SurgiTrack, Wisdom, and PANDAC, an operating room inventory management program. O&M uses these tools to control and reduce inventory costs. Wisdom is a business intelligence tool that helps customers analyze their supply usage. Sales to the top 10 customers constituted 20 percent of revenue, so providing support and linkages are critical to the success of O&M.

In 2010, O&M needed to move its custom ERP system onto new Windows servers. The original system was a custom COBOL application running on mainframe hardware, consisting of about 10 million lines of code. The company decided that the complex business logic built into the code was too valuable to lose and too expensive to recreate. So it used Micro Focus COBOL and tweaked the code to run on Windows-based hardware (Thibodeau 2010). In addition to saving money by not buying and customizing a new ERP system, the company cut its hardware costs in half.

The company had been looking at the project for several years and struggled to find a way to transfer the system without "bringing the business to a grinding halt for six to nine months," as CEO Craig Smith put it (Nash 2010). Transferring to a new ERP system ultimately saved $100 million. The new system simplifies usability. When a customer calls in, sales representatives can see all of the relevant data on one screen instead of searching 11 or 12 screens to get an answer.

Security

With any network system, and particularly with Internet-based systems, security is a critical factor—particularly in the health care industry. At one level, security today is relatively straightforward: identify each resource and user, and then assign the appropriate permissions to users. Except that with thousands of users and hundreds of applications, it is expensive to manage the tasks. Owens & Minor has more than 12 administrators who are dedicated to managing, adding, and deleting employee and customer access rights (Hulme 2003). And, the company has to be extremely careful to monitor access rights to ensure they are correct and that they are updated as employees change jobs or leave the company. To simplify the process, O&M is consolidating the identity databases onto a single centralized Microsoft Active Directory repository. Active Directory (AD) runs on Microsoft servers and holds user credentials. It can be accessed by a variety of applications and services. Users essentially login to AD, and the system authenticates the user to other applications. O&M is also using IdentityMinder from Netegrity Inc. to transfer the user rights into the 20 or 30 applications accessed by each user. Paul Higday, chief architect for O&M, notes that "what this will allow us to do is set up a user based on what they're allowed to access with a single click instead of having to manually add each account" (Hulme 2003).

Questions

1. How does Owens & Minor use networks and information systems to reduce costs?
2. Given the innovations by American Hospital Supply in the 1980s, why are hospitals not even more integrated into the supply chain?
3. How are logistics services different from simply delivering supplies?
4. How will RFID affect the use and purchasing of hospital supplies? Will the technology be widely accepted?

Additional Reading

Ante, Spencer E., "Owens & Minor," *BusinessWeek*, November 24, 2003, p. 92.

Hulme, George V., "Identity Checkpoint," *Information Week*, January 20, 2003.

Kolbasuk McGee, Marianne, "Supply Chain's Missing Link," *Information Week*, September 17, 2001.

Kontzer, Tony, "Owens & Minor Takes Supply Chain Deeper," *Information Week*, September 22, 2003.

Murphy, Chris and Beth Bacheldor, "Tying It All Together," *Information Week*, March 17, 2003.

Nash, Kim S., "What Getting IT Right Means to a CEO," *CIO*, May 18, 2010.

Stahl, Stephanie, "Prescription For Change," *Information Week*, February 2, 2004.

Thibodeau, Patrick, "One Firm's Story: The Mainframe Goes, but Cobol Stays Behind," *CIO*, April 20, 2010.

Vijayan, Jaikumar, "Owens & Minor Axes IBM For Perot," *Computerworld*, July 29, 2002.

Case: W.W. Grainger

Founded in 1927 in Chicago, W.W. Grainger (ticker: GWW) is in an interesting industry. The company supplies products to other businesses for maintenance, repair, and operations (MRO). The products include janitorial supplies, light bulbs, and bolts. But the repair side can include any number of unique parts or tools used to repair equipment, from gaskets to pumps. The company estimates the U.S. market for MRO to be $100 billion a year, giving Grainger about a 5 percent share of the market—the largest of any competitor. About 90 percent of the market is filled by local and regional suppliers, and Grainger's strategy is to capture market share from these fragmented rivals. In 2003, the company had 575 sales outlets backed by 17 distribution centers and carried more than 500,000 different products (2003 Annual Report). In 2006, Grainger had sales of $5.9 billion with net earnings of $383 million. In 2004, Grainger began installing more outlet stores in major metropolitan areas—in a goal to make it more convenient for people to purchase products. The company expected to spend $50 to $80 million on the project, as well as $10 to $15 million on information technology. (2006 Annual Report). In 2010, Grainger had sales of $7.2 billion with 607 branches, including new outlets in Asia in Latin America. By 2011, Grainger carried 354,000 different products and James Ryan, the CEO, has a goal of reaching the "sweet spot" of 450,000 – 500,000 different products (2010 Annual Report).

