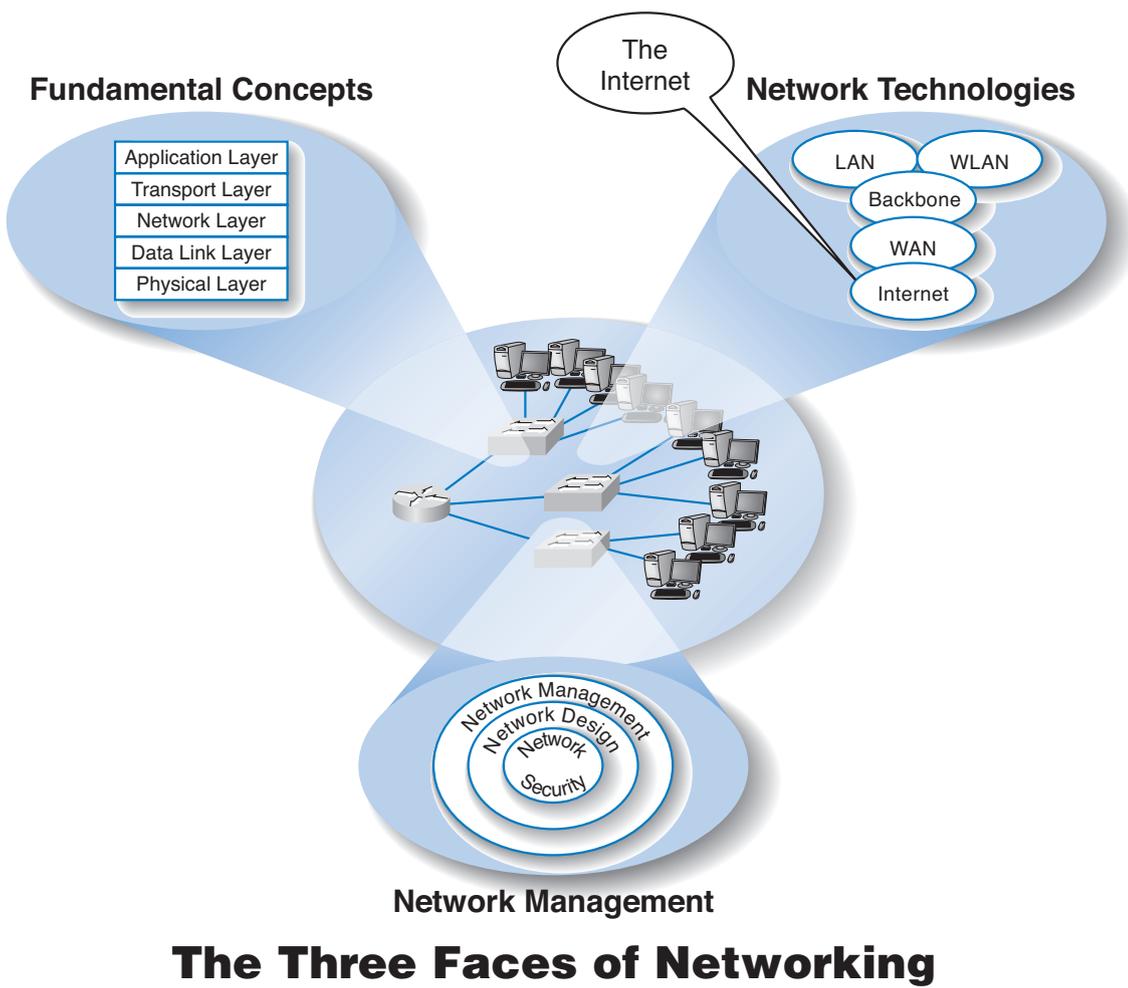


CHAPTER 10

THE INTERNET



THIS CHAPTER examines the Internet in more detail to explain how it works and why it is a network of networks. This chapter also examines Internet access technologies, such as DSL and cable modem, as well as the possible future of the Internet in the form of Internet 2.

OBJECTIVES

- Understand the overall design of the Internet
- Be familiar with DSL, cable modem, and Wireless Application Protocol
- Be familiar with Internet 2

CHAPTER OUTLINE

INTRODUCTION

HOW THE INTERNET WORKS

Basic Architecture

Connecting to an ISP

The Internet Today

INTERNET ACCESS TECHNOLOGIES

DSL

Cable Modems

Fixed Wireless

Mobile Wireless

Future Technologies

INTERNET GOVERNANCE

INTERNET 2

IMPLICATIONS FOR MANAGEMENT

SUMMARY

INTRODUCTION

The Internet is the most used network in the world, but it is also one of the least understood. There is no one network that is *the* Internet. Instead, the Internet is a network of networks—a set of separate and distinct networks operated by various national and state government agencies, nonprofit organizations, and for-profit corporations. The Internet exists only to the extent that these thousands of separate networks agree to use Internet protocols and to exchange data packets among one another.

The Internet is simultaneously a strict, rigidly controlled club in which deviance from the rules is not tolerated and a freewheeling, open marketplace of ideas. All networks that connect to the Internet must rigidly conform to an unyielding set of standards for the transport and network layers; without these standards, data communication would not be possible. At the same time, content and new application protocols are developed freely and without restriction, and quite literally anyone in the world is allowed to comment on proposed changes.

In this chapter, we first explain how the Internet really works and look inside one of the busiest intersections on the Internet, the Chicago network access point, at which about 100 separate Internet networks meet to exchange data. We then turn our attention to how you as an individual can access the Internet and what the Internet may look like in the future.

HOW THE INTERNET WORKS

Basic Architecture

The Internet is hierarchical in structure. At the top are the very large national *Internet service providers (ISPs)*, such as AT&T and Sprint, that are responsible for large Internet networks. These *national ISPs*, sometimes called NSPs, connect together and exchange data at *network access points (NAPs)* (Figure 10.1).

In the early 1990s, when the Internet was still primarily run by the U.S. National Science Foundation (NSF), the NSF established four main NAPs in the United States to connect the major national ISPs. When the NSF stopped funding the Internet, the companies running these NAPs began charging the national ISPs for connections, so today the NAPs in the United States are all commercial enterprises run by various common carriers such as Ameritech and Sprint. As the Internet has grown, so too has the number of NAPs; today there are about a dozen NAPs in the United States with many more spread around the world.

NAPs were originally designed to connect only national ISPs. These national ISPs in turn provide services for their customers and also to *regional ISPs* such as BellSouth and EarthLink. These regional ISPs rely on the national ISPs to transmit their messages to national ISPs in other countries. Regional ISPs, in turn, provide services to their customers and to local ISPs, who sell Internet access to individuals. As the number of ISPs grew, a new form of NAP called a metropolitan area exchange (MAE) emerged. MAEs are smaller versions of NAPs and typically link a set of regional ISPs whose networks come together in major cities (Figure 10.1). Today there are about 50 MAEs in the United States.

Because most NAPs, MAEs, and ISPs now are run by commercial firms, many of the early restrictions on who could connect to whom have been lifted. Indiana University, for

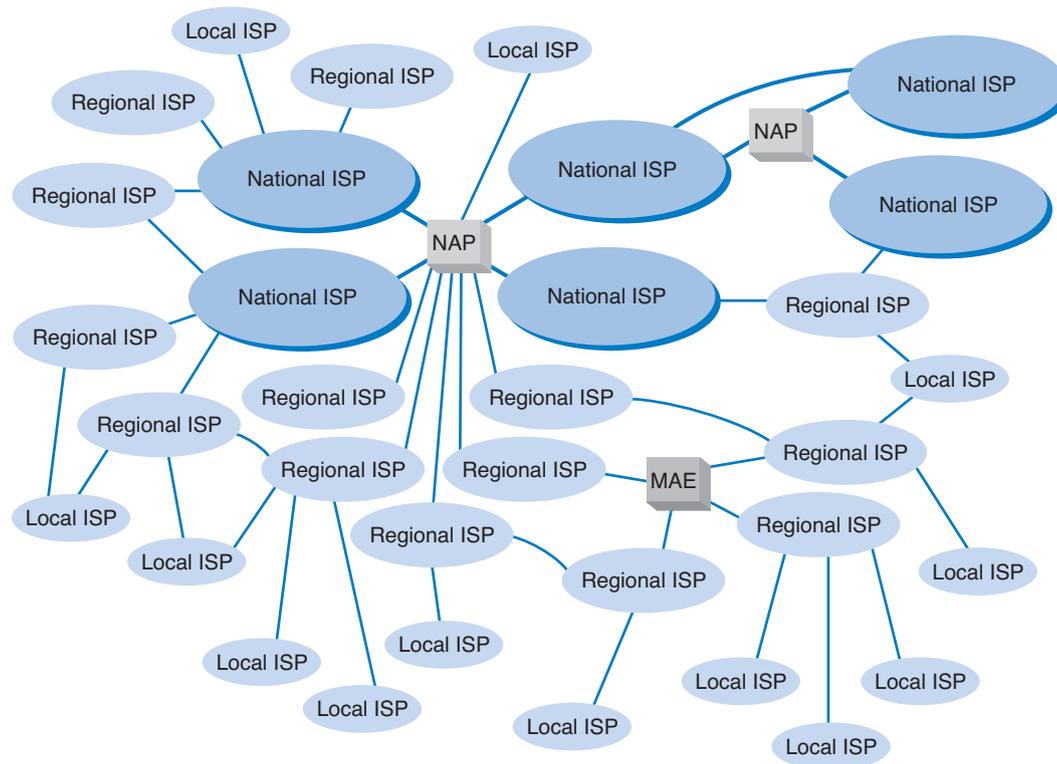


FIGURE 10.1 Basic Internet architecture. ISP = Internet service provider; MAE = metropolitan area exchange; NAP = network access point.

example, which might be considered a local ISP because it provides Internet access for about 40,000 individuals, has a direct connection into the Chicago NAP, as do several other universities and large corporations. Regional and local ISPs often will have several connections into other national, regional, and local ISPs to provide backup connections in case one Internet connection fails. In this way, they are not dependent on just one higher-level ISP.

In general, ISPs at the same level do not charge one another for transferring messages they exchange across a NAP or MAE. That is, a national ISP does not charge another national ISP to transmit its messages, and a regional ISP does not charge another regional ISP. This is called *peering*. Figure 10.1 shows several examples of peering. It is peering that makes the Internet work and has led to the belief that the Internet is free. This is true to some extent, but higher-level ISPs normally charge lower-level ISPs to transmit their data (e.g., a national will charge a regional and a regional will charge a local). And of course, a local ISP will charge individuals like us for access!

