

Basic Switching Concepts and Configuration

2.0 Basic Switching Concepts and Configuration

2.0.1.1 Introduction

Switches are used to connect multiple devices together on the same network. In a properly designed network, LAN switches are responsible for directing and controlling the data the flow at the access layer to networked resources.

Cisco switches are self-configuring and no additional configurations are necessary for them to function out of the box. However, Cisco switches run Cisco IOS, and can be manually configured to better meet the needs of the network. This includes adjusting port speed, bandwidth and security requirements.

Additionally, Cisco switches can be managed both locally and remotely. To remotely manage a switch it needs to have an IP address and default gateway configured. These are just two of the configurations discussed in this chapter.

Switches operate at the access layer where client network devices connect directly to the network and IT departments want uncomplicated network access for the users. It is one of the most vulnerable areas of the network because it is so exposed to the user. Switches need to be configured to be resilient to attacks of all types while they are protecting user data and allowing for high speed connections. Port security is one of the security features Cisco managed switches provide.

This chapter examines some of the basic switch configuration settings required to maintain a secure, available, switched LAN environment.

Refer to
Lab Activity
for this chapter

2.0.1.2 Class Activity – Stand By Me

Stand By Me

Scenario

When you arrived to class today, you were given a number by your instructor to use for this introductory class activity.

When class begins, your instructor will ask certain students with specific numbers to stand. Your job is to record the standing students' numbers for each scenario.

Scenario 1

Students with numbers starting with the number 5 should stand. Record the numbers of the standing students.

Scenario 2

Students with numbers ending in B should stand. Record the numbers of the standing students.

Scenario 3

Students with the number 504C should stand. Record the number of the standing student.

At the end of this activity, divide into small groups and record answers to the Reflection questions on the PDF for this activity.

Save your work and be prepared to share it with another student or the entire class.

2.1 Basic Switch Configuration

2.1.1 Configure a Switch with Initial Settings

Refer to
Interactive Graphic
in online course.

2.1.1.1 Switch Boot Sequence

After a Cisco switch is powered on, it goes through the following boot sequence:

1. First, the switch loads a power-on self-test (POST) program stored in ROM. POST checks the CPU subsystem. It tests the CPU, DRAM, and the portion of the flash device that makes up the flash file system.
2. Next, the switch loads the boot loader software. The boot loader is a small program stored in ROM and is run immediately after POST successfully completes.
3. The boot loader performs low-level CPU initialization. It initializes the CPU registers, which control where physical memory is mapped, the quantity of memory, and its speed.
4. The boot loader initializes the flash file system on the system board.
5. Finally, the boot loader locates and loads a default IOS operating system software image into memory and hands control of the switch over to the IOS.

The boot loader finds the Cisco IOS image on the switch as follows: the switch attempts to automatically boot by using information in the BOOT environment variable. If this variable is not set, the switch attempts to load and execute the first executable file it can by performing a recursive, depth-first search throughout the flash file system. In a depth-first search of a directory, each encountered subdirectory is completely searched before continuing the search in the original directory. On Catalyst 2960 Series switches, the image file is normally contained in a directory that has the same name as the image file (excluding the .bin file extension).

The IOS operating system then initializes the interfaces using the Cisco IOS commands found in the configuration file, startup configuration, which is stored in NVRAM.

In the figure, the BOOT environment variable is set using the `boot system` global configuration mode command. Use the `show bootvar` command (`show boot` in older IOS versions) to see what the current IOS boot file is set to.

Refer to
Online Course
for Illustration

2.1.1.2 Recovering From a System Crash

The boot loader provides access into the switch if the operating system cannot be used because of missing or damaged system files. The boot loader has a command-line that provides access to the files stored in flash memory.

The boot loader can be accessed through a console connection following these steps:

- Connect a PC by console cable to the switch console port. Configure terminal emulation software to connect to the switch.
- Unplug the switch power cord.
- Reconnect the power cord to the switch and, within 15 seconds, press and hold down the **Mode** button while the System LED is still flashing green.
- Continue pressing the **Mode** button until the System LED turns briefly amber and then solid green; then release the **Mode** button.
- The boot loader **switch:** prompt appears in the terminal emulation software on the PC.

The **boot loader** command line supports commands to format the flash file system, reinstall the operating system software, and recover from a lost or forgotten password. For example, the **dir** command can be used to view a list of files within a specified directory as shown in the figure.

Refer to
Online Course
for Illustration

2.1.1.3 Switch LED Indicators

Cisco Catalyst switches have several status LED indicator lights. You can use the switch LEDs to quickly monitor switch activity and its performance. Switches of different models and feature sets will have different LEDs and their placement on the front panel of the switch may also vary.

