chapter 11

Architecting the Internet

"... barring total disaster, all elements are eventually acknowledged, even if they require retransmission."

—TCP inventor Vint Cerf, 1973

In this chapter, you will learn how to:

- Define how the four layers of the Internet Architecture map to the seven layers of the OSI Reference Model.
- Explain the Internet addressing rules and configure TCP/IP on a personal computer.
- List the network utilities used to analyze, troubleshoot, and optimize Web sites for maximum performance.
- Explain how domain names map to IP addresses and define the roles that different kinds of Internetworking servers play in transmitting information to these addresses.
- List the content and delivery services that the most popular server products provide to end users via the Internet.

In the previous chapter, you studied the theory behind the seven layers of the OSI Reference Model (OSI/RM). This chapter puts theory into practice by showing you how the Internet implements the OSI/RM through a suite of protocols called TCP/IP. Because it powers the Internet, TCP/IP is the most famous protocol suite in the world.

This chapter begins by explaining the process through which TCP/IP implements the OSI/RM protocols for packet creation, addressing, and routing. Then you learn how to configure TCP/IP on PCs and workstations and optimize them for maximum network performance.

Every network must have one or more servers to respond to requests and provide the services that are the reason why you created the network in the first place. The Internet has many kinds of servers. Web servers, for example, respond to requests from browsers. Mail servers, news servers, FTP servers, and streaming media servers provide communication, entertainment, and information resources that are very popular among end users.

Behind the scenes are servers that are not so well known. Some enhance network performance, while others provide catalog, directory, and back-end database services. This chapter concludes by taking you behind the scenes to the Internetworking servers that provide the infrastructure for the more well-known services that end users crave.

Understanding TCP/IP

Emerging from the pioneering efforts of Vint Cerf and others in the 1970s, TCP/IP became the U.S. military's preferred networking protocol in 1980. When the major networks adopted it in 1983, TCP/IP officially became the Internet's protocol suite. Understanding TCP/IP begins with a comparison of the Internet's Architecture to the OSI Reference Model.

Internet Architecture

As you learned in the previous chapter, the OSI/RM has seven layers. The **Internet Architecture** model, on the other hand, has four layers. Figure 11-1 compares the seven OSI/RM layers to the four Internet Architecture layers. Each Internet Architecture Layer corresponds to one or two of the OSI/RM layers.

Internet Engineering Task Force (IETF)

The **Internet Engineering Task Force (IETF)** is the standards body in charge of defining Internet protocols. Any interested person or organization can join the IETF by following the link to membership at www.ietf.org.

The IETF defines protocols through a **Request for Comments (RFC)** process. Most of the RFCs are generated by working committees of the IETF, but any individual or organization can submit an RFC for the IETF to consider. Once an RFC has been accepted for consideration, the IETF assigns it an **RFC number** and an **RFC maturity state**. An RFC moves through three maturity states on its way to becoming a standard:

- **Proposed** The protocol has been accepted for consideration and may possibly advance through the rest of the stages to become an Internet standard. The IETF invites interested parties and researchers to test the proposed protocol and to identify any issues or problems that may need to be solved prior to standardization.
- **Draft** The protocol passed its preliminary testing at the proposal stage. The IETF is now seriously considering adopting the protocol as an Internet standard. The IETF invites large-scale testing and requires that at least two independent trials prove the interoperability of the proposed protocol. If major problems arise, the protocol's maturity level may be downgraded to the proposed state.
- **Internet standard** The protocol passed the testing and proved its merit on the Internet. The RFC is considered to be an established Internet standard.

In special circumstances, the IETF may assign three other maturity states to an RFC. These other three states are

- **Experimental** The protocol is not ready for testing outside carefully controlled laboratory situations. Testing on the public Internet could cause protocol conflicts.
- **Historic** The standard has become obsolete or unnecessary, or it has been superseded by a newer standard.
- **Informational** The standard is from a non-IETF vendor or standards body. The IETF provides access to the standard for the benefit of site visitors, but it is not being proposed as an Internet standard.

You can bring any RFC onscreen by typing its number into the search engine at www.ietf.org/rfc.html. The complete list of RFCs is at www.ietf.org/iesg/1rfc_index.txt, which indexes the protocols in numerical order. If you take a moment to scroll down through this list, it provides an interesting historical perspective on the evolutionary nature of the Internet.

Internet Architecture Layers	*OSI/RM Layers*
Application	7. Application 6. Presentation
Transport	5. Session 4. Transport
Internet	3. Network
Network Access	2. Data link 1. Physical

FIGURE 11-1 *The Internet Architecture model implements in four layers the seven separate layers of the OSI Reference Model. The color-coding in this figure illustrates that each layer of Internet Architecture corresponds to one or two of the OSI/RM layers.* ■

Application Layer Protocols

At the top of the Internet Architecture's protocol stack is the **Application Layer**, which encompasses OSI/RM Layers 7 and 6. Some of the Internet's most well-known protocols reside at the Application Layer. The Application Layer protocols and their corresponding RFC standard numbers are listed as follows:

- **Hypertext Transfer Protocol (HTTP)** True to its name, HTTP is the protocol that transfers hypertext Web pages across the Internet. There are two versions known as HTTP 1.0 and HTTP 1.1. Version 1.0 opens a separate connection for each downloaded file. HTTP 1.1 avoids the overhead of creating separate sessions for each file by keeping the connection open so multiple files associated with a page can be downloaded all at once. Both the client and the server must be running HTTP 1.1 for the multiple download feature to work. Otherwise, the files get downloaded separately via HTTP 1.0. The IETF standards are RFC 1945 for HTTP 1.0, and RFC 2616 for HTTP 1.1.
- **File Transfer Protocol (FTP)** RFC 959 specifies the File Transfer Protocol, which you use when you log on to an FTP server to transfer files over the Internet from one computer to another. Chapter 13 teaches you how to use an FTP client that encrypts the transfers to keep them secure as your documents wind their way across the Internet.
- **Trivial File Transfer Protocol (TFTP)** RFC 1350 defines a simpler form of FTP called Trivial File Transfer Protocol, which diskless workstations and some routers use to get their configuration files during startup. Unlike FTP, TFTP uses the connectionless UDP protocol that does not have a logon process for user authentication.
- **Telnet** RFC 854 defines telnet, which is the terminal emulation protocol that enables users to log on to remote host computers over the Internet.
- **Gopher** RFC 1436 specifies the Gopher protocol for distributed document search and retrieval. As you learned in Chapter 2, Gopher was the rage prior to the invention of the World Wide Web.
- **Simple Mail Transfer Protocol (SMTP)** RFC 821 defines the Simple Mail Transfer Protocol, which specifies the rules for transferring e-mail over the Internet. A related standard that is not part of TCP/IP is RFC 1939, which specifies how version 3 of the Post Office Protocol (POP3) stores mail on a server until users log on and download it.
- **Network News Transfer Protocol (NNTP)** RFC 977 specifies the NNTP protocol that powers USENET newsgroups. The newsgroup servers to which you connected to read news in Chapter 2, for example, are NNTP servers.

- **Domain Name System (DNS)** RFCs 1034 and 1035 define the DNS protocol that translates a fully qualified domain name (e.g., www.loc.gov) into a numeric IP address (e.g., 140.147.249.7) needed to route information across the Internet.
- **Simple Network Management Protocol (SNMP)** RFC 1157 specifies SNMP, which network administrators use to remotely manage TCP/IP network devices that are SNMP compliant. An administrator who needs to reconfigure or get statistics from a router, for example, can use an SNMP utility to contact the router. Network administrators like how SNMP enables them to administer all of their SNMP network devices from a central location.
- **Bootstrap Protocol (BOOTP)** RFC 951 defines BOOTP, which is a startup protocol that enables a workstation to discover configuration information including its IP address, router address, and DNS server address.
- **Dynamic Host Configuration Protocol (DHCP)** RFC 2131 defines DHCP, which works in conjunction with BOOTP to assign during workstation initialization an IP address, router address, and other configuration parameters. During startup, the client computer sends a DHCP message to which a DHCP server, if present, responds. This automatic IP configuration process is very popular among network administrators, because it saves the time they would otherwise need to spend manually configuring each workstation.

Transport Layer Protocols

The Internet Architecture's **Transport Layer** encompasses OSI/RM Layers 5 and 4. It is the responsibility of the Transport Layer to divide into packets the data received from the Application Layer. Depending on the kind of session being serviced, the Transport Layer uses one of the following two protocols:

- **Transmission Control Protocol (TCP)** The vitally important role of the Transmission Control Protocol (TCP) is evident from the inclusion of its acronym in the name of the TCP/IP protocol suite. As defined by RFC 793, TCP establishes and manages the connection between the computers that are exchanging data, numbers the packets on the sending computer, reassembles the packets on the receiving computer, and ensures that all of the data is intact with no omissions or duplications.
- **User Datagram Protocol (UDP)** As defined by RFC 768, UDP is a connectionless protocol that does not require the negotiation and establishment of a session between the sending and receiving computers. Instead, the sending computer simply puts each output from the Application Layer into a packet. Because there is no provision for the resending of lost packets, UDP is considered to be an unreliable transport protocol that should be used only by applications that transmit relatively small amounts of data without establishing a session between the sending and receiving computers.

Internet Layer Protocols

In the Internet Architecture, the **Internet Layer** corresponds to layer 3, which is the Network Layer, of the OSI/RM. It is the responsibility of the Internet Layer to take the packet from the Transport Layer, determine the best way to route it across the Internet, and transform it into an IP packet containing an IP header and trailer. To accomplish this, the Internet Layer uses the following protocols:

- **Internet Protocol (IP)** As defined by RFC 791, the Internet Protocol (1) determines the best path for routing the packet to its destination address, (2) addresses the packet accordingly, and (3) fragments the packet if it is too long for the network segment. The critical importance of these tasks earned IP its place in the name of the TCP/IP protocol suite.
- **Address Resolution Protocol (ARP)** In order to transmit a packet from one node to another, the Physical Layer must know the nodes' physical MAC addresses. RFC 826 defines how ARP translates IP addresses into physical MAC addresses.
- **Reverse Address Resolution Protocol (RARP)** Given the MAC address of a network device, RARP determines its IP address. RFC 903 defines this process as Reverse ARP (RARP), because it reverses the translation performed by ARP.
- **Internet Group Management Protocol (IGMP)** RFC 1112 defines the IGMP group management protocol used for **multicasting**, which is the sending of a message from one computer to a group of IP addresses belonging to users who subscribe to the group. Videoconferencing is an example of multicasting.
- **Internet Control Message Protocol (ICMP)** RFC 792 defines the ICMP messaging protocol that TCP/IP uses for troubleshooting. Routers and servers normally send error messages in ICMP packets when things fail on a TCP/IP network.

Network Access Layer Protocols

In the Internet Architecture, the **Network Access Layer** corresponds to OSI/RM Layers 2 and 1, which transform the packets into a binary encoded stream of 0's and 1's for transmission over the physical network. Then the NIC transforms the 0's and 1's into the signals that get transmitted physically over the network. The specific protocols are determined by the device drivers and the physical connections from the workstation to the network cable or wireless transmission medium.

Figure 11-2 summarizes the Internet Architecture by showing the protocols that come into play as the messages work their way down from the Application Layer to the Network Access Layer.

Demultiplexing

This chapter presented the four layers of the Internet Architecture in the order in which a message on the sending computer would pass down through the Application, Transport, Internet, and Network Access Layers to the physical transmission medium that carries the packet across the Internet. When the packets reach their destination on the receiving computer, the message must ascend back up through the same four layers, which decode, unpack, and reassemble the message and make sure it arrives without error. The process of unpacking the message by processing and removing the headers added to the packets at each layer is called **demultiplexing**. There are four stages in the demultiplexing process:

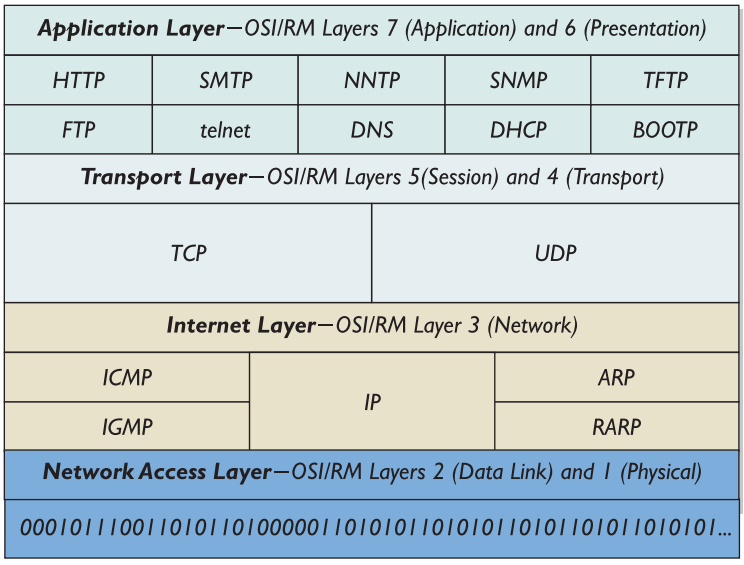

FIGURE 11-2 *Protocols that come into play as messages work their way through the four layers of the TCP/IP protocol suite.* ■

1. On the receiving computer, the Network Access Layer takes a look at the packet and uses the MAC address to determine whether it should be processed here. Packets that do not get processed here are ignored and are passed on to other network nodes.
2. Packets that belong here pass up to the Internet Layer, which takes a look at the IP addressing in the packet to determine whether any further routing is required.
3. The Transport Layer takes a look at the TCP or UDP port number to determine which service needs to receive the message.
4. The Application Layer passes the message to the service or application that will act on the message. If the message contains an HTTP request for a Web page, for example, the server answers by sending the page.

