CHAPTER

8

The Upper Layers

Protocol is everything.
Francois Giuliani[1]

The above quote is truly succinct, a real economy of words. This quote is not only true at the United Nations but also is easily applied to the networking environment. When you think of the mix of various equipment, wiring, networking operating systems, computer operating systems, programs running on servers as multiuser platforms, programs running on local computer workstations (which includes pretty much anything a person can hang off a network segment), the ability to communicate is essential. The United Nations uses translators to ensure that all the representatives from the many varied nations can understand the procedures. A network protocol also acts as a translator between the many subcomponents that we lump together under the word ''network.''

We would hate to think what a General Assembly meeting of the United Nations would look and sound like without the translators they employ. There is only one word that comes to mind: chaos. How would you ever be able to get anything done? The same goes for networks, except things move much faster than the world's fastest talker can utter even a single word. So protocol is truly everything in the networking world.

[1]Francois Giuliani worked at the United Nations for 25 years. At the time of his departure in March 1996, he was the director of the Media Division of the Department of Public Information (DPI).

RANDOM BONUS DEFINITION

hardware address — Synonymous with MAC address, physical address, and unicast address.

This chapter investigates the upper layers of the OSI reference model: the Application layer, the Presentation layer, and the Session layer. We will identify the ''translators'' being used so that information can flow smoothly and without error between these layers and eventually be sent over the network media to another network node and the device servicing that node. This is a top-down approach where users attempt to interact with the device they are using to communicate with another device and/or users somewhere over the net.[2]

8.1 Background

Software programs use the upper layers of the OSI reference model to send and receive data over a network. Normally such programs are called applications and although they may interface with the Application layer of the OSI reference model, it does not necessarily need to be the case. In this chapter, ''application program'' and ''Application layer'' are not synonymous and refer to different aspects of computer usage.

A computer user purchases an application program and loads it on to his or her computer's hard drive. Basically, programs can be divided into two broad categories: locally run application programs and client/server-based application programs. As the name implies, a locally run application program executes program instructions and all data is maintained within the local computer, so there is never a need to utilize a network connection. A client/server application implies that a client computer and a server need to communicate if the application program is to run successfully. A client/server application in most cases requires a degree of interconnectivity for the application program to communicate with its counterpart server-based program. As this book is concerned with networking, the only application programs that have relevance are application programs that follow the client/server model. Figure 8-1 illustrates a client/server application program scenario.

As you can see in the figure, a client computer communicates over the network with a server. Although they are working in conjunction within a certain application program, they run within their own realms. The server listens on the network, awaiting requests from client computers. When the server receives a request from a client, it fulfills it. The communication between a particular client computer and the server is considered a *session*. Servers only respond to

[2]The ''net'' is in reference to any and all segments of a network, which can include in part or in whole any of the following: local network segment, the local LAN, intranet, or the Internet.

session requests in this environment; they do not initiate the start of session. Once a data transfer to or from the server is complete, it may request to terminate the session. Depending on the server application being run on the server, the server may be capable of maintaining a number of simultaneous sessions with multiple client computers. Server applications that can maintain multiple sessions are usually referred to as *multiuser applications*.

POP QUIZ

True or false: The Application layer is where all the application programs you load on your PC are stored.

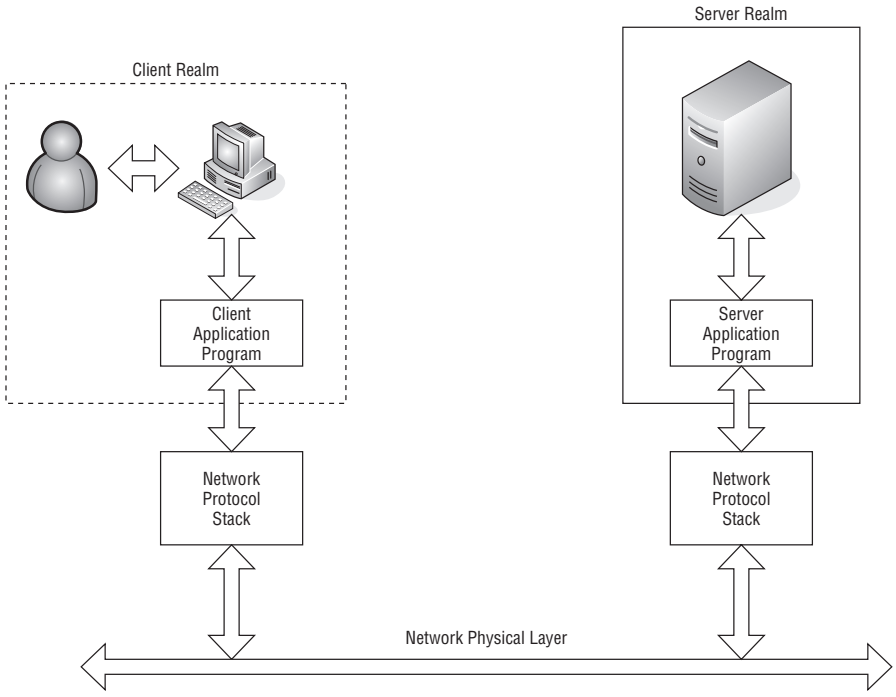

Figure 8-1 A client/server application

The client realm involves not only the client computer and application program, but a user as well. The user initiates requests to the client computer via an input device (usually a keyboard, mouse, or both). The application responds back to the user in graphic images or text displayed on a screen or tone signals played back through the computer's audio system. The application program requires user input in the form of commands and data in order for it to interact with the server application it is working in conjunction with in a particular client/server application.

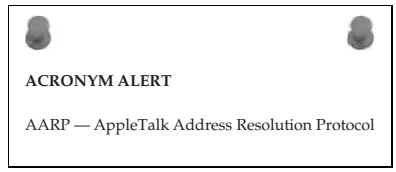

Although client/server applications work in conjunction with each other, they are autonomous until a session is established between a particular client application workstation and the application server. The server application, in most cases, is constantly running on a server that is rarely shut down. For instance, a mail server is always available to receive messages from client workstations, process them, and direct them to another mail server where the recipient of that message has an account. Received messages from other mail servers destined for users on a particular mail server are stored on the server until the mail server is queried by a user to see if there are any messages.

Mail servers or other application servers may also have to perform user authentication to ensure security and user privacy. An example of this would be when users launch a particular application on their client workstation, such as a mail reader. They may be first presented with a dialog message box to enter their user ID and password. Unbeknownst to the users, when they launched the client application it went out over the network and requested to establish a session with the server. The server at that point returned a response that security is required and requested that a user ID and password be provided for the connection to be established and maintained over the length of the session. Users at the client workstation enter their user ID and password, and if it matches the authentication parameters that the mail server is using for authentication, a mail session is opened between the client workstation and the mail server. The simple process of just logging on to a mail server requires interaction of the application program and the network stack[3] to ensure that messages are properly transmitted over the network between the client workstation and the server within a predetermined protocol.

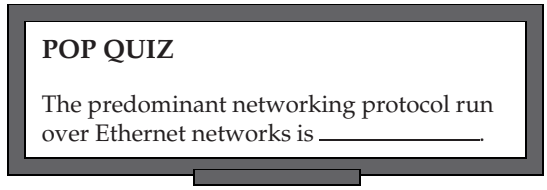

Since TCP/IP (Transmission Control Protocol/Internet Protocol) is the predominant network protocol in use within today's networking world, the remainder of this chapter will refer to the network stack in terms of how it relates to the TCP/IP protocol suite. Most, if not all, of today's computer operating systems provide a network stack that is compatible and easily interacts with applications that use TCP/IP to communicate over a network.

[3]Usually in reference to the OSI model, ''network stack'' or simply the ''stack'' refers to layers within the OSI reference model that, in most cases, have been embedded within the particular operating system running on the computer in use.

8.2 The TCP/IP Model

The TCP/IP model consists of four layers: an Application layer, a Transport layer, an Internet layer, and a Link layer. To accommodate a wide range of application programs that need to communicate over a network structure, encapsulation is performed between the layers to allow data to be moved independently of the application that produced the data. Figure 8-2 illustrates a conceptual view of the TCP/IP network stack.

