Global Knowledge®

Expert Reference Series of White Papers

Switching Operations

Switching Operations

Alan Thomas, CCNA, CCSI, Global Knowledge Instructor

Introduction

Ethernet Local Area Network (LAN) switches play an integral role in moving data from one device on the network to another. They also represent a huge leap forward in terms of capabilities and efficiency over their predecessor . . . the Ethernet hub.

This white paper describes how a switch moves data through the network, and what makes a switch efficient. To start, there are a few definitions, then a look at the process of data forwarding with a switch. Commands and output of commands are included throughout.

Definitions

OSI Network Reference Model	The Open Systems Interconnection (OSI) model is a set of guidelines that define the communications process. It is comprised of seven layers and 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, and so on.
Broadcast Domain	A broadcast domain is a collection of devices that receive broadcast frames from each other. A broadcast frame is destined (addressed) to every device within a broadcast domain. The destination Ethernet address is all ones. 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 locally connected and can exchange data with each other through a switch using a Layer 2 Media Access Control (MAC) address. Devices located in different broadcast domains are remotely connected and require a router to exchange data using Layer 3 (IP) addresses. The term *Virtual Local Area Network* (VLAN) is also used to describe a broadcast domain.

<table>
<tr><td>Ethernet</td><td>Ethernet is a Layer 2 protocol. It defines how devices access the physical segment to which they are connected. Ethernet uses 48-bit addresses, or MAC addresses, to uniquely identify devices on the network.</td></tr>
<tr><td>MAC Address</td><td>The MAC address is the address used for data exchanged between devices that are 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 usually 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 address.</td></tr>
<tr><td>Ethernet Frame</td><td>A frame is a Layer 2 Protocol Data Unit (PDU). The chart below shows the fields of an Ethernet frame.
 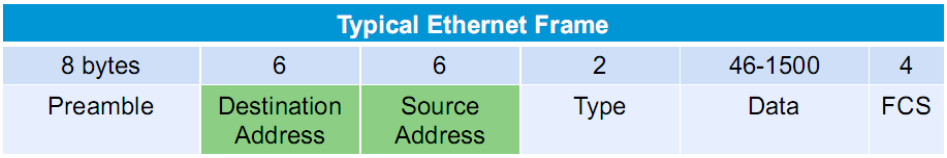

For this paper, the key fields are the destination and source address fields, as they are used in the frame forwarding process.</td></tr>
<tr><td>Switch</td><td>A switch is a Layer 2 device that forwards data based on the destination MAC address. A switch contains a MAC address table, which is an association 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.</td></tr>
<tr><td>Hub</td><td>A hub is a Layer 1 device that provides physical connectivity to multiple endpoint devices. It has no "intelligence" in the way a switch does, so it has no mechanism for learning which devices are connected to which interface. This means a hub always floods received data out of every interface.

Hubs are also considered shared bandwidth. This requires the use of the slower half-duplex connectivity method, as only one device can access the shared bandwidth at a time.</td></tr>
</table>

Half-Duplex	Half-duplex is a connectivity method in which a device can send or receive data, but cannot do both simultaneously. It is similar in concept to using a walkie-talkie, where one person can talk while the other person listens, but both people cannot talk at the same time. Half-duplex is an inefficient connectivity method because devices spend a lot of time waiting to send data since data can only be sent one direction at a time. Half-duplex was required when hubs were used extensively in networks, but today, it is considered to be a legacy connectivity method. However, if necessary, switches can support half-duplex connectivity.
Full-Duplex	Full-duplex is a connectivity method in which a device can send and receive data simultaneously. It is similar to an in-person conversation where both people can talk and be heard at the same time. Full-duplex is much more efficient than half-duplex because devices do not have as much wait time. Full-duplex requires that both devices support full-duplex, and that the connectivity be point-to-point. Point-to-point connectivity requires the device be connected to a dedicated switch port, or the device must be physically connected to the other device with a cable.
Collision	On a shared bandwidth segment, such as hub, a collision occurs when two devices simultaneously attempt to send data. When a collision occurs, both devices wait a random amount of time before attempting to send data again. The two devices should wait different amounts of time, thereby avoiding another collision between the same two devices. Devices that are connected to the same bandwidth segment are said to be in the same collision domain.
Bandwidth	Bandwidth, also referred to as speed, is the rate at which the network can move data between two endpoints. It is represented in bits per second.

