Switching & ARP
Week 3

Module : Computer Networks
Lecturer: Lucy White lbwhite@wit.ie
Office : 324

Many Slides courtesy of Tony Chen

Ethernet – Using Switches

- In the last few years, switches have quickly become a fundamental part of most networks.
  - •Switches allow the segmentation of the LAN into separate collision domains.
  - •Each port of the switch represents a separate collision domain and provides the full media bandwidth to the node or nodes connected on that port.
  - •With fewer nodes in each collision domain, there is an increase in the average bandwidth available to each node, and collisions are reduced.
- In a LAN where a hub is connected to a switch port, there is still shared bandwidth, which may result in collisions within the shared environment of the hub.
  - •However, the switch will isolate the segment and limit collisions to traffic between the hub's ports.

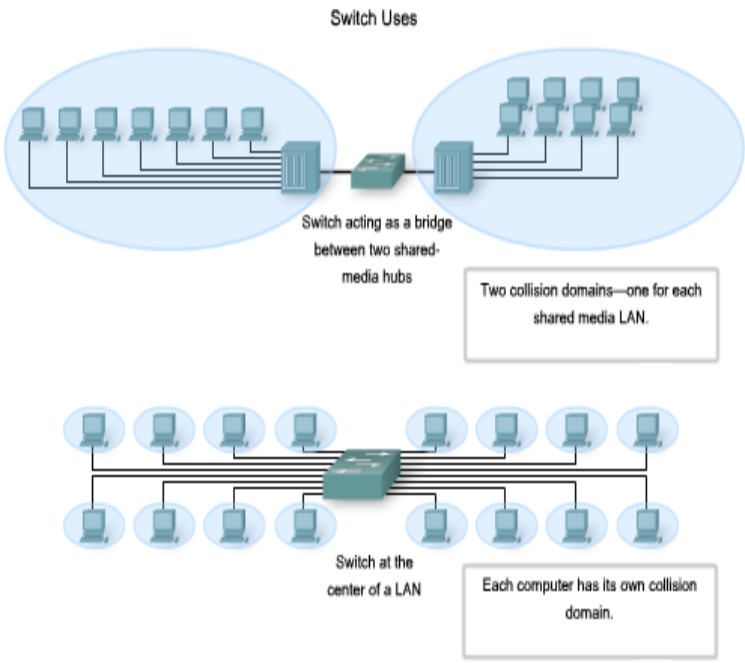

Ethernet – Using Switches

- In a LAN where all nodes are connected directly to the switch, the throughput of the network increases dramatically. These physical star topologies are essentially point to point links.
  - •Dedicated bandwidth to each port
    - With switches, each device effectively has a dedicated point-to-point connection between the device and the switch, without media contention.
  - •Collision-free environment
    - A dedicated point-to-point connection to a switch also removes any media contention between devices, allowing a node to operate with few or no collisions.
      - •In a moderately-sized classic Ethernet network using hubs, approximately 40% to 50% of the bandwidth is consumed by collision recovery.
  - •Full-duplex operation
    - With full-duplex enabled in a switched Ethernet network, the devices connected directly to the switch ports can transmit and receive simultaneously, at the full media bandwidth.

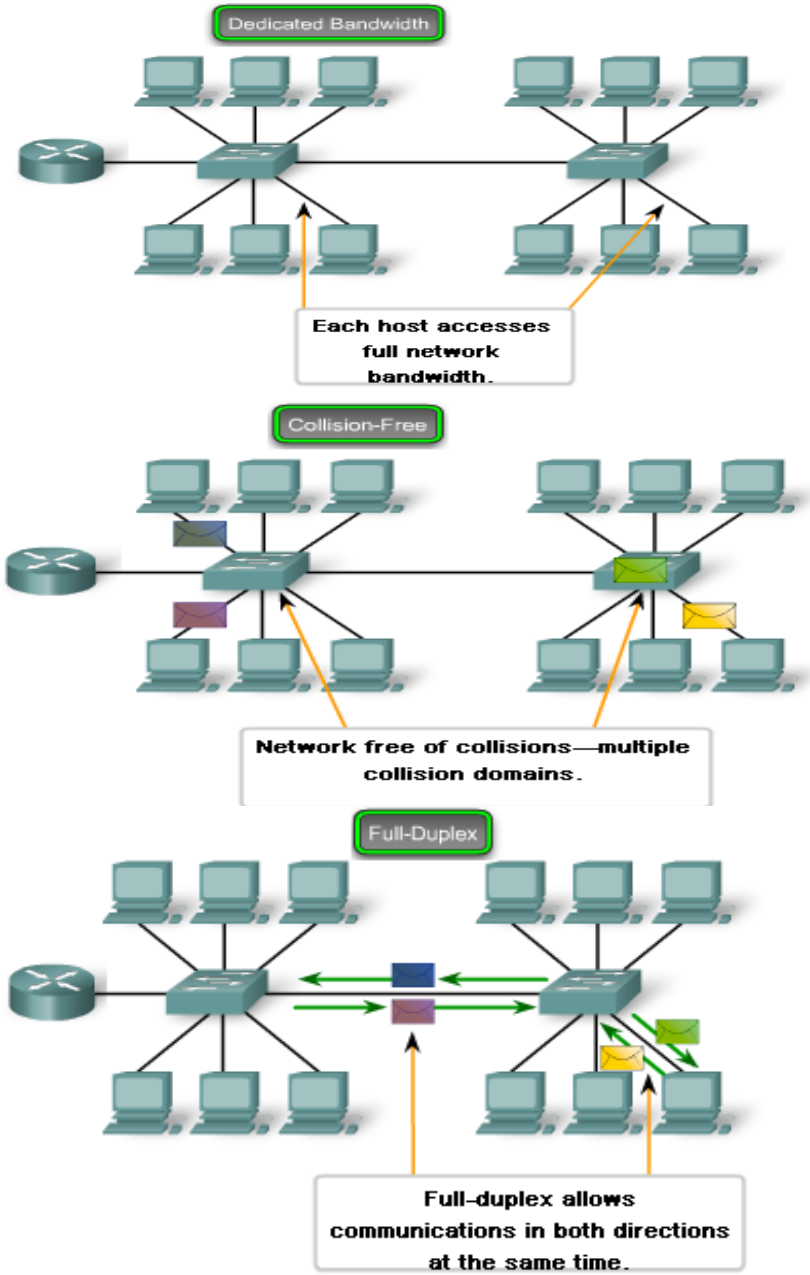

Switches – Selective Forwarding

- Switch forwarding is based on the **Destination** MAC
  - The switch maintains a table, called a MAC table. that matches a destination MAC address with the port used to connect to a node.
  - For each incoming frame, the destination MAC address in the frame header is compared to the list of addresses in the MAC table.
  - If a match is found, the port number in the table that is paired with the MAC address is used as the exit port for the frame.
- The MAC table can be referred to by many different names.
  - It is often called the switch table.
  - Because switching was derived from transparent bridging, the table is sometimes called the bridge table.