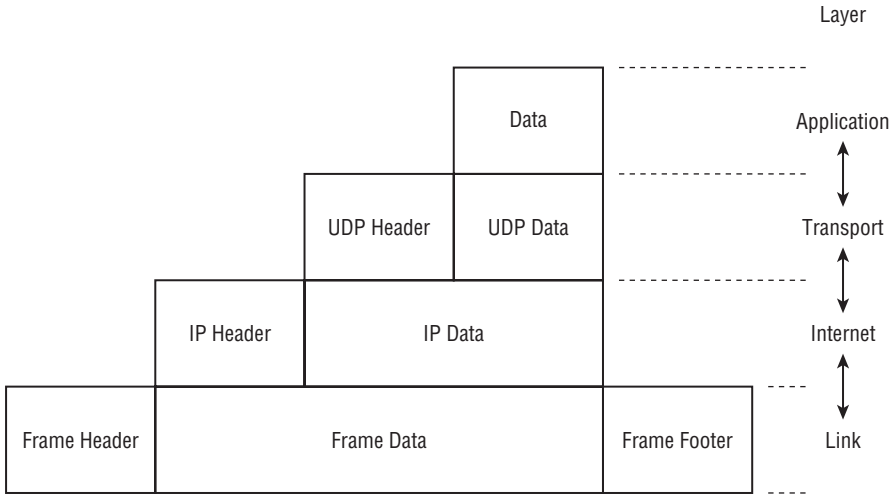

Figure 8-2 The TCP/IP network stack/model

The top level Application layer is the data portion of the network stack. It contains the upper level protocols that allow application programs to encapsulate data so that it can be passed down to the Transport layer. Since the OSI model Presentation layer and Session layer are combined with the OSI model Application layer to make up the TCP/IP network stack's Application layer, any protocols needed within the OSI model for these layers are accomplished via the use of libraries[4] within the TCP/IP model's Application layer.

The TCP/IP model Transport layer maps directly to the Layer 4 Transport layer of the OSI model, and the TCP/IP model Internet layer is usually mapped directly to the OSI model's Network layer. However, the TCP/IP model's Link layer covers both the OSI model's Physical layer and Data Link layer.

Application layer data is passed to the Transport layer, where a UDP header is applied and is framed with the data, as shown in Figure 8-3.

[4]Libraries are collections of protocol routines for various protocol functions.

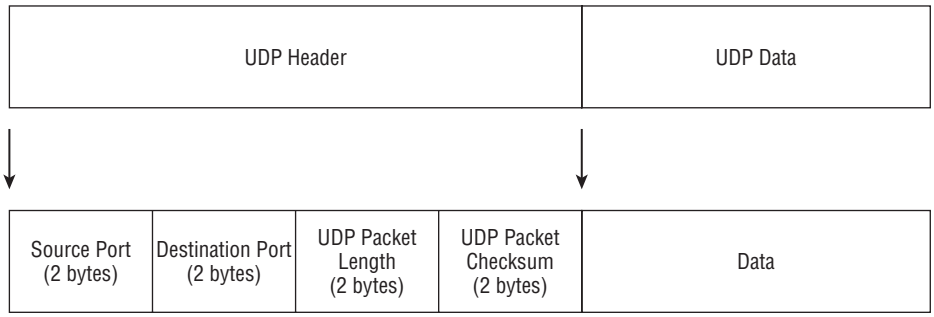

Figure 8-3 A UDP packet

As you can see, there is no address information other than the ports that that are being accessed. Since there is a lack of addressing and control, UDP is referred to as a *connectionless protocol*.[5] With 2 bytes allocated for both the source and destination port addresses, this accommodates up to 65,536 port numbers. However, the lower 1,024 port address values are reserved for defined services and are considered to be the *well-known port values*.[6]

The UDP Packet Length field is 2 bytes in length and contains the number of bytes of the whole packet, including header and data. The UDP Packet Checksum field is also 2 bytes in length and is the checksum of the whole packet, including header and data. Unlike TCP, the Checksum field is optional, which brings into question its use for packet transport over the network. The choice between using UDP and TCP depends on the transport mode selected by the application program developers. A deciding factor may be speed, since UDP does not require further encapsulation and the overall packet size is smaller than TCP by 12 bytes. On a single packet basis, this seems like a small price to pay; however, in applications where large amounts of data are transferred over the network, there can be noticeable performance differences. A software developer may choose not to use UDP where reliability of the transfer is required. UDP has no means of guaranteeing packet delivery. To guarantee delivery requires further encapsulation and the packet is then passed to the Internet layer of the TCP/IP network stack.

POP QUIZ

True or false: UDP is a connection-based protocol.

[5]A connectionless protocol means that packets are streamed onto the network without any relation to one another. There is no means to connect packets that may have been fragmented or to determine if packets have been received out of order.

[6]Well-known port addresses are reserved; however, the range above 1024 also has some predetermined services using a high-numbered port. An example would be radius server authentication using port 1812.

At the Internet layer, the UDP packet is encapsulated as data within the IP packet. Figure 8-4 illustrates the applied IP header.

Bit	0-3	4-7	8-15	16-18	19-31
0	Version	Header Length	Type of Service	Total Length	
32	Identification			Flags	Fragment Offset
64	Time to Live		Protocol	Header Checksum	
96	Source Address				
128	Destination Address				
160	Options				
160 or 192 +	**Data**				

Figure 8-4 The IP packet header

You can see that additional information is added to the packet that can affect its delivery over the network. The bit order of the packet delivery begins with bit position 0. Streaming from left to right across the header, the first field encountered is the Version field. Since this packet complies with IP version 4 (IPv4), the value contained in this field is 4.[7]

The next field is the Header Length of the IP header. The value contained in this field is the number of 32-bit words that are contained in the header. This value also indicates the bit position of where the Data field begins. The minimum value for this field is 5. So, in a header containing five 32-bit words, the start of data will begin at bit position 160 (5×32 bits $= 160$ bits). The beginning of the Data field will be pushed back an additional 32 bits if the Options field is present.

The Type of Service field was allocated to provide control over the packet's delivery priority. In the past, this field was not utilized; in recent days,

[7]Because this is a 4-bit binary field, the value in binary 4-bit notation would appear as 0100.

RANDOM BONUS DEFINITION

flow control — A function that prevents a sender of traffic from sending faster than the receiver is capable of receiving.

it has evolved into a Differentiated Services field (DiffServ). DiffServ provides a method of classifying network traffic for manageability and provides quality of service (QoS) guarantees across an IP network. This ability is essential for delivering time-sensitive packets for applications that require real-time performance. An example of a real-time application in wide use today is Voice over IP (VoIP).

The Total Length field contains the value in the number of bytes of the total length of the IP packet datagram, which also includes the header. The minimum value this field can contain is 20, which is the minimum number of bytes in an IP header without any data. Since this is a 16-bit field, the maximum amount of bytes in the datagram is restricted to a theoretical limit of 65,536 bytes. However, most networks do not permit the transfer of super-sized packets without fragmentation. The customary size restriction for TCP/IP on an Ethernet network is 1500 bytes. Larger packets would need to be fragmented and delivered reliably so they can be reconstructed on the receiving network node.

The next three fields, Identification, Flags, and Fragment Offset, are all used when fragmentation of a packet is required. A packet that is too large is broken into fragments, which are placed within a collection of packets to transfer the information within the original unfragmented packet. The Identification field is used to uniquely identify all IP packets that are fragments of a packet that needed to be fragmented before being placed on the network. The Flags field consists of 3 bits. The value of each field may either be a 0 or a 1, where 0 indicates ''no flag'' being present and 1 indicates ''flag bit set.'' In order of precedence, the most significant bit is reserved and always must be set to 1. The next bit is the do-not-fragment bit. When set, this bit signals that the packet is not to be fragmented. This can lead to packets being dropped if they exceed the overall packet size permitted by a receiving node. The only reason for use of the do-not-fragment flag is that the network node sending the packet knows that the network node that is to receive the packet does not have the capability to reassemble fragmented packets and sets the flag so upstream routers will not fragment the packet. The next flag bit is the more-frames bit, which indicates that more fragment packets are to follow this particular packet. The last packet containing a packet fragment segment will have this bit set to 0 to indicate that no other fragments are to follow this fragment. This bit is always set to 0 for all packets that don't contain fragmented packet segments.[8]

[8]If a packet does not contain fragmented packet segments, it is a packet unto itself and is considered an unfragmented packet. Whether to fragment a packet is determined by the amount of data that is be transmitted, since the header is for the most part of fixed length.

The Fragmentation Offset field contains the number of 8-byte blocks that the fragment data is offset from where it was located in the original unfragmented packet. The field is 13 bits long, so the maximum number of offset is 65,528.[9] Since the maximum packet size is fixed at 65,536, the values of the offset, plus the 20 bytes required for the IP header, is greater than the maximum size of a packet. Thirteen bytes are more than adequate for this field.

The Time to Live field is an 8-bit field that indicates how many seconds a packet can live on the Internet. With that many bits, it would equate to 255 seconds as a maximum or four and a quarter minutes. Imagine waiting more than four minutes per packet to see if they had arrived. Needless to say, the reason for the TTL timer is to prevent lost packets from traversing the Internet into infinity if they cannot find a home or until they end up being dropped somewhere along the way. These days this field is not used to display the amount of seconds but is a hop count.[10] As a packet travels across the Internet, each network forwarding device it passes through decrements the TTL field by one before forwarding the packet along to the next network hop. The packet will continue to travel until the packet with a TTL set to zero arrives at the input of a network forwarding device. When a packet with TTL equal to zero is received by a network forwarding device, it will simply not forward the packet and it is dropped.[11] When a packet is dropped, an ICMP (Internet Control Message Protocol) error is sent to the sender alerting it that the packet has been dropped. The typical message is that the TTL has been exceeded, which means the destination was not found. ICMP utilities include `ping` and `traceroute` and use error messages to allow a sender to know if a target address is reachable over the Internet.