The figure shows the switch LEDs and the Mode button for a Cisco Catalyst 2960 switch. The Mode button is used to toggle through port status, port duplex, port speed, and PoE (if supported) status of the port LEDs. The following describes the purpose of the LED indicators, and the meaning of their colors:

- **System LED** - Shows whether the system is receiving power and is functioning properly. If the LED is off, it means the system is not powered on. If the LED is green, the system is operating normally. If the LED is amber, the system is receiving power but is not functioning properly.
- **Redundant Power System (RPS) LED** - Shows the RPS status. If the LED is off, the RPS is off or not properly connected. If the LED is green, the RPS is connected and ready to provide back-up power. If the LED is blinking green, the RPS is connected but is unavailable because it is providing power to another device. If the LED is amber, the RPS is in standby mode or in a fault condition. If the LED is blinking amber, the internal power supply in the switch has failed, and the RPS is providing power.
- **Port Status LED** - Indicates that the port status mode is selected when the LED is green. This is the default mode. When selected, the port LEDs will display colors with different meanings. If the LED is off, there is no link, or the port was administratively shut down. If the LED is green, a link is present. If the LED is blinking green, there

is activity and the port is sending or receiving data. If the LED is alternating green-amber, there is a link fault. If the LED is amber, the port is blocked to ensure a loop does not exist in the forwarding domain and is not forwarding data (typically, ports will remain in this state for the first 30 seconds after being activated). If the LED is blinking amber, the port is blocked to prevent a possible loop in the forwarding domain.

- **Port Duplex LED** - Indicates the port duplex mode is selected when the LED is green. When selected, port LEDs that are off are in half-duplex mode. If the port LED is green, the port is in full-duplex mode.
- **Port Speed LED** - Indicates the port speed mode is selected. When selected, the port LEDs will display colors with different meanings. If the LED is off, the port is operating at 10 Mb/s. If the LED is green, the port is operating at 100 Mb/s. If the LED is blinking green, the port is operating at 1000 Mb/s.
- **Power over Ethernet (PoE) Mode LED** - If PoE is supported; a PoE mode LED will be present. If the LED is off, it indicates the PoE mode is not selected and that none of the ports have been denied power or placed in a fault condition. If the LED is blinking amber, the PoE mode is not selected but at least one of the ports has been denied power, or has a PoE fault. If the LED is green, it indicates the PoE mode is selected and the port LEDs will display colors with different meanings. If the port LED is off, the PoE is off. If the port LED is green, the PoE is on. If the port LED is alternating green-amber, PoE is denied because providing power to the powered device will exceed the switch power capacity. If the LED is blinking amber, PoE is off due to a fault. If the LED is amber, PoE for the port has been disabled.

Refer to
Online Course
for Illustration

2.1.1.4 Preparing for Basic Switch Management

To prepare a switch for remote management access, the switch must be configured with an IP address and a subnet mask. Keep in mind, that to manage the switch from a remote network, the switch must be configured with a default gateway. This is very similar to configuring the IP address information on host devices. In the figure, the switch virtual interface (SVI) on S1 should be assigned an IP address. The SVI is a virtual interface, not a physical port on the switch.

SVI is a concept related to VLANs. VLANs are numbered logical groups to which physical ports can be assigned. Configurations and settings applied to a VLAN are also applied to all the ports assigned to that VLAN.

By default, the switch is configured to have the management of the switch controlled through VLAN 1. All ports are assigned to VLAN 1 by default. For security purposes, it is considered a best practice to use a VLAN other than VLAN 1 for the management VLAN.

Note that these IP settings are only for remote management access to the switch; the IP settings do not allow the switch to route Layer 3 packets.

Refer to
Online Course
for Illustration

2.1.1.5 Configuring Basic Switch Management Access with IPv4

- **Configure Management Interface** An IP address and subnet mask is configured on the management SVI of the switch from VLAN interface configuration mode. As shown in Figure 1, the `interface vlan 99` command is used to enter interface configuration mode. The `ip address` command is used to configure the IP address. The `no shutdown` command enables the interface. In this example, VLAN 99 is configured with IP

address 172.17.99.11. The SVI for VLAN 99 will not appear as “up/up” until VLAN 99 is created and there is a device connected to a switch port associated with VLAN 99. To create a VLAN with the `vlan_id` of 99, and associate it to an interface, use the following commands:

```
S1(config)# vlan vlan_id
S1(config-vlan)# name vlan_name
S1(config)# end
S1(config)# interface interface_id
S1(config-if)# switchport access vlan vlan_id
```

■ Configure Default Gateway

The switch should be configured with a default gateway if it will be managed remotely from networks not directly connected. The default gateway is the router the switch is connected to. The switch will forward IP packets with destination IP addresses outside the local network to the default gateway. As shown in Figure 2, R1 is the default gateway for S1. The interface on R1 connected to the switch has IP address 172.17.99.1. This address is the default gateway address for S1.