Routing

Routing is the process of determining the network path over which packets are sent. It is the responsibility of the Internet Protocol to determine this optimal path. The two main kinds of routing are direct and indirect.

Direct routing occurs when two computers on the same network communicate with each other. Direct routing happens on an Ethernet, for example, when the Address Resolution Protocol (ARP) converts an IP address into a MAC address to transmit packets between sending and receiving computers on the same local network.

Indirect routing happens when the sending and receiving computers are not on the same local network. The packets get sent to the MAC address of a router, which analyzes the destination address and forwards the packet

on toward the appropriate network. The router to which the packets first get sent is called the **default gateway**, which is one of the IP addresses you set when you manually configure a PC for network operation. If the PC is configured for dynamic host addressing, the DHCP server determines the IP address of the default gateway.

The router uses the **routing information table** to keep track of the routes over which it will send packets to different networks to which the router is connected. When a packet arrives, the router examines the packet's destination IP address, consults the routing table to determine the optimum route, and sends the packet to the corresponding router. Each trip between routers is called a **hop**. The number of trips between the packet's origin and destination addresses is called a **hop count**. Figure 11-3 illustrates that routers typically send packets along the path with the lowest hop count, unless network traffic conditions dictate otherwise.

Static versus Dynamic Routers

Network administrators consider a router to be static or dynamic, depending on the nature of the router's information tables. If the tables are fixed and can be updated only by manual changes made by the network administrator, the router is said to be a **statically configured router**. Routers in stable networks that connect to a relatively small number of other networks can use static routing tables effectively. If a new route is added to the network, however, the network administrator has to add the new path to all of the static routing tables to enable them to route traffic along that path.

A dynamic router, on the other hand, communicates with other routers to exchange information about new routes that have been added, or old routes that are no longer available. As this information changes, the router information table updates automatically; hence the term **dynamically configured router**.

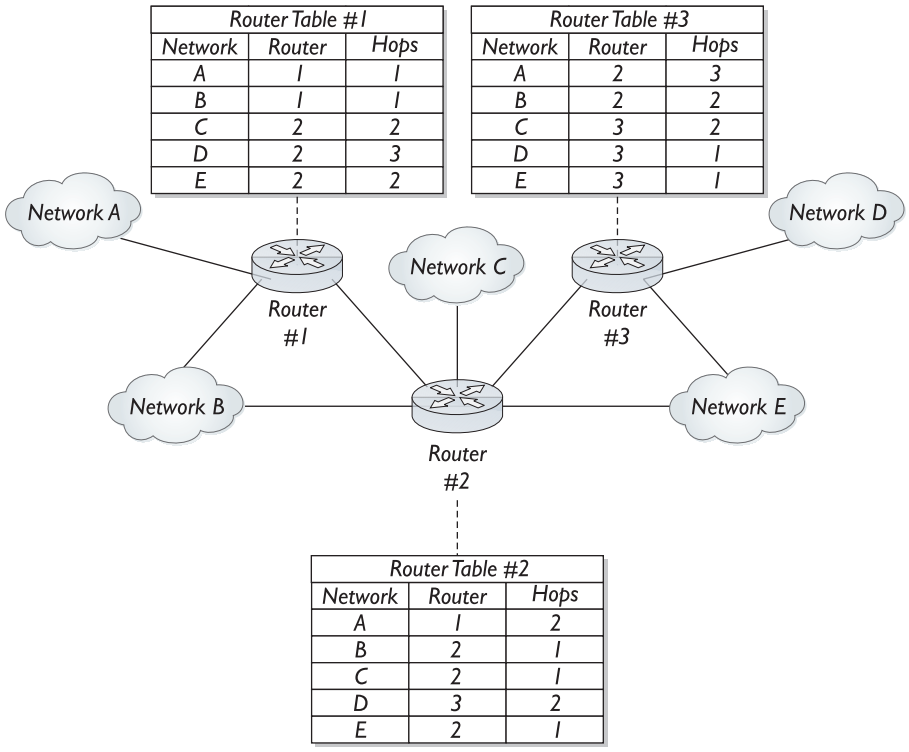

FIGURE 11-3 *Routers consult routing information tables to determine the best path for sending the packet on toward its destination. In this example, the routing follows the paths with the lowest hop counts.* ■

Routing Protocols

When routers share information, they follow routing protocols that define how to communicate changes in the routing tables. There are two basic kinds of routing protocols, namely, exterior and interior. As you might expect, exterior routing protocols are

for communicating with routers that are outside of an organization's network. If you are studying for the CIW exam, you should learn the names of the following two exterior routing protocols:

- **External Gateway Protocol (EGP)** Defined by RFC 904, the EGP defines the protocol used to exchange net-reachability information between Internet gateways belonging to autonomous systems. An **autonomous system (AS)** is a set of routers under a single technical administration, such as an ISP or a backbone such as the NSFNET.
- **Border Gateway Protocol (BGP)** Defined by RFCs 1267 and 1268, the BGP builds on experience gained with EGP on the NSFNET backbone. A border router keeps track of the status of neighboring AS's and uses a pruning process to select optimum routes.

Interior routing protocols are for communicating with routers inside of an organization's network. You should know the following two interior routing protocols:

- **Routing Information Protocol (RIP)** Defined by RFCs 1058 and 2453, RIP is a protocol whereby routers periodically send their information tables every thirty seconds across their network connections to their neighboring routers. Based on this information, dynamic routers update their tables to reflect any changes on the neighboring routers and route packets over the path with the lowest hop count.
- **Open Shortest Path First (OSPF)** Defined by RFC 2328, OSPF is a routing information protocol that improves upon RIP in three ways. First, changes in router tables get exchanged as soon as they happen, instead of having to wait thirty seconds. Second, only the changes get sent, instead of the whole table, thereby saving bandwidth. Third, and most importantly, OSPF exchanges statistics on the transmission speed of multiple possible routings, enabling the router to take advantage of faster routes that would otherwise be unused under the hop count rule. Hence the name, Open Shortest Path First (OSPF), which can use routes that have faster transmission times even though they may have higher hop counts.

Port Numbers

Servers on the Internet typically run many services at once. It is common, for example, to have HTTP Web serving, SMTP mail transfer, DNS name resolving, telnet remote logon, FTP file transfer, and POP3 mail delivery services running on a single server. To provide the Transport Layer with a fast way of determining which application should receive an incoming request, each packet contains a destination **port number**, which is a 16-bit number indicating the service that the packet is using. The range of a 16-bit

Port Number	Protocol
21	FTP
23	Telnet
25	SMTP
53	DNS
80	HTTP
110	POP3
194	IRC

TABLE 11-1 *Well Known Port Numbers for Selected Internet Services* ■

number is from 0 to 65535. Of these, the ports numbered from 0 to 1023 are known as **Well Known Ports**. The Internet Corporation for Assigned Names and Numbers (ICANN) is in charge of assigning these well-known, or reserved, port numbers. Table 11-1 lists the port numbers for some of the Internet's most popular services.

Ports numbered from 1024 to 49151 are called **registered ports**. Many well-known applications use these registered ports. Microsoft SQL Server, for example, uses port 1433, Macromedia Shockwave uses 1626, and Cisco License Management uses 1986.

The rest of the port numbers from 49152 to 65535 are called **Dynamic and/or Private Ports**. ICANN does not control these port numbers, and any process can use them. Because they are intended for short-lived processes, they are also known as ephemeral or temporary ports. The complete list of assigned ports is at www.iana.org/assignments/port-numbers.

nternet Addressing

To participate as a node on the Internet, every computer or network device must have a unique IP address. As you learned in Chapter 1, an IP address consists of four bytes separated by periods. Each byte is an eight-bit number that ranges in value from 0 to 255. The smallest address is 0.0.0.0 and the largest is 255.255.255.255. The number of IP addresses this scheme allows is 256^4, which is 4,294,967,296.

The Internet Corporation for Assigned Names and Numbers (ICANN) is in charge of assigning IP addresses. Every IP address consists of two basic parts: the Network ID and the host ID. The Network ID always comes first, followed by the host ID. Depending on the size of the network, the Network ID occupies the first one, two, or three bytes in the IP address. The remainder of the address is the host ID.

Internet Address Classes

When an organization applies for a range of IP addresses to serve the computers in its network, ICANN considers the size of the organization and determines the size of an address space to allocate. Five **Internet address classes** can be assigned. These classes are named A, B, C, D, and E.

- **Class A** Extremely large organizations that have more than 16 million hosts receive Class A Internet addresses. In a Class A address, the first byte in the dotted quad IP address is the Network ID, and the other three bytes are the host ID.
- **Class B** Medium to large organizations with up to 65,534 hosts get Class B addresses, in which the first two bytes are the Network ID, and the other two bytes are the host ID.

- **Class C** Small organizations with up to 254 hosts get Class C addresses, in which the first three bytes are the Network ID, and the last byte is the host ID.
- **Class D** Multicast groups receive Class D addresses, which are set aside for multicasting. In a Class D address, all four bytes are the Network ID. There is no host ID because everyone in the group receives the multicast.
- **Class E** The Internet has reserved some addresses for future use. These are the Class E addresses.

Figure 11-4 shows the range of IP addresses set aside for each of the five Internet address classes. Studying this figure will enable you to recognize the class of a network from the first byte of its IP address.

IP Addressing Rules

If you do the math, you may wonder why some of the network numbers in the last two columns of Figure 11-4 do not seem to encompass all of the possible network addresses. The reason is because of special rules whereby the Internet reserves certain addresses for special purposes. The special cases are loopback, broadcast, network source, and private IP addresses. These four cases are described in the sections that follow.

IP Loopback Address

The Internet reserves the Network ID 127 as the **loopback address**, which is a diagnostic IP address reserved for testing purposes that redirects packets to the same computer that sent them. When a developer wants to test a Web server running on the local host (i.e., the same computer that the developer is using), the developer typically uses the address 127.0.0.1. You can find out whether a Web server is running on your computer, for example, by using your browser to go to http://127.0.0.1. If you get no response, there is no Web server on your local host.

Address Class	IP Address Range	IP Structure (Network ID versus Host ID)	Potential Number of Networks	Potential Number of Hosts per Network
A	0.0.0.0 to 127.255.255.255	24.131.47.114	126	16,777,214
B	128.0.0.0 to 191.255.255.255	134.123.174.201	16,384	65,534
C	192.0.0.0 to 223.255.255.255	201.113.241.196	2,097,152	254
D	224.0.0.0 to 239.255.255.255	230.148.32.157	1,048,560 multicast goups (no networks or hosts)	
E	240.0.0.0 to 255.255.255.255	Class E is reserved for future use.		

FIGURE 11-4 *The five Internet address classes are named A, B, C, D, and E. In a Class A Internet address, the first byte can be 0 to 127. In Class B addresses, the first byte ranges from 128 to 191. The first byte in Class C ranges from 192 to 223. Multicast addresses in Class D have first bytes ranging from 224 to 239. Class E comprises IP addresses that have first bytes ranging from 240 to 255, which are reserved for future use.* ■

IP Broadcast Addresses

A broadcast is a message that gets sent to all of the hosts on a network. The **IP broadcast address byte** is 255, which sets all eight bits in the address byte to 1. There are four kinds of broadcast addresses:

- **Limited broadcast** The **limited broadcast** address is 255.255.255.255. Routers block this address, keeping it inside the local network—hence the name, *limited broadcast*. Computers that do not have IP addresses typically send a 255.255.255.255 broadcast on startup to find a DHCP or BOOTP server that can respond with an IP address assignment.
- **Net-directed broadcast** On a Class A network, the **Net-directed broadcast** address is *netid*.255.255.255, where *netid* is the Network ID. This broadcasts a message to all the hosts on that network. On a Class B network, the Net-directed broadcast address is *netid.netid*.255.255, where *netid.netid* is the Network ID. On a Class C network, the Net-directed broadcast address is *netid.netid.netid*.255.
- **Subnet-directed broadcast** When a large network is divided into subnets, you send a **subnet-directed broadcast** to a subnet by giving as much of the Network ID as needed to identify the subnet, followed by 255 or 255.255, depending on how many bytes were needed to identify the subnet. You learn how to create subnets later in this chapter.
- **All-subnets-directed broadcast** The all-subnets-directed broadcast sends a message to all hosts on a network. It is now considered obsolete, because multicasting using Class D addresses has replaced it.

Special Case IP Addresses Containing Zeros

IP addresses in which the Network ID or the host ID are all zeros are special network addresses that cannot be assigned as a host's IP address. Any zeros used in IP addresses must observe the following rules:

- In a Class A network address, the first byte cannot be zero, and the last three bytes cannot all be zeros.
- In a Class B address, the first two bytes cannot both be zeros, and the last two bytes cannot both be zeros.
- In a Class C address, the first three bytes cannot all be zeros, and the last byte cannot be zero.
- A special case is the IP address 0.0.0.0, which is the source address that dynamically configured hosts use when broadcasting a request for an IP address. You may never assign a host a permanent IP address of 0.0.0.0.

Reserved or Private IP Addresses

In the previous chapter, you learned that an organization can use a NAT to enable many private internal IP addresses to access the Internet through one or more ICANN-assigned external IP addresses. Whenever an organization creates private, internal IP addresses, they must fall within the following ranges:

- **Class A** 10.0.0.0 to 10.255.255.255
- **Class B** 172.16.0.0 to 172.31.255.255
- **Class C** 192.168.0.0 to 192.168.255.255

Routers on the public Internet reject packets that use these private addresses. When users dial up to the Internet, their ISP usually assigns them one of these private addresses. Out on the public Internet, the ISP uses one of its public IP addresses on behalf of the user. The switch happens behind the scenes, so to speak, and the vast majority of end users are happy to be totally unaware of it.

Subnet Masks

A **subnet mask** is a dotted quad number that enables the local network to determine whether any given IP address is internal or external to the local network. If a network node makes a request from an IP address that is internal, the local network knows to handle it without routing the request over the Internet. If the request is an external address, on the other hand, the routers send the request on its way over the Internet.