In October, 2005, an argument between two national ISPs, Level 3 and Cogent, shut down 45 million Web sites for a week. The two ISPs have a peering agreement but Level 3 complained that Cogent was sending it more traffic than it should and demanded payment. Cogent refused, so Level 3 stopped accepting Cogent's traffic leaving large portions of Co-

gent's network isolated from the rest of the Internet. The dispute was resolved and Level 3 began accepting traffic from Cogent, connecting it to the rest of the Internet again.

In Figure 10.1, each of the ISPs are *autonomous systems*, as defined in Chapter 5. Each ISP is responsible for running its own interior routing protocols and for exchanging routing information via the BGP exterior routing protocol (see Chapter 5) at NAPs and MAEs and any other connection points between individual ISPs.

Connecting to an ISP

Each of the ISPs is responsible for running its own network that forms part of the Internet. ISPs make money by charging customers to connect to their part of the Internet. Local ISPs charge individuals for broadband or dial-up access whereas national and regional ISPs (and sometimes local ISPs) charge larger organizations for higher-speed access.

Each ISP has one or more *points of presence (POP)*. A POP is simply the place at which the ISP provides services to its customers. To connect into the Internet, a customer must establish a circuit from his or her location into the ISP POP. For individuals, this is often done using a DSL modem, cable modem, or dial-up modem over a traditional telephone line (Figure 10.2). This call connects to the modem pool at the ISP and

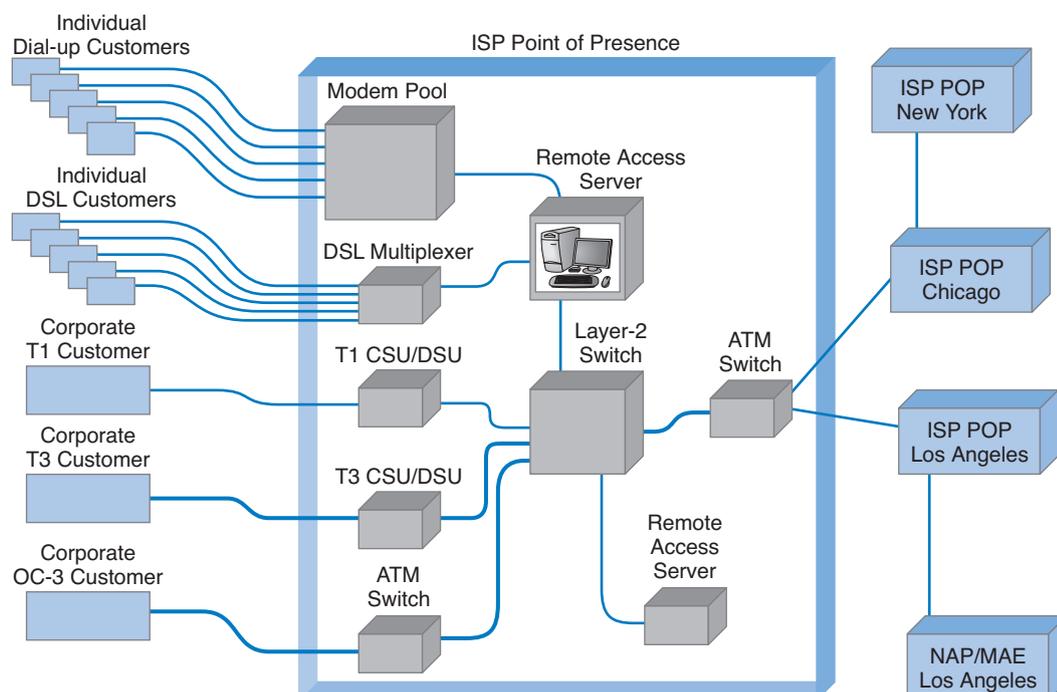


FIGURE 10.2 Inside an Internet service provider (ISP) point of presence (POP). ATM = asynchronous transfer mode; CSU = channel service unit; DSU = data service unit; MAE = metropolitan area exchange; NAP = network access point.

from there to a *remote-access server (RAS)*, which checks the user ID and password to make sure the caller is a valid customer. Once logged in, the user can begin sending TCP/IP packets from his or her computer over the phone to the POP. Figure 10.2 shows a POP using a collapsed backbone with a layer-2 switch. The POP backbone can take many forms, as we discussed in Chapter 8.

In the next section, we will discuss Internet access technologies such as DSL, cable modem, and Wireless Application Protocol (WAP) in more detail. Customers who need more network capacity simply lease a higher-capacity circuit. Figure 10.2 shows corporate customers with T1, T3, and OC-3 connections into the ISP POP. It is important to note that the customer must pay for both Internet access (paid to the ISP) and for the circuit connecting from their location to the POP (usually paid to the local exchange carrier [e.g., BellSouth, Ameritech], but sometimes the ISP also can provide circuits). For a T1 connection, for example, a company might pay the local exchange carrier \$400 per month to provide the T1 circuit from its offices to the ISP POP and *also* pay the ISP \$600 per month to provide the Internet access.

As Figure 10.2 shows, the ISP POP is connected in turn to the other POPs in the ISP's network. Any messages destined for other customers of the same ISP would flow

MANAGEMENT**10-1 INSIDE THE CHICAGO NETWORK ACCESS POINT****FOCUS**

The Chicago network access point (NAP) is one of the busiest NAPs in the world. As we write this, it processes an average of about 4 gigabits of data per second.

More than 140 different Internet service providers (ISPs), including national ISPs (e.g., BBN Planet and Sprint), regional ISPs (e.g., Michigan's Merit network), and local ISPs (e.g., Indiana University), as well as ISPs in other countries (e.g., Germany's Tiscali network and the Singapore Advanced Research and Education Network), exchange traffic at the Chicago NAP. At present, most connections are asynchronous transfer mode (ATM) OC-3, or ATM OC-12, and the rest are T3. Pricing starts at about \$4,000 per month for T3 and about \$4,700 per month for OC-3. (Remember, this is only for Internet access; the ISPs must also lease a T3 or OC-3 circuit from their closest point-of-presence [POP] to the NAP.)

The NAP currently uses a large Cisco ATM switch that connects the more than 140 separate ISP networks (Figure 10.3). The ISP networks exchange IP packets through the NAP. They also ex-

change routing information through the Border Gateway Protocol (BGP) exterior routing protocol. Normally, the border router at each ISP simply generates BGP packets and sends them to the border routers at the other ISPs connected to the NAP. The Chicago NAP has so many ISPs that this is impossible. Because there are about 140 ISPs, each ISP would send messages to about 140 other ISPs, meaning a total of about 1 million BGP packets moving through the NAP every few minutes.

Instead, the Chicago NAP uses a route server in much the same way large networks based on OSPF (Open Shortest Path First) used designated routers (see "Routing on the Internet" in Chapter 5). The border router in each ISP sends BGP packets just to the NAP route server. The route server consolidates the routing information and then sends BGP packets back to each border router. This results in more efficient processing and only 200 messages every few minutes.

SOURCE: www.aads.net/main.html.

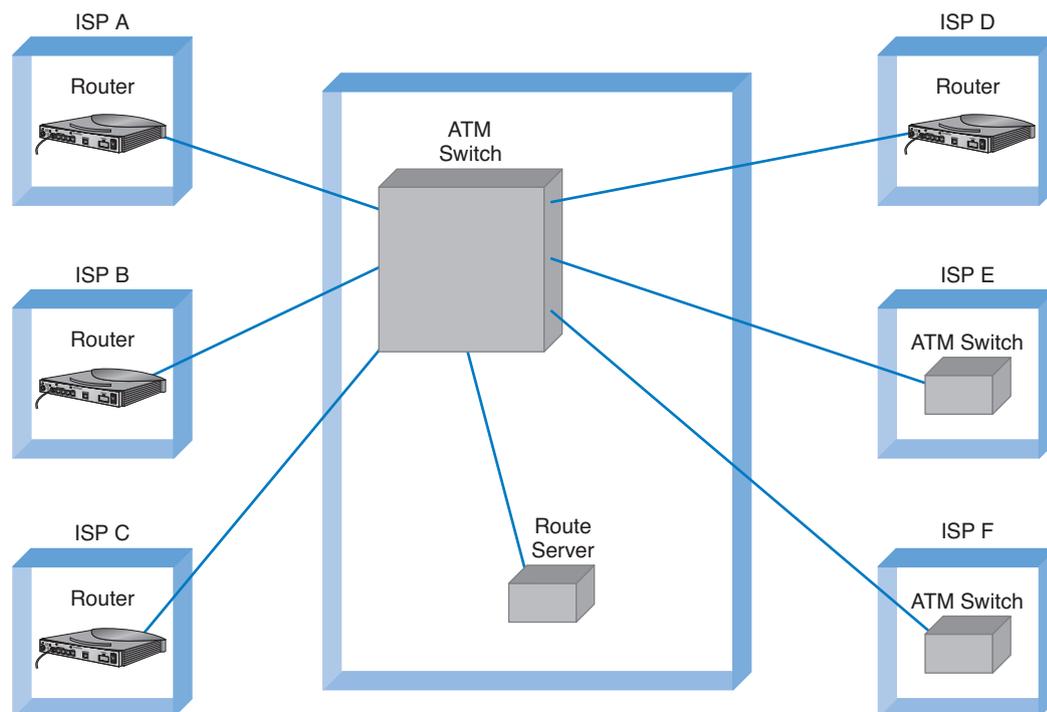


FIGURE 10.3 Inside the Internet's Chicago network access point. ATM = asynchronous transfer mode; ISP = Internet service provider.

within the ISP's own network. In most cases, the majority of messages entering the POP are sent outside of the ISP's network and thus must flow through the ISP's network to the nearest NAP/MAE, and from there, into some other ISP's network.

This can be less efficient than one might expect. For example, suppose you are connected to the Internet via a local ISP in Minneapolis and request a Web page from another organization in Minneapolis. A short distance, right? Maybe not. If the other organization uses a different local ISP, which in turn uses a different regional ISP, the message may have to travel all the way to the Chicago NAP before it can move between the two separate parts of the Internet.

The Internet Today

Sprint is one of the national ISPs in North America. Figure 10.4 shows Sprint's North American backbone as it existed while we were writing this book; it will have changed by the time you read this. As you can see, Sprint has a number of Internet circuits across the United States and Canada. Many interconnect in Chicago where Sprint connects into the Chicago NAP. Sprint also connects into major NAPs and MAEs in Reston, Virginia; Miami; Los Angeles; San Jose; Palo Alto; Vancouver; Calgary; Toronto; and Montreal. Most of the circuits are ATM OC-12, but a few are OC-48 and OC-192.

