Global Knowledge®

Expert Reference Series of White Papers

The Packet Delivery Process: Remotely Connected Hosts

The Packet Delivery Process: Remotely Connected Hosts

Alan Thomas, CCNA, CCSI, Global Knowledge Instructor

Introduction

Moving data from one networked device to another requires several different functions. Each function has its own protocol or protocols that define how it is accomplished. Also, the process of delivering data from one device to another can vary. The main factor in data delivery is determining whether the two devices are directly connected or remotely connected.

This paper is focused on the packet delivery process when two devices are remotely connected. Definitions are first, followed by a detailed look at the actual process of exchanging data.

Definitions

Before looking at the actual process of delivering data from one computer to another, a few terms need to be defined:

OSI Network Model

The Open Systems Interconnect (OSI) model is a set of guidelines that define the communications process. It is comprised of seven layers. Each layer has a specific function. How that function is accomplished can vary from one vendor to another, but the basic function must be performed. This layered approach ensures interoperability between vendors, makes the learning process more manageable, and provides a logical framework for troubleshooting.

The seven layers of the OSI model are:

7) Application
6) Presentation
5) Session
4) Transport
3) Network
2) Data Link
1) Physical

In order for communications to take place, each lower layer must be functioning properly. In other words, in order for the Layer 2 function to operate, Layer 1 must be in place and functioning properly. In order for Layer 4 to function properly, Layers 1, 2, and 3 must be in place and functioning properly.

TCP/IP Network Model

The Transmission Control Protocol (TCP)/Internet Protocol (IP) model is a specific implementation of a communications process. It uses four layers instead of seven. TCP/IP is the model used in network communications today.

The four layers of the TCP/IP model are:

4) Application (it combines the functions of Layers 5-7 of the OSI model)
3) Transport (the same as Layer 4 of the OSI model)
2) Network (the same as Layer 3 of the OSI model)
1) Network Access (also called the Link Layer) (combines the functions of Layers 1-2 of the OSI model)

TCP/IP is the set of protocols that are used in network communications. However, it is common to refer to devices and functions based on the OSI layer functions. For example, TCP is a protocol that operates at the Transport Layer, so TCP is often referred to as a Layer 4 protocol, because the Transport Layer is the fourth layer of the OSI model.

Broadcast Domain

A broadcast domain is a collection of devices that receive broadcast packets and broadcast frames from each other. A broadcast packet or frame is one that is destined (addressed) to every device within a broadcast domain.

A broadcast domain is the most basic of networks and may contain only two devices or hundreds of devices. Devices located within the same broadcast domain are considered to be directly connected and can exchange data with each other through a switch. Devices located in different broadcast domains are considered to be remotely connected and require a router to exchange data.

The terms Virtual Local Area Network (VLAN) and subnet are also used to describe a broadcast domain.

Ethernet

Ethernet is a Layer 2 protocol. It defines how devices access the physical segment to which they are connected. Ethernet uses 48-bit addresses to uniquely identify devices on the network.

MAC Address

The Media Access Control (MAC) address is the address used for data exchanged between devices connected to the same broadcast domain. It is a 48-bit number represented in hexadecimal format. A MAC address is assigned to a network adaptor by the manufacturer and does not change. This means the MAC address of a device will be the same regardless of the broadcast domain to which the device might be connected. The MAC address is also called the hardware address, the physical address, or the Layer 2 (OSI Model) address.

IP Address	The Internet Protocol (IP) address is the address used for data exchanged between devices connected to different broadcast domains. If using IPv4, the IP address is a 32-bit number represented in dotted decimal format. If using IPv6, the IP address is a 128-bit number represented in hexadecimal format. The IP address is assigned to a network adaptor by network administrators. If a device is moved from one broadcast domain to another, the IP address assigned to the device will need to be changed. The IP address is also called the logical address or the Layer 3 (OSI Model) address.
Subnet Mask	The subnet mask serves three main purposes: a) It delineates network bits from host bits within an IP address. b) It is used by a source device to determine if the destination device is in the same IP network or IP subnet as itself. c) It is used by a router to determine to which IP network a destination IP address belongs so the router can forward the packet out of the appropriate interface.
Default Gateway	The default gateway is the device that provides another device connectivity outside of its own broadcast domain. Think of the default gateway as being like a door to a room. Suppose you are in a room and you need to access another part of the building, or you need to completely leave the building—you must go through the door of the room. A default gateway is the "door" to a broadcast domain. The default gateway is a Layer 3 (OSI Model) device, meaning it is a router or a multilayer switch.

Address Resolution Protocol

Address Resolution Protocol (ARP) is a protocol used to associate, or map, IP addresses to MAC addresses.

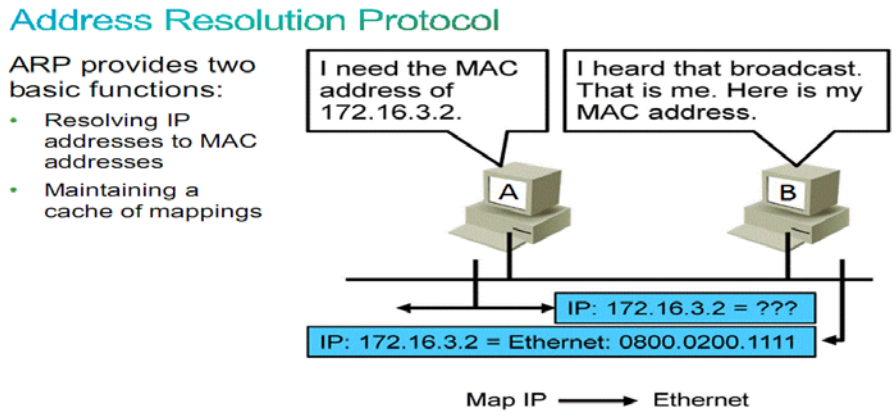

The figure above shows this process. Host A needs to send some data to host B. Host A knows the IP address of host B—172.16.3.2—but in order to deliver the data, host A also needs the MAC address of host B. Host A sends an ARP request that says, "What is the AC address of the device with the IP address 172.16.3.2?" Host B responds with an ARP reply that says, "I have the IP address 172.16.3.2. My MAC address is 0800.0200.1111." Now, host A knows the MAC address for host B, and the data can be delivered.

ARP also maintains a listing of those associations called a cache. By caching the IP address to MAC address associations, if the sending host needs to send more data to the same destination host, there is no need to perform another ARP request. This speeds up the delivery of the data.

