

Telecommunications Essentials

The Complete Global Source
for Communications Fundamentals,
Data Networking and the Internet,
and Next-Generation Networks

Lillian Goleniewski

◆ Addison-Wesley

Boston • San Francisco • New York • Toronto • Montreal
London • Munich • Paris • Madrid
Capetown • Sydney • Tokyo • Singapore • Mexico City

Many of the designations used by manufacturers and sellers to distinguish their products are claimed as trademarks. Where those designations appear in this book, and Addison-Wesley, Inc. was aware of a trademark claim, the designations have been printed with initial capital letters or in all capitals.

Lido Telecommunications Essentials® is the registered trademark of The Lido Organization, Inc.

The author and publisher have taken care in the preparation of this book, but make no expressed or implied warranty of any kind and assume no responsibility for errors or omissions. No liability is assumed for incidental or consequential damages in connection with or arising out of the use of the information or programs contained herein.

The publisher offers discounts on this book when ordered in quantity for special sales. For more information, please contact:

Pearson Education Corporate Sales Division
201 W. 103rd Street
Indianapolis, IN 46290
(800) 428-5331
corpsales@pearsoned.com

Visit AW on the Web: www.aw.com/cseng/

Library of Congress Cataloging-in-Publication Data

Goleniewski, Lillian.

Telecommunications essentials : the complete global source for communications fundamentals, data networking and the Internet, and next-generation networks / Lillian Goleniewski.

p. cm.

Includes bibliographical references and index.

ISBN 0-201-76032-0

1. Telecommunication. I. Title.

TK5101 G598 2002

621.382—dc21

2001053752

Copyright © 2002 by Pearson Education, Inc.

All rights reserved. No part of this publication may be reproduced, stored in a retrieval system, or transmitted, in any form, or by any means, electronic, mechanical, photocopying, recording, or otherwise, without the prior consent of the publisher. Printed in the United States of America. Published simultaneously in Canada.

For information on obtaining permission for use of material from this work, please submit a written request to:

Pearson Education, Inc.
Rights and Contracts Department
75 Arlington Street, Suite 300
Boston, MA 02116
Fax: (617) 848-7047

ISBN 0-201-76032-0

Text printed on recycled paper

1 2 3 4 5 6 7 8 9 10—CRS—0504030201

First printing, December 2001

The Internet: Infrastructure and Service Providers

This chapter explores the Internet, including how the Internet actually works, the structure of the Internet in terms of the various levels of service providers, and the organization and characteristics of the growing variety of service providers.

■ Internet Basics

Figure 9.1 is an astounding graph that speaks to the pace of Internet development. It shows the number of years it took a number of technologies to reach 50 million users worldwide. As you can see, whereas it took 74 years for the telephone to reach 50 million users, it took the World Wide Web only 4.

What forces are propelling our interest in the Internet? One main force is that usage is increasing dramatically; today some 250 million people worldwide have Internet access, and that number is growing by leaps and bounds. The Internet is very useful and easy to use, and for a growing number of people in the developed world, it is now the first place to look for information. As one colleague recently told me, in the past week, besides the numerous times he had used the Internet to get information for my work, he'd used the Internet to look up hotels for a weekend break, to determine what concerts are on in Dublin, to check the specification of a car, to transfer funds between bank accounts, to find the address of an old friend, and to obtain sheet music. Electronic commerce (e-commerce) is also growing, in both the business-to-consumer and business-to-business sectors. Another




Figure 9.1 Internet pace: Years to reach 50 million users worldwide

contributor is the major shift toward the use of advanced applications, including pervasive computing, which introduces a wide range of intelligent appliances that are ready to communicate through the Internet, as well as applications that include the more captivating visual and sensory streams. Finally, the availability of broadband, or high-speed access technologies, further drives our interest in and our ability to interact with Web sites that involve the use of these advanced applications and offer e-commerce capabilities.

A Brief History of the Internet

To help understand the factors that contributed to the creation of the Internet, let's look very briefly at the history of the Internet. In 1969 the Advanced Research Projects Agency (ARPA) of the U.S. Department of Defense initiated a project to develop a distributed network. There were several reasons for doing this. First, the project was launched during the Cold War era, when there was an interest in building a network that had no single point of failure, and that could sustain an attack yet continue to function. Second, four supercomputer centers were located in four universities throughout the United States, and we wanted to connect them together so that we could engage in some more intensive processing feats. So, the Internet started as a wide area, packet-switching network called the ARPANET.

Toward the mid-1970s, ARPA was renamed the Defense Advanced Research Projects Agency (DARPA), and while it was working on the distributed, or packet-switched, network, it was also working on local area networks (LANs), paging networks, and satellite networks. DARPA recognized that there was a need for some form of internetworking protocol that would allow open communications between

Jonathan Postel and the Internet

Jonathan Postel played a pivotal role in creating and administering the Internet. He was one of a small group of computer scientists who created the ARPANET, the precursor to the Internet. For more than 30 years he served as editor of the Request for Comments (RFC) series of technical notes that began with the earliest days of the ARPANET and continued into the Internet. Although intended to be informal, RFCs often laid the foundation for technical standards governing the Internet's operation. Nearly 2,500 RFCs have been produced.

Also for 30 years, Postel handled the administrative end of Internet addresses, under the auspices of the Internet Assigned Numbers Authority (IANA), a U.S. government-financed entity. As part of the effort to hand over administration of the Internet to an international private corporation, Postel delivered a proposal to the government for transforming IANA into a nonprofit corporation with broad representation from the commercial and academic sectors. That organization is today known as the Internet Corporation for Assigned Names and Numbers (ICANN).

disparate networks. So, Internet Protocol (IP) was created to support an open-architecture network that could link multiple disparate networks via gateways—what we today refer to as *routers*.

In 1980, Transmission Control Protocol/Internet Protocol (TCP/IP) began to be implemented on an experimental basis, and by 1983, it was required in order for a subnetwork to participate in the larger virtual Internet.

The original Internet model was not based on the telephone network model. It involved distributed control rather than centralized control, and it relied on cooperation among its users, which initially were largely academicians and researchers. With the original Internet, there's no regulation, no monopoly, and no universal service mandate (although these issues are being considered seriously now).

Today, no one agency is in charge of the Internet, although the Internet Society (ISOC) is a nonprofit, nongovernmental, international organization that focuses on Internet standards, education, and policy issues. ISOC is an organization for Internet professionals that serves as the organizational home of the Internet Engineering Task Force (IETF), which oversees various organizational and coordinating tasks. ISOC is composed of a board of trustees, the Internet Architecture Board, the IETF, the Internet Research Task Force, the Internet Engineering Steering Group, and the Internet Research Steering Group.

The IETF is an international community of network designers, operators, vendors, and researchers, whose job is to evolve the Internet and smooth its operation by creating technical standards through consensus. Other organizations that are critical to the functioning of the Internet include American Registry for Internet Numbers (ARIN) in the United States, Asia Pacific Network Information Center

Prevailing Conditions and Exchange Points

Since the beginning of the Internet's history, we've been trying to prevent having a single point of failure. We have distributed nodes throughout the network so that if one node goes down or a series of links goes down, there can still be movement between the other devices, based on a wide variety of alternative nodes and links.

But we're doing a turnaround now because these very interconnection points that provide interconnection between ISPs can also act as vulnerable points for the network and even for a nation. If the exchange points are taken down within a given country, the Internet activity within that country may cease or fail altogether, with great economic consequences. Always remember that, in the end, the prevailing conditions dictate whether an architecture is truly good, reliable, and high performance.

(APNIC) in Asia-Pacific, and RIPE NCC (Reseaux IP Europeens Network Coordination Center) in Europe. These organizations manage and sell IP addresses and autonomous system numbers. IANA manages and assigns protocol and port number, and ICANN (formed in 1998) is responsible for managing top-level domain names and the root name servers. ICANN also delegates control for domain name registry below the top-level domains. (Domain names and the role of IP addresses are discussed later in this chapter.)

What the Internet Is and How It Works

To understand the Internet, it's important to first understand the concept of a computer network (see Figure 9.2). A network is formed by interconnecting a set of computers, typically referred to as *hosts*, in such a way that they can interoperate with one another. Connecting these hosts involves two major components: hardware (that is, the physical connections) and software. The software can be run on the same or dissimilar host operating systems, and it is based on standards that define its operation. These standards, referred to as *protocols*, provide the formats for passing packets of data, specify the details of the packet formats, and describe how to handle error conditions. The protocols hide the details of network hardware and permit computers of different hardware types, connected by different physical connections, to communicate, despite their differences. (Protocols are discussed in detail later in this chapter.)

In the strictest sense, the Internet is an internetwork composed of a worldwide collection of networks, routers, gateways, servers, and clients, that use a common set of telecommunications protocols—the IP family—to link them together (see Figure 9.3). The term *client* is often used to refer to a computer on the network




Figure 9.2 Network components




Figure 9.3 An internetwork

that takes advantage of the services offered by a server. It also refers to a user running the client side of a client/server application. The term *server* describes either a computer or a software-based process that provides services to network users or Web services to Internet users.

Networks connect servers and clients, allowing the sharing of information and computing resources. Network equipment includes cable and wire, network adapters, hubs, switches, and various other physical connectors. In order for the network to be connected to the Internet, the network must send and retrieve data by using TCP/IP and related protocols. Networks can also be connected to form their own internets: Site-to-site connections are known as *intranets*, internal networks that are generally composed of LANs interconnected by a WAN that uses IP. Connections between partnering organizations, using IP, are known as *extranets*.