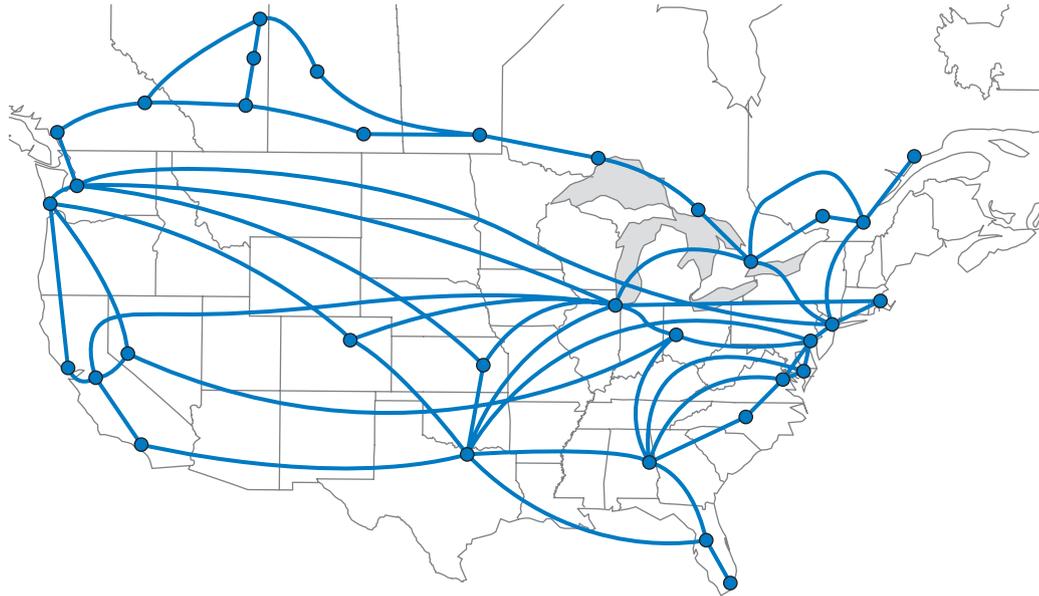


FIGURE 10.4 Sprint's North American Internet backbone.

Today, the backbone circuits of the major U.S. national ISPs operate at ATM OC-48 and OC-192. Most of the largest national ISPs (e.g., Sprint, Cable & Wireless) plan to convert their principal backbones to OC-192 (10 Gbps) by the end of 2005. A few are now experimenting with OC-768 (80 Gbps), and several are in the planning stages with OC-3072 (160 Gbps). This is good because the amount of Internet traffic has been growing rapidly. The Internet traffic in the U.S. is expected to reach 40 Tbps (40 trillion bits per second) by 2007.

As traffic increases, ISPs can add more and faster circuits relatively easily, but where these circuits come together at NAPs and MAEs, bottlenecks are becoming more common. Network vendors such as Cisco and Juniper are making larger and larger switches capable of handling these high-capacity circuits, but it is a daunting task. When circuit capacities increase by 100 percent, switch manufacturers also must increase their capacities by 100 percent. It is simpler to go from a 622 Mbps circuit to a 10 Gbps circuit than to go from a 20 Gbps switch to a 200 Gbps switch.

The Internet is constantly changing, so by the time you read this, CAIS, CompuServe, and iSTAR will likely have added extra circuits. Up-to-date maps of the major ISPs whose networks make up large portions of the Internet are available at www.caida.org and at navigators.com/isp.html.

INTERNET ACCESS TECHNOLOGIES

There are many ways in which individuals and organizations can connect to an ISP. Some individuals use 56-Kbps dial-up modems over telephone lines; some use DSL or cable

modem. As we discussed in the preceding section, many organizations lease T1 or T3 lines into their ISPs. DSL and cable modem technologies are commonly called *broadband technologies* because they provide higher-speed communications than traditional modems.¹

It is important to understand that Internet access technologies are used only to connect from one location to an ISP. Unlike the MAN and WAN technologies in the previous chapter, Internet access technologies cannot be used for general-purpose networking from any point to any point. In this section, we discuss four principal Internet access technologies (DSL, cable modem, fixed wireless, and mobile wireless) and also discuss some future technologies that may become common.

DSL

Digital subscriber line (DSL) is a family of point-to-point technologies designed to provide high-speed data transmission over traditional telephone lines.² The reason for the limited capacity on traditional telephone circuits lies with the telephone and the switching equipment at the end offices. The actual cable in the *local loop* from a home or office to the telephone company end office is capable of providing much higher data transmission rates. So conversion from traditional telephone service (POTS) to DSL usually requires just changing the telephone equipment, not rewiring the local loop, which is what has made it so attractive.

Architecture DSL uses the existing local loop cable but places different equipment on the customer premises (i.e., the home or office) and in the telephone company end office. The equipment that is installed at the customer location is called the *customer premises equipment (CPE)*. Figure 10.5 shows one common type of DSL installation. (There are other forms.) The CPE in this case includes a *line splitter* that is used to separate the traditional voice telephone transmission from the data transmissions. The line splitter directs the telephone signals into the normal telephone system so that if the DSL equipment fails, voice communications are unaffected.

The line splitter also directs the data transmissions into a *DSL modem*, which is sometimes also called a DSL router. As you will recall from Chapter 3, this is both a modem and an FDM multiplexer. The DSL modem produces Ethernet 10Base-T packets so it can be connected directly into a computer or to a router and hub and can serve the needs of a small network.

Figure 10.5 also shows the architecture within the local carrier's end office (i.e., the telephone company office closest to the customer premises). The local loops from many customers enter and are connected to the *main distribution facility (MDF)*. The MDF works like the CPE line splitter; it splits the voice traffic from the data traffic and directs

¹*Broadband* is a technical term that means "analog transmission" (see Chapter 3). The new broadband technologies often use analog transmission, so they were called broadband. However, the term *broadband* has been corrupted in common usage so that to most people it usually means "high speed."

²DSL is rapidly changing because it is so new. More information can be found from the DSL forum (www.adsl.com, www.dsllife.com) and the ITU-T under standard G.992.

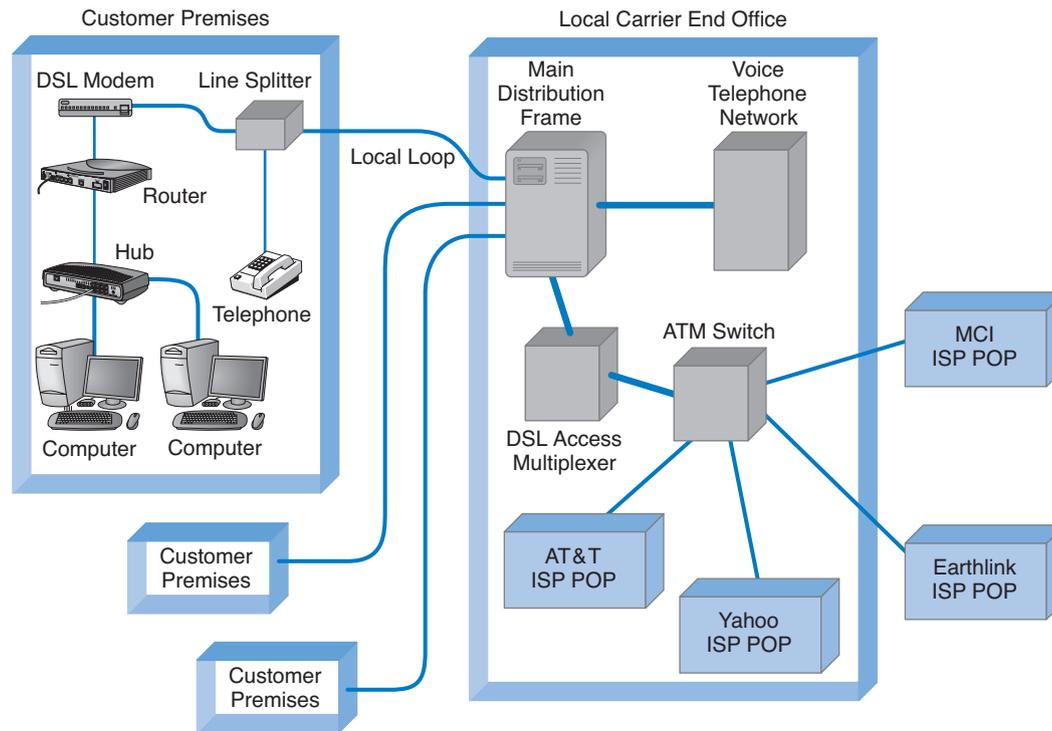


FIGURE 10.5 Digital subscriber line (DSL) architecture. ATM = asynchronous transfer mode; ISP = Internet service provider; POP = point of presence.

the voice traffic to the voice telephone network and the data traffic to the *DSL access multiplexer (DSLAM)*. The DSLAM demultiplexes the data streams and converts them into ATM data, which are then distributed to the ISPs. Some ISPs are collocated, in that they have their POPs physically in the telephone company end offices. Other ISPs have their POPs located elsewhere.

Types of DSL DSL services are not available in all locations. In general, DSL services have advanced more quickly in Canada, Europe, Australia, and Asia than in the United States, owing to their newer telephone networks from the end offices to the customer.

There are many different types of DSL. The most common type of DSL in use today is *asymmetric DSL (ADSL)*. ADSL uses frequency division multiplexing (see Chapter 3) to create three separate channels over the one local loop circuit. One channel is the traditional voice telephone circuit. A second channel is a relatively high-speed simplex data channel downstream from the carrier's end office to the customer. The third channel is a slightly slower duplex data channel primarily used for upstream from the customer to the carrier's end office.³ ADSL is called asymmetric because its two data channels have dif-

³Because the second data channel is intended primarily for upstream data communication, many authors imply that this is a simplex channel, but it is actually a set of half-duplex channels.

Type	Maximum Length of Local Loop	Maximum Downstream Rate	Maximum Upstream Rate
ADSL T1 (G.Lite)	18,000 feet	1.5 Mbps	384 Kbps
ADSL E1*	16,000 feet	2.0 Mbps	384 Kbps
ADSL T2	12,000 feet	6 Mbps	640 Kbps
SDSL	18,000 feet	1.5 Mbps	1.5 Mbps

*E1 is the European standard services similar to T1 services in North America.

FIGURE 10.6 Digital subscriber line data rates.

ferent speeds. Each of the two data channels are further multiplexed using time division multiplexing so they can be further subdivided.

The size of the two digital channels depends on the distance from the CPE to the end office. The shorter the distance, the higher the speed, because with a shorter distance, the circuit suffers less attenuation and higher-frequency signals can be used, providing a greater bandwidth for modulation. Figure 10.6 lists the common types of ADSL.

ADSL providers face a challenge in selecting what type of ADSL to offer in a given market. On one hand, customers want the highest speed access possible. However, because there is a trade-off between speed and distance, if an ADSL provider chooses a high-speed version, they have just limited the number of customers they can serve because a significant proportion of households in the United States are long distances from the nearest end office. Most ADSL providers have therefore chosen the T1 level of ADSL and offer it under the trademarked name of *G.Lite ADSL*. Higher speed versions are also available.

