CHAPTER

10

The Network Layer

It's not what you know but who you know that makes the difference.

— Anonymous

There is not much difference between human networking and computer networking. You can be the most gifted human or the highest powered computer, but lacking the ability to share those resources, you can do nothing as far as the progression of humankind is concerned. The power of information is in its capacity to be shared. Since the evolution of wireless networking, information can be shared not only globally but beyond this world into outer space.[1]

The number of connected computers on the World Wide Web is staggering. Two computers are able to share information between them without concern about how that information is to navigate over the Internet. This is the ''who you know that makes the difference'' portion of what networking is about. Networking is about being able to route information to a particular computer and receive requested information from that computer without a need to know the path it travels over the Internet.

Think of the Internet as a giant matrix with routing devices at every crossing point to aid in the movement of a packet of information along the cables connecting to the next crossing point. The route a packet of information travels can be different each time another packet of information is sent. The routing device's responsibility is to make sure that the packet will arrive at the destination it is intended for.

[1] Amateur Radio on the International Space Station (ARISS) has been experimenting with packet mail from amateur radio operators from around the world to the International Space Station. Although this is not conventional wireless networking, it may be a precursor of things to come when there is a manned base on the moon.

A computer is concerned only with its locally connected default gateway. A *default gateway* is where network traffic is sent when a computer wants to send information to a computer that it knows does not reside on its local network. Every computer and network-connected device has a default gateway set within their network configuration parameters. When information comes in via the Internet, it is accepted by the default gateway and routed on to the local network, directed toward the computer the received data is intended for.

Routing or network traffic forwarding devices need not know every other device that is connected to the Internet. They just need to have a good working relationship with their immediate peers. It is dependent upon networking through these other peer routing devices to know other devices that they also have a working relationship with. It is essential that networks know the right entities to network to.

RANDOM BONUS DEFINITION

mirror port — A switch port configured to reflect the traffic appearing on another of the switch's ports.

The Network layer occupies Layer 3 of the OSI model. It receives network requests from the Transport layer and, in turn, issues network requests to the Data Link layer. It is the layer that is responsible for end-to-end information transfer. The delivery of information is within a *datagram*, also known as a *frame* or *packet*. The Network layer loosely maps to the Internet layer of the TCP/IP model, but the Internet layer deals only with the Internet Protocol (IP), whereas the OSI Network layer encompasses a broader range of both connection-oriented and connectionless network services.

10.1 Network Connection Types

What does a connection-oriented service versus a connectionless network service really mean? All network-enabled devices[2] are connected to a network, right? So they must be connected, right? Well, in the physical sense that is true. However, as far as a network service is concerned, it does make a difference how information is delivered between network nodes. The easy way to differentiate between the two types of network services is that a connected-oriented network service is one where the endpoint network nodes know who a session was established with, whereas in a connectionless network service, the two network nodes do not need to establish a direct connection in order to share information

[2] A ''network-enabled'' device is simply any computer or packet-forwarding device with the right network interface for the network medium connecting the device, along with the appropriate network software.

10.1.1 Connectionless Network Services

How can two network nodes exchange information if they do not have a connection established between them? This is where connectionless network services come into play. A great example of a connectionless network service is e-mail. E-mail is addressed to a particular user residing in a particular domain. It has no relation to a particular computer or geographical location.

The following is an example of a typical e-mail address:

```
john.doe@hishome.com
```

The recipient of this e-mail is `john.doe` who resides in the network domain of `hishome.com`. This brings in the concept once again of domain names and their relationship to network services. There is a hierarchy to network addressing, and the domain name is the highest level. Figure 10-1 illustrates the network addressing hierarchy.

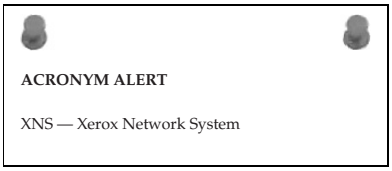

ACRONYM ALERT

XNS — Xerox Network System

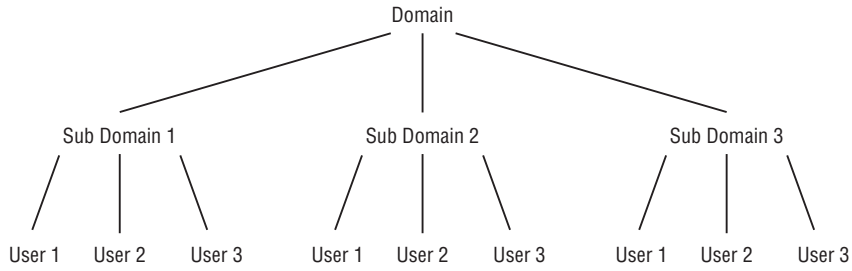

Figure 10-1 The network addressing hierarchy

As shown in Figure 10-1, the top level of addressing is the domain.[3] A domain can contain subdomains that have a varying number of users assigned. For example, the Widget Company has various departments with varying groups of users assigned to those departments. Figure 10-2 could be a method the Widget Company uses to set up their domain.

The Widget Company is a family-owned business founded in the mid-1800s. It prides itself on being wholly American owned and its operations being located only within the geographic boundaries of the United States. Although their products are shipped globally, they support sales and customer service from within the good old USA. Even though they face fierce price cutting from

[3]Domains are named by the organization that wants to create a domain for its network infrastructure. Domain names are usually classified with either a company name or some other meaningful words or acronyms for the easy identification of domain ownership.

manufacturers that off-shored their operations, the Widget family of products have maintained their competitive edge due to superior product reliability and what is considered to be best-in-class customer service.

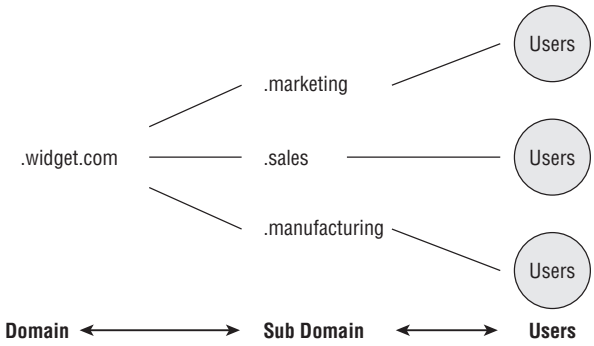

Figure 10-2 The Widget Company's domain hierarchy

The Widget Company wants to create three subdomains for its marketing, sales, and manufacturing departments. These departments have control of various servers that service the users of each department. These users may be either internal or external users over the Internet. The Widget Company domain does not require that all the entities of the domain be located within a single building, city, state, or country. Components that are not only for the overall domain but also for the subdomains may be located in geographically distant locations.[4] However, the network nodes that are part of the domain can still be reached using domain names without the need for absolute address locations. Figure 10-3 illustrates what the overall network topology of Widget Company might look like on a top level network map.

The top level drawing of the Widget Company network shows locations that are solely contained within the United States in various distantly located sites. The various sites are interconnected using the

RANDOM BONUS DEFINITION

link aggregation — The process of combining multiple physical links into a single logical link for use by higher layer link clients.

[4]Geographically distant locations can be in the building next door, down the street, in the next town, in the next state, or in the next country. If they are not on the same local network, they are considered to be distant and require special handling to ensure information is transmitted reliably.

Internet as a transportation medium for the domain's network infrastructure. Because these sites are connected over the Internet, they utilize IP for the transmission from site to site.