A Brief History

A broadcast domain is a collection of devices that can be reached via a broadcast frame. A broadcast frame is delivered to each device. Address Resolution Protocol (ARP) and Dynamic Host Configuration Protocol (DHCP) are two examples of communications protocols that utilize broadcasts.

A broadcast domain is the most basic of computer networks. Terms such as *LAN*, *VLAN*, and *subnet* are all ways of describing a broadcast domain. A LAN is a network, or broadcast domain, that covers a relatively small geographic area. Typically, a LAN covers a room or group rooms, possibly an entire floor, or even an entire building, if the building is small enough.

When LANs were first introduced, they were made of up computers connected to a single coaxial cable that was usually placed around the outer edge of a room, and the computers all tapped into the coaxial cable. The coaxial cable was the physical media or segment (Layer 1) that connected the devices, and the bandwidth it provided was shared, meaning all devices used the same bandwidth.

Ethernet was the Layer 2 protocol used to format data for transmission on the coaxial cable. Ethernet was also the protocol used to control access to the media. In a shared bandwidth environment, Ethernet limits use of the media to one device at a time. This required devices to connect at half-duplex.

If two devices attempted to access the media at the same time, meaning two devices attempted to send data at the same time, a collision occurred. When a collision occurred, both devices would wait a random amount time before trying again to access the media. They would wait different amounts of time in an attempt to prevent further collisions. However, while the two devices that collided both waited, a third device could see the segment as available and begin to send data. This further increased the wait time hosts experienced.

This type of LAN had many limitations. One of the most significant is that they were slow, as devices had to wait relatively long periods of time to transmit data. Also, a failure of one device usually meant that all devices lost their ability to send and receive data.

Eventually, Ethernet was updated to operate over twisted pair copper cabling. This led to the development of Ethernet hubs. Hubs provide a dedicated physical connection for every device, which helps reduce the possibility that a failure of one computer will cause all computers to lose connectivity. However, because a hub is still a shared bandwidth device, connectivity is limited to half-duplex. Additionally, collisions remain an issue as well, so hubs do not help improve the performance of the network.

Today, switches are used extensively—almost exclusively—throughout enterprise networks. Switches are a huge leap forward in efficiency and performance. One reason is that each physical interface on a switch is an independent Ethernet segment. As a result, an end device no longer has to compete to access a segment, which allows devices to use full-duplex connectivity. Additionally, the collision domain is reduced to just two devices, so collisions are virtually eliminated. These advantages result in greatly increased throughput, as the amount of waiting a device must do to send data is significantly reduced.

A second reason switches are efficient is they keep track of the interface on which an end device is connected. With this knowledge, a switch can be much more efficient in its forwarding of data, because data can be forwarded out of a single interface instead of flooded out of all interfaces, which is how a hub always forwards data.

Today's enterprise networks typically consist of multiple LANs (or VLANs/subnets). LANs might group devices together based on physical location or functions within the organization. In either case, switches are the key device for providing access to a LAN, and so a good understanding of how switches operate is crucial for the networking professional.

Duplex

A host that is connected at full-duplex has the ability to send and receive data simultaneously. A host that is connected at half-duplex can send and receive data, but it cannot do both simultaneously. Figure 1 below shows the concept. Red indicates connectivity does not work and green means it does.

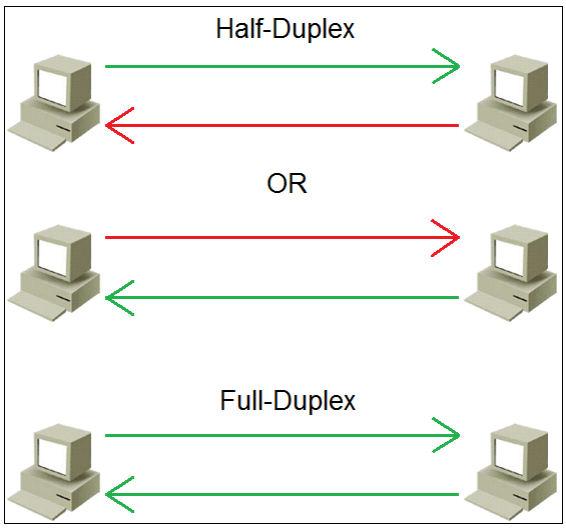

Figure 1: Duplex

Half-duplex connectivity is considered to be legacy, and should be avoided if possible. However, it is even more important to ensure there is not a duplex mismatch, which will cause a loss of connectivity.

