

Application Layer Functionality and Protocols

Objectives

Upon completion of this chapter, you will be able to answer the following questions:

- How do the functions of the three upper OSI model layers provide network services to end-user applications?
- How do the TCP/IP application layer protocols provide the services specified by the upper layers of the OSI model?
- How do people use the application layer to communicate across the information network?
- What are the functions of well-known TCP/IP applications, such as the World Wide Web and e-mail, and their related services (HTTP, DNS, DHCP, STMP/POP, and Telnet)?
- What are the file-sharing processes that use peer-to-peer applications and the Gnutella protocol?
- How do protocols ensure that services running on one kind of device can send to and receive from many different network devices?
- How can you use network analysis tools to examine and explain how common user applications work?

Key Terms

This chapter uses the following key terms. You can find the definitions in the Glossary.

data page 67

source device page 67

Domain Name System (DNS) page 68

Request for Comments (RFC) page 68

syntax page 70

session page 71

client page 72

server page 72

daemon page 73

peer page 75

scheme page 76

IP address page 77

domain name page 77

network address page 78

resource record page 78

DNS resolver page 78

nslookup page 78

query page 78

cache page 79

authoritative page 81

plug-in page 82

HTTP page 82

distributed page 82

collaborative page 82

encryption page 82

Post Office Protocol (POP) page 83

Simple Mail Transfer Protocol (SMTP) page 83

Mail User Agent (MUA) page 83

spam page 85

gateway page 85

Dynamic Host Configuration Protocol (DHCP)
page 87

subnet mask page 87

broadcast page 88

Server Message Block (SMB) page 89

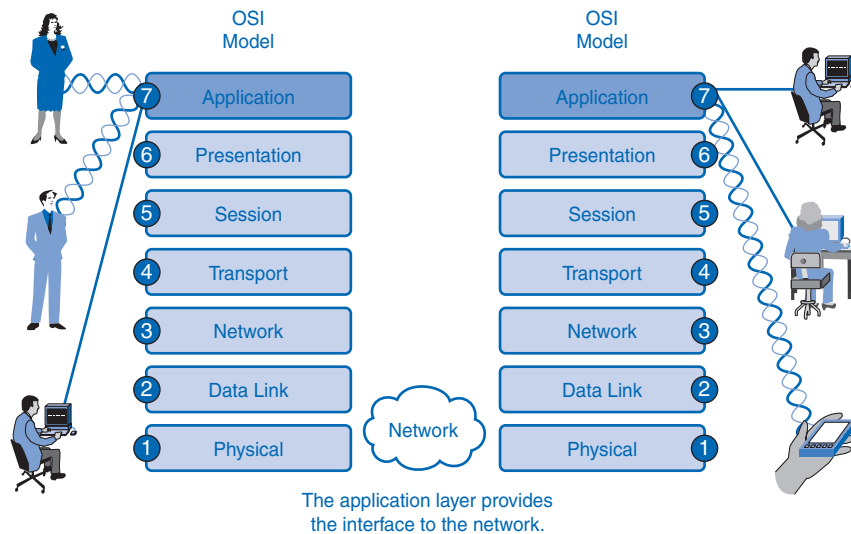
UNIX page 89

Interpret as Command (IAC) page 91

The world experiences the Internet through the use of the World Wide Web, e-mail, and file-sharing programs. These applications, as well as others, provide the human interface to the underlying network, allowing you to send and receive information with relative ease. Most of the applications are intuitive; they can be accessed and used without the need to know how they work. As you continue to study the world of networking, it becomes more important to know how an application is able to format, transmit, and interpret messages that are sent and received across the network.

Visualizing the mechanisms that enable communication across the network is made easier if you use the layered framework of the Open System Interconnection (OSI) model. Figure 3-1 depicts that framework. The OSI model is a seven-layer model, designed to help explain the flow of information from layer to layer.

Figure 3-1 Interfacing Human and Data Networks



This chapter focuses on the role of Layer 7, the application layer, and its components: applications, services, and protocols. You explore how these three elements make the robust communication across the information network possible.

Applications: The Interface Between the Networks

This section introduces two important concepts:

- **Application layer:** The application layer of the OSI model provides the first step of getting data onto the network.

- Application software:** Applications are the software programs used by people to communicate over the network. Examples of application software, including HTTP, FTP, e-mail, and others, are used to explain the differences between these two concepts.

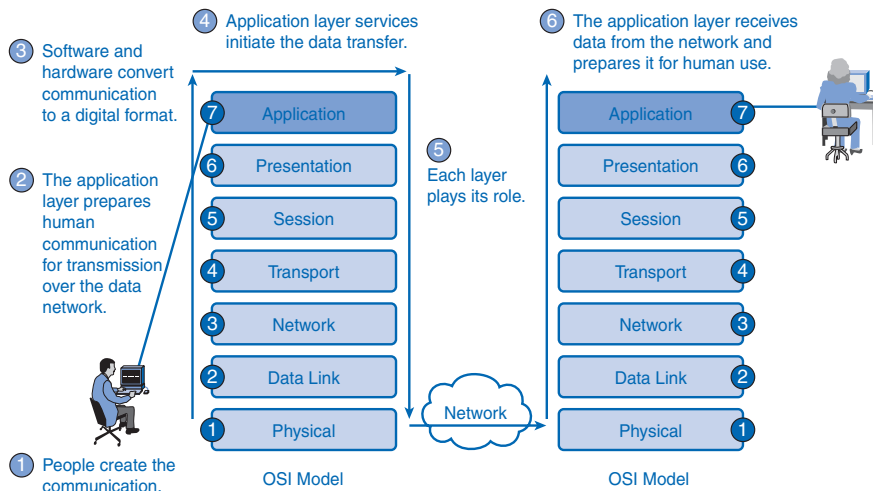
OSI and TCP/IP Model

The OSI reference model is a layered, abstract representation created as a guideline for network protocol design and instruction. The OSI model divides the networking process into seven logical layers, each of which has unique functionality and to which are assigned specific services and protocols.

In the OSI model, information is passed from one layer to the next, starting at the application layer on the transmitting host and proceeding down the hierarchy to the physical layer, then passing over the communications channel to the destination host, where the information proceeds back up the hierarchy, ending at the application layer. Figure 3-2 depicts the steps in this process. The following explains the six steps:

1. People create the communication.
2. The application layer prepares human communication for transmission over the data network.
3. Software and hardware convert communication to a digital format.
4. Application layer services initiate the data transfer.
5. Each layer plays its role. The OSI layers encapsulate data down the stack. Encapsulated data travels across the media to the destination. OSI layers at the destination unencapsulate the data up the stack.
6. The application layer receives data from the network and prepares it for human use.

Figure 3-2 OSI Encapsulation Process

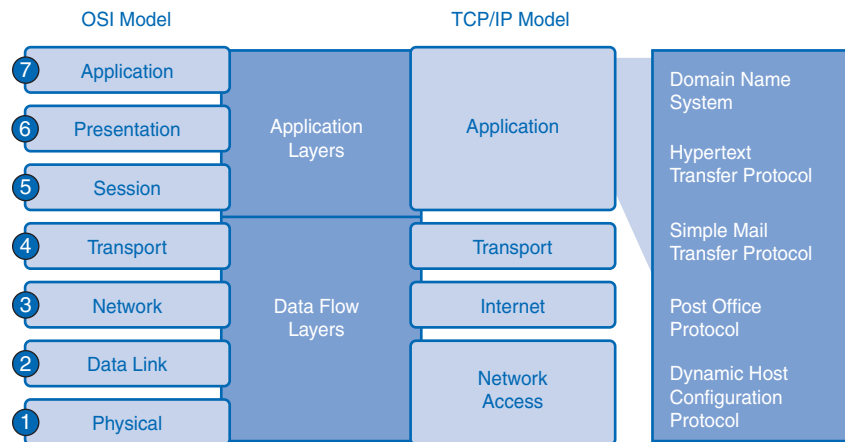


The application layer, Layer 7, is the top layer of both the OSI and TCP/IP models. (Refer to the section “Protocol and Reference Models” in Chapter 2, “Communicating over the Network,” for more information about the TCP/IP model.) Layer 7 provides the interface between the applications you use to communicate and the underlying network over which your messages are transmitted. Application layer protocols are used to exchange *data* between programs running on the source and destination hosts. There are many application layer protocols, and new protocols are always being developed. (Refer to the section “User Applications, Services, and Application Layer Protocols,” later in this chapter, for examples.)

Although the TCP/IP protocol suite was developed prior to the definition of the OSI model, the functionality of the TCP/IP application layer protocols fits roughly into the framework of the top three layers of the OSI model: application, presentation, and session.