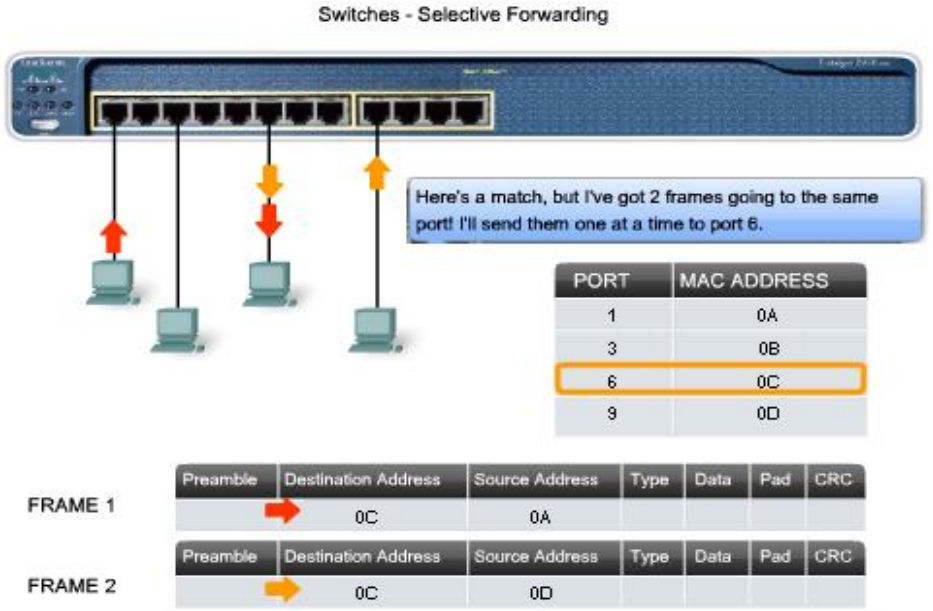

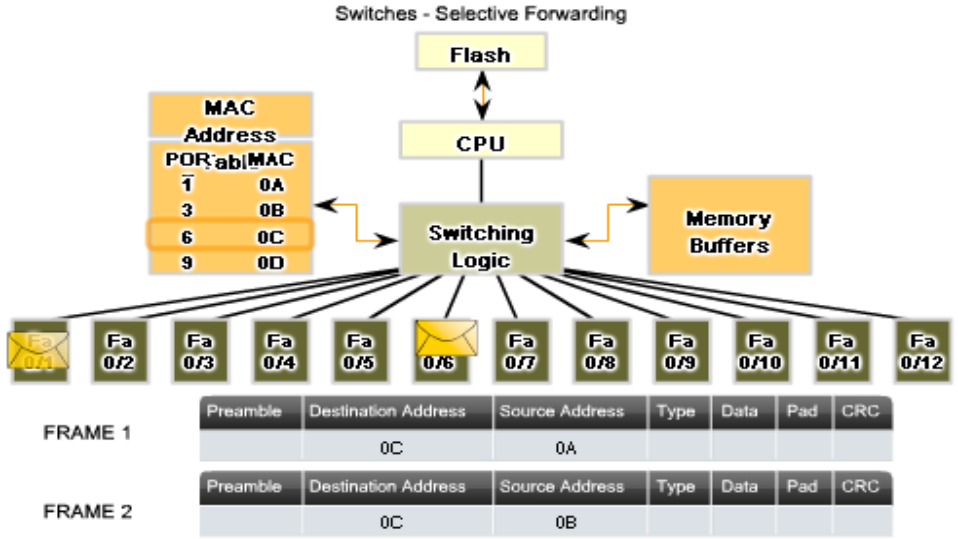

Switches – store and forward

- Any node operating in full-duplex mode can transmit anytime it has a frame, without regard to the availability of the receiving node.
  - •This is because a LAN switch will buffer an incoming frame and then forward it to the proper port when that port is idle.
  - •This process is referred to as store and forward.
- With store and forward switching, the switch receives the entire frame, checks the FCS for errors, and forwards the frame to the appropriate port for the destination node.
  - •Because the nodes do not have to wait for the media to be idle, the nodes can send and receive at full media speed without losses due to collisions or the overhead associated with managing collisions.