Interfaces on Cisco switches are configured to auto-negotiate duplex connectivity by default, but auto-negotiation can create unpredictable results. It is Cisco best practice for an interface that connects to another switch, a router, or a critical endpoint to be manually configured for the desired duplex. Other interfaces can be left at the default of auto-negotiation, although this practice is not accepted by all network administrators.

To verify duplex settings, use the **show interface <*type*> <*number*>** command. Type can be Ethernet, fastethernet, gigabitethernet, or tengigabitethernet. Number is the number of the interface. Figure 2 shows the output for interface Ethernet 0/1. The interface is configured for the default of auto-negotiation (highlighted). The figure also shows the command fully typed out.

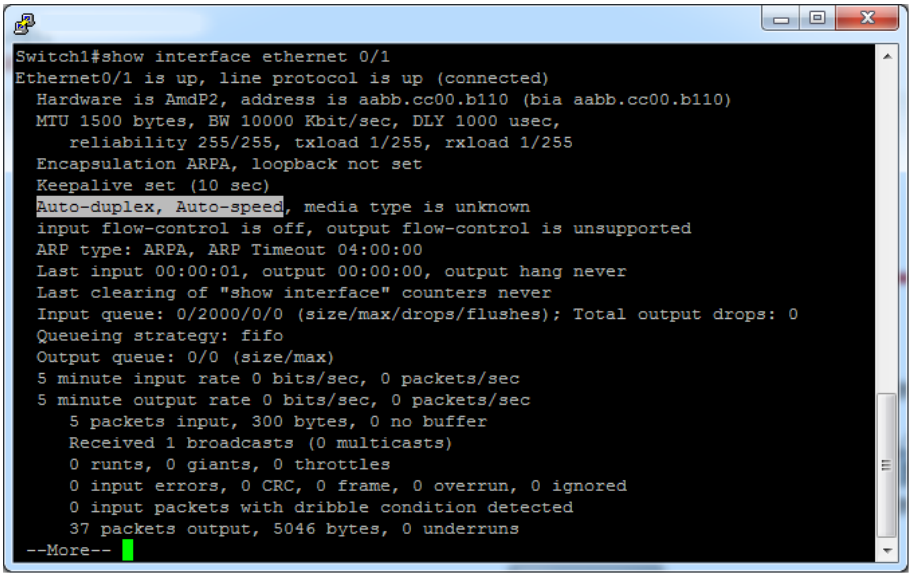

Figure 2: Verify Auto-Duplex

Figure 3 shows the output when the interface is manually configured for full-duplex (highlighted) and the command in abbreviated form.

