

Chapter 16: Networking Fundamentals

Chapter 16 Objectives

At the conclusion of this chapter, the reader will be able to:

- € Describe the function of network hardware devices such as hubs, switches, and routers
- € Describe the differences between local, metropolitan, and wide-area networks
- € Describe the various media used in local area networks
- € List the layers in the OSI model and give the purpose of each
- € Work with IP addresses, unroutable blocks, subnet masks, and other IP features
- € Divide IP address blocks using subnetting
- € Describe the operation of DHCP and ARP in a local area network
- € Explain the operation of UDP and TCP transport mechanisms
- € Apply the principles of the OSI model to troubleshooting network systems
- € Use the *ping* and *tracert* commands to troubleshoot a network

Computer networks are crucial to the functioning of our modern society. They're found everywhere -- in businesses, homes, factories, and even in vehicles. The Internet is the world's largest network, providing global communication to anyone who can access it. By itself, a computer is an island; connected to a network, it can share in all the resources an organization has to offer. Networks both large and small operate on a small number of basic principles. A technician who has mastered these principles can take care of network problems with relative ease.

16-1 Network Hardware and Media

Networks are built for many purposes. *The primary purpose of a network is the sharing of information and resources.* You may have already used a network to share files, a printer, or some other item that several people needed to access. You've probably used e-mail and instant messaging. These applications are just the beginning of what is possible. For example, you can use a network to control the automation in a factory, check the security of your home while you're away, or remotely operate a robot in a hazardous environment. Thanks to the advancement of wireless networks, and the introduction of the "Internet of Things" (IoT), you can do these things from nearly anywhere in the world. Networks come in many shapes and sizes and use a variety of equipment to accomplish their jobs.

Types of Networks: LANs, MANs, and WANs

There are three basic classifications of networks. These are *local area networks* or *LANs*, *metropolitan area networks* or *MANs*, and *wide area networks* or *WANs*. They are well named. A LAN is a network contained primarily within a single building (or part of a building). Most organizations have one or more LAN systems, usually interconnected. A LAN features simple, inexpensive hardware and can operate at very high speeds (usually 100 Mbps or better; gigabit Ethernet itself can operate at 1000 Mbps or 1 Gbps). In a medium sized company there may be a separate LAN for each department or office, which improves security (no one can "sniff" the adjacent department's network for sensitive information) and reliability (if one department's network breaks, the rest of the organization remains operative). Most LANs use Ethernet and category 5 or 5e twisted-pair cable to carry the data signals. In general, LANs tend to have a small group of similar computers. Figure 16-1 shows a small LAN.

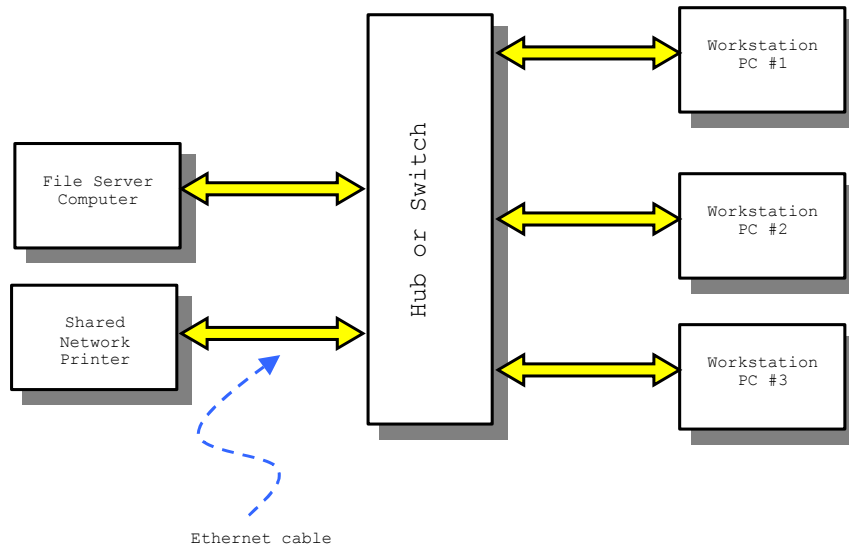


Figure 16-1: A Small Ethernet LAN with Workstations, a Server, and a Printer

A MAN is a network that is contained in more than one building, or is spread across a community. It is often used on college campuses to connect all the departments. Medium and large companies use MANs to connect their offices within the same city. Because of the larger distances in a MAN, twisted-pair copper cable is usually inadequate to carry the signals. The three most popular choices include *FDDI*, *fiber distributed data interface*, a fiber optic ring network; *leased lines*, which are provided by a telephone company; and wireless radio frequency links, which are used for point-to-point communications between two locations. Because of the specialized hardware (or the use of leased telephone lines), MANs are more expensive to operate than LANs. They may still operate at a high data rate as a LAN, but are more likely to be slower if leased lines are involved, or if the volume of traffic is high. Figure 16-2 shows a MAN utilizing an FDDI backbone. The *backbone* of a network is the part that carries the majority of traffic between network destinations.