As a distributor and supplier, Grainger makes money on the price differential by purchasing items in bulk and selling them at a markup to customers. But the markup is small. However, Grainger is expanding into the logistics market. The company estimates that about 40 percent of the total procurement cost comes from trying to locate and purchase MRO products. One of Grainger's competitive strengths is the number of products it carries and makes available almost immediately at any of its stores. Customers need to know that Grainger will carry the parts and supplies they need, regardless of the manufacturer. Parts that are not stocked in the local branches can usually be shipped for next-day delivery out of the distribution centers. To strengthen its local presence, Grainger is adding more showrooms and warehouse capacity across the United States. The goal is to have an outlet within 20 minutes of most businesses and institutions (annual report).

Of course, it does not make sense to carry products that rarely sell, so Grainger has to carefully choose the products that it holds in inventory. The company also works closely with manufacturers to reduce cycle times so that they can quickly refill warehouses or custom-order products if needed. Grainger manages the inventory problem by using multiple channels. Its 3,700-page CD-ROM catalog describes 82,000 products. Its Web site offers more than 200,000 products. When customers need more specialized repair parts, Grainger's Parts division can obtain more than 2.5 million parts from its suppliers. The FindMRO division is even more specialized and can track down over 5 million facilities maintenance products (annual report). George Rimnac, vice president and chief technologist at Grainger, observes that "many of the products we sell are products our customers didn't know they needed until today" (Pratt 2001).

The FindMRO division demonstrates a key element of Grainger's success: its knowledge base. Between its systems and workers, the company is extremely successful at finding parts, including chemicals to fuel Universal Studios' fog machines. This knowledge has a huge value to Grainger. A typical online order is worth $250 with small profit margins. Orders on FindMRO average $1,200 and 80 percent of the items are shipped directly from the manufacturer—giving Grainger an immediate profit (Sviokla 2001).

Grainger has been adversely affected by the transfer of manufacturing from the United States to other nations. The company has countered this shift somewhat by focusing on health care and government agencies—which cannot easily move offshore. At the same time, the company is expanding its purchases from lower-cost offshore suppliers—particularly for house brands ("Top Distributor" June 2003).

Logistics Services

Selling to businesses and government agencies is a key component to growth for Grainger. But Grainger knows that simply offering them products is not good enough. Grainger has developed sophisticated information tools to help customers analyze and replace their purchasing systems. The company's Integrated Supply division is a professional services group that will reengineer a customer's stockroom and provide a just-in-time inventory system.

Like other distributors, Grainger was hurt by the economic downturn of 2001 and 2002. Although sales revenue declined, the company used its supply chain software to boost its profit margins, resulting in an increase in net income (Konicki 2002). The company also expanded its investments in warehouses and information technology. Rimnac observes that "if you have financial resources, a recession is a good time to invest so when the rebound comes you can excel" (Konicki 2002). After another decline in sales in 2002, yet a continued improvement in profits, CEO Richard L. Keyser commented that "we remain committed to improving service to our customers as they continue to look for ways to reduce costs. Enhancements to our logistics network and local availability of the right products will provide higher levels of service. Our initiatives should accelerate sales growth as more customers experience this improved service" ("Lower Sales" August 2003).

To combat these changes, Grainger is expanding into additional services. In 2001, the company opened an on-site branch at Florida State University. It followed by opening a second on-site branch at Langley Air Force Base in 2002.