Switch

A switch is a Layer 2 (OSI Model) device that forwards data based on the destination MAC address. A switch contains a MAC address table, which is an association of MAC addresses to interfaces. When data arrives at the switch, the destination MAC address is identified. The MAC address table is then consulted. If the destination MAC address is in the switch's MAC address table, the data is forwarded out of the appropriate interface. If the destination MAC address is NOT in the switch's MAC address table, the data is flooded out of every interface EXCEPT for the interface on which the data was received.

Router

A router is a Layer 3 (OSI Model) device that forwards data based on the destination IP address. A router contains a routing table, which is an association of an IP network to an interface. When data arrives at the router, the destination IP address is identified. The routing table is then consulted. If the destination IP address belongs to an IP network in the router's routing table, the data is forwarded out of the appropriate interface. If the router cannot identify a path for the data to take, the data is discarded.

Data Encapsulation

Computer networks use the Transmission Control Protocol/Internet Protocol (TCP/IP) model to exchange data between devices. TCP/IP is a four-layer communications model, with each layer having a specific function.

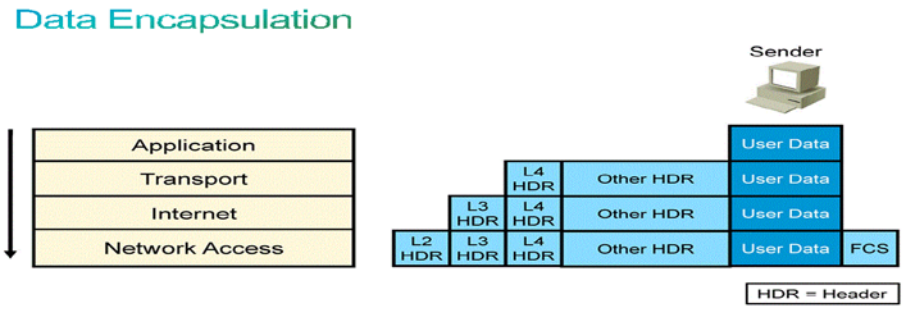

The figure above shows the data encapsulation process. On a sending device, data is generated by the application. It is then sent to the Transport Layer where a header (L4 HDR) is added. The header contains information specific to, and necessary for, the Transport Layer to perform its function. The data continues to move down through the layers, with each layer adding its own header.

Data De-Encapsulation

Computer networks use the Transmission Control Protocol/Internet Protocol (TCP/IP) model to exchange data between devices. TCP/IP is a four-layer communications model, with each layer having a specific function.

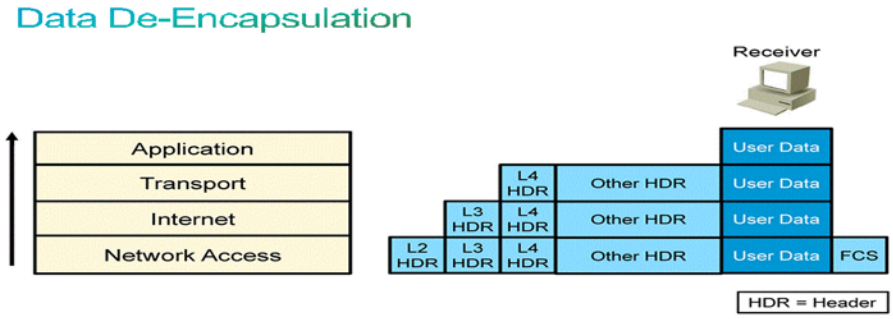

The figure above shows the de-encapsulation process. On a receiving device, data is received from the network and processed by the Network Access Layer (also called the Link Layer). Part of that process is the removal of the header (L2 HDR). The data continues to move up through the layers, with each layer removing the appropriate header.

PDU

A Protocol Data Unit (PDU) is a data construct formed by each layer of the communications process.

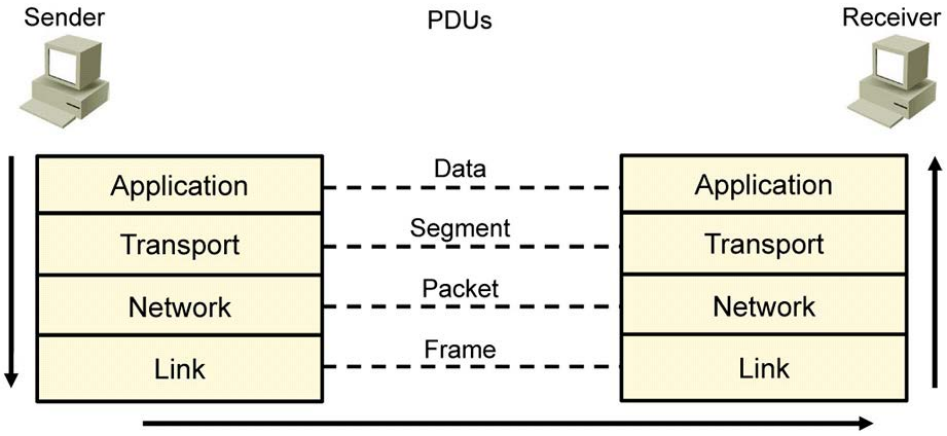

The figure above shows each layer's PDU, and its name. At the Application Layer, the PDU is called Data. At the Transport Layer, the PDU is called a segment. At the Internet Layer, the PDU is called a packet. At the Link Layer (also called the Network Access Layer), the PDU is called a frame.

TCP/UDP

Transmission Control Protocol (TCP) and User Datagram Protocol (UDP) are Transport Layer (Layer 4 of the OSI Model) protocols. Both protocols provide segmentation of data, meaning a large piece of data is broken down into smaller pieces, or segments, of data. Segmentation of data is needed for two reasons:

1) Lower layer protocols (such as Ethernet) limit how much data a frame can carry. If the piece of data to be exchanged is larger than this limit, the data needs to be put into segments small enough to fit the limits set by the lower layer protocol.
2) It is easier and more reliable to send many small segments of data rather than sending one large segment of data.

TCP has more features than UDP. TCP is considered a reliable protocol because it can recover lost segments. Also, TCP is considered a connection-oriented protocol because it establishes a connection to the destination device before sending any data. UDP does not provide these features, so it is considered to be unreliable and connectionless.

The Packet Delivery Process

The process of delivering data from one host to another depends on whether or not the sending and receiving hosts are in the same broadcast domain. (Remember, the terms VLAN and subnet describe the term broadcast domain.) If the sending and receiving devices are connected to the same broadcast domain, data can be delivered using switch and MAC addresses. However, if the sending and receiving devices are connected to different broadcast domains, the use of router and IP addresses is required to deliver data. Below are the steps for data delivery between remotely connected hosts.