The Internet is a complex, highly redundant network of telecommunications circuits connected together with internetworking equipment, including routers, bridges, and switches. In an environment consisting of several network segments with different protocols and architectures, the network needs a device that not only knows the address of each segment but can also determine the best path for sending data and filtering broadcast traffic to the local segment. The Internet moves data by relaying traffic in packets from one computer network to another. If a particular network or computer is down or busy, the network is smart enough to reroute the traffic automatically. This requires computers (that is, routers) that are able to send packets from one network to another. Routers make decisions about how to route the data or packets, they decide which pipe is best, and then they use that best pipe. Routers work at the network layer, Layer 3, of the OSI model, which allows them to switch and route packets across multiple networks. Routers read complex network addressing information in the packet; they can share status and routing information with one another and use this information to bypass slow or malfunctioning connections.

Routing is the main process that the Internet host uses to deliver packets. The Internet uses a hop-by-hop routing model, which means that each host or router that handles a packet examines the destination address in the packet's IP header, computes the next hop that will bring the packet one step closer to its destination, and delivers that packet to the next hop, where the process is repeated. To make this happen, routing tables must match destination addresses with next hops, and routing protocols must determine the content of these tables. Thus, the Internet and the public switched telephone network (PSTN) operate quite differently from one another. The Internet uses packet switching, where there's no dedicated connection and the data is fragmented into packets. Packets can be delivered via different routes over the Internet and reassembled at the ultimate destination. Historically, "back-office" functions such as billing and network management have not been associated with Internet. But the Internet emphasizes flexibility—the capability to route packets around congested or failed points.

Recall from Chapter 5, "The PSTN," that the PSTN uses circuit switching, so a dedicated circuit is set up and taken down for each call. This allows charging based on minutes and circuits used, which, in turn, allows chain-of-supply dealings. The

major emphasis of the PSTN is on reliability. So, the Internet and the PSTN have different models and different ways of managing or routing traffic through the network, but they share the same physical foundation in terms of the transport infrastructure, or the types of communication links they use. (Chapter 4, "Establishing Communications Channels," discusses packet switching and circuit switching in detail.)

Internet Protocols

The Internet is a collection of networks that are interconnected logically as a single, large, virtual network. Messages between computers are exchanged by using packet switching. Networks can communicate with one another because they all use an internetworking protocol. Protocols are formal descriptions of messages to be exchanged and of rules to be followed in order for two or more systems to exchange information in a manner that both parties will understand. The following sections examine the Internet's protocols: TCP/IP, User Datagram Protocol (UDP), Internet Control Message Protocol (ICMP), Internet Group Management Protocol (IGMP), Address Resolution Protocol (ARP)/Reverse Address Resolution Protocol (RARP), routing protocols, and network access protocols.

TCP/IP

The IETF has technical responsibility for TCP/IP, which is the most popular and widely used of the internetworking protocols. All information to be transmitted over the Internet is divided into packets that contain a destination address and a sequence number. Packets are relayed through nodes in a computer network, along the best route currently available between the source and destination. Even though the packets may travel along different routes and may arrive out of sequence, the receiving computer is able to reassemble the original message. Packet size is kept relatively small—at 1,500 bytes or less—so that in the event of an error, retransmission is efficient. To manage the traffic routing and packet assembly/disassembly, the networks rely on intelligence from the computers and software that control delivery.

TCP/IP, referred to as the *TCP/IP suite* in Internet standards documents, gets its name from its two most important protocols, TCP and IP, which are used for interoperability among many different types of computers. A major advantage of TCP/IP is that it is a nonproprietary network protocol suite that can connect the hardware and operating systems of many different computers.

TCP Network applications present data to TCP. TCP divides the data into packets and gives each packet a *sequence number* that is not unique, but which is nonrepeating for a very long time. These packets could represent text, graphics, sound, or video—anything digital that the network can transmit. The sequence numbers

help to ensure that the packets can be reassembled correctly at the receiving end. Thus, each packet consists of content, or data, as well as the *protocol header*, the information that the protocol needs to do its work. TCP uses another piece of information to ensure that the data reaches the right application when it arrives at a system: the *port number*, which is within the range 1 to 65,535. Port numbers identify running applications on servers, applications that are waiting for incoming connections from clients. Port numbers identify one listening application from another. Ports 1 to 1,023 are reserved for server applications, although servers can use higher port numbers as well. Numbers between 1 and 1,023 are reserved for “well-known” applications (for example, Web servers run on port 80, FTP runs on port 21). Also, many recent protocols have been assigned well-known port numbers above 1,023. Ports with higher numbers, called “ephemeral” ports, are dynamically assigned to client applications as needed. A client obtains a random ephemeral port when it opens a connection to a well-known server port.

Data to be transmitted by TCP/IP has a port from which it is coming and a port to which it is going, plus an IP source and a destination address. Firewalls can use these addresses to control the flow of information. (Firewalls are discussed in Chapter 11, “Next-Generation Network Services.”)

TCP is the protocol for sequenced and reliable data transfer. It breaks the data into pieces and numbers each piece so that the receipt can be verified and the data can be put back in the proper order. TCP provides Layer 4 (transport layer) functionality, and it is responsible for virtual circuit setup, acknowledgments, flow control, and retransmission of lost or damaged data. TCP provides end-to-end, connection-oriented, reliable, virtual circuit service. It uses virtual ports to make connections; ports are used to indicate where information must be delivered in order to reach the appropriate program, and this is how firewalls and application gateways can filter and direct the packets.

IP IP handles packet forwarding and transporting of datagrams across a network. With packet forwarding, computers can send a packet on to the next appropriate network component, based on the address in the packet’s header. IP defines the basic unit of data transfer, the datagram, also referred to as the packet, and it also defines the exact format of all data as it travels across the Internet. IP works like an envelope in the postal service, directing information to its proper destination. With this arrangement, every computer on the Internet has a unique address. (Addressing is discussed later in this chapter.)

IP provides software routines to route and to store and forward data among hosts on the network. IP functions at Layer 3 (the network layer), and it provides several services, including host addressing, error notification, fragmentation and reassembly, routing, and packet timeout. TCP presents the data to IP in order to provide basic host-to-host communication. IP then attaches to the packet, in a pro-

protocol header, the address from which the data comes and the address of the system to which it is going.

Under the standards, IP allows a packet size of up to 64,000 bytes, but we don't transmit packets that large because they would cause session timeouts and big congestion problems. Therefore, IP packets are segmented into 1,500-byte-maximum chunks.

IP always does its best to make the delivery to the requested destination host, but if it fails for any reason, it just drops the packet. As such, upper-level protocols should not depend on IP to deliver the packet every time. Because IP provides connectionless, unreliable service and because packets can get lost or arrive out of sequence, or the messages may take more than 1,500 bytes, TCP provides the recovery for these problems.

UDP

Like TCP, UDP is a Layer 4 protocol that operates over IP. UDP provides end-to-end, connectionless, unreliable datagram service. It is well suited for query-response applications, for multicasting, and for use with Voice over IP (VoIP). (VoIP is discussed in Chapter 11.) Because UDP does not request retransmissions, it minimizes what would otherwise be unmanageable delay; the result is that sometimes the quality is not very good. For instance, if you encounter losses or errors associated with a voice packet, the delays that would be associated with retransmitting that packet would render the conversation unintelligible. In VoIP, when you lose packets, you do not request retransmissions. Instead, you hope that the user can recover from the losses by other means. Unlike TCP, UDP does not provide for error correction and sequenced packet delivery; it is up to the application itself to incorporate error correction if required.

ICMP

ICMP provides error handling and control functions. It is tightly integrated with IP. ICMP messages, delivered in IP packets, are used for out-of-band messages related to network operation or misoperation. Because ICMP uses IP, ICMP packet delivery is unreliable. ICMP functions include announcing network errors, announcing network congestion, assisting in troubleshooting, and announcing timeouts.

IGMP

Another Layer 3 protocol is Internet Group Management Protocol (IGMP), whose primary purpose is to allow Internet hosts to participate in multicasting. The IGMP standard describes the basics of multicasting IP traffic, including the format of multicast IP addresses, multicast Ethernet encapsulation, and the concept of a host group (that is, a set of hosts interested in traffic for a particular multicast address). IGMP enables a router to determine which host groups have members on a given

network segment, but IGMP does not address the exchange of multicast packets between routers.

ARP and RARP

At Layer 3 you also find ARP/RARP. ARP determines the physical address of a node, given that node's IP address. ARP is the mapping link between IP addresses and the underlying physical (MAC) address. RARP enables a host to discover its own IP address by broadcasting its physical address. When the broadcast occurs, another node on the LAN answers back with the IP address of the requesting node.

Routing Protocols

Routing protocols are protocols that allow routers to communicate with each other. They include Routing Information Protocol (RIP), Interior Gateway Protocol (IGP), Open Shortest Path First (OSPF), Exterior Gateway Protocol (EGP), and Border Gateway Protocol (BGP).

There are several processes involved in router operation. First, the router creates a routing table to gather information from other routers about the optimum paths. As discussed in Chapter 8, "Local Area Networking," the routing tables can be static or dynamic; dynamic routing tables are best because they adapt to changing network conditions. Next, when data is sent from a network host to a router en route to its destination, the router breaks open the data packet and looks at the destination address to determine the most efficient path between two endpoints. To identify the most efficient path, the router uses algorithms to evaluate a number of factors (called *metrics*), including distance and cost. Routing protocols consider all the various metrics involved when computing the best path.

Distance-Vector Versus Link-State Protocols Two main types of routing protocols are involved in making routing decisions:

- **Distance-vector routing protocols**—These routing protocols require that each router simply inform its neighbors of its routing table. For each network path, the receiving router picks the neighbor advertising the lowest cost, and then the router adds this into its routing table for readvertisement. Common distance-vector routing protocols are RIP, Internetwork Packet Exchange (IPX) RIP, AppleTalk Routing Table Management Protocol (RTMP), and Cisco's Interior Gateway Routing Protocol (IGRP).
- **Link-state routing protocols**—Link-state routing protocols require that each router maintain at least a partial map of the network. When a network link changes state—up to down or vice versa—a notification is flooded throughout the network. All the routers note the change and recompute the routes accordingly. This method is more reliable, easier to debug, and less

bandwidth-intensive than distance-vector routing, but it is also more complex and more computer- and memory-intensive. OSPF, Intermediate System to Intermediate System (IS-IS), and Network Link Services Protocol (NLSP) are link-state routing protocols.