In 1996, Grainger took on a more in-depth role at the American Airlines facility at the Dallas/Fort Worth airport. The company essentially took over all janitorial and MRO services for the facility. With its success, two years later, the contract was extended to cover the nearby American Airlines headquarters building. In 2003, the company ran integrated supply programs for more than 40 customers. Large customers with over $2 million a year in MRO purchases from Grainger are eligible for the program. Grainger customizes the service for each company, but essentially, Grainger identifies the inventory needs and handles everything from ordering to stocking and might even include an on-site center distributing products to employees. At American Airlines, Grainger was able to substantially reduce the amount of MRO inventory sitting around. At the same time, facility worker complaints about not having the necessary tools available disappeared (Fraza 2003).

Grainger experimented with several Web sites and e-business approaches before reaching its current configuration. In 1999, the company had multiple Web divisions, including OrderZone (a marketplace), FindMRO (a search site for hard-to-find parts), and MROverstocks (an auction site). It also had relationships with logistics sites such as Ariba and Commerceone as well as several other e-marketplaces. Carol Rozewell, vice president research director at research firm Gartner Inc., noted that "Grainger customers were confused. They offered such a wide variety of products, customers needed guidance to navigate [the Web sites]" (Maddox 2002). Ultimately, Grainger killed off all of the sites except the company's main site (grainger.com) and the FindMRO site that is accessed only through the main site. But Grainger spent more than $180 million on Internet technologies and took a $23.2 million charge, followed by a $13.4 million write-off of digital properties (Maddox 2002). The main Web site was redesigned to make it easier and faster to use.

Internal Systems

Despite the growing importance of Web-based sales and in-house ties to customers, Grainger is also emphasizing increased sales through its local branches. Pushing products more efficiently through the local stores is the main reason for the $200 million redesign of Grainger's distribution system (Buss 2002).

One step Grainger took to improve efficiency was to install an SAP R/3 ERP system in 1999. Unfortunately, Grainger had several problems installing and configuring the system. For example, the ERP software miscounted the items on hand—partly because of problems in the transaction-processing subsystem. In the first year alone, the system cost Grainger $19 million in lost sales (Stedman 2000). Ultimately, Grainger got the system fixed and consolidated its financial data onto a single system. One of the problems Grainger had was trying to connect too many outside products to the SAP system. In 2003, the company began phasing out the old systems and using standard SAP components instead. Jarnail Lail, vice president of business systems, observed that the move enabled the company to reduce the number of outside consulting firms from 100 down to 10 (Colkin Cuneo 2003).

Despite the integrated enterprise resource planning system, Vice President of Finance Laura Brown wants more. In particular, she wants to be able to match sales figures with expense data in real time (Colkin and Whiting 2002). The principles of the Sarbanes-Oxley Act are scaring financial managers and executives who want to ensure that the numbers the systems are spitting out are true representations of the business finances.

In January 2006, Grainger implemented an Enterprise Resource Planning System from SAP. The goal was to integrate the information across the channels, products, and logistics network. CEO Richard L. Keyser in a September 2004 press release stated that "for us, technology is not a mere indulgence or a nice to have. It is at the very heart of our business and helps us drive our multi-channel strategy. Providing customers access to our broad product line through several channels has been a competitive weapon for us because it spells speed and convenience for our customers." He also observed that 60 percent of the company's orders came in via telephone, so the company also upgraded its telephone system to make it simpler to reach people and place orders. The early-phase implementation of SAP resulted in a $115 million reduction in inventory and a 25 percent increase in productivity, so Grainger was finally happy with the system. According to SAP, the system supports 2 million customer contacts, 800,000 products and 115,000 orders per day (SAP Web site).

Grainger has repeatedly been high in the rankings of "Best Places to Work in IT," compiled by Computerworld. In 2010, the company had 586 IT employees (and 13,000 total employees), with a turnover rate of only two percent. It probably helps that James Ryan the CEO was once the CIO at Grainger.

Questions

1. How is Grainger's business different from that faced by Owens & Minor?
2. Why did Grainger's initial Web approach using multiple sites fail?
3. Grainger's problems implementing its ERP system are often cited for being unusually severe. Why did Grainger experience so many problems?
4. How does Grainger support local stores as well as online sales? Why is that approach successful for Grainger but not for Dell?

Additional Reading

Buss, Dale, "The New Deal," *Sales and Marketing Management*, June 2002, Vol 154(6), pp. 24-29.