Most subnet masks have one of the following forms:

- **Class A: 255.0.0.0** The 255, which is a byte with all bits on, signifies that the first byte of this host's IP address is its local Network ID. The zeros in the other three bytes indicate that all of the addresses within that Network ID are internal to that network.
- **Class B: 255.255.0.0** The two leading 255s mean that the first two bytes of this host's IP address are its local Network ID. The zeros in the other two bytes indicate that all of the addresses within that Network ID are part of the local network.
- **Class C: 255.255.255.0** The three leading 255s mean that the first three bytes of this host's IP address are its local Network ID. The zero in the fourth byte indicates that all of the addresses within that Network ID are part of the local network.

True to its name, the subnet mask can also be used to mask out parts of the local Network ID that are not on the local host's network segment. At my home, for example, the subnet mask is 255.255.255.240. That mask means that sixteen local addresses are on the subnet in my home (240–255). All other addresses are external to my subnet, making my router handle requests directed to those other addresses.

The subnet mask is a critically important part of a network host's configuration. If the subnet mask is incorrect, the errant host's packets can be misdirected and lost. Therefore, you want your organization's subnet masks to be set by network administrators who know what they are doing.

Configuring TCP/IP on a Personal Computer

In order for a personal computer to communicate on a TCP/IP network, you must first configure the computer's TCP/IP settings. At a minimum, the computer must have an IP address and a subnet mask, which enable communications with hosts on the local network. To communicate with computers on a WAN, such as the Internet, the computer must also have a default gateway.

The Internet uses the Domain Name System (DNS) to enable end users to access resources by name, such as www.loc.gov, instead of requiring users to type the site's numeric IP address. In addition to DNS, the Windows operating system has a naming system called **Windows Internet Naming Service (WINS)**. You must configure PCs to use either DNS or WINS. Because DNS is so important on the Internet, you should configure DNS if the computer has access to the Internet. WINS enables a Windows computer to be known by its NetBIOS name, which is the "computer name" on the Computer Name tab of the Windows Control Panel's System settings.

TCP/IP settings can be configured in two ways, namely, statically or dynamically. True to its name, a static configuration consists of predetermined settings that you type by hand into the computer's TCP/IP window. Dynamic configuration uses DHCP to configure the network settings automatically during the computer's booting process on startup. Happily, the settings for static and dynamic PC configuration are in the same place. Follow these steps:

1. Click the Windows Start button and choose Control Panel. When the Control Panel window appears, double-click Network Connections.
2. When the Network Connections window appears, double-click Local Area Connection to bring up the Local Area Connection Properties window, which displays in a listbox the protocols that have been assigned to your computer's NIC. **Protocol binding** is the act of assigning a protocol to a network interface card.
3. Scroll the listbox down to reveal the Internet Protocol (TCP/IP), as illustrated in Figure 11-5.
4. Click once on Internet Protocol (TCP/IP) to select it, and then click the Properties button to bring up the TCP/IP Properties window.
5. If you are dynamically configuring the PC, click to select the options entitled Obtain an IP address automatically and Obtain DNS server address automatically. Click OK to close the windows, and you are finished.

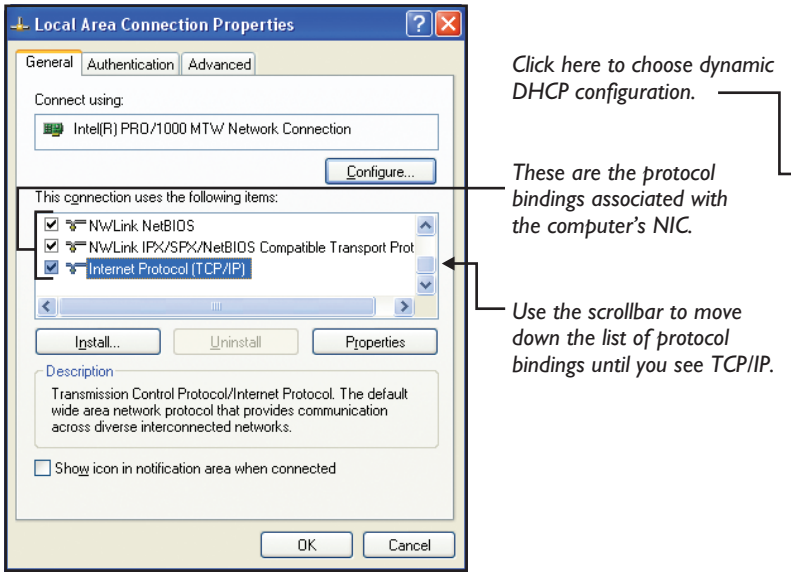

FIGURE 11-5 *When the Network Connections window appears, it displays in a listbox the network protocols on your computer. If the TCP/IP protocol is not visible, you need to scroll the list down to reveal TCP/IP.* ■

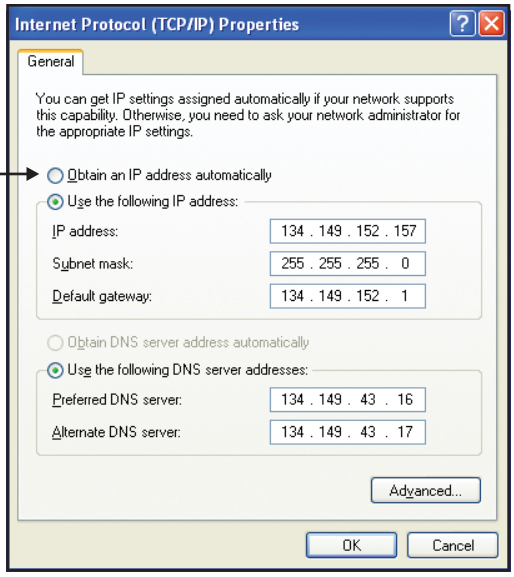

FIGURE 11-6 *The Internet Protocol Properties window lets you set a PC for static or dynamic TCP/IP configuration. If you also want to configure the PC for the Windows Internet Naming Service, click the Advanced button, then click the WINS tab.* ■

6. If you are statically configuring the PC, click to select the options entitled Use the following IP address, and Use the following DNS server addresses. Clicking these options activates the fields illustrated in Figure 11-6.
7. Fill out the fields by typing the IP address, subnet mask, default gateway (i.e., router), and DNS server addresses.
8. If you want to configure the PC for WINS, click the Advanced button, select the WINS tab, and use the Add button to add the network's WINS server address.
9. Click OK to close the windows.

Configuring Networks for Optimum Performance

Network technicians and troubleshooters use diagnostic tools that help solve problems and tune networks for optimum performance. This part of the book introduces seven such tools and provides hands-on experience monitoring network performance. The seven tools are (1) ping, (2) traceroute, (3) netstat, (4) ipconfig, (5) winipcfg, (6) arp, and (7) network analyzers.

ping

The most basic network troubleshooting utility is the Packet Internet Groper (**ping**), which sends ICMP echo request packets to a destination IP address. When the destination returns the echo, the ping utility measures the response time and displays a message indicating the duration of the

round trip. You use the ping utility when you want a quick test of whether a PC has connectivity with the network, or whether a certain server is responding. To try the ping utility, follow these steps:

1. Get a command prompt onscreen. If you need help doing this, click the Windows Start button and choose Programs | Accessories | Command Prompt.
2. At the command prompt, type: **ping 127.0.0.1**
3. Press ENTER, and see how your computer responds. Remember that 127.0.0.1 is the special IP address that lets you run tests in which your local computer issues commands to itself. By executing the command ping 127.0.0.1, you are testing whether your computer can ping itself. If this test fails, your computer's NIC is not configured properly or has failed. The response you get should look something like this:

```
Pinging 127.0.0.1 with 32 bytes of data:
Reply from 127.0.0.1: bytes=32 time<1ms TTL=128
Reply from 127.0.0.1: bytes=32 time<1ms TTL=128
Reply from 127.0.0.1: bytes=32 time<1ms TTL=128
Reply from 127.0.0.1: bytes=32 time<1ms TTL=128

Ping statistics for 127.0.0.1:
    Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
Approximate round trip times in milli-seconds:
    Minimum = 0ms, Maximum = 0ms, Average = 0ms
```

4. Now try pinging the Library of Congress. When I tried this, I got an average response time of 48 milliseconds. To ping the LOC, type one of the following commands, press ENTER, and watch what happens:

```
ping  140.147.249.7
```

or

```
ping www.loc.gov
```

A few hosts on the Internet have disabled ping. If you try to ping www.microsoft.com, for example, you will get no reply because Microsoft has disabled ping on that server.

traceroute

The **traceroute** networking utility reports the path data follow as a packet winds its way over the network from the source to the destination computer. You use the traceroute utility when you need to isolate the source of a network connectivity problem. The best way to see how traceroute works is to try it. On UNIX and Linux systems, the command is **traceroute**. Under Windows, however, the command is **tracert**. Assuming you have Windows, these instructions use tracert:

1. Get a command prompt onscreen and type:

```
tracert www.mcgraw-hill.com
```

2. Press ENTER to execute the command. When I tried this, I got the following report of 15 hops the packet took on its way from the University of Delaware to McGraw-Hill. For each hop, tracert reports three timings as it tries each hop three times:

```
Tracing  ^lroute to www.elb.mcgraw-hill.com [198.45.19.151]
over a maximum of 30 hops:

  1     5 ms    10 ms     5 ms  128.175.__.__ [the author's machine]
  2    37 ms    38 ms    38 ms  host-214-65.nss.udel.edu [128.175.214.65]
  3    70 ms    39 ms    38 ms  chp-rt2-v-7.nss.udel.edu [128.175.13.254]
  4    37 ms    39 ms    37 ms  spare-7206-g0-3-9.nss.udel.edu [128.175.111.11]
  5    40 ms    39 ms    39 ms  g1.ba21.b003003-1.phl01.atlas.cogentco.com [38.112.7.61]
  6    39 ms    39 ms    39 ms  g9-0.core01.phl01.atlas.cogentco.com [38.112.34.141]
  7    43 ms    46 ms    43 ms  p5-0.core02.jfk02.atlas.cogentco.com [66.28.4.1]
  8    42 ms    41 ms    41 ms  p6-0.pr01.jfk05.atlas.psi.net [154.54.1.166]
  9    41 ms    42 ms    42 ms  204.255.169.13
 10    41 ms    42 ms    42 ms  0.so-6-0-0.XL2.NYC4.ALTER.NET [152.63.21.82]
 11    42 ms    41 ms    42 ms  0.so-3-1-0.XL2.NYC9.ALTER.NET [152.63.21.14]
 12    41 ms    41 ms    42 ms  POS7-0.GW6.NYC9.ALTER.NET [152.63.24.69]
 13    42 ms    43 ms    42 ms  mgh-t3-gw.customer.ALTER.NET [157.130.18.70]
 14    45 ms    44 ms    45 ms  gw2.mcgraw-hill.com [198.45.19.20]
 15    43 ms    44 ms    43 ms  198.45.19.151
Trace complete.
```

netstat

The **netstat** utility displays information about the connections that are open and the protocol processes that are currently running on a network host. The name netstat stands for network statistics. To run the netstat utility, follow these steps:

1. Get a command prompt onscreen and type: **netstat**
2. Press ENTER to execute the command. You will get a list of the connections currently open on your machine.
3. The netstat utility can do much more, however. To see a list of the netstat options, type the following command and press ENTER: **netstat ?**
4. Try the netstat options that interest you. When you type an option, you must precede it with a hyphen. To see the router table, for example, you type**: netstat -r**

ipconfig

As you can tell from its name, **ipconfig** is a TCP/IP configuration utility that runs on computers with Windows NT/2000/XP/2003 or later operating systems. You use ipconfig whenever you want to inspect the current IP configuration. Furthermore, if the computer's IP settings are dynamically configured, you can use ipconfig to release, renew, or refresh the DHCP leases. To use ipconfig, follow these steps:

1. At a command prompt, type **ipconfig** and press ENTER. You will see a report revealing the computer's IP address, subnet mask, and default gateway.
2. If you want a more detailed report that includes the computer's name and physical MAC address, execute the command followed by the /all switch, as follows: **ipconfig /all**
3. To release the leases on IP addresses obtained from a DHCP server, type the following command: **ipconfig /release**
4. To renew the leases, type the following command: **ipconfig /renew**
5. To learn about other options, type the command: **ipconfig ?**

winipcfg

winipcfg is an older version of ipconfig for Windows 95/98/Me. If you want to run winipcfg on one of these older Windows versions, follow these steps:

1. Click the Windows Start button and choose Run to make the Run dialog appear.
2. In the Run dialog, type **winipcfg** and press ENTER.
3. When the IP Configuration window appears, use the onscreen controls to inspect the IP addresses. If the computer is dynamically configured, you can release and renew the DHCP leases.

If you want a more detailed report, click the More info button.

arp

Earlier in this chapter, you learned how the Address Resolution Protocol (ARP) translates IP addresses into physical MAC addresses. You can use a command-line utility called **arp** to inspect the current contents of your computer's ARP table, which contains the MAC addresses of computers with which you have communicated recently. The arp utility can also delete entries or add permanent entries into the ARP table. To use the arp utility, follow these steps:

1. Get a command prompt onscreen and type: **arp -a**
2. Press ENTER to execute the command, which will list the current contents of your computer's ARP table.
3. The arp command can also query the ARP table of a specific IP address. You can use this feature, for example, to see the ARP table in your default gateway. Execute the following command, replacing *xxx.xxx.xxx.xxx* by your router's IP address:

   ```
   arp  -a   xxx.xxx.xxx.xxx
   ```

4. To delete an entry from the ARP table, you run the following command, replacing *xxx.xxx.xxx.xxx* with the IP address of the entry you want removed:

   ```
   arp  -d   xxx.xxx.xxx.xxx
   ```

5. To find out about other arp command options, execute the arp command by itself, without typing any parameters. The arp utility will reply with a help screen describing all the options.