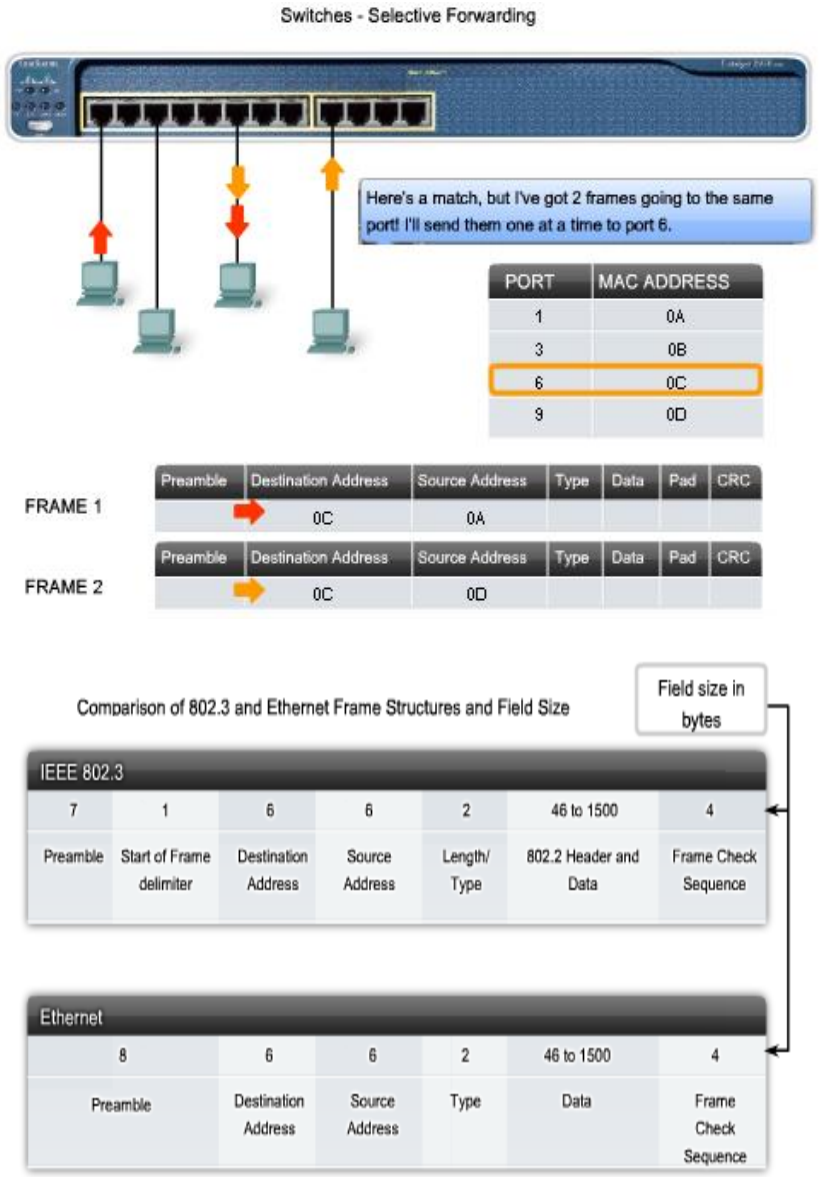

MAC Address Tables on Connected Switches

- A switch can have multiple MAC addresses associated with a single port.
- This occurs when the switch is connected to another switch.
- Video demonstration 4.2.1.2

Sending a Frame to the Default Gateway

When a device has an IP address that is on a remote network, the Ethernet frame cannot be sent directly to the destination device.

The Ethernet frame is sent to the MAC address of the default gateway, which is the router.

Switch Operation

- Ethernet LAN switches use 5 basic operations:

1. Learning

- The MAC table must be populated with MAC addresses and their corresponding ports.
- The Learning process allows these mappings to be dynamically acquired during normal operation.
- As each frame enters the switch, the switch examines the source MAC address.
  - If no entry exists, the switch creates a new entry in the MAC table using the source MAC address and pairs the address with the port on which the entry arrived.
- The switch now can use this mapping to forward frames to this node.

2. Aging

- The entries in the MAC table are time stamped.
- After entry made in MAC table, a countdown begins.
- After the value reaches 0, the entry in the table will be removed.

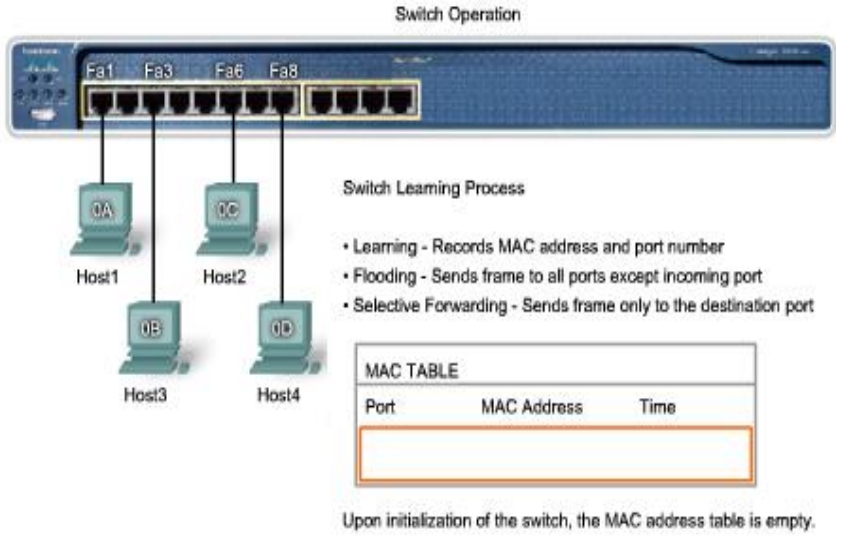

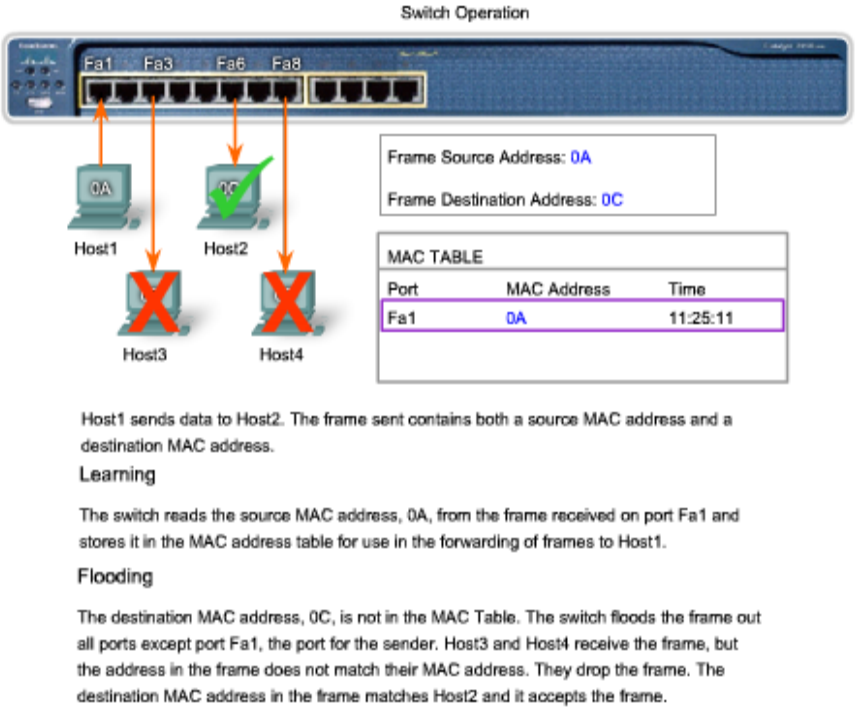

Switch Operation

- Flooding
  - •If the switch does not know to which port to send a frame because the destination MAC address is not in the MAC table, the switch sends the frame to all ports except the port on which the frame arrived.
  - •Flooding is also used for frames sent to the broadcast MAC address.
- Selective Forwarding
  - •Selective forwarding is the process of examining a frame's destination MAC address and forwarding it out the appropriate port.
- Filtering
  - •One use of filtering has already been described: a switch does not forward a frame to the same port on which it arrived.
  - •A switch will also drop a corrupt frame. If a frame fails a CRC check, the frame is dropped.
  - •An additional reason for filtering a frame is security.
  - •A switch has security settings for blocking frames to and/or from selective MAC addresses.