Step 1

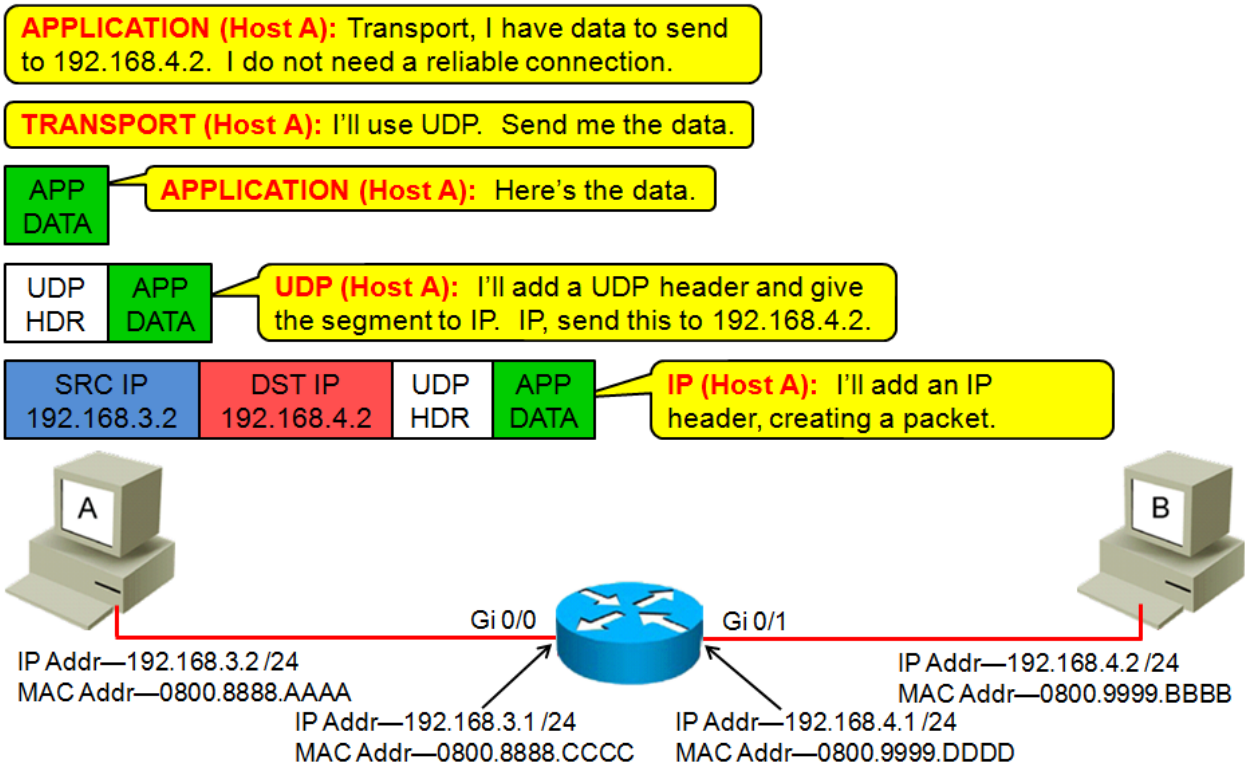

In step 1, Host A generates data destined for Host B. The application tells the Transport Layer (Layer 4) that it has data to send and that it does not need a reliable connection. The Transport Layer chooses UDP and adds a UDP header, creating a Segment. The UDP header contains the source port, which is randomly chosen by Host A, and the destination port, which is based on the application being used. These ports are used to keep track of the communications session and to identify which application is being used. The Segment is then given to the Internet Layer (Layer 3).

Layer 3 (IP) adds an IP header, creating a Packet. The IP header contains the source and destination IP address. In this example, Host A's IP address is the source, and Host B's IP address is the destination.

The process of adding each layer adding its own header onto the existing data is called encapsulation.