Network Analyzers

A **network analyzer** is a tool that enables a network administrator to capture and analyze packets crossing a network. Network analyzers come in handy when you need to:

- **Test connections** By sending test packets, a network analyzer enables you to troubleshoot network connections and identify faulty cables or malfunctioning network devices.
- **Send alerts** A network analyzer can be configured to send an alert to an operator when something fails on the network.
- **Sniff packets for analysis** Network analyzers are sometimes called *packet sniffers*, because they can grab packets and save them for later analysis. Chapter 13 teaches techniques you can use to keep sniffers from deciphering sensitive information, such as passwords or credit card information, that may be crossing the Internet.
- **Generate reports** Network analyzers can keep statistics and generate reports that can help identify peak usage patterns or bottlenecks on different segments of the network.

Some of the leading network analyzer products include Microsoft's Network Monitor, Network Associates' Sniffer product line, and Agilent's family of network analyzers. For more on network analyzers, search Google or Yahoo for *network analyzer*.

Try This!

Create the World's Smallest Network

You can learn a lot about networking by cabling two PCs together and exploring how the TCI/IP protocols and diagnostic tools function on those PCs. Because two nodes are the minimum number needed to create a network, you will by definition create the world's smallest network. To complete this exercise, you need two PCs with Ethernet jacks. To cable them together, you need a crossover cable. You can either buy a crossover cable from an electronics store such as Radio Shack, or you can make your own cable following the Try This! instructions for creating 10/100baseT cables in the previous chapter. Once you have the two PCs and the crossover cable, you can create and test the world's smallest network by following these steps:

1. On each of the two PCs, click the Windows Start button and choose Control Panel. When the Control Panel window appears, double-click Network Connections.
2. When the Network Connections window appears, double-click Local Area Connection to bring up the Status window, then click the Properties button to bring up the Local Area Connection Properties window, which displays in a listbox the protocol bindings on each computer's NIC.
3. Scroll the listbox down to reveal the Internet Protocol (TCP/IP), click it once to select it, then click the Properties button to bring up the TCP/IP Properties window.
4. On a piece of paper, write down the current settings that you see in the TCP/IP Properties window. You make note of the current settings so you can reset them when you are finished with this exercise.
5. Click to select the option entitled Use the following IP address. Fill out the IP addresses on the two computers as follows, which assigns each computer an IP address within the ICANN address range reserved for private networks:

   Computer A: IP address 192.168.0.1, subnet mask 255.255.255.0
   Computer B: IP address 192.168.0.2, subnet mask 255.255.255.0
6. Click OK to close the windows.
7. Use the crossover cable to connect the two PCs. Now you have the world's smallest network. In the rest of these instructions, you run tests to make sure you have the network configured properly.
8. On each computer, click the Windows Start button and choose Programs | Accessories | Command Prompt. At the command prompt, type the following command:

   Computer A: ping 192.168.0.1
   Computer B: ping 192.168.0.2
9. Press ENTER, and see how the computers respond. If each ping succeeds in obtaining an echo from the other computer, you have succeeded in creating the world's smallest network. Congratulations!
10. To return the computers to their previous configurations, repeat steps 1, 2, and 3. Then repeat step 5 using the settings you wrote down in step 4.

Internetworking Servers

Now that you know how to configure a PC to get on a network as a TCP/IP client, it is time to learn how Internetworking servers run utility services that provide basic infrastructure and enhance the network by making it more efficient. We begin by considering the critical role that DNS servers play in the infrastructure of the Internet.

DNS Servers

The Domain Name System (DNS) was invented because people prefer to go to Internet sites by domain names, such as www.loc.gov, instead of numeric IP addresses, such as 140.147.249.7. Before a client workstation can send a request that is addressed to a site's domain name, the Application Layer must convert the domain name to its corresponding IP address. The servers that perform this conversion are DNS Servers. The process they perform is known as *domain name resolution*. To resolve a Domain Name is to convert it into the corresponding IP address.

Because DNS name resolution is such an integral part of the Internet Architecture, it is important that a client be configured to use a DNS server that runs fast. Large organizations, for example, may dedicate one or more servers to performing the task of domain name resolution. As you noticed previously in this chapter in Figure 11-6, the TCP/IP configuration settings let you specify an alternate DNS server that can resolve the names if the primary DNS server goes down or is not responding.

DNS Name Space Hierarchy

Figure 11-7 shows how the DNS system is powered by a hierarchically distributed database called the *name space*, which is organized according to three levels: (1) the root level, (2) the top level, and (3) the second level. Each level contains DNS servers that are in charge of keeping track of the domains in the next lower level.

At the highest level of the hierarchy are the **root servers**, which keep track of the top-level domains, such as .com, .net, .edu, .gov, .org, and .us. The top-level DNS servers keep track of the second-level domains, such as mcgraw-hill.com, loc.gov, w3.org, and ny.us. The second-level DNS servers keep track of the names assigned by the organizations that own the domains. Examples of host names assigned in second-level domains are www.mcgraw-hill.com and mail.mcgraw-hill.com. Large organizations can further subdivide second-level domains into subdomains, such as investor.mcgraw-hill.com.

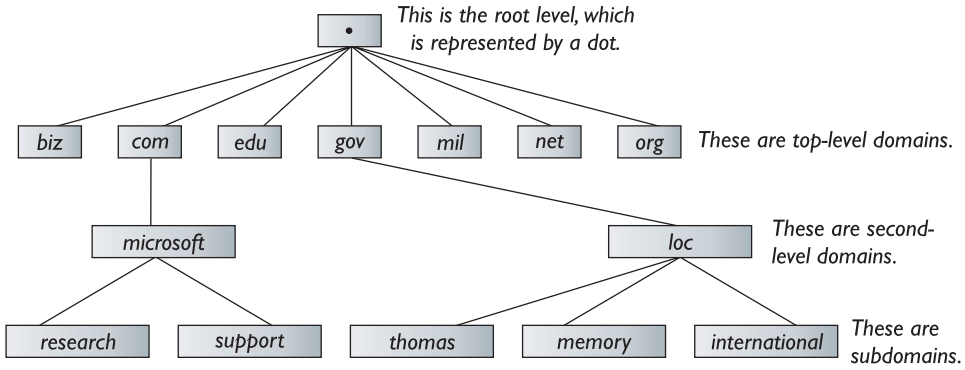

FIGURE 11-7 *The DNS name space has three levels: (1) the root level, (2) the top level, and (3) the second level. At each level, DNS servers keep track of the names assigned in that level. Root level servers, for example, keep track of the top-level domains. In this illustration, notice how the root level domain has a dot (.) for its designation. In a fully qualified domain name, you can include this dot, although in practice it is seldom used. Try going to thomas.loc.gov., however, to see how the Internet accepts the root dot as part of the name.* ■

DNS Name Servers and Resolvers

The DNS service has two components that work together to streamline the process of resolving names. These components are the (1) name server and (2) name resolver.

- **Name server** When a client workstation requests to resolve a domain name, the request first goes to a **name server**, which is in charge of responding with the IP address that corresponds to the domain name in the request. To speed the process of looking up these addresses, the name server maintains a cache containing the IP addresses for domain names that the server has already looked up. When a request comes in, the name server looks first in the cache and responds instantly if the requested address is there.
- **Name resolver** If the name server does not already know the IP address of the domain name in the request, the name server calls upon a **name resolver**, which goes out on the Internet and consults the necessary name servers in the DNS hierarchy. If those name servers do not know, they consult name resolvers further up the hierarchy. The cache that is kept on the name servers, however, minimizes the amount of actual name resolving. As a result, most names on the Internet can be resolved within a second or two.

Three Types of DNS Servers

Something all DNS servers have in common is that if they do not already know the IP address for a requested domain name, they become a client and request the address from the nearest name resolver. Thus, the DNS system follows the client-server model. Included in this model are the following DNS server types:

- **Root server** At the top of the DNS hierarchy, root servers can resolve all of the top-level domains on the Internet. If none of the name servers closer to the requesting workstation can resolve the name, one of the DNS root servers assumes responsibility for finding the name within the requested domain's name space.
- **Primary server** The first DNS server in a domain is called its **primary server**. As the domain's naming authority, the primary server maintains the master copy of the database containing the assigned names and IP addresses that are in the primary server's domain.
- **Secondary server** A domain can have one or more **secondary servers**, which help share the name-serving load and provide backup in case the primary server goes down. The secondary server contains a copy of the database from the primary server. In large networks, the primary server can delegate authority for different parts of the database to multiple secondary servers that resolve names in the network's subdomains.

Common DNS Record Types

Many types of resource records are in the DNS database. If you are studying for the CIW Foundations exam, you should learn to recognize the most common resource record types, which Table 11-2 describes. Only DNS programmers and troubleshooters need to know the other types.

DNS Record Type	Purpose
Address (A)	Identifies the IP address of a domain name. This record occurs most often in the DNS database.
Canonical Name (CNAME)	An alias that lets a host with one name be accessed also by another name. A computer fulfilling the dual role of Web and news server, for example, could answer to requests addressed to both www.mydomain.com and news.mydomain.com.
Mail Exchanger (MX)	Identifies the IP address of a mail exchanger on a domain.
Name Server (NS)	Identifies the IP addresses of the primary and secondary name servers for a domain.
Start of Authority (SOA)	Identifies the IP address of the DNS server that is the primary authority for a domain.

TABLE 11-2 *Important DNS Resource Record Types*

Host Tables and Files

Back in the good old days before 1983 when DNS began evolving into an Internet standard, there was no hierarchically distributed database of domain names and IP addresses. Instead, one huge file called the **hosts table** contained the name and IP address of every named host on the Internet. The Stanford Research Institute's Network Information Center (SRI-NIC) managed and updated this file. Server administrators downloaded this file periodically to keep their local systems updated. As the Internet grew, however, the hosts table became too large for server administrators to download regularly.

Why, then, am I telling you this story? Because to this day, personal computers still have a hosts table, which you can use in the following situations:

1. You would like to refer to a computer on another network by a nickname instead of having to type its complete domain name.
2. You want to enhance performance on your local network, so you copy onto each PC on the network a hosts table identifying each node's name and IP address.
3. You have an isolated internal network with no DNS server, and you want each computer to have a fully qualified domain name.

To see the hosts table on your computer, follow these steps:

1. Get the Notepad running, pull down the Notepad's Files of Type menu, and set it to look for files of all types.
2. Use the Look in menu to navigate to the system32\drivers\etc subfolder of your Windows system folder. The complete path probably will be one of the following:

   c:\windows\system32\drivers\etc

   or

   c:\winnt\system32\drivers\etc

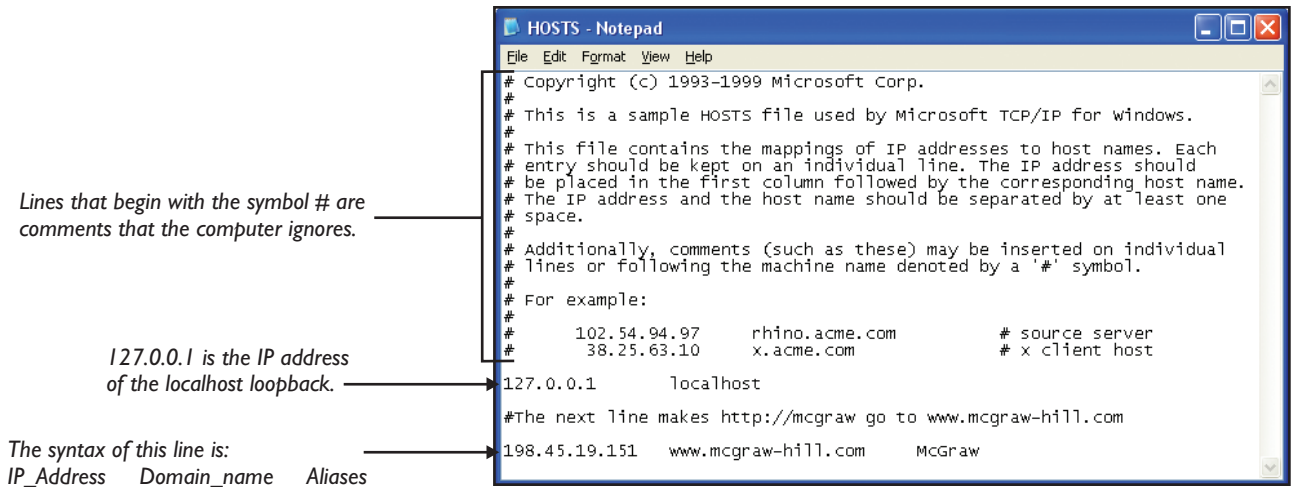

FIGURE 11-8 *The HOSTS file associates IP addresses with domain names. An important duty of the HOSTS file is to define the IP address of the computer's localhost loopback address. You can also use the HOSTS file to set up aliases. In this example, I created the alias McGraw and directed it to 198.45.19.151, thereby saving the time required to type www.mcgraw-hill.com.* ■

3. One of the files in this folder is named HOSTS. Use the Notepad to open the HOSTS file.
4. Figure 11-8 shows my HOSTS file. Notice that I created an alias that enables me to browse to www.mcgraw-hill.com simply by typing mcgraw.

Proxy Servers

In everyday life, people send a proxy when they cannot go somewhere themselves. On the Internet, proxy servers work somewhat like that. In computer networking, a **proxy server** is a computer that serves as an intermediary between client workstations and the external network.

At the University of Delaware, for example, students who live off campus configure their browsers to use a proxy server that enables them to access the same network privileges as on-campus students. The proxy server receives the off-campus requests and reissues them to the campus network. The campus network sees the request as coming from the proxy, to which the campus servers respond. Anyone from off campus who does not come in through the proxy is denied access to secured student resources.