Most applications, such as web browsers or e-mail clients, incorporate functionality of the OSI Layers 5, 6, and 7. A comparison of the OSI and TCP/IP model is shown in Figure 3-3.

Figure 3-3 OSI and TCP/IP Model



Most TCP/IP application layer protocols were developed before the emergence of personal computers, GUIs, and multimedia objects. As a result, these protocols implement little of the functionality that is specified in the OSI model presentation and session layers. The next sections describe the OSI presentation and session layers in more detail.

Presentation Layer

The presentation layer has three primary functions:

- Coding and conversion of application layer data to ensure that data from the *source device* can be interpreted by the appropriate application on the destination device

- Compression of the data in a manner that can be decompressed by the destination device
- Encryption of the data for transmission and decryption of data upon receipt by the destination

Presentation layer implementations are not typically associated with a particular protocol stack. The standards for video and graphics are examples. Some well-known standards for video include QuickTime and Motion Picture Experts Group (MPEG). QuickTime is an Apple Computer specification for video and audio, and MPEG is a standard for video compression and coding.

Among the well-known graphic image formats are Graphics Interchange Format (GIF), Joint Photographic Experts Group (JPEG), and Tagged Image File Format (TIFF). GIF and JPEG are compression and coding standards for graphic images, and TIFF is a standard coding format for graphic images.

Session Layer

Functions at the session layer create and maintain dialogs between source and destination applications. The session layer handles the exchange of information to initiate dialogs and keep them active, and to restart sessions that are disrupted or idle for a long period of time.

TCP/IP Application Layer Protocols

The most widely known TCP/IP application layer protocols are those that provide the exchange of user information. These protocols specify the format and control information necessary for many of the common Internet communication functions. Among these TCP/IP protocols are the following:

- *Domain Name System (DNS)* is used to resolve Internet names to IP addresses.
- Hypertext Transfer Protocol (HTTP) is used to transfer files that make up the web pages of the World Wide Web.
- Simple Mail Transfer Protocol (SMTP) is used for the transfer of mail messages and attachments.
- Telnet, a terminal emulation protocol, is used to provide remote access to servers and networking devices.
- File Transfer Protocol (FTP) is used for interactive file transfer between systems.

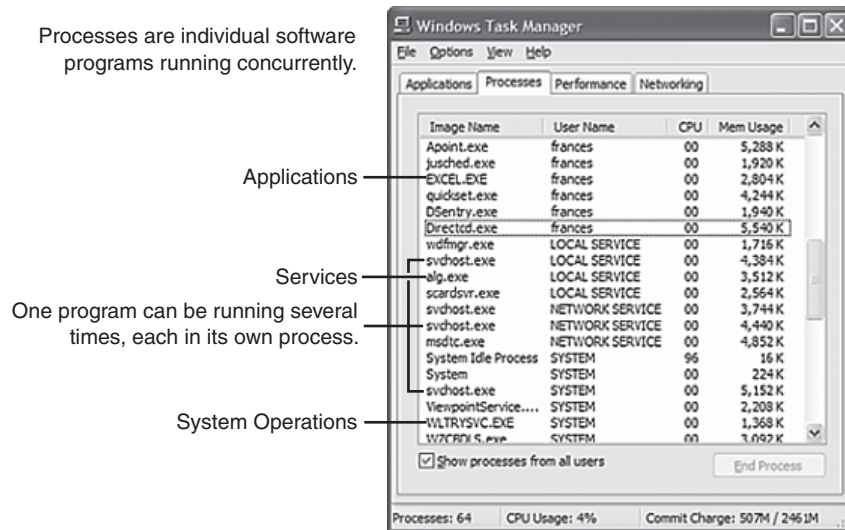
The protocols in the TCP/IP suite are generally defined by *Requests for Comments (RFC)*. The Internet Engineering Task Force (IETF) maintains the RFCs as the standards for the TCP/IP suite.

Application Layer Software

The functions associated with the application layer protocols in both the OSI and the TCP/IP models enable the human network to interface with the underlying data network. When you open a web browser or an instant message window, an application is started, and the program is put into the device memory, where it is executed. Each executing program loaded on a device is referred to as a *process*.

Within the application layer, there are two forms of software programs or processes that provide access to the network: applications and services. This concept is shown in Figure 3-4.

Figure 3-4 Software Processes



Network-Aware Applications

Some end-user applications are network aware, meaning that they implement the application layer protocols and are able to communicate directly with the lower layers of the protocol stack. E-mail clients and web browsers are examples of these types of applications.

Application Layer Services

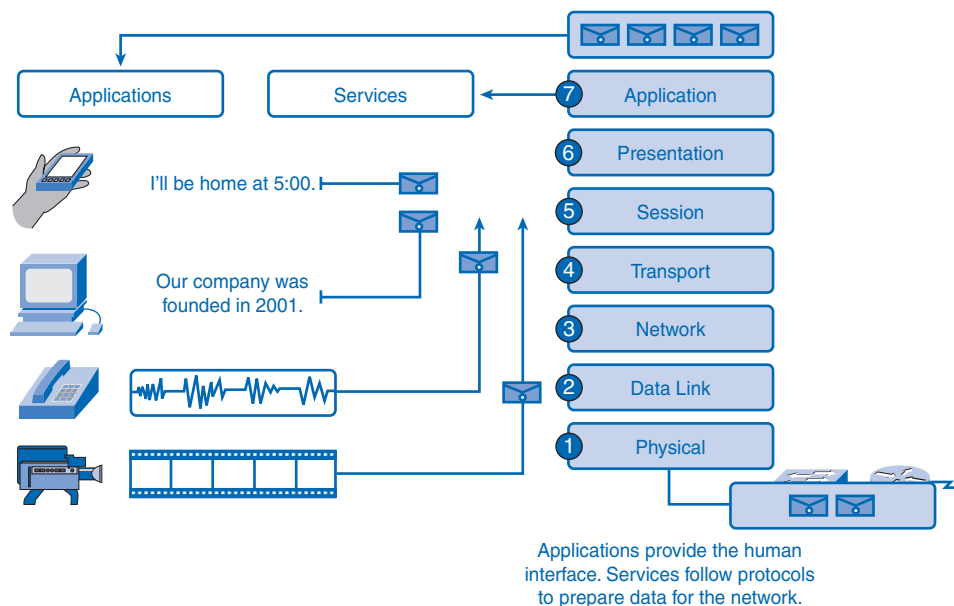
Other programs, such as file transfer or network print spooling, might need the assistance of application layer services to use network resources. Although transparent to the user, these services interface with the network and prepare the data for transfer. Different types of data—whether it is text, graphics, or video—require different network services to ensure that it is properly prepared for processing by the functions occurring at the lower layers of OSI model.

Each application or network service uses protocols that define the standards and data formats to be used. A service provides the function for doing something, and a protocol provides the rules the service uses. To understand the *function* of various network services, you need to become familiar with the underlying protocols that govern their operation.

User Applications, Services, and Application Layer Protocols

The application layer uses protocols that are implemented within applications and services. Applications provide people with a way to create messages, application layer services establish an interface to the network, and protocols provide the rules and formats that govern how data is treated, as shown in Figure 3-5. A single executable program can use all three components. For example, when discussing “Telnet,” you could be referring to the Telnet application, the Telnet service, or the Telnet protocol.

Figure 3-5 Interfacing Human and Data Networks



In the OSI model, applications that interact directly with people are considered to be at the top of the stack, as are the people themselves. Like all layers within the OSI model, the application layer relies on the functions of the lower layers to complete the communication process. Within the application layer, protocols specify what messages are exchanged between the source and destination hosts, the *syntax* of the control commands, the type and format of the data being transmitted, and the appropriate methods for error notification and recovery.

Application Layer Protocol Functions

Both the source and destination devices use application layer protocols during a communication *session*. For the communications to be successful, the application layer protocols implemented on the source and destination host must match.

Protocols perform the following tasks:

- Establish consistent rules for exchanging data between applications and services loaded on the participating devices.
- Specify how data inside the messages is structured and the types of messages that are sent between source and destination. These messages can be requests for services, acknowledgments, data messages, status messages, or error messages.
- Define message dialogues, ensuring that a message being sent is met by the expected response and that the correct services are invoked when data transfer occurs.

Many different types of applications communicate across data networks. Therefore, application layer services must implement multiple protocols to provide the desired range of communication experiences. Each protocol has a specific purpose and contains the characteristics required to meet that purpose. The right protocol details in each layer must be followed so that the functions at one layer interface properly with the services in the lower layer.

Applications and services can also use multiple protocols in the course of a single conversation. One protocol might specify how to establish the network connection, and another might describe the process for the data transfer when the message is passed to the next lower layer.