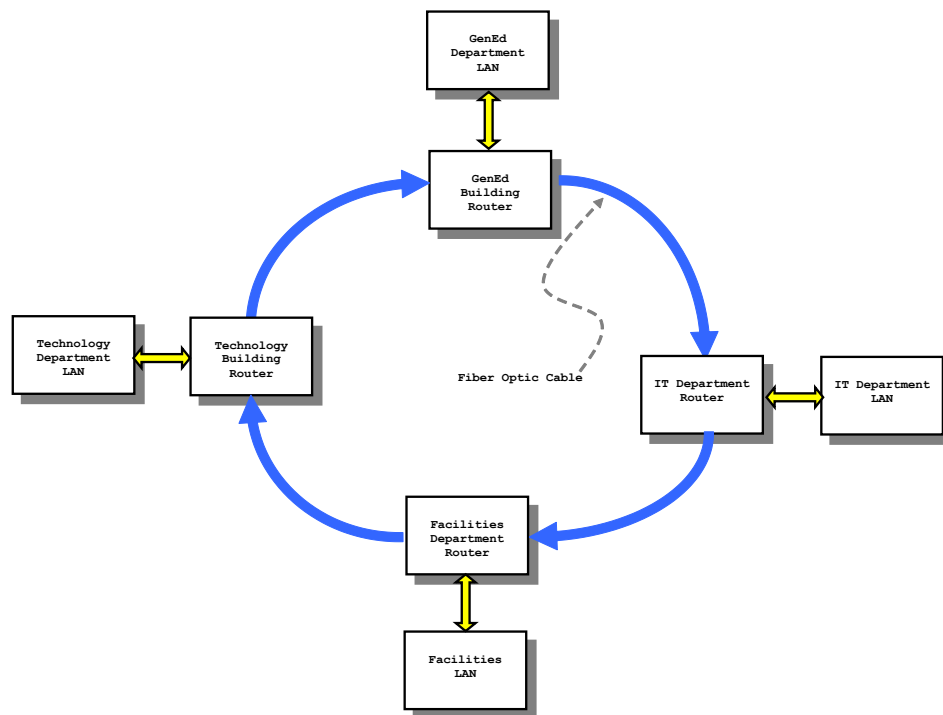


Figure 16-2: A MAN Employing an FDDI Backbone

The MAN in the figure effectively allows each department to communicate with other departments when needed, but also provides important isolation between them. A *router* with an FDDI interface is used at each department to decide what messages should be passed onto the FDDI ring, and which should not. Generally, *broadcast* messages are not repeated by the routers, which helps to prevent a condition called a *broadcast storm* from appearing on the MAN. A broadcast is a message intended to be received by all computers on a network. In a LAN, this is useful, but in a larger network, repetition of broadcasts by routers (or other equipment) can cause the same broadcast message (or a group of messages) to be endlessly repeated on the network; hence a “storm” of broadcast messages appears, which clogs the network with useless traffic. The FDDI ring operates just like the token ring network discussed in Chapter 15. It uses a back ring to help ensure reliability.

As you might have guessed, a *WAN* is a network covering most of a city, several cities, or any larger area. The Internet is the ultimate WAN! A WAN is the most expensive of all networks to operate. Individuals can afford to access the Internet because the cost is divided by the millions of users. It contains many different types of computers and a variety of connections between them, ranging from humble 56 kbps dial-up lines in the backwaters to gigabit-speed OC-768 (optical carrier) fiber optic links capable of operating at 40 Gbps. Even with high-speed backbones, WANs are almost always much slower than LANs and the delays in traffic are quite variable. Most of the structure of the Internet is provided by major telecommunications providers such as Qwest, Sprint, and the regional Bell operating companies (RBOCs).

A *firewall*, *gateway*, or *router* is needed to connect a LAN or MAN to the Internet. A firewall is a specialized hardware or software device that examines that traffic being passed to and from the WAN and halts messages that could cause harm (for example, an unauthorized connection requested to a file server within an organization). A gateway performs the basic function of placing messages onto the WAN.