The Protocol field is an 8-bit field used to indicate the protocol of the data portion of the IP packet. These are pre-assigned values maintained by the Internet Assigned Numbers Authority (IANA). Some of the most common protocols found in IP headers are a value of 1 for ICMP messages, a value of 6 for TCP messages, and a value of 17 for UDP messages.

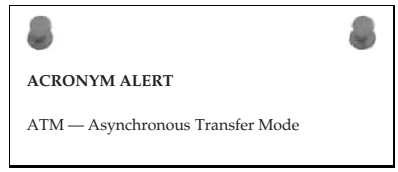

The Header Checksum field is a 16-bit field that contains the checksum of the header portion of the IP packet. The data portion carries the checksum of the protocol that is contained within it. When the packet is received, the checksum is calculated and compared to the value contained within the field. If

[9]This values is derived by $(2^{13}-1) \times 8$ bytes per block, or 65,528 bytes.
[10]*Hop count* is a method of counting the hops a packet traverses. As a packet is passed through a network forwarding device (e.g., a router), it is considered as a single hop.
[11]What is meant by ''dropped''? Simply that the packet is ignored and not forwarded or analyzed any further. It just ends up in the sky, where all lost packets go. However, network administrators always like to know why a packet is dropped.

there is a checksum mismatch, the packet is dropped. Since the header includes the Time to Live field, which is decremented each time the packet crosses a network hop, the header checksum will need to change if it is to remain valid at the next receiving network node. Because of known decrementing of the TTL field and the possibility that a network forwarding device may fragment the packet before passing it to the next network hop, each network forwarding device must insert the new valid checksum value in order to not create a checksum mismatch at the next receiving network node.

The Source Address field contains 32 bits of address information. The address is represented as four octets. Normally, IP addresses are annotated in what is called *dot-decimal notation*, such as:

```
192.168.16.1
```

Converting each octet into binary is represented as follows:

```
11000000.10101000.00010000.00000001
```

Binary address information in the Source Address field is represented as follows:

```
11000000101010000001000000000001
```

There are times when the source address of a packet is not the address of the sending network node. Various packet-forwarding network devices can perform a NAT function. Figure 8-5 illustrates a user workstation behind a router that is providing a NAT function.

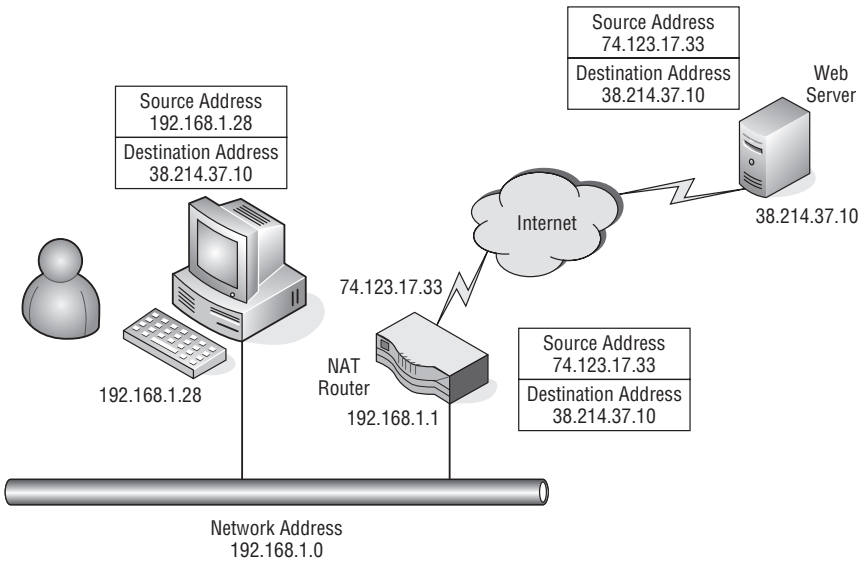

Figure 8-5 A private network behind a NAT router

In the figure, there is a private network[12] with a network address of 192.168.1.0, and on that network is a router with NAT capability of taking packets from a device on the 192.168.1.0 network and routing them out to the Internet. A user workstation at 192.168.1.28 wants to access a web page from a web server over the Internet at 38.214.37.10. Since the NAT router is the default gateway for the 192.168.1.0 network, all traffic that is not destined for the local LAN is sent to it. The user workstation in its TCP/IP settings has the default gateway address of 192.168.1.1, which is the NAT router's local network interface. The user workstation sends a request packet with a destination address of 38.214.37.10 with its own address of 192.168.1.28 in the Source Address field. Since the destination address is not on the local LAN, it is sent to the default gateway at 192.168.1.1.

The NAT router accepts the packet from the workstation at 192.168.1.28 and determines that it is destined to another network device over the Internet. The router replaces the user workstation's IP address with its own public interface[13] address in the Source Address field of the packet. After the address is replaced, it computes a new checksum for the header and inserts it into the checksum field before sending the packet out its public interface at 74.123.17.33.

RANDOM BONUS DEFINITION

Fast Ethernet — 100 Mbps Ethernet.

The packet is routed over the Internet and arrives at the web server residing at the public IP address of 38.214.37.10. The server determines that the request is destined for its address and notes that the source address is 74.123.17.33. The web server has no knowledge of the user workstation IP address of 192.168.1.28. The web server prepares a response using the public IP address of the NAT router as the destination address.

When the response packet arrives at the NAT router from the web server, it uses its NAT translation table to send the packet to the requesting workstation. It accomplishes this by modifying the destination address to the workstation address of 192.168.1.28 and computing a new checksum for the IP header before sending the packet out its private address interface onto the local LAN. For all intents and purposes, the user workstation believes it is interacting directly with the web server. NAT has some advantages and disadvantages, but for most small local networks it works well and offers

[12]Certain network addresses spaces have been determined by the Internet community to remain private. What this really means is that network forwarding devices on the Internet are not to forward any packet with a destination address that falls into the following ranges: 192.168.X.X, 172.16.X.X, and 10.X.X.X, where X denotes any number between 0 and 255.

[13]There are two sides to every router that interfaces a private local LAN network and the Internet. Normally, the interface that is accessible over the Internet is referred to as the *public interface* or *public interface address*.

protection against unsolicited network traffic ever making it through the NAT router to the local private network. If a packet's parameters do not match the translation table's known sessions, the packet is not processed and is dropped.

POP QUIZ

Describe what happens to a packet when it is passed through a NAT-enabled router.

The Destination Address field is pretty much self-explanatory. It is a 32-bit (4-byte) field containing the address information in the same format as the Source Address field. There is no difference in how the destination address is presented. In most circumstances the destination address is not messed with as the source address is with NAT. However, there are instances where the destination address may be translated and that is in special cases involving some sort of NAT router or a firewall. Actually, most routers used for the NAT function on outbound network traffic also have some capability to perform a port forwarding NAT. Notice that the web server in Figure 8-5 is directly connected to the Internet. That is certainly a possibility but is rarely found in today's networking environment because of possible attacks on the server via the Internet. Figure 8-6 illustrates a network that offers services available on the Internet but is protected and hidden from users.

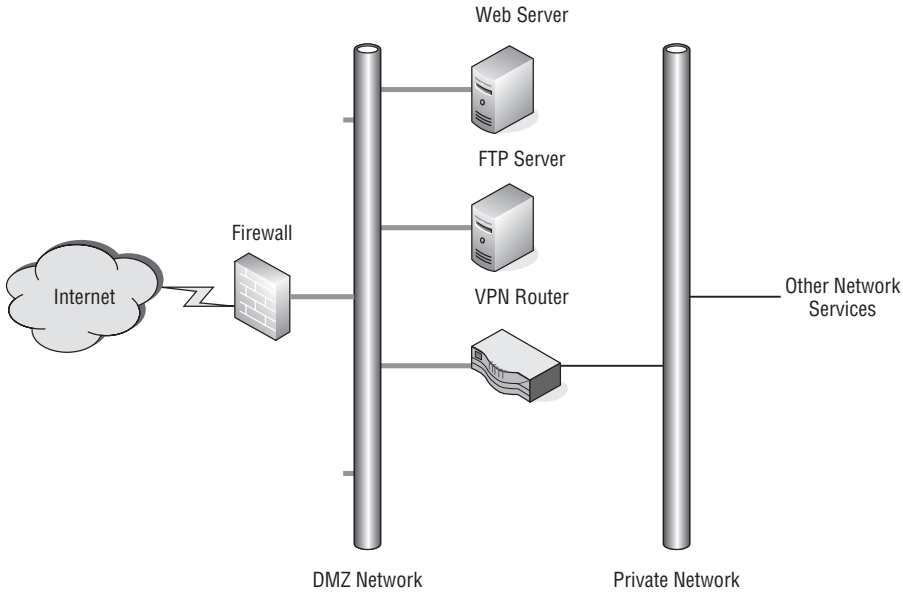

Figure 8-6 Port forwarding NAT

As you can see, the network located behind the firewall is shielded to prevent users on the Internet from accessing these services directly. A firewall may be

a network device that is designed as a firewall for the inspection of packets as they are received, or it may be a router running a firewall application on it that provides the packet inspection. In any case, the firewall function requires packet inspection and a determination by the policies put in place by the network administrators of what to do with the received packet. If a packet is received and does not match any of the existing policies, it is dropped.