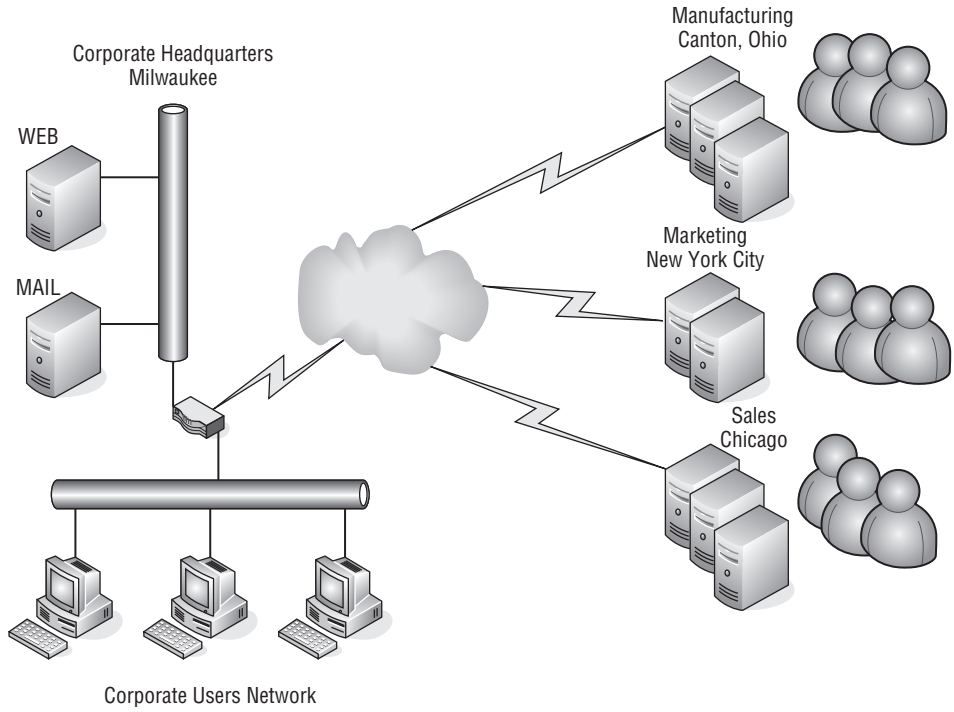

Figure 10-3 The Widget Company's top level network diagram

It was already mentioned that the TCP/IP model's Internet layer is a subset of the OSI model's Network layer. There will be places in this chapter where we discuss the aspects of TCP/IP where it is relevant within the OSI Network layer. The domain aspect can be used for either connection-oriented or connectionless network services. However, the world of TCP/IP uses IP to move information along the world's information highway. To bridge between domain names and IP protocol addresses requires domain name resolution, commonly referred to as DNS (Domain Name System). Further discussion of DNS can be found in Section 10.1.3.

ACRONYM ALERT

Telnet — Teletype Network

As you can see in Figure 10-3, the corporate offices located in Milwaukee have multiple networks, various computer systems, and a number of servers. This diagram is simplistic in its presentation for a large corporate network, which is far more complex. However, the base principles of network interoperability[5] are fairly similar due to the scalability[6] of networking technologies. The figure shows two servers: a mail server and a web server. The remote offices also have servers located at their sites that are able to pass information from other servers and users located either locally or over the Internet. Using domain names to reach various servers has the following format:

```
Host name.domain name.sub domain name.domain name suffix
```
[7]

The mail server named `mail` located at the corporate office would have a domain name that appears as follows:

```
mail.widget.com
```

If the marketing group located in New York City also has a mail server that gathers its mail from the corporate mail server, its name could be:

```
mail.marketing.widget.com
```

Mail shared between users is connectionless[8] because the computer sending the mail does not need a connection directly to the mail server the recipient of the e-mail is connected to. There are differences in e-mail, and perhaps there is some confusion due to the type of e-mail service being used. A local mail program on a computer is capable of creating a mail program entirely independent from any other computer. When it is ready to send the e-mail message, it does so by forwarding the mail to a Simple Mail Transport Protocol (SMTP) server where the user has an account. The message is forwarded by the SMTP server without any further action by the user to aid in the delivery of the message.

[5]''Interoperability'' is just a fancy name for network node devices to play nice with all the other network node devices connected on the same network.

[6]''Scalability'' simply means that networks can start small and grow larger as needed. However, larger networks usually require higher capacity network devices able to handle the amount of information that is to traverse the network within a fixed period of time.

[7]Domain names as illustrated in this example do not have spaces within the name. So, using the above example as a domain name would actually appear as `hostname.domainname.subdomainname.domainnamesuffix`.

[8]A computer connected to its local mail server uses the POP or POP3 protocol to receive mail and SMTP to send mail. These protocols are connection-based because the PC has a direct session with its local mail server. However, mail user to mail user is connectionless because a user-to-user PC session is not needed to send or read mail.

RANDOM BONUS DEFINITION

learning state — A transition state in the Spanning Tree Protocol state machine where a bridge port is learning address-to-port mappings to build its filtering database before entering the forwarding state.

If a user is using web-based mail, the session established by the browser to create the e-mail is a connection-oriented network service. In using web mail, the user establishes a connection to the server serving his or her account to create and forward the message. However, the type of service is still connectionless since the user is not required to provide any further action to ensure delivery of the e-mail message. This illustrates that even connectionless processes may require some elements of a connection-oriented network service.

SMTP mail servers deliver e-mail to the SMTP mail server servicing a particular domain. Although a user name is attached as part of the message, the SMTP server does not deliver the message to the user. A user must have an account on a mail server in order for the mail to be delivered to that user's post office box. In the case of incorrect spelling of a user name or if a user never had an e-mail account or their account had been deleted, the SMTP server would return the original message with an error header[9] stating the cause for the message not being delivered. The most common reason for return is ''user unknown.

E-mail for a user is held on the mail server for a period of time established by the administrator of that server. There are various parameters on most mail servers that allow for a mailbox's size, usually in megabytes, length of time a message is held, and the maximum allowable size of a message. An error message may be returned to an e-mail sender if the recipient is not in compliance with any of the preset parameters. Depending on the mail service provided by the mail server, mail may be read while remaining on a mail server or it may have to be downloaded using the Post Office Protocol (POP or POP3) to the local workstation for reading and any other required action.[10]

POP QUIZ

Mail is what type of network service?

To summarize, a connectionless network service has the capability to prepare information for transmittal to another network node without the creation of a real-time connection to that network node in order to complete the transfer of the information being sent.

[9]In computerese, the header is simply the top of the message. In other words, you do not need to read the whole message to see why it was bounced back.

[10]The required action is usually reading the message and either filing it or discarding it. Unfortunately, just like your postal mailbox, your e-mail mailbox also gets a lot of junk mail.

10.1.2 Connection-Oriented Network Services

A connection-oriented network service is exactly what the name implies. A network connection[11] is established between two computers to transfer information from one computer to the other over the Internet. Many client/server application programs are connection-oriented network services. A good example of this would be the interaction between an FTP client and an FTP server.[12] Figure 10-4 illustrates a user residing on a local network at IP address 192.168.2.13 requesting information from a local FTP server whose IP address is 192.168.2.5.

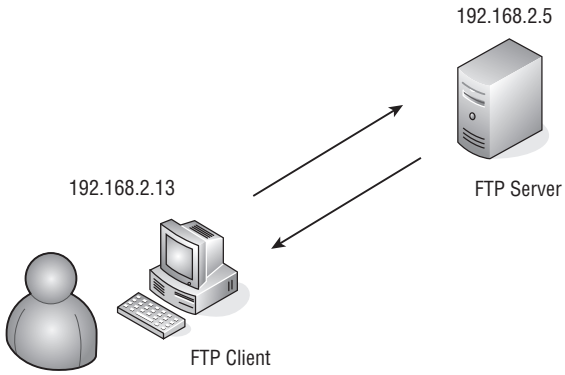

Figure 10-4 An FTP client/server connection-oriented network server

The following portion of the FTP server log illustrates the interaction of the client with the FTP server:

```
Oct 11 20:28:35     Cerberus FTP Server started
Oct 11 20:28:35     Local Host: Rbramant-2

Oct 11 20:28:35     Local Interface 0 located at 192.168.2.5
Oct 11 20:28:35     Listening on Port 21

Oct 11 20:34:39  1  Incoming connection request on interface
                    192.168.2.5
Oct 11 20:34:39  1  Connection request accepted from 192.168.2.13
Oct 11 20:34:52  1  USER anonymous
Oct 11 20:34:52  1  331 User anonymous, password please
Oct 11 20:34:57  1  PASS ***********
```

[11] Although networks use electrical connections for signal transmission, a network connection is when two endpoint network node devices know each other and establish a session that is connected.

[12] Many places within the text server and client are shown and discussed as totally separate network entities. In reality, a computer can be both a server and a client simultaneously for network services.

```
Oct 11 20:34:57  1  230 Password Ok, User logged in
Oct 11 20:34:57  1  Anonymous user ''anonymous'' logged in with
                    password ''guest''
Oct 11 20:35:00  1  PORT 192,168,2,13,19,137
Oct 11 20:35:00  1  200 Port command received
Oct 11 20:35:00  1  LIST
Oct 11 20:35:00  1  150 Opening data connection
Oct 11 20:35:00  1  226 Transfer complete
Oct 11 20:35:08  1  QUIT
Oct 11 20:35:08  1  Connection terminated.
```

You can see that the client initiated the connection to the server. The server forced the client to supply a user ID and a password. The client responded with a user ID and password combination, and is authenticated and allowed to maintain the session with the FTP server.

The FTP client user requested a directory listing from the FTP server. After the listing was received, the user quit the session and thus caused the termination of the connection between the client and the server.

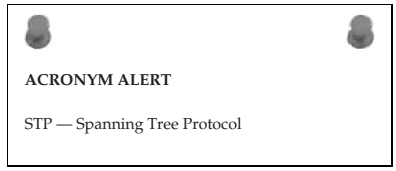

A packet capture of this session was performed at the FTP server, as illustrated in Figure 10-5.