Making Provisions for Applications and Services

When people attempt to access information on their device, whether it is a PC, laptop, PDA, cell phone, or some other device connected to a network, the data might not be physically stored on their device. If that is the case, a request to access that information must be made to the device where the data resides. The following sections cover three topics that will help you understand how the request for data can occur and how the request is filled:

- Client/server model
- Application layer services and protocols
- Peer-to-peer networking and applications

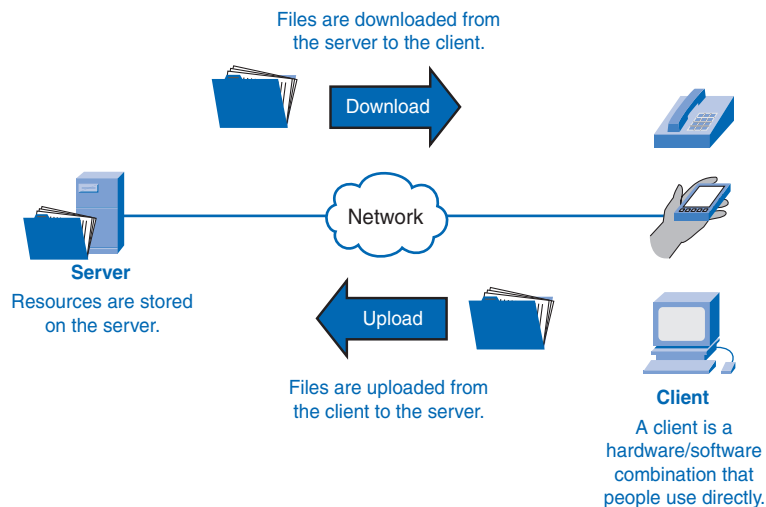
Client/Server Model

In the client/server model, the device requesting the information is called a *client* and the device responding to the request is called a server. Client and *server* processes are considered to be in the application layer. The client begins the exchange by requesting data from the server, which responds by sending one or more streams of data to the client. Application layer protocols describe the design of the requests and responses between clients and servers. In addition to the actual data transfer, this exchange can require control information, such as user authentication and the identification of a data file to be transferred.

One example of a client/server network is a corporate environment where employees use a company e-mail server to send, receive, and store e-mail. The e-mail client on an employee computer issues a request to the e-mail server for any unread mail. The server responds by sending the requested e-mail to the client.

Although data is typically described as flowing from the server to the client, some data always flows from the client to the server. Data flow can be equal in both directions or can even be greater in the direction going from the client to the server. For example, a client might transfer a file to the server for storage purposes. Data transfer from a client to a server is referred to as an *upload*, and data from a server to a client is a *download*. Figure 3-6 shows the client/server model concept.

Figure 3-6 Client/Server Model



Servers

In a general networking context, any device that responds to requests from client applications is functioning as a server. A server is usually a computer that contains information to be shared with many client systems. For example, web pages, documents, databases,

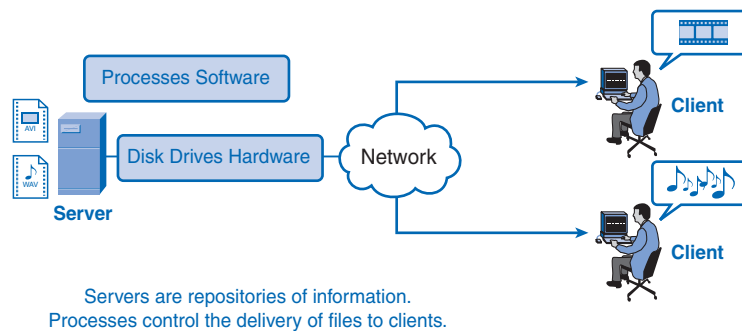
pictures, video, and audio files can all be stored on a server and delivered to requesting clients. In other cases, such as a network printer, the print server delivers the client print requests to the specified printer.

Different types of server applications can have different requirements for client access. Some servers can require authentication of user account information to verify whether the user has permission to access the requested data or to use a particular operation. Such servers rely on a central list of user accounts and the authorizations, or permissions (both for data access and operations), granted to each user. When using an FTP client, for example, if you request to upload data to the FTP server, you might have permission to write to your individual folder but not to read other files on the site.

In a client/server network, the server runs a service, or process, sometimes called a server *daemon*. Like most services, daemons typically run in the background and are not under an end user's direct control. Daemons are described as "listening" for a request from a client, because they are programmed to respond whenever the server receives a request for the service provided by the daemon. When a daemon "hears" a request from a client, it exchanges appropriate messages with the client, as required by its protocol, and proceeds to send the requested data to the client in the proper format.

Figure 3-7 shows the clients requesting services from the server; specifically, one client is requesting an audio file (.wav) and the other client is requesting a video file (.avi). The server responds by sending the requested files to the clients.

Figure 3-7 Servers

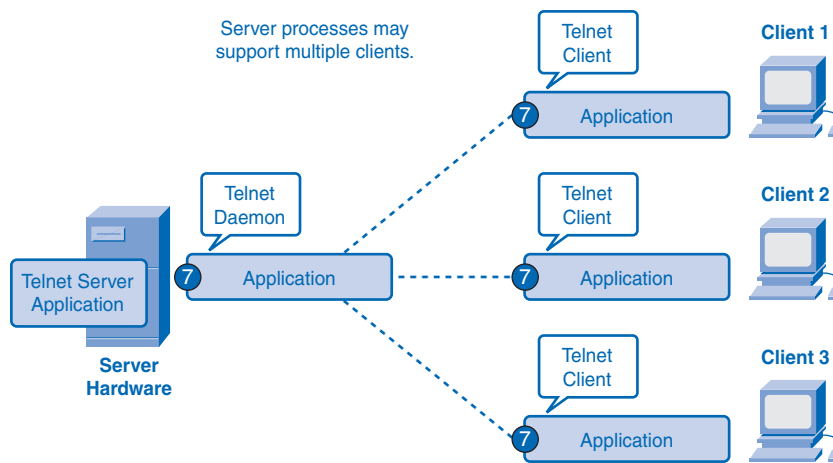


Application Layer Services and Protocols

A single application can employ many different supporting application layer services. Thus, what appears to the user as one request for a web page might, in fact, amount to dozens of individual requests. For each request, multiple processes can be executed. For example, the FTP requires a client to initiate a control process and a data stream process to a server.

Additionally, servers typically have multiple clients requesting information at the same time, as shown in Figure 3-8. For example, a Telnet server can have many clients requesting connections to it. These individual client requests must be handled simultaneously and separately for the network to succeed. The application layer processes and services rely on support from lower-layer functions to successfully manage the multiple conversations.

Figure 3-8 Multiple Clients' Service Requests



Packet Tracer
Activity

Client Server Interaction (3.2.3.2)

In this activity, you will study a simple example of client/server interaction, which can serve as a model for more complex interactions later in the course. Use file e1-3232.pka on the CD-ROM that accompanies this book to perform this activity using Packet Tracer.

Peer-to-Peer (P2P) Networking and Applications

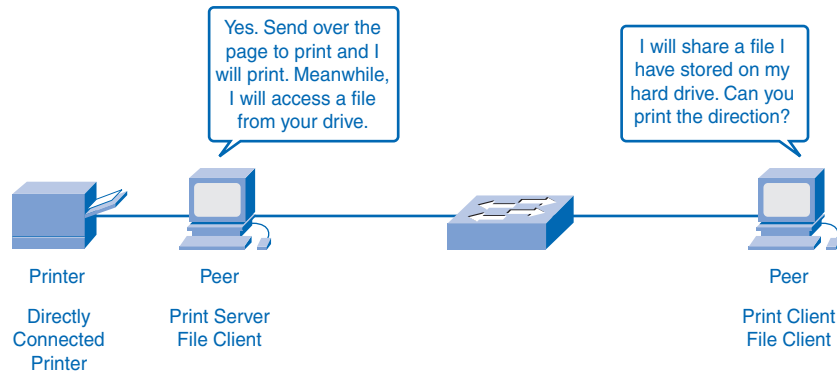
In addition to the client/server model for networking, there is a peer-to-peer (P2P) model. P2P networking involves two distinct forms: peer-to-peer network design and peer-to-peer applications. Both forms have similar features but in practice work very differently.

P2P Networks

In a peer-to-peer network, two or more computers are connected through a network and can share resources such as printers and files without having a dedicated server. Every connected end device, known as a peer, can function as either a server or a client. One computer might assume the role of server for one transaction while simultaneously serve as a client for another. The roles of client and server are set on a per-request basis, as shown in Figure

3-9. The figure shows one peer asking the other peer to provide print services, while at the same time acting as a file server that shares one of its files.