Just as the Delaware students must configure their browsers to use the proxy, so must the clients on private networks that run behind a proxy. To configure Internet Explorer to use a proxy, you pull down the Tools menu, choose Internet Options, click the Connections tab, and click the LAN Settings button to reveal the proxy server settings. Other kinds of client applications, including FTP programs, mail readers, and telnet clients can similarly be configured to access the Internet through a proxy server. Here are some reasons why you may choose to use a proxy server:

- **Private IP address hiding** When the proxy server accesses the public Internet, it uses a different IP address than the private addresses assigned to the workstations running behind the proxy. Thus, the proxy hides the internal addresses on the private network from the public Internet.
- **Public IP address pooling** The proxy server can conserve resources by accessing the Internet from a smaller number of public IP addresses than the number of client workstations running behind the proxy.
- **Enhanced network security** The proxy server can block users on the public Internet from accessing network hosts running on the private network behind the proxy. Thus, the proxy serves as a kind of firewall.
- **Web content caching** The proxy server caches documents that it requests on behalf of users running behind the proxy. If other users running behind the proxy request the same documents, the proxy can serve the content instantly from the cache instead of having to retrieve it from the Internet.
- **Transaction filtering** You may want to block access on a private network to certain kinds of Internet resources. A proxy server can monitor the requests coming from the private network and block access to forbidden resources. This kind of monitoring and blocking is called **transaction filtering**. You can deny access, for example, to certain IP addresses, Web pages, URLs, host names, or computer names. You can even block individual users from accessing specific resources.
- **Transaction logging** A proxy server can record and timestamp the URLs, IP addresses, and external services accessed by clients running behind the proxy. By tracking the number of bytes received along with transmission times, the proxy server can keep vital statistics that network administrators use to monitor network performance.

Caching Servers

A **caching server** speeds access to resources by making a local copy of resources requested from the network so Web content and other kinds of documents and files can be served more quickly to subsequent users who request the same resources. Only if the date on the original document changes will the caching server download a fresh copy of the requested resource.

Caching servers can run either on a standalone computer that is dedicated to caching or as server components alongside other services on the same computer. Proxy servers, for example, normally include a caching server component. Computers that come preconfigured as caching servers are sometimes called *cache-in-a-box* because the server is ready to run as soon as you take it out of the box.

Mirrored Servers

A **mirrored server** is a computer whose data reads and writes are simultaneously executed on another computer that keeps an updated faithful copy of everything on the system from which the data are copied. The purpose of a mirrored server is to provide redundancy and fault tolerance in mission-critical operations in which server failure could cause serious problems.

Another use of mirroring is on a single server that has a redundant array of independent drives (RAID) controller. One of the RAID configurations known as *RAID level one* mirrors the data on multiple disk drives within a single computer. I use RAID level one, for example, to mirror the drive containing the operating system. If the operating system drive fails, the mirror drive automatically takes over, and the RAID controller sends an e-mail message alerting me that the failed drive needs maintenance. End users never notice the problem, because the server keeps operating.

Certificate Servers

There are two main categories of security concerns on the Internet. First, because packet sniffers can inspect the contents of data packets as they wind their way over the Internet, companies need a way to ensure that only intended recipients can read the data. Second, recipients need a way to determine the authenticity of the information to make sure it is coming from a trusted source and has not been modified along the way. A cracker, for example, could be masquerading as a legitimate vendor while sending you malicious code in disguise.

To shore up these security issues, the Internet uses **certificate servers** to issue digital certificates that network hosts use to digitally sign and encrypt messages using public-private key pairs. When you download a new version of the Flash player, for example, you get a message telling you that Macromedia digitally signed the software. The message includes a link to the certificate authority where you can verify that Macromedia is the source of the software. Then you can decide whether you want to proceed with installing the software.

Because of the threat of malicious code, you need to be wary of downloading software that is not digitally signed. Chapter 13 provides detailed coverage of Internet security best practices including certificates, digital signatures, and public-private key encryption.

Directory Servers

A **directory server** is a standalone computer or server component in charge of managing a database that keeps track of all the users, passwords, resources, printers, servers, e-mail addresses, phone numbers, and departmental contacts throughout an organization's network. When a user logs on to the network, the directory server provides access to those resources that the user has permission to use.

To make it possible for directory servers to communicate with each other and distribute the directory database over multiple networks, the International Standards Organization (ISO) in 1984 began work on creating

a **Directory Access Protocol (DAP)** called **X.500**. The X.500 database has a hierarchical design that allows different parts of the database to reside on different directory servers. A large multinational company, for example, can have a directory server on each continent keeping track of employee information and resources for the local country networks. Through DAP, these directory servers can exchange information, thereby creating a worldwide directory of the company's employees and resources. An advantage of this exchange is that each employee's information need be entered and maintained on just one server, thereby avoiding the problems that can be caused by storing employee data in multiple places that all must be updated if the information changes.

In practice, however, X.500 was overly complex. To streamline the process, the University of Michigan invented in 1993 a lighter version called the **Lightweight Directory Access Protocol (LDAP)** that can run over TCP/IP. The streamlined LDAP protocol became an instant success. Netscape adopted LDAP in 1997 to power the Netscape directory server. Bigfoot used LDAP to create its directory of people at www.bigfoot.com, and Yahoo used LDAP to power the people search at people.yahoo.com. Most significantly, Microsoft used LDAP to create Active Directory, which is a core server component that provides directory services on Windows servers.

RFC 1487 defines the LDAP standard. If you are studying for the CIW exam, you should understand three advantages of the X.500 and LDAP directory strategies:

- **Synchronization** The directory on one server can synchronize with the directory on another, thereby keeping the data current.
- **Replication** Part or all of the directory database on one server can copy itself onto another server. This advantage provides backup and fault tolerance in case one of the servers fails, and it speeds access to resources by reducing the number of hops needed to look up a given resource.
- **Scalability** Because network administrators can distribute the database over multiple servers, there is no limit to the size to which the directory can grow.

Catalog Servers

A **catalog server** uses robots called *spiders* that comb through a network's files and create an index of everything they find. When users need to find something on the network, they search this catalog. Because the search looks up the keywords in the prebuilt index, users quickly find what they seek.

Catalog servers can index all kinds of information, including Web pages, newsgroups, word-processed documents, PDF files, mail messages, images, movies, audio, and software applications. In addition to indexing the full text of written documents, the catalog indexes all of the file's properties. Thus, you can search for all of the documents written by a certain author or modified after a specific date. When users conduct a search, they are shown only resources to which they have access.

Transaction Servers

Transaction servers work behind the scenes in the business tier of the multi-tier e-commerce model to ensure that when a financial transaction occurs, all of the necessary databases get updated and related services receive the proper notifications. A transaction is a set of events that must be performed or rolled back simultaneously. For example, when someone transfers money from a savings account to a checking account, a transaction server makes sure that the savings account gets debited when the checking account is credited. A more complex transaction occurs when someone uses a credit card to purchase a product online; the online storefront's transaction server (1) finances the purchase by debiting the customer's account in the mercantile database at the bank that issued the card, (2) updates the inventory database, (3) arranges for shipping and updates the shipping database, and (4) may notify the manufacturing database if product inventory dips too low.

One of the most well-known transaction systems is IBM's Customer Information Control System (CICS). Originally designed as a legacy mainframe application, CICS now is available in a client-server version. Another popular example is Microsoft Transaction Server (MTS), which debuted in Windows NT Server. In Windows 2003 and later Microsoft server products, MTS is an integral part of component services.

Serving Internet Resources

After studying the Internetworking servers that power the network's infrastructure behind the scenes, now it is time to look more deeply into the well-known services that are the reason for the Internet's mass-market popularity. We begin with the Internet's most popular servers, which are Web servers.

Web Servers

Web servers are standalone computers or server components that respond to HTTP requests from browsers and other kinds of Internet clients, including media players and handheld devices. Because Web servers use the HTTP protocol, they are also called HTTP servers.

Home Page Default Filenames

When you visit a site without specifying a filename, the Web server responds by sending you the site's default page. Thus, the end user does not have to know the filename of the default page. Behind the scenes, however, the filename is important, because competing brands of Web servers have different default filename conventions. The default page's author needs to know the name to give the page to make it appear when a user visits the site without specifying a filename. Table 11-3 identifies the default filenames that may be used by various brands of Web servers. Consult your server administrator to find out your Web site's default filename convention.

Web Server Hit Logs

Every time a request hits a Web site, the site's Web server may log certain information about the hit. Depending on the brand of Web server, the specific information logged may vary, but all brands are capable of recording three general categories of information:

- **Client access data** For each hit, the log identifies the IP address of the client that issued the request. Remember that this address can be the one used by a proxy server making the request on behalf of a client. Therefore, this address does not necessarily identify the specific client that made the request.
- **Referrer data** The referrer data reveals the URL that the user typed or clicked to reach the site, the filename or command that may have been appended to this URL, and the HTTP method of the request (i.e., GET or POST).
- **Error data** If errors occur, they appear in status codes that the server administrator can study to improve site operations. Dropped connections, security access violations, and malformed URLs are some of the errors that can occur.

Default Filename	Web Server Brand
default.asp default.htm default.html	Internet Information Server (IIS) on Microsoft Windows servers
index.htm index.html	Apache on UNIX and Linux servers
home.htm home.html	Vendor neutral
main.htm main.html	Vendor neutral
welcome.htm welcome.html	Vendor neutral

TABLE 11-3 *Filenames Typically Used for a Web Site's Default Page*

The Windows IIS Web server keeps logs by default, meaning that logging is on unless the server administrator turns it off. Information recorded in IIS logs includes (1) the date and timestamp, (2) the destination IP address, (3) the request method (i.e., GET or POST), (4) the URL requested, (5) the port (normally 80, which is the Web's default port), (6) the IP address of the requestor, (7) the name and version of the browser or other client issuing the request, and (8) status codes related to the success or failure of the server's handling of the request.

Web Root Folders

On a Web server, the **Web root** is the physical folder that represents the beginning of the server's Web space. On Windows IIS servers, the Web root is typically located at c:\inetpub\wwwroot. When someone browses to an IIS server's Web address without specifying a filename, the server responds by sending the file c:\inetpub\wwwroot\default.htm, default.html, or default.asp. Sites hosted on that server typically reside in subfolders off the root. If Santa Claus has a site on this server, for example, it could be located in the folder c:\inetpub\wwwroot\santa.

Depending on the nature of the site, however, it may not be possible to store all of the site's files on a single drive. Even if all the files could fit on one drive, it may be more efficient to put some of the site's folders on other drives. To accomplish this, the server administrator can create virtual directories and aliases.

Virtual Directories and Aliases

A virtual directory is the name of a path to a Web folder that may reside anywhere in the host computer's file space, even in a physical location outside the scope of the server's root Web space. A virtual directory's physical location is often on a different drive from that of the server's root Web space. Locating a virtual folder on a different drive enables the server to distribute the file-serving workload among multiple drives. Situations arise, however, in which it helps to create virtual directories within the Web root's directory tree. Imagine a situation in which a site's Web address contains many folder and subfolder names, such as:

http://many.sites.com/external/northern/nonprofit/toymakers/santa

Creating a virtual directory named santa for the path c:\inetpub\wwwroot\external\northern\nonprofit\toymakers\santa enables users to access the same site via the simpler Web address:

http://many.sites.com/santa

File Permissions and Access Control

Many public Web sites allow anonymous access to all of the Web pages at the site. Other sites restrict access to authorized users. A restricted site can decide whether a user should be permitted to access it in two ways. The first way is via file permissions, and the second way is through authenticated access control.

- **File permissions** Most Web servers, including the most popular Windows, UNIX, and Linux Web servers, observe the operating system's file permission settings that the server administrator can configure for any individual file or folder at the site. Typical permissions include the ability to read a file, write (create or delete a file), execute (run a program), modify (edit a file), or deny access. Figure 11-9 shows that the Windows operating system enables you to set different file permissions depending on the role of the person who is accessing the site. Using this role-based model, for example, you can permit network managers to read, write, execute, and modify, while granting ordinary employees permission only to read and execute.
- **Authenticated user access control** To authenticate means to have a user log on by entering a user name and password, which the site looks up in a database to find out whether or not the user should be granted access. The database of users can be application specific, or the site can use the operating system's database of assigned user names and passwords. You have no doubt visited Web portals that invite you to register by choosing your own user name and password. Web applications that power such portals maintain a separate database of user names and passwords. You learn how to create this kind of user database in the next chapter.

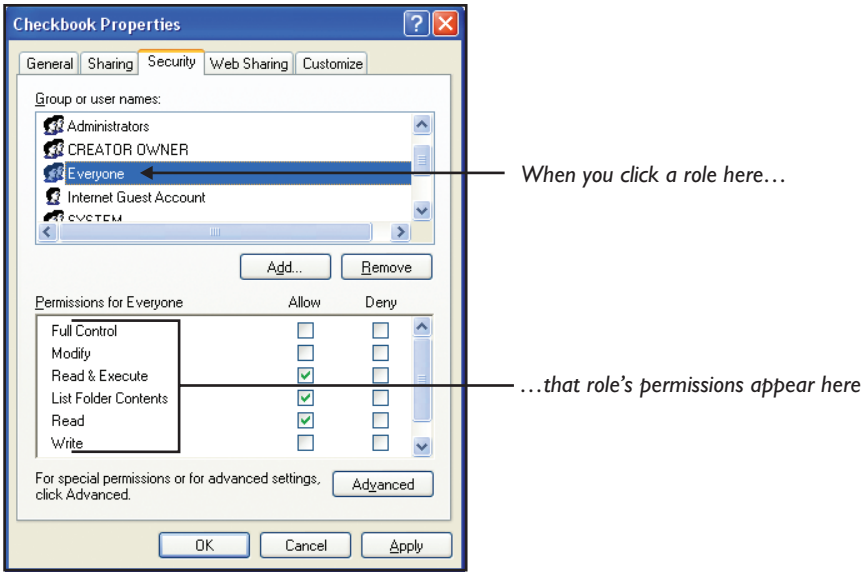

FIGURE 11-9 *In the Windows operating system, the Security tab of a file or folder's Properties window enables the server administrator to assign role-based permissions that provide different levels of access depending on the category of users visiting the site. Clicking a role in the top window displays that role's permissions in the bottom window.* ■

Web Gateways

When lay people think of a Web site, they conjure a collection of Web pages that the user can bring onscreen by browsing to the desired page. Certain sites called *gateways*, however, consist of no pages whatsoever. Instead of serving traditional HTML pages, a Web gateway runs a computer application consisting of one or more programs, or scripts, which generate the HTML response that the end user sees onscreen. Because the browser displays this HTML as though it were an ordinary Web page, many users are unaware that they are communicating with a Web application as opposed to viewing a traditional Web page.