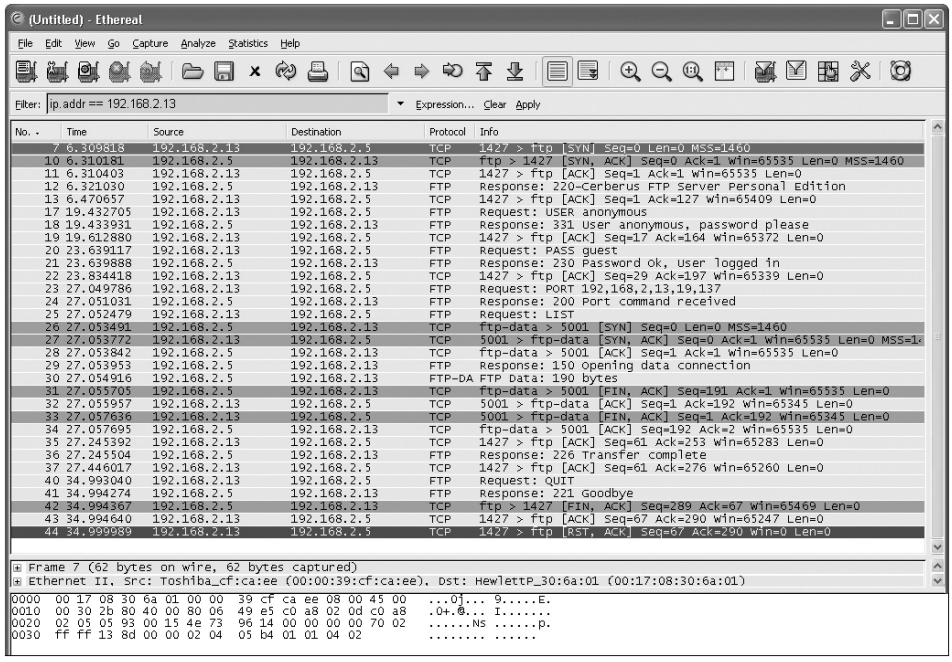

Figure 10-5 A packet capture of an FTP session

The FTP session uses the TCP/IP protocol to establish the session and complete the transfer of information from the FTP server to the FTP client. Packet number 7 shows the client requesting a session with the FTP server. Packet 10 is the FTP server acknowledging the session request. Packets 12 through 21 are the packets showing the interaction between the FTP client and FTP server to authenticate the FTP client and establish the FTP session. Packets 22 through 36 are the directory listing request and the transfer of the directory contents information to the FTP client. Packets 37 through 44 are the packets showing that the FTP client is terminating the FTP session and thus terminating the network connection.

An FTP[13] session does involve layers above the Network layer, but FTP helps illustrate the concept of a connection-oriented network service. The two computers establish a connection session and transfer information between them. The Network layer is responsible only for the end-to-end connections and is not involved with the hop-to-hop[14] transfer of the packets over the network.

POP QUIZ

Name the ports used by an FTP client to request an FTP session with an FTP server. Which port is used for data transmission?

WANT TO TRY SOMETHING?

You are encouraged to reproduce the FTP session as illustrated in this section. It requires two computers and software that can be obtained by a free download from the Internet. The FTP session was accomplished by using FTP server software from `www.cerberusftp.com` and using the ftp command from the command prompt of a Windows XP PC. You can obtain packet capture software for free from `www.wireshark.org`. The computers can either be on the same network segment or on different segments with network routing devices between the network segments.

10.1.3 Domain Name Services

Many of you are probably familiar with the term URL (uniform resource locator). A typical URL would appear as follows:

```
http://www.mydomainname.com
```

[13]The FTP protocol uses two ports for control and data transfer. Control is dedicated to port 21, and port 20 is dedicated to data transfer. An FTP server would listen on port 21 for FTP requests, and the FTP session is negotiated and controlled using this port.

[14]A network hop is any network node a data packet needs to be forwarded through on its journey to the requested destination.

The `http` indicates this is a request for port 80 on a computer with the host name `www` located in the domain `mydomain.com`. In the TCP/IP world, computer addresses take the following form:

```
XXX.XXX.XXX.XXX
```

where `XXX` can be a numeric decimal value between 0 and 255. We are preconditioned to think of URLs as being as follows:

```
prefix.domainname.suffix
```

We are accustomed to seeing .com, .org, .gov, .edu, or .net being used as a suffix, although many others are in use. Also, a country code may be used as the suffix to denote where the domain and host computer are found. So how does one get from a text-based URL name to an IP address? Someone has to take care of it, like the telephone company has with the use of area codes, exchange numbers, and the last four unique digits to reach a particular telephone. So in the case of finding an IP address for a particular computer by its host name, who would have the super-sized host name book that lists every computer connected to the Internet?

Telephones are basically static devices. They are wired into a particular telephone switch with a fixed number. Computers can be moved or exchanged with other computers, and occasionally IP addresses associated with a particular URL can also be changed. So, host-name-to-IP addresses can be pretty dynamic, and a dynamic system is required to maintain the capability to perform host name resolution. There needs to be some form of registration to enable this to occur. There are many companies that sell domain name registration for a fee. But what does that really mean?

RANDOM BONUS DEFINITION

jam — In Ethernet, the process of sending an additional 32 data bits following the detection of a collision to ensure that all parties to the collision properly recognize the event as such.

As with IP addresses, domain names also need to be unique. Domain names must be registered to ensure that they are not duplicated on the Internet. The Internet Assigned Numbers Authority (IANA) is an organization created to establish standard naming for what is called the top level domain (TLD), or root zone. The suffix portion of a URL is the root zone. It is used to parse a host name URL to establish which root zone the host name is a member of. The Internet Network Information Center (InterNIC) is maintained by the Internet Corporation for Assigned Names and Numbers (ICANN) and is responsible for the registration of domain names through registered domain name hosting companies.

When a domain name is registered, it is associated with an IP address and is maintained on a DNS server. Each DNS server needs to know what the designated authoritative name server is in order to receive DNS updates. Although the service is fairly dynamic, caching[15] is used to save time querying the root name servers each time a request is made for a particular host name. Figure 10-6 illustrates a typical DNS server scenario.

ACRONYM ALERT

SD — Start delimiter

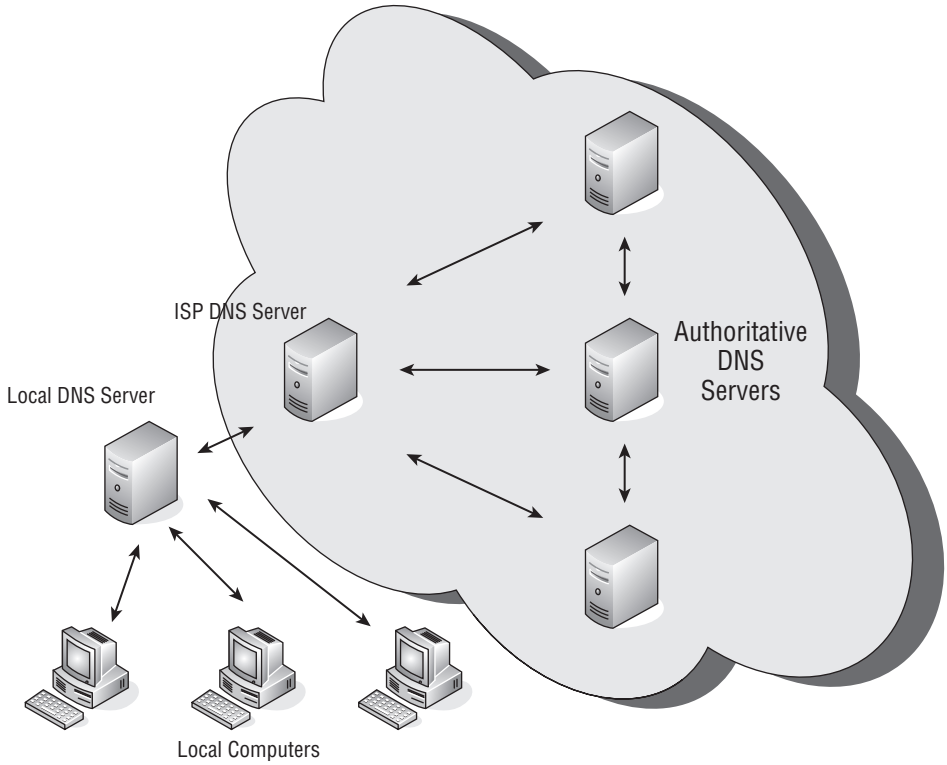

Figure 10-6 A typical DNS server scenario

DNS is part of the TCP/IP protocol suite. The computers on the local network have configured the IP address of the local DNS server into their

[15]Caching is the process of saving information for a predetermined amount of time. In DNS, caching can save time for address resolution. However, to ensure that a name resolution stays ''fresh,'' there is usually an expiration time associated with the entry. Old entries are aged out automatically. When a DNS request is made, if it is not in the cache, name resolution needs to be performed. Although under normal circumstances it is completed fairly rapidly, it does take more time than just pulling it up from the local cache storage.