Figure 3-9 Peer-to-Peer Networking



In a peer-to-peer exchange, both devices are considered equal in the communication process.

A simple home network with two connected computers sharing a printer is an example of a peer-to-peer network. Each person can set his or her computer to share files, enable networked games, or share an Internet connection. Another example of peer-to-peer network functionality is two computers connected to a large network that use software applications to share resources between one another through the network.

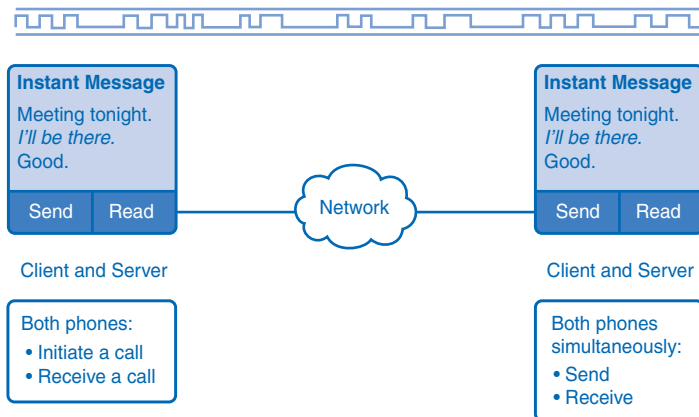
Unlike the client/server model, which uses dedicated servers, peer-to-peer networks decentralize the resources on a network. Instead of locating information to be shared on dedicated servers, information can be located anywhere on any connected device. Most of the current operating systems support file and print sharing without requiring additional server software. Because peer-to-peer networks usually do not use centralized user accounts, permissions, or monitors, it is difficult to enforce security and access policies in networks containing more than just a few computers. User accounts and access rights must be set individually on each *peer* device.

P2P Applications

A P2P application, unlike a peer-to-peer network, allows a device to act as both a client and a server within the same communication session. In this model, every client is a server and every server a client, as shown in Figure 3-10. Figure 3-10 shows two phones belonging to the same network sending an instant message. The blue lines at the top of the figure depict the digital traffic between the two phones. Both can initiate a communication and are considered equal in the communication process. However, peer-to-peer applications require that each end device provide a user interface and run a background service. When you launch a

specific peer-to-peer application, it invokes the required user interface and background services. After that, the devices can communicate directly.

Figure 3-10 Peer-to-Peer Applications



A type of peer-to-peer application is the P2P hybrid system, which utilizes a centralized directory called an index server even though the files being shared are on the individual host machines. Each peer accesses the index server to get the location of a resource stored on another peer. The index server can also help connect two peers, but after they are connected, the communication takes place between the two peers without additional communication to the index server.

Peer-to-peer applications can be used on peer-to-peer networks, in client/server networks, and across the Internet.

Application Layer Protocols and Services Examples

Now that you have a better understanding of how applications provide an interface for the user and provide access to the network, you will take a look at some specific commonly used protocols.

As you will see later in this book, the transport layer uses an addressing *scheme* called a port number. Port numbers identify applications and application layer services that are the source and destination of data. Server programs generally use predefined port numbers that are commonly known by clients. As you examine the different TCP/IP application layer protocols and services, you will be referring to the TCP and UDP port numbers normally associated with these services. Some of these services are

- **Domain Name System (DNS):** TCP/UDP port 53
- **HTTP:** TCP port 80

- **Simple Mail Transfer Protocol (SMTP):** TCP port 25
- **Post Office Protocol (POP):** UDP port 110
- **Telnet:** TCP port 23
- **DHCP:** UDP port 67
- **FTP:** TCP ports 20 and 21

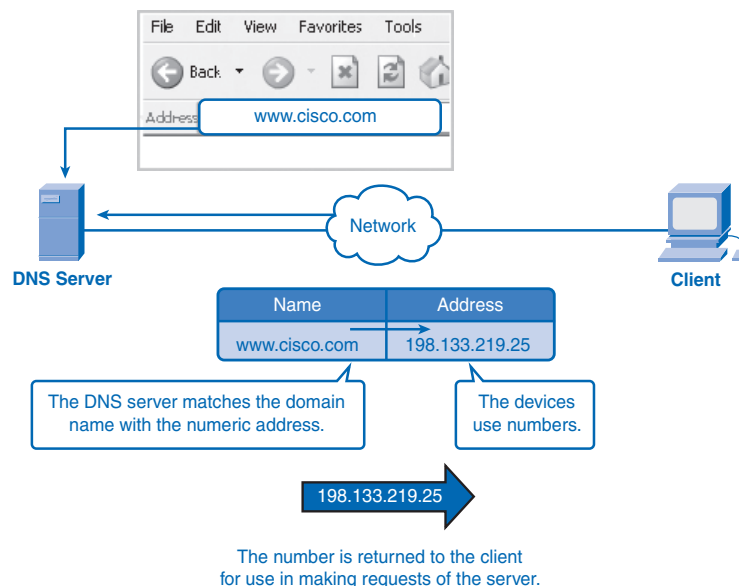
The next sections take a closer look at DNS, world wide web services, and HTTP.

DNS Services and Protocol

In data networks, devices are assigned *IP addresses* so that they can participate in sending and receiving messages over the network. However, most people have a hard time remembering this numeric address. Hence, domain names were created to convert the numeric address into a simple, recognizable name.

On the Internet, these domain names, such as <http://www.cisco.com>, are much easier for people to remember than 198.132.219.25, which, at the time of this writing, is the numeric address for this server. Also, if Cisco decides to change the numeric address, it is transparent to the user, because the *domain name* will remain <http://www.cisco.com>. The new address will simply be linked to the existing domain name and connectivity is maintained, as shown in Figure 3-11. When networks were small, it was a simple task to maintain the mapping between domain names and the addresses they represented. However, as networks began to grow and the number of devices increased, this manual system became unworkable.

Figure 3-11 Resolving DNS Addresses



DNS was created for domain name-to-address resolution for these networks. DNS uses a distributed set of servers to resolve the names associated with these numbered addresses.

How DNS Works

The DNS protocol defines an automated service that matches resource names with the required numeric *network address*. It includes the format for queries, responses, and data formats. DNS protocol communications use a single format called a message. This message format is used for all types of client queries and server responses, error messages, and the transfer of *resource record* information between servers.

DNS is a client/server service; however, it differs from the other client/server services that you are examining. Whereas other services use a client that is an application (web browser, e-mail client, and so on), the DNS client runs as a service itself. The DNS client, sometimes called the *DNS resolver*, supports name resolution for the other network applications and other services that need it.

When configuring a network device, you generally provide one or more DNS server addresses that the DNS client can use for name resolution. Usually the Internet service provider (ISP) gives you the addresses to use for the DNS servers. When a user's application requests to connect to a remote device by name, the requesting DNS client queries one of these DNS servers to resolve the name to a numeric address.

Computer operating systems also have a utility called *nslookup* that allows the user to manually *query* the name servers to resolve a given host name. You also can use this utility to troubleshoot name resolution issues and to verify the current status of the name servers.

In Example 3-1, when the **nslookup** command is issued, the default DNS server configured for your host is displayed. In this example, the DNS server is dns-sjk.cisco.com, which has an address of 171.68.226.120.

Example 3-1 nslookup Command

```
Microsoft Windows XP [Version 5.1.2600]
Copyright 1985-2001 Microsoft Corp.

C:\> nslookup

Default Server: dns-sjk.cisco.com
Address: 171.68.226.120
>www.cisco.com
Server: dns-sj.cisco.com
Address: 171.70.168.183

Name: www.cisco.com
Address: 198.133.219.25
```


You can then type the name of a host or domain for which you want to get the address. In the first query in Example 3-1, a query is made for www.cisco.com. The responding name server provides the address of 198.133.219.25.

Although the queries shown in Example 3-1 are only simple tests, the **nslookup** command has many options available to do extensive testing and verification of the DNS process.

Name Resolution and Caching

A DNS server provides the name resolution using the name daemon, which is often called *named* (pronounced name-dee). The DNS server acts as the phone book for the Internet: It translates human-readable computer host names, for example, <http://www.cisco.com>, into the IP addresses that networking equipment needs for delivering information.

The DNS server stores different types of resource records used to resolve names. These records contain the name, address, and type of record.

Some of these record types are

- **A:** An end device address
- **NS:** An authoritative name server
- **CNAME:** The canonical name (or fully qualified domain name [FQDN]) for an alias; used when multiple services have the single network address but each service has its own entry in DNS
- **MX:** Mail exchange record; maps a domain name to a list of mail exchange servers for that domain

When a client makes a query, the “named” process first looks at its own records to see whether it can resolve the name. If it is unable to resolve the name using its stored records, it contacts other servers to resolve the name.