Interior and Exterior Routing *Interior routing* occurs within an *autonomous system*, which is a collection of routers under a single administrative authority that uses a common interior gateway protocol for routing packets. Most of the common routing protocols, such as RIP and OSPF, are interior routing protocols. The autonomous system number is a unique number that identifies an autonomous system in the Internet. Autonomous system numbers are managed and assigned by ARIN (North America), APNIC (Asia-Pacific), and RIPE NCC (Europe). Exterior routing protocols, such as BGP, use autonomous system numbers to uniquely define an autonomous system. The basic routable element is the IP network or subnetwork, or the Classless Interdomain Routing (CIDR) prefix for newer protocols. (CIDR is discussed a little later in the chapter.)

OSPF, which is sanctioned by the IETF and supported by TCP, is intended to become the Internet's preferred interior routing protocol. OSPF is a link-state protocol with a complex set of options and features. Link-state algorithms control the routing process and enable routers to respond quickly to changes in the network. Link-state routing makes use of the Dijkstra algorithm (which determines routes based on path length and is used in OSPF) to determine routes based on the number of hops, the line speed, the amount of traffic, and the cost. Link-state algorithms are more efficient and create less network traffic than do distance-vector algorithms, which can be crucial in environments that involve multiple WAN links.

Exterior routing occurs between autonomous systems and is of concern to service providers and other large or complex networks. Whereas there may be many different interior routing schemes, a single exterior routing scheme manages the global Internet, and it is based on the exterior routing protocol BGP version 4 (BGP-4). The basic routable element is the autonomous system. Routers determine the path for a data packet by calculating the number of hops between internetwork segments. Routers build routing tables and use these tables along with routing algorithms.

Network Access Protocols

Network access protocols operate at Layer 2. They provide the underlying basis for the transport of the IP datagrams. The original network access protocol was Ethernet, but IP can be transported transparently over any underlying network, including Token Ring, FDDI, Fibre Channel, Wireless, X.25, ISDN, Frame Relay, or ATM.

Both Serial Line Internet Protocol (SLIP) and Point-to-Point Protocol (PPP) were designed specifically for IP over point-to-point connections. PPP provides data-link

layer functionality for IP over dialup/dedicated links. In other words, whenever you dial in to your ISP, you negotiate a PPP session, and part of what PPP does is to provide a mechanism to identify and authenticate the user that is dialing up.

Internet Addressing

To make the Internet an open communications system, a globally accepted method of identifying computers was needed, and IP acts as the formal addressing mechanism for all Internet messaging.

Each host on the Internet is assigned a unique 32-bit Internet address, called the *IP address*, which is placed in the IP header and which is used to route packets to their destinations. IP addresses are assigned on a per-interface basis, so a host can have several IP addresses if it has several interfaces (note that a single interface can have multiple addresses, too). Therefore, an IP address refers to an interface, not to the host. A basic concept of IP addressing is that some of the bits of the IP address can be used for generalized routing decisions because these bits indicate which network (and possibly which subnet) the interface is a member of. Addressing is performed on the basis of network/subnet and host; routing is performed based on the network/subnet portion of the address only. When a packet reaches its target network, the host portion of the address is then examined for final delivery.

The current generation of IP is called IP version 4 (IPv4). IP addresses have two parts: The first is the network ID and the second is the host ID. Under IPv4, there are five classes, which differ in how many networks and hosts are supported (see Figure 9.4):

- **Class A**—With Class A, there can be a total of 126 networks, and on each of those networks there can be 16,777,214 hosts. Class A address space is largely exhausted, although there is some address space reserved by IANA.
- **Class B**—Class B addresses provide for 16,384 networks and each of which can have 65,534 hosts. Class B space is also largely exhausted, with a few still available, albeit at a very high cost.
- **Class C**—Class C allows 2,097,152 networks, each of which can have 254 hosts.
- **Class D**—Class D belongs to a special aspect of the Internet called the multicast backbone (MBONE). *Singlecast*, or unicast, means going from one transmitter to one receiver. *Multicast* implies moving from one transmitter to multiple receivers. Say, for instance, that you are in San Francisco and you want to do a videoconferencing session that involves three offices located in London. In the unicast mode, you need three separate IP connections to London from the conferencing point in San Francisco. With multi-




Figure 9.4 IPv4 32-bit addressing

cast, however, you need only one IP connection. A multicast router (mrouter) would enfold your IP packets in special multicast packets and forward those packets on to an mrouter in London; in London that mrouter would remove the IP packets, replicate those packets, and then distribute them to the three locations in London. The MBONE system therefore conserves bandwidth over a distance, relieves congestion on transit links, and makes it possible to address a large population in a single multicast.

- **Class E**—Class E is reserved address space for experimental purposes.

The digits in an IP address tell you a number of things about the address. For example, in the IP address 124.29.88.7, the first set of digits, 124, is the network ID, and because it falls in the range of numbers for Class A, we know that this is a Class A address. The remaining three sets, 29.88.7, are the host ID. In the address 130.29.88.7, the first two sets, 130.29, comprise the network ID and indicate that this is a Class B address; the second two sets in this address, 88.7, comprise the host ID. Figure 9.5 shows an example of IP addressing.

Network IDs are managed and assigned by ARIN, APNIC, and RIPE NCC. Host IDs are assigned locally by the network administrator. Given a 32-bit address field, we can achieve approximately 4.3 billion different addresses with IPv4. That seems like a lot, but as we began to experience growth in the Internet, we began to worry about the number of addresses left. In the early 1990s, the IETF began to consider the potential of IP address space exhaustion. The result




Figure 9.5 An example of IP network addressing

was the implementation of CIDR, which eliminated the old class-based style of addressing. The CIDR address is a 32-bit IP address, but it is classless. The CIDR addressing scheme is hierarchical. Large national and regional service providers are allocated large blocks of contiguous Internet addresses, which they then allocate to other smaller ISPs or directly to organizations. Networks can be broken down into subnetworks, and networks can be combined into supernet, as long as they share a common network prefix. Basically, with CIDR a route is no longer an IP address broken down into network and host bits according to its class; instead, the route becomes a combination of an address and a mask. The mask indicates how many bits in the address represent the network prefix. For example, the address 200.200.14.20/23 means that the first 23 bits of the binary form of this address represent the networks. The bits remaining represent the host. In binary form, the prefix 23 would look like this: 255.255.254.0. Table 9.1 lists the most commonly used masks represented by the prefix, and the number of host addresses available with a prefix of the type listed. CIDR defines address assignment and aggregation strategies designed to minimize the size of top-level Internet routing tables. The national or regional ISP needs only to advertise its single supernet address, which represent an aggregation of all the subnets within that supernet. Routers in the Internet no longer give any credence to class—it's entirely based on the CIDR prefix. CIDR does require the use of supporting routing protocols, such as RIP version 2, OSPF version 2, Enhanced Interior Gateway Routing Protocol (EIGRP), and BGP-4.

Table 9.1 CIDR Masking Scheme

Mask as Dotted-Decimal Value	Mask as Prefix Value	Number of Hosts
255.255.255.224	/27	32
255.255.255.192	/26	64
255.255.255.128	/25	128
255.255.255.0 (Class C)	/24	256
255.255.254.0	/23	512
255.255.252.0	/22	1,024
255.255.248.0	/21	2,048
255.255.242.0	/20	4,096
255.255.240.0	/19	8,192
255.255.224.0	/18	16,384
255.255.192.0	/17	32,768
255.255.0.0 (Class B)	/16	65,536
255.254.0.0	/15	131,072
255.252.0.0	/14	262,144
255.248.0.0	/13	524,288

Subnetting is a term you may have heard in relationship to addressing. It once referred to the subdivision of a class-based network into subnetworks. Today, it generally refers to the subdivision of a CIDR block into smaller CIDR blocks. Subnetting allows single routing entries to refer either to the larger block or to its individual constituents, and this permits a single general routing entry to be used through most of the Internet, with more specific routes being required only for routers in the subnetted block.

Researchers are predicting that even with CIDR and subnetting in place, we will run out of IPv4 address space by the year 2010. Therefore, several years ago the IETF began developing an expanded version of IP called IPv6 (originally called IPng—for IP Next Generation). IPv6 uses a 128-bit address, which allows a total of 340 billion billion billion billion unique addresses, which equates to approximately

IPv6 Address Allocation and Assignment

ARIN has published a draft policy document on IPv6 address allocation and assignment that is available at www.arin.net.

70 IP addresses for every square inch of the earth's surface, including oceans. That should hold us for a while and be enough for each and every intelligent appliance, man, woman, child, tire, pet, and curb to have its own IP address. Along with offering a greatly expanded address space, IPv6 also allows increased scalability through multicasting and includes increased Quality of Service (QoS) capabilities. (QoS is discussed in detail in Chapter 10, "Next-Generation Networks.")

The IPv6 specification includes a flow label to support real-time traffic and automated connectivity for plug-and-play use. In addition, IPv6 provides improved security mechanisms. It incorporates Encapsulated Security payload (ESP) for encryption, and it includes an authentication header, to make transactions more secure. Although IPv6 offers many benefits, it requires a major reconfiguration of all the routers out there, and hence we haven't seen the community jump at the migration from IPv4 to IPv6. But in order to meet the demands of the growing population of not just human users but smart machines that are tapping into the Internet, the transition will be necessary. An experimental network called the 6Bone network is being used as an environment for IPv6 research. So far more than 400 networks in more than 40 countries are connected to the 6Bone IPv6 network.