To configure the default gateway for the switch, use the `ip default-gateway` command. Enter the IP address of the default gateway. The default gateway is the IP address of the router interface to which the switch is connected. Use the `copy running-config startup-config` command to back up your configuration.

- **Verify Configuration** As shown in Figure 3, the `show ip interface brief` command is useful when determining the status of both physical and virtual interfaces. The output shown in the figure confirms that interface VLAN 99 has been configured with an IP address and subnet mask, and Fast Ethernet port F0/18 has been assigned to the VLAN 99 management interface. Both interfaces are now “up/up” and operational.

Refer to
Lab Activity
for this chapter

2.1.1.6 Lab - Basic Switch Configuration

In this lab, you will complete the following objectives:

- Part 1: Cable the Network and Verify the Default Switch Configuration
- Part 2: Configure Basic Network Device Settings
- Part 3: Verify and Test Network Connectivity
- Part 4: Manage the MAC Address Table

Refer to
Online Course
for Illustration

2.1.2 Configure Switch Ports

2.1.2.1 Duplex Communication

The figure illustrates full-duplex and half-duplex communication.

Full-duplex communication improves the performance of a switched LAN. Full-duplex communication increases effective bandwidth by allowing both ends of a connection to transmit and receive data simultaneously. This is also known as bidirectional. This method of optimizing network performance requires micro-segmentation. A micro-segmented

LAN is created when a switch port has only one device connected and is operating at full-duplex. This results in a micro size collision domain of a single device. Because there is only one device connected, a micro-segmented LAN is collision free.

Unlike full-duplex communication, half-duplex communication is unidirectional. Sending and receiving data does not occur at the same time. Half-duplex communication creates performance issues because data can flow in only one direction at a time, often resulting in collisions. Half-duplex connections are typically seen in older hardware, such as hubs. Full-duplex communication has replaced half-duplex in most hardware.

Most Ethernet and Fast Ethernet NICs sold today offer full-duplex capability. Gigabit Ethernet and 10Gb NICs require full-duplex connections to operate. In full-duplex mode, the collision detection circuit on the NIC is disabled. Frames that are sent by the two connected devices cannot collide because the devices use two separate circuits in the network cable. Full-duplex connections require a switch that supports full-duplex configuration, or a direct connection using an Ethernet cable between two devices.

Standard, shared hub-based Ethernet configuration efficiency is typically rated at 50 to 60 percent of the stated bandwidth. Full-duplex offers 100 percent efficiency in both directions (transmitting and receiving). This results in a 200 percent potential use of the stated bandwidth.

Refer to
Online Course
for Illustration

2.1.2.2 Configure Switch Ports at the Physical Layer Duplex and Speed

Switch ports can be manually configured with specific duplex and speed settings. Use the `duplex` interface configuration mode command to manually specify the duplex mode for a switch port. Use the `speed` interface configuration mode command to manually specify the speed for a switch port. In Figure 1, port F0/1 on switch S1 and S2 are manually configured with the `full` keyword for the `duplex` command, and the `100` keyword for the `speed` command.

The default setting for both duplex and speed for switch ports on Cisco Catalyst 2960 and 3560 switches is auto. The 10/100/1000 ports operate in either half- or full-duplex mode when they are set to 10 or 100 Mb/s, but when they are set to 1000 Mb/s (1 Gb/s), they operate only in full-duplex mode. Cisco recommends only using the `auto` command for duplex and the `speed` command to avoid connectivity issues between devices. When troubleshooting switch port issues, the duplex and speed settings should be checked.

Note Mismatched settings for the duplex mode and speed of switch ports can cause connectivity issues. Auto negotiation failure creates mismatched settings.

All fiber optic ports, such as 100BASE-FX ports, operate only at one preset speed and are always full-duplex.

Use the Syntax Checker in Figure 2 to configure port F0/1 of switch S1.

Refer to
Online Course
for Illustration

2.1.2.3 Auto-MDIX

Until recently, certain cable types (straight-through or crossover) were required when connecting devices. Switch-to-switch or switch-to-router connections required using different Ethernet cables. Using the automatic medium-dependent interface crossover (auto-MDIX) feature on an interface eliminates this problem. When auto-MDIX is enabled, the interface

automatically detects the required cable connection type (straight-through or crossover) and configures the connection appropriately. When connecting to switches without the auto-MDIX feature, straight-through cables must be used to connect to devices such as servers, workstations, or routers and crossover cables must be used to connect to other switches or repeaters.

With auto-MDIX enabled, either type of cable can be used to connect to other devices, and the interface automatically corrects for any incorrect cabling. On newer Cisco routers and switches, the `mdix auto` interface configuration mode command enables the feature. When using auto-MDIX on an interface, the interface speed and duplex must be set to `auto` so that the feature operates correctly.

The commands to enable auto-MDIX are shown in Figure 1.

Note The auto-MDIX feature is enabled by default on Catalyst 2960 and Catalyst 3560 switches, but is not available on the older Catalyst 2950 and Catalyst 3550 switches.