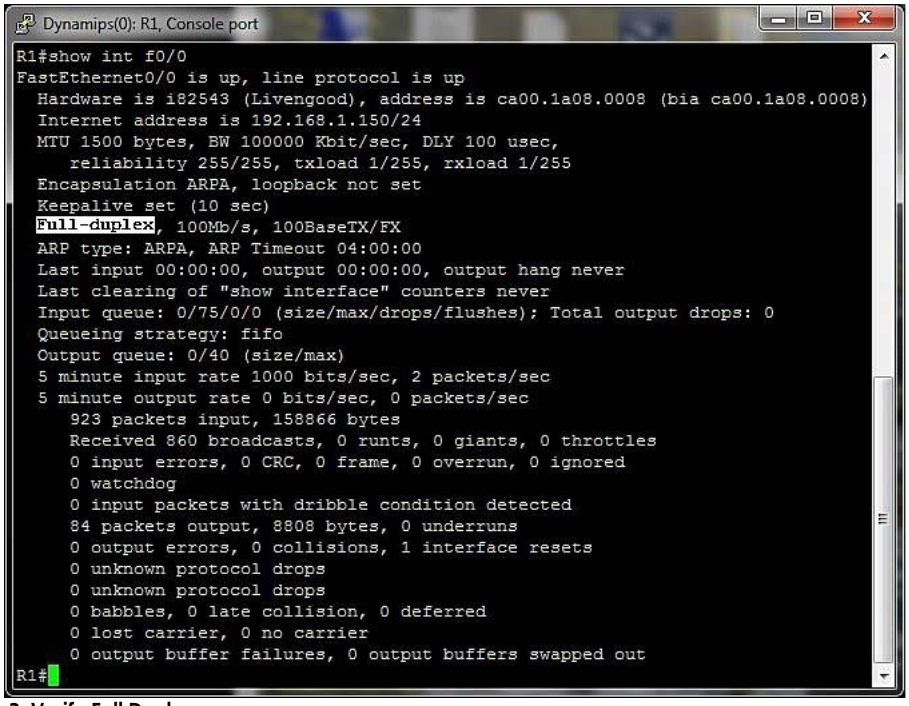

Figure 3: Verify Full-Duplex

MAC Address

Every device on a LAN must have a unique MAC address, which is used to identify the location of the device on the LAN. The manufacturer of a device assigns the MAC address, which is a 48-bit number represented using hexadecimal numbering.

MAC addresses are comprised of two parts. The first 24 bits are called the Organizationally Unique Identifier (OUI), and are assigned by the Institute of Electrical and Electronics Engineers (IEEE) to a manufacturer of network equipment. The last 24 bits are vendor assigned and create the unique MAC address for the piece of equipment.

Different manufacturers of network devices display MAC addresses in different ways. For example, Cisco displays MAC addresses as three sets of four hexadecimal digits, with each set separated by a dot, while Microsoft displays MAC addresses as six sets of two hexadecimal digits, with each set separated by a hyphen. So, the MAC address **0000.0c12.acdc** is in Cisco display format, while **00-00-0C-12-AC-DC** is the same MAC address in Microsoft display format.

MAC Address Table

A MAC address table associates the MAC address of a device with an interface on the switch. The switch uses the information in the MAC address table to determine how the frame should be forwarded. The options a switch has are to forward, flood, or filter the frame. These options are discussed later in the paper.

Figure 4 shows a simple LAN consisting of one switch and three PCs. By default, the MAC address table of the switch is empty.

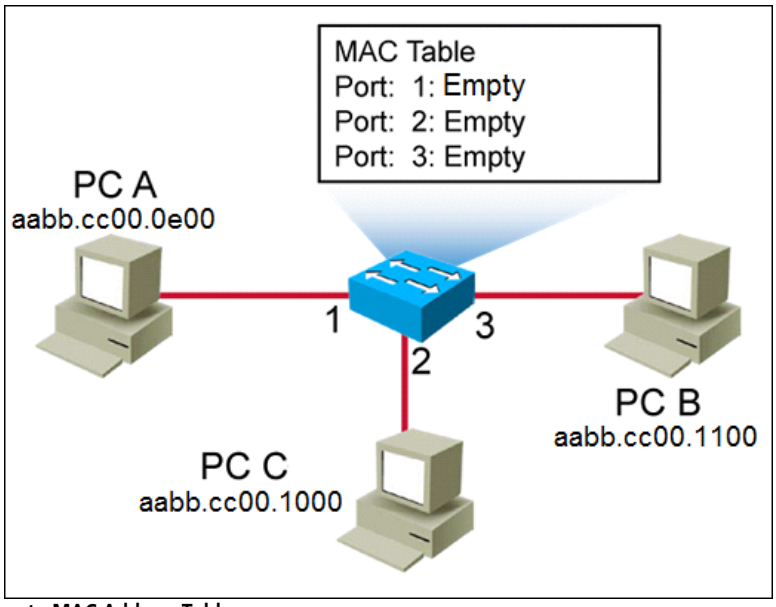

Figure 4: Empty MAC Address Table

To view the MAC address table, use the **show mac address-table** command. Figure 5 shows the output of an empty MAC address table on a Cisco switch.