The Domain Name System

The *Domain Name System* (DNS) is a distributed database system that operates on the basis of a hierarchy of names. DNS provides translation between easy-for-humans-to-remember host names (such as www.telecomwebcentral.com or www.lidoorg.com) and the physical IP addresses, which are harder for humans to remember. It identifies a domain's mail servers and a domain's name servers. When you need to contact a particular URL, the host name portion of the URL must be *resolved* to the appropriate IP address. Your Web browser goes to a local name server, maintained either by your ISP, your online service provider, or your company. If the IP address is a local one—that is, it's on the same network as the one you are on—then the name server will be able to resolve that URL with the IP address right away. In this case, the name server sends the true IP address to your computer, and because your Web browser now has the real address of the place you're trying to locate, it contacts that site, and the site sends the information you've requested.

If the local name server determines that the information you have requested is not on the local network, it must get the information from a name server on the Internet. The local name server contacts the root domain server, which contains a list of the top-level domain name servers managed by ICANN. The root domain server tells the local server which top-level domain name server contains the domain specified in the URL. The top-level domain name server then tells the local server which primary name server and secondary name server have the information about the requested URL. The local name server can then contact the primary name server. If the information can't be found in the primary name server, then the local name server contacts the secondary name server. One of those name servers will have the proper information, and it will then pass that information back to the local name server. The local name server sends the information back to your browser, which then uses the IP address to contact the proper site.

Top-Level Domains

For some time, there have been seven generic top-level domains:

.com	commercial
.gov	government
.mil	military
.edu	education
.net	for network operation
.org	nonprofit organizations
.int	international treaty organizations

Seven new top-level domains have been operational since 2001:

.aero	air-transport industry
.biz	businesses
.coop	cooperatives
.info	any use
.museum	museums
.name	individuals
.pro	Accountants, lawyers, and physicians

There are also country code top-level domains. Each of these is a two-letter country code (for example, .au, .ca), and there are 245 country code top-level domains, including a .us domain for the United States. So if you wanted to protect

your domain name in the .com domain, for example, you would actually have to register it in 246 domains—.com and then .com with the appropriate two-letter country code after that—and if you really want to get serious about branding, you'd probably want to register another 246 each in the .net, .org, and .biz domains! Of course, very few organizations actually do this.

The Importance of Domain Names

Many new domain names are registered every minute, and it seems that all the simple one- and two-word .com names have already been taken. Therefore, there's a call for new domains to be added. Originally IANA, which was funded by the U.S. government, administrated the DNS. Since 1993, Network Solutions had been the sole provider of direct domain name registration services in the open generic top-level domains, and registration authority over the country code top-level domains has been relegated to the individual countries and bodies within them.

In September 1998 ICANN was formed to take over. ICANN is now introducing competition into the administration of the DNS through two attempts. One is a policy for the accreditation of registrars, and the other is a shared registry system for the .com, .net, and .org domains. In 2001 ICANN operationalized seven new top-level domains, and it must still negotiate with the winning applicants the terms under which they will operate this registry. The future of ICANN is still a bit tenuous; it is a bit political, to say the least.

A story illustrates what value domain names have. There's a small Pacific Islands country, with a population of 10,600 people, known as Tuvalu, and it was assigned the country code .tv. Naturally, .tv is a very appealing domain. It has reference to entertainment, streaming media, and screaming multimedia, and it also has a global context: Once you register something as .tv, you would no longer be able to alter it by appending another country code because it already is a country code. Of course, many entrepreneurs developed an interest in Tuvalu, and many companies approached the country, trying to acquire its domain name; Tuvalu auctioned the name. A company called .tv bought the name for roughly US\$1 million quarterly—adjustable for inflation—with a US\$50 million cap over 10 years. In addition, Tuvalu holds a 20% stake in the company. This auctioning of the country's domain name produced four times the country's GDP. Needless to say, the island is richly developing its transportation, educational, and health care facilities.

On the .tv domain, some domain names are quite expensive, with bidding starting at US\$250,000 for broadband.tv, for instance. On the other hand, some creative and descriptive domains haven't yet been registered, and you'd be able to acquire those for as little as US\$50. A lot of money is tied up in domain names, and the process of creating new domains will further challenge identifying the best branding strategy.

■ The Organization of the Internet

It's important to understand what the Internet infrastructure is composed of and how it's structured in terms of the large variety of players that are represented in the Internet space. It's also important to keep in mind that similarly to the PSTN, the Internet was not originally structured for what we're asking it to do now.

Initially, the Internet was designed to support data communications—bursty, low-speed text data traffic. It was structured to accommodate longer hold times while still facilitating low data volumes, in a cost-effective manner. (That was the introduction of the packet-switching technique, whereby through statistical multiplexing long hold times don't negatively affect the cost structure because you're sharing the channel with other users as well.) The capacities of the links initially dedicated to the Internet were very narrowband: 56Kbps or 64Kbps. The worldwide infrastructure depended on the use of packet switches (that is, routers), servers (that is, repositories for the information), and clients (that is, the user interfaces into the repositories). The Internet was composed of a variety of networks, including both LANs and WANs, with internetworking equipment such as routers and switches designed for interconnection of disparate networks. The Internet relied on TCP/IP to move messages between different subnetworks, and it was not traditionally associated with strong and well-developed operational support systems, unlike the PSTN, where billing systems, provisioning systems, and network management systems are quite extensive, even if they are not integrated.

The traditional Internet relied on the PSTN for subscriber access to the Internet. So the physical framework, the roadways over which a package travels on what we know as the Internet, is the same type of physical infrastructure as the PSTN—it uses the same types of communications, links, and capacities. And in order for users to actually access this public data network, they had to rely on the PSTN. So, two types of access were facilitated: dialup for consumers and small businesses (that is, the range of analog modems, Narrowband ISDN) and dedicated access in the form of leased lines, ISDN Primary Rate Interface (PRI), and dedicated lines based on T-1/E-1 capacities for larger enterprises, and, in some cases, even T-3/E-3.

The Evolution of the POP Architecture

The early Internet point of presence (POP) architecture was quite simple, as illustrated in Figure 9.6. You would have either 56Kbps or 64Kbps lines coming in to access ports on a router. Out of that router, T-1/E-1 trunks would lead to a UNIX host. This UNIX environment was, for most typical users, very difficult to navigate. Until there was an easier way for users to interface—the World Wide Web—the




Figure 9.6 POP architecture in the 1980s

Internet was very much the province of academicians, engineers, and computer scientists.

The architecture of the Internet today is significantly different from what it was in the early days. Figure 9.7 shows some of the key components you would find in a higher-level network service provider's (NSP's) or a high-tier ISP's POP today. (Of course, a local service provider with just one POP or one node for access purposes, perhaps to a small community, looks quite different from this.)

First, let's look at the support for the dialup users. Today, we have to facilitate a wide range of speeds; despite our admiration of and desire for broadband access, it's not yet widely available. In the next several years, we should see more activity in terms of local loop modernization to provide broadband access to more users. But for the time being, we have to accommodate a wide range of analog modems that operate at speeds between 14.4Kbps and 56Kbps. Therefore, the first point of entry at the POP requires an analog modem pool of modems that complement the ones that individuals are using. Also, as we add broadband access alternatives, additional access devices are required, for instance, for DSL modems or cable modems. The analog modem pool communicates with a terminal server, and the terminal server establishes a PPP session. PPP does two things: It assigns an IP address to a dialup user's session, and it authenticates that user and authorizes entry. By dynamically allocating an IP address when needed, PPP enables us to reuse IP addresses, helping to mitigate the problem of the growing demand for IP addresses. A user is allocated an address when she dials in for a session; when she




Figure 9.7 POP architecture today

terminates the session, the IP address can be assigned to another user. PPP supports both Password Authentication Protocol (PAP) and Challenge Handshake Authentication Protocol (CHAP) to provide link-level security. PAP uses a two-way handshake for the peer to establish its identity upon link establishment. The peer repeatedly sends the password to the authenticator until verification is acknowledged or the connection is terminated. CHAP uses a three-way handshake to periodically verify the identity of the peer throughout the life of the connection. The server sends to the remote workstation a random token that is encrypted with the user's password and sent back to the server. The server performs a lookup to see if it recognizes the password. If the values match, the authentication is acknowledged; if not, the connection is terminated. A different token is provided each time a remote user dials in, which provides additional robustness.

The terminal server resides on a LAN, which would typically be a Gigabit Ethernet network today. Besides the terminal server, the ISP POP houses a wide range of other servers:

- **E-mail servers**—These servers house the e-mail boxes.
- **Domain name servers**—These servers resolve the uniform resource locaters (URLs) into IP addresses.
- **Web servers**—If the ISP is engaged in a hosting business, it needs a Web server.

- **Security servers**—Security servers engage in encryption, as well as in authentication and certification of users. Not every ISP has a security server. For example, those that want to offer e-commerce services or the ability to set up storefronts must have them. (Security is discussed in detail in Chapter 11.)
- **Newsgroup servers**—Newsgroup servers store the millions of messages that are posted daily, and they are updated frequently throughout the day.
- **Proxy servers**—A proxy server provides firewall functionality, acting as an intermediary for user requests, establishing a connection to the requested resource either at the application layer or at the session or transport layer. Proxy servers provide a means to keep outsiders from directly connecting to a service on an internal network. Proxy servers are also becoming critical in support of edge caching of content. People are constantly becoming less tolerant of lengthy downloads, and information streams (such as video, audio, and multimedia) are becoming more demanding of timely delivery. You want to minimize the number of hops that a user has to go through. For example, you could use a tracing product to see how many hops you've gone through to get to a Web site. You'll see that sometimes you need to go through 17 or 18 hops to get to a site. Because the delay at each hop can be more than 2,000 milliseconds, if you have to make 18 hops when you're trying to use a streaming media tutorial, you will not be satisfied. ISPs can also use proxy servers to cache content locally, which means the information is distributed over one hop rather than over multiple hops, and that greatly improves your streaming media experience. Not all proxy servers support caching, however.

The ISP POP also contains network management systems that the service providers can use to administer passwords and to monitor and control all the network elements in the POP, as well as to remotely diagnose elements outside the POP.