Step 2

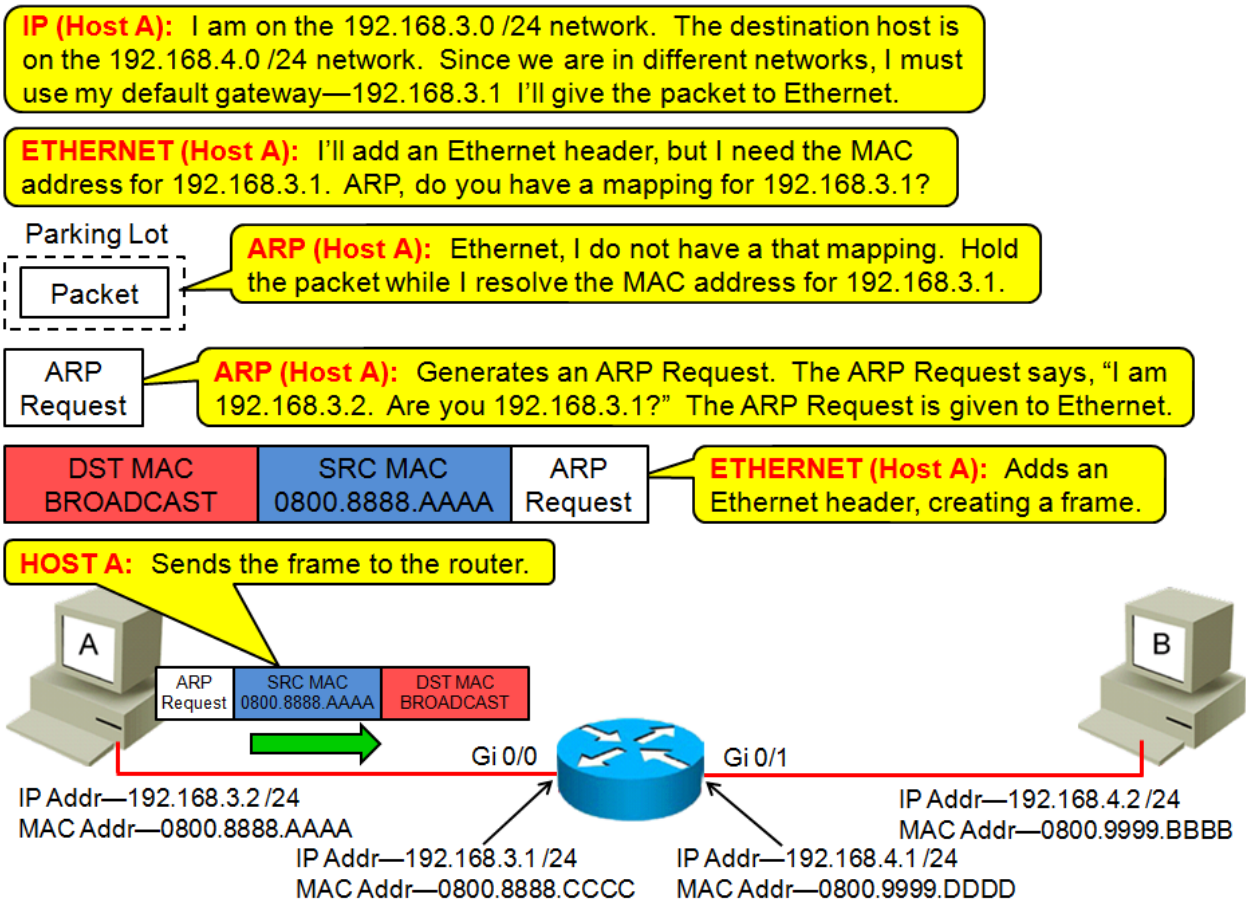

In step 2, IP on Host A determines that Host A is not in the same network (broadcast domain) as Host B. This was accomplished by comparing the IP address and subnet mask of Host A to the destination IP address. Because Hosts A and B are in different IP networks, Host A must send its data to the default gateway. The Internet Layer (Layer 3) gives the Packet to the Link Layer (Layer 2).

At this point, Ethernet wants to add an Ethernet header, but in order to do so, Ethernet needs the MAC address for the router because the router is the default gateway. So Ethernet asks ARP if ARP has the required mapping. ARP does not have the required mapping, so the Packet is placed into the "parking lot"—or queue—and held while ARP resolves the MAC address for the IP address 192.168.3.1.

Next, ARP on Host A generates an ARP request. In the ARP request, the MAC address for 192.168.3.1 is requested. Additionally, 192.168.3.2 is identified as the requesting device. ARP gives the ARP request to the Network Access layer (Layer 2), where the Ethernet header is added, creating a Frame.

The Ethernet header contains the source and destination MAC address. In this instance, the destination MAC address field contains the broadcast address, while the source MAC address field contains Host A's MAC address. The reason the destination MAC address is the broadcast address is because by addressing the frame to broadcast, every device on the segment must process the frame. This provides the highest chance that the device with the requested IP address will receive the ARP Request.

The frame containing the ARP request is sent to the router, while the original packet remains in Host A's "parking lot."

Step 3

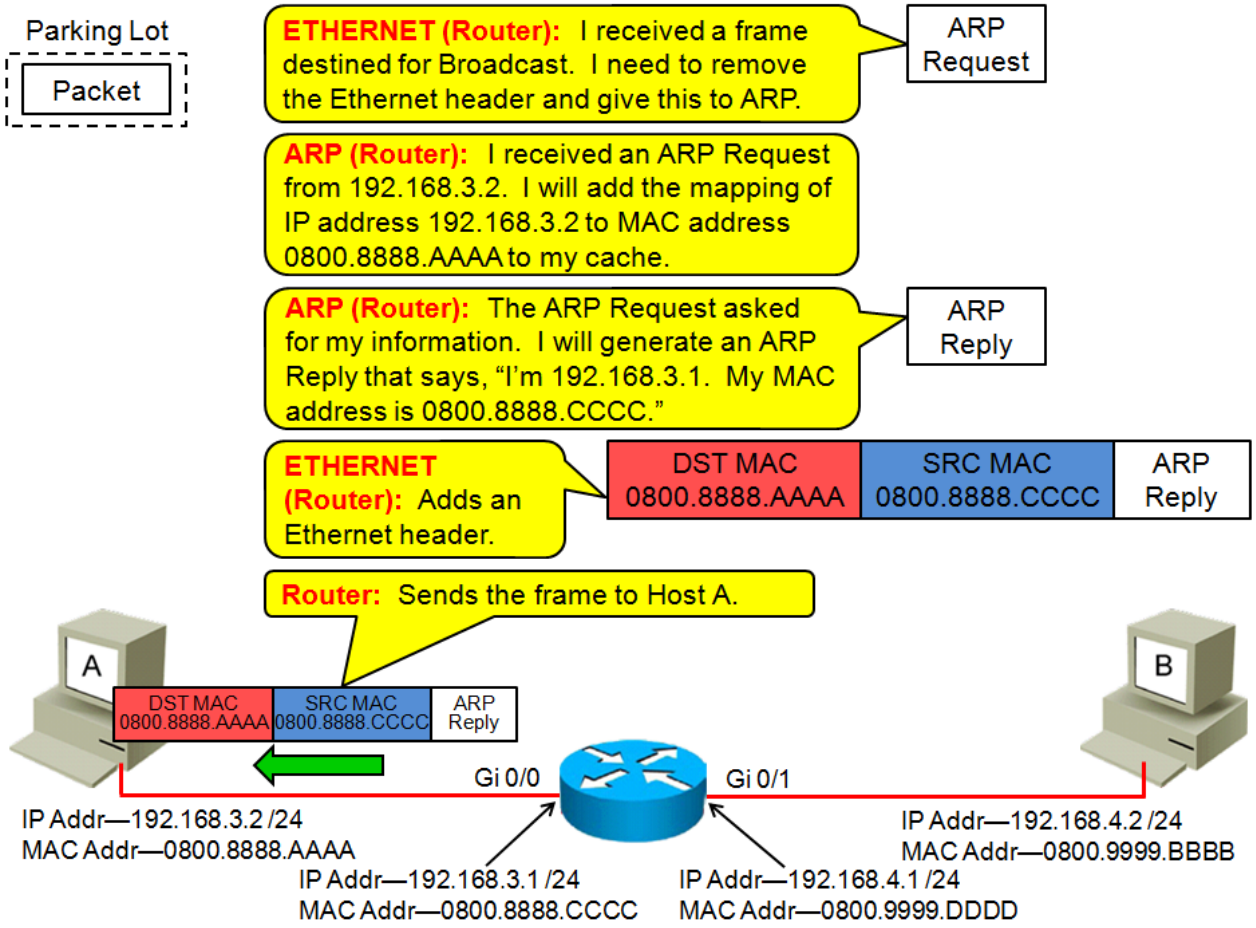

In Step 3, as the router receives the frame, Ethernet identifies the destination MAC address is the broadcast address so it removes the Ethernet header (de-encapsulation) and gives the ARP Request to ARP.

ARP on the router updates its cache to reflect that IP address 192.168.3.2 is associated with MAC address 0800:8888.AAAA. Additionally, ARP determines that the router has the IP address for which the MAC address is being requested, so an ARP Reply is generated.

The ARP Reply says that MAC address 0800:8888:CCCC is associated with IP address 192.168.3.1. The ARP Reply is then given to Layer 2, where the Ethernet header is added. This time, the destination MAC address is 0800:8888:AAAA (Host A's MAC address) and the source MAC address is 0800:8888:CCCC (the router's MAC address).