To provide a standard way for Web gateways to communicate with browsers and other kinds of clients, the National Center for Supercomputing Applications (NCSA) created the **Common Gateway Interface (CGI)** protocol. The CGI protocol defines the manner in which forms data, cookies, and other kinds of information in a Web request get submitted to the program or script that processes and responds to the request. The programs that respond to CGI requests are often referred to as **CGI scripts**. CGI is language neutral, meaning that CGI scripts can be authored in any computer language. Perl and Python are different brands of CGI-scripting languages. I write my CGI programs in C# and Java.

CGI is an open protocol. Any application that can open an Internet socket can implement the CGI interface. The latest version of Flash, for example, supports CGI. This enables the Flash author to create shows that can display Web forms and interact with users in a browser-like manner. Thus, the Flash author can use CGI to create a custom Web client that can replace the browser in situations requiring a custom user interface.

Mail Servers

True to its name, the Simple Mail Transfer Protocol (SMTP) defines the manner in which e-mail gets sent over the Internet. As you learned in Chapter 3, either clients can use POP3 to deliver the mail post-office style by downloading the messages to the user's PC, or you can use IMAP to leave the mail on the server, which enables you to read the mail from different client workstations. I keep an IMAP folder, for example, for each of my student advisees. Whenever I need to look up something about a student, such as during an office meeting, over coffee in a café, or while traveling, I can easily consult the IMAP folder from my desktop computer or wireless PDA. Thus, I am a huge fan of IMAP.

It is now time for you to learn about some of the more subtle distinctions related to the manner in which mail gets formatted for transmission over the Internet. We begin with MIME.

MIME

RFCs 2045 through 2049 define **Multipurpose Internet Mail Extensions (MIME)**, which are a set of standards that specify the formatting of Internet message bodies, the media types of Internet files and message bodies, and the method for attaching files that do not consist of plain ASCII text. Although the "ME" in "MIME" stands for Mail Extensions, MIME is used for a lot more than mail. Each file that gets transmitted over the Web, for example, has a MIME header that identifies the file's media type. Browsers, e-mail programs, and other kinds of Internet clients rely on MIME headers when deciding how to handle files. Table 11-4 identifies the most common MIME types.

Uuencoding

Prior to the invention of MIME and its adoption by modern mail clients, users had to go through a manual encoding process to send binary files in a mail message. This process was UNIX to UNIX encoding (**uuencoding**). It worked so well that uuencoding became popular under Windows as well as the UNIX operating system.

Uuencoding works by translating the 8-bit character stream of a binary (i.e., non-ASCII) file attachment into a stream of 7-bit characters. Because the resulting 7-bit file stream is all ASCII characters, you can simply paste a uuencoded file into the mail message's text body. When an end user receives a message containing a uuencoded file, the user must copy the uuencoded portion into a separate file and run the uudecode program to decode it. The result of the uudecode process is an exact copy of the binary file as it existed prior to being uuencoded.

For most practical purposes, MIME happily obsoletes manual uuencoding, although you could encounter some uuencoded files in legacy situations.

BinHex Encoding

BinHex encoding is to the Macintosh what uuencoding is to UNIX. To decode a BinHex file received from a Macintosh, you need to use the appropriate decoder. There are different versions of BinHex encoding, so you need to make sure you use the proper decoder. The first line of a BinHex file identifies its version.

Filename Extension	MIME Type	Media Type
.css	text/css	Cascading style sheet
.doc	application/msword	Microsoft Word document
.exe	application/octet-stream	Executable file
.gif	image/gif	GIF image
.html	text/html	Web page
.jpg	image/jpeg	JPEG image
.midi	audio/x-midi	MIDI music file
.mp3	audio/mpeg	MP3 music file
.mpeg	video/mpeg	MPEG movie
.pdf	application/pdf	Adobe PDF
.png	image/png	PNG image
.ppt	application/vnd.ms-powerpoint	PowerPoint presentation
.qt	video/quicktime	QuickTime video
.rtf	application/rtf	Rich text document
.sit	application/x-stuffit	Stuffit archive
.swf	application/x-shockwave-flash	Flash animation
.txt	text/plain	ASCII text
.wav	audio/wav, audio/x-wav	Waveform audio
.xls	application/vnd.ms-excel	Spreadsheet
.xml	application/xml	XML file
.zip	application/zip application/x-compressed-zip	Zipped archive

TABLE 11-4 *The Most Common MIME Types* ■

Mailing List Servers

Chapter 3 provided a detailed tutorial on joining and using a listserv. Behind the scenes, mailing list servers are powered by SMTP. When a member of the list e-mails a message to the list, the mailing list server e-mails the message via SMTP to everyone on the list.

The listserv instructions in Chapter 3 are based on a specific brand of mailing list server that is called *listserv*, which is distributed by L-Soft at www.lsoft.com. On the Windows platform, Microsoft Exchange supports the creation and management of mailing lists. Another popular brand of mailing list manager is the freely distributed UNIX-based Majordomo list server at www.greatcircle.com. GreatCircle hosts Majordomo listservs for the purpose of supporting and developing Majordomo.

Streaming Media Servers

Microsoft, Apple, and Real Networks are the primary vendors of the streaming media services you studied in Chapter 3. Behind the scenes, the streaming media servers use UDP, as opposed to TCP, to transfer their packets. UDP transmits more continuously by putting each output from the Application Layer directly into a packet, thereby avoiding the overhead of TCP session management. The fact that no provision has been made for the resending of lost packets does not cause as much of a problem for video as it would for a mission-critical financial transaction. In the case of a lost video packet, the retransmission arrives out of sequence, making it too late to view anyway. To paraphrase the old Broadway saying, "the stream must go on."

note *If you are watching a stream and the video appears choppy due to lost packets, use your media player's settings to configure the stream to play at a lower bit rate. Figure 11-10 shows the typical bit rate settings.*

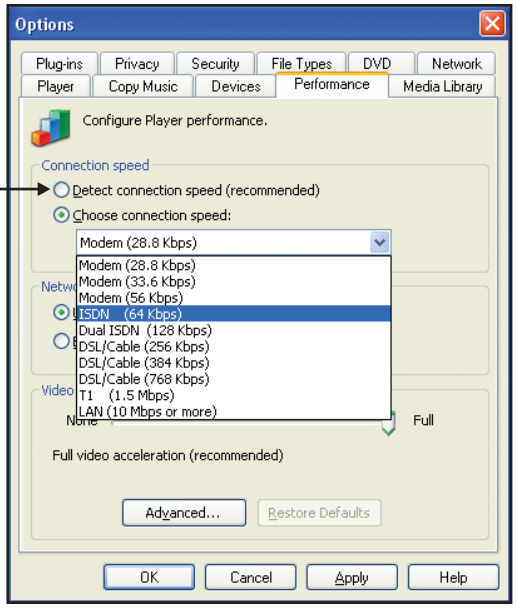

Microsoft recommends that you let the media player auto detect the connection speed.

FIGURE 11-10 *If a media stream plays choppily due to dropped packets, use your media player's options to select a slower connection speed. Shown here are the Windows Media Player connection speeds.* ■

FTP Servers

Chapter 2 touted the many advantages of the File Transfer Protocol (FTP). You need to be careful setting up an FTP server on your network for two reasons. First, many brands of FTP servers transmit user names and passwords in clear text, making them prone to detection by packet sniffing. One way to avoid this problem is to let users log on anonymously, but this solution is not always practical. Chapter 13 teaches you how to set up an FTP server that encrypts user names, passwords, and file transfers.

Second, you may unknowingly expose files that you would not want the public to access. When you install an FTP server on your computer, you must carefully configure it so as to expose only the files and folders to which you want to provide remote access.

If FTP is prone to all these problems, you may justifiably ask, why should you install an FTP server when you could just e-mail the files instead? There are three reasons. First, most mail servers impose a file attachment limit of 2 or 3 megabytes per message. If a file is larger than the e-mail server's limit, you cannot e-mail it. Second, FTP servers let users download a file at their convenience, instead of having to ask you to e-mail it. Third, FTP is faster than e-mail and makes more efficient use of Internet bandwidth. These are the reasons why Chapter 13 teaches you how to use FTP securely.

News Servers

Chapter 3 taught you how to participate in USENET newsgroups. Behind the scenes, USENET newsgroup servers run on port 119, following the Network News Transport Protocol (NNTP) as defined by RFC 1036.

As an alternative to reading news with NNTP newsgroup clients, Web interfaces let you read USENET newsgroups via HTTP. For example, the newsgroup gateways at groups.google.com and www.mailgate.org enable you to participate in USENET newsgroups via your browser using the Web's HTTP protocol.

Popular Server Products

Two leading brands of server products are UNIX/Linux and Microsoft Windows Server. Both of these product families can host all of the Internet services and run all of the Internetworking servers covered in this chapter, either on stand-alone computers dedicated to running one type of service on a large network, or as server components on computers that serve multiple functions on smaller networks.

UNIX and Linux

UNIX and Linux listen for Internet requests through a program called **inetd**, which stands for **Internet daemon**. In Greek mythology, daemons are supernatural guardian spirits that serve as intermediaries between gods and humans. On UNIX and Linux systems, the term *daemon* refers to any process that runs in the background, waiting to respond to certain kinds of requests. The role of the Internet daemon is to dispatch requests coming from the Internet to the server components that will handle them. The network administrator uses a configuration file named *inetd.conf* to define the different kinds of Internet requests to which the server will respond, and the name of the service daemon that will handle them. If a mail message comes in, for example, inetd routes it to smtpd, which is the SMTP mail daemon. Similarly, inetd routes telnet requests to telnetd, which is the telnet daemon. FTP requests go to ftpd, the FTP daemon.

To permit a UNIX/Linux system to appear in the Windows Network Neighborhood, a UNIX/Linux file server called Samba enables drag-and-drop file transfers between Windows and UNIX/Linux systems. Samba further enables UNIX/Linux users to send printed output to printers attached to a Windows system. When the Internet daemon sees a NetBIOS packet, inetd routes it to smbd, which stands for Samba daemon. There is also a NetBIOS name server daemon called nmbd. Figure 11-11 shows a sample *inetd.conf* file configured to participate in a NetBIOS network as well as respond to mail, telnet, FTP, and HTTP requests.

The Internet daemon runs with root privileges, which are the UNIX/Linux equivalent to Administrator status on a Windows server. Because inetd has root privileges, it is important to restrict access to the *inetd.conf* file so that only the network or system administrators can modify it. If a cracker should gain access to the *inetd.conf* file, system security is seriously compromised.

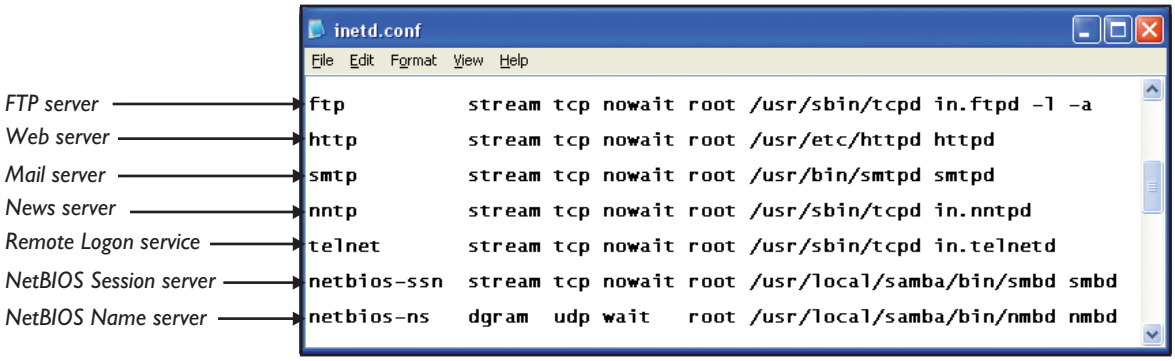

FIGURE 11-11 *The network administrator uses the inetd.conf file to identify the kinds of requests to which the Internet daemon responds, and the process handler to which the daemon passes those requests. In addition to the popular Internet services of FTP, HTTP, SMTP, NNTP, and telnet, this example includes the Samba support for NetBIOS, which enables the UNIX/Linux server to appear in the Network Neighborhood on a Windows client workstation.* ■

Microsoft Windows Server

When a Microsoft Windows server boots up, the computer goes through a startup process that starts the Internet services. This preloading makes the services ready to respond to various kinds of Internet requests as soon as they come in.