An *access router* filters local traffic. If a user is simply checking his e-mail, working on his Web site, or looking up newsgroups, there's no reason for the user to be sent out over the Internet and then brought back to this particular POP. An access router keeps traffic contained locally in such situations. A distribution router, on the other hand, determines the optimum path to get to the next hop that will bring you one step closer to the destination URL, if it is outside the POP from which you are being served. Typically, in a higher-level ISP, this distribution router will connect into an ATM switch, which enables the ISP to guarantee QoS; this is especially necessary for supporting larger customers on high-speed interfaces and links and for supporting virtual private networks, VoIP, or streaming media applications. The ATM switch, by virtue of its QoS characteristics, enables us to map the packets into

the appropriate cells, which guarantee that the proper QoS is administered and delivered. (ATM QoS is discussed further in Chapter 10.) The ATM switch then is front-ended by a data service unit (DSU), the data communications equipment on which the circuit terminates, which performs signal conversion and provides diagnostic capabilities. The network also includes a physical circuit, which, in a larger higher-tier provider, would generally be in the optical carrier levels.

An *access concentrator* can be used to create the appearance of a virtual POP. For instance, if you want your subscribers to believe that they're accessing a local node—that is, to make it appear that you're in the same neighborhood that they are in—you can use an access concentrator. The user dials a local number, thinking that you're located down the street in the business park, when in fact, the user's traffic is being hauled over a dedicated high-speed link to a physical POP located elsewhere in the network. Users' lines terminate on a simple access concentrator, where their traffic is multiplexed over the T-1s or T-3s, E-1s or E-3s, or perhaps ISDN PRI. This gives ISPs the appearance of having a local presence when, in fact, they have none. I talk later in this chapter about the advantages of owning the infrastructure versus renting the infrastructure, but clearly, if you own your infrastructure, backhauling traffic allows you to more cost-effectively serve remote locations. If you're an ISP that's leasing facilities from a telco, then these sorts of links to backhaul traffic from more remote locations will add cost to your overall operations.

You can see that the architecture of the POP has evolved and become incredibly more sophisticated today than it was in the beginning; the architecture has evolved in response to and in preparation for a very wide range of applications.

Internet Challenges and Changes

Despite all its advances over the past couple of decades, the Internet is challenged today. It is still limited in bandwidth at various points. The Internet is composed of some 10,000 service providers. Although some of the really big companies have backbone capacities that are 50Gbps or greater, there are still plenty of small backbones worldwide that have only a maximum of 1.5 or 2Mbps. Overall, the Internet still needs more bandwidth.

One reason the Internet needs more bandwidth is that traffic is increasing at an alarming rate. People are drawn to Web sites that provide pictures of products in order to engage in demonstrations and in order to be able to conduct multimedia communications. Those greater capacities required by these visual objects also demand greater bandwidth. This means that we frequently have bottlenecks at the ISP level, at the backbone level (that is, the NSP level), and at the network access points (NAPs) where backbones interconnect to exchange traffic between providers. These bottlenecks greatly affect our ability to roll out new time-sensitive,

loss-sensitive applications, such as Internet telephony, VoIP, VPNs, streaming media, and TV over Internet.

Therefore, we are redefining the Internet as we are redefining the PSTN. In both cases, we're trying to support more real-time traffic flows, real audio, real video, and live media. This requires the introduction of QoS into the Internet. There are really two types of metrics that we loosely refer to as QoS: class of service (CoS) and true QoS. CoS is a prioritization scheme in which you can prioritize streams and thereby facilitate better performance. QoS, however, deals with very strict traffic measurements, where you can specify the latencies end to end (that is, the jitter or variable latencies in the receipt of the packets, the tolerable cell loss, and the mechanism for allocating the bandwidth continuously or on a bursty basis). Thus, QoS is much more stringent than CoS, and what we are currently introducing into the Internet is really more like CoS than QoS.

Techniques such as DiffServ (as discussed in Chapter 10) allow us to prioritize the traffic streams, but they really do not allow us to control the traffic measurements. That is why, as discussed in Chapter 7, "Wide Area Networking," we tend to still rely on ATM within the core: ATM allows the strict control of the traffic measurements, and it therefore enables you to improve performance, quality, reliability, and security. Efforts are under way to develop QoS standards for IP, but we're still a couple years away from clearly defining the best mechanism. In the meantime, we are redesigning the Internet core, moving away from what was a connectionless router environment that offered great flexibility and the ability to work around congestion and failures, but at the expense of delays. We're moving to a connection-oriented environment in which we can predefine the path and more tightly control the latencies, by using techniques such as Frame Relay, ATM, and MPLS, each of which allow you to separate traffic types, prioritize the time-sensitive traffic, and, ultimately, to reduce access costs by eliminating leased-lines connections.

The other main effort in redesigning the Internet core is directed at increasing its capacity, moving from OC-3 and OC-12 (that is, 155Mbps and 622Mbps) at the backbone level to OC-48 (that is, 2.5Gbps) and even OC-192 (that is, 10Gbps). But remember that the bits per second that we can carry per wavelength doubles every year, and the number of wavelengths we can carry per fiber also doubles every year. So, the migration beyond 10Gbps is also under way in the highest class of backbones, and it will continue at a rapid pace.

The emergent generation of Internet infrastructure is quite different from the traditional foundation. First, it's geared for a new set of traffic and application types: high-speed, real-time, and multimedia. It must be able to support and guarantee CoS and QoS. It includes next-generation telephony, which is a new approach to providing basic telephony services, but it uses IP networks. (These types of next-generation network services are discussed in Chapter 11.)

The core of the Internet infrastructure, like the PSTN, will increasingly rely on SDH/SONET, DWDM, and optical networking. It will require the use of ATM, MPLS, and MPλS (Multiprotocol Lambda Switching) networking protocols to ensure proper administration of performance. New generations of IP protocols are being developed to address real-time traffic, CoS, QoS, and security. Distributed network intelligence is being used to share the network functionality.

We are working on providing the capability to rely on multiple broadband access options, not just the PSTN. You may be able to access the Internet on a dial-up basis through the new generation of xDSL facilities, through a cable TV company and a cable modem, through a satellite TV company, via direct broadcast satellites, or through point-to-point microwave solutions such as MMDS and LMDS. (These solutions are discussed in Chapter 13, "Broadband Access Solutions.") For the dedicated environment, largely we're seeing a migration to higher bandwidth (that is, T-1 moving to T-3, E-1 moving to E-3, early adopters and high-bandwidth consumers in the optical carrier levels), and we're seeing increased reliance on Frame Relay and ATM as the access technique.

■ Service Providers

There is a wide range of service providers in the Internet space. One way they differ is in their coverage areas. Some providers focus on serving a local area, others are regionally based, and others offer national or global coverage. Service providers also vary in the access options that they provide. All ISPs offer plain old telephone service (POTS), and some offer ISDN, xDSL, Frame Relay, ATM, cable modem service, satellite, and wireless as well. Providers also differ in the services that they support. Almost all providers support e-mail (but not necessarily at the higher-tier backbone level). Some also offer FTP hosting, Web hosting, name services, VPNs, VoIP, application hosting, e-commerce, and streaming media. Providers could service a very wide variety of applications, and as a result, there is differentiation on this front as well. Two other important issues are customer service and the number of hops a provider must take in order to get to the main point of interconnection into the Internet.

It is pretty easy to become an ISP: pick up a router, lease a 56Kbps/64Kbps line, and you're in business. This is why there are some 10,000 such providers, of varying sizes and qualities, worldwide. There is a service provider pecking order. Research backbones have the latest technology. Top-tier providers focus on business-class services; lower-tier providers focus on rock-bottom pricing. Consequently, there are large variations in terms of available capacity, the performance you can expect, the topology of the network, the levels of redundancy, the numbers of connections with other operators, and the level of customer service and the

extent of its availability (that is, whether it's 24/7 or whether it's a Monday-through-Friday, 9-to-5 type of operation). Ultimately, of course, ISPs vary greatly in terms of price.

Figure 9.8 shows an idealized model of the service provider hierarchy. At the top of the heap are research backbones. For example, Internet 2 replaces what the original Internet was for—the academic network. Some 85% of traffic within the academic domain stays within the academic domain, so there's good reason to have a separate backbone for the universities and educational institutions involved in research and learning. Internet 2 will, over time, contribute to the next commercialized platform. It acts as a testbed for many of the latest and greatest technologies, so the universities stress test Internet 2 to determine how applications perform and which technologies suit which applications or management purposes best. Other very sophisticated technology platforms exist, such as the Abilene Project and the Interplanetary Internet (IPNSIG), the first Internet being constructed in space for purposes of deep space communications. The objective of IPNSIG is to define the architecture and protocols necessary to permit interoperation of the Internet resident on earth with other remotely located internets resident on other planets or spacecraft in transit.