The network behind the firewall may be a private network, but in this example it is shown as a DMZ[14] network. Connected to this network are services that the Internet community is permitted to reach. In this example, we have a web server, an FTP server, and a VPN router. Obviously, the web server is where web pages can be accessed and is generally used only for queries to obtain information. The FTP server may be only for file downloads but if allowed may also be a place where users can upload files. An example where users from the Internet community at large can upload files to an FTP server is a website that allows user posting on the site or a photo lab site that prints users' digital JPEG files on photographic paper.

In the figure, there is a VPN[15] router between the DMZ network and the private network. This device may be used as a remote access device for users who are remotely located but have permission to use the network service located on the private network. Usually VPN routers require user authentication, which can be performed locally on the VPN router, although it may depend on other authentication servers. For more information on this topic, see Chapter 14, ''Network Security.''

Back to our lovely red-brick firewall. We said that the firewall is responsible for inspecting the packets and using the policies installed by the network administrators to make a determination on what to do with the packet. To ensure that traffic is routed to the proper services, there must be port forwarding policies in place on the firewall. There are two ways this may be accomplished: either by changing the destination address and forwarding the packet on to the DMZ network, or, if the DMZ network addresses are routable Internet addresses, the packet may be inspected to ensure that only certain traffic is permitted to pass through the firewall. If the DMZ network uses addresses that are classified as nonroutable addresses, the only way traffic can be directed to the servers providing the requested services is by changing the packet's destination address. In this example, the web and FTP services

[14]DMZ is the acronym for *demilitarized zone*. In networking parlance, it refers to a network that may have some access by the public at large. The private network is protected by some sort of authentication process to only allow users with the proper credentials to reach the private network.

[15]VPN is the acronym for *virtual private network*. Usually the acronym is applied to the device, but in reality it is not the network in itself. It provides access to the network using security authentication and encryption processes to ensure that the private network is accessed only by those authorized to use its services.

only receive traffic for those particular services. Although these services are shown as separate computers, many services can be supplied by a single server running multiple protocols. In this example, packets directed to port 80 for web services would be directed to the web server, while packets using ports 20 and 21 would be directed to the FTP server. Lastly, VPN requests would be directed to the VPN router, and there are a few VPN protocols that may be used, so for now we will just say any VPN service requests will be directed to it.

The next field in the IP header is the Options field. As the name connotes, this is an optional field that follows the Destination Address field but is not used often. The last field in IP packet is the Data field, which is not part of the IP header so it is not used in the computation of the header checksum. The contents of the Data field are specified within the protocol header and can be any one of the IP protocols. Some of the most common protocols used in an IP packet are ICMP, TCP, UDP, and OSPF. OSPF (Open Shortest Path First) is a routing protocol used to route IP packets over the network.

The last layer of the TCP/IP Model is the Link layer. This is a combination of physical hardware and software to frame the IP packet to transport it over whatever network medium is being used. So frame information depends on the type of network connectivity that is being used. In the case of Ethernet, the IP packet is encapsulated within the Ethernet frame. Figure 8-7 shows Ethernet encapsulation of an IP packet.

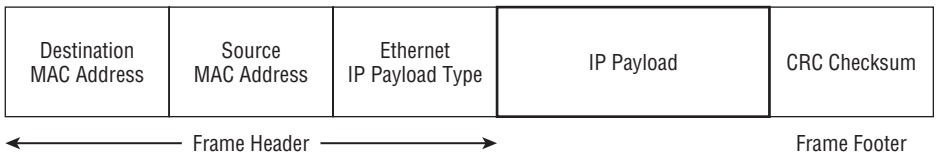

Figure 8-7 Ethernet encapsulation of an IP packet

The Ethernet frame header contains both the MAC (Media Access Control) destination and source addresses, each containing 12 bytes of addressing information. These addresses are unique and are directly associated with the physical network device. The last field in the Ethernet frame header is the Ethernet IP payload type. This is a 2-byte field and indicates the type of IP payload being transmitted by the Ethernet frame. Two of the most common IP payload types are 0x0800 for an IPv4 datagram and 0x0806[16] indicating that the frame is an ARP[17] (Address Resolution Protocol).

[16]The numeric representation with an ''x'' contained within it signifies that the number is a hexadecimal number. Each unit position is 4 binary bits in width. Thus, four hexadecimal numbers would contain 16 binary bits, or 2 bytes. If you still have difficulty grasping the concept of hexadecimal in relation to binary numbers, it is time for a review of number systems.

[17]ARP is a mechanism for a transmitting network node to determine which network node is associated with a particular IP address. The network node assigned that IP address responds with its MAC address.

POP QUIZ

At which layer of the TCP/IP model can the physical component of a network node be found?

The Ethernet frame footer contains the CRC checksum for the entire Ethernet frame. It contains 4 bytes of checksum data, which is used to validate that the frame was received correctly by the network node it was forwarded to. So, if the minimum size of an IP packet is 46 bytes, the minimum size if an Ethernet frame is 64 bytes, with the addition of the 18 bytes of Ethernet header and footer. The maximum size of an IP packet is 1500 bytes, which makes the maximum Ethernet frame allowed onto an Ethernet to be 1518 bytes in total. For large data payloads, fragmentation must be used.

We have worked our way down the TCP/IP model and now it is time to put the frame on the wire. Figure 8-8 conceptually illustrates the relationship between actual network elements and the TCP/IP network stack.

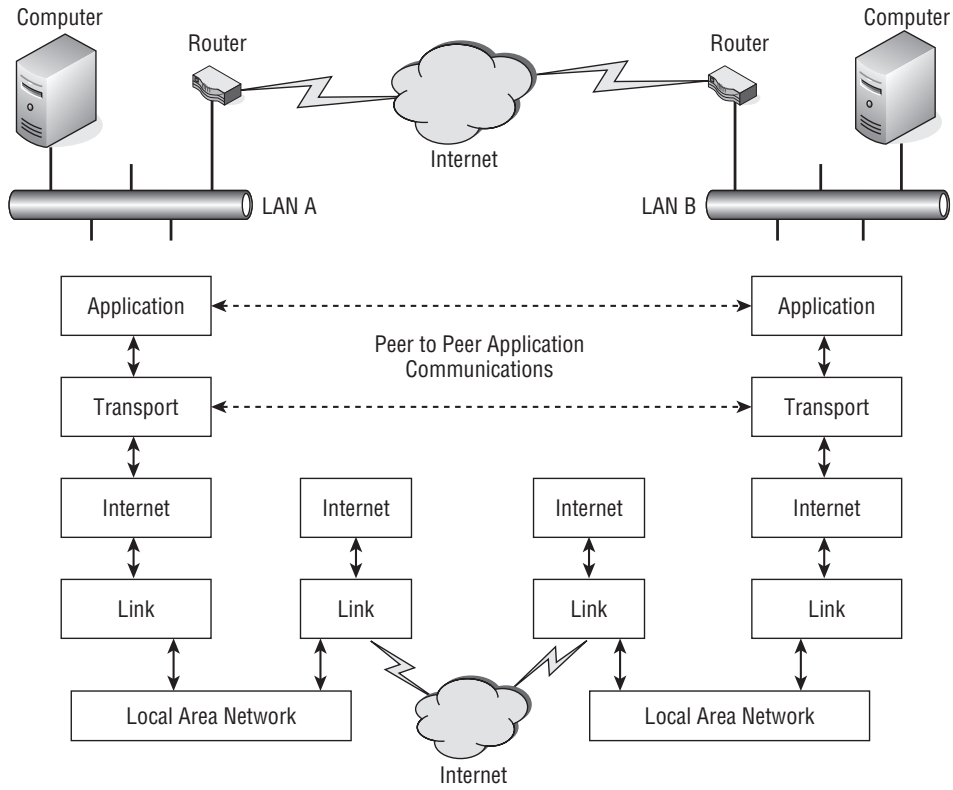

Figure 8-8 The relationship between network elements and the TCP/IP network stack

Two LANs, LAN A and LAN B, have computers that want to communicate with each other using an application program that supports their capability to establish a session and communicate effectively. This is shown as a computer and router connected to each LAN. Each router is connected to the Internet, shown as a cloud since there is an unknown amount of network devices that may be in the path between the router on LAN A and the router on LAN B. The assumption is that if a frame[18] is constructed properly, it can travel across many networks and through many devices in its path to reliably arrive at its predetermined destination.