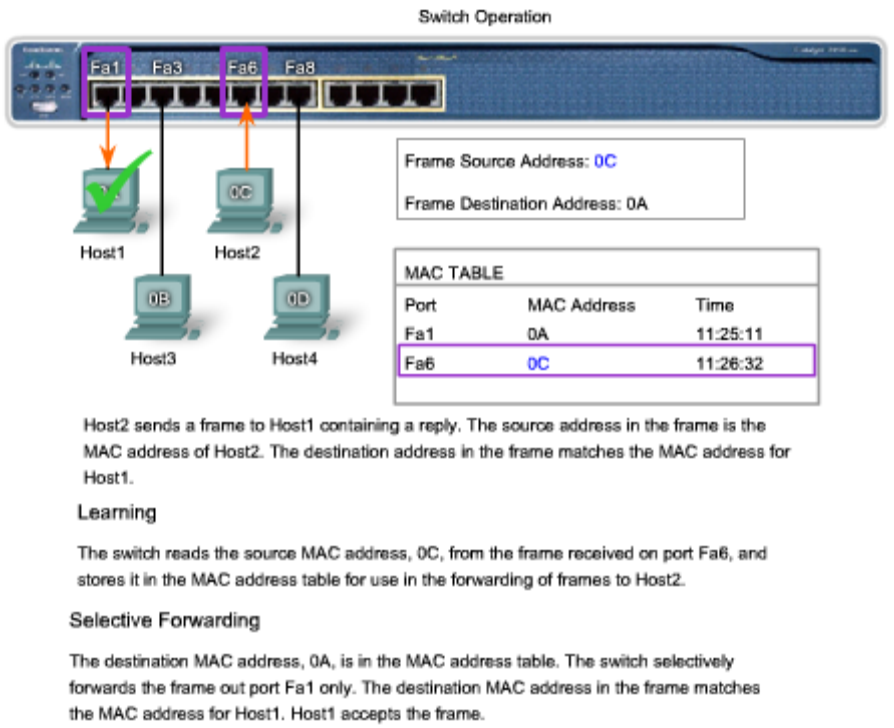

Switches – Activity: page 4.2.1.7

Please go to the page and do more exercises, until you competently understand the topics.

Activity
Determine how the switch forwards a frame based on the Source MAC and Destination MAC addresses and information in the switch MAC table.
Answer the questions below using the information provided.

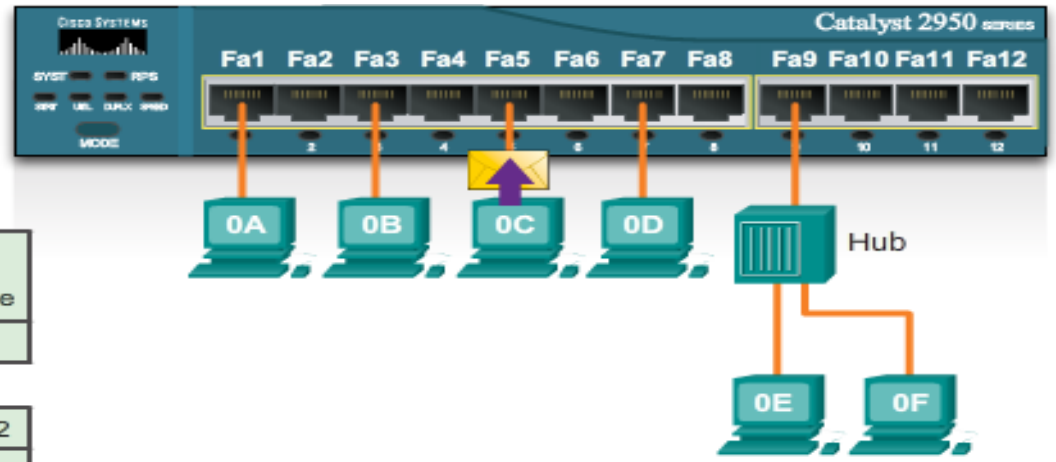

Frame

Preamble	Destination MAC	Source MAC	Length Type	Encapsulated Data	End of Frame
	FF	0C			

MAC Table

Fa1	Fa2	Fa3	Fa4	Fa5	Fa6	Fa7	Fa8	Fa9		Fa10	Fa11	Fa12
0A												

Question 1 - Where will the switch forward the frame?

☐ Fa1 ☐ Fa2 ☐ Fa3 ☐ Fa4 ☐ Fa5 ☐ Fa6 ☐ Fa7 ☐ Fa8 ☐ Fa9 ☐ Fa10 ☐ Fa11 ☐ Fa12

Question 2 - When the switch forwards the frame, which statement(s) are true?

- ☐ Switch adds the source MAC address to the MAC table.
- ☐ Frame is a broadcast frame and will be forwarded to all ports.
- ☐ Frame is a unicast frame and will be sent to specific port only.
- ☐ Frame is a unicast frame and will be flooded to all ports.
- ☐ Frame is a unicast frame but it will be dropped at the switch.