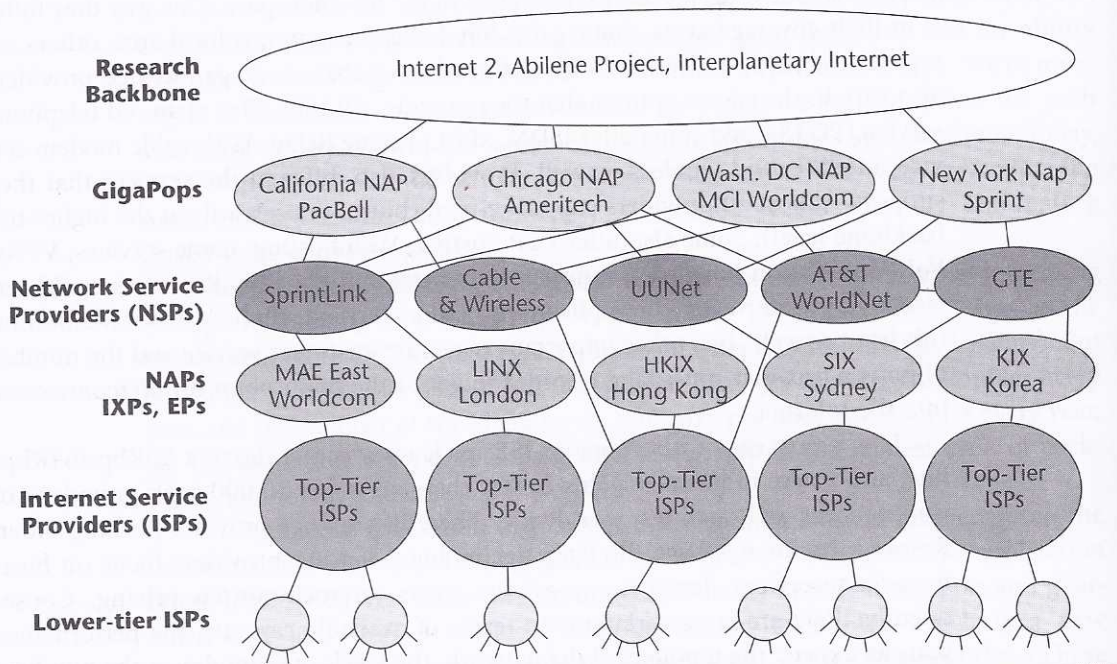


Figure 9.8 Service provider hierarchy

In the commercial realm, the highest tier is the NSP. NSPs are very large backbones, global carriers that own their own infrastructures. The top providers at this level are AT&T, Worldcom, UUnet, Sprint, Verizon, Cable & Wireless, and Qwest. The NSPs can be broken down into three major subsets:

- **National ISPs**—These ISPs have national coverage. They include the incumbent telecom carriers and the new competitive entrants.
- **Regional ISPs**—These ISPs are active regionally throughout a nation, and these service providers own their equipment and lease their lines from the incumbent telco or competitive operator.
- **Retail ISPs**—These ISPs have no investment in the network infrastructure whatsoever. They're basically using their brand name and outsourcing all the infrastructure to perhaps an NSP or a high-tier ISP, but they're building from a known customer database that's loyal and that provides an opportunity to offer a branded ISP service.

These various levels of NSPs interconnect with one another in several ways. First, they can connect at the National Science Foundation-supported NAPs. These NAPs are used to provide connection into the Internet 2 project, largely in the United States. This occurs between the largest of the NSPs as well as the research backbones. Second, there are commercial interconnection and exchange points, which people refer to as NAPs as well, although they are also called metropolitan area exchanges (MAEs) and interconnection and exchange points. These NAPs are typically run by consortiums of ISPs, telcos, entrepreneurs, and others seeking to be in the business of the Internet; these consortiums build public exchange points for traffic between the various Internet backbones. Third, service providers can connect using bilateral arrangements (called *private peering*) between one another to pass traffic over each others' backbones. These NAPs are discussed in more detail in the following section, but for now, suffice it to say that local service providers, the lower-tier ISPs, typically connect to NAPs through the top-tier ISPs, so they are a greater number of hops away from the point of interconnection, which can have a big impact on the performance of time-sensitive applications.

Evaluating Service Providers

The following sections describe some of the important characteristics you should expect in the different levels of service providers.

Evaluating NSPs

NSPs provide at least a national backbone, but it's generally international. Performance depends on the total network capacity, the amount of bandwidth, and the

total number of customers contending for that bandwidth. If you divide the total capacity by the number of direct access customers, you get an idea of how much bandwidth is available on average per customer; however, that doesn't take into account the additional traffic that consumer or dialup traffic may add. If you were to perform such an exercise based on your ISP, you'd probably be shocked at how little bandwidth is actually available to each individual user.

The NSPs should operate at the higher speeds, ranging today from OC-3 to OC-48, and looking to move on to OC-192 and beyond. They should have engineered the network such that during peak periods the network has at least 30% to 40% spare bandwidth. They should have nodes in all the major cities. NSPs that own their facilities—that is, facilities-based ISPs—have an added advantage of being able to provide backhauled circuits at no extra charge, whereas those that are leasing facilities from interexchange carriers or telcos incur additional charges to the overall operation. There should be redundancy applied to the circuits, routers, and switches. Most NSPs have at least two redundant paths from each node. Some may have anywhere from three to five redundant paths. An NSP should practice diversity in the local loops and long-haul circuits—that is, it should get them from different carriers so that in the event of a disaster, there is a fallback plan. Facilities-based carriers often offer redundancy, but rarely do they provide diversity. NSPs should have generators at each node. Remember that it doesn't matter how much redundancy you build into your equipment or your circuits if you lose power. You need redundant power, and this is a major differentiator for the very high-tier operators.

NSPs should have implemented BGP in order to filter, or block, any faulty messages from the backbone that can replicate themselves and cause major brown-outs or blackouts on the Internet. NSPs should connect to one another through multiple NAPs and interexchange points and also have multiple private peering agreements, again to cover all their bases. They should provide redundant paths to the NAPs as well as to their peers, and they should be able to articulate a very clear reason for their architectural decisions (Why are they using IP routing and connectionless networks? Why does their core consist of ATM? Are they using ATM for its QoS benefits?). They might want to speak to issues of speed, overhead, or QoS, but you want to work with an NSP that actually has a clear-cut architectural reason for its decision.

When comparing backbones, look at the backbone speeds. Look at the underlying transport technology and what that means in terms of QoS. Look at the number of nodes, the number of redundant paths per node, the availability of power backup at each node, the availability of BGP filtering, the number of NAPs or ISPs that they interconnect through, the number of private peering arrangements, and a disclosure of who those peers are. Look at whether performance guarantees are offered, whether you have the potential for any online monitoring, and the access price.

Evaluating ISPs

The top-tier ISPs have the greatest coverage in a region of a particular country. They are associated with multiple high-speed links, generally in the T-1/T-3, E-1/E-3 range, up to perhaps the OC-3 level. They have redundant routers and switches at each node, and they have multiple high-speed connections into their NAPs (and more and more NAPs are becoming discriminating, as discussed later in this chapter, in the section "NAPs"). They require two connections, for example, into the Internet, so you have redundancy, and the larger NAPs may also require that for purposes of transit links, the ISP have arrangements with alternative international providers. Their main focus is to provide high levels of service, to address the business community.

The lower-tier ISPs are generally associated with providing local access—to a mountain village that has a ski community in the winter or to a remote lake where there is summer traffic, or to any other neighborhood. There is generally only one rather low-speed connection, either 56Kbps or 64Kbps, or perhaps up to T-1/E-1, and it leads into a top-tier ISP; these lower-tier ISPs generally do not have direct connections into the NAPs or high-tier ISPs. Lower-tier ISPs focus on offering the lowest prices possible, so they offer the least amount of redundancy of any providers—no power backups and fairly minimal capacities.

As you can see in the idealized model of the Internet shown in Figure 9.9, information flows from a user through a local ISP, through a regional ISP, to a national ISP, through transit ISPs, and back down the hierarchy. Some companies operate on a vertically integrated basis, so they may be represented as local ISPs to their consumers but also offer national backbones to their business customers. Therefore, the relationships can be a little less defined than the figure indicates, but it helps provide a summary view of how this works.

An issue in ISP selection is the level of coverage (How many countries are served? How many cities are served within that country? What's the total number of backbone connections into the countries that it serves?). Another consideration is the number of exchange points you have to go through and what type and number of peering relationships are in practice. You also need to consider the total amount of bandwidth and therefore what the level of oversubscription is if everyone tries to use it at the same time. And you need to consider the transit delays being experienced; ideally we want to see networks evolve to latencies of less than 80 milliseconds, but today 800 to 1,000 milliseconds is much more realistic. Similarly, you want less than 5% packet loss, but today at peak hours you'll see up to 30% or 40% packet loss. Data packets most often are retransmitted to correct for these losses. However, with real-time traffic, such as voice or video, packet retransmission would add too much delay, with the result that conversations can be rendered unintelligible. Again, you need to think about redundancy—the number of lines into the network, the level of network diversity, and the amount of power backup involved.

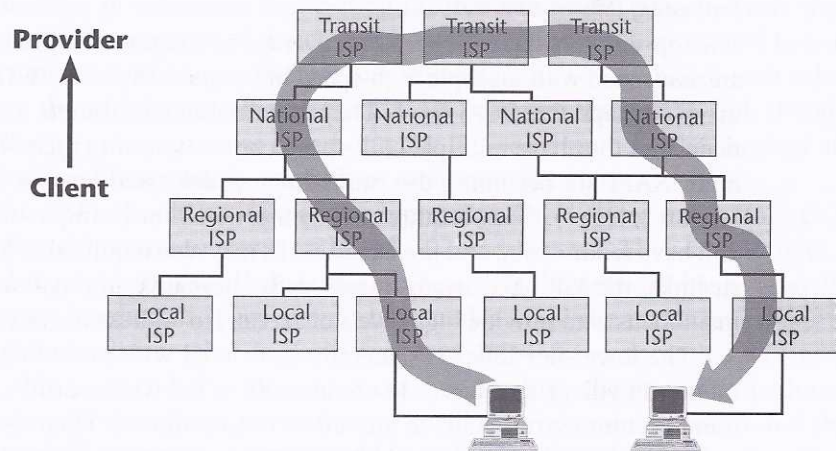


Figure 9.9 Information flow in an idealized model of the Internet

Evaluating Emerging Service Providers

In the past couple years, there have been some exciting developments in niched applications for service providers. This section talks about content delivery networks, application service providers (ASPs), management service providers (MSPs), online service providers (OSPs), and virtual ISPs (VISPs). Each of these serves a different purpose. Some of them have an immediate future, some of them perhaps may last a little bit longer, and some of them are quite unknown.