Hardware Devices

Networks use a variety of devices to handle the traffic between computers. The three main devices used are *hubs*, *switches*, and *routers*. A hub is a device used to connect the workstations together on one segment of a LAN. A hub creates a common “party line” (bus) connection between all the devices sharing it. If any device transmits, all other devices on the hub’s ports will get the message. Hubs come in various sizes -- they may have anywhere from 4 to 32 I/O connections for computers. Hubs also come in two primary speeds, 10 Mbps (for 10-BaseT Ethernet application), and 100 Mbps (for 100-baseT Ethernet). The latter hubs will usually be marked “10/100-BaseT” and can support both speeds. Figure 16-3 shows a small 10-BaseT hub that might be a handy addition to a technician’s toolkit. Note that ports 1, 2, and 3 are for direct connection to personal computers. Port 4 is the expansion port. To expand the LAN to include more workstations, a second hub can be connected to port 4, making sure to move the switch to the “MDI” position.

The hub of Figure 16-3 is old and slow by today’s standards, yet you may want to include something like it in your toolkit. A hub repeats traffic on all of its outputs, unlike a switch. It can therefore be a useful tool for temporary insertion into a network for monitoring (“sniffing”) its activity.



Figure 16-3: A 10-BaseT Hub

Hubs create a common bus connection for computers on a LAN. Although this is certainly a simple way to connect them, it creates two problems. First, since any workstation can intercept messages meant for any other workstation, the security of such a network is not exactly optimal. Second, if two devices transmit at the same time, a *collision* will occur on the wire. You may recall from Chapter 15 that a bus network usually employs CSMA/CD (carrier sense multiple access with collision detection) protocol. Each device must listen to make sure the bus is quiet before transmitting. However, it is possible and very likely that two or more devices will need to talk at the same time when the network is quiet. Since no unit is transmitting, both devices see that it is OK to send and they both begin transmitting. This results in a *collision* on the wire and lost data. Both devices must back off and try again (using a pseudo-random “roll of the dice” to decide which one will go first). Collisions waste valuable time on a network. Excessive collisions can cause a network to grind almost to a halt. For this reason, it is not feasible to connect more than 20 to 25 workstations on a hub.

For larger LANs (or when security is more critical), *switches* are used as shown in Figure 16-4.



Figure 16-4: A Gigabit LAN Switch

The switch in the figure doesn't look much different than the hub of Figure 16-3, but its internal circuitry works in a very different way. When a computer sends a message frame to a switch, the switch examines the destination MAC address contained within the message. (We will study MAC addresses later. Each computer has a unique MAC address on a LAN.) The switch then forwards the message *only* to the computer with the correct MAC address. The ports for the other computers on the LAN remain silent. This improves security, since a message now flows only to the intended receiver. It also improves speed greatly; since the uninvolved switch ports are silent, the other computers can communicate on them (through the switch) with little or no risk of collisions. A switch is a much better choice for a LAN, even if there are only a few workstations on it. Simple switches require no programming and are simply connected to a network just like a hub. More advanced switches may be programmable and require specialized training to operate.

A *router* is a device that examines the network (usually IP) address of a message and decides which network it should be forwarded to. Routers always have two or more I/O ports and require programming to make them active on a network. Routers have slots where “blades” (I/O cards) can be installed to increase the number of interfaces (and hence, number of networks they can bridge). Routers use *routing tables* to determine where messages should be moved, and *access lists* to decide which messages are allowed through. A routing table is a list of network address destinations with instructions on how to get the message to the destination (which port on the router to send the message out). Routing tables can be configured *statically* (manually, by a network engineer) or *dynamically* (using a routing protocol), or a mixture of both. Routers form the backbone of the Internet; without them, the Internet couldn't exist!

Networking Media

A variety of cables are used to connect devices in a network. By far the most popular is UTP (unshielded twisted pair) Ethernet because of its low cost and ease of installation. Table 16-1 summarizes the characteristics of the most common media types. The naming system gives the maximum data capacity and signaling method. For example, “10Base2” media operates at a maximum of 10 Mbps and uses baseband signaling (the data signal is sent directly onto the wire without the use of any modulation method).

Media Type	Maximum Data Rate	Nodes per Segment	Maximum Length
Coaxial 10Base2, 10Base5	10 Mbps	30 (10Base2) 100 (10Base5)	500 m (10Base5) 200 m (10Base2)
10BaseT (UTP) Category 3 or 5	10 Mbps	2	100 m
100BaseT (UTP) Category 5	100 Mbps	2	100 m
1000BaseTX (UTP) Category 5e	1 Gbps	2	100 m
Fiber Optic	> 2 Gbps	2	> 10 km typical
Wireless 802.11(b) [2.4 GHz] 802.11(a) [5.4 GHz]	11 Mbps (802.11b) 54 Mbps (802.11a)	Limited by application; typically 50 or less	Range determined by terrain, antennas; typically 1 km or less.