A second common type of DSL is *very-high-data-rate digital subscriber line (VDSL)*. VDSL is asymmetric DSL service designed for use over very short local loops of at most 4,000 feet, with 1,000 feet being more typical. It also uses frequency division multiplexing (FDM) to provide three channels: the normal analog voice channel, an upstream digital channel, and a downstream digital channel. Figure 10.7 lists the types of VDSL we anticipate will become common.

VDSL has not yet been standardized, and five separate standards groups are working on different standards. Therefore, the exact data speeds and channels are likely to

Type	Maximum Length of Local Loop	Maximum Downstream Rate	Maximum Upstream Rate
1/4 OC-1	4,500 feet	13 Mbps	1.6 Mbps
1/2 OC-1	4,000 feet	26 Mbps	2.3 Mbps
OC-1	4,000 feet	52 Mbps	16 Mbps

FIGURE 10.7 Data rates for very-high-data-rate digital subscriber line. OC = optical carrier.

change as manufacturers, telephone companies, and ITU-T gain more experience and as the standards groups attempt to merge competing standards. Several companies are also developing symmetric versions of VDSL in which upstream and downstream channels have the same capacity. We expect major changes to VDSL.

Cable Modems

One alternative to DSL is the *cable modem*, a digital service offered by cable television companies. There are several competing standards, but the *Data over Cable Service Interface Specification (DOCSIS)* standard is the dominant one. DOCSIS is not a formal standard but is the one used by most vendors of *hybrid fiber coax (HFC)* networks (i.e., cable networks that use both fiber-optic and coaxial cable). As with DSL, these technologies are changing rapidly.⁴

Architecture Cable modem architecture is very similar to DSL—with one very important difference. DSL is a point-to-point technology whereas cable modems use *shared* multipoint circuits. With cable modems, each user must compete with other users for the available capacity. Furthermore, because the cable circuit is a multipoint circuit, all messages on the circuit go to all computers on the circuit. If your neighbors were hackers, they could use pocket sniffers such as Ethereal (see Chapter 4) to read all messages that travel over the cable, including yours.

Figure 10.8 shows the most common architecture for cable modems. The cable TV circuit enters the customer premises through a cable splitter that separates the data transmissions from the TV transmissions and sends the TV signals to the TV network and the data signals to the cable modem. The cable modem (both a modem and frequency division multiplexer) translates from the cable data into Ethernet packets, which then are directed into a computer to a router and hub for distribution in a small network.

The cable TV cable entering the customer premises is a standard coaxial cable. A typical segment of cable is shared by anywhere from 300 to 1,000 customers, depending on the cable company that installed the cable. These 300 to 1,000 customers share the available data capacity, but of course, not all customers who have cable TV will choose to install cable modems. This coax cable runs to a *fiber node*, which has an *optical-electrical (OE) converter* to convert between the coaxial cable on the customer side and fiber-optic cable on the cable TV company side. Each fiber node serves as many as half a dozen separate coaxial cable runs.

The fiber nodes are in turn connected to the cable company *distribution hub* (sometimes called a headend) through two separate circuits: an upstream circuit and a downstream circuit. The upstream circuit, containing data traffic from the customer, is connected into a *cable modem termination system (CMTS)*. The CMTS contains a series of cable modems/multiplexers and converts the data from cable modem protocols into protocols needed for Internet traffic, before passing them to a router connected to an ISP POP. Often, the cable company is an Internet regional ISP, but sometimes it just provides Internet access to a third-party ISP.

⁴More information can be found at www.cablemodem.com and www.cable-modems.org.

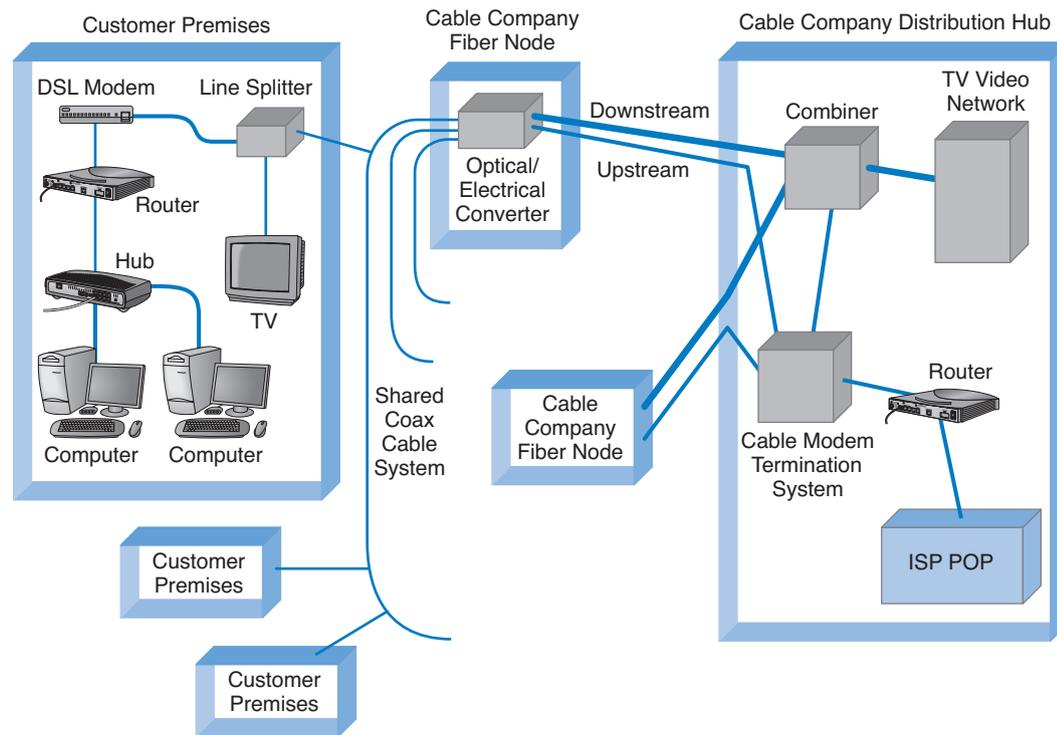


FIGURE 10.8 Cable modem architecture. ISP = Internet service provider; POP = point of presence.

The downstream circuit to the customer contains both ordinary video transmissions from the cable TV video network and data transmissions from the Internet. Downstream data traffic enters the distribution hub from the ISP POP and is routed through the CMTS, which produces the cable modem signals. This traffic is then sent to a *combiner*, which combines the Internet data traffic with the ordinary TV video traffic and sends it back to the fiber node for distribution.

Types of Cable Modems There are few widely used standards in the cable modem industry because, unlike the telephone system, each cable TV company was able to build very different HFC cable plants because each cable company was a separate entity with no need to connect to other cable TV networks. In theory, cable modems can provide downstream speeds of 27 to 55 Mbps and upstream speeds of 2 to 10 Mbps, depending on the exact nature and quality of the HFC cable plant. In practice, most cable systems do not offer speeds at this rate. Today, typical downstream speeds range between 768 Kbps and 1.5 Mbps and typical upstream speeds range between 200 Kbps and 1 Mbps. However, as some cable modem standards emerge as dominant standards, we should see a consolidation in the types of cable modem services offered.

MANAGEMENT**10-2 INTERNET SPEED TEST****FOCUS**

The speed of your Internet connection depends upon many things, such as your computer's settings, the connection from your computer to your ISP, and the connections your ISP has into the Internet. There are many Internet

sites that enable you to test how fast your Internet connection actually is. Our favorite is operated by CNET: reviews.cnet.com/Bandwidth_meter/7004-7254_7-0.html

Fixed Wireless

The most popular type of *fixed wireless* is *wireless DSL*, which requires a line of sight between the communicating transmitters. For this reason, it has limited application because it requires tall buildings or towers to be effective. The most common use today is to provide Internet access to multitenant buildings such as remote office buildings, apartment buildings, and hotels. Transmitters are used to connect the building to the ISP, and DSL is used inside the building to connect to the wireless transceiver (Figure 10.9).

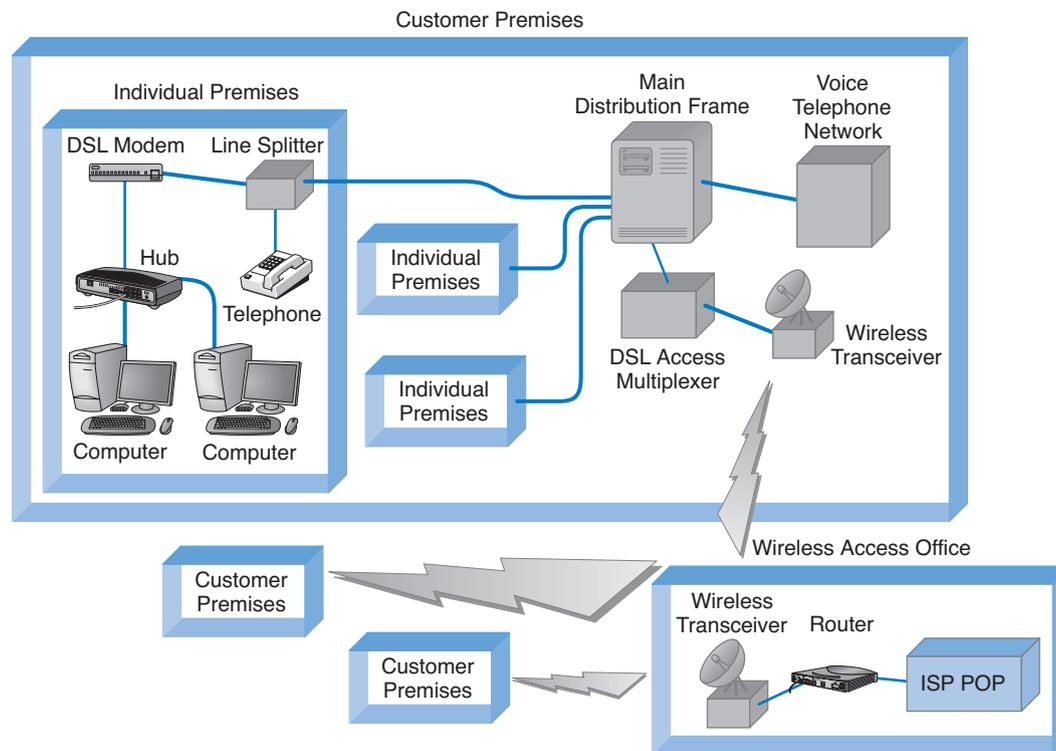


FIGURE 10.9 Fixed wireless architecture. DSL = digital subscriber line; ISP = Internet service provider; POP = point of presence.