Content Delivery Networks Content delivery networks can be structured to support exactly the kind of application you need them to support. For example, you may use Web-delivered training, which is designed to have the types of requirements that streaming media applications have. Content delivery services are becoming essential to the development of e-commerce, making them a requirement for business-class Web sites. They are delivery services that are focused on streaming audio, video, and media, as well as the supporting e-commerce applications. Currently, the major clients for content delivery networks are ISPs and content providers because they stand to reduce their need for bandwidth and get better profit margins on services to customers.

Content delivery and hosting service providers aim to deliver Web-site content to customers faster by storing the content in servers located at the edges of Internet networks, rather than in the network's central location. This reduces the number of hops, thereby reducing the latencies and improving performance. This works because of a combination of technologies, including caching and *load balancing*,

which means spreading the content across multiple servers so that at peak periods, you don't overload the server but can still provide access on a balanced basis. Content delivery providers will also use enhanced Internet routing and probably proprietary algorithms that facilitate the administration of QoS. For instance, Enron has a proprietary intelligent call agent that knows how to prioritize streaming content over bursty data applications that may be supported. As another example, IBasis supports VoIP, and when the Internet gets congested, the routing algorithm switches the traffic over to the circuit-switched environment, thereby guaranteeing the high quality that the customers are expecting for the telephony. Also, these content delivery services will be able to deliver content to a user from the nearest servers, or at least from a server located at the edge of the network.

All the changes in the content delivery networks, of course, are driven by the need for faster and faster online transactions. Humans get addicted to speed. We have a physiological speed center in our brain, and each time we complete tasks at a certain pace, we resynchronize that speed center. Consequently, as we've used the Web over the years, improvements in network infrastructure have allowed us to experience faster downloads. Customer loyalty can be increasingly affected by time frames as small as a second.

There are a number of content-delivery providers in the market. Akamai Technologies has more than 4,000 servers and is growing. Digital Island plans to install 8,000 servers by 2002, increasing its current capacity by a factor of 30. Enron recently signed a deal with Blockbuster, and this is an indicator of the importance of the content as part of the overall picture.

ASPs ASPs have been receiving a great deal of press. There is not really one major business model for ASPs; there's quite a niching opportunity. An ASP is a supplier that makes applications available on a subscription basis. ASPs provide hosting, monitoring, and management of a variety of well-known software applications on a world-class data center infrastructure. A great deal of money is being spent in the ASP arena, and ASPs are increasingly becoming application infrastructure providers (AIPs), which is discussed later in this section.

ASPs are most useful when the majority of an application's users reside outside the enterprise network. The more external touch points there are, the more sense it makes to use an ASP. An ASP is viable for e-commerce, customer relations management, human resources, and even e-mail and listserv applications. An ASP is not really viable for productivity tools, such as word processing or spreadsheets. And an ASP is not really good for financial applications, which benefit from being maintained in-house, because of the more stringent requirements for security and confidentiality of data.

You might want to use an ASP if there's a need for additional bandwidth; if you lack technical resources in-house for reliable 24/7 application support; when you feel

that a third party could do the job better; when you need a large and readily available applications inventory; when scalability demands dynamic increases; or when you're seeking performance reliability.

With ASPs, you pay setup fees; on a low-end Web server these fees start at around US\$2,000 and on a low-end database server they start at around US\$10,000. Setup fees on a high-end Oracle cluster, for instance, could run from US\$5,000 for the Web servers to US\$40,000 for database servers. These customers are generally ISPs or major international corporations. Ongoing ASP fees also range anywhere from US\$2,000 to US\$40,000 per month for the software licensing, the applications, and the equipment maintenance, along with perhaps the broadband connectivity. So an ASP paying US\$1 million to buy the license for an application may charge its customers between US\$200,000 and US\$500,000 per year for three years (the typical life span of most of the applications contracts). ASPs that concentrate on small- to medium-size businesses typically offer application hosting in the range of US\$200 to US\$500 per month, sometimes with an additional US\$5 to US\$30 per month charge per user for enterprise applications. Thus, ASPs need to invest a lot of money into automating customer setup and maintenance, which then helps them reduce cost to the customer and ensure an attractive margin. Enterprises may not have the resources to realize these types of cost benefits.

Although small and medium-sized businesses may today appear to be the best targets for using ASPs, predictions show that very large enterprise clients will be the most fruitful by 2004. It is likely that large enterprises will use ASPs for e-commerce applications and for internal applications, such as e-mail, data management, office automation, and basic business applications. The emerging "skills drought" may drive larger companies to ASPs as well.

The ASP model comprises no single entity or service. Instead, it's a complex supply chain that includes the following:

- **Independent software vendors (ISVs)**—ISVs develop the applications that the ASPs then put up for sale or for rent.
- **AIPs**—AIPs manage the data center servers, databases, switches, and other gears on which the applications run. It's reminiscent of the gold miner analogy: It wasn't the gold miners who got the gold; it was the guys who sold the picks and shovels. The AINs are the segment of the ASP market that will be seeing the greatest run of success at first.
- **MSPs**—MSPs take over the actual management and monitoring of the network. (They are discussed in more detail later in this section.)
- **NSPs**—NSPs are network access providers.
- **Value-added resellers (VARs)**—VARs deal with distribution and sales.

- **Systems integrators (SIs)**—Like VARs, SIs deal with distribution and sales.
- **E-business infrastructure providers (eBIPs)**—This group saves small businesses time and money with Web-based solutions for human resources, accounting, marketing, group collaboration, and other services. Some of these eBIPs offer their services for free, making money from ads and partnerships with the VARs. Others charge affordable fees that range from \$30 to \$200 per month. Online business center sites that offer targeted small-business content and community are good partnership candidates for such eBIPs.

ASPs form a complex supply chain. Today, more than 650 companies call themselves ASPs, and a great number of these are telcos. A good resource for information on ASPs is the Application Service Provider Industry Consortium (ASPIC) at www.aspindustry.org.

MSPs MSPs specialize in providing software for network management and for monitoring applications, network performance, and network security. MSPs also take over the actual management and monitoring of the network. They operate similarly to ASPs, in that they use a network to deliver services that are billed to their clients. They differ from ASPs in that they very specifically address network management rather than business process applications. MSPs deliver system management services to IT departments and other customers that manage their own technology assets.

The appeal of the MSP model is that it eliminates the need for companies and individuals to buy, maintain, or upgrade information technology infrastructure management systems, which typically require a major capital expense, are highly technical in terms of the expertise they mandate, and require a considerable investment of time. This model appeals in particular to enterprises that manage e-commerce applications, such as ASPs and ISPs, whose expertise lies in the applications or network infrastructure they provide to customers—not necessarily in their management. It also appeals to small- and medium-size companies that prefer not to invest in large IT staffs. As with the ASP model, using specialists to deploy and maintain complex technology enables companies to focus on their own core competencies and to immediately tap into high-quality expertise as needed. There are variations in the model: some MSPs provide tools and services, whereas others provide services only; and some target corporations, whereas others are designed for consumers.

The MSP Association (www.mspassociation.com) was formed in June 2000, and it aims to be at the forefront of creating new standards for network management and for defining the best practices for the network management market. Its first working group has the job of defining what the whole MSP market looks like. Market analysts expect the demand for MSPs to grow exponentially as an attractive alternative to internally run IT management applications.

OSPs The OSP is a provider that organizes the content for you and provides intuitive user navigation. An example of an OSP is America Online. You can liken an OSP to a department store in a shopping mall in the United States. You enter the mall through strategically placed doors. At that point, you can browse all the different content providers, or shops, in that shopping mall. You may get a little assistance through a directory board, but you must take the initiative to go to each of these locations to determine whether the content you seek is really there. A shopping mall usually has a major tenant, a large department store, that provides its own doors from the parking lot. When you enter through the department store's doors, the department store tries to hold you captive with its content. It has some of everything else you could find in all the individual boutiques throughout the mall, but at the department store, there's a more exclusive or tailored selection. It's been organized into meaningful displays, and assistants are available to guide you through the store to your selected content. So, an ISP is the shopping mall. It gives you access to all the Web sites out there, but you have to use search engines to narrow your search. An OSP, such as America Online, is like the department store, which gives you a more cozy and exclusive space in which to browse the same content.

VISPs The VISP offers outsourced Internet service, running as a branded ISP. It is a turnkey ISP product aimed at affinity groups and mass marketers that want to add Internet access to their other products and services. Services of a VISP could include billing, Web-site maintenance, e-mail and news services, customized browsers, and a help desk. VISPs today include AT&T, Cable & Wireless, GTE, IConnect, and NaviNet. Early customers include SurfFree.com and Liberty-Bay.com.

The Service Provider Value Chain

Figure 9.10 is a simplistic diagram of the current value chain, from a technology standpoint. The lower-tier ISP is getting fees, essentially subscription fees and perhaps hosting fees, from a retail end user—and that is the lower-tier ISP's cash flow in. The ISP's cash flow out is the fees that it pays for connection into the higher-tier ISP, as well as the money associated with the links it's leasing from the telecom operator. The higher-tier ISP is getting fees for access in from the lower-tier ISP, and it's also getting subscription fees from higher-end business customers. The higher-tier ISP's outflow is the money it pays to connect into the backbone provider, as well as the money it may be paying for its communication links from a network operator. The backbone provider, then, is getting money from the higher-tier ISPs, as well as from customers that want to host their content—that is, from their Web farms or media servers.




Figure 9.10 The current value chain

Regulatory Decisions

Regulatory decisions could cause some major shifts in the next several years. These decisions will affect who ends up being considered the major backbone provider by virtue of how it regulates this environment. Decisions that are made on unbundling the local loops will have a major impact on the availability of broadband access. The Internet is not going to grow without growth in broadband access, so these decisions play a very important role. Decisions will be made about things such as privacy and whose laws apply in e-commerce. When I buy something in the United Kingdom from the United States, does the U.K. law or U.S. law apply? Further decisions will affect content sensitivity, censorship, and the acceptance of digital signatures as being valid. Interconnection issues need to be decided as well (for example, in Figure 9.10, how moneys are exchanged between the backbone providers).