The frame containing the ARP Reply is sent to host A.

Notice that the original packet is still sitting in Host A's "parking lot."

Step 4

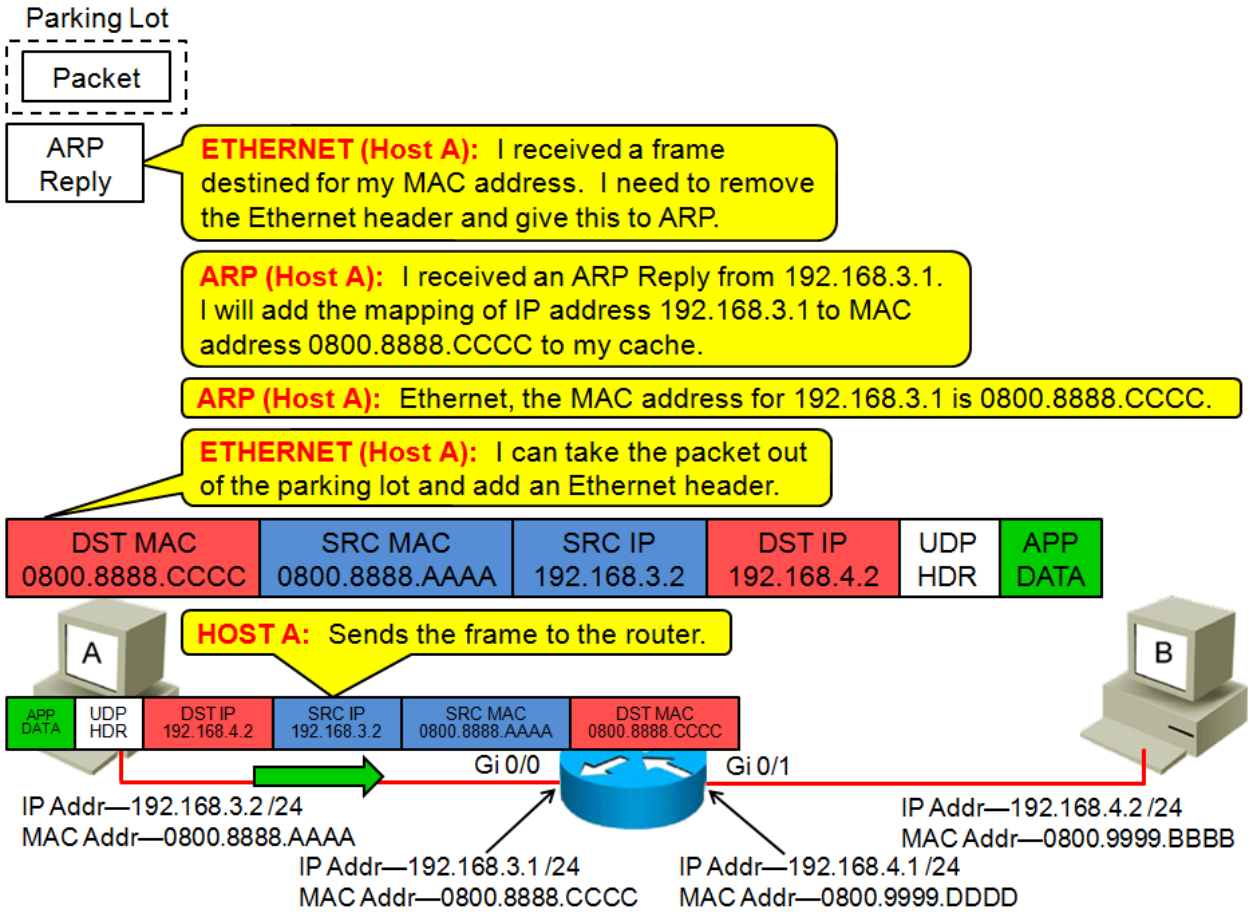

In step 4, Host A receives the frame and Ethernet determines the destination MAC address is Host A's MAC address, so it removes the Ethernet header and gives the ARP Reply to ARP.

ARP on Host A updates its cache to reflect that IP address 192.168.3.1 is associated with MAC address 0800:8888.CCCC. Now ARP is able to provide Ethernet with the router's MAC address.

The original packet is taken out of the "parking lot," and Ethernet adds an Ethernet header, where the router's MAC address is the destination MAC address, and Host A's MAC address is the source MAC address.

Notice that the destination IP address is still Host B's IP address. This is because Host A must deliver the data to a directly connected device. Since the router is directly connected to Host A, Host A can send data to the router, thus the router's MAC address is the destination MAC address. However, the data's ultimate destination is Host B, so Host B's IP address is still the destination IP address in the IP header.

Host A sends the frame to the router.

