


What You Will Learn

On completing this chapter, you will be able to:

- ✓ Differentiate among unicast, multicast, and broadcast transmission methods
- ✓ Describe store-and-forward, cut-through, and fragment-free switching mechanisms
- ✓ Describe Layer 2 and Layer 3 switching operation



CHAPTER 6


How a Switch Works

Up to this point, frames going in and out of the LAN switch have been discussed, but not what those frames are doing while in the switch and what the switch is doing with the frames. As you might have surmised by now, this chapter discusses these very points, and a few more. To understand how a switch processes the frames that it receives and forwards, you will first learn about the three types of transmission methods found in a local-area network (LAN): unicast, multicast, and broadcast.

Frames Revisited

Recall from Chapter 1, “Networking Basics,” that frames carry data across the network and are made up of three parts: the header, the data itself (payload), and the trailer, as illustrated in the Figure 6-1.

Figure 6-1 Complete Frame (Header, Data [Payload], Trailer)



These three frame components—the header, data, and trailer—combine in making up a complete frame. The header identifies the destination data-link address of the




frame, the payload is data from upper-layer protocols (such as packets from the network layer), and the trailer signifies the end of the frame.

Recall from Chapter 5, “Ethernet LANs,” that the MAC address (Media Access Control address or physical address) is the unique serial number burned into network adapters that differentiates that network card from all others on the network. To be a part of any network, you must have an address so that others can reach you. There are two types of addresses found in a network: the logical network address and the physical data-link address. In LAN bridging and switching environments, you are concerned with the physical address (MAC address), and the MAC address is found in the frame header.

A MAC address is the physical address of the device and is 48 bits (6 bytes) long. It is made up of two parts: the organizational unique identifier (OUI) and the vendor-assigned address, as illustrated in Figure 6-2.

Figure 6-2 MAC Address



Recall that the MAC address on a computer might look like this: 00-06-0f-08-b4-12. This MAC address is used for the Fast Ethernet adapter on the computer in question—the OUI is 00-06-0f, and the vendor-assigned number is 08-b4-12.

Transmission Methods


LAN data transmissions at Layer 2 fall into three classifications: unicast, multicast, and broadcast. In each type of transmission, a single frame is sent to one node on the network. If the frame is to be sent to more than one node on the network, the sender must send individual unicast data streams to each node.

In a *unicast* transmission, a single frame or packet is sent from a single source to a single destination on a network. In a *multicast* transmission environment, a single data frame or a single source to multiple destinations packet is copied and sent to a specific subset of nodes on the network. In a *broadcast* transmission environment from a single source to all nodes, a single data frame or packet is copied and sent to all nodes on the network.

Unicast

Unicast is a one-to-one transmission method in which the network carries a message to one receiver, such as from a server to a LAN workstation. In a unicast environment, even though multiple users might ask for the same information from the same server at the same time, such as a video clip, duplicate data streams are sent. One stream is sent to each user, as illustrated in the Figure 6-3.

Figure 6-3 Unicast Operation




Unicast sends separate data streams to each computer requesting the data, in turn flooding the network with traffic. Unicast might be compared to an after-work gathering. You and several of your co-workers might be going to the same destination, but each taking his own vehicle, flooding the streets with cars. (So the next time you go to an after-work gathering, and each person drives his own car, tell them you're "unicasting.")



Multicast

Multicast is a one-to-many transmission method in which the network carries a message to multiple receivers at the same time. Multicast is similar to broadcasting, except that multicasting means sending to a specific group, whereas broadcasting implies sending to everybody, whether they want the traffic or not. When sending large amounts of data, multicast saves considerable network bandwidth because the bulk of the data is sent only once. The data travels from its source through major backbones and is then multiplied, or distributed out, at switching points closer to the end users (see Figure 6-4). This is more efficient than a unicast system, in which the data is copied and forwarded to each recipient.

Figure 6-4 Multicast Operation




Multicast conserves network bandwidth by sending a single data stream across the network, much as you and others might carpool to and from work, thereby reducing the traffic on the roads. For example, a few of you might ride together to some point, such as a drop-off point in the city, and then disperse from there. Multicasting works in the same way by using the concept of shared transmission across a network. Multicasting sends the data to a predetermined endpoint, such as a switch, where the traffic is sent to each intended recipient, instead of each traffic stream being sent from start to finish across the network, independent of others.

Broadcast

Broadcast is a one-to-all transmission method in which the network carries a message to all devices at the same time, as illustrated in Figure 6-5.

