Chapter 1:
Fundamentals Review

CCNP SWITCH: Implementing Cisco IP Switched Networks

Cisco | Networking Academy®
Mind Wide Open™

Chapter 1 Objectives

This chapter covers the following basic switching topics as a review to CCNA and serves as a teaser for topics covered later in chapter:

- Hubs and switches
- Bridges and switches
- Switches of today
- Broadcast domains
- MAC addresses
- The basic Ethernet frame format
- Basic switching function
- VLANs
- The Spanning Tree Protocol
- Trunking
- Port channels
- Multilayer switching (MLS)

Fundamentals Review

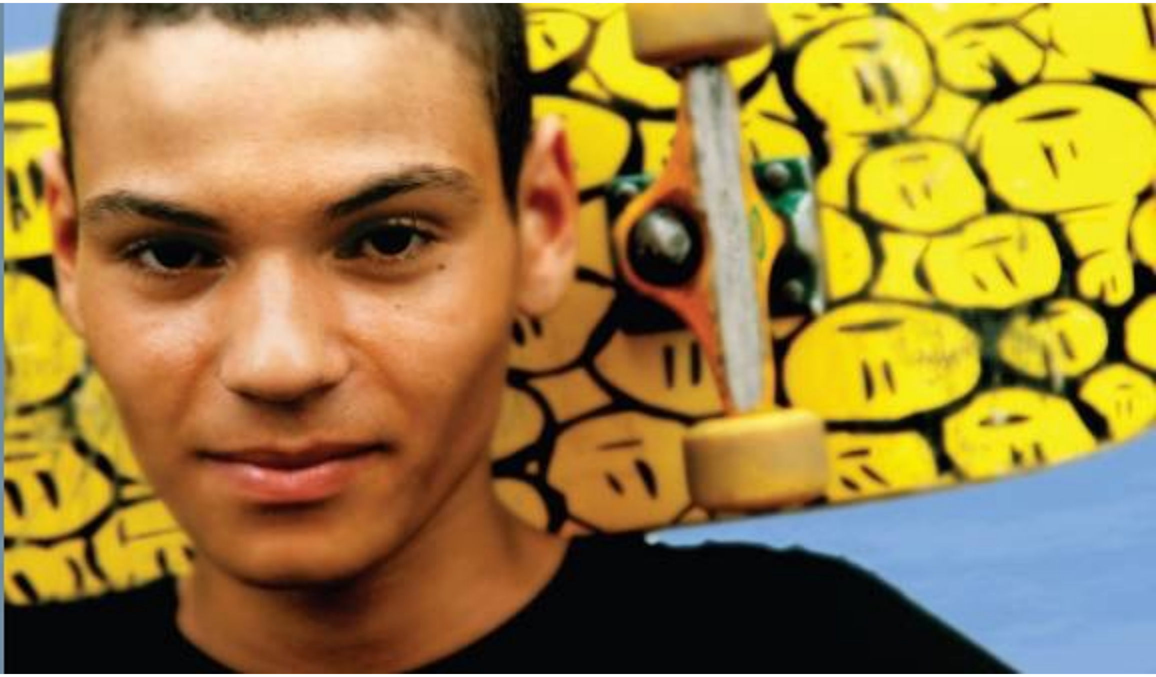

Hubs and Switches

- Hubs are archaic, and the terminology should be avoided. Even the simplest multiport Ethernet devices for the home are switches.
- In review, hubs died off as a product because they are shared-bandwidth devices.
- Switches introduced dedicated bandwidth. A hub allows multiple devices to be connected to the same network segment. The devices on that segment share the bandwidth with each other.
- A switch allows multiple devices to be connected to the same network, just like a hub does, but this is where the similarity ends.
- A switch allows each connected device to have dedicated bandwidth instead of shared bandwidth.
- Switches also support additional capabilities beyond what hubs support.

Bridges and Switches

- A basic switch is considered a Layer 2 device. When we use the word *layer*, we are referring to the seven-layer OSI reference model.
- A switch does not just pass electrical signals along, like a hub does; instead, it assembles the signals into a frame (Layer 2), and then decides what to do with the frame.
- A switch determines what to do with a frame by borrowing an algorithm from a previously common networking device: a transparent bridge.
- Logically, a switch acts just like a transparent bridge would, but it can handle frames much faster than a transparent bridge could
- Once a switch decides where the frame should be sent, it passes the frame out the appropriate port (or ports). You can think of a switch as a device creating instantaneous connections between various ports, on a frame-by-frame basis.

Switches of Today

Today’s switches have evolved beyond just switching frames. Most modern switches can actually route traffic. In addition, switches can prioritize traffic, support no downtime through redundancy, and provide convergence services around IP telephony and wireless networks.