To examine the auto-MDIX setting for a specific interface, use the `show controllers ethernet-controller` command with the `phy` keyword. To limit the output to lines referencing auto-MDIX, use the `include Auto-MDIX` filter. As shown in Figure 2, the output indicates On or Off for the feature.

Use the Syntax Checker in Figure 3 to configure the FastEthernet 0/1 interface on S2 for auto-MDIX.

Refer to
Online Course
for Illustration

2.1.2.4 Verifying Switch Port Configuration

Figure 1 describes some of the options for the `show` command that are helpful in verifying common configurable switch features.

Figure 2 shows sample abbreviated output from the `show running-config` command. Use this command to verify that the switch has been correctly configured. As seen in the output for S1, some key information is shown:

- Fast Ethernet 0/18 interface configured with the management VLAN 99
- VLAN 99 configured with an IP address of 172.17.99.11 255.255.0.0
- Default gateway set to 172.17.99.1

The `show interfaces` command is another commonly used command, which displays status and statistics information on the network interfaces of the switch. The `show interfaces` command is frequently used when configuring and monitoring network devices.

Figure 3 shows the output from the `show interfaces fastEthernet 0/18` command. The first line in the figure indicates that the FastEthernet 0/18 interface is up/up meaning that it is operational. Further down the output shows that the duplex is full and the speed is 100 Mb/s.

Refer to
Online Course
for Illustration

2.1.2.5 Network Access Layer Issues

The output from the `show interface` command can be used to detect common media issues. One of the most important parts of this output is the display of the line and data

link protocol status. Figure 1 indicates the summary line to check the status of an interface.

The first parameter (FastEthernet0/1 is up) refers to the hardware layer and, essentially, reflects whether the interface is receiving the carrier detect signal from the other end. The second parameter (line protocol is up) refers to the data link layer and reflects whether the data link layer protocol keepalives are being received.

Based on the output of the **show interface** command, possible problems can be fixed as follows:

- If the interface is up and the line protocol is down, a problem exists. There could be an encapsulation type mismatch, the interface on the other end could be error-disabled, or there could be a hardware problem.
- If the line protocol and the interface are both down, a cable is not attached or some other interface problem exists. For example, in a back-to-back connection, the other end of the connection may be administratively down.
- If the interface is administratively down, it has been manually disabled (the **shutdown** command has been issued) in the active configuration.

Figure 2 shows an example of **show interface** command output. The example shows counters and statistics for the FastEthernet0/1 interface.

Some media errors are not severe enough to cause the circuit to fail, but do cause network performance issues. Figure 3 explains some of these common errors which can be detected with using the **show interface** command.

“Input errors” is the sum of all errors in datagrams that were received on the interface being examined. This includes runts, giants, CRC, no buffer, frame, overrun, and ignored counts. The reported input errors from the **show interface** command include the following:

- **Runt Frames** - Ethernet frames that are shorter than the 64-byte minimum allowed length are called runts. Malfunctioning NICs are the usual cause of excessive runt frames, but they can be caused by the same issues as excessive collisions.
- **Giants** - Ethernet frames that are longer than the maximum allowed length are called giants. Giants are caused by the same issues as those that cause runts.
- **CRC errors** - On Ethernet and serial interfaces, CRC errors usually indicate a media or cable error. Common causes include electrical interference, loose or damaged connections, or using the incorrect cabling type. If you see many CRC errors, there is too much noise on the link and you should inspect the cable for damage and length. You should also search for and eliminate noise sources, if possible.

“Output errors” is the sum of all errors that prevented the final transmission of datagrams out of the interface that is being examined. The reported output errors from the **show interface** command include the following:

- **Collisions** - Collisions in half-duplex operations are completely normal and you should not worry about them, as long as you are pleased with half-duplex operations. However, you should never see collisions in a properly designed and configured network that uses full-duplex communication. It is highly recommended that you use full-duplex unless you have older or legacy equipment that requires half-duplex.

- **Late collisions** - A late collision refers to a collision that occurs after 512 bits of the frame (the preamble) have been transmitted. Excessive cable lengths are the most common cause of late collisions. Another common cause is duplex misconfiguration. For example, you could have one end of a connection configured for full-duplex and the other for half-duplex. You would see late collisions on the interface that is configured for half-duplex. In that case, you must configure the same duplex setting on both ends. A properly designed and configured network should never have late collisions.

Refer to
Online Course
for Illustration

2.1.2.6 Troubleshooting Network Access Layer Issues

Most issues that affect a switched network are encountered during the original implementation. Theoretically, after it is installed, a network continues to operate without problems. However, cabling gets damaged, configurations change, and new devices are connected to the switch that require switch configuration changes. Ongoing maintenance and troubleshooting of the network infrastructure is required.