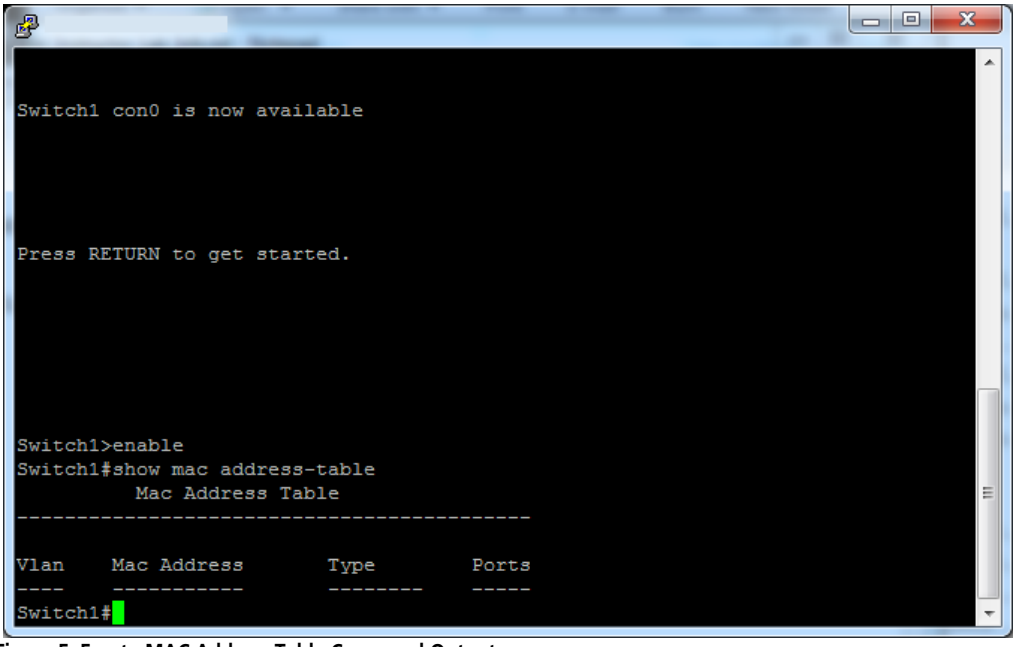

Figure 5: Empty MAC Address Table Command Output

The MAC address table is only populated when a device sends a frame. In figure 6, PC A sends a frame destined for PC C. When the switch receives the frame, the switch reads the source address field in the frame. The source MAC address, the inbound interface number, and the VLAN ID are placed into the MAC address table. When PC B and PC C send a frame, the MAC address table will be updated with their MAC address, inbound interface number, and VLAN ID.

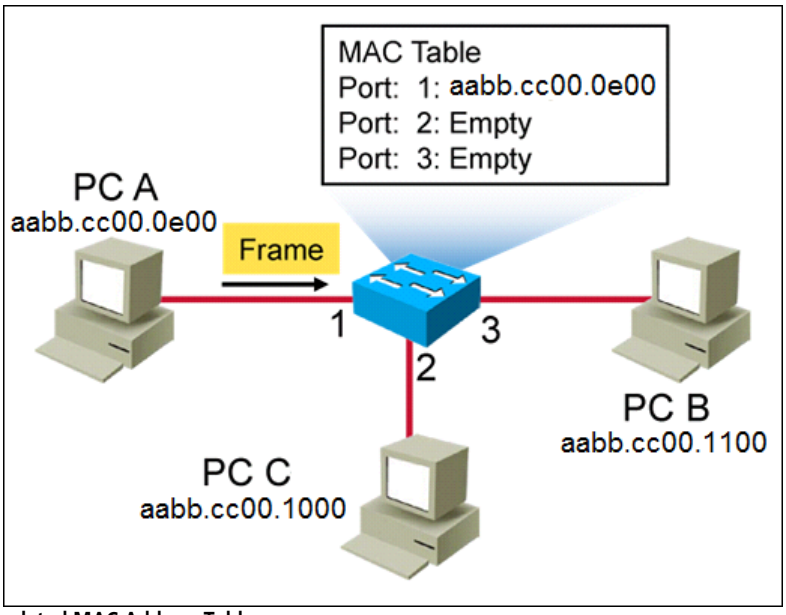

Figure 6: Populated MAC Address Table

Figure 7 shows the output of a populated MAC address table on a Cisco switch. The way a switch utilizes the MAC address table in the forwarding process is covered later in this paper.