Figure 6-5 Broadcast Operation



Broadcast message traffic is sent out to every node on the network where the broadcast is not filtered or blocked by a router. Broadcasts are issued by the **Address Resolution Protocol (ARP)** for address resolution when the location of a user or server is not known. For example, the location could be unknown when a network client or server first joins the network and identifies itself. Sometimes broadcasts are a result of network devices continually announcing their presence in the network, so that other devices don't forget who is still a part of the network. Regardless of the reason for a broadcast, the broadcast must reach all possible stations that might potentially respond.

Frame Size

Frame size is measured in bytes and has a minimum and maximum length, depending on the implemented technology. For example, the minimum frame size for an Ethernet LAN is 64 bytes with a 4-byte **cyclic redundancy check (CRC)**, and the maximum frame size is 1518 bytes. The minimum/maximum for a Token Ring LAN is 32 bytes/16 kilobytes (KB), respectively.




Why is it important to know the minimum and maximum frame sizes your network can support? Knowing the sizes enables you to ensure that your users' message traffic gets to where it needs to go quickly and accurately.

Suppose your corporate mailroom is equipped only to handle letter- and business-sized envelopes and is not equipped to handle postcards or larger legal-sized envelopes. The letter-sized envelope is the minimum size, and the business-sized envelope is the maximum sized "frame" allowed by your mailroom. Anything smaller than the letter-sized envelope, such as a postcard, might be considered a *runt*, and anything larger than the business-sized envelope might be considered a *giant*.

Figure 6-6 illustrates the concept of a minimum and maximum frame size, and the result, in a corporate mailroom. (Let's hope this doesn't really happen, although it might explain a few missing pieces of mail.)

Figure 6-6 Mailroom in Action



In this mailroom (switch) scenario, both the postcards (runts) and legal-sized envelopes (giants) would not be accepted by the mailroom (the switch) and therefore would be dropped into the trash.

**note**

The maximum frame size is also known as the maximum transmission unit, or MTU. When a frame is larger than the MTU, it is broken down, or fragmented, into smaller pieces by the Layer 3 protocol to accommodate the MTU of the network.

Layer 2 Switching Methods

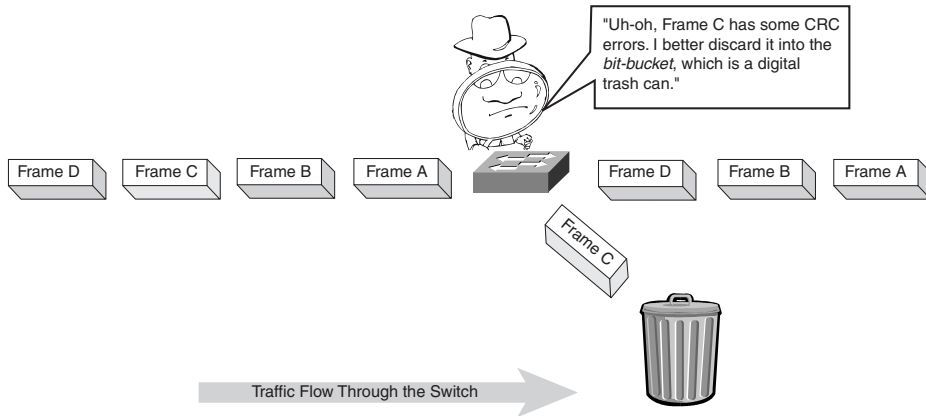
LAN switches are characterized by the forwarding method that they support, such as a store-and-forward switch, cut-through switch, or fragment-free switch. In the store-and-forward switching method, error checking is performed against the frame, and any frame with errors is discarded. With the cut-through switching method, no error checking is performed against the frame, which makes forwarding the frame through the switch faster than store-and-forward switches.

Store-and-Forward Switching

Store-and-forward switching means that the LAN switch copies each complete frame into the switch memory buffers and computes a cyclic redundancy check (CRC) for errors. CRC is an error-checking method that uses a mathematical formula, based on the number of bits (1s) in the frame, to determine whether the received frame is errored. If a CRC error is found, the frame is discarded. If the frame is error free, the switch forwards the frame out the appropriate interface port, as illustrated in Figure 6-7.



Figure 6-7 Store-and-Forward Switch Discarding a Frame with a Bad CRC



An Ethernet frame is discarded if it is smaller than 64 bytes in length, a runt, or if the frame is larger than 1518 bytes in length, a giant, as illustrated in Figure 6-8.



note

Some switches can be configured to carry giant, or jumbo, frames.

If the frame does not contain any errors, and is not a runt or a giant, the LAN switch looks up the destination address in its forwarding, or switching, table and determines the outgoing interface. It then forwards the frame toward its intended destination.