The application program running on both computers may be aware of the other's network parameters, such as address and type of service, but it does not concern itself with the actual delivery of the data between the two peer computers running the application program. The only concern of the application program running within the TCP/IP's Application layer is preparing the data so it can present it to the Transport layer in anticipation of having the data delivered to the computer residing on the other LAN. So a peer-to-peer application session between two computers over the Internet appears as if they communicate with each other only using the Application layer and the Transport layer of TCP/IP model.

If application programs only concern themselves with getting the data properly packaged for the Transport layer, who does the rest of the actual delivery of the information? As illustrated in Figure 8-8, the lower two layers of the TCP/IP model are the Internet layer and the Link layer, which are directly responsible for reliably transporting the packet of information over the Internet. Since routing devices only need to be aware of addressing information, they only need to use the two lower layers of the model to effect the proper transmission of the information on its journey over the Internet. They are not concerned with data content since routing decisions are made on address and type of service.

The Internet and Link layers are normally part of the operating system and the hardware that is installed on the computer. If we assume an Ethernet-based LAN, then the computer would require an NIC that is capable of providing an Ethernet connection to the LAN. This is what would be Layer 1 or the Physical layer of the OSI reference model. However, it is a portion of the TCP/IP model

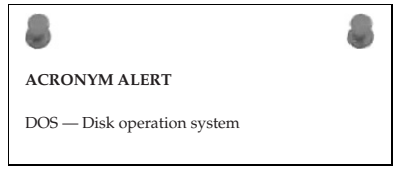

[18] *Frame* and *packet* are terms that are used interchangeably and are pretty much synonymous. Another term that may be tossed about from time to time is *datagram*. All these terms refer to some sort of encapsulation that includes the data to be transferred along with addressing and type of service being requested. It is how data can traverse the Internet from one computer to another.

Link layer. In order for the operating system to communicate properly with the NIC, device drivers are required that allow the software operating system to configure and control the physical components of the NIC. In a Microsoft Windows environment, this may be transparent to the user due to the capability of the operating system to recognize various pieces of computer hardware and automatically load the appropriate driver to communicate with the installed device. This portion of the TCP/IP model Link layer that includes device drivers maps to the Data Link layer of the OSI reference model.

Once an NIC is installed in a computer and the device drivers are loaded so that the operating system is able to communicate with the device on a physical level,[19] a network operating protocol needs to be bound to the card for it to communicate over the network with another network-connected device. In the case of TCP/IP, this is the address applied to the computer network interface along with its default gateway[20] and the location of at least one DNS server. Most operating systems allow these parameters to be set manually, or the computer requesting the values can apply them automatically from a DHCP server that is servicing that network segment.

POP QUIZ

What determines the type of framing that is to be used on a particular network segment?

Note that the routers illustrated in Figure 8-8 have their Link layers connected to both the LAN and the Internet. In reality these would be two different interfaces and also of differing types of network connectivity. More than likely the router will have an Ethernet interface allowing it to be interconnected to an Ethernet-based LAN. The interface to the Internet is dependent upon the type of service the router is connected to. It may be a point-to-point T1 interface, a FDDI interface, or some other form of high-speed service to the Internet. So a router's Link layer may consist of differing network hardware, device drivers, and Internet layer parameters to effectively transmit a data packet from the LAN to the Internet.

[19]Physical level kind of implies actual hardware but includes software that allows the hardware registers be written to for data and control. It is the device driver that makes the translation from hardware-specific elements to the standardized routines within the operating system controlling network-based communications.

[20]Default gateway has been mentioned more than once in this chapter. In a simple network, as illustrated in Figure 8-3, the address applied to the router on the LAN side would be considered to be a default gateway address. Basically, any packet with a destination address that is not located on the local LAN segment is forwarded to the address that is programmed into the default gateway address parameter field.

8.2.1 TCP/IP Application Layer

The Application layer of the TCP/IP model contains the upper level protocols of the TCP/IP protocol suite, such as FTP (File Transport Protocol) and SMTP (Simple Mail Transfer Protocol). Data is encapsulated and passed to the Transport Control Protocol for actual transmission on the network. The Application layer is dependent upon the lower layers to provide an effective and reliable means of network communications. The Application layer may be aware of the IP addresses and port numbers that are being used by the Transport layer, but it is that layer's responsibility to encapsulate this information as it is passed to the Internet layer below it. Some of the more common Application layer protocols are listed in Table 8-1.

RANDOM BONUS DEFINITION

AppleTalk — A protocol suite developed by Apple Computer.

8.2.2 TCP/IP Transport Layer

The two predominant protocols found in the TCP/IP Transport layer are UDP (User Datagram Protocol) and TCP (Transmission Control Protocol). The main difference between these protocols is that UDP does not guarantee delivery, and packets can arrive at the receiving network node out of order or duplicated, or not arrive at all. UDP is considered an unreliable delivery protocol whereas TCP is considered a reliable delivery protocol. TCP has the capability to detect missing, duplicated, and out of order packets and possesses mechanisms to request a packet be retransmitted if necessary. UDP relies on the use of ports for application-to-application communications. Since the port number is a 16-bit field in the UDP datagram, it can be anything between 0 and 65,535 or $(2^{16}-1)$.[21]

RANDOM BONUS DEFINITION

endpoint node — A node that interfaces with the user and the user's communication within a LAN.

Port numbers may range from 0 to 65,535, but for the most part the first 1024 (0 to 1023 decimal or 0x03FF hexadecimal) are considered to be the well-known ports. The ports from 1024 to 49,151 (0x0400 to 0xBFFF) are registered ports

[21] Why would the max port number would be $2^{16}-1$? True, the number 2 raised to the 16th power is equal to 65,536, so that is the maximum number of combinations that can be found when using 16 binary bits. However, one of those combinations is zero, so the −1 from the maximum value for the zero value and you end up the highest numeric value that can be attained with 16 binary bits is 65,535.

with the Internet Corporation for Assigned Names and Numbers (ICANN). Ports 49,152 to 65,535 (0xC000 to 0xFFFF) are considered to be temporary ports that clients can use when they communicate with servers.

Table 8-1 Common Application Layer Protocols

PROTOCOL	PORT(S)	DESCRIPTION
DHCP	67 and 68	Dynamic Host Configuration Protocol provides the means for network clients to obtain an IP address, default gateway IP address, and Domain Name System server addresses.
DNS	53	Domain Name System server requests are used to convert a host name to an IP address so it may be found on the Internet.
FTP	20 and 21	File Transfer Protocol is used to transfer files between an FTP client workstation and an FTP server. Port 20 is for data and port 21 is used for control signaling between server and client.
HTTP	80	Hypertext Transfer Protocol is used to transfer hypertext information over the Internet. The most familiar application use for hypertext information retrieval is a web browser.
IRC	194[22]	Internet Relay Chat is used for group communications over the Internet. Groups are referred to as channels and can also provide direct client-to-client chats and file transfers.
POP3	110	Post Office Protocol version 3 is used to retrieve mail from a mail server by a mail reader application program.
SMTP	25	Simple Mail Transport Protocol is used to send and receive mail messages between mail servers over the Internet.
SNMP	161	Simple Network Management Protocol is used to manage and monitor network devices over the local network and Internet.
Telnet	23	Telecommunications Network protocol is used over local networks and the Internet to establish terminal sessions between a client computer and a server.
NTP	123	Network Time Protocol is used to synchronize time on a network by synchronizing network devices to a time standard found on the local network or over the Internet,
BGP	179	Border Gateway Protocol is the main routing protocol of the Internet. It is responsible for maintaining a table of IP networks and makes routing decisions on path networking policies and rules.
RIP	520	Routing Information Protocol is routing protocol run on local network segments to advertise route gateway addresses within the local network.

[22]IRC runs on the de facto standard port of 6667 and other nearby ports in the range of 6665 to 6669.

HELPFUL HINT

As you'll recall from the discussion on Network Address Translation (NAT,) a device that has NAT capability keeps a translation table. The device uses its own public interface address as the source address, while maintaining a cross-reference to the actual address of the requesting workstation. A technique known as *port mapping* maps the hidden source address to an unused port number. A workstation that requests a page from a web server must access the server using port 80 for the server to respond to the request. When the server receives the request, its only concern is the destination port, which must be port 80. So, what the source port number is makes no difference when servicing the request. The server simply sends the packet back to the requesting IP address, even though it is of a NAT-enabled router and not the actual workstation making the request. When the packet arrives at the NAT-enabled router, it examines the packet and finds that the destination port address correlates to a workstation on its private LAN in its NAT translation table. It modifies the packet with a new destination IP and port address, recalculates a new checksum, and then transmits it on to the private LAN. Therefore, knowing those temporary port addresses are available can come in handy when you're using NAT.