Check

New Problem

Help

5.3.1.9

MAC and IP

MAC address

- This address does not change
- Similar to the name of a person
- Known as physical address because physically assigned to the host NIC

IP address

- Similar to the address of a person
- Based on where the host is actually located
- Known as a logical address because assigned logically
- Assigned to each host by a network administrator

Both the physical MAC and logical IP addresses are required for a computer to communicate just like both the name and address of a person are required to send a letter

Introduction to ARP

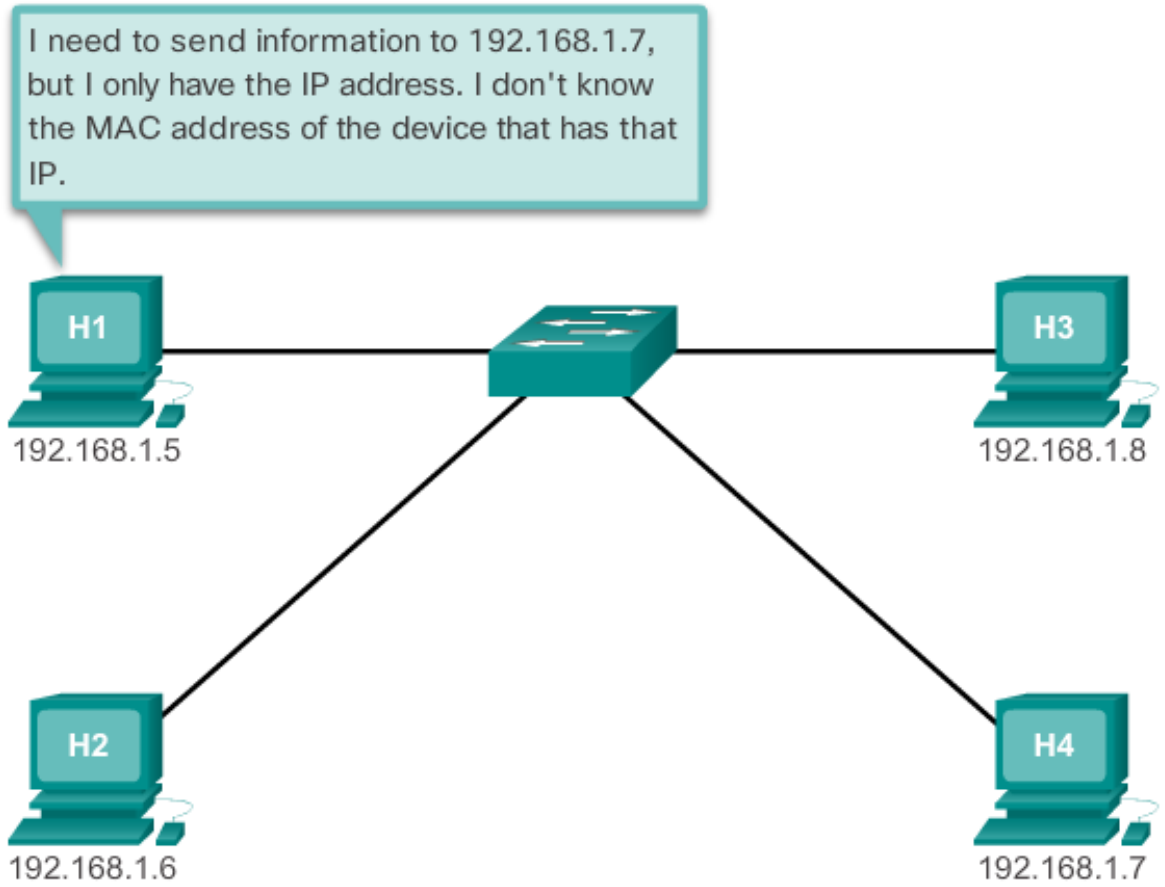

The ARP Process – Mapping IP to MAC Address

- The ARP protocol provides two basic functions:
- Resolving IPv4 Addresses to MAC Addresses
  - •For a frame to be placed on the LAN media, it must have a destination MAC address.
  - •When a packet is sent to the Data Link layer to be encapsulated into a frame, the node refers to a table in its memory to find the Data Link layer address that is mapped to the destination IPv4 address.
  - •This table is called the ARP table or the ARP cache.
  - •The ARP table is stored in the RAM of the device.
- Maintaining the ARP Table
  - •There are 2 ways that a device can gather MAC addresses.
    - **One way is to monitor the traffic occurs on the local segment.**
    - **Another way is to broadcast an ARP request.**
  - •ARP sends a Layer 2 broadcast to all devices on the Ethernet LAN. The frame contains an ARP request packet with the IP address of the destination host.
    - **The node receiving the frame that identifies the IP address as its own responds by sending an ARP reply packet back to the sender as a unicast frame. This response is then used to make a new entry in the ARP table.**
  - •These dynamic entries in the MAC table are timestamped.

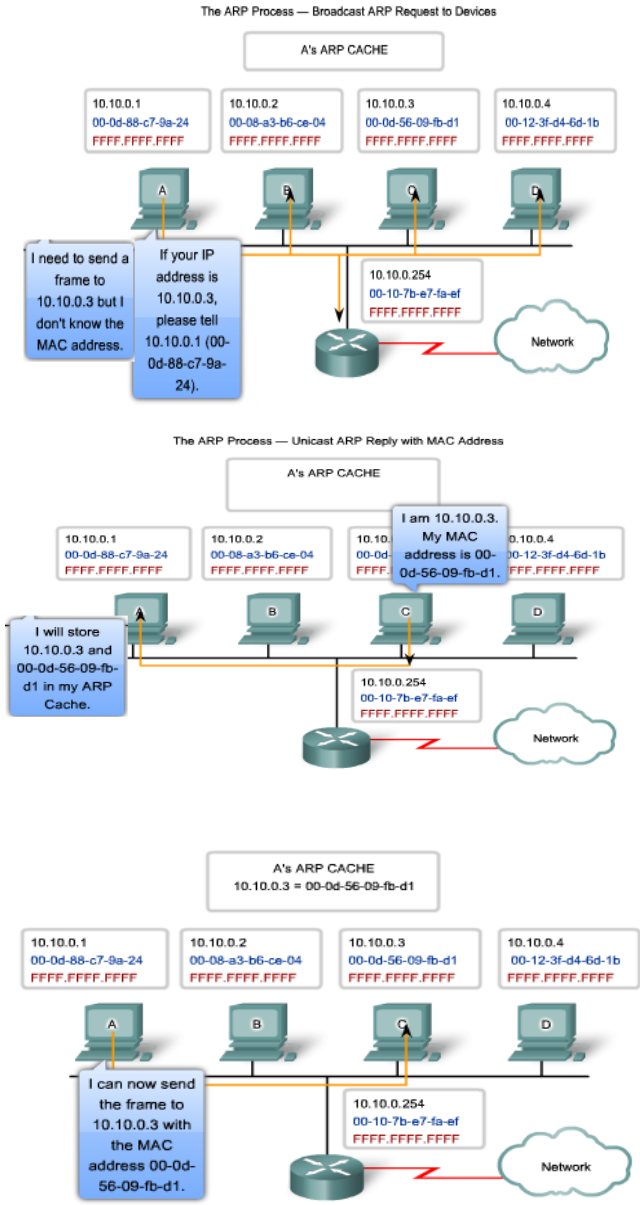

ARP Process – Destinations not on the local Network

- If the destination IPv4 host is not on the local network, the source node needs to deliver the frame to the router interface that is the gateway or next hop used to reach that destination.
  - The source node will use the MAC address of the gateway as the destination address for frames containing an IPv4 packet addressed to hosts on other networks.
  - In the event that the gateway entry is not in the table, the normal ARP process will send an ARP request to retrieve the MAC address associated with the IP address of the router interface.