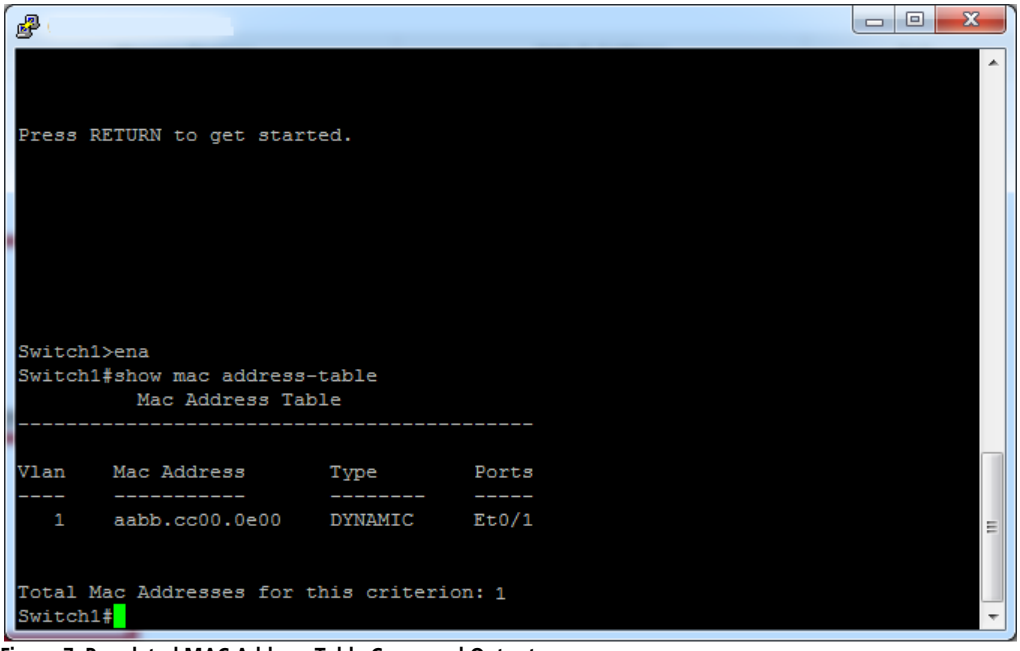

Figure 7: Populated MAC Address Table Command Output

Content-Addressable Memory (CAM)

Another feature that makes Cisco switches very efficient at forwarding frames is the ability to make forwarding decisions at the inbound interface instead of using the CPU. This process is often called hardware switching. Cisco switch interfaces use an Application-Specific Integrated Circuit (ASIC), which is an integrated circuit customized for a specific purpose. In this case, the special purpose is the forwarding of frames.

The ASIC utilizes a special memory table called CAM. The CAM table—also called the MAC address table—contains a listing of all known MAC addresses that are connected to the interface and VLAN. A CAM table is ideal for frame forwarding, because CAM table searches yield only a true or false result. In order for the result to be true, there must be an *exact* match of the search criteria. If a frame is destined to a specific device, an exact MAC address match in the CAM table is required to perform the most efficient type of frame forwarding. (Frame forwarding methods are discussed later in this paper.)

Also, utilizing an ASIC with a CAM table is a very fast process for forwarding frames. One reason is that the entire contents of a CAM table can be searched in a single CPU cycle, which is much faster than searching RAM. Another reason is that when the inbound interface determines the proper outbound interface, which is based on the CAM table lookup, the frame can be sent across the backplane of the switch directly from one interface to another. In other words, the CPU of the switch is not needed to actually move the frame from one interface to another, which greatly reduces the amount of time needed to move the frame.

Forwarding Actions

When a frame arrives on an interface, a Cisco switch has to decide what forwarding action to take on the frame. Cisco switches have three actions they can take with the frame: forward, flood, or filter.

1. Forward: The switch sends the frame out of a single, specific interface. This is the most efficient action a switch can take. The forwarding action is taken when the destination MAC address field of the frame contains the MAC address of a specific host, *and* that MAC address has an entry in the MAC address table. In Figure 8, PC A sends a frame to PC B. Since there is an entry in the MAC address table for PC B, the frame is forwarded out of port 3 only.