Store-and-Forward Switching Operation

Store-and-forward switches store the entire frame in internal memory and check the frame for errors before forwarding the frame to its destination. Store-and-forward switch operation ensures a high level of error-free network traffic, because bad data frames are discarded rather than forwarded across the network, as illustrated in Figure 6-9.

Figure 6-8 Runts and Giants in the Switch

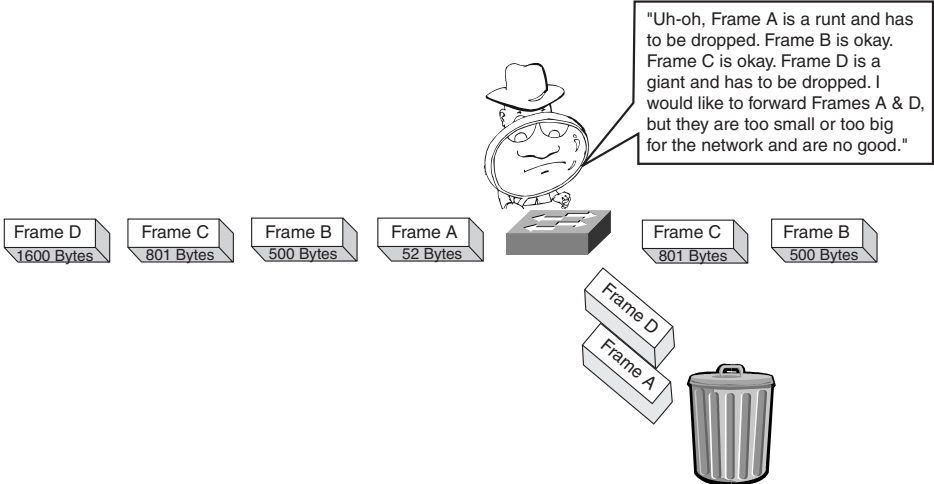



Figure 6-9 Store-and-Forward Switch Examining Each Frame for Errors Before Forwarding to Destination Network Segment





The store-and-forward switch shown in Figure 6-9 inspects each received frame for errors before forwarding it on to the frame's destination network segment. If a frame fails this inspection, the switch drops the frame from its buffers, and the frame is thrown in to the proverbial bit bucket.


A drawback to the store-and-forward switching method is one of performance, because the switch has to store the entire data frame before checking for errors and forwarding. This error checking results in high switch latency (delay). If multiple switches are connected, with the data being checked at each switch point, total network performance can suffer as a result. Another drawback to store-and-forward switching is that the switch requires more memory and processor (central processing unit, CPU) cycles to perform the detailed inspection of each frame than that of cut-through or fragment-free switching.

Cut-Through Switching

With cut-through switching, the LAN switch copies into its memory only the destination MAC address, which is located in the first 6 bytes of the frame following the preamble. The switch looks up the destination MAC address in its switching table, determines the outgoing interface port, and forwards the frame on to its destination through the designated switch port. A cut-through switch reduces delay because the switch begins to forward the frame as soon as it reads the destination MAC address and determines the outgoing switch port, as illustrated in Figure 6-10.

The cut-through switch shown in Figure 6-10 inspects each received frame's header to determine the destination before forwarding on to the frame's destination network segment. Frames with and without errors are forwarded in cut-through switching operations, leaving the error detection of the frame to the intended recipient. If the receiving switch determines the frame is errored, the frame is thrown out to the bit bucket where the frame is subsequently discarded from the network.

Figure 6-10 Cut-Through Switch Examining Each Frame Header Before Forwarding to Destination Network Segment



Cut-through switching was developed to reduce the delay in the switch processing frames as they arrive at the switch and are forwarded on to the destination switch port. The switch pulls the frame header into its port buffer. When the destination MAC address is determined by the switch, the switch forwards the frame out the correct interface port to the frame's intended destination.

Cut-through switching reduces *latency* inside the switch. If the frame was corrupted in transit, however, the switch still forwards the bad frame. The destination receives this bad frame, checks the frame's CRC, and discards it, forcing the source to resend the frame. This process wastes bandwidth and, if it occurs too often, network users experience a significant slowdown on the network. In contrast, store-and-forward switching prevents errored frames from being forwarded across the network and provides for *quality of service (QoS)* managing network traffic flow.



note

Today's switches don't suffer the *network latency* that older (legacy) switches labored under. This minimizes the effect switch latency has on your traffic. Today's switches are better suited for a store-and-forward environment.