TCP/IP configuration settings. You can verify these settings by issuing an `ipconfig /all` command at the command window of a Windows-based PC. The response would be similar to the following:

```
Ethernet adapter Local Area Connection:

        Connection-specific DNS Suffix  . :
        Description . . . . . . . . . . . : Broadcom NetXtreme
                                            Gigabit Ethernet
        Physical Address. . . . . . . . . : 00-17-08-30-6A-01
        DHCP Enabled. . . . . . . . . . . : Yes
        Autoconfiguration Enabled . . . . : Yes
        IP Address. . . . . . . . . . . . : 192.168.2.5
        Subnet Mask . . . . . . . . . . . : 255.255.255.0
        Default Gateway . . . . . . . . . : 192.168.2.1
        DHCP Server . . . . . . . . . . . : 192.168.2.4
        DNS Servers . . . . . . . . . . . : 192.168.2.1
        Lease Obtained. . . . . . . . . . : Sunday, October 12, 2008
                                            8:08:02 AM
        Lease Expires . . . . . . . . . . : Monday, October 13, 2008
                                            8:08:02 AM
```

In this example, there is only one DNS server, and it is the same as the device that is acting as the default gateway. In this particular setup, the router is capable of running a DNS service, and its DNS servers are the upstream servers at the ISP, as shown in Figure 10-6.

Using the example of a browser attempting to reach a particular URL, if the computer does not have the resolved host name stored in its local DNS cache, it will request it from its assigned DNS server. Figure 10-7 shows a packet capture of a DNS request from a local PC to its local DNS server.

The user is calling the URL `www.imagesbybramante.com` and, not having the host name cached, it places the request to its local DNS server. If the local DNS server does not have the host name cached, it makes a DNS request to its upstream server and would eventually work its way back to a root authoritative server until the name is resolved. If the name cannot be resolved, an error message is returned. When the name is resolved, it is passed back through the servers until it reaches the computer that made the original request. Figure 10-8 shows a successful host name lookup for the query used in this example.

This has been a top-level discussion of DNS to give you a basic understanding of name resolution in regard to IP. You are encouraged to explore literature dedicated solely to DNS concepts for additional, in-depth information.

POP QUIZ

Name some top-level domain names.

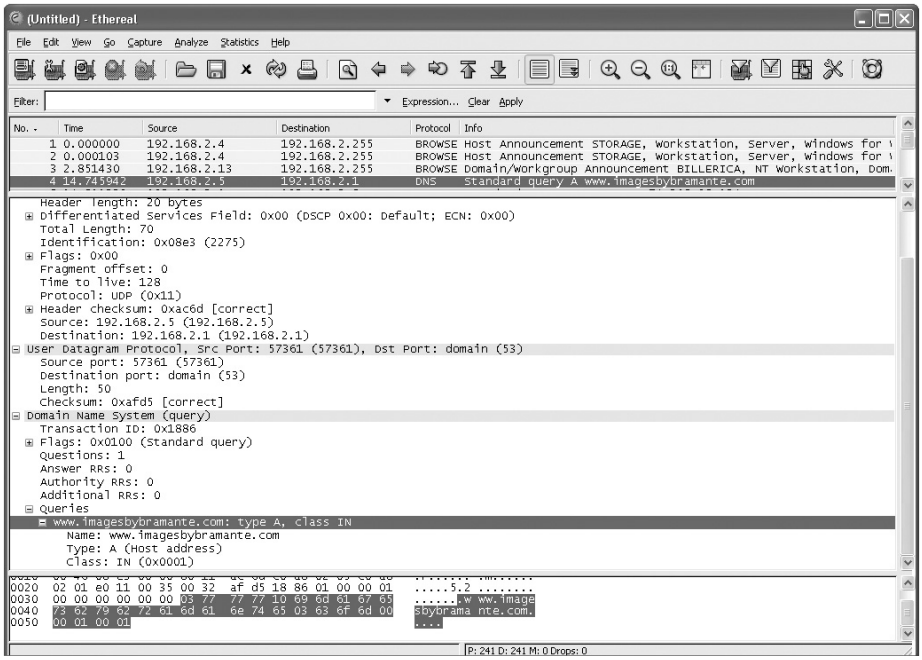

Figure 10-7 A packet capture of a DNS request

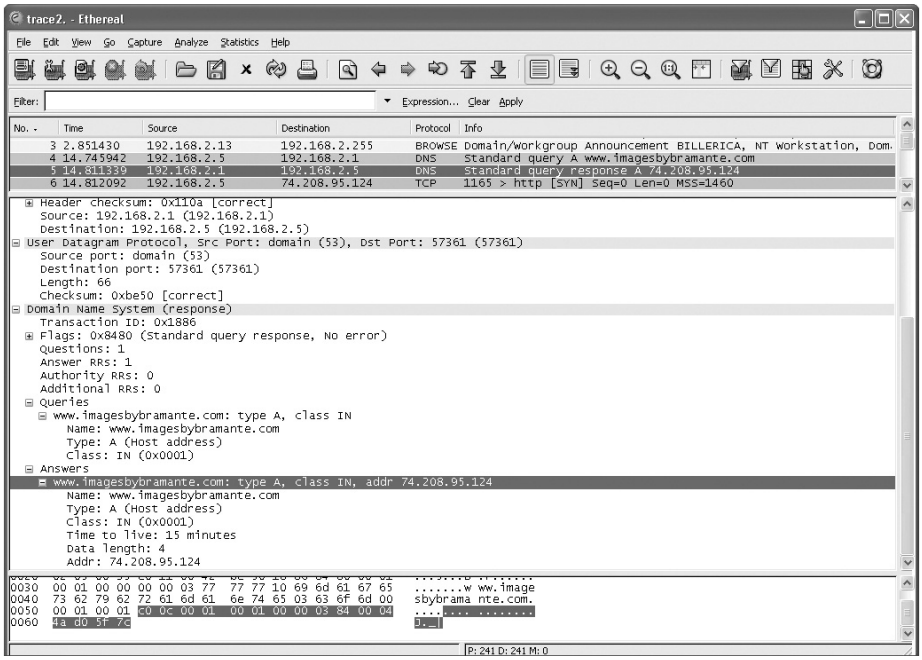

Figure 10-8 A packet capture of a DNS response

SOMETHING TO TRY

We suggested earlier that you download a freeware version of Wireshark. It is a useful tool not only for troubleshooting but to give added insight to what is occurring on your computer in terms of network communications. It can be loaded on either a desktop or laptop with your other Windows-based applications. It may be launched prior to opening any application and allowed to capture the packets of that application. This will help build familiarity with the Wireshark application itself and aid in increasing your understanding of TCP/IP and the protocols supported within the TCP/IP protocol suite.

10.2 TCP/IP Network Layer Protocols

The Network layer of the OSI model provides for both connectionless network services and connection-oriented services. It encompasses the protocols of the TCP/IP model's Internet layer. However, the OSI model's Network layer is broader in scope than TCP/IP's Internet layer, and at times, it includes other TCP/IP protocols from its Link layer. Due to this difference, the two layers should not be considered mirror images of each other, although they do have some protocols in common.

RANDOM BONUS DEFINITION

filtering — The process of inspecting frames received on an input port of a switch and deciding whether to discard or forward them.

10.2.1 Internet Protocol

The Internet Protocol (IP) is primarily a method of moving packets of data across networks comprising various media, seamlessly delivering these packets solely based on the destination address. This is accomplished by encapsulating data from the upper layers into packets[16] in preparation for delivery over the network. IP is a connectionless protocol since packets can be transmitted without the establishment of a circuit to the destination network node. Because IP is a best-effort delivery service, it makes no guarantee that a packet will be delivered. Therefore, data can become corrupted, packets can

[16]The word ''packet'' is synonymous with ''datagram'' or ''frame.'' These three words are used interchangeably and refer to the structure containing all the pertinent information for the proper construct so that the data can be reliably transmitted and that it can be properly unencapsulated when received at the intended network node.

arrive out of order, duplicate packets can be received, and packets can be lost or discarded.

The mainstay for many years has been Internet Protocol version 4 (IPv4), but due to its limitation of addressing, Internet Protocol version 6 (IPv6) is currently being deployed worldwide. To work around the limited address space of IPv4, the development of Network Address Translation (NAT) helped delay the need to deploy IPv6 any sooner.