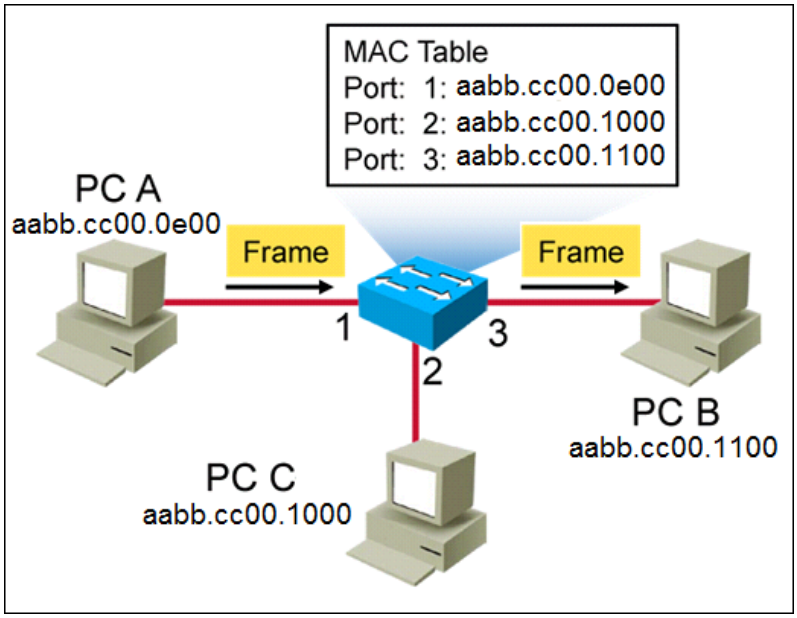

Figure 8: -Frame Forward

2. Flood: This means the switch sends the frame out of every interface except the one on which the frame was received. This is the least efficient action a switch can take because the frame must be re-generated for every active interface on the switch. Another inefficiency is every device receives a copy of the frame, even if that device is not the target device. There are two situations in which a switch will flood a frame:

    a. If the destination MAC address field of the frame contains the broadcast MAC address (ffff.ffff.ffff)

    b. If the destination MAC address field of the frame contains the MAC address of a specific host, but that MAC address does *not* have an entry in the MAC address table. This situation is called an "unknown unicast MAC address" because the switch has not yet learned the interface to which the destination host is connected.

In figure 9, PC A sends a frame to PC B. Since there is not an entry in the MAC address table for PC B, the frame is flooded out of every interface except the one on which the frame was received.

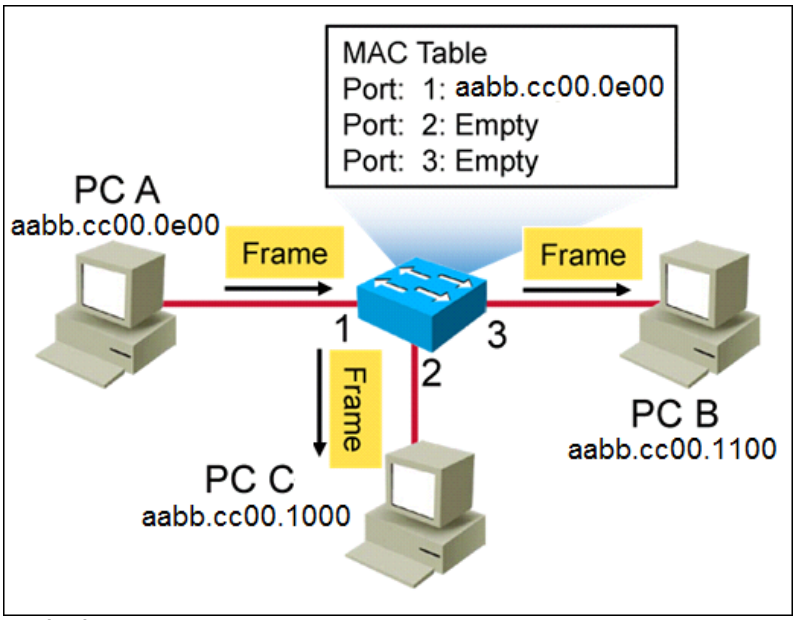

Figure 9: Frame Flood

3. Filter: This means the switch drops the frame and does not forward it at all. The filtering action is taken when the destination MAC address and the source MAC address are connected to the same interface on the switch. This condition occurs when a hub is connected to the switch and in wireless implementations where a wireless controller is not used.