Cut-Through Switching Operation

Cut-through switches do not perform any error checking of the frame because the switch looks only for the frame's destination MAC address and forwards the frame out the appropriate switch port. Cut-through switching results in low switch latency. The drawback, however, is that bad data frames, as well as good frames, are sent to their destinations. At first blush, this might not sound bad because most network cards do their own frame checking by default to ensure good data is received. You might find that if your network is broken down into workgroups, the likelihood of bad frames or collisions might be minimized, in turn making cut-through switching a good choice for your network.

Fragment-Free Switching

Fragment-free switching is also known as *runtless switching* and is a hybrid of cut-through and store-and-forward switching. Fragment-free switching was developed to solve the late-collision problem.



note

Recall that when two systems' transmissions occur at the same time, the result is a collision. Collisions are a part of Ethernet communications and do not imply any error condition. A late collision is similar to an Ethernet collision, except that it occurs after all hosts on the network should have been able to notice that a host was already transmitting.

A late collision indicates that another system attempted to transmit after a host has transmitted at least the first 60 bytes of its frame. Late collisions are often caused by an Ethernet LAN being too large and therefore needing to be segmented. Late collisions can also be caused by faulty network devices on the segment and duplex (for example, half-duplex/full-duplex) mismatches between connected devices.

Fragment-Free Switching Operation

Fragment-free switching works like cut-through switching with the exception that a switch in fragment-free mode stores the first 64 bytes of the frame before forwarding. Fragment-free switching can be viewed as a compromise between store-and-forward switching and cut-through switching. The reason fragment-free switching stores only the first 64 bytes of the frame is that most network errors and collisions occur during the first 64 bytes of a frame.



note

Different methods work better at different points in the network. For example, cut-through switching is best for the network core where errors are fewer, and speed is of utmost importance. Store-and-forward is best at the network access layer where most network problems and users are located.


Layer 3 Switching

Layer 3 switching is another example of fragment-free switching. Up to now, this discussion has concentrated on switching and bridging at the data link layer (Layer 2) of the Open System Interconnection (OSI) model. When bridge technology was first developed, it was not practical to build wire-speed bridges with large numbers of high-speed ports because of the manufacturing cost involved. With improved technology, many functions previously implemented in software were moved into the hardware, increasing performance and enabling manufacturers to build reasonably priced *wire-speed* switches.

Whereas bridges and switches work at the data link layer (OSI Layer 2), routers work at the network layer (OSI Layer 3). Routers provide functionality beyond that offered by bridges or switches. As a result, however, routers entail greater complexity. Like early bridges, routers were often implemented in software, running on a special-purpose processing platform, such as a personal computer (PC) with two network interface cards (NICs) and software to route data between each NIC, as illustrated in Figure 6-11.



Figure 6-11 PC Routing with Two NICs



The early days of routing involved a computer and two NIC cards, not unlike two people having a conversation, but having to go through a third person to do so. The workstation would send its traffic across the wire, and the routing computer would receive it on one NIC, determine that the traffic would have to be sent out the other NIC, and then resend the traffic out this other NIC.



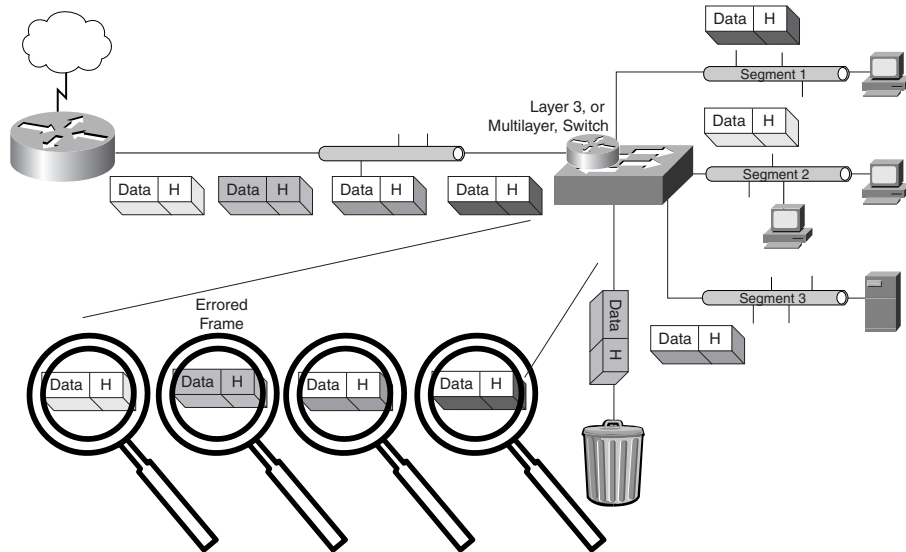
note

In the same way that a Layer 2 switch is another name for a bridge, a Layer 3 switch is another name for a router. This is not to say that a Layer 3 switch and a router operate the same way. Layer 3 switches make decisions based on the port-level Internet Protocol (IP) addresses, whereas routers make decisions based on a map of the Layer 3 network (maintained in a routing table).