Step 5

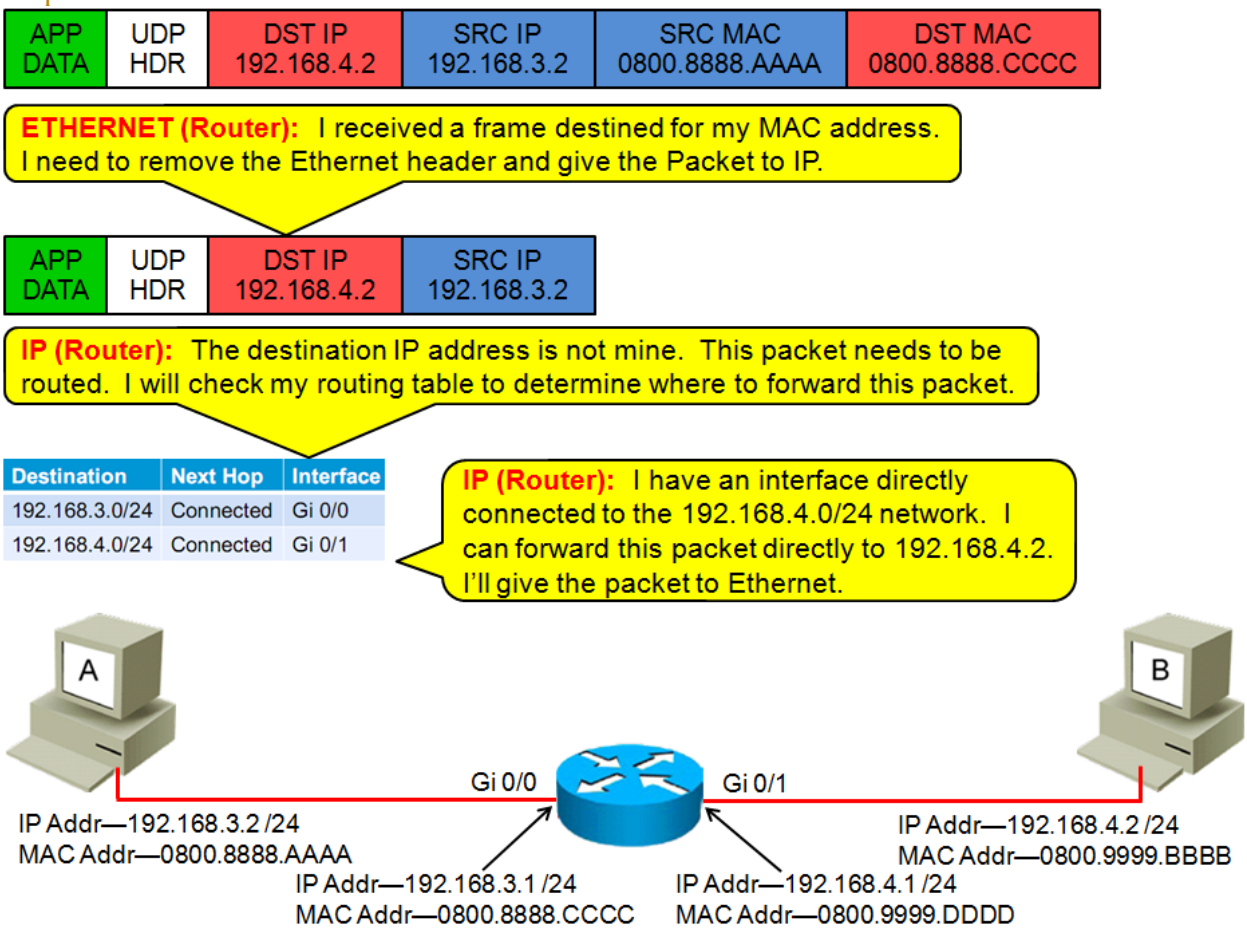

In step 5, the router receives the frame and Ethernet determines the destination MAC address is the router's MAC address, so it removes the Ethernet header and gives the packet to IP.

IP determines that the destination IP address does not belong to the router. This means the packet needs to be routed. So the router evaluates its routing table and discovers it has an interface directly connected to the network to which the destination IP address belongs. This means the router can deliver the data directly to the destination device, so IP gives the packet to Layer 2.

Step 6

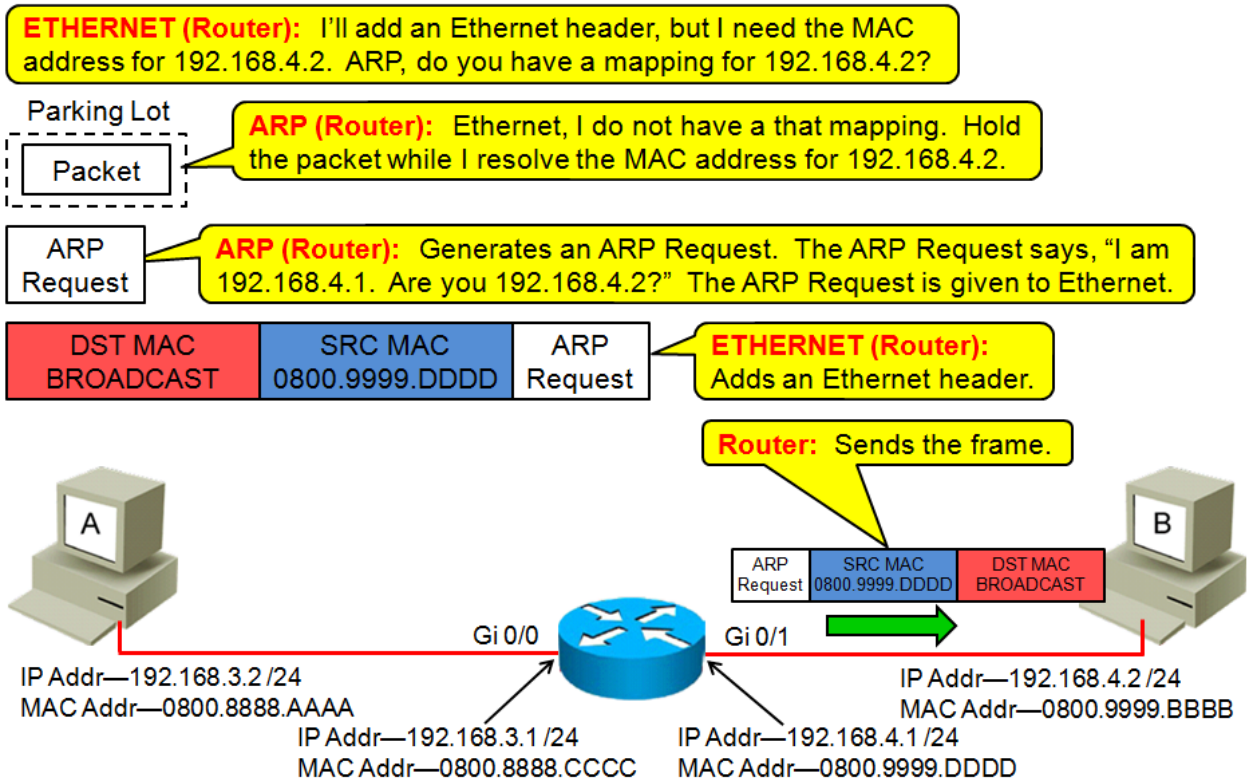

In step 6, Ethernet wants to add an Ethernet header, but in order to do so, Ethernet needs the MAC address for Host B. So Ethernet asks ARP if ARP has the required mapping. ARP does not have the required mapping, so the Packet is placed into the router's "parking lot" and held while ARP resolves the MAC address for the IP address 192.168.4.2.

ARP on the router generates an ARP Request. In the ARP Request, the MAC address for IP address 192.168.4.2 is requested. Additionally, 192.168.4.1 is identified as the requesting device. ARP gives the ARP Request to Layer 2, where the Ethernet header is added.

The Ethernet header contains the source and destination MAC address. In this instance, the destination MAC address field contains the broadcast address, while the source MAC address field contains the router's MAC address.

The frame containing the ARP Request is sent to Host B, while the original packet remains in the router's "parking lot."