Colkin, Eileen and Rick Whiting, "Inadequate IT?" *Information Week*, September 9, 2002.

Colkin Cuneo, Eileen, "Less Is More In Software," *Information Week*, January 13, 2003.

Fraza, Victoria, "Making Integrated Supply Work," *Industrial Distribution*, April 2003, Vol 92(4) p. 69.

Industrial Distribution, "Top Distributor Talks Business," June 2003, Vol 92(6), p. 46.

Industrial Distribution, "Lower Sales, Higher Earnings," August 2003, Vol 92(8), p. 34.

Konicki, Steve, and Eileen Colkin, "Optimism's Back!" *Information Week*, March 25, 2002.

Maddox, Kate, "Growing Wiser," *B to B*, September 9, 2002, Vol 87(9), pp. 1,2.

Pratt, Mary K., "Maintenance, Repair And Operations," *Computerworld*, February 26, 2001.

Stedman, Craig, "ERP Woes Cut Grainger Profits," *Computerworld*, January 7, 2000.

Sviokla, John, "Knowledge Pays," *Information Week*, February 15, 2001.

Case: Univar: Van Waters & Rogers

Van Waters & Rogers is a respected name in the chemical supply industry. Established in Seattle in 1924 by George Van Waters and Nat S. Rogers, the company grew through acquisitions to become the largest wholesale provider of chemicals in the United States and Canada. Through international mergers, the company is also a major (but not the largest) supplier in Europe. In 1973, the company changed its official name to Univar, but kept the esteemed Van Waters & Rogers name on its local units. In 1986, the company acquired McKesson Chemical to extend its reach across the United States. Starting in 1991, the company began a series of mergers and acquisitions in Europe, forming Univar Europe. The rest of the chemical supply industry was consolidating at the same time. In 1996, Royal Pakhoed, a major worldwide player in liquid chemicals, purchased the shares of Univar it did not already own. In 1999, Royal Pakhoed merged with its major competitor (Royal Van Ommeren) and created Royal Vopak. The growth continued with the purchase of Ellis & Everard in 2001, giving Vopak a substantial market share in the United Kingdom and Ireland. However, with the global economic downturn after 2001, all chemical suppliers were suffering. In 2002, Vopak spun off the chemical and distribution assets and reestablished Univar as an independent company. Univar is a Dutch company, listed on the Euronext Amsterdam stock exchange as UNIVR. Its headquarters is in Rotterdam, but major administrative offices remain in Bellevue, Washington (company history Web site).

Largely because of the economic recession of 2003 and the high prices for oil (a major ingredient for many chemicals as well as a factor in transportation costs), all of the major chemical suppliers have suffered. Analysts estimate that Vopak has invested about $1 billion in Univar. Shortly after stock was issued in the new Univar company, its shares fell to €5.6, which gave the company a valuation of only about €150 million, less than one-fifth of that billion dollars (Tilton 2002).

Univar has progressed through considerable turmoil with the various mergers and the flat economies. In the early 1990s, the company tried to centralize management and run its distribution centers using a hub-and-spoke system. Analysts have pointed out that the system worked poorly and sales suffered. With the purchase by Vopak, the Van Waters & Rogers division was moved back to a decentralized approach, with more responsibility pushed to regional directors and a focus on specific product segments. In the late 1990s, the company moved more operations out to multiple local warehouses (Morris 1996). The multiple acquisitions during the late 1990s and early 2000s did not really help sales in the short run. Managers and investors faced huge uncertainties about what was going to happen next.

Initially, splitting the company back from Vopak appears to have helped. Sales for 2003 increased 9 percent to 1.62 billion eruos. However, much of that gain was due to currency effects from the 40 percent Euro appreciation. (In the chart at the top of this section, Univar's sales have been converted to dollars using a single exchange rate to eliminate the exchange-rate effects.) Sales in North America were flat and hampered by the manufacturing slowdown and the rise in oil prices (Van Arnum 2003). A key step in the mergers was to consolidate the European operations under a single brand. Managing the company globally instead of regionally is a second imperative. As part of Vopak, the strategy had been to balance regional storage with distribution. John Phillpotts, president of Univar Europe, notes that "we never operated commercially as a global business before. For the last 18 months, we've had no distractions about where the company is going or who's going to own it. Chemical distribution is all we talk about now" (Young 2004).