Multilayer switching is a switching technique that switches at both the data link (OSI Layer 2) and network (OSI Layer 3) layers. To enable multilayer switching, LAN switches must use store-and-forward techniques because the switch must

receive the entire frame before it performs any protocol layer operations, as illustrated in Figure 6-12.

Figure 6-12 Layer 3 (Multilayer) Switch Examining Each Frame for Error Before Determining the Destination Network Segment (Based on the Network Address)



Similar to a store-and-forward switch, with multilayer switching the switch pulls the entire received frame into its memory and calculates its CRC. It then determines whether the frame is good or bad. If the CRC calculated on the packet matches the CRC calculated by the switch, the destination address is read and the frame is forwarded out the correct switch port. If the CRC does not match the frame, the frame is discarded. Because this type of switching waits for the entire frame to be received before forwarding, port latency times can become high, which can result in some latency, or delay, of network traffic.



Layer 3 Switching Operation

You might be asking yourself, “What’s the difference between a Layer 3 switch and a router?” The fundamental difference between a Layer 3 switch and a router is that Layer 3 switches have optimized hardware passing data traffic as fast as Layer 2 switches. However, Layer 3 switches make decisions regarding how to transmit traffic at Layer 3, just as a router does.



note

Within the LAN environment, a Layer 3 switch is usually faster than a router because it is built on switching hardware. Bear in mind that the Layer 3 switch is not as versatile as a router, so do not discount the use of a router in your LAN without first examining your LAN requirements, such as the use of *network address translation (NAT)*.


Before going forward with this discussion, recall the following points:

- A switch is a Layer 2 (data link) device with physical ports and that the switch communicates via frames that are placed on to the wire at Layer 1 (physical).
- A router is a Layer 3 (network) device that communicates with other routers with the use of packets, which in turn are encapsulated inside frames.


Routers have interfaces for connection into the network medium. For a router to route data over the Ethernet, for instance, the router requires an Ethernet interface, as illustrated in Figure 6-13.

A serial interface is required for the router connecting to a wide-area network (WAN), and a Token Ring interface is required for the router connecting to a Token Ring network.

A simple network made up of two network segments and an internetworking device (in this case, a router) is shown in Figure 6-14.

Figure 6-13 Router Interfaces

The router in Figure 6-14 has two Ethernet interfaces, labeled E0 and E1. The primary function of the router is determining the best network path in a complex network. A router has three ways to learn about networks and make the determination regarding the best path: through locally connected ports, static route entries, and dynamic routing protocols. The router uses this learned information to make a determination by using routing protocols. Some of the more common routing protocols used include Routing Information Protocol (RIP), Open Shortest Path First (OSPF), Interior Gateway Routing Protocol (IGRP), and Border Gateway Protocol (BGP).

Figure 6-14 Two-Segment Network with a Layer 3 Router

**note**

Routing protocols are used by routers to share information about the network. Routers receive and use the routing protocol information from other routers to learn about the state of the network. Routers can modify information received from one router by adding their own information along with the original information, and then forward that on to other routers. In this way, each router can share its version of the network.

Packet Switching

Layer 3 information is carried through the network in packets, and the transport method of carrying these packets is called packet switching, as illustrated in Figure 6-15.

Figure 6-15 Packet Switching Between Ethernet and Token Ring Network Segments




Figure 6-15 shows how a packet is delivered across multiple networks. Host A is on an Ethernet segment, and Host B on a Token Ring segment. Host A places an Ethernet frame, encapsulating an *Internet Protocol (IP)* packet, on to the wire for transmission across the network.

The Ethernet frame contains a source data link layer MAC address and a destination data link layer MAC address. The IP packet within the frame contains a source network layer *IP address* (TCP/IP network layer address) and a destination network layer IP address. The router maintains a routing table of network paths it has learned, and the router examines the network layer destination IP address of the packet. When the router has determined the destination network from the destination IP address, the router examines the routing table and determines whether a path exists to that network.

In the case illustrated in Figure 6-15, Host B is on a Token Ring network segment directly connected to the router. The router peels off the Layer 2 Ethernet encapsulation, forwards the Layer 3 data packet, and then re-encapsulates the packet inside a new Token Ring frame. The router sends this frame out its Token Ring interface on to the segment where Host B will see a Token Ring frame containing its MAC address and process it.