Fixed wireless comes in both point-to-point and multipoint versions. The point-to-point version is designed to connect only two locations and is often used as backbone between buildings owned by the same organization. The multipoint version is sometimes called point-to-multipoint because there is one central receiver and all other locations communicate only with it. The multipoint version is designed as an alternative to DSL and cable modems and is intended for use by an ISP supporting a small number of customers. Like cable modems, the circuit is a shared circuit, so users must compete for the shared capacity, but most installations are limited to a few dozen users. Data transmission for both versions ranges from 1.5 to 54 Mbps, depending on the vendor.

Other fixed wireless technologies such as *satellite* are also available. Satellite technologies use the satellite for downstream transmissions (from the ISP to the customer) but use traditional dial-up modems for upstream transmissions. Although satellite technology has been available for several years, it has never become really popular.

Mobile Wireless

Mobile wireless technologies enable users to access the Internet from any location where there is mobile wireless service. Widespread mobile wireless Internet access is probably the next major change in networking. Mobile wireless Internet access technologies exist today (e.g., cell phone connections), but most are slow compared with wired access, whether DSL, cable modem, or simply a dial-up modem. The WLAN technologies discussed in Chapter 7 (e.g., 802.11g) are primarily intended for use inside one organization although they are being installed in public places such as airports for open access to the Internet.

Wireless Application Protocol (WAP) provides a set of application and network protocols called the *Wireless Application Environment (WAE)* to support mobile wireless Internet applications. WAP is designed to enable the use of normal Web applications on

MANAGEMENT

10-3 BANKING ON WIRELESS APPLICATION PROTOCOL

FOCUS

SkandiaBanken, a leading Swedish Internet bank, provides retail banking services to over 350,000 clients. The bank operates entirely via the Internet, having no traditional branches. Customers communicate with the bank through the Web, e-mail, telephones, and now Wireless Application Protocol (WAP) technology.

SkandiaBanken chose to implement WAP to provide customers with safe and easy access to their financial data from anywhere in the world. Customers can securely view their deposit and credit card accounts, execute their credit card

payments, make balance inquiries, and pay their bills. Users of the mobile service may also check the foreign exchange, gold, and treasury bill rates, as well as access information on the bank's range of financial products. There is even location-based information available, such as city guides, restaurant reviews, movies, theaters, museums, art galleries, libraries, and other important facilities.

SOURCE: "Financial Institutions Worldwide Use Infinite WAP Server to Offer Mobile Banking," Infinite.com.

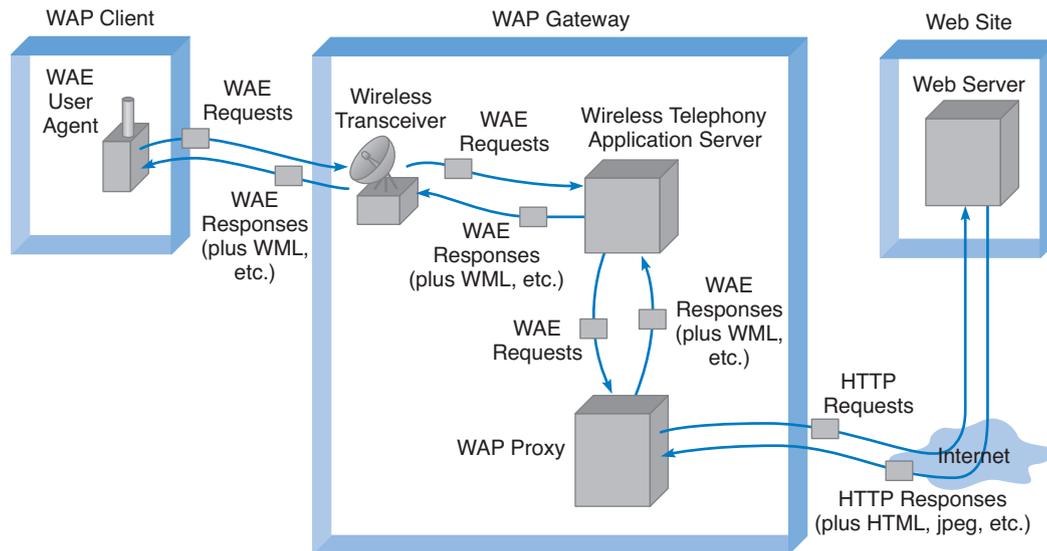


FIGURE 10.10 Mobile wireless architecture for Wireless Application Protocol (WAP) applications. HTML = Hypertext Markup Language; HTTP = Hypertext Transfer Protocol; WAE = Wireless Application Environment; WML = Wireless Markup Language.

computers and devices with small display screens operating over low-speed wireless connections. Figure 10.10 shows the basic WAP architecture.

The WAP client (a mobile phone, palm computer, or laptop computer) runs special WAP software called a WAE user agent. This software generates WAE requests that are similar in many ways to HTTP requests and transmits them wirelessly to a WAP gateway. A transceiver at the WAP gateway passes the requests to a *wireless telephony application (WTA) server*. This server responds to the requests and, if the client has requested a Web page on the Internet, sends a WAE request to a *WAP proxy*. The WAP proxy translates the WAE request into HTTP and sends it over the Internet to the desired Web server.

This Web server responds to the request and sends back to the WAP proxy an HTTP response that contains HTML, JPEG, and other Internet application protocols. The WAP proxy in turn translates these into their WAE equivalents and sends them to the WTA server, which sends them to the client.

Future Technologies

Internet access technologies are one of the fastest growth areas in networking, so there are several new technologies that have the potential to become important alternatives to DSL, cable modems, and wireless technologies. In this section, we focus on two up-and-coming technologies: passive optical networking (PON) and Ethernet.

Passive Optical Networking *Passive optical networking (PON)*, sometimes called *fiber-to-the-home (FTTH)*, is exactly what it sounds like: running fiber-optic cable

into the home. The traditional set of hundreds of copper telephone lines that run from the telephone company switch office is replaced by one fiber-optic cable that is run past each house or office in the neighborhood. Data is transmitted down the signal fiber cable using wavelength division multiplexing (WDM), providing hundreds or thousands of separate channels. At each subscriber location, a fiber splitter separates the channels belonging to that location and runs them into an optical electrical converter, which then connects to an Ethernet switch.

This approach is called passive optical because the splitters require no electrical current and thus are quicker and easier to install than traditional electrical-based hubs and repeaters. However, because they are passive, the optical signal fades quickly, giving a maximum length of about 10 miles.

Each single fiber has a capacity of about 155 Mbps, which must be allocated among the subscribers. This means about 1.5 Mbps if there are 100 subscribers per fiber, or 15 Mbps if there are only 10 subscribers. At present, there are no standards for PON and FTTH, but several vendors have joined together to develop standards. The larger problem, of course, is the cost of laying miles and miles of fiber-optic cable.

Ethernet to the Home Perhaps the most exciting possibility is Ethernet to the home. If we were to start over and design an entirely new network for Internet access from home, we would probably start with Ethernet because of its low cost and popularity in organizational LANs. Using common protocols would make the whole task of networking much simpler for everyone involved.

Pioneered by *Yipes.com*, such an approach is exactly what is being used in several major U.S. cities. With this approach, the common carrier installs a TCP/IP router with 10Base-T or 100Base-T connections into the customer's network and an Ethernet fiber on the other. The IP/Ethernet traffic moves from the router into the carrier's Ethernet MAN and then onto the Internet.

Although this approach is also limited because of the cost of providing Ethernet fiber to the customer, we believe this has great potential. Because conversions between protocols are not required at the customer site, connecting to the network is much simpler than with other Internet access technologies.

INTERNET GOVERNANCE

Because the Internet is a network of networks, no one organization operates the Internet. The closest thing the Internet has to an owner is the *Internet Society (ISOC)* (www.isoc.org). ISOC is an open-membership professional society with more than 175 organizational and 8,000 individual members in over 100 countries, including corporations, government agencies, and foundations that have created the Internet and its technologies. Because membership in ISOC is open, anyone, including students, is welcome to join and vote on key issues facing the Internet.

The ISOC mission is to ensure "the open development, evolution and use of the Internet for the benefit of all people throughout the world."⁵ ISOC works in three general areas:

⁵See www.isoc.org/isoc/mission.

public policy, education, and standards. In terms of public policy, ISOC participates in the national and international debates on important issues such as censorship, copyright, privacy, and universal access. ISOC delivers training and education programs targeted at improving the Internet infrastructure in developing nations. The most important ISOC activity lies in the development and maintenance of Internet standards. ISOC works through four interrelated standards bodies: Internet Engineering Task Force (IETF), Internet Engineering Steering Group (IESG), Internet Architecture Board (IAB), and Internet Research Task Force (IRTF).

The *Internet Engineering Task Force (IETF)* (www.ietf.org) is a large, open international community of network designers, operators, vendors, and researchers concerned with the evolution of the Internet architecture and the smooth operation of the Internet. IETF works through a series of working groups, which are organized by topic (e.g., routing, transport, security). The *requests for comment (RFCs)* that form the basis for Internet standards are developed by the IETF and its working groups.

Closely related to the IETF is the *Internet Engineering Steering Group (IESG)*. The IESG is responsible for technical management of IETF activities and the Internet standards process. It administers the process according to the rules and procedures that have been ratified by the ISOC trustees. The IESG is directly responsible for the actions associated with entry into and movement along the Internet “standards track,” including final

TECHNICAL**10-1 REGISTERING AN INTERNET DOMAIN NAME****FOCUS**

Until the 1990s, there was only a moderate number of computers on the Internet. One organization was responsible for registering domain names (sets of application layer addresses) and assigning IP addresses for each top-level domain (e.g., .COM). Network Solutions, for example, was the sole organization responsible for domain name registrations for the .COM, .NET, and .ORG domains. In October 1998, the *Internet Corporation for Assigned Names and Numbers (ICANN)* was formed to assume responsibility for the IP address space and domain name system management.

In spring 1999, ICANN established the Shared Registration System (SRS) that enabled many organizations to perform domain name registration and address assignment using a shared database. More than 80 organizations are now accredited by ICANN as registrars and are permitted to use the SRS. Each registrar has the right to assign names and addresses in one or

more top-level domains. For a list of registrars and the domains they serve, see www.internic.com.

If you want to register a new domain name and obtain an IP address, you can contact any accredited registrar for that top-level domain. One of the oldest privately operated registrars is register.com. Each registrar follows the same basic process for registering a name and assigning an address, but each may charge a different amount for their services. In order to register a name, you must first check to see if it is available (i.e., that no one else has registered it). If the name has already been registered, you can find out who owns it and perhaps attempt to buy it from them.