Table 16-1: Common Networking Media Types and Characteristics

Coaxial cable is obsolete in LAN installations, but you may see it in some very old legacy installations. It is much more expensive and difficult to install when compared to twisted pair Ethernet. Coaxial cable is installed as shown in Figure 16-5. You may find it in older installations. With coax, a special device called a *vampire tap* is placed on the cable wherever a PC is to be connected. The cable is marked with the acceptable locations for tapping. Cables must also be terminated at the ends with a resistor plug in order to prevent reflections, which will cause false collisions and other problems on the line. It's easy to "break" a coaxial cable network; a moment's carelessness with a connector is all it takes.

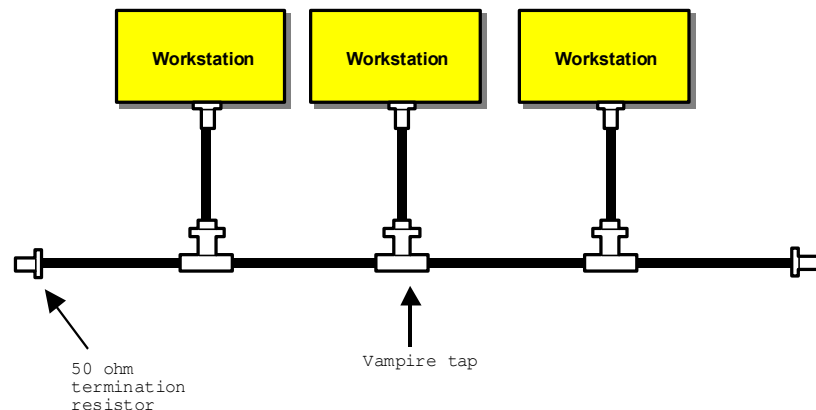


Figure 16-5: Coaxial Cable Connection in a LAN

UTP is by far the most commonly used media for LANs. It's inexpensive and relatively easy to work with. It comes in various grades, as shown in Table 16-1. Category 3 cable is conventional telephony-grade wire. When LANs were first being installed for businesses in the 1990s, quite a few shortsighted people decided that the existing telephone wiring in their buildings would be great for hooking up the network. However, this type of cable is actually poorly suited for data application, and those same people discovered this the hard way (their networks were unreliable!). For current installations, category 5 or even better, category 5e, is preferred. Category 5e can support data rates up to 1 Gbps with modern interface cards.

UTP cable uses a type RJ-45 telephone connector with eight conductors. Four of the conductors are not used except for 1000BaseTX applications. You'll likely be installing these connectors, so make sure to learn the standard color code shown in Figure 16-6. Figure 16-7 shows how to install RJ-45 connectors. In general, UTP cable segments should not exceed 100 meters in length.

UTP cables can be described as *straight*, *crossover*, or *rolled*. A straight cable has all wires directly connected to the same pin numbers on both ends of the cable. If you wire a Ethernet patch cable on both ends according to Figure 16-6, you'll have a straight cable. Straight cables connect PCs and other devices to hubs, switches, and routers. However, straight cables can't directly connect two PCs. For this application, a *crossover* cable is needed. To build a crossover cable, wire one end according to Figure 16-6; on the other end,


swap the orange and green pairs. A rolled cable is sometimes required to interface two devices. This cable has all the wires in normal order at one end, and in *reverse* order at the other end.

Some UTP cabling has very faint, pastel coloring on the conductors. You may need to look carefully to discriminate between the wires.

Most organizations use jacket color codes to help identify cables in racks. Blue, green, or gray jackets normally signify *straight* cables, while yellow is typically used for *crossover* cables. Other colors are also used; for example, red is sometimes used for jumpers between switches, or busses passing to router ports. Using a consistent color code makes it easier to understand where circuits are being routed when you're troubleshooting a network equipment rack.