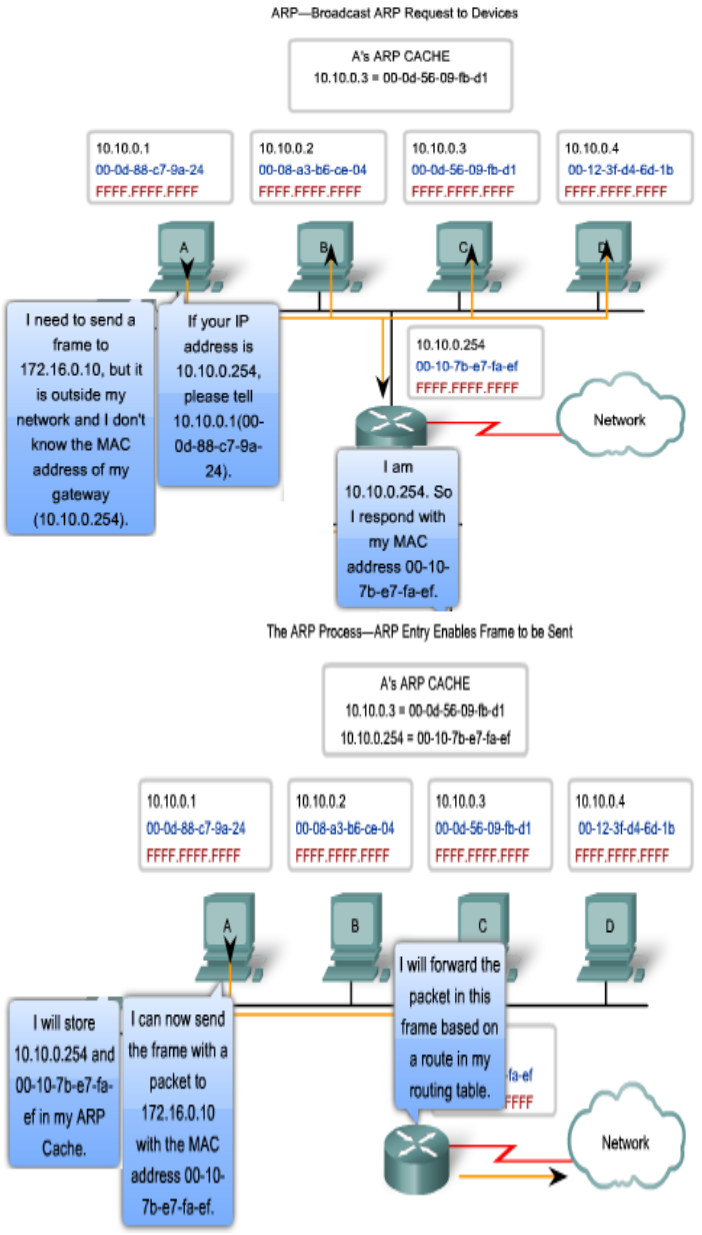

Proxy ARP– Destinations not on the local Network

- There are circumstances under which a host might send an ARP request seeking to map an IPv4 address outside of the range of the local network.
  - In these cases, the device sends ARP requests for IPv4 addresses not on the local network instead of requesting the MAC address with the IPv4 address of the gateway.
  - To provide a MAC address for these hosts, a router use a proxy ARP to respond on behalf of remote hosts.
  - This means that the ARP cache of the requesting device will contain the MAC address of the gateway mapped to any IP addresses not on the local network.
  - If proxy ARP is disabled on the router interface, these hosts cannot communicate out of the local network.
- One such use of this process is
  - IPv4 cannot determine whether the destination host is on the same network as the source.
  - When a host believes that it is directly connected to the same network as the destination host. This generally occurs when a host is configured with an improper mask.
  - When a host is not configured with a default gateway. Proxy ARP can help devices on a network reach remote subnets.

http://www.cisco.com/en/US/tech/tk648/tk361/technologies_tech_note09186a0080094adb.shtml

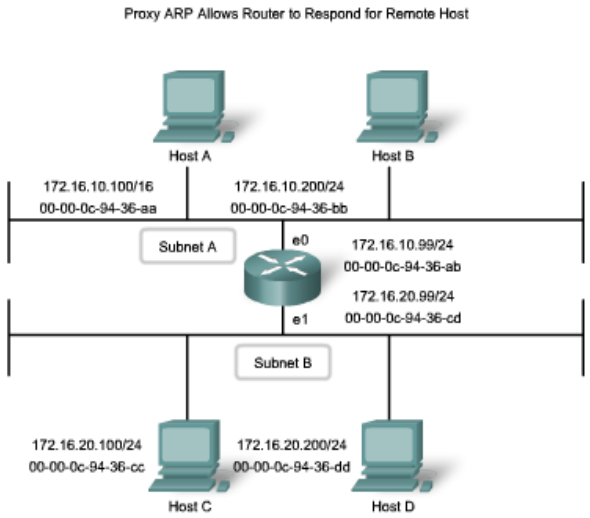

The ARP cache of Host A is shown in this table:

IP Address	MAC Address
172.16.20.200	00-00-0c-94-36-ab
172.16.20.100	00-00-0c-94-36-ab
172.16.10.99	00-00-0c-94-36-ab
172.16.10.200	00-00-0c-94-36-bb