The request can be passed along to a number of servers, which can take extra time and consume bandwidth. When a match is found and returned to the original requesting server, the server temporarily stores the numbered address that matches the name in the *cache*.

If that same name is requested again, the first server can return the address by using the value stored in its name cache. Caching reduces both the DNS query data network traffic and the workloads of servers higher up the hierarchy. The DNS client service on Windows PCs optimizes the performance of DNS name resolution by storing previously resolved names in memory, as well. The **ipconfig/displaydns** command displays all the cached DNS entries on a Windows XP or 2000 computer system.

DNS Hierarchy

DNS uses a hierarchical system to create a name database to provide name resolution. The hierarchy looks like an inverted tree with the root at the top and branches below.

At the top of the hierarchy, the root servers maintain records about how to reach the top-level domain servers, which in turn have records that point to the secondary-level domain servers and so on.

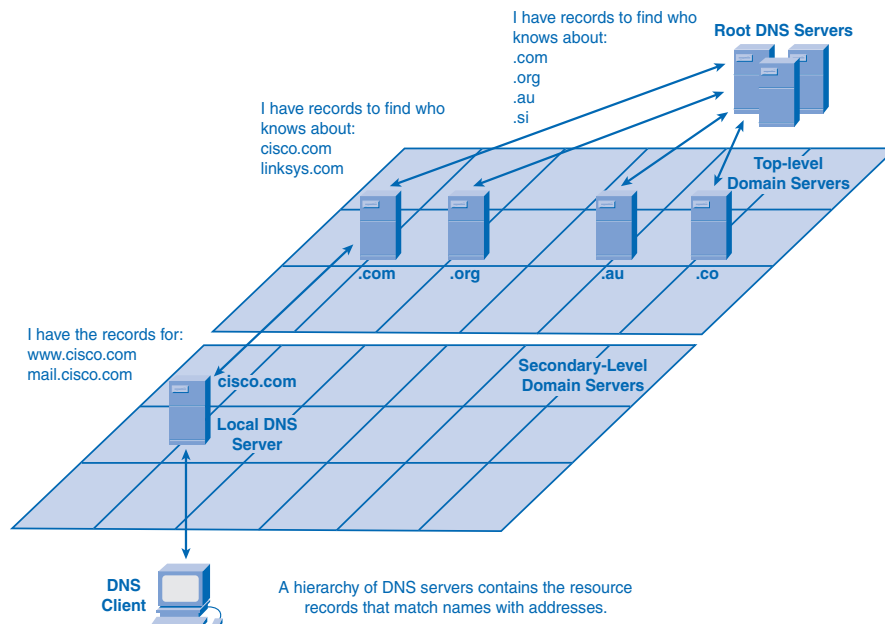
The different top-level domains represent either the type of organization or the country of origin. The following are examples of top-level domains are:

- **.au:** Australia
- **.co:** Colombia
- **.com:** A business or industry
- **.jp:** Japan
- **.org:** A nonprofit organization

After top-level domains are second-level domain names, and below them are other lower-level domains. A great example of that is the domain name <http://www.cisco.netacad.net>. The **.net** is the top-level domain, **.netacad** is the second-level domain, and **.cisco** is at the lower level.

Each domain name is a path down this inverted tree starting from the root. For example, as shown in Figure 3-12, the root DNS servers might not know exactly where the e-mail server mail.cisco.com is located, but they maintain a record for the **.com** domain within the top-level domain. Likewise, the servers within the **.com** domain might not have a record for mail.cisco.com, but they do have a record for the **cisco.com** secondary-level domain. The servers within the **cisco.com** domain have a record (an MX record to be precise) for mail.cisco.com.

Figure 3-12 DNS Server Hierarchy



DNS relies on this hierarchy of decentralized servers to store and maintain these resource records. The resource records list domain names that the server can resolve and alternative servers that can process requests. If a given server has resource records that correspond to its level in the domain hierarchy, it is said to be *authoritative* for those records.

For example, a name server in the cisco.netacad.net domain would not be authoritative for the mail.cisco.com record because that record is held at a higher-domain-level server, specifically the name server in the cisco.com domain.

Note

Two links to the DNS protocol RFCs are

- <http://www.ietf.org/rfc/rfc1034.txt>
- <http://www.ietf.org/rfc/rfc1035.txt>

Request for Comments (RFC) are standards documents encompassing new research, innovations, and methodologies applicable to Internet technologies. These RFCs are very technical in nature, but they can provide you with some insight to how detailed these standards really are.

WWW Service and HTTP

When a web address (or URL) is typed into a web browser, the web browser establishes a connection to the web service running on the server using HTTP. URLs and URIs (uniform resource identifiers) are the names most people associate with web addresses.

The URL <http://www.cisco.com/index.html> refers to a specific resource—a web page named [index.html](http://www.cisco.com/index.html) on a server identified as [cisco.com](http://www.cisco.com).

Web browsers are the client applications computers use to connect to the World Wide Web and access resources stored on a web server. As with most server processes, the web server runs as a background service and makes different types of files available.

To access the content, web clients make connections to the server and request the desired resources. The server replies with the resources and, upon receipt, the browser interprets the data and presents it to the user.

Browsers can interpret and present many data types, such as plain text or HTML, the language in which web pages are constructed). Other types of data, however, might require another service or program, typically referred to as a *plug-in* or add-on. To help the browser determine what type of file it is receiving, the server specifies what kind of data the file contains.

To better understand how the web browser and web client interact, you can examine how a web page is opened in a browser. For this example, consider the URL <http://www.cisco.com/web-server.htm>.

First, the browser interprets the three parts of the URL:

- **http**: The protocol or scheme
- **www.cisco.com**: The server name
- **web-server.htm**: The specific filename requested

The browser then checks with a name server to convert <http://www.cisco.com> into a numeric address, which it uses to connect to the server. Using the HTTP requirements, the browser sends a GET request to the server and asks for the file `web-server.htm`. The server in turn sends the HTML code for this web page to the browser. Finally, the browser deciphers the HTML code and formats the page for the browser window.

HTTP, one of the protocols in the TCP/IP suite, was originally developed to publish and retrieve HTML pages and is now used for *distributed, collaborative* information systems. HTTP is used across the world wide web for data transfer and is one of the most used application protocols.

HTTP specifies a request/response protocol. When a client, typically a web browser, sends a request message to a server, the HTTP protocol defines the message types the client uses to request the web page and the message types the server uses to respond. The three common message types are:

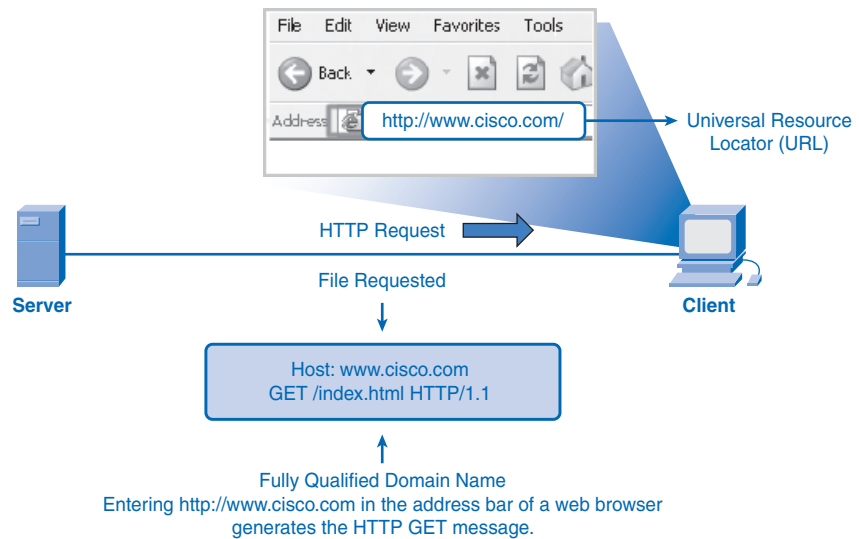
- GET
- POST
- PUT

GET is a client request for data. A web browser sends the GET message to request pages from a web server. As shown in Figure 3-13, when the server receives the GET request, it responds with a status line, such as HTTP/1.1 200 OK, and a message of its own, the body of which can be the requested file, an error message, or some other information.

POST and PUT are used to send messages that upload data to the web server. For example, when the user enters data into a form embedded in a web page, POST includes the data in the message sent to the server. PUT uploads resources or content to the web server.