If the domain name is available, you will need to provide the IP address of the DNS server that will be used to store all IP addresses in the domain. Most large organizations have their own DNS servers, but small companies and individuals often use the DNS of their ISP.

approval of specifications as Internet standards. Each IETF working group is chaired by a member of the IESG.

Whereas the IETF develops standards and the IESG provides the operational leadership for the IETF working groups, the *Internet Architecture Board (IAB)* provides strategic architectural oversight. The IAB attempts to develop conclusions on strategic issues (e.g., top-level domain names, use of international character sets) that can be passed on as guidance to the IESG or turned into published statements or simply passed directly to the relevant IETF working group. In general, the IAB does not produce polished technical proposals but rather tries to stimulate action by the IESG or the IETF that will lead to proposals that meet general consensus. The IAB appoints the IETF chairperson and all IESG members, from a list provided by the IETF nominating committee. The IAB also adjudicates appeals when someone complains that the IESG has failed.

The *Internet Research Task Force (IRTF)* operates much like the IETF through small research groups focused on specific issues. Whereas IETF working groups focus on current issues, IRTF research groups work on long-term issues related to Internet protocols, applications, architecture, and technology. The IRTF chairperson is appointed by the IAB.

INTERNET 2

The Internet is changing. New applications and access technologies are being developed at lightning pace. But these innovations do not change the fundamental structure of the Internet. It has evolved more slowly because the core technologies (TCP/IP) are harder to change gradually; it is difficult to change one part of the Internet without changing the attached parts.

Many organizations in many different countries are working on dozens of different projects in an attempt to design new technologies for the next version of the Internet.⁶ The two primary American projects working on the future Internet got started at about the same time in 1996. The U.S. National Science Foundation provided \$100 million to start the *Next Generation Internet (NGI)* program, which developed the *very-high-performance Backbone Network Service (vBNS)* now run by MCI WorldCom, and 34 universities got together to start what turned into the *University Corporation for Advanced Internet Development (UCAID)*, which developed the *Abilene network*, commonly called *Internet 2*. In 1997, the Canadian government established the *Advanced Research and Development Network Operations Center (ARDNOC)*, which developed *CA*net*, the Canadian project on the future Internet.⁷

Figure 10.11 shows the major high-speed circuits in the Internet 2 Abilene network and the CA*net network. All the major circuits in these networks are OC-192 (10 Gbps). The two networks peer in Seattle, Chicago, and New York. National Lambda Rail (www.nlr.net) is another major high-speed network that is experimenting with long distance Ethernet (10 GbE) running over fiber-optic circuits.

Each of the networks has a set of access points called *gigapops*, so named because they provide a point of presence at gigabit speeds. Although traditional Internet NAPs

⁶For a listing of several major international projects, see www.startup.net.

⁷For more information on these projects, see www.internet2.org and www.canarie.ca.

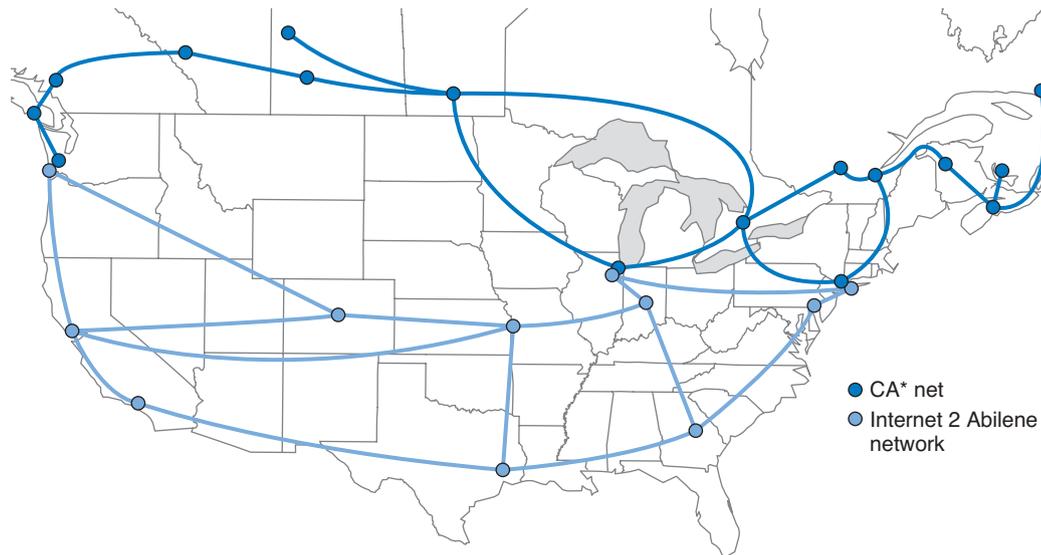


FIGURE 10.11 Selected high-speed Internet backbones.

provide connections between networks at T1, T3, OC-1, OC-3, and—occasionally—OC-12 speeds, gigapops are designed to provide access at much higher speeds so that different networks can exchange data at much higher rates of speed, usually OC-198 or 10 Gbps. Gigapops also usually provide a wider range of services than traditional NAPs that are primarily just data exchange points.

Besides providing very high-speed Internet connections, these networks are intended to experiment with new protocols that one day may end up on the future Internet.

MANAGEMENT

10-4 INSIDE THE PACIFIC/NORTHWEST GIGAPOP

FOCUS

The Pacific/Northwest Gigapop is located in Seattle, Washington, and is run by the University of Washington and University Corporation for Advanced Internet Development (i.e., Internet 2). It provides gigabit Ethernet and SONET OC-192 (10 Gbps) connections to several high-speed networks such as Abilene, CA*net, Microsoft, and the Defense Research and Engineering Network, which is funded by the U.S. Department of Defense. It also provides a network access point for these high-speed networks to connect to lower-speed networks of the tradi-

tional Internet, such as those run by Sprint, AT&T, Singapore's SingAREN, and Australia's AARNet, as well as a number of universities in the Pacific Northwest.

The basic core of the gigapop is a set of two high-speed switches, connected to two high-speed routers. High-speed networks, such as Abilene, connect directly into the core devices whereas lower-speed networks connect into the core via a set of routers.

SOURCE: www.pnw-gigapop.net.

For example, most of these networks run IPv6 as the primary network layer protocol, rather than IPv4. Most are also working on new ways to provide quality of service (QoS) and multicasting. Some, such as Internet 2, are also working on developing new applications for a high-speed Internet, such as tele-immersion and videoconferencing.

IMPLICATIONS FOR MANAGEMENT

Several years ago, there was great concern that the Internet would reach capacity. The growth of traffic on the Internet was increasing significantly faster than the construction of new Internet circuits; several experts predicted the collapse of the Internet. It did not happen for the simple reason that companies could make money by building new circuits and charging for their use. Today, there are a large number of fiber-optic circuits that have been built but not yet been turned on. Wavelength division multiplexing technologies mean that 10–20 times more data can now be transmitted through the same fiber-optic cable (see Chapter 3). Many countries, companies, and universities are now building the Next Generation Internet using even newer, experimental, very high-speed technologies. The Internet will not soon run out of capacity.

In recent years, there has been a blossoming of new “broadband” technologies for higher speed Internet access. Individuals and organizations can now access the Internet at relatively high speeds—much higher speeds than we would have even considered reasonable 5–10 years ago. This means that it is now simple to move large amounts of data into most homes and businesses in North America. As a result, software applications that use the Internet can provide a much richer multimedia experience than ever before.

In previous chapters, we have described how there has been a significant reduction in a number of different technologies in use in LANs, backbones, MANs, and WANs over the past few years. We are about to enter that stage with regard to Internet access technologies. Today there are many choices; over the next two years a few dominant standards will emerge, and the market will solidify around those standards. Organizations that invest in the technologies that ultimately become less popular will need to invest significant funds to replace those technologies with the dominant standards. The challenge, of course, is to figure out which technology standards will become dominant. Will it be cable modem and DSL, or Ethernet to the home? Only time will tell.

SUMMARY

How the Internet Works The Internet is a set of separate networks, ranging from large national ISPs to midsize regional ISPs to small local ISPs, that connect with one another at NAPs and MAEs. NAPs and MAEs charge the ISPs to connect, but similar-sized ISPs usually do not charge each other to exchange data. Each ISP has a set of points of presence through which it charges its users (individuals, businesses, and smaller ISPs) to connect to the Internet. Users connect to a POP to get access to the Internet. This connection may be via a dial-up modem over a telephone line or via a higher-speed circuit such as a T1.

DSL DSL enables users to connect to an ISP POP over a standard point-to-point telephone line. The customer installs a DSL modem that connects via Ethernet to his or her computer system. The modem communicates with a DSLAM at the telephone company office, which sends the data to the ISP POP.

ADSL is the most common type of DSL and often provides 1.5 Mbps downstream and 384 Kbps upstream. VDSL is a faster version that runs over short distances and has speeds up to 51.8 Mbps.

Cable Modem Cable modems use a shared multipoint circuit that runs through the cable TV cable. They also provide the customer with a modem that connects via Ethernet to his or her computer system. The modem communicates with a CMTS at the cable company office, which sends the data to the ISP POP. The DOCSIS standard is the dominant standard, but there are no standard data rates today. Typical downstream speeds range between 768 Kbps and 1.5 Mbps, and typical upstream speeds range between 200 Kbps and 1.5 Mbps.

Wireless Fixed wireless systems provide DSL-like speeds over a single line-of-sight wireless circuit to a multitenant building. Inside the building, DSL is used to provide service to a large number of users over the existing phone lines. Mobile wireless uses cellular telephone technologies to provide access to small hand-held devices using WAP. WAP translates from traditional Internet protocols such as HTTP and HTML into their WAE equivalents for use in the small devices.

Internet Governance The closest the Internet has to an owner is the ISOC, which works on public policy, education, and Internet standards. Standards are developed through four related organizations governed by ISOC. The IETF develops the actual standards through a series of working groups. The IESG manages IETF activities. The IAB sets long-term strategic directions, and the IRTF works on future issues through working groups in much the same way as the IETF.

Internet 2 There are many different organizations currently working on the next generation of the Internet, including the Abilene network, vBNS, and CA*net. Although each is working in a slightly different fashion, all join together with one another and parts of the regular Internet at gigapops (gigabit points of presence).