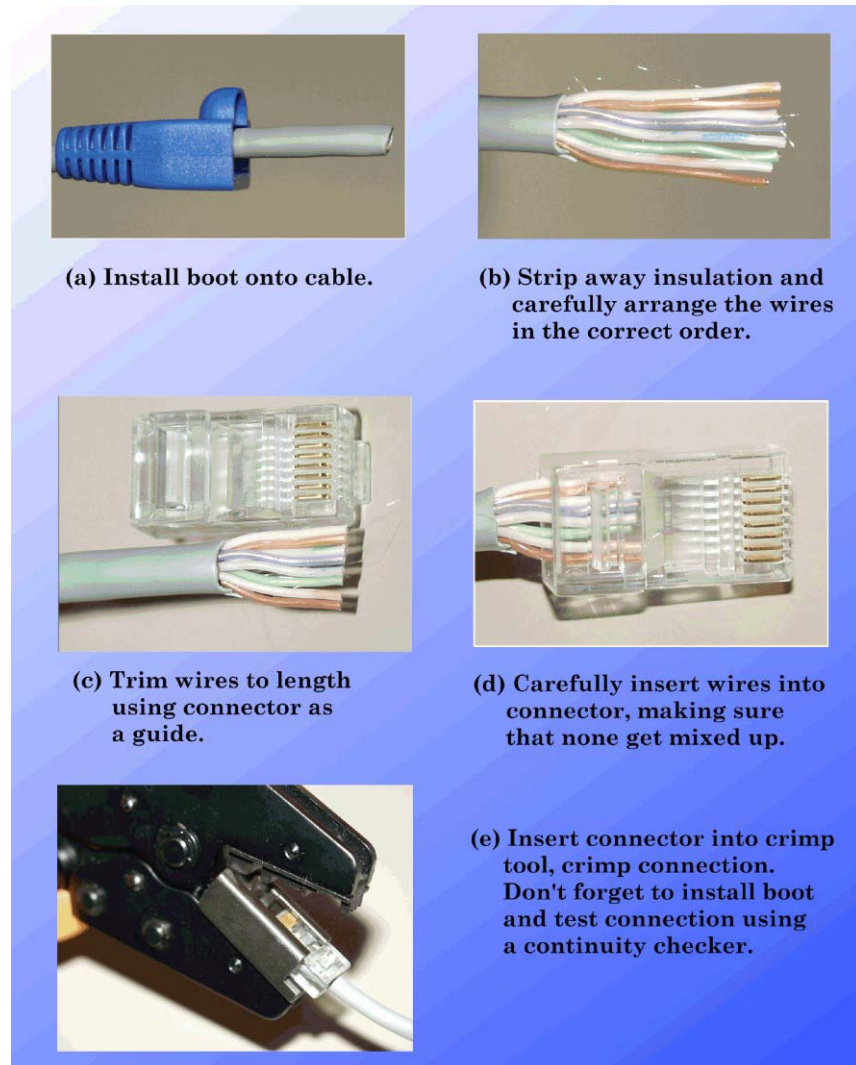
Fiber optic cable (discussed in detail in Chapter 18) offers the highest data rates of all the common media and also the largest distances. Special techniques are needed for installation of fiber optic connectors. A fiber optic network interface requires *two* fibers, one for receiving and the other for sending. Fiber is practically immune to eavesdropping, so it provides the highest level of security of any network media. You'll often find it used to connect a LAN to a backbone, or for use as a backbone in a FDDI network.

Wireless (sometimes called *WiFi*) network components have advanced rapidly in the last few years. The two predominant types are 802.11(b) and 802.11(g). To connect a computer to an existing 802.11 network, a compatible wireless network card must be installed in it. If the network is using encryption, the encryption key must be programmed into the workstation so that it can access the LAN. Wireless LANs can operate in two modes, *ad-hoc* and *access point*. In an ad-hoc wireless LAN, each computer directly talks to the other computers. This is useful for sharing files between computers; several computers can form an ad-hoc LAN anywhere at any time. The ad-hoc LAN has no connection to a wired network unless one or more of the PCs has a separate wired network interface. In access point mode, the wireless cards in the PCs can only talk to a central wireless access point (AP), which grants each device access to the network. The AP is usually connected to a wired LAN so that the computers can access servers, printers, and the Internet.




RJ45 Standard Wiring Pattern
Viewed with spring clip down, contacts facing up

Figure 16-6: Ethernet RJ-45 Standard Wiring Practice




(a) Install boot onto cable.

(b) Strip away insulation and carefully arrange the wires in the correct order.



(c) Trim wires to length using connector as a guide.



(d) Carefully insert wires into connector, making sure that none get mixed up.



(e) Insert connector into crimp tool, crimp connection. Don't forget to install boot and test connection using a continuity checker.

Figure 16-7: How to Install RJ-45 Connectors on UTP Cable

Section Checkpoint

- 16-1 List the three types of networks and compare them in terms of speed, cost, and types of hosts likely to found on them.
- 16-2 What type of media is used by FDDI?
- 16-3 What is a good application of FDDI?
- 16-4 Explain the difference between a switch and a hub. Why would it be a poor practice from a security standpoint to use a hub as a permanent part of a LAN?
- 16-5 What are routers used for?
- 16-6 Why must CSMA/CD be used on an Ethernet network?
- 16-7 What is the most popular media for connecting Ethernet networks? Why?
- 16-8 What media would be used for forming a high-speed connection to a backbone?
- 16-9 Explain the difference between straight, crossover, and rolled Ethernet cables.
- 16-10 What are the two most popular WiFi standards? What data rates can they support?

16-2 The ISO/OSI Model

The OSI model is a way of representing the function of the various portions of a telecommunication system. It divides systems into seven levels as shown in Figure 16-8.