10.2.1.1 Internet Protocol Version 4

Because IPv4 utilizes 4 bytes to express an address, it has only 32 bits that can be used for its address. This allows for a maximum combination of addresses that can be supported of 2^{32}, or 4,294,967,296 unique addresses. Since some of the addresses are within reserved address space, the total space is not available as public Internet addresses. IPv4 addresses are mostly expressed in what is referred to as *dot-decimal notation*, for example:

```
192.168.15.85
```

Each dotted section is representative of the decimal value of the byte. So it would look as follows in binary:

```
11000000.10101000.00001111.01010101
```

There is a multitude of variations to express IP addresses, but the dot-decimal notation is the most widely used.

REMEDIAL EXERCISE

For those of you who are not proficient in manipulating numbers between various number systems, try to convert the above dotted binary number to a hexadecimal-dotted notation. Hint: The bits of a byte are equally divided to create two hexadecimal numbers for each dotted binary section. A hexadecimal number is usually represented by 0x<hex valued number>. If you want the answer, wait until you give it an honest try, and then look at the footnote below.[17]

If the upper layers present the TCP/IP Internet layer with data that is too large to transmit within a single packet, the data will be fragmented and transmitted over the network in

POP QUIZ

What is a maximum transmission unit?

[17] xC0.0xA0.0x0F.0x55, or in not-dotted notation, 0xC0A00F55

multiple packets. IP performs the fragmentation since it is host dependent, not machine dependent. The maximum transmission unit (MTU) is the number of bytes of data that a particular network medium is capable of handling. It is determined by the largest packet the medium is capable of handling, minus any number of bytes required as a header to transmit the packet over the medium. In the case of Ethernet, which has a maximum packet size of 1500 bytes, the MTU is the maximum packet size minus the number of bytes required for the header. Ethernet normally requires 20 bytes for a header, which provides for an MTU of 1480 bytes. The data is fragmented into the number of packets needed, with each packet tagged indicating it contains a fragment. The receiving network nodes unencapsulate the received fragmented packets and reassemble the data before passing it up to the layer above it.

10.2.1.1.1 Network Address Translation

It was mentioned that Network Address Translation (NAT) was developed to provide a method of using addresses designated as private address space behind a network device, such as a router, that is able to perform the translation from a nonroutable IP address to a publicly known IP address. Table 10-1 shows the reserved addresses for private networking.[18]

Table 10-1 Private Networking Reserved Addresses

ADDRESS RANGE	CIDR	NETWORK CLASS	ADDRESSES
10.0.0.0 to 10.255.255.255	10.0.0.0/8	Single Class A	16,777,216
172.16.0.0 to 172.31.255.255	172.16.0.0/12	16 Contiguous Class B	1,048,576
192.168.0.0 to 192.168.255.255	192.168.0.0/16	Single Class B	65,536

USEFUL NOTE

Generally speaking, the first address of a subnet range ending in 0 is used to designate the network address, and the last address ending in 255 is used to designate the broadcast address of the subnet. It depends on the subnet mask being used whether 0 or 255 is assigned to a host. For example, if the whole

(continued)

[18] Although these address ranges are shown to be contiguous, they may be subdivided (subnetted) into smaller subnet ranges within a private network space. Often, these private network classes can be found using class C 24-bit subnet masks of 255.255.255.0 to form smaller network ranges.

USEFUL NOTE *(continued)*

class A subnet of 10.0.0.0 with a subnet mask of 255.0.0.0 is used for a network, then 10.0.0.0 would be the network address and 10.255.255.255 would be the broadcast address for that subnet. This would permit 10.0.1.0, 10.0.1.255, 10.0.2.0, 10.0.2.255, etc., to be used for host addresses. It is important to know the subnet mask that is assigned to a particular IP address range.

USEFUL NOTE #2

The CIDR is the number that indicates the number of mask bits being assigned to a subnet. So, /8 would have a subnet mask of 255.0.0.0, /12 would have a subnet mask of 255.240.0.0, and /16 would have a subnet mask of 255.255.0.0. Have you guessed what the CIDR represents yet? The CIDR indicates the number of bits for the subnet mask starting at the highest significant position and working its way down to the least significant position.[19]

USEFUL NOTE #3

When you see /32 or a subnet mask of 255.255.255.255, it is called a *host route*. This means there is only one network node connected to that address; there is no network, just one device — and that is it. Host routes are used more frequently than you would think, but be aware of this when someone talks of a "slash 32" or "32-bit route."

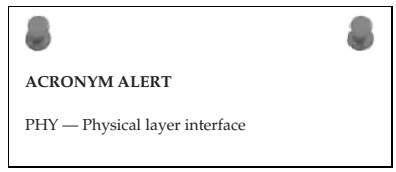

Figure 10-9 illustrates three separate networks all performing NAT on the 192.168.0.0/16 network address space.

You will notice that the IP addresses assigned to the public interface[20] are addresses that are assigned by an ISP (Internet service provider). These addresses can be either statically or dynamically assigned IP addresses. The type of installation usually dictates how IP address assignment is handled. DSL (digital subscriber line), PPPoE (Point-to-Point Protocol over Ethernet), and dialup network circuits are usually configured to have a dynamically assigned IP address. Dynamically assigned addresses

[19]Okay, for the readers who fell asleep during math class: The higher the power of the number, the more significance it has. In our number system, the number to the left of another number has a higher power, thus more significance. The leftmost number is always the most significant number.

[20]A public interface is one that has a publicly routable IP address assigned to it. Private IP addresses cannot be routed over the Internet.

change each time a connection is made. However, the ISP can usually assign a static IP address, if requested to do so.

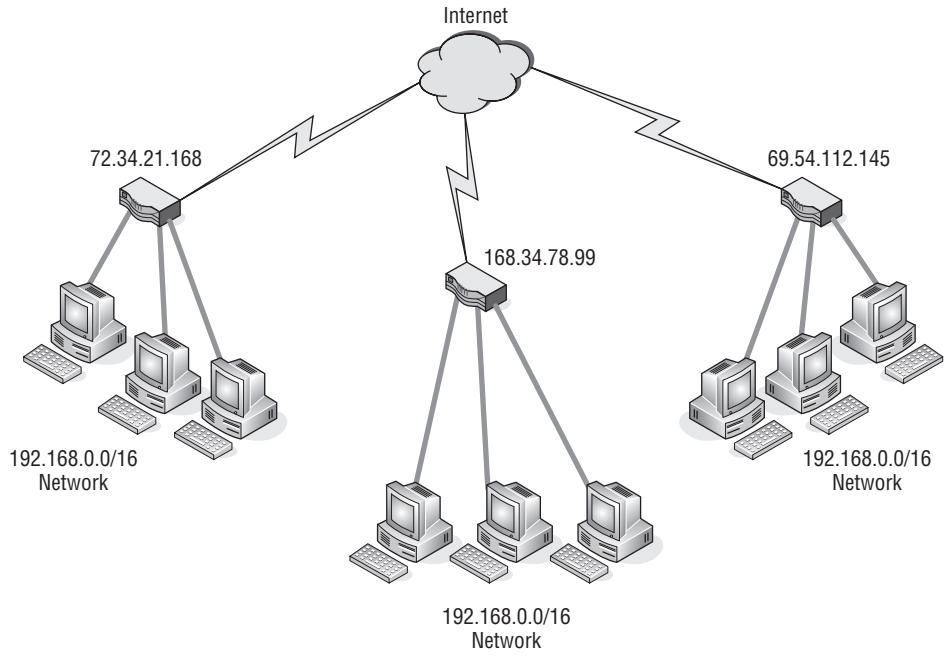

Figure 10-9 A NAT example

Notice that the private IP addresses for all three are in the 192.168.0.0/16 network and that the NAT-enabled router will translate those addresses to its public IP address. The receiving node over the Internet will see the public IP address in the packet's source address. The sending NAT-enabled router keeps a translation table in order to recall which sending workstation on its private IP address space has initiated the session. When the receiving node sends a reply back to the NAT-enabled router, it removes its address from the destination address field of the packet, replaces the IP address from its translation table of the workstation that started the session, and passes the packet into the private network.

ACRONYM ALERT

CSMA/CA — Carrier Sense, Multiple Access with Collision Avoidance

Since workstations are on a private IP address space, they are not reachable from the Internet unless a policy to allow this is embedded within the NAT-enabled router. Such policies are called *port forwarding policies*. Usually servers offering web services, e-mail services, or FTP services are located on private networks behind a NAT-enabled router. Figure 10-10 illustrates

services being offered to users over the Internet while being located behind a NAT-enabled router on a private IP address space.