Note the original frame was Ethernet, and the final frame is Token Ring encapsulating an IP packet. This is called media transition and is one of the features of a network router. When the packet arrives on one interface and is forwarded to another, it is called Layer 3 switching or routing.

Routing Table Lookup

Routers (and Layer 3 switches) perform table lookups determining the next hop (next router or Layer 3 switch) along the route, which in turn determines the output port over which to forward the packet or frame. The router or Layer 3 switch makes this decision based on the network portion of the destination address in the received packet.

This lookup results in one of three actions:

- **The destination network is not reachable**—There is no path to the destination network and no default network. In this case, the packet is discarded.
- **The destination network is reachable by forwarding the packet to another router**—There is a match of the destination network against a known table entry, or to a default route if a method for reaching the destination network is unknown. The first lookup tells the next hop. Then a second lookup is performed to determine how to get to the next hop. Then a final determination of the exit port is reached. The first lookup can return multiple paths, so the port is not known until after the determination of how to get there is made. In either case, the lookup returns the network (Layer 3) address of the next-hop router, and the port through which that router can be reached.



- **The destination network is known to be directly attached to the router**—The port is directly attached to the network and reachable. For directly attached networks, the next step maps the host portion of the destination network address to the data link (MAC) address for the next hop or end node using the ARP table (for IP). It does not map the destination network address to the router interface. It needs to use the MAC of the final end node so that the node picks up the frame from the medium. Also, you are assuming IP when stating that the router uses the *ARP table*. Other Layer 3 protocols, such as Internetwork Packet Exchange (IPX), do not use ARP to map their addresses to MAC addresses.

Routing table lookup in an IP router might be considered more complex than a MAC address lookup for a bridge, because at the data link layer addresses are 48-bits in length, with fixed-length fields—the OUI and ID. Additionally, data-link address space is flat, meaning there is no hierarchy or dividing of addresses into smaller and distinct segments. MAC address lookup in a bridge entails searching for an exact match on a fixed-length field, whereas address lookup in a router looks for variable-length fields identifying the destination network.

IP addresses are 32 bits in length and are made up of two fields: the network identifier and the host identifier, as illustrated in Figure 6-16.

Both the network and host portions of the IP address can be of a variable or fixed length, depending on the hierarchical network address scheme used. Discussion of this hierarchical, or subnetting, scheme is beyond the scope of this book, but suffice to say you are concerned with the fact that each IP address has a network and host identifier.

The routing table lookup in an IP router determines the next hop by examining the network portion of the IP address. After it determines the best match for the next hop, the router looks up the interface port to forward the packets across, as illustrated in Figure 6-17.


Figure 6-16 IP Address Space


Figure 6-17 shows that the router receives the traffic from Serial Port 1 (S1) and performs a routing table lookup determining from which port to forward out the traffic. Traffic destined for Network 1 is forwarded out the Ethernet 0 (E0) port. Traffic destined for Network 2 is forwarded out the Token Ring 0 (T0) port, and traffic destined for Network 3 is forwarded out Serial Port 0 (S0).

**note**

In terms of the Cisco Internet Operating System (IOS) interface, port numbers begin with zero (0), such as serial port 0 (S0). Not all vendors, including Cisco, use ports; some use slots or modules, which might begin with zero or one.



Figure 6-17 Routing Table Lookup Operation



The host identifier portion of the network address is examined only if the network lookup indicates that the destination is on a locally attached network. Unlike data-link addresses, the dividing line between the network identifier and the host identifier is not in a fixed position throughout the network. Routing table entries can exist for network identifiers of various lengths, from 0 bits in length, specifying a default route, to 32 bits in length for host-specific routes. According to IP routing procedures, the lookup result returned should be the one corresponding to the entry that matches the maximum number of bits in the network identifier. Therefore, unlike a bridge, where the lookup is for an exact match against a fixed-length field, IP routing lookups imply a search for the longest match against a variable-length field.

For example, a network host might have both the IP address of 68.98.134.209 and a MAC address of 00-0c-41-53-40-d3. The router makes decisions based on the IP address (68.98.134.209), whereas the switch makes decisions based on the MAC address (00-0c-41-53-40-d3). Both addresses identify the same host on the network, but are used by different network devices when forwarding traffic to this host.