So, remember that you need to consider more than just the progress and the technology. There are some big human and political issues to think about as well.

The backbones until now have largely practiced Sender Keep All (SKA). Those in an SKA arrangement assume that there is an even exchange of traffic between the two peers, and hence they don't pay each other any money. This is likely to change. The vast majority of content still currently resides in the United States, and that's made some of the U.S. ISPs rather cocky. They tell their overseas colleagues, "If you want a transit link into my backbone, you have to pick up the entire cost of

that transit link because basically your people are coming to get our content. There's nothing we really want on your side." One of two things will happen. Either this will become regulated or market forces will take over and we will allow these other countries to say, "We'll develop our own content, and now when your people want to get at it, you can cover the cost of access into *our* backbone."

■ Internet Interconnection and Exchange

As mentioned earlier, NSPs can interconnect at one of several points. First, four NAPs in the United States are partially funded by the National Science Foundation (Ameritech, PacBell, Sprint, and MFS Datanet), and the reason they get some support from the U.S. government is that they also provide interconnection into Internet 2, the academic research backbone. Second, connection can be made via agreements between ISPs (that is, peering). Third, there are commercial interconnection and exchange points, which can be used through private NAPs.

NAPs

Remember that the local service providers typically connect into the NAPs, or exchange points, through top-tier ISPs rather than directly. So what's the definition of a NAP? A NAP is the place where NSPs and/or ISPs exchange traffic with their counterparts. It's a public meeting point that houses cables, routers, switches, LANs, data communications equipment, and network management and telecommunications links. NAPs enable the NSPs and top-tier ISPs to exchange Internet traffic without having to send the traffic through a main transit link. This translates to decreased costs associated with the transit links, and it reduces congestion on the transit links.

Figure 9.11 shows the inside of one of the NAPs, Pacific Bell's point in California. On the left are the various ISPs and NSPs that are connecting to the NAP; note the various ISPs and their connection over a DSU into the Pac Bell ATM switch. The ATM switches are redundant and connected by redundant optical carrier levels (OC-3 initially, but they are being upgraded all the time). In addition, there are route servers, which maintain the databases of the appropriate path to take, based not just on the routing algorithm but also on the policy that the ISPs want to observe. In other words, there may be relationships about how traffic is passed between two providers—not just a next-hop or lowest-cost scenario—that determine what the path is. The route servers reside on a LAN, in this case, a fiber-based LAN called Fiber Distributed Data Interface (FDDI). FDDI was the first of the 100Mbps LAN backbones, and today most of these LANs would be Gigabit Ethernet. Resources connect into DSUs that lead to the routes of the selected NSPs.

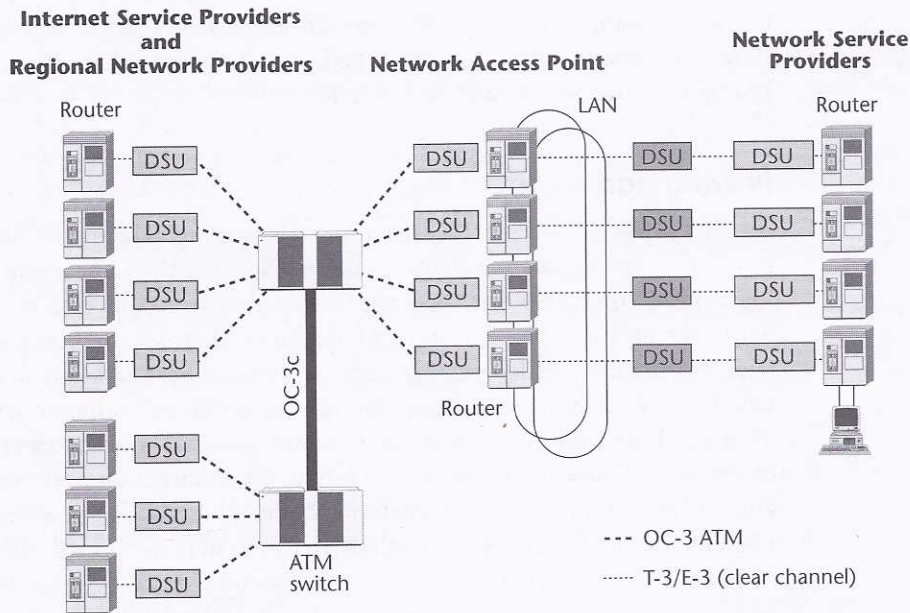


Figure 9.11 The Pacific Bell NAP

Routing arbiters are in place so that if there's dispute about how traffic has been routed, the arbiters can step in and arbitrate the situation.

The NAPs have equipment that's quite similar to what you find inside an Internet POP, which is quite similar to what you find inside a telco POP. This speaks to the convergence between the information types and their supporting infrastructures.

NAPs are becoming more discriminating. For instance, some require that the minimum connections be T-3 or E-3 and that they be redundant as well. That eliminates a lot of the smaller players from connecting to NAPs. If the NAPs were not so discriminating, there would be tremendous congestion at the public exchange points. For any traffic that's time sensitive, such as Internet telephony, VoIP applications, networked interactive games, multimedia, and streaming video, this congestion will cause problems with reliability and predictability. NAPs are also increasing in number; there are about 175 NAPs worldwide today. (For more information on NAPs, go to www.ep.net.)

Besides the four government-funded NAPs in the United States, other NAPs are for profit, and they charge per connection into the switch or router, initially in the US\$4,000 to US\$6,000 range. The cost varies with the economic times as well as the speed of the connection. And the NAPs, again, demand that you be able to guarantee a level of QoS before you connect at that exchange point.

Despite the fact that NAPs are now becoming more discriminating, they have become a point of congestion. Losses and delays are negatively affecting applications everyone wants to see developed, and therefore other alternatives have been brought about, as discussed in the following sections.

Peering Agreements

An alternative to NAPs is the use of private peering agreements. In a peering solution, operators agree to exchange with one another the same amount of traffic over high-speed lines between their routers so that users on one network can reach addresses on the other. This type of agreement bypasses public congestion points, such as NAPs. It's called *peering* because there's an assumption that the parties are equal in that they have an equal amount of traffic to exchange with one another. That's an important point because it obviously makes a difference in how money is exchanged. This is an issue of concern at the moment and one that many people are studying to determine ultimately whether there has to be a regulatory mechanism or whether market forces will drive it, but in general, with the first generation of peering agreements there was an understanding that peers were equals. Newer agreements often call for charges to be applied when traffic levels exceed what was agreed to in negotiations.

The most obvious benefit of peering is that because two parties agree on working with one another, and therefore exchanging information about the engineering and performance of their networks, the overall performance of the network is increased, including better availability, the ability to administer service-level agreements, and the ability to provide greater security. Major backbone providers are very selective about international peering, where expensive international private line circuits are used to exchange international routes. Buying transit provides the same benefits as peering, but at a higher price. Exchanging traffic between top-tier providers basically means better performance and fewer routers involved. And again, these types of arrangements are critical to seeing the evolution and growth in IP telephony, VoIP, and multimedia.

One problem with peering is that it can be limited. Under peering arrangements, ISPs often can have access only to each other's networks. In other words, I'll agree to work with you, but you can work only with me; I don't want you working with anyone else. Exclusivity types of demands might arise.

Private NAPs

The alternative to NAPs and peering arrangements is to use a private NAP, also called an overnet. Private NAPs are connected directly, via private lines, to the IP

Considerations for the Future

We need to be careful what we consider to be ready for primetime business, especially when it comes to the public Internet, that uncontrolled subcollection of networks that number more than 150,000. For example, one of the U.S. NAPs, MAE East, which handles something like 33% of the world's Internet traffic, for a long time resided in an underground parking garage in an unprotected space. Since that information was published, the problem surely has been resolved, but this was a condition whereby a person who was going in reverse with a very heavy foot could have taken down a third of the Internet.

There's a NAP in California as well: MAE West. MAE West is one of the busiest of the NAPs, connecting points between various NSPs. It's so crowded with equipment that the air conditioning can't keep up at times.

There was an instance not too long ago, in a small Florida ISP, in which a command entry error during router setup redirected 25% of all Internet traffic across a single T-1 link, which caused a two-hour brownout on the global Internet. The ISP hadn't used BGP filtering, but if it had, that message may have been caught and so much of the world's traffic would not have been redirected to what amounted to a local service provider backbone.

Finally, there was a software glitch at Network Solutions, which managed the top-level domain names for the entire Internet. A file containing more than one million Internet addresses got corrupted. Worse yet, a technician ignored the alarm that was issued and allowed the files to be replicated across the Internet, thereby blacking out Web sites for several hours.

So, I very much appreciate the candor of Mike O'Dell, chief scientist for UUnet Technologies, who used to say about the public Internet, "If you're not scared, you just don't understand."

But the Internet is not all bad and scary. The problems described here mainly affect business-class applications, where you need stringent controls. We also see a great many benefits and cost-efficiencies as a result of the public Internet.

backbones of several of the major NSPs or backbone providers. This design means that customers can get access to the major backbones without having to peer at the congested NAPs. The private exchange enables second-tier providers to connect to several first-tier providers, all in one location as well. Examples of private NAPs are InterNAP, Savvis Communications, and Digital Island. Savvis opened a private exchange point in London, and it expects to connect two first-tier ISPs there. InterNAP is opening exchanges in London and Amsterdam, and it also has plans for Frankfurt and Paris, expecting to connect with six first-tier ISPs, as well as two major in-country providers. Thus, InterNAP currently has some 50 NAPs that are

live or about to open. Both Savvis and InterNAP run exchanges in the United States, and each also claims to get great deals on transit relationships or even to negotiate peering relationships that are unavailable to smaller ISPs.

For more learning resources, quizzes, and discussion forums on concepts related to this chapter, see www.telecomessentials.com/learningcenter.