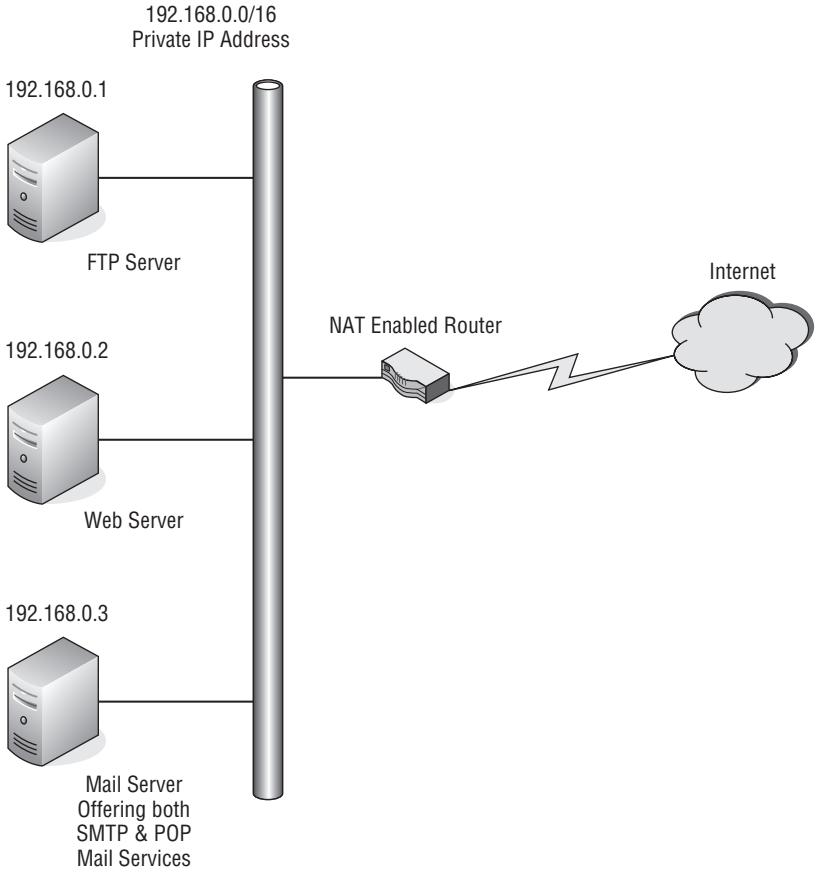

Figure 10-10 Servers behind a NAT-enabled router

Table 10-2 is representative of a NAT port forwarding table a NAT-enabled router would have to accept service requests on its public IP address interface.

When a packet arrives at the public interface with the destination address set to its public IP address and a port service request that matches a port address in the NAT port forwarding table, the packet is modified and passed on to the network, directed toward the server that supplies that service. An example of this is a web page request that arrives at the public IP address of the NAT-enabled server. The NAT-enabled router sees the requested port is port 80, so it replaces its address in the destination field with the IP address of the web server that is at 192.168.0.2, recalculates the checksum for the packet, and passes it on to the private network.

Table 10-2 NAT Port Forwarding Table

SERVICE	PORT	SERVER ADDRESS
SMTP Mail	25	192.168.0.3
POP Mail	109	192.168.0.3
POP3 Mail	110	192.168.0.3
FTP Control	21	192.168.0.1
FTP Data	20	192.168.0.1
HTTP	80	192.168.0.2

NAT does offer some firewall protection since the addresses used on the private IP address are not routed over the Internet. Unsolicited connection requests are dropped by the NAT-enabled router since there is no entry in its NAT translation table. However, when port forwarding policies are enabled within the NAT-enabled router, there is an possibility that one of the servers may be hacked and compromised. A prudent measure would be to have a DMZ (demilitarized zone) by using a router that has multiple private IP address interfaces. Place the servers on one interface isolated by different network addresses and policies. This will prevent the servers from initiating connections into the private network where other users and devices are protected behind the NAT-enabled router.

POP QUIZ

What is the type of address translation that is used to keep track of sessions initiated by a computer on a private network to a service on the Internet?

Due to the development of NAT-enabled devices and the capability to use private network IP address space, the stress of coming up with a new standard to replace IPv4 was lessened. This has allowed the life span of IPv4 to be extended and a gradual transition made to the newer IP address standard IPv6. Although current devices are IPv6-capable, they are still able to be installed and used within the IPv4 environment.

10.2.1.2 Internet Protocol Version 6

The real thrust of moving to IPv6 is the larger address space that it provides, with 128 bits dedicated to address space. The number is so large that it exceeds the national debt, which is pretty hard to do these days. Our scientific calculator claims it is 3.4028 e^{38}, give or take a few addresses. It is so great that each person alive on the face of the earth can have multiple devices using

IP addresses and there would still be addresses left over. Although these numbers are staggering, the real intent of IPv6 is to increase the efficiency of network management and routing. There is a high probability that only a small percentage of the address space will actually be used.

Figure 10-11 illustrates the IPv6 header, which is 40 bytes in total length.

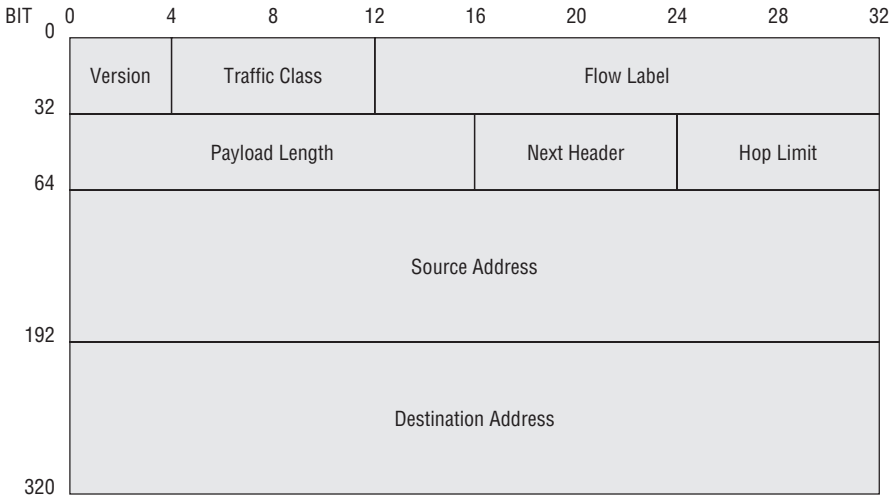

Figure 10-11 The IPv6 header

The first 4 bits are the Version field for IPv6. The next 8 bits are the Traffic Class field, which adds in the control options. The next 20 bits are the Flow Label field, allocated for QoS (quality of service). The Payload Length field indicates the packet length in bytes. When this field is set to 0, the packet contains a jumbo-sized payload. The Next Header field indicates which encapsulated protocol follows. These protocol values are compatible with the IPv4 protocol field values. The Hop Limit field replaces the TTL (time to live) field of IPv4. Both the Source Address and Destination Address fields contain 128 bits of address data. The standard sized payload can be 65,536 bytes, and if the option is set, the payload can be jumbo-sized.[21] All data fragmentation is controlled by the sending network node since routers will never fragment a packet. However, IPv6 sending network nodes are expected to use a technique known as *path*

RANDOM BONUS DEFINITION

encapsulating bridge — A bridge that encapsulates LAN frames.

[21] Techno-geeks like to use jargon and you may hear a variety of words used to describe an entity. A ''jumbo-sized'' payload merely refers to a payload that is very large.

MTU discovery (PMTUD) to determine the MTU that can be used to send packets over the network.

The address notation used in IPv6 is quite different from that of IPv4. IPv6 addresses consist of eight groups of four hexadecimal numbers, where each field is separated by a colon. For example:

```
113A:00AB:8900:0000:0000:7EA3:0034:3347
```

A shorthand notation would be to reduce the fields containing zeros to just a pair of colons (`::`). IPv6 address notation has a variety of rules that allow for various methods of displaying the same address. As IPv6 begins to be deployed more widely, there is sure to be a particular notation format that will become the more widely used and accepted format.

POP QUIZ

What is the difference between how IPv4 IP addresses are denoted and how IPv6 IP addresses are denoted?

10.2.2 Internet Control Message Protocol

Internet Control Message Protocol (ICMP) is an essential part of the TCP/IP protocol suite. It provides a means of messaging when a sent datagram is unable to be received by the intended network node. The `ping` and `traceroute` networking tools are also part of this protocol. ICMP error messages are generated in response to detected errors in the IP datagram, routing, or diagnostics. The ICMP protocol suite is part of IPv4, but there is an equivalent to ICMP that is a protocol within IPv6, referred to as ICMPv6. For the most part, computer users are unaware of network problems until a network error message is triggered by ICMP. When a user suspects that a network problem may exist, he or she can use the `ping` and `traceroute`[22] commands to aid in troubleshooting the problem.