Figure 16-8: Open Systems Interconnect Model

Many people use the phrase "*Please Do Not Throw Sausage Pizza Away*" in order to help them remember the names of the seven layers. In fact, many systems may not even have or need all seven layers. (The TCP/IP protocol suite is an example of a system employing only five of the layers.) Notice that these layers are stacked. The reason for this is that information passes through these layers in sequence in both directions. This will become more clear as we study the purpose of each layer.

Physical Layer

The physical layer includes all data communication *hardware* that is necessary to carry a message. This includes UARTs, RS232 interfaces, modems, transmission lines, fiber optics, and so on. The physical layer also defines the electrical (or optical) representation of binary information and the procedural events that place the encoded signal onto the medium. All information must eventually pass through the physical layer - - so if it fails, all layers above it also stop working. At the physical layer, information is referred to as *bits*.

Data Link Layer

This layer is concerned with getting data to flow from one device to another on a network segment. (A *network segment* is a communications channel that is directly shared by two or more computers, such as a local area network or LAN.) This involves assembling data into message units called *frames* (think of assembling a collection of pages into an envelope), as well as providing error-checking and error-recovery mechanisms (such as checksums and CRCs). Note that the data link layer merely gets the data moved between two points on a local network. It just oversees the communications, correcting errors as needed.

The data link layer can be divided into two sub-layers called *MAC* and *LLC*. The MAC (medium access control) sub-layer controls physical access to the communications channel. It provides the mechanisms for controlling who can transmit messages, and when. It also provides local-level addressing in the form of *MAC addresses* (also known as *hardware addresses*). A MAC address is a 48-bit number that represents the local address of the device sending a message. All devices have unique MAC addresses, usually contained in a ROM (read-only memory). The LLC (logical link control) sub-layer supervises the transmission of frames and participates in the error recovery process.

At the data link layer, information is referred to as *frames*.

TIP: MAC addresses are usually expressed as a set of six hexadecimal byte values. The first three bytes represent the manufacturer's ID, assigned to the manufacturer by the institute of electrical and electronic engineers (IEEE), and the last three are the serial number of the network interface card.

Network Layer

Since a network can be composed of hundreds or thousands of computers, *addresses* are needed on all messages so that data arrives at the proper remote destination when it is sent. This is the function of the *network* layer. This layer forms logical addresses, and takes care of *routing* messages over paths that may require that the message be repeated by many computers along the way.

Addresses at the network layer (normally IPv4 or IPv6) consist of two portions, a *host* address and a *network address*. The network address specifies which physical network segment (anywhere in the world) is to be used as the destination for the message, and the host portion of the address tells which computer on that network should get the message. We will discuss how IP works in greater detail soon. At the network layer, information is referred to as *packets*.

Transport Layer

This layer is responsible for delivering error-free communication, and it ensures that message packets aren't lost or "dropped." Note that the error-checking provided in this layer is in addition to that provided by the data link layer. The transport layer breaks up large messages from the *session* layer (above it, to be discussed next) for "digestion" by the lower network layers (network, data link, physical). Incoming fragments of messages are reassembled into complete (large) messages by this layer. The three primary Internet protocols that operate at this layer are UDP (user datagram protocol), TCP (transmission control protocol), and ICMP (Internet control message protocol).

At the transport layer, data are referred to as *segments* if TCP is being used, or *datagrams* if UDP is employed.

Session Layer

This mechanisms within this layer provide a virtual (abstract) connection between two computers called a *session*. It provides needed functions such as name lookup (finding the address of a remote computer by its assigned name using domain name system (DNS) services), security (deciding who can establish a connection, and who can't), and session management. Session management allows several computer processes to share the same communications channel without interference (multiplexing). The session layer uses *source and destination port numbers* to identify and separate communication streams between computers.

Presentation Layer

This layer translates data between the format needed for applications (programs typically operated by an end-user) and the session layer (which delivers a virtual connection between two computers). An example of this layer is the *network redirector* in a local area network. The network redirector is a program that is part of the operating system. It makes remote files on a server computer visible to the client computer; it also allows the sharing of printers and other resources over the network. Another example is the translation of graphics (drawing) commands between different computer systems.

Application Layer

This last layer provides services that support end-user applications, such as electronic mail, database access, Internet access, and so on. When a computer program (such as a web browser or e-mail package) accesses network services, it interacts directly with this layer. The services provided by the application layer are called APIs, or *application programming interfaces*. The most common API implementation for networking is the Berkeley sockets model. In this model, a *socket* is an abstract (virtual) point of connection for a computer program to communicate through, just like a physical socket on a piece of equipment can be used for passing electrical signals.

How the OSI Layers Work Together

The concepts involved with the seven OSI layers are fairly abstract. It helps to see an example of what they really do. Imagine that Jim wants to send an electronic mail message to his friend Gerry, who lives in Texas. Jim connects his computer to the Internet, using an Internet service provider and launches an electronic mail *application program*.