KEY TERMS

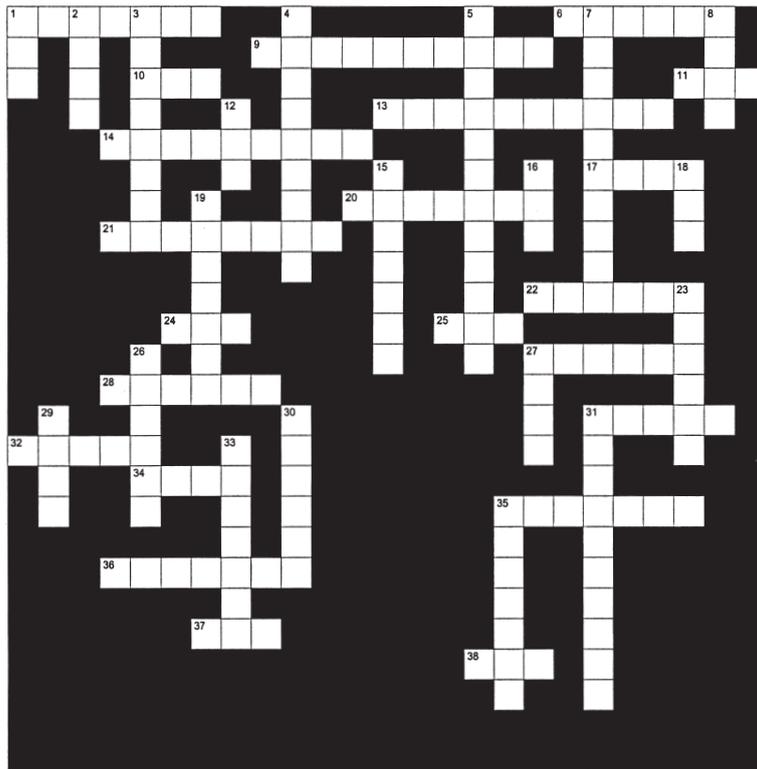
Abilene network	DSL modem	local loop	University Corporation for
Advanced Research and	fiber-to-the-home (FTTH)	main distribution facility	Advanced Internet De-
Development Network	fixed wireless	(MDF)	velopment (UCAID)
Operations Center	G.Lite ASDL	metropolitan area ex-	very-high-data-rate digital
(ARDNOC)	hybrid fiber coax (HFC)	change (MAE)	subscriber line (VDSL)
asymmetric DSL (ADSL)	Internet Architecture	mobile wireless	very-high-performance
autonomous systems	Board (IAB)	national ISP	Backbone Network Ser-
broadband technologies	Internet Corporation for	network access point	vice (vBNS)
cable modem	Assigned Names and	(NAP)	WAP proxy
cable modem termination	Numbers (ICANN)	Next Generation Internet	Wireless Application En-
system (CMTS)	Internet Engineering	(NGI)	vironment (WAE)
CA*net	Steering Group (IESG)	optical-electrical (OE)	Wireless Application Pro-
customer premises	Internet Engineering Task	converter	toocol (WAP)
equipment (CPE)	Force (IETF)	passive optical networking	wireless DSL
Data over Cable Service	Internet Research Task	(PON)	wireless telephony appli-
Interface Specification	Force (IRTF)	peering	cation (WTA) server
(DOCSIS)	Internet service provider	point of presence (POP)	Yipes.com
digital subscriber line	(ISP)	regional ISP	
(DSL)	Internet Society (ISOC)	remote-access server	
distribution hub	Internet 2	(RAS)	
DSL access multiplexer	line splitter	request for comment	
(DSLAM)	local ISP	(RFC)	

QUESTIONS

1. What is the basic structure of the Internet?
2. Explain how the Internet is a network of networks.
3. Compare and contrast an NAP and a MAE.
4. What is a POP?
5. Explain one reason why you might experience long response times in getting a Web page from a server in your own city.
6. What type of circuits are commonly used to build the Internet today? What type of circuits are commonly used to build Internet 2?
7. Compare and contrast cable modem and DSL.
8. Explain how DSL works.
9. How does a DSL modem differ from a DSLAM?
10. Explain how ADSL works.
11. Explain how VDSL works.
12. Compare and contrast ADSL and VDSL.
13. Explain how a cable modem works.
14. What is an OE converter? A CMTS?
15. Which is better, cable modem or DSL? Explain.
16. Explain how one type of fixed wireless called wireless DSL works.
17. Compare and contrast mobile wireless and fixed wireless.
18. Explain how WAP works.
19. What are some future technologies that might change how we access the Internet?
20. What is PON, and how does it work?
21. Explain how Ethernet to the home works.
22. What are the principal organizations responsible for Internet governance, and what do they do?
23. How is the IETF related to the IRTF?
24. What are two principal American organizations working on the future of the Internet?
25. What is Internet 2?
26. What is a gigapop?
27. There are many different organizations working on their vision of a high-speed Internet. Is this good or bad? Would we be better off just having one organization working on this and coordinating the work?
28. Today, there is no clear winner in the competition for higher-speed Internet access. What technology or technologies do you think will dominate in 2 years' time? Why?
29. Some experts believe that in 5 years, the modem will have disappeared. What do you think?
30. Many experts predicted that small, local ISPs would disappear as regional and national ISPs began offering local access. This hasn't happened. Why?

EXERCISES

- 10-1. Describe the current network structure of the Abilene network, the vBNS network, and the CA*net network.
- 10-2. Provide the service details (e.g., pricing) for at least two high-speed Internet access service providers in your area.
- 10-3. Many people are wiring their homes for 10Base-T or 100Base-T. Suppose a friend who is building a house asks you what—if any—network to put inside the house and what Internet access technology to use. What would you recommend?
- 10-4. Explore the products available to install wireless Internet access in your home. How much would it cost to install and how much would it cost each month?
- 10-5. See puzzle on page 362.



Across

- | | |
|--|---|
| 1. One of the Internet 2 networks | 24. One approach to WLAN security |
| 6. The backbone layer closet to the users | 25. A place you can connect into the Internet |
| 9. Ethernet's media access control approach | 27. A layer 3 device connecting devices over a backbone |
| 10. Device used to connect computers together in a LAN | 28. A "smart" device that connects computers in a LAN |
| 11. A point where regional ISPs connect | 31. A common type of WAN cable |
| 13. Media access control is most WLANs used both contention and _____ access | 32. A fast dedicated circuit WAN technology |
| 14. High speed Internet access technologies are often called this | 34. A common type of LAN cable |
| 17. A common WAN architecture | 35. T1 speed |
| 20. A common speed in LANs | 36. OC-12 speed |
| 21. The most common LAN technology in use today | 37. A common Internet access technology |
| 22. A common WLAN speed | 38. Internet-based WAN circuits |

Down

1. A fast packet switching WAN technology
2. I still don't need it
3. A leading contender for future Internet access
4. A _____ backbone is a common design
5. The middle layer in a backbone
7. A shared circuit Internet access technology
8. Ethernet's physical topology
12. Main distribution facility
15. A common type of WLAN
16. Another name for T3
18. A layer 1 device connecting devices over a backbone
19. breaking a LAN into several of these may improve performance
23. You should never put a _____ on a WLAN
26. A layer 2 device connecting devices over a backbone
27. Most backbones today use these to hold equipment
29. The inner most layer in a backbone
30. T3 speed
31. A common packet switching WAN technology
33. OC-3 speed
35. A typical cable modem speed

MINI-CASES

I. Cathy's Collectibles

Your cousin Cathy runs a part-time business out of her apartment. She buys and sells collectibles such as antique prints, baseball cards, and cartoon cells and has recently discovered the Web with its many auction sites. She has begun buying and selling on the Web by bidding on collectibles at lesser-known sites and selling them at a profit at more well-known sites. She downloads and uploads lots of graphics (pictures of the items she's buying and selling). She is getting frustrated with the slow Internet access she has with her 56-Kbps dial-up modem and asks you for advice. DSL is available at a cost of \$60 per month for 1.5 Mbps down and 384 Kbps up. Cable modem service is available for a cost of \$50 per month for 1.5 Mbps down and 640 Kbps up. Wireless DSL is available in her apartment building for \$45 per month for 1.5 Mbps down and 256 Kbps up. Explain the differences in these services and make a recommendation.

II. Surfing Sam

Sam likes to surf the Web for fun, to buy things, and to research for his classes. Suppose the same Internet access technologies are available as in mini-case I above. Explain the differences in these services and make a recommendation.

III. Cookies Are Us

Cookies Are Us runs a series of 100 cookie stores across the midwestern United States and central Canada. At the end of each day, the stores express-mail a diskette or two of sales and inventory data to headquarters, which uses the data to ship new inventory and plan marketing campaigns. They have decided to move data over a WAN or the Internet. What type of a WAN topology and service (see Chapter 9) or Internet connection would you recommend? Why?

IV. Organic Foods

Organic Foods operates organic food stores in Toronto. The store operates like a traditional grocery store but offers only organically grown produce and meat, plus a wide array of health food products. Organic Foods sells memberships, and its 3,000 members receive a discount on all products they buy. There are also special member events and sales promotions each month. Organic Foods wants to open a new Internet site that will enable it to e-mail its members monthly and provide up-to-date information and announcements about new products, sales promotions, and member events on its Web site. It has two options. First, it could develop the software on its own server in its office and connect the office (and the server) to the Internet via an ISDN, DSL, T1, or similar connection from its offices to an ISP. Alternately, it could pay the ISP to host the Web site on its servers and just connect the office to the ISP for Internet service. Costs for several Internet access options are present in mini-case I above. In addition, ISDN service costs \$120 per month for BRI and \$1,500 per month (plus \$1,000 to install) for PRI; T1 service would cost \$1,000 to install and \$1,200 per month to operate; frame relay would cost \$1,000 to install and \$500 per month for 256Kps or \$750 for 1.5 Mbps. Web hosting would cost \$100–400 per month, depending upon the traffic. Which would you recommend and what size of an Internet connection would you recommend? Justify your choice.

CASE STUDY

NEXT-DAY AIR SERVICE

See the Web site.

HANDS-ON ACTIVITY

Seeing the Internet

The Internet is a network of networks. One way to see this is by using the VisualRoute software. VisualRoute is a commercial package, but provides a demonstration on its Web site. Go to visualroute.visualware.com and register to use their free service. Then enter a URL and watch as the route from your computer to the destination is traced and graphed. Figure 10.12 shows the route from my house in Indiana to the City University of Hong Kong.

Another interesting site is the Internet Traffic Report (www.internettrafficreport.com). This site shows how busy the parts of the Internet are in real time. The main page enables you to see the current status of the major parts of the world, including a "traffic index" that rates performance on a 100 point scale. You can also see the average response time at key Internet NAPs, MAEs, and peering points (at least those that have agreed to be monitored), which is an average of 135 milliseconds as I write this. It also shows the global packet loss rates—the per-

cent of packets discarded due to transmission errors—(an average of 3 percent today).

By clicking on a region of the world you can see the same statistics for routers in that region. If you click on a specific router you can see a graph of its performance over the past 24 hours. Figure 10.13 shows the statistics for one router operated by Sprint.