Step 7

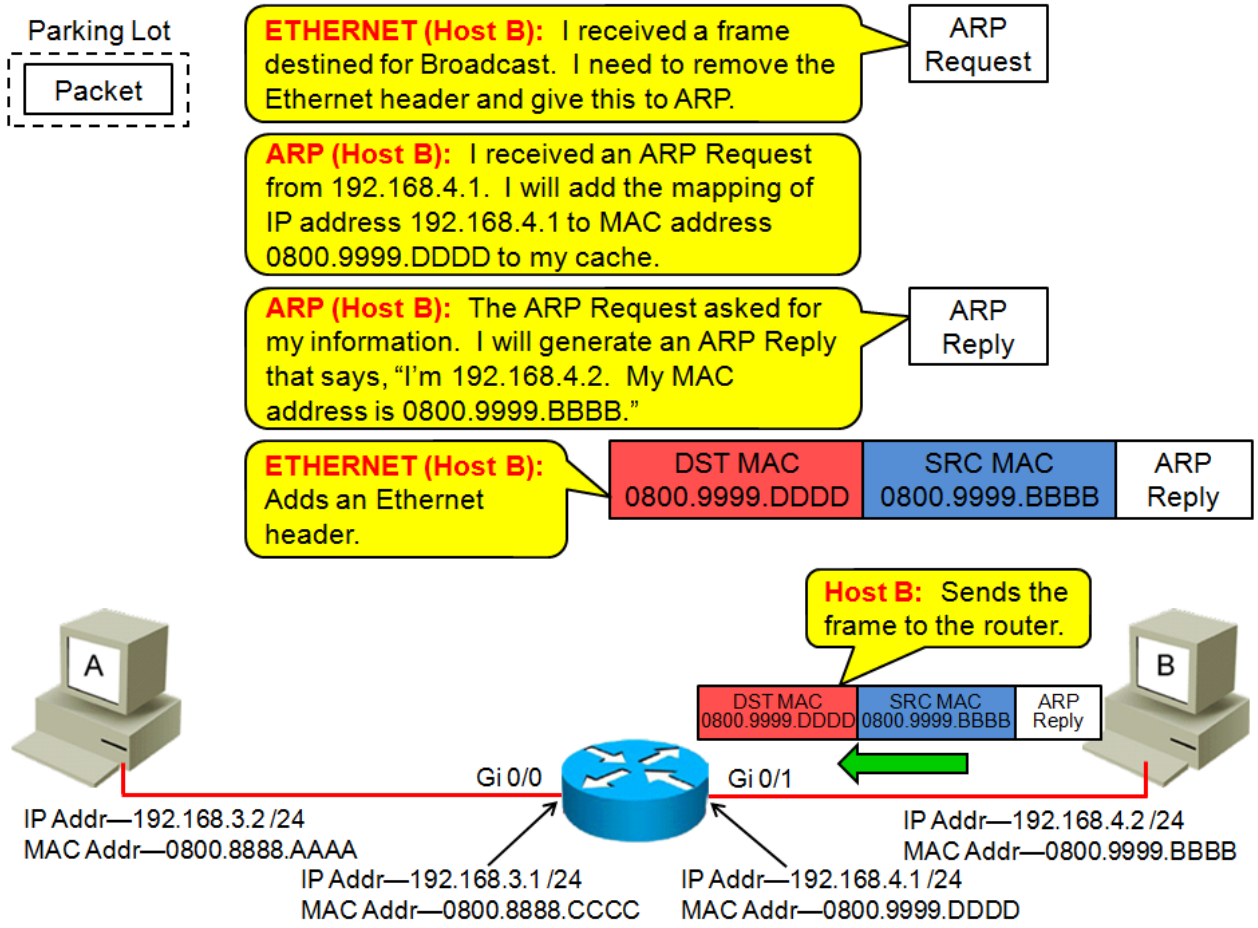

In step 7, as Host B receives the frame, Ethernet identifies the destination MAC address is the broadcast address, so it removes the Ethernet header and gives the ARP Request to ARP.

ARP on Host B updates its cache to reflect that IP address 192.168.4.1 is associated with MAC address 0800:9999.DDDD. Additionally, ARP determines that Host B has the IP address for which the MAC address is being requested, so an ARP Reply is generated.

The ARP Reply says that MAC address 0800:9999:BBBB is associated with the IP address 192.168.4.2. The ARP Reply is then given to Layer 2, where the Ethernet header is added. This time, the destination MAC address is 0800:9999:DDDD (the router's MAC address) and the source MAC address is 0800:9999:BBBB (Host B's MAC address).

The frame containing the ARP Reply is sent to the router.

Notice that the original packet is still sitting in the router's "parking lot."