10.2.2.1 Ping

The `ping` command within the Microsoft Windows operating system has the following syntax:

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

[22] `traceroute` is the normal command for many various operating systems. However, in the Microsoft Windows-based world, the actual command for `traceroute` is truncated to `tracert`.

```
C:\>ping

Usage: ping [-t] [-a] [-n count] [-l size] [-f] [-i TTL] [-v TOS]
            [-r count] [-s count] [[-j host-list] | [-k host-list]]
            [-w timeout] target_name

Options:
    -t             Ping the specified host until stopped.
                   To see statistics and continue - type Control-Break;
                   To stop - type Control-C.
    -a             Resolve addresses to hostnames.
    -n count       Number of echo requests to send.
    -l size        Send buffer size.
    -f             Set Don't Fragment flag in packet.
    -i TTL         Time To Live.
    -v TOS         Type Of Service.
    -r count       Record route for count hops.
    -s count       Timestamp for count hops.
    -j host-list   Loose source route along host-list.
    -k host-list   Strict source route along host-list.
    -w timeout     Timeout in milliseconds to wait for each reply.
```

The most common use of `ping` is to verify that a particular network is available over the network. In the following example, a computer has issued a `ping` command for its default gateway.

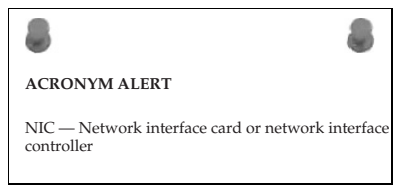


```
C:\>ping 192.168.2.1

Pinging 192.168.2.1 with 32 bytes of data:

Reply from 192.168.2.1: bytes=32 time<1ms TTL=64
Reply from 192.168.2.1: bytes=32 time<1ms TTL=64
Reply from 192.168.2.1: bytes=32 time<1ms TTL=64
Reply from 192.168.2.1: bytes=32 time<1ms TTL=64

Ping statistics for 192.168.2.1:
    Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
Approximate round trip times in milli-seconds:
    Minimum = 0ms, Maximum = 0ms, Average = 0ms
```

There are many ways to use the `ping` command for troubleshooting various network issues. Chapter 16, ''Troubleshooting,'' will go into greater detail about how this

POP QUIZ

What option would be used to modify the size of a `ping` packet?

tool can be used as an aid in determining what is causing certain network issues.

10.2.2.2 Traceroute

The `traceroute` command is used to trace the path from the sending network node to the receiving network node on a hop-to-hop basis. It reports back on each hop as it is traversed over the network on the path to the destination network node. The following is the syntax for the `traceroute` command for the Microsoft Windows operating system:

```
C:\>tracert

Usage: tracert [-d] [-h maximum_hops] [-j host-list] [-w timeout]
                          target_name

Options:
    -d                    Do not resolve addresses to hostnames.
    -h maximum_hops       Maximum number of hops to search for
                          target.
    -j host-list          Loose source route along host-list.
    -w timeout            Wait timeout milliseconds for each reply.
```

The options are self-explanatory. The `-d` option is often used to save the time that is required to resolve IP addresses to host names for each hop along the path. The following is an example of a successful completion of a `traceroute` command:

RANDOM BONUS DEFINITION

common and internal spanning tree — A collection of the internal spanning trees in a multiple spanning tree region, combined with the common spanning tree that connects MST regions to form a single spanning tree that ensures all LANs in the bridge network are fully connected and loop-free.

```
C:\>tracert www.richardbramante.com

Tracing route to www.richardbramante.com [68.180.151.74]
over a maximum of 30 hops:

  1    <1ms    <1ms    <1ms    192.168.2.1
  2    <1ms    <1ms    <1ms    192.168.0.1
  3    4ms     4ms     4ms      L100.VFTTP-12.BSTNMA.verizon-
                                gni.net [72.74.235.1]
  4    3ms      4ms     4ms     P4-1.LCR-04.BSTNMA.verizon-
                                gni.net [130.81.60.226]
  5    26ms     27ms    27ms    so-7-0-0-0.ASH-PEER-
```

```
                                        RTR2.verizon-gni.net
                                        [130.81.17.179]
  6    28 ms     27 ms     27 ms     130.81.14.98
  7    65 ms     64 ms     64 ms     so-2-0-0.pat2.dax.yahoo.com
                                        [216.115.96.21]
  8    91 ms     92 ms     92 ms     as1.pat2.pao.yahoo.com
                                        [216.115.101.130]
  9    92 ms     92 ms     92 ms     ae1-p151.msr2.sp1.yahoo.com
                                        [216.115.107.79]
 10   100 ms     92 ms     92 ms     ge-1-41.bas-b2.sp1.yahoo.com
                                        [209.131.32.33]
 11    92 ms     94 ms     92 ms      www.richardbramante.com
                                        [68.180.151.74]

Trace complete.
```

If a target node does not allow ICMP, a `traceroute` would not end normally and would appear as follows:

```
C:\>tracert www.wiley.com

Tracing route to www.wiley.com [208.215.179.146]
over a maximum of 30 hops:

  1            <1 ms     <1 ms    <1 ms       192.168.2.1
  2            <1 ms     <1 ms    <1 ms       192.168.0.1
  3    3 ms     4ms       4ms      L100.VFTTP-12.BSTNMA.verizon-
                                    gni.net   [72.74.235.1]
  4    5 ms     4ms       4ms      P4-1.LCR-04.BSTNMA.verizon-
                                    gni.net     [130.81.60.226]
  5             4 ms      4 ms     4 ms      130.81.29.170
  6    8 ms     7ms       7ms      0.so-1-0-0.XL2.BOS4.ALTER.NET
                                    [152.63.16.141]
  7    15ms      14ms      14ms     0.so-7-0-0.XL4.NYC4.ALTER.NET
                                    [152.63.17.97]
  8    13ms      14ms      14ms     0.ge-5-1-0.BR3.NYC4.ALTER.NET
                                    [152.63.3.118]
  9             16 ms     17 ms    17 ms  192.205.34.49
 10             18 ms     17 ms    17 ms      tbr1.n54ny.ip.att.net
                                    [12.122.105.14]
 11             16 ms     17 ms    17 ms      gar3.nw2nj.ip.att.net
                                    [12.122.105.49]
 12             18 ms     17 ms    19 ms  12.88.61.178
 13             *         *         *         Request timed out.
 14             *         *         *         Request timed out.
 15             *         *         *         Request timed out.
 16   ^C
```

The `traceroute` command was truncated with the Ctrl+C key combination to shorten the number of hops, as the command would have continued with ''Request timed out'' for the default number of 30 hops.

RANDOM BONUS DEFINITION

10BASE-T — A baseband Ethernet system operating at 10 Mbps over two pairs of Category 3 UTP cable.

For further testing, the `ping` command was then issued with the following results:

```
C:\>ping www.wiley.com

Pinging www.wiley.com [208.215.179.146] with 32 bytes of data:

Request timed out.
Request timed out.
Request timed out.
Request timed out.

Ping statistics for 208.215.179.146:
    Packets: Sent = 4, Received = 0, Lost = 4 (100% loss),
```

So, it first appears that the target network node is not available, but we know that is not true since our browser displays the following image for that network node address, as shown in Figure 10-12.

Although the `ping` and `traceroute` commands are very useful tools, they are not 100 percent accurate in predicting what is going on in the network. From the above indications it appears that the Wiley website is dropping ICMP packets.

POP QUIZ

What is the default maximum hop count for the `traceroute` command?

10.2.3 Internet Group Management Protocol

Internet Group Management Protocol (IGMP) is a protocol for handling multicast group memberships, which are required in situations with streaming video or multiplayer games. The protocol is used by the client computer to establish a connection to a local multicast[23] router. With the use of local and

[23]Multicast is when you have many users connected to a single service simultaneously. It is analogous to the broadcast of a radio or TV program.

multicast routers, these streaming applications are able to provide service to many multicast clients simultaneously.