Depending on the size and complexity of the network, the system administrator decides which Internet services to run on each server. Figure 11-12 shows the Windows Components Wizard, with which the system adminis-

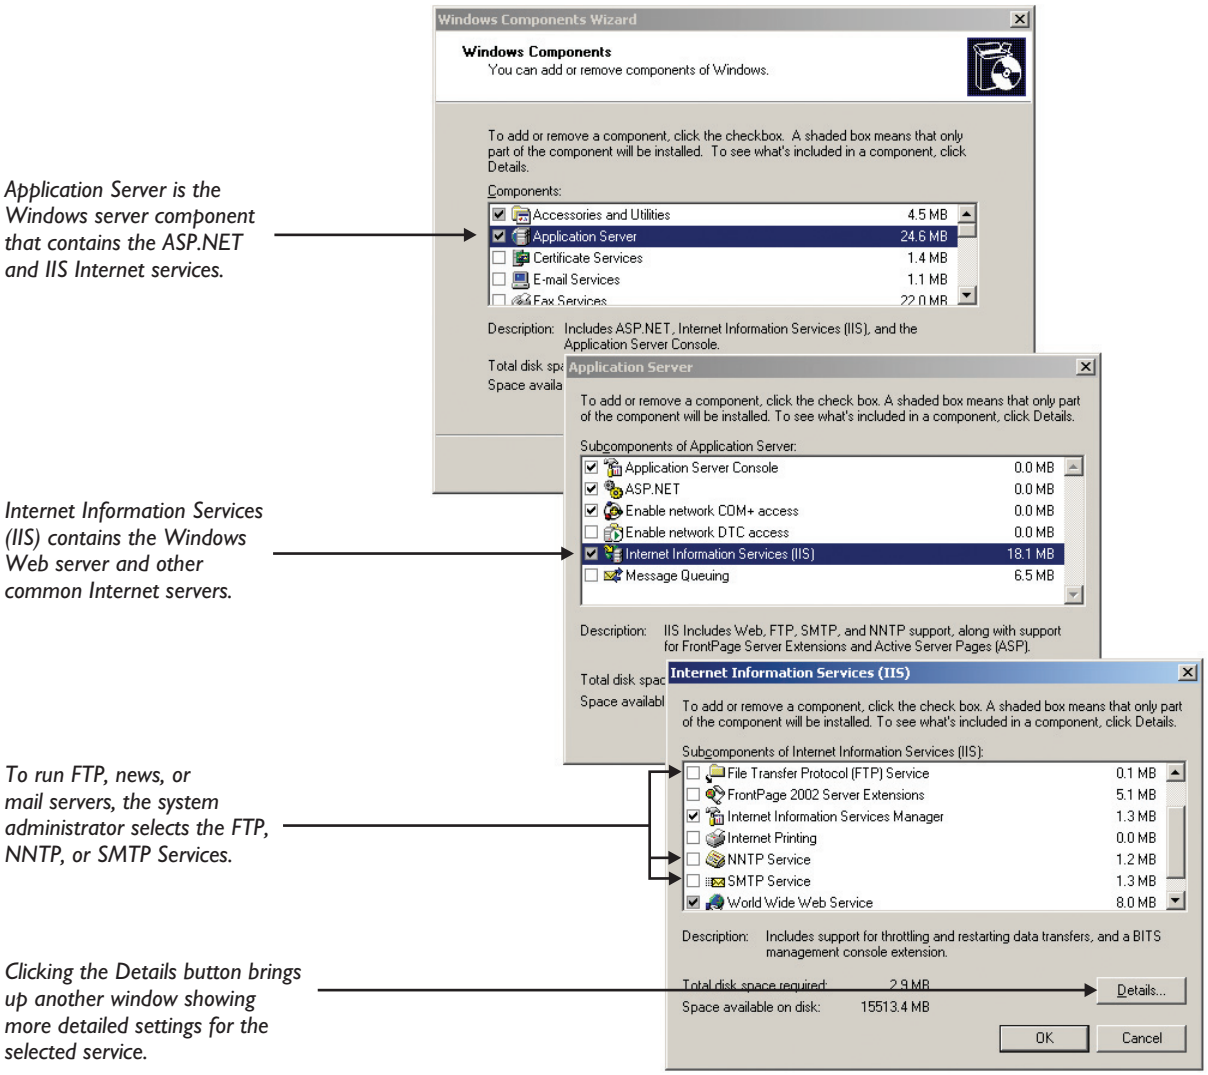

FIGURE 11-12 *The system administrator uses the Windows Components Wizard to select the Internet components that the server will run. The configuration shown here is dedicated to serving Web pages and running ASP.NET Web applications.* ■

trator selects the services to run. Most versions of Windows have this wizard, although the Windows Server version has more options than the typical client computer. If you want to launch the Windows Components Wizard to see how it appears on your computer, click the Windows Start button, choose Control Panel, double-click Add or Remove Programs, and click Add/Remove Windows Components.

After selecting the Windows components that will run on a given server, the system administrator uses the Microsoft Management Console (MMC) to configure them. In Figure 11-13, the MMC displays the IIS Manager, with which the system administrator configures Internet services. Thus, the Windows operating system provides graphical property windows for settings that UNIX and Linux administrators configure manually in the *inetd.conf* file.

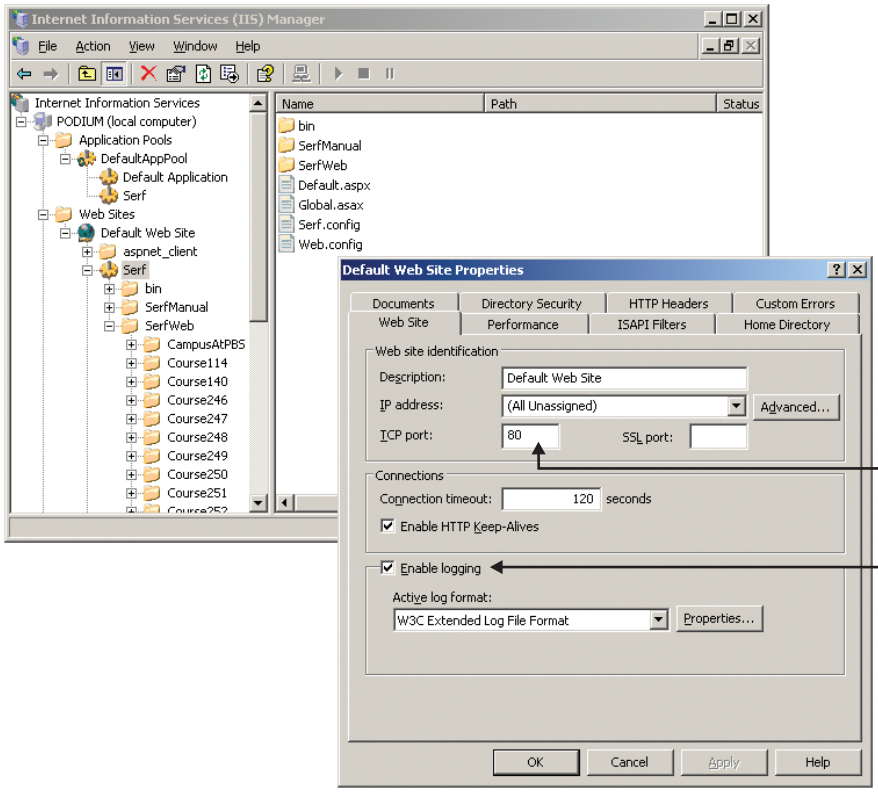

FIGURE 11-13 *On a Windows server, the system administrator uses the Microsoft Management Console (MMC) to configure Windows components. You right-click any component to pop out its quick menu, and then choose Properties to bring up the settings. In this example, the system administrator is reviewing the settings for the server's default Web site.* ■

Try This!

Inspect Your Computer's Services

The Windows operating system runs many services that listen for requests to be received on various Internet ports. Not all computers run all services. To find out what services are running on your computer, follow these steps:

1. Click the Windows Start button and choose Control Panel. When the Control Panel opens, double-click Administrative Tools.
2. When the Administrative Tools window appears, double-click Services. When the Services window appears, click the Extended tab near the bottom of the window.
3. Scroll the window up or down to see all the services running on your computer.
4. To find out what a service does, click to select it and read the service description at the left of the window. Another way to read the service description is to double-click the name of the service and see its service description in the properties window.
5. You should not make any changes to the services unless you know what you are doing. Many services depend on each other to perform interrelated tasks. If you stop one service, other services that depend on it will also stop working.
6. If you notice any suspicious processes running on your computer, particularly unsigned services that do not have names, it is possible that your computer may have been hacked. Chapter 13 discusses methods crackers use to install rogue services.

Chapter 11 Review

Chapter Summary

After reading this chapter and completing the step-by-step tutorials and Try This! exercises, you should understand the following facts about the Internet:

Understanding TCP/IP

- The Internet Architecture implements in four layers the seven separate layers of the OSI Reference Model.
- At the top of the Internet Architecture's protocol stack is the Application Layer, which encompasses OSI/RM Layers 7 and 6. It is at the Application Layer that some of the Internet's most well-known protocols reside, such as HTTP, FTP, TFTP, telnet, Gopher, SMTP, NNTP, DNS, SNMP, BOOTP, and DHCP.
- The Internet Architecture's Transport Layer encompasses OSI/RM Layers 5 and 4. It is the responsibility of the Transport Layer to divide into packets the data received from the Application Layer. Depending on the kind of session being serviced, the Transport Layer uses either TCP or UDP.
- The Internet Architecture's Internet Layer corresponds to OSI/RM Layer 3. It is the responsibility of the Internet Layer to take the packet from the Transport Layer, determine the best way to route it across the Internet, and transform it into an IP packet containing an IP header and trailer.
- The Internet Architecture's Network Access Layer corresponds to OSI/RM Layers 2 and 1, which transform the packets into a binary encoded stream of 0's and 1's for transmission over the physical network. Then the NIC transforms the 0's and 1's into the signals that get transmitted physically over the network.
- Routing is the process of determining the network path over which the packets will be sent. It is the responsibility of the Internet Protocol to determine this optimal path. Associated routing protocols include EGP, BGP, RIP, and OSPF.
- At the destination computer on the receiving end, the process of unpacking the message by processing and removing the headers added to the packets at each layer is called *demultiplexing*.
- To provide the Transport Layer with a fast way of determining which application should receive an incoming request, each packet contains a destination port number, which is a 16-bit number indicating the service that the packet is using.
- The Internet Engineering Task Force (IETF) defines the Internet Architecture protocols through a Request for Comments (RFC) process. The maturity states that an RFC moves through on its way to becoming a standard are called (1) Proposed, (2) Draft, and (3) Internet standard.

Internet Addressing

- The Internet Corporation for Assigned Names and Numbers (ICANN) is in charge of assigning IP addresses. Every IP address consists of two basic parts: the Network ID and the host ID. The Network ID always comes first, followed by the host ID. Depending on the size of the network, the Network ID occupies the first one, two, or three numbers in the IP address. The remainder of the address is the host ID.
- The five classes of Internet addresses are named A, B, C, D, and E.
- Extremely large organizations that have more than 16 million hosts receive class A Internet addresses. In a class A address, the first byte in the dotted quad IP address is the Network ID, and the other three bytes are the host ID.
- Medium to large organizations with up to 65,534 hosts get class B addresses, in which the first two bytes are the Network ID and the other two bytes are the host ID.
- Small organizations with up to 254 hosts get class C addresses, in which the first three bytes are the Network ID and the last byte is the host ID.

- Multicast groups receive class D addresses, which are set aside for multicasting. In a class D address, all four bytes are the Network ID. There is no host ID because everyone in the group receives the multicast.
- The Internet has reserved some addresses for future use. These are the class E addresses.
- The Network ID 127 is the loopback address, which is a diagnostic IP address reserved for testing purposes that redirects packets to the same computer that sent them.
- A broadcast is a message that gets sent to all of the hosts on a network. The IP broadcast address byte is 255, which sets all eight bits in the address byte to 1. Computers that do not have IP addresses typically send a 255.255.255.255 broadcast on startup to find a DHCP or BOOTP server that can respond with an IP address assignment.
- A subnet mask is a dotted quad number that enables the local network to determine whether any given IP address is internal or external to the local network.
- In order for a personal computer to communicate on a TCP/IP network, you must first configure the computer's TCP/IP settings. At a minimum, the computer must have an IP address and a subnet mask, which enable communications with hosts on the local network. To communicate with computers on a WAN, such as the Internet, the computer must also have a default gateway.

Configuring Networks for Optimum Performance

- The most basic network troubleshooting utility is the Packet Internet Groper (ping), which sends ICMP echo request packets to a destination IP address. When the destination returns the echo, the ping utility measures the response time and displays a message telling the duration of the round trip.
- The traceroute program is a networking utility that reports the path data followed as a packet winds its way over the network from the source to the destination computer. You use the traceroute utility when you need to isolate the source of a network connectivity problem.
- The program netstat, which stands for network statistics, is a utility that displays information about the connections that are open and the protocol processes that are currently running on a network host.
- The ipconfig program is a TCP/IP configuration utility that runs on computers with Windows NT/2000/XP/2003 or later operating systems. You use ipconfig whenever you want to inspect the current IP configuration. Furthermore, if the computer's IP settings are dynamically configured, you can use ipconfig to release, renew, or refresh the DHCP leases. winipcfg is an older version of ipconfig for Windows 95/98/Me.
- The arp program is a command-line utility that you can use to inspect the current contents of your computer's Address Resolution Protocol (ARP) table, which contains the MAC addresses of computers with which you have communicated recently. The arp utility can also delete entries or add permanent entries into the ARP table.
- A network analyzer is a tool that enables a network administrator to capture and analyze packets crossing a network. Network analyzers are sometimes called *packet sniffers*, because they can grab packets and save them for later analysis.

Internetworking Servers

- The DNS system is powered by a hierarchically distributed database called the *name space*, which is organized according to three levels: (1) the root level, (2) the top level, and (3) the second level. Each level contains DNS servers that are in charge of keeping track of the domains in the next lower level.
- When a client workstation requests to resolve a domain name, the request first goes to a name server, which is in charge of responding with the IP address that corresponds to the domain name in the request. If the name server does not already know the IP address of the domain name in the request, the name server calls upon a name resolver that goes out on the Internet and consults the necessary name servers in the DNS hierarchy.