Step 8

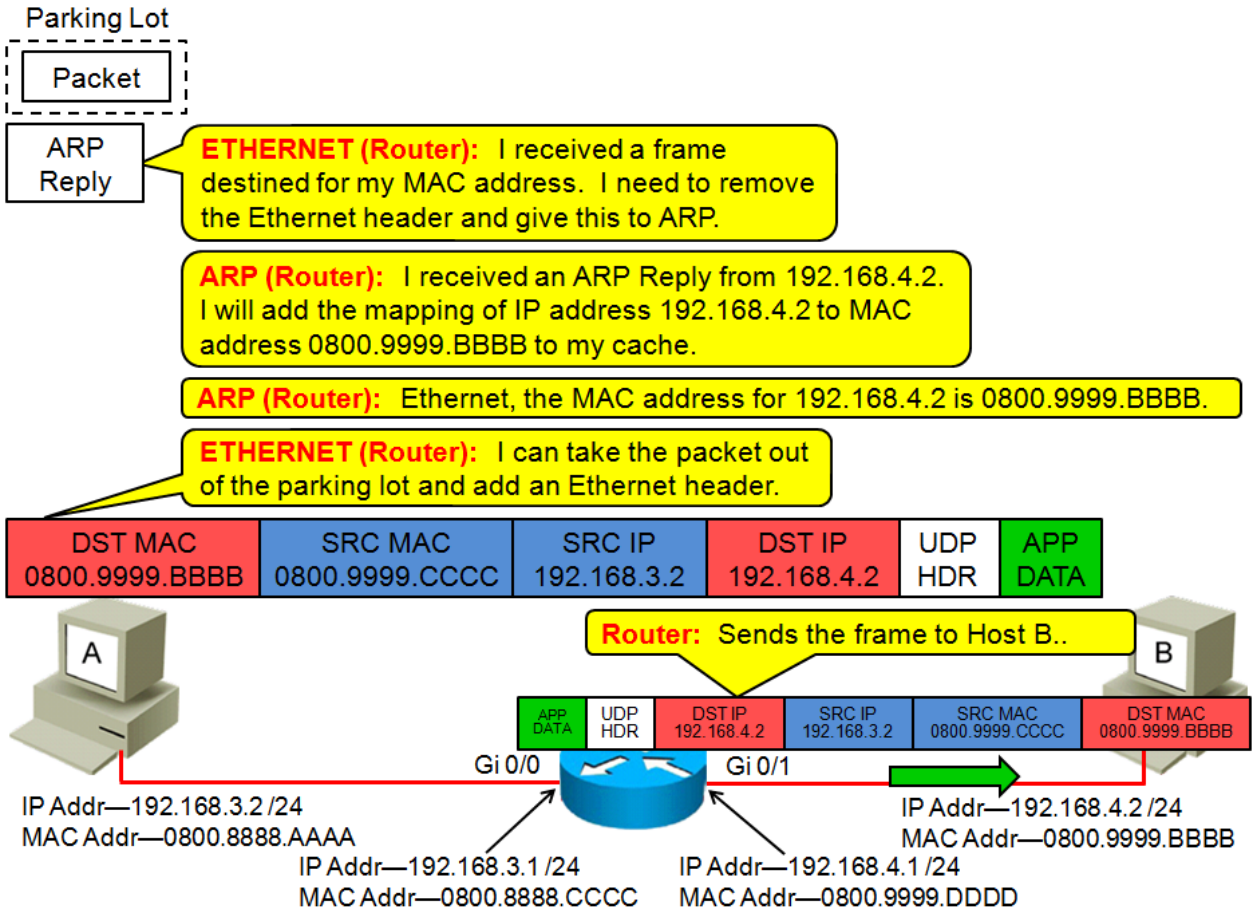

In step 8, as the router receives the frame, Ethernet determines the destination MAC address is the router's MAC address, so it removes the Ethernet header and gives the ARP Reply to ARP.

ARP on the router updates its cache to reflect that IP address 192.168.4.2 is associated with MAC address 0800:9999.BBBB. Now ARP is able to provide Ethernet with Host B's MAC address.

The original packet is taken out of the router's "parking lot," and Ethernet adds an Ethernet header, where Host B's MAC address is the destination MAC address and the router's MAC address is the source MAC address.

Notice that the destination IP address is still Host B's IP address, and the source IP address is still Host A's IP address.

The router sends the frame to Host B.

Step 9

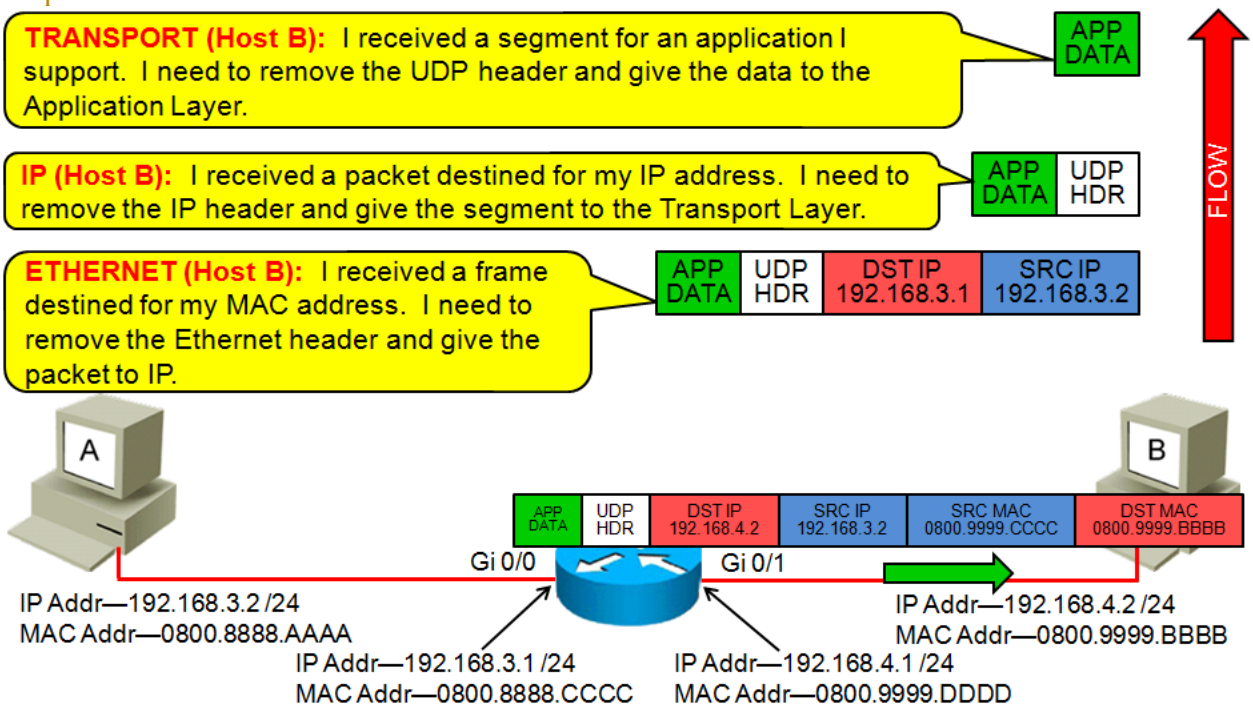

In step 9, read this figure, bottom to top.

As Host B receives the frame, Ethernet identifies the destination MAC address is Host B's MAC address, so it removes the Ethernet header and gives the packet to IP.

IP determines that the destination IP address is Host B's IP address, so the IP header is removed and the Segment is given to the Transport Layer.

The Transport Layer determines the destination port is for an application Host B supports, so the UDP header is removed and the Data is given to the Application.

Conclusion

The process described in the preceding steps is often times called the frame re-write process, because a new frame is created by every Layer 3 device in the path. Notice, however, that the packet remains unchanged end-to-end.

The ARP resolution process only needs to occur once. However, the encapsulation process, the router operations, and the de-encapsulation process continue until all the data is transferred.

The packet delivery process for devices that are remotely connected is an interesting process. It requires the use of multiple protocols, each of which performs a specific function or service, along with a router to make the forwarding decisions based on IP addresses.

Learn More

Learn more about how you can improve productivity, enhance efficiency, and sharpen your competitive edge through training.

- ICND1 v2.0 - Interconnecting Cisco Networking Devices, Part 1
- ICND2 v2.0 - Interconnecting Cisco Networking Devices, Part 2
- SWITCH - Implementing Cisco IP Switched Networks v2.0
- ROUTE - Implementing Cisco IP Routing v2.0
- TSHOOT - Troubleshooting and Maintaining Cisco IP Networks v2.0

Visit **www.globalknowledge.com** or call **1-800-COURSES (1-800-268-7737)** to speak with a Global Knowledge training advisor.

About the Author

Alan Thomas holds a Bachelor of Science degree in technical management and has been a network professional in several capacities for over 20 years. Alan is a Global Knowledge Instructor and has received the Quality Instructor Award.