Jim types his message to Gerry:

```
Hi Gerry. I just wanted to let you know that we got the fruitcake you and Lisa sent us
for Christmas. By the way, I am still having trouble with my El Camino. It runs really rough
when the engine is cold, and it stalls every time I stop at an intersection. After the engine
warms up it smooths right out and runs perfectly. Any ideas?
Jim
```

Jim presses the *SEND* button, and in just a few seconds, the message is on its way to Gerry. What has actually happened here? The *overall* action is that the letter Jim typed ended up on the electronic mail (e-mail) server computer across town at Jim's ISP. That letter then got passed across the country to Gerry's ISP, so that Jerry would be able to retrieve it later. Figure 16-9 shows the activity in terms of the OSI model.

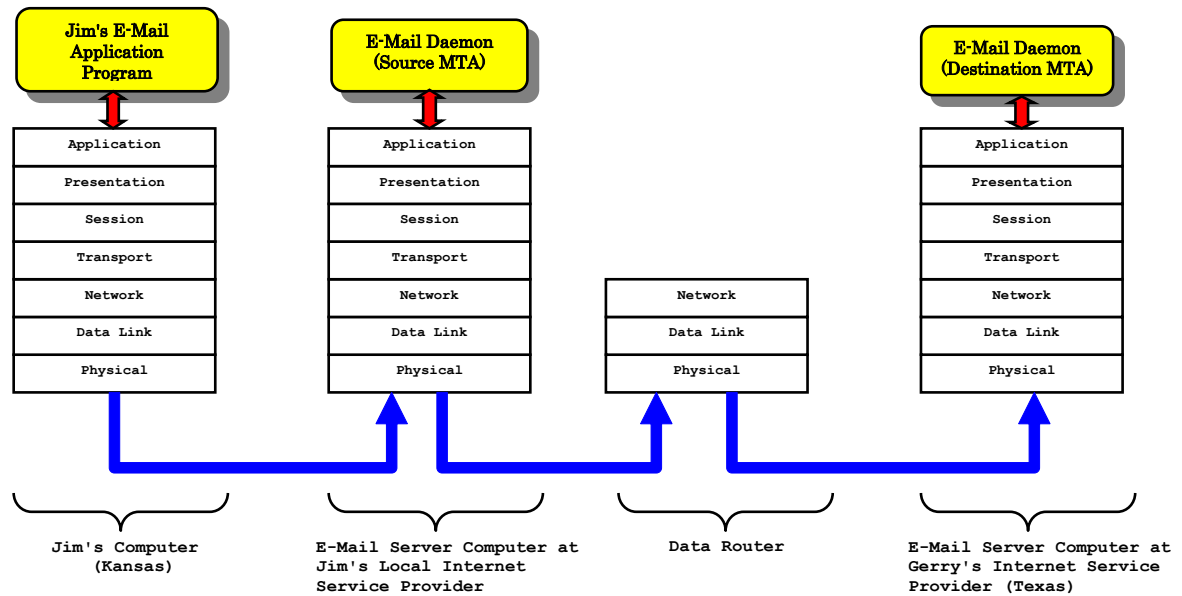


Figure 16-9: Sending Electronic Mail

The basic electronic sequence of events is as follows (and starts at the top layer, *application* in Jim's computer at left):

1. Jim's E-mail application program opened a "connection" to the Internet, though the *application* layer within Jim's computer, using subroutine calls to the operating system entry points. The Internet connection was requested with the *e-mail server computer* located at Jim's ISP. This server acts as the mail transfer agent (MTA) for all e-mail leaving and entering Jim's ISP.
2. The connection being established, Jim's E-mail application began passing the information and necessary commands to the ISP e-mail server computer. These commands were actually passed to the operating system on Jim's computer. They were translated by the *presentation layer* into the proper format for talking to the ISP computer. The *session layer* simplified matters by allowing the e-mail program to establish a virtual "connection" to the e-mail server computer, and helped find the address of the e-mail server. The session layer also identified which process on the e-mail server would receive the e-mail data by examining the *destination port number* (25, SMTP) supplied by the E-mail application on Jim's computer. The *transport* and *network* layers made sure the message and commands got to the right computer (the e-mail server). The *transport* layer also divided the message into bite-sized chunks for the lower layers. After the message passed through the *transport* layer, it looked a little like this:

```
01 Hi Gerry. I just wanted to...
02 let you know that we got the...
[ . . . More pieces of the message . . . ]
37 ? Jim
```