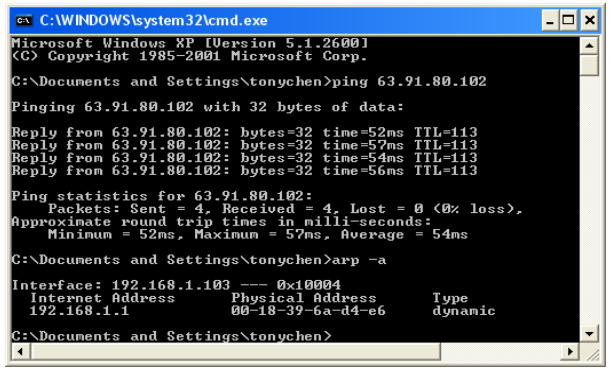

ARP Process – Removing Address Mapping

- For each device, an ARP cache timer removes ARP entries that have not been used for a specified period of time.
  - •The times differ depending on the device and its operating system.
  - •For example, some Windows operating systems store ARP cache entries for 2 minutes. If the entry is used again during that time, the ARP timer for that entry is extended to 10 minutes.
- Commands may also be used to manually remove all or some of the entries in the ARP table.
- After an entry has been removed, the process for sending an ARP request and receiving an ARP reply must occur again to enter the map in the ARP table.

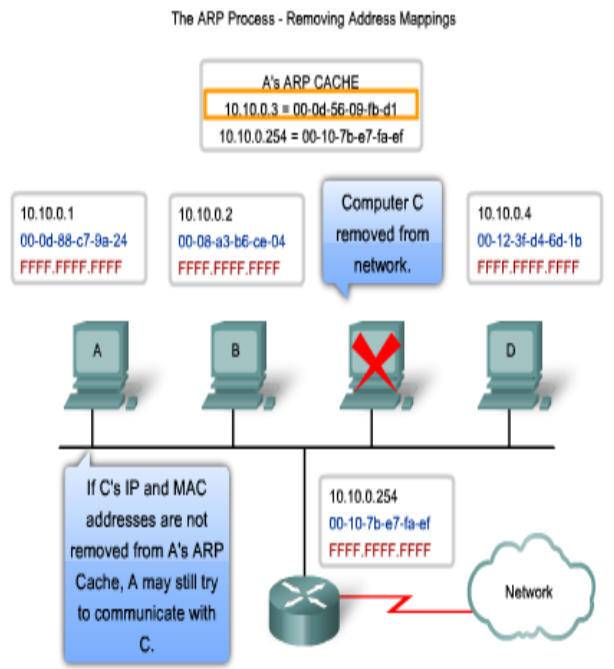

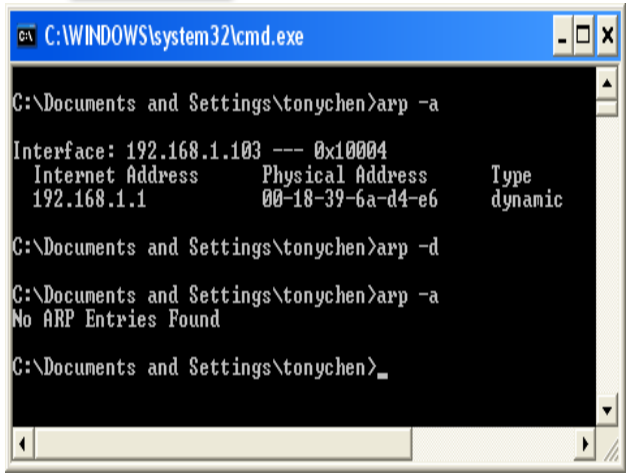

ARP Broadcasts Issues

- Overhead on the Media
  - As a broadcast frame, an ARP request is received and processed by every device on the local network.
  - On a typical business network, these broadcasts would probably have minimal impact on network performance.
- Security
  - In some cases, the use of ARP can lead to a potential security risk. ARP spoofing, or ARP poisoning, is a technique used by an attacker to inject the wrong MAC address association into a network by issuing fake ARP requests.
    - An attacker forges the MAC address of a device and then frames can be sent to the wrong destination.
  - Manually configuring static ARP associations is one way to prevent ARP spoofing.
  - Authorized MAC addresses can be configured on some network devices to restrict network access to only those devices listed.

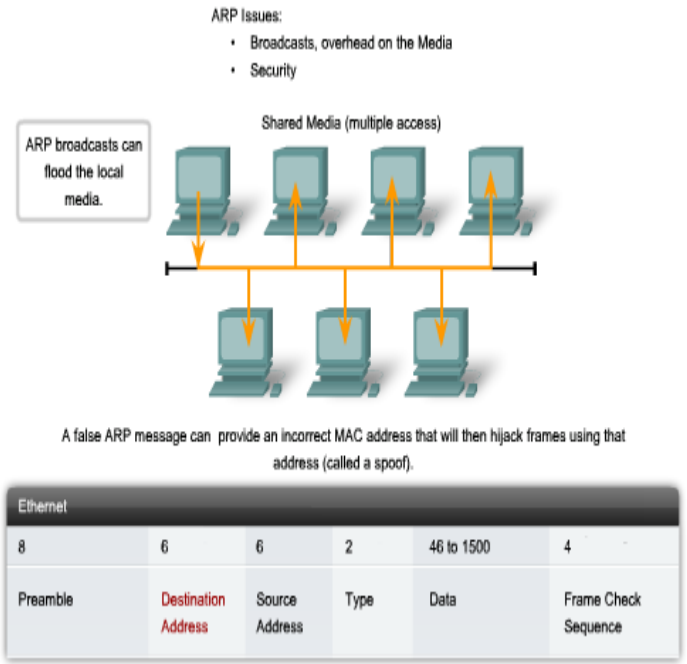

ARP Functions

ARP Table

- Used to find the MAC address that is mapped to the destination IPv4 address.
- If the destination IPv4 address is on the same network as the source IPv4, the device will search the ARP table for the destination IPv4 address.
- If the destination IPv4 address is on a different network, the device will search for the IPv4 address of the default gateway.
- If the device locates the IPv4 address, its corresponding MAC address is used as the destination MAC address in the frame.
- If no entry is found, then an ARP request is sent.