Port 0 is normally reserved, but its use is allowed as a valid source port in transmissions where the transmitting network node does not require a response from the receiving network node, which would be true in a case of a streaming application. Some common UDP network applications that are considered streaming applications are video teleconferencing, gaming, telephone using voice over IP (VoIP), and Trivial File Transfer Protocol (TFTP). Domain Name Services (DNS), an essential component of the Internet for the resolution of IP addresses to domain names, also utilizes UDP for its transmissions.

Whereas UDP is connectionless, TCP is considered a connection-oriented protocol. This means that an end-to-end communication is required with the use of handshaking between client and server. Once the connection is established between the client and server, data can flow across that connection. Servers provide a multitude of services, including web, FTP, and Telnet.

TCP utilizes a three-way handshake in establishing a connection. The server first must bind to a particular service and be available to all connections. This listening on a port is considered to be a passive open. Establishing a connection requires an active open on the server port. To do this the client sends a SYN (synchronization) packet with a random packet sequence number to the server. In response the client's SYN the server replies with a SYN-ACK (acknowledgment) with the initial sequence number received from the client but incremented by one for the next sequence number it is expecting

to receive. Also in the packet is the server's initial sequence number. The client then replies back to the server an ACK that contains its initial sequence number incremented by one along with the server's acknowledgment number, which is the server's sequence number incremented by one. After a successful SYN, SYN-ACK, ACK sequence between client and server, a connection[23] is established.

With the use of sequence numbers, it is very easy to determine packet order, duplicate packets, or missing packets. This provides TCP with the capability to provide error-free transmission. Applications requiring a high degree of reliability work best when they use TCP to set up communications over the network between a client and a server running that application program.

HELPFUL HINT

This section noted that certain applications utilize UDP for their transmission of data. An example of this is VoIP. However, telephone conversations are somewhat forgiving for lost audio packets. Voice quality can degrade rapidly when packet loss begins to increase. Depending on bandwidth usage on networks and with the addition of quality of service (QoS) for some traffic, UDP traffic may be affected because of its best-effort delivery method. With VoIP, this is manifested in choppy voice quality and dead air, which some users find intolerable. One way around this issue is further encapsulation, although it does add a degree of overhead to each packet. Some users opt for sending their VoIP data through a tunneling protocol, which is delivered using TCP/IP.

POP QUIZ

Which TCP/IP model Transport layer protocol is connection based?

To terminate a TCP connection, the protocol uses a FIN, ACK sequence. When a network node desires to terminate the connection it sends a FIN packet, and the receiving network node sends an ACK in acknowledgment of receiving the FIN packet. This is considered a half open connection. The network node that has terminated its connection can no longer use the connection for data transmission, but the network node that has not sent its FIN packet can remain open and transmitting data. This sequence of FIN, ACK, FIN, ACK from both nodes is termed a four-way handshake sequence.

Perhaps the most commonly used connection termination sequence is one network node sends a FIN packet and the other network node responds with a FIN-ACK combining the two handshakes into one. The network node that

[23] *Connection* is sometimes synonymous with the word *session*, as in client server session. These words are sometimes used interchangeably to represent the SYN, SYN-ACK, ACK sequence.

initiated the termination sequence just responds with an ACK. This type of termination sequence is considered a three-way handshake.

There is a possibility that both network nodes may send a FIN packet simultaneously and also will send their ACK packets at the same time. Since this sequence is done in parallel it is considered a two-termination sequence.

8.2.3 TCP/IP Internet Layer

Some of the common services found at the Internet layer of the TCP/IP model are IP (Internet Protocol), ICMP (Internet Control Message Protocol), and IPSec (Internet Protocol Security). The primary protocol of the Internet layer suite of protocols is IP. Its main purpose is the delivery of packets between network nodes based solely on source and destination addresses since it is a connectionless protocol. Data from the upper layers is encapsulated within the IP datagram for delivery. IP is a best-effort delivery method and has no provision for out of order, duplicate, or missing packets. IP does not guarantee that the data payload has not been corrupted since the checksum it carries is only for the header, ensuring that it is error free. However, this does allow for quick discarding of packets whose headers have been corrupted.

IP is responsible for fragmentation into multiple packets if the data load it receives from the upper layers is too large to send within a single packet. When fragmentation is involved, the IP layer uses flags and offset to aid in the determination of packet sequence and their order. However, IP depends on the upper layers to ensure that the end-to-end integrity of the connection is maintained.

POP QUIZ

True or false: The TCP/IP model Internet layer IP protocol is a connectionless protocol.

ICMP is another integral protocol of the Internet layer. Its chief responsibility is to send a message to the operating system of a computer when a network error has been detected. These messages usually report that a requested service is not available or the other host could not be reached. Normally ICMP is a single-ended protocol since it not used to transmit messages between network nodes. However, there are some exceptions and the most common of these are

RANDOM BONUS DEFINITION

bottleneck — A point in a data communications path or computer processing flow that limits overall throughput or performance.

the `ping` and `traceroute`[24] commands. These two tools require a reply from a receiving network node. If no reply is received, an error message is displayed.

The `ping` utility is used to determine if a target network node is available over the network. If it replies, the assumption[25] is that the path is good between two network nodes. The `traceroute` utility returns replies from each hop that it crosses to reach a particular targeted network node. Usually, it will try to reach a target in a given number of hops. The customary maximum hop count is 30 hops. It is a good indication if the packet is traveling in the right direction or not.

POP QUIZ

What two ICMP applications can be used to verify the presence of an IP address on the Internet or local network?

8.2.4 TCP/IP Link Layer

We already mentioned that the TCP/IP model's Link layer maps to the OSI model's Data Link layer and Physical layer. The Physical layer components are the tangible pieces of hardware required to connect a computer to the network. It consists of the cabling, connectors, and NIC, which in most cases is installed in the computer. The hardware pieces are the lowest level of the TCP/IP model and make up the first level of the OSI model.

Normally we do not think of hardware in terms of protocols. However, there are standards and specifications that hardware from different manufacturers must meet to be considered compliant with a standard. An example of this would be the electrical characteristics of cabling used for networking. There are also mechanical considerations such as size and form factor. The interconnection world is large, and manufacturers from all over the globe produce various components that all need to interconnect with products from other networking products manufacturers. So the protocol of the Physical layer is the standards and specifications that define various networking components.

However, we know that the demarcation line between the Physical layer and the Data Link layer of the OSI model is at the Link layer of the TCP/IP model. It is the Network Interface Card.

[24] `traceroute` is found mostly in Unix-based systems. In the Microsoft Windows world, the command is `tracert`. This is an accommodation to its predecessor MS-DOS, since commands and filenames could not be longer than eight characters.

[25] The word ''assumption'' is used here since the fact that a reply is received does not guarantee that the host you desire to reach is actually the host that is replying. There is always a possibility of a duplicate address on a network. You will read more about this in Chapter 16, ''Troubleshooting.''

An NIC card is a piece of hardware with electrical capabilities of sending intelligent electrical signals to another NIC card on the same network. The intelligence is contained within the bits and order that it places over the network medium, which in a lot of cases is wire based but may also be either fiber or air, in the case of wireless networking. The NIC contains registers and buffer space where the data and network control signals from the computer operating system are written to while sending packets to or reading packets from the network medium. Figure 8-9 shows a block diagram of a generic NIC.

The diagram in Figure 8-9 is a representation of the basics of any type of NIC card. It is drawn to indicate that the card is capable of full-duplex operation because it contains both send and receive paths that are independent from one another, which would allow for simultaneous receive and transmit capability. To send a frame, the computer operating system needs to communicate with the card. Since these cards are functionally the same, the method used to communicate with a network interface is fixed by the operating system's developer. It is up to the card manufacturer to either manufacture the card so it can be installed in a computer using generic N driver software or provide a tailored driver that would perform this function. Hardware interface software drivers[26] are the link between operating system and the actual network hardware.

Reviewing the block diagram, the computer bus interface component has to adhere to the architecture of the bus structure used within the computer. There have been many bus structures used since the spawning of PCs. In the earlier days, many were proprietary designs. As the industry evolved so did bus standards. One of the earlier standards was S-100, and cards of this type can be found in computer museums and in the cellars of computer aficionados. With IBM's development of the IBM-PC, the bus standard that was rapidly adopted was ISA (Industry Standard Architecture). As computer capabilities began to expand so too did the bus architecture. The next evolution of the bus was the Extended ISA card or, simply, EISA card. Today's bus standard is PCI (Peripheral Component Interconnect). So a network card or any sort of peripheral card needs the capability to be inserted into the internal bus of the computer it is being installed in.

[26] *Device driver* is the common name for software that performs the hardware interface to the operating system. It is a piece of software code that allows the addressing and control of a hardware card installed in a computer.