To troubleshoot these issues when you have no connection or a bad connection between a switch and another device, follow this general process:

Use the **show interface** command to check the interface status.

If the interface is down:

- Check to make sure that the proper cables are being used. Additionally, check the cable and connectors for damage. If a bad or incorrect cable is suspected, replace the cable.
- If the interface is still down, the problem may be due to a mismatch in speed setting. The speed of an interface is typically auto-negotiated; therefore, even if it is manually configured on one interface, the connecting interface should auto-negotiate accordingly. If a speed mismatch does occur through misconfiguration or a hardware or software issue, then that may result in the interface going down. Manually set the same speed on both connection ends if a problem is suspected.

If the interface is up, but issues with connectivity are still present:

- Using the **show interface** command, check for indications of excessive noise. Indications may include an increase in the counters for runts, giants, and CRC errors. If there is excessive noise, first find and remove the source of the noise, if possible. Also, verify that the cable does not exceed the maximum cable length and check the type of cable that is used. For copper cable, it is recommended that you use at least Category 5.
- If noise is not an issue, check for excessive collisions. If there are collisions or late collisions, verify the duplex settings on both ends of the connection. Much like the speed setting, the duplex setting is usually auto-negotiated. If there does appear to be a duplex mismatch, manually set the duplex on both connection ends. It is recommended to use full-duplex if both sides support it.

Refer to
Online Course
for Illustration

2.2 Switch Security: Management and Implementation

2.2.1 Secure Remote Access

2.2.1.1 SSH Operation

Secure Shell (SSH) is a protocol that provides a secure (encrypted) management connection to a remote device. SSH should replace Telnet for management connections. Telnet is an older protocol that uses insecure plaintext transmission of both the login authentication (username and password) and the data transmitted between the communicating devices. SSH provides security for remote connections by providing strong encryption when a device is authenticated (username and password) and also for the transmitted data between the communicating devices. SSH is assigned to TCP port 22. Telnet is assigned to TCP port 23.

In Figure 1, an attacker can monitor packets using Wireshark. A Telnet stream can be targeted to capture the username and password.

In Figure 2, the attacker can capture the username and password of the administrator from the plaintext Telnet session.

Figure 3 shows the Wireshark view of an SSH session. The attacker can track the session using the IP address of the administrator device.

However, in Figure 4, the username and password are encrypted.

To enable SSH on a Catalyst 2960 switch, the switch must be using a version of the IOS software including cryptographic (encrypted) features and capabilities. In Figure 5, use the **show version** command on the switch to see which IOS the switch is currently running, and IOS filename that includes the combination “k9” supports cryptographic (encrypted) features and capabilities.

Refer to
Online Course
for Illustration

2.2.1.2 Configuring SSH

Before configuring SSH, the switch must be minimally configured with a unique hostname and the correct network connectivity settings.

- **Verify SSH support.** Use the **show ip ssh** command to verify that the switch supports SSH. If the switch is not running an IOS that supports cryptographic features, this command is unrecognized.

- **Configure the IP domain.**

Configure the IP domain name of the network using the **ip domain-name domain-name** global configuration mode command. In Figure 1, the *domain-name* value is **cisco.com**.

- **Generate RSA key pairs.**

Generating an RSA key pair automatically enables SSH. Use the **crypto key generate rsa** global configuration mode command to enable the SSH server on the switch and generate an RSA key pair. When generating RSA keys, the administrator is prompted to enter a modulus length. Cisco recommends a minimum modulus size of 1,024

bits (see the sample configuration in Figure 1). A longer modulus length is more secure, but it takes longer to generate and to use.

Note To delete the RSA key pair, use the `crypto key zeroize rsa` global configuration mode command. After the RSA key pair is deleted, the SSH server is automatically disabled.

- **Configure user authentication.**

The SSH server can authenticate users locally or using an authentication server. To use the local authentication method, create a username and password pair using the `username username password password` global configuration mode command. In the example, the user `admin` is assigned the password `ccna`.

- **Configure the vty lines.** Enable the SSH protocol on the vty lines using the `transport input ssh` line configuration mode command. The Catalyst 2960 has vty lines ranging from 0 to 15. This configuration prevents non-SSH (such as Telnet) connections and limits the switch to accept only SSH connections. Use the `line vty` global configuration mode command and then the `login local` line configuration mode command to require local authentication for SSH connections from the local username database. Use the Syntax Checker in Figure 2 to configure SSH on switch S1.

Refer to
Online Course
for Illustration

2.2.1.3 Verifying SSH

On a PC, an SSH client, such as PuTTY, is used to connect to an SSH server. For the examples in Figures 1 to 3, the following have been configured:

- SSH enabled on switch S1
- Interface VLAN 99 (SVI) with IP address 172.17.99.11 on switch S1
- PC1 with IP address 172.17.99.21

In Figure 1, the PC initiates an SSH connection to the SVI VLAN IP address of S1.