Although it is remarkably flexible, HTTP is not a secure protocol. The POST messages upload information to the server in plain text that can be intercepted and read. Similarly, the server responses, typically HTML pages, are unencrypted.

For secure communication across the Internet, the Secure HTTP (HTTPS) protocol is used for accessing and posting web server information. HTTPS can use authentication and *encryption* to secure data as it travels between the client and server. HTTPS specifies additional rules for passing data between the application layer and the transport layer.

Figure 3-13 HTTP Protocol Using GET

Packet Tracer
Activity

Network Representations (3.3.2.3)

In this activity, you will configure DNS and HTTP services, and then study the packets that result when a web page is requested by typing a URL. Use file e1-3323.pka on the CD-ROM that accompanies this book to perform this activity using Packet Tracer.

E-Mail Services and SMTP/POP Protocols

E-mail, the most popular network service, has revolutionized how people communicate through its simplicity and speed. Yet to run on a computer or other end device, e-mail requires several applications and services. Two examples of application layer protocols are *Post Office Protocol (POP)* and *Simple Mail Transfer Protocol (SMTP)*. As with HTTP, these protocols define client/server processes.

POP and POP3 (Post Office Protocol, version 3) are inbound mail delivery protocols and are typical client/server protocols. They deliver e-mail from the e-mail server to the client (MUA).

SMTP, on the other hand, governs the transfer of outbound e-mail from the sending client to the e-mail server (MDA), as well as the transport of e-mail between e-mail servers (MTA). (These acronyms are defined in the next section.) SMTP enables e-mail to be transported across data networks between different types of server and client software and makes e-mail exchange over the Internet possible.

When people compose e-mail messages, they typically use an application called a *Mail User Agent (MUA)*, or e-mail client. The MUA allows messages to be sent and places

received messages into the client mailbox, both of which are distinct processes, as shown in Figure 3-14.

Figure 3-14 E-Mail Client (MUA)



To receive e-mail messages from an e-mail server, the e-mail client can use POP. Sending e-mail from either a client or a server uses message formats and command strings defined by the SMTP protocol. Usually an e-mail client provides the functionality of both protocols within one application.

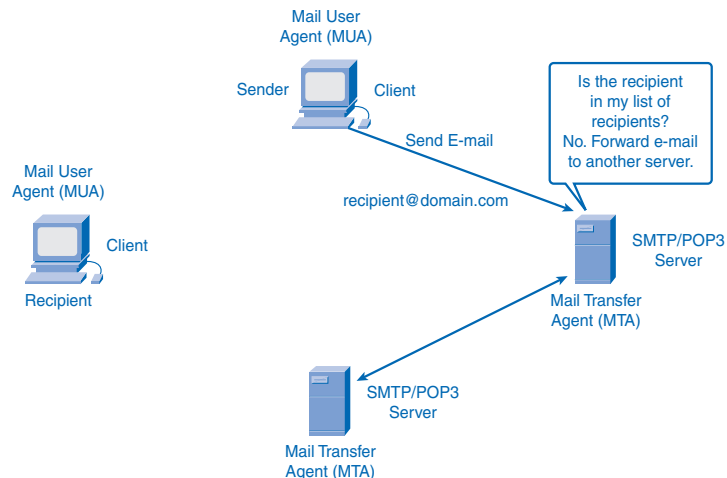
E-Mail Server Processes: MTA and MDA

The e-mail server operates two separate processes:

- Mail Transfer Agent (MTA)
- Mail Delivery Agent (MDA)

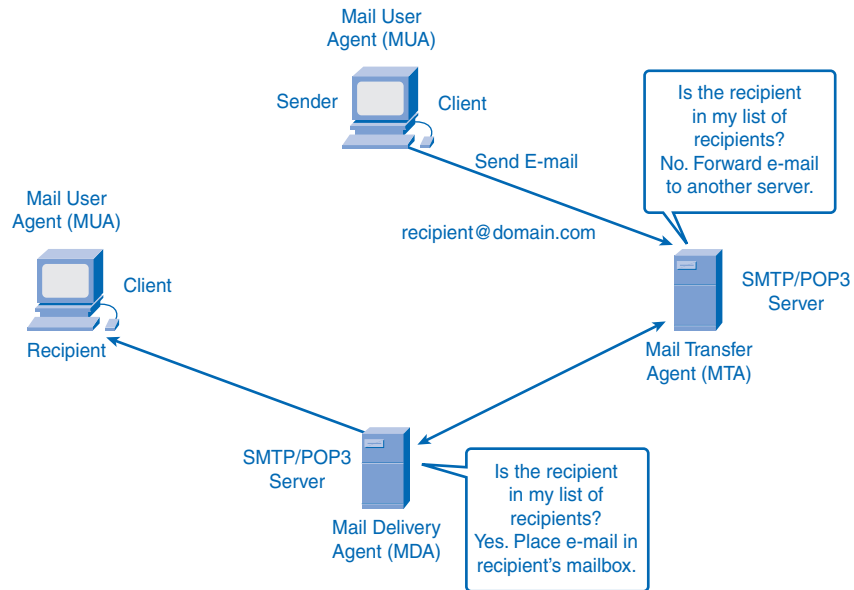
The Mail Transfer Agent (MTA) process is used to forward e-mail. As shown in Figure 3-15, the MTA receives messages from the MUA or from another MTA on another e-mail server. Based on the message header, it determines how a message has to be forwarded to reach its destination. If the mail is addressed to a user whose mailbox is on the local server, the mail is passed to the MDA. If the mail is for a user not on the local server, the MTA routes the e-mail to the MTA on the appropriate server.

Figure 3-15 E-Mail Server: MTA



In Figure 3-16, you see that the Mail Delivery Agent (MDA) accepts a piece of e-mail from a Mail Transfer Agent (MTA) and performs the delivery. The MDA receives all the inbound mail from the MTA and places it into the appropriate users' mailboxes. The MDA can also resolve final delivery issues, such as virus scanning, *spam* filtering, and return-receipt handling.

Figure 3-16 E-Mail Server: MDA



Most e-mail communications use the MUA, MTA, and MDA applications. However, there are other alternatives for e-mail delivery. A client can be connected to a corporate e-mail system, such as IBM Lotus Notes, Novell Groupwise, or Microsoft Exchange. These systems often have their own internal e-mail format, and their clients typically communicate with the e-mail server using a proprietary protocol.

The server sends or receives e-mail through the Internet through the product's Internet mail *gateway*, which performs any necessary reformatting. If, for example, two people who work for the same company exchange e-mail with each other using a proprietary protocol, their messages can stay completely within the corporate e-mail system of the company.

As another alternative, computers that do not have an MUA can still connect to a mail service on a web browser to retrieve and send messages in this manner. Some computers can run their own MTA and manage interdomain e-mail themselves.

The SMTP protocol message format uses a rigid set of commands and replies. These commands support the procedures used in SMTP, such as session initiation, mail transaction,

forwarding mail, verifying mailbox names, expanding mailing lists, and the opening and closing exchanges.

Some of the commands specified in the SMTP protocol are:

- **HELO:** Identifies the SMTP client process to the SMTP server process
- **EHLO:** Is a newer version of HELO, which includes services extensions
- **MAIL FROM:** Identifies the sender
- **RCPT TO:** Identifies the recipient
- **DATA:** Identifies the body of the message

FTP

FTP is another commonly used application layer protocol. FTP was developed to allow file transfers between a client and a server. An FTP client is an application that runs on a computer that is used to push and pull files from a server running the FTP daemon (FTPD).

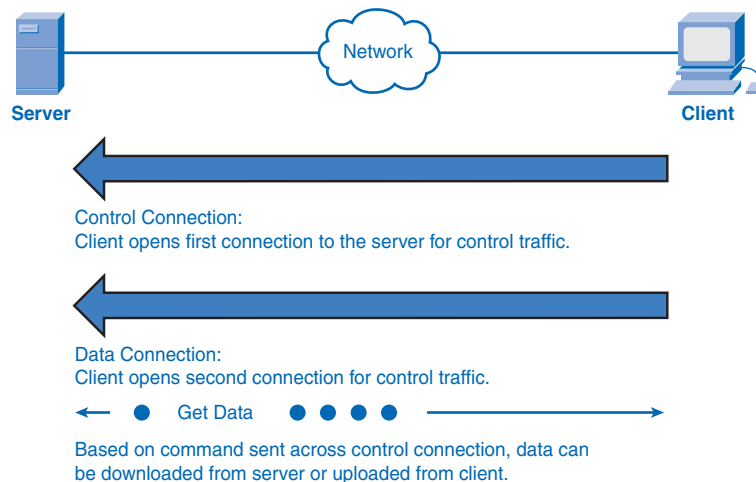
To successfully transfer files, FTP requires two connections between the client and the server: one for commands and replies, and the other for the actual file transfer.