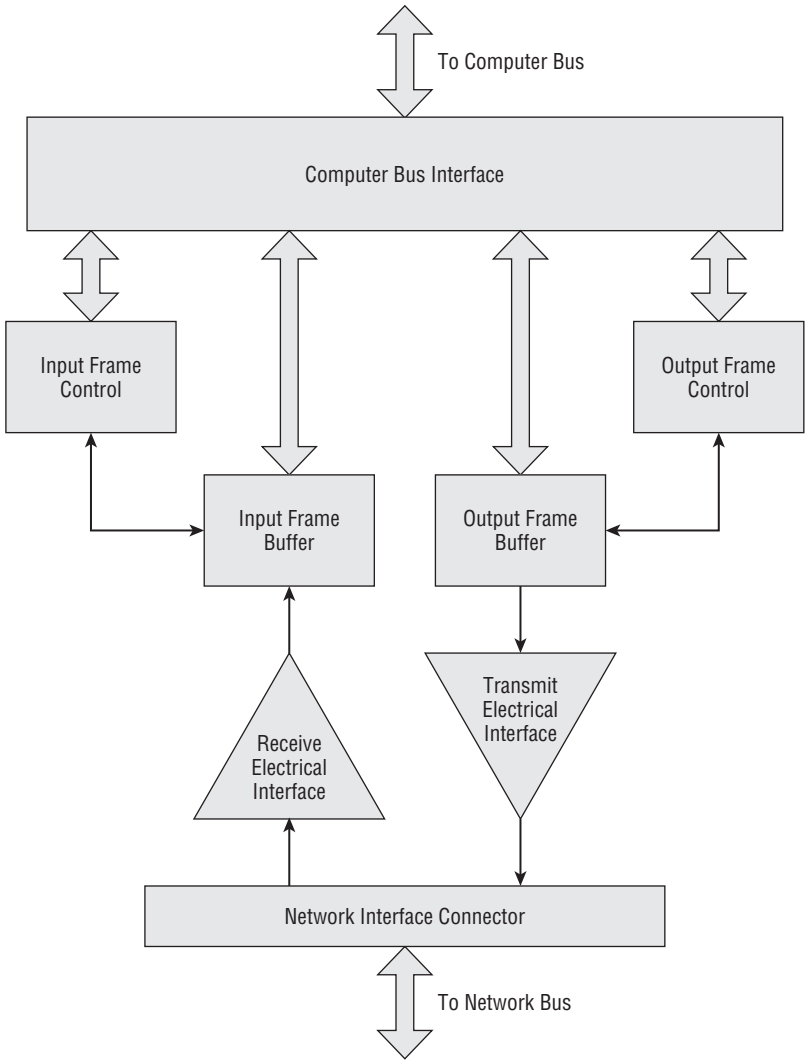

Figure 8-9 A block diagram of a generic NIC

With the network card installed in the computer chassis and the appropriate device driver installed into the computer so that the operating system knows how to communicate with the NIC, the next step is to bind a network protocol to the card so data can be moved to and from the network. Depending on the operating system, differing methods can be used; consult your computer documentation. When all of that is completed, data can be sent and received from the network transparent to the workstation's user.

RANDOM BONUS DEFINITION

collision domain — A set of nodes connecting to a shared medium among which a collision can occur. Stations on the same shared LAN are in the same collision domain.

Outgoing packets from an application program flow down the network stack with each layer encapsulated within the proper protocol. Once the frame that is to be transmitted is assembled and loaded into the output frame buffer, the output frame control prevents any further packets from being written into the output buffer until the frame has been completely sent. When the output frame buffer is cleared, the output frame control (through the device driver associated with this card) alerts the operating system that the card is ready to transmit another frame. On the receive side, the card monitors the network medium. When it has received a frame and it is completely in the input frame buffer and passes the checksum validation, the operating system is alerted (again via the device driver for the card) that a frame is ready to be passed up the network stack. As the packet passes through each layer, it is verified and checked as it is de-encapsulated. The input frame control is alerted that the frame is read and that another frame can be received.

POP QUIZ

List what is required for a network card to have full-duplex capability.

The last component to be discussed from the block diagram of the NIC is the connector. Many people are already familiar with the UTP RJ-45 connectors and plugs that are fairly commonplace on PCs, hubs, switches, and routers. However, depending on the medium being used, the connector will be different and adhere to the standards governing the usage of that type of medium.

8.2.4.1 TCP/IP Link Layer Protocols

The three common protocols residing at the Link layer of the TCP/IP model are ARP (Address Resolution Protocol), RARP (Reverse Address Resolution Protocol), and OSPF (Open Shortest Path First). ARP and RARP are the complement of each other in resolving network addresses. ARP is used to find what hardware MAC address is associated with a particular IP address. It accomplishes this by sending out an ARP request packet as a broadcast to all nodes on its local network segment. The packet contains the IP address that the transmitting network node is seeking. The receiving nodes on the network that do not have the IP address being requested simply ignore the packet. The

network node that does have that IP address bound to its network interface responds with its MAC hardware address.

RARP is a protocol that attempts to determine its IP address by broadcasting on the local network segment with its MAC address. It expects a receiving network node to have an entry in its ARP cache that matches that MAC address with an IP address to transmit back a packet containing the IP address. With DHCP now in wide use, RARP has fallen into disuse. However, DHCP is a TCP/IP model Application layer protocol and does not reside at the Link layer.

POP QUIZ

What is ARP used for?

OSPF is a dynamic routing protocol used to move packets from network segment to network segment. Two network segments with a router in each that have a path between them can build and interchange route information. Figure 8-10 illustrates a network utilizing OSPF to pass network routing information.

ACRONYM ALERT

RSTP — Rapid Spanning Tree Protocol

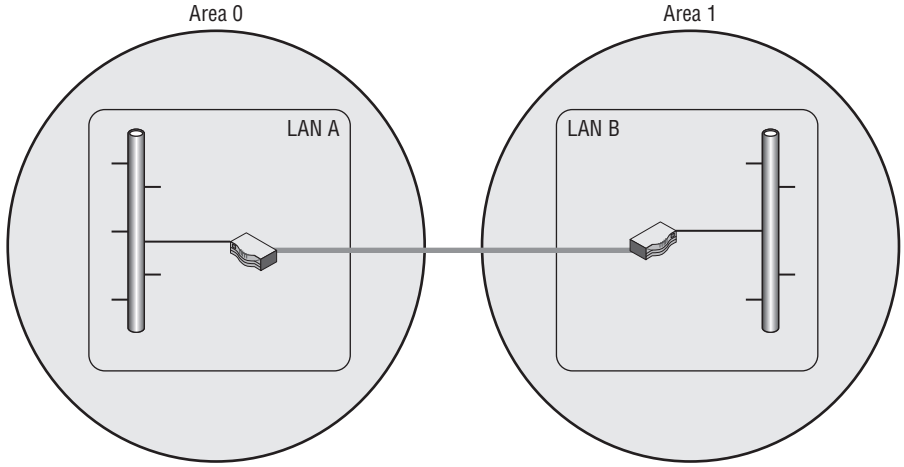

Figure 8-10 OSPF passing network routing information

Notice there are two areas: Area 0 and Area 1. An area is a collection of network segments with routers and other network forwarding devices. For the sake of simplicity, these are shown as two large circles. Within each area there is a router to route traffic from that area to another area. Routers that border a network and pass routing information to another router within another area

POP QUIZ

What is OSPF?

are called area border routers (ABRs). You will recall in the earlier discussion of routers in this chapter we said they resided within the lower level of the TCP/IP model. The OSPF information passed between routers is used to update their routing information tables. The two routers only communicate OSPF information between them and do not pass that information into the network they control. So if a workstation on LAN A wanted to pass data to a server or another workstation on LAN B, it would send the packet to its default gateway. The packet will ultimately end up at the ABR for Area 0 and finding that the targeted address when compared to its learned routes in its routing table is destined for network node in Area 1, forwards the adjacent ABR for Area 1. The information that is used is the Link State Database (LSDB) routing data that is passed between the Area 0 and Area 1 ABR router.

HELPFUL HINT

The OSPF example used is very simplistic. Large networks have multiple areas where one ABR may be interconnected to many other ABR routers. The key to OSPF is to know that the updates exchanged between routers can be found within the router's LSDB. Since this is a dynamic routing protocol, routes may pop up or age out as network nodes are inserted or removed from the network.

8.3 OSI Application Layer

The OSI Application layer resides at Layer 7 at the top of the OSI model. It was mentioned that the TCP/IP Application model directly links to this layer. So the protocols listed in the discussion of the TCP/IP model Application layer are also contained within this layer of the OSI model. This is the layer that is directly responsible for interfacing with the application program a user is using on the computer. The most common use of a computer with Internet access is e-mail. The e-mail protocols residing at this layer are POP (Post Office Protocol), POP3 (Post Office Protocol version 3), and SMTP (Simple Mail Transfer Protocol). POP and POP3 are mail client–based in the form of user e-mail reader programs. SMTP is

e-mail server–based and is used to transfer mail from one mail server to another, so this layer is keenly aware of its communication peers. Mail clients know where their mail server is, and mail servers can establish a connection for the transfer of mail between them. Using the example of e-mail at the Application layer, the information the layer is concerned with is the identity of the sender and the identity of the recipient of that message and what application is available to assist in preparing the message to be sent. All e-mail users are pretty familiar with the address format used, e.g., john.doe@his_company.com.