In Figure 2, the user has been prompted for a username and password. Using the configuration from the previous example, the username `admin` and password `ccna` are entered. After entering the correct combination, the user is connected via SSH to the CLI on the Catalyst 2960 switch.

To display the version and configuration data for SSH on the device that you configured as an SSH server, use the `show ip ssh` command. In the example, SSH version 2 is enabled. To check the SSH connections to the device, use the `show ssh` command (see Figure 3).

Refer to Packet
Tracer Activity
for this chapter

2.2.1.4 Packet Tracer - Configuring SSH

SSH should replace Telnet for management connections. Telnet uses insecure plaintext communications. SSH provides security for remote connections by providing strong encryption of all transmitted data between devices. In this activity, you will secure a remote switch with password encryption and SSH.

Refer to
Online Course
for Illustration

2.2.2 Security Concerns in LANs

2.2.2.1 Common Security Attacks: MAC Address Flooding

Basic switch security does not stop malicious attacks. Security is a layered process that is essentially never complete. The more aware the team of networking professionals within an organization are regarding security attacks and the dangers they pose, the better. Some types of security attacks are described here, but the details of how some of these attacks work are beyond the scope of this course. More detailed information is found in the CCNA WAN Protocols course and the CCNA Security course.

MAC Address Flooding

The MAC address table in a switch contains the MAC addresses associated with each physical port and the associated VLAN for each port. When a Layer 2 switch receives a frame, the switch looks in the MAC address table for the destination MAC address. All Catalyst switch models use a MAC address table for Layer 2 switching. As frames arrive on switch ports, the source MAC addresses are recorded in the MAC address table. If an entry exists for the MAC address, the switch forwards the frame to the correct port. If the MAC address does not exist in the MAC address table, the switch floods the frame out of every port on the switch, except the port where the frame was received.

The MAC address flooding behavior of a switch for unknown addresses can be used to attack a switch. This type of attack is called a MAC address table overflow attack. MAC address table overflow attacks are sometimes referred to as MAC flooding attacks, and CAM table overflow attacks. The figures show how this type of attack works.

In Figure 1, host A sends traffic to host B. The switch receives the frames and looks up the destination MAC address in its MAC address table. If the switch cannot find the destination MAC in the MAC address table, the switch then copies the frame and floods (broadcasts) it out of every switch port, except the port where it was received.

In Figure 2, host B receives the frame and sends a reply to host A. The switch then learns that the MAC address for host B is located on port 2 and records that information into the MAC address table.

Host C also receives the frame from host A to host B, but because the destination MAC address of that frame is host B, host C drops that frame.

As shown in Figure 3, any frame sent by host A (or any other host) to host B is forwarded to port 2 of the switch and not broadcast out every port.

MAC address tables are limited in size. MAC flooding attacks make use of this limitation to overwhelm the switch with fake source MAC addresses until the switch MAC address table is full.

As shown in Figure 4, an attacker at host C can send frames with fake, randomly-generated source and destination MAC addresses to the switch. The switch updates the MAC address table with the information in the fake frames. When the MAC address table is full of fake MAC addresses, the switch enters into what is known as fail-open mode. In this mode, the switch broadcasts all frames to all machines on the network. As a result, the attacker can see all of the frames.

Some network attack tools can generate up to 155,000 MAC entries on a switch per minute. Depending on the switch, the maximum MAC address table size varies.

As shown in Figure 5, as long as the MAC address table on the switch remains full, the switch broadcasts all received frames out of every port. In this example, frames sent from host A to host B are also broadcast out of port 3 on the switch and seen by the attacker at host C.

One way to mitigate MAC address table overflow attacks is to configure port security.

Refer to
Online Course
for Illustration

2.2.2.2 Common Security Attacks: DHCP Spoofing

DHCP is the protocol that automatically assigns a host a valid IP address out of a DHCP pool. DHCP has been in use for nearly as long as TCP/IP has been the main protocol used within industry for allocating clients IP addresses. Two types of DHCP attacks can be performed against a switched network: DHCP starvation attacks and DHCP spoofing.

In DHCP starvation attacks, an attacker floods the DHCP server with DHCP requests to use up all the available IP addresses that the DHCP server can issue. After these IP addresses are issued, the server cannot issue any more addresses, and this situation produces a denial-of-service (DoS) attack as new clients cannot obtain network access. A DoS attack is any attack that is used to overload specific devices and network services with illegitimate traffic, thereby preventing legitimate traffic from reaching those resources.

In DHCP spoofing attacks, an attacker configures a fake DHCP server on the network to issue DHCP addresses to clients. The normal reason for this attack is to force the clients to use false Domain Name System (DNS) or Windows Internet Naming Service (WINS) servers and to make the clients use the attacker, or a machine under the control of the attacker, as their default gateway.

DHCP starvation is often used before a DHCP spoofing attack to deny service to the legitimate DHCP server, making it easier to introduce a fake DHCP server into the network.