The client establishes the first connection to the server on TCP port 21. This connection is used for control traffic, consisting of client commands and server replies.

The client establishes the second connection to the server over TCP port 20. This connection is for the actual file transfer and is created every time a file is transferred.

The file transfer can happen in either direction, as shown in Figure 3-17. The client can download (pull) a file from the server or upload (push) a file to the server.

Figure 3-17 FTP Process



DHCP

The *DHCP* enables clients on a network to obtain IP addresses and other information from a DHCP server. The protocol automates the assignment of IP addresses, subnet masks, gateway, and other IP networking parameters.

DHCP allows a host to obtain an IP address dynamically when it connects to the network. The DHCP server is contacted by sending a request, and an IP address is requested. The DHCP server chooses an address from a configured range of addresses called a *pool* and assigns it to the host client for a set period.

On larger networks, local networks, or where the user population changes frequently, DHCP is preferred. New users might arrive with laptops and need a connection. Others have new workstations that need to be connected. Rather than have the network administrator assign IP addresses for each workstation, it is more efficient to have IP addresses assigned automatically using DHCP.

When a DHCP-configured device boots up or connects to the network, the client broadcasts a DHCP DISCOVER packet to identify any available DHCP servers on the network. A DHCP server replies with a DHCP OFFER, which is a lease offer message with an assigned IP address, *subnet mask*, DNS server, and default gateway information as well as the duration of the lease.

DHCP-distributed addresses are not permanently assigned to hosts but are only leased for a period of time. If the host is powered down or taken off the network, the address is returned to the pool for reuse. This is especially helpful with mobile users who come and go on a network. Users can freely move from location to location and re-establish network connections. The host can obtain an IP address after the hardware connection is made, either through a wired or wireless LAN.


DHCP makes it possible for you to access the Internet using wireless hotspots at airports or coffee shops. As you enter the area, your laptop DHCP client contacts the local DHCP server through a wireless connection. The DHCP server assigns an IP address to your laptop.

Various types of devices can be DHCP servers when running DHCP service software. The DHCP server in most medium to large networks is usually a local dedicated PC-based server.

With home networks, the DHCP server is usually located at the ISP, and a host on the home network receives its IP configuration directly from the ISP.

Many home networks and small businesses use an Integrated Services Router (ISR) device to connect to the ISP. In this case, the ISR is both a DHCP client and a server. The ISR acts as a client to receive its IP configuration from the ISP and then acts a DHCP server for internal hosts on the local network.

Figure 3-18 shows the different ways of having DHCP servers arranged.

Figure 3-18 DHCP Servers

DHCP can pose a security risk because any device connected to the network can receive an address. This risk makes physical security an important factor when determining whether to use dynamic or static (manual) addressing.

Dynamic and static addressing have their places in network designs. Many networks use both DHCP and static addressing. DHCP is used for general-purpose hosts such as end-user devices, and static, or fixed, addresses are used for network devices such as gateways, switches, servers, and printers.

The client can receive multiple DHCP OFFER packets if the local network has more than one DHCP server. The client must choose between them and *broadcast* a DHCP REQUEST packet that identifies the explicit server and lease offer that it is accepting. A client can choose to request an address that it had previously been allocated by the server.

Assuming that the IP address requested by the client, or offered by the server, is still valid, the chosen server would return a DHCP ACK (acknowledgment) message. The ACK message lets the client know that the lease is finalized. If the offer is no longer valid for some reason, perhaps because of a timeout or another client allocating the lease, the chosen server must respond to the client with a DHCP NAK (negative acknowledgment) message. When the client has the lease, it must be renewed prior to the lease expiration through another

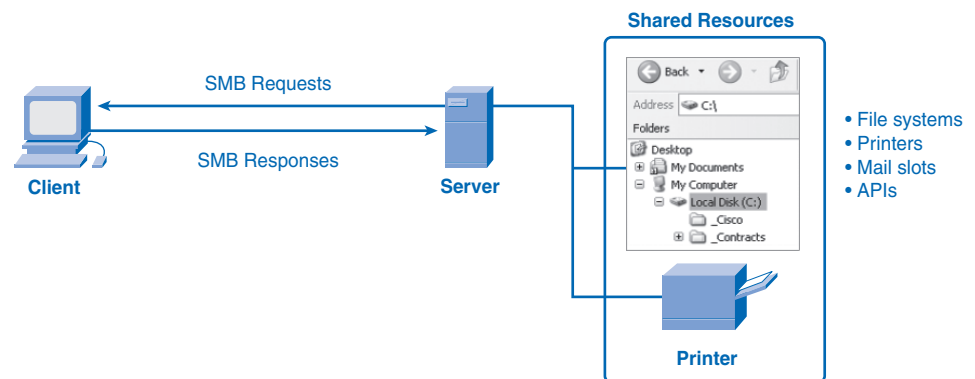
DHCP REQUEST message. The DHCP server ensures that all IP addresses are unique. (An IP address cannot be assigned to two different network devices simultaneously.)

File-Sharing Services and SMB Protocol

Server Message Block (SMB) is a client/server file-sharing protocol. IBM developed SMB in the late 1980s to describe the structure of shared network resources, such as directories, files, printers, and serial ports. It is a request/response protocol. Unlike the file sharing supported by FTP, clients establish a long-term connection to servers. After the connection is established, the user of the client can access the resources on the server as if the resource is local to the client host.

SMB file-sharing and print services have become the mainstay of Microsoft networking. With the introduction of the Windows 2000 series of software, Microsoft changed the underlying structure for using SMB. In previous versions of Microsoft products, the SMB services used a non-TCP/IP protocol to implement name resolution. Beginning with Windows 2000, all subsequent Microsoft products use DNS naming. This allows TCP/IP protocols to directly support SMB resource sharing, as shown in Figure 3-19.

Figure 3-19 File Sharing Using the SMB Protocol



The Linux and *UNIX* operating systems also provide a method of sharing resources with Microsoft networks using a version of SMB called SAMBA. The Apple Macintosh operating systems also support resource sharing using the SMB protocol.

The SMB protocol describes file system access and indicates how clients can make requests for files. It also describes the SMB protocol interprocess communication. All SMB messages share a common format. This format uses a fixed-sized header followed by a variable-sized parameter and data component.

SMB messages can perform the following tasks:

- Start, authenticate, and terminate sessions
- Control file and printer access
- Allow an application to send or receive messages to or from another device

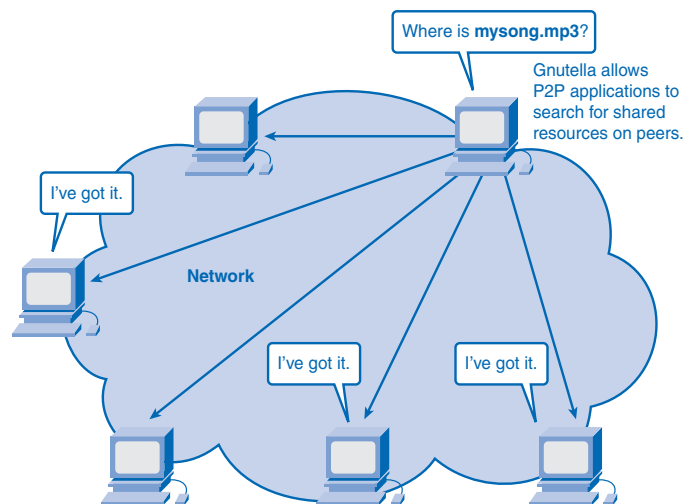
P2P Services and Gnutella Protocol

You learned about FTP and SMB as ways of obtaining files. This section describes another application protocol, Gnutella. Sharing files over the Internet has become extremely popular. With P2P applications based on the Gnutella protocol, people can make files on their hard disks available to others for downloading. Gnutella-compatible client software allows users to connect to Gnutella services over the Internet and to locate and access resources shared by other Gnutella peers.

Many client applications are available for accessing the Gnutella network, including BearShare, Gnutelux, LimeWire, Morpheus, WinMX, and XoloX. Although the Gnutella Developer Forum maintains the basic protocol, application vendors often develop extensions to make the protocol work better on their applications.

Many P2P applications do not use a central database to record all the files available on the peers. Instead, the devices on the network each tell the other what files are available when queried and use the Gnutella protocol and services to support locating resources, as shown in Figure 3-20. When a user is connected to a Gnutella service, the client applications will search for other Gnutella nodes to connect to. These nodes handle queries for resource locations and replies to those requests. They also govern control messages, which help the service discover other nodes. The actual file transfers usually rely on HTTP services.