There are two parts to the recipient address: the user name ''john.doe'' and the domain name ''his_company.com''. The e-mail is formatted with sender address, recipient address, and message and passed on to the local mail server servicing that sender. The mail server is concerned with both the domain name portion of the recipient's address and the recipient's name. The recipient's name is used to identify the local mailbox for that user on the server. The application on the mail server is designed to use SMTP to send and receive mail from other mail servers. Most mail servers run a local post office where local users communicate locally over the local network using either POP or POP3 to retrieve mail from the local mail server. To send mail, users direct their outgoing messages to the SMTP service running on the mail server. Mail clients run POP at the Application layer to read mail and use SMTP to send mail. A mail server also runs two protocols at that layer, SMTP and POP and/or POP3. These protocols rely on the layers below them to actually get the message delivered and alert them when there is a message to pass up from the network.

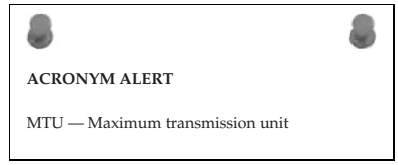

The Application layer is concerned with any syntax restraints such as the ''@'' sign in an e-mail message being required as a delimiter between recipient address and domain address. It is also the layer where security is applied for user identification and privacy. If quality of service is being applied to network communications, this is the layer concerned with determining the priority of a packet by its QoS[27] tagging.

NOTE Although there are many devices that are capable of QOS tagging of packets, there is no support for it over the Internet. The Internet is still a best-effort network.

[27]QoS is the acronym for quality of service. We mentioned that the DiffServ field or the Type of Service field of the IP header is used for tagging packets to allow them to be transmitted along the network with a priority determined by how they are tagged.

This chapter covers only a handful of the most familiar Layer 7 protocols. Many more protocols are available, considering that the combination for port numbers is 65,536. Even with some protocols using more than one port, there is still a lot of them. You can obtain information on many protocols by reading their RFCs. RFCs are available over the Internet at `www.ietf.org/rfc.html`.

POP QUIZ

True or false: The maximum number of protocols the TCP/IP Application layer can have at any one time is two.

8.4 OSI Presentation Layer

The middle layer of the OSI model upper layers is the Presentation layer, which occupies Layer 6 of that model. It has been mentioned that within the TCP/IP model, this OSI layer resides within its top Application layer. In the OSI model, it takes service requests from the Application layer and then issues requests to the Session layer below.

Although we said that this layer resides within the TCP/IP model Application layer, its components are more likely to be found within the computer's operating system. Within this layer, incoming and outgoing data can be translated from one data format to another. This layer also offers the capability for data encryption and compression as well as decrypting and uncompressing data received.

POP QUIZ

True or false: The OSI model Presentation layer maps directly to the Transport layer of the TCP/IP model.

8.5 OSI Session Layer

The lowest layer of the upper layers of the OSI model is Layer 5, the Session layer. Like the OSI model's Application and Presentation layers, it too can be found within the Application layer of the TCP/IP model. True to its name, it is the layer that is responsible for opening, managing, and closing a session between applications. It also provides the capability of restoring a session. It is the layer where authentication and permissions are granted.

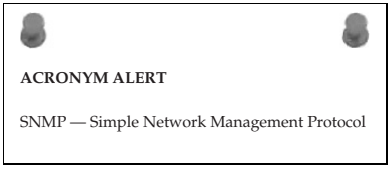

The Session layer is where TCP SYN handshake sequences are provided for. Although the Session layer is responsible for checkpointing and recovery within the OSI model, it is seldom used by protocols of the Internet Protocol suite. Some of the protocols found within the Session layer are

- **L2F** (Layer 2 Forwarding Protocol) — Used to provide virtual private networks (VPN) over the Internet.
- **L2TP** (Layer 2 Tunneling Protocol) — Used to provide virtual private networks (VPN) over the Internet.
- **NetBIOS** (Network Basic Input/Output System) — In today's networks is usually run over TCP/IP on the local network. It is a naming convention used to identify hosts on a Windows-based network. Although it is run over TCP/IP, its host name is not to be confused with the host domain name a computer may be given to resolve its name to an IP address. Those host names are registered with a DNS server and are not associated at all with a computer's NetBIOS host name, which on larger networks is resolved by a WINS (Windows Internet Name Service) server. In small networks where WINS may not be available, WINS name resolution can be accomplished by editing the `LMHOSTS` file on the computer to correlate the NetBIOS name to an IP address.
- **PAP** (Password Authentication Protocol) — A simple authentication protocol to allow users access to network services. A major drawback to PAP is that passwords are passed in cleartext and can be easily captured. Since PAP is not secure, network administrators have been making use of CHAP (Challenge Handshake Authentication Protocol), which uses a hashing function to secure the password. MS-CHAP is Microsoft's implementation of CHAP.
- **PPTP** (Point-to-Point Tunneling Protocol) — Provides a means of creating a VPN over the Internet. PPTP uses a standard PPP (Point-to-Point Protocol) session to its peer endpoint using the Generic Routing Encapsulation (GRE) protocol. A second session is then opened using TCP port 1723 to initiate and control the GTE session. Due to the need to have two simultaneous sessions opened, PPTP is not easily passed through a firewall. PPTP has lost favor and is being replaced by the L2TP and IPSec tunneling protocols.

RANDOM BONUS DEFINITION

catenet — A collection of networks connected together at the Data Link layer level.

- **SSH** (Secure Shell) — Allows for the secure exchange of data between two network nodes. It was designed as a replacement for Telnet and

other insecure protocols that were used for remote access over the Internet. These shells sent communications in cleartext, and passwords were easily compromised. SSH makes use of public key cryptography for authentication of the remote computer and allows the remote computer to also authenticate the user establishing the session.

The Session layer provides for either half-duplex or full-duplex operation, synchronization points in the message stream, and error checking.

POP QUIZ

At which layer of the TCP/IP model is the OSI Session layer found?

LAST BUT NOT LEAST

As mentioned previously, you are encouraged to review the RFC documentation for any further information on protocols. Be aware, however, that any RFC is subject to variations in interpretation, and one implementation of a protocol may not be identical to another. A network administrator or member of the support staff must always be aware of this when integrating network pieces from different manufacturers. When there are interoperability issues, performance degradation issues, or functional issues, you may have to draw on the RFC to find which way to point the finger.

8.6 Chapter Exercises

1. List in order from highest to lowest the upper layers of the OSI model, also indicating their layer number.
2. An application that runs on a user's workstation and communicates over a network with an appropriate application that is running on a server is considered to be what type of application?
3. Which protocol is considered to be a connection-based protocol?
4. What functionality can be used to disguise addresses from a private address space to be seen on the Internet?
5. List the three private address spaces that may be used and are considered to be not routable over the Internet.
6. Name an Application layer protocol that may be used to perform file transfers over the network.
7. What is the protocol that resolves IP addresses to hardware addresses?

8.7 Pop Quiz Answers

1. True or false: The Application layer is where all the application programs you load on your PC are stored.

   False

2. The predominant networking protocol run over Ethernet networks is TCP/IP

3. True or false: UDP is a connection-based protocol.

   False

4. Describe what happens to a packet when it is passed through a NAT-enabled router.

   A technique known as *port mapping* maps the hidden source address to an unused port number. A workstation that requests a page from a web server must access the server using port 80 for the server to respond to the request. When the server receives the request, its only concern is the destination port, which must be port 80. So what the source port number is makes no difference when servicing the request. The server simply sends the packet back to the requesting IP address, even though it is of a NAT-enabled router and not the actual workstation making the request. When the packet arrives at the NAT-enabled router, it examines the packet and finds that the destination port address correlates to a workstation on its private LAN in its NAT translation table. It modifies the packet with a new destination IP and port address, recalculates a new checksum, and then transmits it on to the private LAN. Therefore, knowing those temporary port addresses are available can come in handy when you're using NAT.

5. At which layer of the TCP/IP model can the physical component of a network node be found?

   Layer 1

6. What determines the type of framing that is to be used on a particular network segment?

   The media being used for that network segment.

7. Which TCP/IP model Transport layer protocol is connection based?

   TCP

8. True or false: The TCP/IP model Internet layer IP protocol is a connectionless protocol.

   True

9. What two ICMP applications can be used to verify the presence of an IP address on the Internet or local network?

   ping and/or traceroute

10. List what is required for a network card to have full-duplex capability.
    - input frame control
    - input frame buffer
    - receive circuit
    - output frame control
    - output frame buffer
    - transmit circuit

11. What is ARP used for?

    Address resolution

12. What is OSPF?

    A routing protocol

13. True or false: The maximum number of protocols the TCP/IP Application layer can have at any one time is two.

    False

14. True or false: The OSI model Presentation layer maps directly to the Transport layer of the TCP/IP model.

    True

15. At which layer of the TCP/IP model is the OSI Session layer found?

    Layer 5