ARP Mapping

Address Resolution Protocol (ARP) is a network layer protocol used in IP to convert IP addresses into MAC addresses. A network device looking to learn a MAC address broadcasts an ARP request onto the network. The host on the network that has the IP address in the request replies with its MAC (hardware) address. This is called ARP mapping, the mapping of a Layer 3 (network) address to a Layer 2 (data link) address.



note

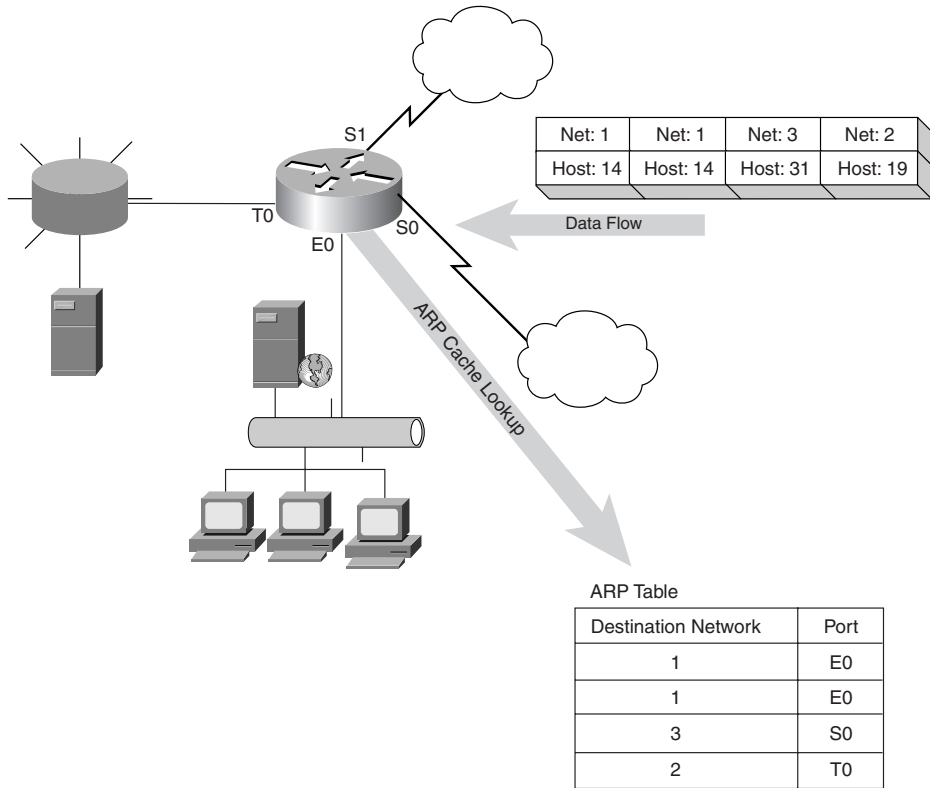
Some Layer 3 addresses use the MAC address as part of their addressing scheme, such as IPX.

Because the network layer address structure in IP does not provide for a simple mapping to data-link addresses, IP addresses use 32 bits, and data-link addresses use 48 bits. It is not possible to determine the 48-bit data-link address for a host from the host portion of the IP address. For packets destined for a host not on a locally attached network, the router performs a lookup for the next-hop router's MAC address. For packets destined for hosts on a locally attached network, the router performs a second lookup operation to find the destination address to use in the data-link header of the forwarded packet's frame, as illustrated in Figure 6-18.

After determining for which directly attached network the packet is destined, the router looks up the destination MAC address in its ARP cache. Recall that ARP enables the router to determine the corresponding MAC address when it knows the network (IP) address. The router then forwards the packet across the local network in a frame with the MAC address of the local host, or next-hop router.



Figure 6-18 Router ARP Cache Lookup



note

Note in Figure 6-18 that Net 3, Host: 31 is not part of the ARP cache, because during the routing table lookup, the router determined that this packet is to be forwarded to another, remote (nonlocally attached) network.

The result of this final lookup falls into one of the three following categories:

- **The packet is destined for the router itself**—The IP destination address (network and station portion combined) corresponds to one of the IP addresses of the router. In this case, the packet must be passed to the appropriate higher-layer entity within the router and not forwarded to any external port.

- **The packet is destined for a known host on the directly attached network**—This is the most common situation encountered by a network router. The router determines the mapping from the ARP table and forwards the packet out the appropriate interface port to the local network.
- **The ARP mapping for the specified host is unknown**—The router initiates a discovery procedure by sending an ARP request determining the mapping of network to hardware address. Because this discovery procedure takes time, albeit measured in milliseconds, the router might drop the packet that resulted in the discovery procedure in the first place. Under *steady-state* conditions, the router already has ARP mappings available for all communicating hosts. The address discovery procedure is necessary when a previously unheard-from host establishes a new communication session.

**note**

The current version of Cisco IOS (12.0) Software drops the first packet for a destination without an ARP entry. The IOS does this to handle denial of service (DoS) attacks against incomplete ARPs. In other words, it drops the frame immediately instead of awaiting a reply.