ARP Request

- Sent when a device needs a MAC address associated with an IPv4 address, and it does not have an entry in its ARP table.
- The ARP request message includes:
  - Target IPv4 address – This is the IPv4 address that requires a corresponding MAC address.
  - Target MAC address – This is the unknown MAC address and will be empty in the ARP request message.
- The ARP request is encapsulated in an Ethernet frame using the following header information:
  - Destination MAC address – This is a broadcast address requiring all Ethernet NICs on the LAN to accept and process the ARP request.
  - Source MAC address – This is the sender's MAC address.
  - Type – ARP messages have a type field of 0x806.
- See VIDEO DEMONSTRATION

ARP Reply

- The device with the target IPv4 address in the ARP request will respond with an ARP reply. The ARP reply message includes:
  - Sender's IPv4 address – This is the IPv4 address of the sender, the device whose MAC address was requested.
  - Sender's MAC address – This is the MAC address of the sender, the MAC address needed by the sender of the ARP request.
- The ARP reply is encapsulated in an Ethernet frame using the following header information:
  - Destination MAC address – This is the MAC address of the sender.
  - Source MAC address – This is the sender of the ARP reply's MAC address.
  - Type – ARP messages have a type field of 0x806.
- See VIDEO DEMONSTRATION

Video Demonstration – ARP Role in Remote Communication

- When the destination IPv4 address is not on the same network as the source IPv4 address, the source device needs to send the frame to its default gateway.
- The source checks its ARP table for an entry with the IPv4 address of the default gateway.
- If there is not an entry, it uses the ARP process to determine the MAC address of the default gateway.
- See VIDEO DEMONSTRATION

Removing Entries from an ARP Table

- ARP cache timer removes ARP entries that have not been used for a specified period of time.
- Commands may also be used to manually remove all or some of the entries in the ARP table.

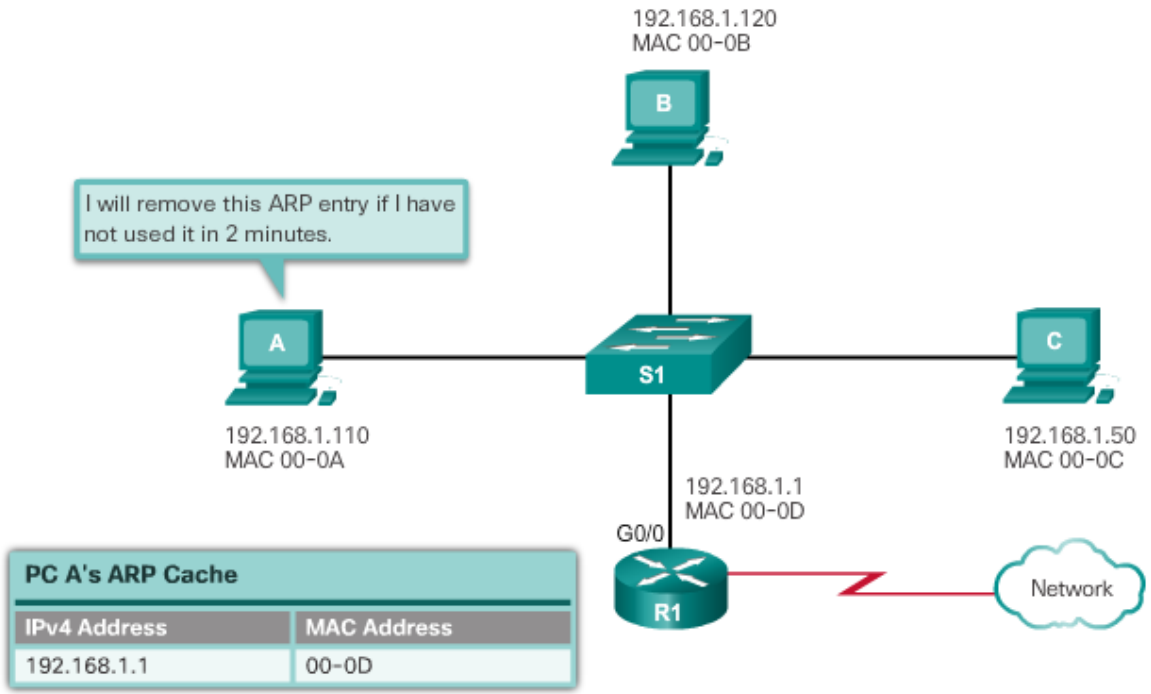

ARP Tables on Networking Devices

Router ARP Table

```
Router# show ip arp

                     Age
Protocol  Address          (min)  Hardware Addr    Type   Interface
Internet  172.16.233.229    -    0000.0c59.f892   ARPA   Ethernet0/0
Internet  172.16.233.218    -    0000.0c07.ac00   ARPA   Ethernet0/0
Internet  172.16.168.11     -    0000.0c63.1300   ARPA   Ethernet0/0
Internet  172.16.168.254    9    0000.0c36.6965   ARPA   Ethernet0/0
```

ARP Tables on Networking Devices (cont.)

Host ARP Table

```
C:\> arp -a

Interface: 192.168.1.67 --- 0xa
  Internet Address      Physical Address      Type
  192.168.1.254         64-0f-29-0d-36-91     dynamic
  192.168.1.255         ff-ff-ff-ff-ff-ff     static
  224.0.0.22            01-00-5e-00-00-16     static
  224.0.0.251           01-00-5e-00-00-fb     static
  224.0.0.252           01-00-5e-00-00-fc     static
  255.255.255.255       ff-ff-ff-ff-ff-ff     static

Interface: 10.82.253.91 --- 0x10
  Internet Address      Physical Address      Type
  10.82.253.92          64-0f-29-0d-36-91     dynamic
  224.0.0.22            01-00-5e-00-00-16     static
  224.0.0.251           01-00-5e-00-00-fb     static
  224.0.0.252           01-00-5e-00-00-fc     static
  255.255.255.255       ff-ff-ff-ff-ff-ff     static
```

Summary

- Ethernet uses the Address Resolution Protocol to determine the MAC addresses of destinations and map them against known Network layer addresses.
- Each node on an IP network has both a MAC address and an IP address.
- The ARP protocol resolves IPv4 addresses to MAC addresses and maintains a table of mappings.
- A Layer 2 switch builds a MAC address table that it uses to make forwarding decisions.
- Layer 3 switches are also capable of performing Layer 3 routing functions, reducing the need for dedicated routers on a LAN.
- Layer 3 switches have specialized switching hardware so they can typically route data as quickly as they can switch.