To mitigate DHCP attacks, use the DHCP snooping and port security features on the Cisco Catalyst switches. These features are covered in a later topic.

Refer to
Online Course
for Illustration

2.2.2.3 Common Security Attacks: Leveraging CDP

The Cisco Discovery Protocol (CDP) is a proprietary protocol that all Cisco devices can be configured to use. CDP discovers other Cisco devices that are directly connected, which allows the devices to auto-configure their connection. In some cases, this simplifies configuration and connectivity.

By default, most Cisco routers and switches have CDP-enabled on all ports. CDP information is sent in periodic, unencrypted broadcasts. This information is updated locally in the CDP database of each device. Because CDP is a Layer 2 protocol, CDP messages are not propagated by routers.

CDP contains information about the device, such as the IP address, software version, platform, capabilities, and the native VLAN. This information can be used by an attacker to find ways to attack the network, typically in the form of a denial-of-service (DoS) attack.

The figure is a portion of a Wireshark capture showing the contents of a CDP packet. The Cisco IOS software version discovered via CDP, in particular, would allow the attacker to determine whether there were any security vulnerabilities specific to that particular version of IOS. Also, because CDP is not authenticated, an attacker could craft bogus CDP packets and send them to a directly-connected Cisco device.

It is recommended that you disable the use of CDP on devices or ports that do not need to use it by using the `no cdp run` global configuration mode command. CDP can be disabled on a per port basis.

Telnet Attacks

The Telnet protocol is insecure and can be used by an attacker to gain remote access to a Cisco network device. There are tools available that allow an attacker to launch a brute force password-cracking attack against the vty lines on the switch.

Brute Force Password Attack

The first phase of a brute force password attack starts with the attacker using a list of common passwords and a program designed to try to establish a Telnet session using each word on the dictionary list. If the password is not discovered by the first phase, a second phase begins. In the second phase of a brute force attack, the attacker uses a program that creates sequential character combinations in an attempt to guess the password. Given enough time, a brute force password attack can crack almost all passwords used.

To mitigate against brute force password attacks use strong passwords that are changed frequently. A strong password should have a mix of upper and lowercase letters and should include numerals and symbols (special characters). Access to the vty lines can also be limited using an access control list (ACL).

Telnet DoS Attack

Telnet can also be used to launch a DoS attack. In a Telnet DoS attack, the attacker exploits a flaw in the Telnet server software running on the switch that renders the Telnet service unavailable. This sort of attack prevents an administrator from remotely accessing switch management functions. This can be combined with other direct attacks on the network as part of a coordinated attempt to prevent the network administrator from accessing core devices during the breach.

Vulnerabilities in the Telnet service that permit DoS attacks to occur are usually addressed in security patches that are included in newer Cisco IOS revisions.

Note It is a best practice to use SSH, rather than Telnet for remote management connections.

Refer to
Interactive Graphic
in online course.

2.2.2.4 Activity - Identify Common Security Attacks

Refer to
Online Course
for Illustration

2.2.3 Security Best Practices

2.2.3.1 Best Practices

Defending your network against attack requires vigilance and education. The following are best practices for securing a network:

- Develop a written security policy for the organization.
- Shut down unused services and ports.
- Use strong passwords and change them often.
- Control physical access to devices.

- Avoid using standard insecure HTTP websites, especially for login screens; instead use the more secure HTTPS.
- Perform backups and test the backed up files on a regular basis.
- Educate employees about social engineering attacks, and develop policies to validate identities over the phone, via email, and in person.
- Encrypt and password-protect sensitive data.
- Implement security hardware and software, such as firewalls.
- Keep software up-to-date by installing security patches weekly or daily, if possible.

These methods are only a starting point for security management. Organizations must remain vigilant at all times to defend against continually evolving threats. Use network security tools to measure the vulnerability of the current network.

Refer to
Online Course
for Illustration

2.2.3.2 Network Security Tools and Testing

Network security tools help a network administrator test a network for weaknesses. Some tools allow an administrator to assume the role of an attacker. Using one of these tools, an administrator can launch an attack against the network and audit the results to determine how to adjust security policies to mitigate those types of attacks. Security auditing and penetration testing are two basic functions that network security tools perform.

Network security testing techniques may be manually initiated by the administrator. Other tests are highly automated. Regardless of the type of testing, the staff that sets up and conducts the security testing should have extensive security and networking knowledge. This includes expertise in the following areas:

- Network security
- Firewalls
- Intrusion prevention systems
- Operating systems
- Programming
- Networking protocols (such as TCP/IP)

Refer to
Online Course
for Illustration

2.2.3.3 Network Security Audits

Network security tools allow a network administrator to perform a security audit of a network. A security audit reveals the type of information an attacker can gather simply by monitoring network traffic.