Fragmentation

Each output port on a network device has an associated maximum transmission unit (MTU). Recall from earlier in this chapter that the MTU indicates the largest frame size (measured in bytes) that can be carried on the interface. The MTU is often a function of the networking technology in use, such as Ethernet, Token Ring, or Point-to-Point Protocol (PPP). PPP is used with Internet connections. If the frame being forwarded is larger than the available space, as indicated by the MTU, the frame is fragmented into smaller pieces for transmission on the particular network.

Bridges cannot fragment frames when forwarding between LANs of differing MTU sizes because data-link connections rarely have a mechanism for fragment reassembly at the receiver. The mechanism is at the network layer implementation, such as with IP, which is capable of overcoming this limitation. Network



layer packets can be broken down into smaller pieces if necessary so that these packets can travel across a link with a smaller MTU.

Fragmentation is similar to taking a picture and cutting it into pieces so that each piece will fit into differently sized envelopes for mailing. It is up to the sender to determine the size of the largest piece that can be sent, and it is up to the receiver to reassemble these pieces. Fragmentation is a mixed blessing; although it provides the means of communication across different link technologies, the processing accomplishing the fragmentation is significant and could be a burden on each device having to fragment and reassemble the data. Further, pieces for reassembly can be received out of order and may be dropped by the switch or router.

As a rule, it is best to avoid fragmentation in your network if at all possible. It is more efficient for the sending station to send packets not requiring fragmentation anywhere along the path to the destination, instead of sending large packets requiring intermediate routers to perform fragmentation.



note

Hosts and routers can learn the maximum MTU available along a network path through the use of MTU discovery. MTU discovery is a process by which each device in a network path learns the MTU size that the network path can support.

Chapter Summary

One of three transmission methods is used to move frames from source to destination: unicast, multicast, or broadcast. Unicast transmission occurs when there is a direct path from source to destination, a “one-to-one” relationship. Multicast has a one-to-many relationship in which the frame is delivered to multiple destinations that are identified as part of a multicast group. Broadcast is a one-to-all relationship in which the frame is delivered to all the hosts on the network segment, whether or not they want the traffic.

Frame size is measured in bytes and has a minimum and maximum length, depending on the implemented technology, such as Ethernet, Token Ring, or with WAN technologies (such as Frame Relay or IP VPN). The maximum frame length supported by a technology is called the maximum transmission unit, or MTU, and is measured in bytes. A frame received by the switch that is less than the minimum frame length for that technology is called a runt, and a frame greater than the maximum frame length is called a giant. Giant frames must be fragmented into smaller frames, smaller than the acceptable MTU, before these frames can be forwarded across the switch's or router's network interface.

There are two common categories of switches: store-and-forward switches and cut-through switches. Store-and-forward switching accepts the complete frame into the switch buffers for error checking before forwarding on to the network. Cut-through switching reads just the destination MAC address (the first 6 bytes of the frame following the preamble) to determine the switch port to forward the traffic. Store-and-forward switching adds some delay to the time it takes for the frame to get from source to destination; unlike cut-through switching, however, store-and-forward switching does not forward a frame with errors. The delay added by store-and-forward switching is minimal and should not be a determining factor when deciding between using cut-through and store-and-forward switching. Store-and-forward has an advantage over cut-through switching by virtue of its error-handling mechanisms.

A third switching category is fragment-free switching, which accepts the first 64 bytes of the frame and checks for errors. Fragment-free switching works on the precept that if there are any errors on the line, they are detectable within the first 64 bytes of the frame.

The fundamental difference between Layer 2 and Layer 3 switch operation is the layer at which each forwarding decision is made. Layer 2 switches make their forwarding decisions based on tables that store the mapping between MAC addresses and switch ports. Layer 3 switches build a table of network addresses and switch



ports, making the forwarding decisions based on the network address information found in Layer 3, rather than just the MAC address found in Layer 2. Layer 3 switches function like routers because of the similar Layer 3 forwarding decision handling. However, Layer 3 switches tend to have better throughput because of the hardware processing of the address tables rather than the software.

Chapter Review Questions

1. What is unicast and how does it work?
2. What is multicast and how does it work?
3. What is broadcast and how does it work?
4. What is fragmentation?
5. What is MTU? What's the MTU for traditional Ethernet?
6. What is a MAC address?
7. What is the difference between a runt and a giant, specific to traditional Ethernet?
8. What is the difference between store-and-forward and cut-through switching?
9. What is the difference between Layer 2 switching and Layer 3 switching?
10. What is the difference between Layer 3 switching and routing?