3. The *data link* layer checked for errors as the pieces of the message passed through it. The *physical* layer was the serial port and modem that connected Jim's computer to the telephone line (he used a dial-up connection at the ISP).
4. At the server computer, the e-mail message and commands moved back up through the layers, being error-checked, reassembled, and transformed as needed until they reached the e-mail *daemon* program. A *daemon* is a program that runs at all times, without the need for human interaction. (A Windows "service" is an example of this kind of program.)
5. The e-mail daemon program processed Jim's message and saw that the e-mail needed to be forwarded to Gerry's ISP in Texas, so the message was again disassembled into its parts and sent down through the layers on the e-mail server computer, where it reached the *physical* layer again. This time, the message went out not on a phone line, but onto a Internet line, where it reached a *router*.
6. The *router* examined the logical (IP) address on the e-mail message (it could not tell it was an e-mail, it could only see fragments of text), and found the correct route to the Texas computer. The message went out through *another* branch of the Internet to the Gerry's ISP server computer in Texas.

7. The e-mail message arrived in pieces at the Texas computer. It bubbled up through the OSI layers, and was reassembled into one continuous message by the *transport* layer, and it then made its way up to the e-mail daemon program, which recognized the incoming message, and stored it in Gerry's in-box. The message would be retrieved when Gerry connected to the Texas server computer.

There are a *lot* of events taking place during network communications! Notice that every device on a network does not need all the OSI layers. In particular, a router only needs the first three layers, since it merely decides where a message is going to go. The primary activity of a router is at the *network* layer, since routers decide where to send messages based on the IP address within the message packet.

Section Checkpoint

- 16-11 List the seven ISO/OSI layers, and explain what each one does.
 - 16-12 What three layers are used by routers? What layer does a router work with primarily?
 - 16-13 What is the name of the abstract (virtual) network connection provided by the application layer?
 - 16-14 What is a network redirector?
 - 16-15 Give several examples of programs that interact with the application layer.
-

16-3 The Internet and Internet Protocol (IP) Addressing

The *Internet* is a global network of computers. Its origins go back to 1957, the year that Russia launched *Sputnik*, the first human-made satellite. The United States government was shocked by this event, and in response, formed *ARPA*, the Advanced Research Projects Agency, as a part of the department of defense. In 1962, the US Air Force commissioned Paul Baran of RAND (a government-controlled corporation) to study how control could be maintained over missiles and bombers in the event of a nuclear attack. This was to be a computer network that could survive a nuclear strike. Baran's proposal involved *decentralizing* the computing power, and developing a *packet-switched network*.

Decentralizing the computers would distribute the burden geographically, so that if a key city was destroyed, the rest of the communication network would be intact to coordinate a counterstrike. A packet-switched network breaks messages up into small packets called *datagrams*, which contain origin and destination addresses, and are forwarded through the network of computers until they reach their destination, where they are reassembled into a complete message.

The first actual physical network, called *ARPANET*, wasn't finished until 1969, and it linked four locations: University of California at Los Angeles, SRI (in Stanford), University of California at Santa Barbara, and University of Utah. The data rates were limited to 50 kbps on the links. In the period between 1969 and the present, many technical improvements and innovations came to pass that improved the efficiency of communications. The number of computers connected to the network increased (as of 2019, there were over 1.1 billion host computers in worldwide DNS), and the speed of the communications backbones increased (155 Mbps backbones are now common.)

In 1992, CERN released the standards that would form the World Wide Web (WWW), and in 1993, commercial use was allowed. Growth of the Internet (as it was now called) literally exploded, with the number of online hosts doubling every 6 months. Internet access is now considered by many to be a basic utility service, much like the telephone. In fact, most local telephone companies are also ISPs; hard-wired telephones are on their way to becoming a part of history, having been replaced by wireless digital communications services.

Internet Addressing System - IPv4

Computers connected to the Internet are assigned a 32-bit *Internet Address*, or *IPv4 address*, or a 128-bit *IPv6 address*. IPv6 is a response to the shortage of IPv4 addresses. An IPv4 address is traditionally written as four decimal numbers separated by periods, even though each of the numbers represents an *octet* or *byte* of information. This format is sometimes called *dotted decimal* notation. IPv4 addresses are still commonly used within LANs; IPv6 is now the mainstay of the open Internet.

An IP address consists of two portions, the *network address* and *host address*. The *network* portion of the address identifies which network segment (worldwide) a message is destined for, and the *host* portion identifies which computer on that network should get the message. Organizations are normally issued *blocks* or groups of IP addresses.

Because some organizations have many computers and others have few, the IPv4 addressing system has some peculiar characteristics. Internet addresses are issued in three groups, class A (for the largest organizations), class B (for large organizations), and class C (for small organizations). This system is called *classful addressing*, and it is the basis of IPv4 addressing. However, it is a wasteful system (about 97% of the available IP addresses are wasted in classful addressing), and is now being supplemented by *CIDR*, *classless*