Connecting a hub to a switch interface is not recommended. A hub shares the bandwidth of a segment among different devices. This results in sub-optimal performance, as half-duplex connectivity must be used and collisions will be an issue.

In figure 10, PC D and PC E are connected to a hub that is connected to interface 4 of the switch. The MAC address table shows that MAC addresses for both PCs are associated with port 4. When PC D sends a frame to PC E, the hub sends the frame out all of its interfaces. When the switch receives the frame and determines that the source and destination MAC addresses are connected to the same port, the switch filters the frame; that is, the switch discards the frame.

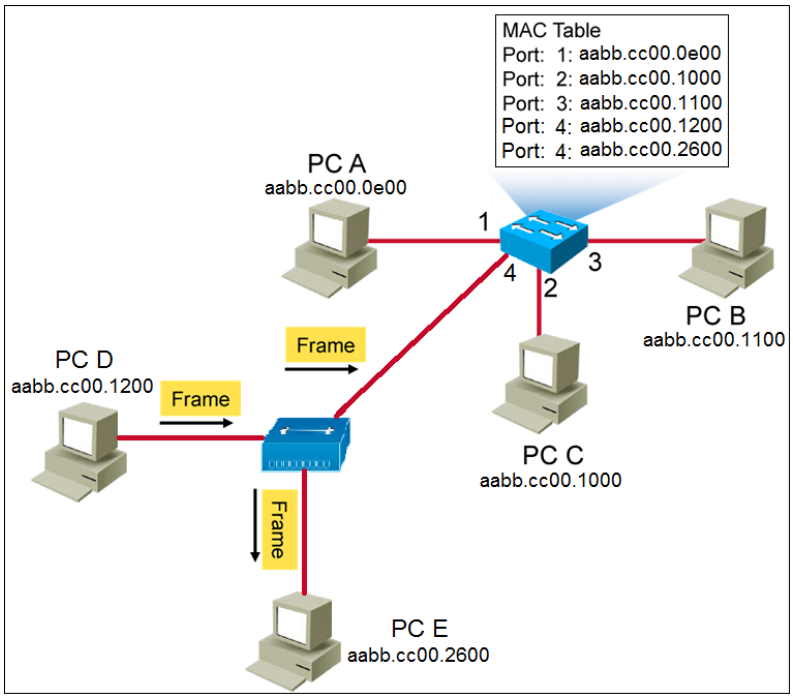

Figure 10: Frame Filter

The Complete Process

The following steps describe what a switch does when a frame arrives on an interface.

1. The switch calculates a Frame Check Sequence (FCS) and compares its own calculation with the value found in the FCS field of the frame. If the values match, the switch knows the frame has arrived error free. If the values do not match, the switch knows an error has occurred during transmission, and the frame is discarded.
2. The switch enters the source MAC address, inbound interface number, and VLAN number into the MAC address table. Additionally, a time stamp is added. If the source MAC address has an existing entry, then the time stamp is refreshed.
3. The switch determines the destination MAC address.
    a. If it is the broadcast address, the frame is flooded out of every interface except the interface on which the frame was received.
    b. If the destination MAC address is not the broadcast address, the switch looks for a match in the MAC address table.
        i. If a match is not found, the frame is flooded out of every interface except the interface on which the frame was received.
        ii. If a match is found, *and* the source and destination MAC addresses are connected to the *same* interface, the frame is filtered.
        iii. If a match is found, *and* the source and destination MAC addresses are connected to *different* interfaces, the frame is forwarded out of the proper interface.

Conclusion

Switches play a vital role in moving data from one device to another. Specifically, switches greatly improve network performance, compared to hubs, by providing dedicated bandwidth to each end device, supporting full-duplex connectivity, utilizing the MAC address table to make forwarding decisions, and utilizing ASICs and CAM tables to increase the rate at which frames can be processed.

Learn More

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

- CCNAX v2.0 - CCNA Routing and Switching Boot Camp
- SWITCH - Implementing Cisco IP Switched Networks v2.0
- CCNP Routing and Switching e-Camp

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's 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.