Figure 3-20 Gnutella Protocol



The Gnutella protocol defines five different packet types:


- **ping:** For device discovery
- **pong:** As a reply to a ping
- **query:** For file location
- **query hit:** As a reply to a query
- **push:** As a download request

Telnet Services and Protocol

Long before desktop computers with sophisticated graphical interfaces existed, people used text-based systems that were often just display terminals physically attached to a central computer. After networks were available, people needed a way to remotely access the computer systems in the same manner that they did with the directly attached terminals.

Telnet was developed to meet that need. It dates back to the early 1970s and is among the oldest of the application layer protocols and services in the TCP/IP suite. Telnet is a client/server protocol that provides a standard method of emulating text-based terminal devices over the data network. Both the protocol itself and the client software that implements the protocol are commonly referred to as Telnet. The Telnet service is depicted in Figure 3-21.

Figure 3-21 Telnet Service



Appropriately enough, a connection using Telnet is called a *VTY (Virtual Terminal) session*, or *connection*. Telnet specifies how a VTY session is established and terminated. It also provides the syntax and order of the commands used to initiate the Telnet session, and it provides control commands that can be issued during a session. Each Telnet command consists of at least 2 bytes. The first byte is a special character called the *Interpret as Command (IAC)* character. As its name implies, the IAC character defines the next byte as a command rather than text. Rather than using a physical device to connect to the server, Telnet uses software to create a virtual device that provides the same features of a terminal session with access to the server command-line interface (CLI).

To support Telnet client connections, the server runs a service called the Telnet daemon. A virtual terminal connection is established from an end device using a Telnet client application. Most operating systems include an application layer Telnet client. On a Microsoft Windows PC, Telnet can be run from the command prompt. Other common terminal applications that run as Telnet clients are HyperTerminal, Minicom, and TeraTerm.

When a Telnet connection is established, users can perform any authorized function on the server, just as if they were using a command-line session on the server itself. If authorized, they can start and stop processes, configure the device, and even shut down the system.

The following are some sample Telnet protocol commands:

- **Are You There (AYT):** Enables the user to request that a response, usually a prompt icon, appear on the terminal screen to indicate that the VTY session is active.
- **Erase Line (EL):** Deletes all text from the current line.
- **Interrupt Process (IP):** Suspends, interrupts, aborts, or terminates the process to which the virtual terminal is connected. For example, if a user started a program on the Telnet server through the VTY, he or she could send an IP command to stop the program.

Although the Telnet protocol supports user authentication, it does not support the transport of encrypted data. All data exchanged during a Telnet session is transported as plain text across the network. This means that the data can be intercepted and easily understood.

The Secure Shell (SSH) protocol offers an alternate and secure method for server access. SSH provides the structure for secure remote login and other secure network services. It also provides stronger authentication than Telnet and supports the transport of session data using encryption. As a best practice, network professionals should use SSH in place of Telnet, whenever possible.

Summary

The application layer is responsible for directly accessing the underlying processes that manage and deliver communication to the human network. This layer serves as the source and destination of communications across data networks. The application layer applications, protocols, and services enable users to interact with the data network in a way that is meaningful and effective.

Applications are computer programs with which the user interacts and that initiate the data transfer process at the user's request.

Services are background programs that provide the connection between the application layer and the lower layers of the networking model.

Protocols provide a structure of agreed-upon rules, much like grammar and punctuation provide "rules" in a language. These protocol rules ensure that services running on one particular device can send and receive data from a range of different network devices.

Delivery of data over the network can be requested from a server by a client. In a peer-to-peer arrangement, either device can function as a client or server, and data is delivered depending on the client/server relationship established. Messages are exchanged between the application layer services at each end device in accordance with the protocol specifications to establish and use these relationships.

Protocols like HTTP, for example, support the delivery of web pages to end devices. SMTP/POP protocols support sending and receiving e-mail. SMB enables users to share files. DNS resolves the human-legible names used to refer to network resources into numeric addresses usable by the network. Telnet provides remote, text-based access to devices. DHCP provides dynamic allocation of IP addresses and other network-enabling parameters. P2P allows two or more computers to share resources over the network.

Activities and Labs

The activities and labs available in the companion *Network Fundamentals, CCNA Exploration Labs and Study Guide* (ISBN 1-58713-203-6) provide hands-on practice with the following topics introduced in this chapter:



Activity 3-1: Data Stream Capture (3.4.1.1)

In this activity, you will use a computer that has a microphone and Microsoft Sound Recorder or Internet access so that an audio file can be downloaded.

**Lab 3-1: Managing a Web Server (3.4.2.1)**

In this lab, you will download, install, and configure the popular Apache web server. You will use a web browser to connect to the server and Wireshark to capture the communication. Analyzing the capture will help you understand how HTTP operates.

**Lab 3-2: E-mail Services and Protocols (3.4.3.1)**

In this lab, you will configure and use an e-mail client application to connect to eagle-server network services. You will then monitor the communication with Wireshark and analyze the captured packets.



Many of the hands-on labs include Packet Tracer companion activities where you can use Packet Tracer to complete a simulation of the lab. Look for this icon in the *Network Fundamentals, CCNA Exploration Labs and Study Guide* (ISBN 1-58713-203-6) for hands-on labs that have Packet Tracer companion activities.

Check Your Understanding

Complete all the review questions listed here to test your understanding of the topics and concepts in this chapter. The appendix, “Check Your Understanding and Challenge Questions Answer Key,” lists the answers.

1. The application layer is _____ of the OSI model.
 - A. Layer 1
 - B. Layer 3
 - C. Layer 4
 - D. Layer 7
2. The TCP/IP application layer consists roughly of which three OSI layers?
 - A. Application, session, transport
 - B. Application, presentation, session
 - C. Application, transport, network
 - D. Application, network, data link
3. HTTP is used to do which of the following?
 - A. Resolve Internet names to IP addresses
 - B. Provide remote access to servers and networking devices
 - C. Transfer files that make up the web pages of the World Wide Web
 - D. Transfer the mail messages and attachments

4. Post Office Protocol (POP) uses which port?
 - A. TCP/UDP port 53
 - B. TCP port 80
 - C. TCP port 25
 - D. UDP port 110

5. What is GET?
 - A. A client request for data
 - B. A protocol that uploads resources or content to the web server
 - C. A protocol that uploads information to the server in plain text that can be intercepted and read
 - D. A response from a server

6. Which is the most popular network service?
 - A. HTTP
 - B. FTP
 - C. Telnet
 - D. E-mail

7. FTP requires _____ connection(s) between client and server to successfully transfer files.
 - A. 1
 - B. 2
 - C. 3
 - D. 4

8. DHCP enables clients on a network to do which of the following?
 - A. Have unlimited telephone conversations
 - B. Play back video streams
 - C. Obtain IP addresses
 - D. Track intermittent denial of service attacks

9. The Linux and UNIX operating systems use SAMBA, which is a version of which protocol?
 - A. SMB
 - B. HTTP
 - C. FTP
 - D. SMTP

10. Which of the following is a connection using Telnet?
 - A. File Transfer Protocol (FTP) session
 - B. Trivial File Transfer Protocol (TFTP) session
 - C. Virtual Terminal (VTY) session
 - D. Auxiliary (AUX) session
11. Is eBay a peer-to-peer or client/server application?
12. In the client/server model, the device requesting the service is referred to as the _____.
13. HTTP is referred to as a request/response protocol. What are three typical message formats?
14. DHCP allows the automation of what?
15. What does FTP stand for, and what is it used for?

Challenge Questions and Activities

These questions require a deeper application of the concepts covered in this chapter. You can find the answers in the appendix.

1. List the six-step process for converting human communications to data.
2. Describe the two forms of application software and the purpose of each.
3. Elaborate on the meaning of the terms server and client in the context of data networks.
4. Compare and contrast client/server with peer-to-peer data transfer over networks.
5. List five general functions that application layer protocols specify.
6. Give the specific purposes of the DNS, HTTP, SMB, and SMTP/POP application layer protocols.
7. Compare and contrast the messages that application layer protocols such as DNS, HTTP, SMB, and SMTP/POP exchange between devices to enable data transfers to occur.

To Learn More

The following questions encourage you to reflect on the topics discussed in this chapter. Your instructor might ask you to research the questions and discuss your findings in class.

1. Why is it important to distinguish between a particular application layer application, the associated service, and the protocol? Discuss this in the context of network reference models.
2. What if it was possible to include all application layer services with a single all-encompassing protocol? Discuss the advantages and disadvantages of having one such protocol.
3. How would you develop a new protocol for a new application layer service? What would have to be included? Who would have to be involved in the process, and how would the information be disseminated?