- **Application intelligence**
  - This helps networks recognize many types of applications and secure and prioritize those applications to provide the best user experience.
- **Unified network services**
  - Combining the best elements of wireless and wired networking allows you to consistently connect to any resource or person with any device. 10 Gigabit Ethernet technology and Power over Ethernet (PoE) technology support new applications and devices.
- **Nonstop communications**
  - Features such as redundant hardware, and nonstop forwarding and stateful switchover (NSF/SSO) technology support more-reliable connections.
- **Integrated security**
  - LAN switches provide the first line of defense against internal network attacks and prevent unauthorized intrusion.
- **Operational manageability**
  - To more easily manage the network, IT staff must be able to remotely configure and monitor network devices from a central location.

Broadcast Domains

- A broadcast domain is a set of network devices that receive broadcast frames originating from any device within the group.
- Routers typically bound broadcast domains because routers do not forward broadcast frames.
- VLANs are an example of broadcast domain.
- Broadcast domains are generally limited to a specific Layer 2 segment that contains a single IP subnet.

MAC Addresses

- MAC addresses are standardized data link layer addresses that are required for every port or device that connects to a LAN.
- Other devices in the network use these addresses to locate specific ports in the network and to create and update routing tables and data structures.
- MAC addresses are 6 bytes long and are controlled by the IEEE.

The Basic Ethernet Frame Format

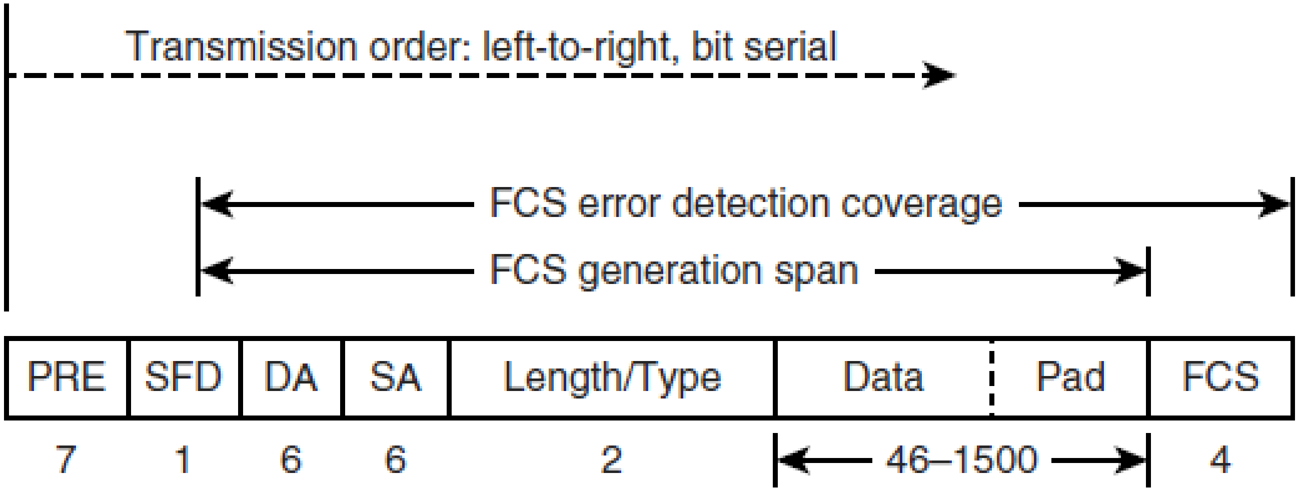

Transmission order: left-to-right, bit serial

FCS error detection coverage

FCS generation span

PRE	SFD	DA	SA	Length/Type	Data	Pad	FCS
7	1	6	6	2	46–1500		4

Field length in bytes

PRE = Preamble
SFD = Start-of-frame delimiter
DA = Destination address
SA = Source address
FCS = Frame check sequence

The Basic Ethernet Frame Format

- **Preamble (PRE)**
  - Consists of 7 bytes. The PRE is an alternating pattern of 1s and 0s that tells receiving stations that a frame is coming, and that provides a means of synchronization
- **Start-of-frame delimiter (SOF)**
  - Consists of 1 byte. The SOF is an alternating pattern of 1s and 0s, ending with two consecutive 1 bits, indicating that the next bit is the leftmost bit in the leftmost byte of the destination address.
- **Destination address (DA)**
  - Consists of 6 bytes. The DA field identifies which station(s) should receive the frame.
- **Source addresses (SA)**
  - Consists of 6 bytes. The SA field identifies the sending station.

The Basic Ethernet Frame Format

- **Length/Type**
  - Consists of 2 bytes. This field indicates either the number of MAC-client data bytes that are contained in the data field of the frame, or the frame type ID if the frame is assembled using an optional format.
- **Data**
  - Is a sequence of *n* bytes of any value, where *n* is less than or equal to 1500.
  - Note that jumbo frames up to 9000 bytes are supported on the current-generation Cisco Catalyst switches.
- **Frame check sequence (FCS)**
  - Consists of 4 bytes. This sequence contains a 32-bit cyclic redundancy check (CRC) value, which is created by the sending MAC and is recalculated by the receiving MAC to check for damaged frames. The FCS is generated over the DA, SA, Length/Type, and Data fields.