In 2007, Univar began the process of selling to CVC Capital Partners, a London-based private equity firm (Berman 2007). In 2007, Univar acquired CHEMCENTRAL Corporation and expanded into new industries including coatings, personal care, and cleaning. Univar servers more than 250,000 client customers worldwide and leads the industry in chemical distribution (2006 Annual Report). Much of Univar's profit margins come from processing, packaging, and distribution. Most of the inventory turns over every 40 days. Most of the costs are fixed, particularly tied up in processing and storage facilities such as the 160 distribution centers in the U.S., Canada, and Europe. Net sales in 2006 were $6.6 billion with a net income of $134 million. For four years, sales grew at an 11 percent compound growth rate and earnings grew at 21 percent due to declining costs (2006 Annual Report). Net sales in 2009 were $7.2 billion and Univar had 179 distribution facilities and 6,900 employees around the world (Web site/annual report). The company continued its acquisition of smaller distributors in Europe.

E-Business

Univar has followed two approaches to e-business. The company is part of the consortium that drives Elemica, the chemical industry Web site. Univar also runs its own B2B Web portal, ChemPoint. The ChemPoint site is also designed to support smaller chemical manufacturers and highlight new products. In particular, the site is aimed at providing distribution services to chemical producers. Customers place orders through the site, and Univar handles the distribution—linking electronically to Yellow Transportation, its main less-than-truckload (LTL) carrier. Chad Steigers, managing director of ChemPoint.com, notes that there are still some difficulties getting companies to adopt electronic transactions. The goal of many firms is to increase productivity. His firm is looking to more technology to encourage the use of e-business. "We think, for example, RFID will be a major factor especially in the chemical arena" (Cottrill 2003). Ultimately, buyers should be able to place an order, and then track it all the way through shipping, storing, and delivery.

ChemPoint focused on ensuring a broad market by signing up at least one supplier for each type of chemical, with the goal of providing buyers with a dependable source. Reliability and reduction of search costs are important factors for many of the hard-to-find chemicals (Cuny 2001).

When the ChemPoint site was started, it faced intense competition. As Chad Steigers, managing director, pointed out in 2000, "There are 25 trading companies that offer auctions of chemicals. We think that will drop down to two or three." At the time, the Web site had 127 customers and sold 1.3 million products provided by 2,200 suppliers, amounting to $24 million in transactions a year (Seideman, 2000). The emphasis on specialty chemicals and distribution are key elements that have kept the site active today. Steigers observes that "e-distribution is a new segment for the industry. We're the only ones focused on the less-than-truckload segment of the specialty and fine chemicals market, which is about a $70-billion/year niche" (Fuller 2000).

Univar also participates in Elemica, an electronic marketplace for the huge chemical producers and distributors. The site was set up at the end of 2000, primarily to facilitate long-term contracts. It was also designed to assist firms with financing and transportation issues. The decision to avoid spot market auctions was probably due to the competing systems already in development (most of them have since failed) (Rosencrance 2000). Because of the emphasis on large sales, Elemica is emphasizing the ability to connect ERP systems between buyers and sellers. Larger purchasers can transmit orders, receive confirmations, and transfer the data directly into their accounting ERP systems (www.elemica.com).

Univar also runs specialty Web sites for certain uses. For instance, the pestweb.com site is used to provide information and sell products to pest and weed control companies. To provide wider access to data and ordering, the site created an iPhone app to provide information on products and the industry. It enables large and small firms to obtain data and order chemicals directly from their nearest Univar provider (pestweb.com).

The company also wants to build in more relationship-based features to its Web sites. The goal is to help customers find related information and products that might be useful to them. As CIO Cummings points out, it is challenging for customers to understand the full breadth of products provided by Univar, so "If we knew a bit more about what the customer does, we could direct them towards more of our capabilities" (Beacham 2010).