You can also get traffic reports for Internet 2 (see loadrunner.uits.iu.edu/weathermaps/Abilene). Figure 10.14 shows the "weathermap" on the Internet 2 Abilene network. Each circuit is color coded (although it's hard to see in this two-color figure). The weathermap shows traffic in both directions because the circuits are full duplex. The circuit from Atlanta to Washington, for example, is running at 5 percent of capacity, while the circuit from Washington to Atlanta is running at 10 percent of capacity. You can also click on any circuit to see a graph of traffic over the last 24 hours.

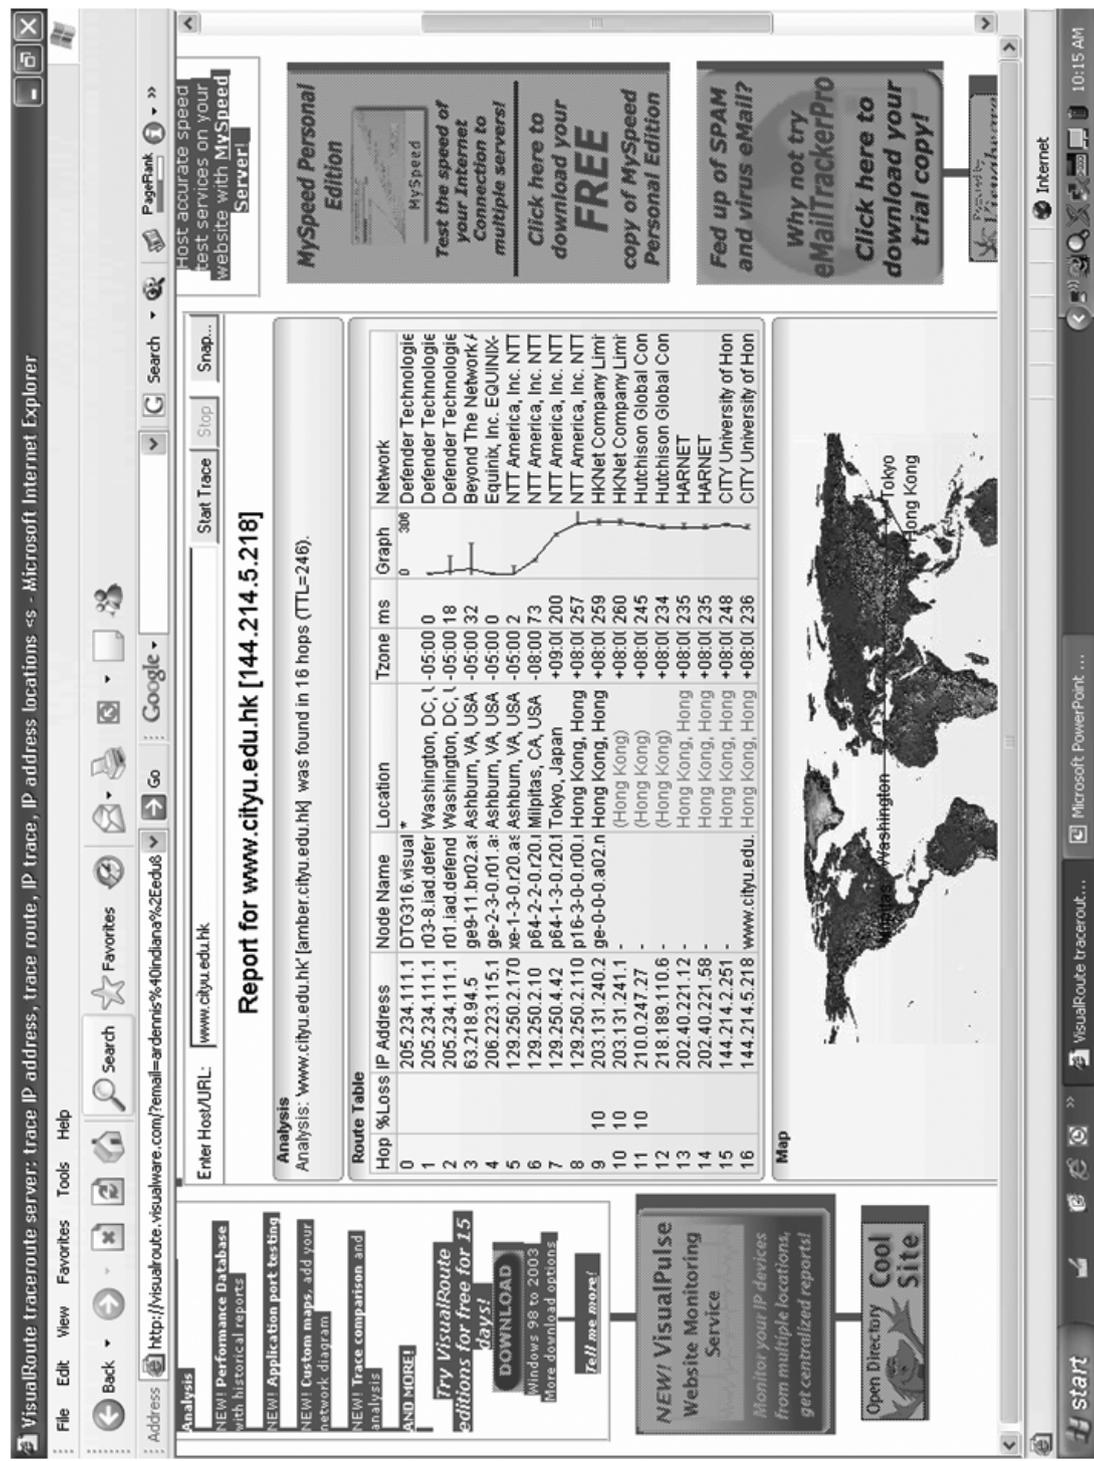


FIGURE 10.12 Visual trace route.

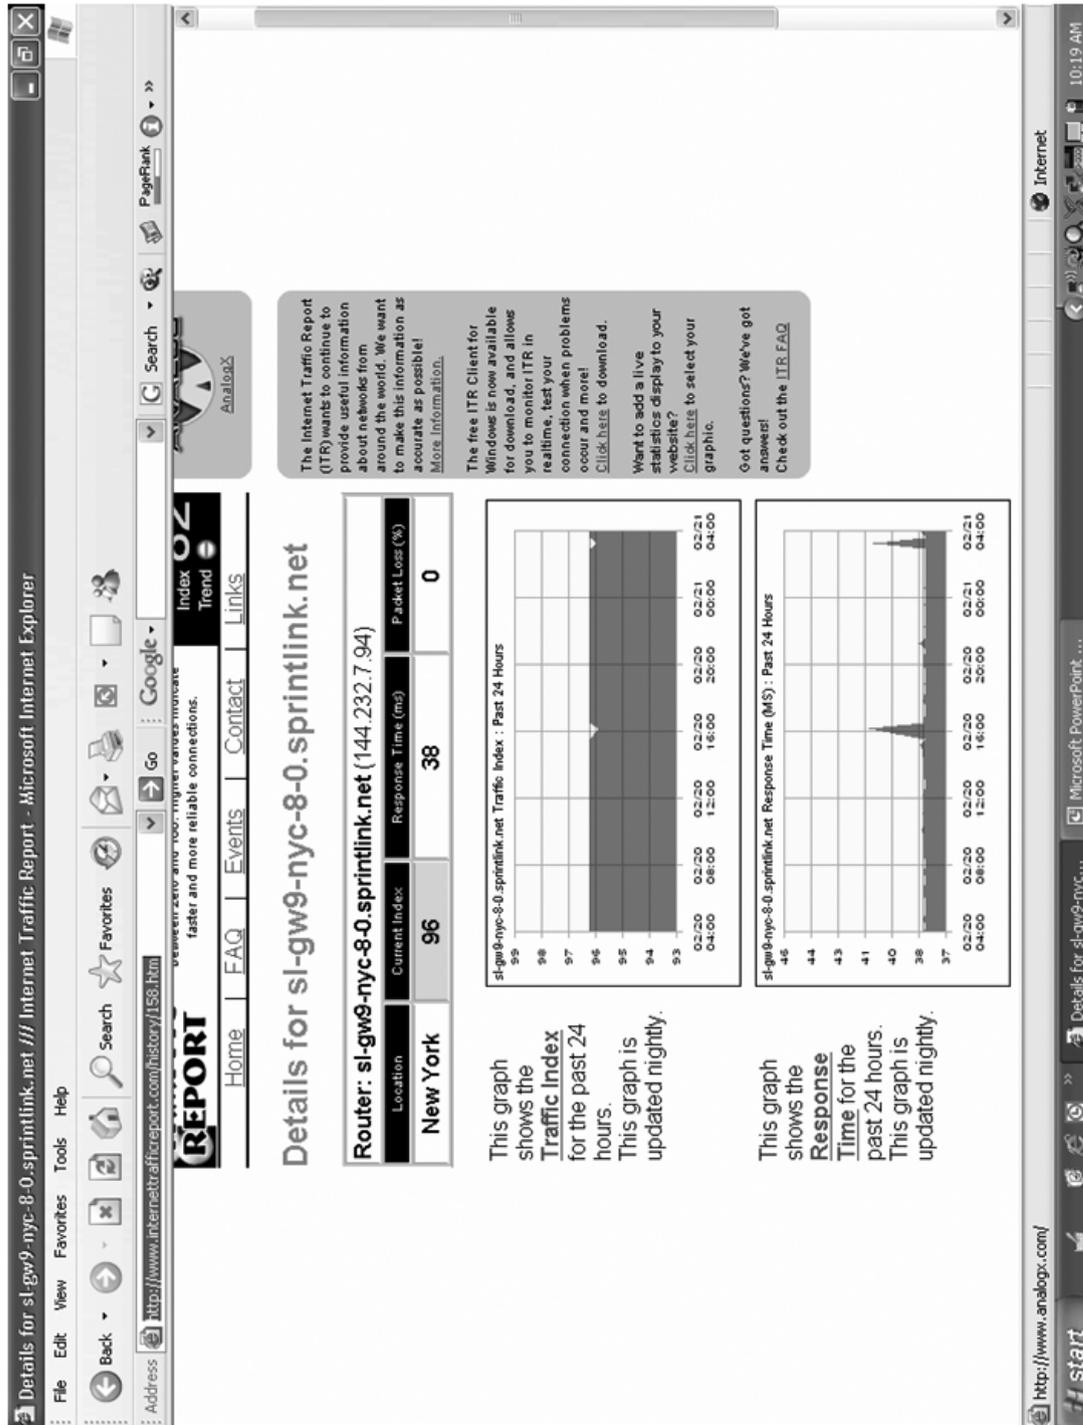


FIGURE 10.13 Internet traffic reports.