Basic Switching Function

- In brief, the basic switching function at Layer 2 adheres to these rules for determining forwarding responsibility:
  - If the destination MAC address is found in the CAM table, the switch sends the frame out the port that is associated with that destination MAC address in the CAM table. This process is called ***forwarding*** .
  - If the associated port to send the frame out is the same port that the frame originally came in on, there is no need to send the frame back out that same port, and the frame is ignored. This process is called ***filtering*** .
  - If the destination MAC address is not in the CAM table (that is, unknown unicast), the switch sends the frame out all other ports that are in the same VLAN as the received frame. This is called ***flooding*** . It does not flood the frame out the same port on which the frame was received.
  - If the destination MAC address of the received frame is the broadcast address (FFFF.FFFF.FFFF), the frame is sent out all ports that are in the same VLAN as the received frame. This is also called ***flooding*** .

VLANs

- Because the switch decides on a frame-by-frame basis which ports exchange data, it is a natural extension to put logic inside the switch to allow it to choose ports for special groupings. This grouping of ports is called a *virtual local-area network* (VLAN).
- The switch makes sure that traffic from one group of ports never gets sent to other groups of ports (which would be routing).
- These port groups (VLANs) can each be considered an individual LAN segment.
- VLANs are also described as broadcast domains. This is because of the transparent bridging algorithm, which says that broadcast packets (packets destined for the *all devices* address) be sent out all ports that are in the same group (that is, in the same VLAN).
- All ports that are in the same VLAN are also in the same broadcast domain.

The Spanning Tree Protocol

- As discussed previously, the switch forwarding algorithm floods unknown and broadcast frames out of all the ports that are in the same VLAN as the received frame. This causes a potential problem. If the network devices that run this algorithm are connected together in a physical loop, flooded frames (like broadcasts) are passed from switch to switch, around and around the loop, forever.
- There is a benefit to a physical loop in your network: It can provide redundancy. If one link fails, there is still another way for the traffic to reach its destination. To allow the benefits derived from redundancy, without breaking the network because of flooding, a protocol called the *Spanning Tree Protocol* (STP) was created.
- Spanning tree was standardized in the IEEE 802.1D specification.
- The purpose of STP is to identify and temporarily block the loops in a network segment or VLAN. The switches run STP, which involves electing a root bridge or switch.
- The other switches measure their distance from the root switch. If there is more than one way to get to the root switch, there is a loop. The switches follow the algorithm to determine which ports must be blocked to break the loop.
- STP is dynamic; if a link in the segment fails, ports that were originally blocking can possibly be changed to forwarding mode.

Trunking

- Trunking is a mechanism that is most often used to allow multiple VLANs to function independently across multiple switches.
- Routers and servers can use trunking, as well, which allows them to live simultaneously on multiple VLANs.
- If your network only has one VLAN in it, you might never need trunking; but if your network has more than one VLAN, you probably want to take advantage of the benefits of trunking.
- A port on a switch normally belongs to only one VLAN; any traffic received or sent on this port is assumed to belong to the configured VLAN.
- A trunk port, however, is a port that can be configured to send and receive traffic for many VLANs.
- It accomplishes this when it attaches VLAN information to each frame, a process called *tagging* the frame.
- Also, trunking must be active on both sides of the link; the other side must expect frames that include VLAN information for proper communication to occur.

Port Channels

- Utilizing port channels (EtherChannels) is a technique that is used when you have multiple connections to the “same device”.
- Rather than each link functioning independently, port channels group the ports together to work as one unit. Port channels distribute traffic across all the links and provide redundancy if one or more links fail.
- Port channel settings must be the same on both sides of the links involved in the channel.
- Normally, spanning tree would block all of these parallel connections between devices because they are loops, but port channels run *underneath* spanning tree, so that spanning tree thinks all the ports within a given port channel are only a single port.

Multilayer Switching

- Multilayer switching (MLS) is the ability of a switch to forward frames based on information in the Layer 3 and sometimes Layer 4 header. Almost all Cisco Catalyst switches model 3500 or later support MLS. MLS is becoming a legacy term due to the wide support.
- The most important aspect to MLS is recognizing that switches can route or switch frames at wire-rate speeds using specialized hardware. This effectively bundles the routing function into the switch and is specifically useful for routing between VLANs in the core of the network.

Chapter 1 Summary

- Hubs and switches
- Bridges and switches
- Switches of today
- Broadcast domains
- MAC addresses
- The basic Ethernet frame format
- Basic switching function
- VLANs
- The Spanning Tree Protocol
- Trunking
- Port channels
- Multilayer switching (MLS)

Chapter 1 Labs

- **CCNPv7.1 SWITCH Lab1 BASELINE STUDENT**

Cisco | Networking Academy®

Mind Wide Open™

Acknowledgment

- *Some of the images and texts are from Implementing Cisco IP Switched Networks (SWITCH) Foundation Learning Guide: (CCNP SWITCH 300-115)* by Richard Froom and Erum Frahim (1587206641)
- Copyright © 2015 – 2016 Cisco Systems, Inc.
- Special Thanks to *Bruno Silva*