- A proxy server is a computer that serves as an intermediary between client workstations and the external network.
- A caching server speeds access to resources by making a local copy of resources requested from the network so Web content and other kinds of documents and files can be served more quickly to subsequent users who request the same resources. Only if the date on the original document changes will the caching server download a fresh copy of the requested resource.
- A mirrored server is a computer whose data reads and writes are simultaneously executed on another computer that keeps an updated faithful copy of everything on the system from which the data are copied.
- Certificate servers issue digital certificates that network hosts use to digitally sign and encrypt messages using public-private key pairs.
- A directory server is a stand-alone computer or server component in charge of managing a database that keeps track of all the users, passwords, resources, printers, servers, e-mail addresses, phone numbers, and departmental contacts throughout an organization's network. When a user logs on to the network, the directory server provides access to those resources that the user has permission to use.
- A catalog server uses robots called *spiders* that comb through a network's files and create an index of everything that they find. When users want to find something on the network, they search this catalog. Because the search looks up the keywords in the prebuilt index, users quickly find what they seek.
- Transaction servers work behind the scenes in the business tier of the multi-tier e-commerce model to ensure that when a financial transaction occurs, all of the necessary databases get updated and related services receive the proper notifications.

Serveing Internet Resources

- Web servers are the stand-alone computers or server components that respond to HTTP requests from browsers and other kinds of Internet clients, including media players and handheld devices. Because Web servers use the HTTP protocol, they are also called HTTP servers.
- The Common Gateway Interface (CGI) protocol defines the manner in which forms data, cookies, and other kinds of information in a Web request get submitted to the program or script that will process and respond to the request. The programs that respond to CGI requests are often referred to as CGI scripts.
- The Simple Mail Transfer Protocol (SMTP) defines the manner in which e-mail gets sent over the Internet.
- The Multipurpose Internet Mail Extensions (MIME) are a set of standards that specify the formatting of Internet message bodies, the media types of Internet files and message bodies, and the method for attaching files that do not consist of plain ASCII text.
- Streaming media servers use UDP, as opposed to TCP, to transfer their packets. UDP transmits more continuously by putting each output from the Application Layer directly into a packet, thereby avoiding the overhead of TCP session management.
- USENET newsgroup servers run on port 119, following the Network News Transport Protocol (NNTP).
- UNIX and Linux listen for Internet requests through a program called *inetd*, which stands for Internet daemon. The UNIX/Linux system administrator uses the *inetd.conf* file to identify the kinds of requests to which the Internet daemon will respond, and the process handler to which the daemon will pass those requests.
- Windows system administrators use the Microsoft Management Console (MMC) to configure via graphical property windows the settings that UNIX/Linux system administrators configure manually in the *inetd.conf* file.

Key Terms

Address Resolution Protocol (ARP) *(6)*
Application Layer *(4)*
arp *(18)*
autonomous system (AS) *(9)*
BinHex *(33)*
Bootstrap Protocol (BOOTP) *(5)*
Border Gateway Protocol (BGP) *(9)*
caching server *(25)*
Canonical Name (CNAME) *(23)*
catalog server *(27)*
certificate server *(26)*
CGI script *(31)*
Common Gateway Interface (CGI) *(31)*
default gateway *(8)*
demultiplexing *(7)*
direct routing *(7)*
Directory Access Protocol (DAP) *(27)*
directory server *(26)*
Domain Name System (DNS) *(5)*
Dynamic and/or Private Ports *(10)*
Dynamic Host Configuration Protocol (DHCP) *(5)*
dynamically configured router *(8)*
External Gateway Protocol (EGP) *(9)*
File Transfer Protocol (FTP) *(4)*
Gopher *(4)*
hop *(8)*
hop count *(8)*
hosts table *(23)*
Hypertext Transfer Protocol (HTTP) *(4)*
indirect routing *(7)*
inetd *(35)*
Internet address classes *(10)*
Internet Architecture *(2)*
Internet Control Message Protocol (ICMP) *(6)*
Internet daemon *(35)*
Internet Engineering Task Force (IETF) *(3)*
Internet Group Management Protocol (IGMP) *(6)*
Internet Layer *(6)*
Internet Protocol (IP) *(6)*
IP broadcast address byte *(12)*
ipconfig *(18)*
Lightweight Directory Access Protocol (LDAP) *(27)*
limited broadcast *(12)*
loopback address *(11)*
mirrored server *(26)*
multicasting *(6)*
Multipurpose Internet Mail Extensions (MIME) *(32)*
name resolver *(22)*
name server *(22)*
Net-directed broadcast *(12)*
netstat *(17)*
Network Access Layer *(6)*
network analyzer *(19)*
Network News Transfer Protocol (NNTP) *(4)*
Open Shortest Path First (OSPF) *(9)*
ping *(15)*
port number *(9)*
primary server *(22)*
protocol binding *(14)*
proxy server *(24)*
registered ports *(10)*
Request for Comments (RFC) *(3)*
Reverse Address Resolution Protocol (RARP) *(6)*
RFC maturity state *(3)*
RFC number *(3)*
root servers *(21)*
routing *(7)*
Routing Information Protocol (RIP) *(9)*
routing information table *(8)*
secondary server *(22)*
Simple Mail Transfer Protocol (SMTP) *(4)*
Simple Network Management Protocol (SNMP) *(5)*
statically configured router *(8)*
subnet-directed broadcast *(12)*
subnet mask *(13)*
telnet *(4)*
traceroute *(16)*
transaction filtering *(25)*
transaction servers *(28)*
Transmission Control Protocol (TCP) *(5)*
Transport Layer *(5)*
Trivial File Transfer Protocol (TFTP) *(4)*
User Datagram Protocol (UDP) *(5)*
uuencoding *(32)*
Web root *(29)*
Web servers *(28)*
Well Known Ports *(10)*
Windows Internet Naming Service (WINS) *(14)*
winipcfg *(18)*
X.500 *(27)*

Key Terms Quiz

1. The _____ implements in four layers the seven separate layers of the OSI Reference Model.
2. _____ is the process of determining the network path over which the packets will be sent.
3. At the destination computer on the receiving end, the process of unpacking the message by processing and removing the headers added to the packets at each layer is called _____.
4. The Internet Engineering Task Force (IETF) defines the Internet Architecture protocols through a(n) _____ process.
5. The Network ID 127 is the _____, which is a diagnostic IP address reserved for testing purposes that redirects packets to the same computer that sent them.
6. A(n) _____ is a dotted quad number that enables the local network to determine whether any given IP address is internal or external to the local network.
7. I am the most basic network troubleshooting utility that sends ICMP echo request packets to a destination IP address. When the destination returns the echo, I measure the response time and display a message indicating the duration of the round trip. My name is _____.
8. A(n) _____ is a computer that serves as an intermediary between client workstations and the external network.
9. The _____ protocol defines the manner in which forms data, cookies, and other kinds of information in a Web request get submitted to the program or script that will process and respond to the request.
10. The _____ are a set of standards that specify the formatting of Internet message bodies, the media types of Internet files and message bodies, and the method for attaching files that do not consist of plain ASCII text.

Multiple-Choice Quiz

1. Which layer of the Internet Architecture has the responsibility to take the packet from the Transport Layer, figure out the best way to route it across the Internet, and transform it into an IP packet containing an IP header and trailer?
   a. Application Layer
   b. Transport Layer
   c. Internet Layer
   d. Network Access Layer
2. Which part of an IP address always comes first?
   a. Host ID
   b. Loopback address
   c. Network ID
   d. Subnet mask
3. Medium to large organizations with up to 65,534 hosts get what kind of IP addresses?
   a. Class A
   b. Class B
   c. Class C
   d. Class D
   e. Class E
4. Which kind of IP address does a multicasting group receive?
   a. Class A
   b. Class B
   c. Class C
   d. Class D
   e. Class E
5. The IP broadcast address byte is
   a. 0
   b. 127
   c. 128
   d. 255
6. Which network utility reports the path data followed as a packet winds its way over the network from the source to the destination computer?
   a. ipconfig
   b. netstat
   c. traceroute
   d. winipcfg

7. Which kind of Internetworking server speeds access to resources by making a local copy of resources requested from the network so Web content and other kinds of documents and files can be served more quickly to subsequent users who request the same resources?
   a. Caching
   b. Directory
   c. Mirrored
   d. Transaction
8. Which protocol defines the messaging rules that newsgroup servers follow on port 119?
   a. NNTP
   b. SMTP
   c. TCP
   d. UDP
9. Which protocol do streaming media servers use to avoid the overhead of transmission control session management?
   a. NNTP
   b. SMTP
   c. TCP
   d. UDP
10. What is the name of the configuration file in which UNIX and Linux system administrators manually edit the settings that identify the kinds of requests to which the Internet daemon will respond?
   a. autoexec.bat
   b. config.sys
   c. inetd.conf
   d. web.config

Essay Quiz

1. What role does the port number play in determining which service will handle an incoming message from the Internet?
2. Imagine that on port 80, a packet is coming in to its destination server. In chronological order, describe what happens at each of the four layers through which this incoming packet ascends during the demultiplexing process of the Internet Architecture. At the end of this ascent, identify the service that will ultimately handle this packet.
3. Explain the differences among class A, class B, and class C Internet addresses. Of the four bytes in the dotted quad IP addresses, which bytes in each class belong to the Network ID, and which bytes are the host ID? For what size organization is each class intended?
4. Compare the roles that name servers and name resolvers play in translating domain names into IP addresses in the Internet's DNS system.
5. Explain the role that the Internet daemon plays on UNIX and Linux operating systems.

Lab Projects

• Lab Project 11-1: Creating an Internet Address Allocation Plan

Imagine that a lot of turnover occurred in the IT division at your school or workplace, and there are shoddy records of which computers were assigned to different IP addresses. Your superior has asked you to review the situation and create a revised IP address allocation plan for your school or workplace. Use your word processor to write an essay in which you assess the current situation and present your plan for revamping the IP address allocation. In developing this plan, consider the following issues:

- **Domain name space** Does your organization have a class A, B, or C Internet address space? Does this class have a sufficient number of addresses to cover the number of hosts that need external IP addresses?

- **Private IP address range** What is the specified range of private IP addresses for your organization's Internet address class? If your organization has workstations with private IP addresses, are these addresses within the specified range for your address class?
- **Proxy servers** Consider whether your organization could use proxy servers to save costs and bolster security by hiding workstations on the internal network from hackers and crackers on the public Internet. If proxies are already in place, make sure they are configured to assign IP addresses within the correct range.
- **Static versus dynamic configuration** How many of the network hosts need to be statically configured? Could you save time and cost by letting the others use DHCP?
- **Current IP address outline** Create an outline that lists the current names and IP addresses of each network host within your organization. Use indentations to group the hosts under the subdomain or subnet to which they belong. For workstations that obtain temporary IP addresses via dynamic configuration, write DHCP in place of the IP address.
- **Planned IP address outline** Create an outline that shows how you propose to revise the current IP addressing scheme. Depending on how closely this revision follows the current addressing scheme, you may be able to add a "proposed" column to the outline you created in the previous step; otherwise, you will need to create a separate outline of your planned IP addressing scheme.

If your instructor asked you to hand in the IP address allocation plan, make sure you put your name at the top of the essay, then copy it onto a disk or follow the other instructions you may have been given for submitting this assignment.

• Lab Project 11-2: Preparing a Network Troubleshooting Guide

Networking problems can cause a lot of lost time and productivity when workers are unable to perform their jobs in a transaction-oriented workplace. Imagine that out of concern over these kinds of delays, your employer has asked you to prepare a network troubleshooting guide. The purpose of this guide is twofold. First, it will consist of a section intended for end users to diagnose and report network problems that nontechnical users cannot repair on their own. Second, the guide will contain a troubleshooting procedure for IT staff and more technically inclined users, who may be able to ease the burden on the IT staff by learning how to solve simple problems on their own. Use your word processor to write an essay in which you present your network troubleshooting guide. In formulating this guide, consider the following issues:

- **Configuration problems** Many network failures, especially with newly acquired workstations, result from incorrect PC network configuration. Include instructions for checking the PC configuration to ensure that the IP address, subnet mask, DNS server address, and gateway settings are correct.
- **Ping** The ping utility enables users to check for basic connectivity with other network devices. Include instructions for using ping to reach strategic network addresses, such as the nearest router and DNS server. If users can reach these devices, their network segment has connectivity.
- **Problem reporting** Provide users with a format for reporting problems they cannot solve on their own. Ask users to report whether the problem consistently occurs or is intermittent. If it is intermittent, ask whether the user can make the problem recur, and if so, ask what the user does to make the problem occur. Ask whether the problem occurs on just a single workstation or if other workstations have the same issue. If so, ask which specific workstations have the problem. Ask users to differentiate between total outages versus slow response times, and specify whether the slow periods are related to the time of day. Ask the user to describe any hardware or software changes that may have been made recently to the problematic workstation or local network segment.

- **Equipment for testing** To troubleshoot connectivity problems, you need spare network cables with connectors of the type used on your network. Swapping a questionable cable with one you know works is an effective way to determine whether you have a failed cable. Similarly, swapping a hub with a spare that you know works is a good way to determine whether a local hub has failed. More sophisticated problems can be diagnosed by plugging in a laptop running network troubleshooting software that can identify specific problems on the local network segment.
- **Other troubleshooting guides** Several network troubleshooting guides are available on the Internet. Use Google or Yahoo to search for the keywords "network troubleshooting guide." Perusing guidelines from the major networking vendors such as 3Com, Cabletron, and Cisco can provide ideas for techniques and strategies to include in your local troubleshooting guide.

If your instructor asked you to hand in the network troubleshooting guide, make sure you put your name at the top of the essay, then copy it onto a disk or follow the other instructions you may have been given for submitting this assignment.