Internal Network Technologies

The global nature of Univar makes it more difficult to manage the company. Even across the United States, the battles between centralization and decentralization led to sales problems. Dealing with multiple offices around the world, in a company that now sees itself as a single global entity, can be difficult and expensive. To improve communications without increasing travel costs, Van Waters & Rogers turned to PlaceWare's hosted collaboration system (now owned by Microsoft and renamed LiveMeeting). The system provides voice and video to any desktop that is connected to the Internet. It also has shared drawing boards, chat facilities, and the ability to show slides. It can also store a presentation for later playback if some people cannot make the initial meeting. In one example, Van Waters & Rogers used it for a strategy meeting between a team in North America and one in Switzerland to discuss a product rollout (Agnew 2000). The system is also used for training and general meetings. Ron Miazga, the human resources training director, notes that the system has been well-received and cost effective. However, building training sessions takes some practice, since "you have to make the presentations very media intensive. You should ask a polling question every three to five minutes. You have to change the screen every minute or so to hold their attention" (Agnew 2000).

In 2006, Univar signed with Odyssey Logistics to handle in-bound deliveries including transportation and management. Odyssey claims that one reason for the strategic alliance was the power and capabilities of its Global Logistics Platform software (OdysseyLogistics.com). The software integrates with commercial and private carriers of rail and trucks to provide scheduling and tracking. It also determines rates, availability, and scheduling over the Web. Outsourcing these operations simplifies management and provides capabilities that Univar would not be able to create by itself.

Around 2007, Univar replaced its IBM computers with an IBM System z9 enterprise system because of the high growth rates and increased server loads. The company's custom ERP system was designed and built to run on IBM large computers and the company needed a new system that could expand and grow with future needs. The system can also run virtual Linux servers which could be used to reduce the number of physical servers Univar runs.

In 2010, Univar hired Dave Cummings as a new CIO. His main experience was with ConocoPhillips as a program manager with SAP applications (press release). Univar has an online ordering system that supports standard ordering, and tracking, as well as material safety data sheets (MSDS).

Questions

1. Why does decentralization at Univar require more use of networks?
2. Why does Univar pursue a multiple-site e-commerce strategy, and why does it work for Univar and not for Grainger?
3. Why are some customers reluctant to adopt electronic ordering for chemicals?

Additional Reading

Agnew, Marion, "Collaboration On The Desktop," *Information Week*, July 10, 2000.

Beacham, Will, Interview with CIO Dave Cummings, March 10, 2010, http://www.icis.com/Articles/2010/03/10/9341649/e-commerce-technologies-could-revolutionise-chemicals.html.

Berman, Dennis K., "CVC Capital Is Set to Purchase Univar," *The Wall Street Journal*, July 9, 2007.

Corporate History Web site: http://www.univarcorp.com/ourhistory.htm

Cottrill, Ken, "Overcoming Inertia," *Traffic World*, January 27, 2003, p. 1.

Cuny, Tim and Marvin Richardson, "Redefining Business: Proving Grounds," *Information Week*, February 12, 2001.

Fuller, John, "Vopak Launches ChemPoint.com," *Chemical Week*, March 15, 2000, Vol 162(11), p. 56.

IBM Case Study, "Univar Extends Computing Capabilities with IBM System z9," http://www.mysiriuszone.com/component/option,com_docman/task,doc_download/gid,2828/.

Morris, Gregory DL, "Pakhoed's VW&R Buy Accelerates Expansion," *Chemical Week*, October 16, 1996, Vol 158 (39), p. 60.

Rosencrance, Linda, "Twelve Chemical Companies Are Setting Up Online Exchange," *Computerworld*, May 17, 2000.

Seideman, Tony, "E-Marketplace Vision Collides With Reality," *InformationWeek*, June 6, 2000.

Tilton, Helga, "Analysts Positive on More Focused Univar," *Chemical Market Reporter*, November 18, 2002, Vol 262(18), p. 19.

Van Arnum, Patricia, "North American Distributors, Wait and See," *Chemical Market Reporter*, November 24, 2003, Vol 264(18) p. FR3.

Young, Ian, "European Unification Takes Shape at Univar," *Chemical Week*, May 19, 2004, Vol 166(17), p. 22.

Summary Industry Questions

1. What information technologies have helped this industry?
2. Did the technologies provide a competitive advantage or were they quickly adopted by rivals?
3. Which technologies could this industry use that were developed in other sectors?
4. Is the level of competition increasing or decreasing in this industry? Is it dominated by a few firms, or are they fairly balanced?
5. What problems have been created from the use of information technology and how did the firms solve the problems?