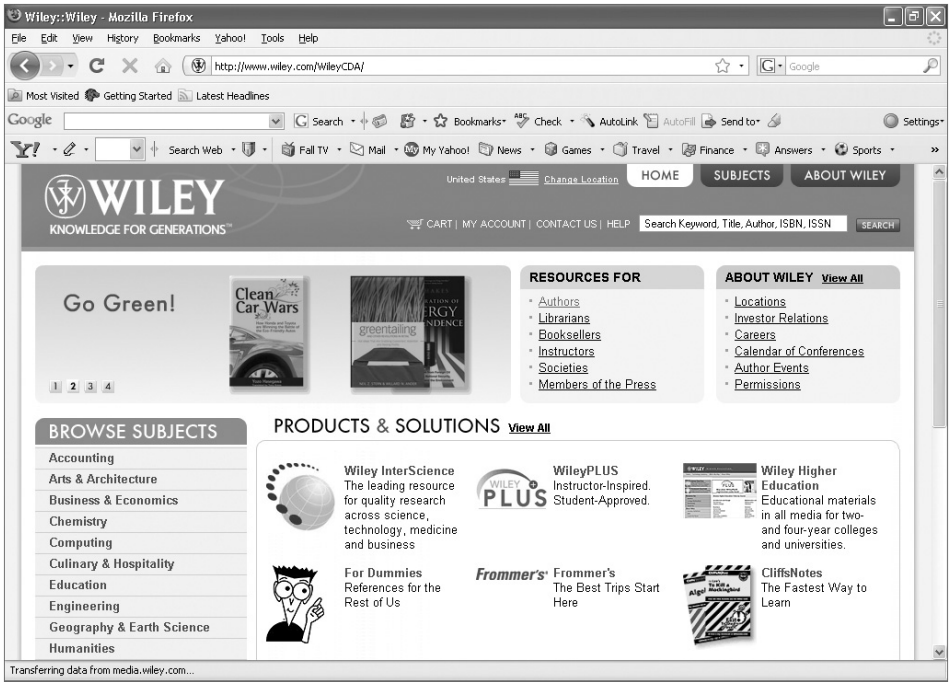

Figure 10-12 The Web page for `www.wiley.com`

IGMP-enabled routers check on the users within the group to determine if they have an active session. As long as there is an active group member, the router will continue to forward the multicast to that subnet. If all members are inactive, the packets are simply dropped.[24]

A computer running an IGMP-based application issues an IGMP report packet to join a group. When the router serving that group determines it has one member of the group, it would forward multicast packets to that subnet. Member computers need not inform the IGMP-enabled router when they leave the group. The IGMP-enabled router will perform member queries at fixed intervals to determine if there are any connected IGMP members. If there are, it continues to forward multicast packets to that subnet. The

ACRONYM ALERT

DSAP — Destination service access point

[24] "Dropped" is a techno-geek word for a packet not being forwarded. The router just discards the packet to work on the next packet that it receives.

reason for this behavior by the IGMP router is to prevent flooding subnets with multicast packets when there are no connected users from that subnet.

10.2.4 Internet Protocol Security

RANDOM BONUS DEFINITION

bit time — The length of time required to transmit 1 bit of information.

Internet Protocol Security (IPSec) uses authentication and encryption to establish a secure connection between endpoint network nodes. The terms *VPN* and *tunneling*[25] are used along with IPSec; however, IPSec is the means that permits the use of these capabilities. Tunnels between endpoint VPN devices normally are point-to-point and use a preshared key (PSK) as part of the authentication process. Once authenticated, the endpoints are able to pass traffic between them that is encapsulated using strong encryption[26] to prevent data from being compromised. Figure 10-13 illustrates the use of IPSec endpoint devices as well as IPSec client workstations establishing VPN network connections over the Internet.

ACRONYM ALERT

GARP — Generic Attribute Registration Protocol

In Figure 10-13, Network A and Network B are connected to VPN-enabled routers. These routers use IPSec to establish a peer-to-peer tunnel to allow data to flow between the private internal network of Network A and Network B. Peer-to-peer networks know each other's statically assigned IP address, and that is part of the security mechanism. The major component of safe data transfer is the use of preshared keys with strong encryption. Depending on the policies established on the VPN routers, users from one network can connect to resources on the remote network the VPN tunnel was established with. Another aspect for consideration when conceptualizing a VPN is determining the permissions that will be allowed for network users. Some users might need access to services on the Internet, whereas users might not require this as part of their jobs. VPN routers act as firewalls and are policy-intensive devices. The normal default state for these

[25] *Tunneling* is the term used to describe a virtual protected conduit between two endpoint network nodes. There is no way to actually ''build'' a real tunnel. The idea is that with strong encryption the packet is undecipherable; thus, it is as if the data stream is traveling within a protected tunnel, unseen to the rest of the Internet. In reality, packets can be snooped, but hacking the real information out of the packet is next to impossible.

[26] Encryption depends on key length. There are two predominant key lengths used within the Data Encryption Standard (DES): 56-bit, referred to as simply DES or single DES, and 128-bit, referred to as triple DES or 3DES.

devices is to allow only tunnel traffic to pass through from one VPN tunnel endpoint to another.

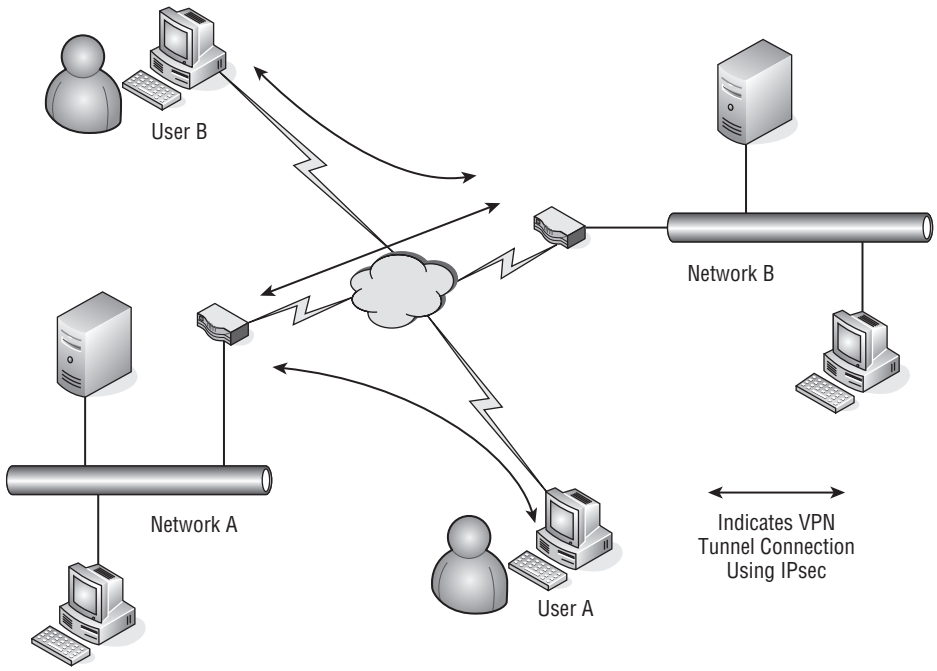

Figure 10-13 VPN networking using IPSec

Remote users using only a computer connected to an Internet access point require a client to be loaded on their PC to establish a VPN tunnel between the computer and a remote VPN router. Once a tunnel session is established, users can launch other applications, which will utilize the connection to gain access to the resources located on the private network protected by the VPN router. The VPN router is capable of setting user policies, either by user or group of users, to limit access to only some of the network's resources. It is also capable of denying remote users the capability to access the private network and then access the remote private network through the established peer-to-peer tunnel.

The use of IPSec to create VPNs using the Internet eliminates the need for direct point-to-point telecommunications between remote network nodes. This is a large cost savings over using directly connected dedicated lines between remote office locations. However, careful planning and thought needs to go

POP QUIZ

Name two components that help make VPNs safe and secure.

into the design of the network, and policies may have to be developed to secure the network from a security breach or the compromise of information as it travels over the Internet.

10.3 Chapter Exercises

1. Name the type of network service being used for each of the following:

   HTTP _____

   FTP _____

   Mail _____

   Telnet _____

2. A client/server application is considered to be what type of network service?
3. What is a TLD and can you name a few?
4. How is the MTU size determined?
5. What does NAT accomplish?
6. Name two network tools that can troubleshoot a network problem?

   _____

   _____

10.4 Pop Quiz Answers

1. Mail is what type of network service?

   Connectionless

2. Name the ports used by an FTP client to request an FTP session with an FTP server. Which port is used for data transmission?

   Ports 20 and 21. Port 20 is used for data.

3. Name some top level domain names.

   .com, .gov, .edu, .net

4. What is a maximum transmission unit?

   The maximum payload size that can be transmitted without the use of fragmentation.

5. What is the type of address translation that is used to keep track of sessions initiated by a computer on a private network to a service on the Internet?

   Port mapping

6. What is the difference between how IPv4 IP addresses are denoted and how IPv6 IP addresses are denoted?

   Dot-decimal notation versus hexadecimal numbers separated by colons.

7. What option would be used to modify the size of a `ping` packet?

   The `-l` option

8. What is the default maximum hop count for the `traceroute` command?

   30 hops

9. Name two components that help make VPNs safe and secure.

   Authentication and encryption