For example, network security auditing tools allow an administrator to flood the MAC address table with fictitious MAC addresses. This is followed by an audit of the switch ports as the switch starts flooding traffic out of all ports. During the audit, the legitimate MAC address mappings are aged out and replaced with fictitious MAC address mappings. This determines which ports are compromised and not correctly configured to prevent this type of attack.

Timing is an important factor in performing the audit successfully. Different switches support varying numbers of MAC addresses in their MAC table. It can be difficult to determine the ideal amount of spoofed MAC addresses to send to the switch. A network administrator also has to contend with the age-out period of the MAC address table. If the spoofed MAC addresses start to age out while performing a network audit, valid MAC addresses start to populate the MAC address table, and limiting the data that can be monitored with a network auditing tool.

Network security tools can also be used for penetration testing against a network. Penetration testing is a simulated attack against the network to determine how vulnerable it would be in a real attack. This allows a network administrator to identify weaknesses within the configuration of networking devices and make changes to make the devices more resilient to attacks. There are numerous attacks that an administrator can perform, and most tool suites come with extensive documentation detailing the syntax needed to execute the desired attack.

Because penetration tests can have adverse effects on the network, they are carried out under very controlled conditions, following documented procedures detailed in a comprehensive network security policy. An off-line test bed network that mimics the actual production network is the ideal. The test bed network can be used by networking staff to perform network penetration tests.

Refer to
Online Course
for Illustration

2.2.4 Switch Port Security

2.2.4.1 Secure Unused Ports

Disable Unused Ports

A simple method that many administrators use to help secure the network from unauthorized access is to disable all unused ports on a switch. For example, if a Catalyst 2960 switch has 24 ports and there are three Fast Ethernet connections in use, it is good practice to disable the 21 unused ports. Navigate to each unused port and issue the Cisco IOS `shutdown` command. If a port later on needs to be reactivated, it can be enabled with the `no shutdown` command. The figure shows partial output for this configuration.

It is simple to make configuration changes to multiple ports on a switch. If a range of ports must be configured, use the `interface range` command.

```
Switch(config)#interface rangetype module/first-number - last-number
```

The process of enabling and disabling ports can be time-consuming, but it enhances security on the network and is well worth the effort.

Refer to
Online Course
for Illustration

2.2.4.2 DHCP Snooping

DHCP snooping is a Cisco Catalyst feature that determines which switch ports can respond to DHCP requests. Ports are identified as trusted and untrusted. Trusted ports can source all DHCP messages; untrusted ports can source requests only. Trusted ports host a DHCP server or can be an uplink toward the DHCP server. If a rogue device on an untrusted port attempts to send a DHCP response packet into the network, the port is shut down. This feature can be coupled with DHCP options in which switch information, such as the port ID of the DHCP request, can be inserted into the DHCP request packet.

As shown in Figures 1 and 2, untrusted ports are those not explicitly configured as trusted. A DHCP binding table is built for untrusted ports. Each entry contains a client MAC address, IP address, lease time, binding type, VLAN number, and port ID recorded as clients make DHCP requests. The table is then used to filter subsequent DHCP traffic. From a DHCP snooping perspective, untrusted access ports should not send any DHCP server responses.

These steps illustrate how to configure DHCP snooping on a Catalyst 2960 switch:

- Enable DHCP snooping using the `ip dhcp snooping` global configuration mode command.
- Enable DHCP snooping for specific VLANs using the `ip dhcp snooping vlan number` command.
- Define ports as trusted at the interface level by defining the trusted ports using the `ip dhcp snooping trust` command.
- (Optional) Limit the rate at which an attacker can continually send bogus DHCP requests through untrusted ports to the DHCP server using the `ip dhcp snooping limit rate rate` command.

Refer to
Online Course
for Illustration

2.2.4.3 Port Security: Operation

Port Security

All switch ports (interfaces) should be secured before the switch is deployed for production use. One way to secure ports is by implementing a feature called port security. Port security limits the number of valid MAC addresses allowed on a port. The MAC addresses of legitimate devices are allowed access, while other MAC addresses are denied.

Port security can be configured to allow one or more MAC addresses. If the number of MAC addresses allowed on the port is limited to one, then only the device with that specific MAC address can successfully connect to the port.

If a port is configured as a secure port and the maximum number of MAC addresses is reached, any additional attempts to connect by unknown MAC addresses will generate a security violation. Figure 1 summarizes these points.

Secure MAC Address Types

There are a number of ways to configure port security. The type of secure address is based on the configuration and includes:

- **Static secure MAC addresses** - MAC addresses that are manually configured on a port by using the `switchport port-security mac-address mac-address` interface configuration mode command. MAC addresses configured in this way are stored in the address table and are added to the running configuration on the switch.
- **Dynamic secure MAC addresses** - MAC addresses that are dynamically learned and stored only in the address table. MAC addresses configured in this way are removed when the switch restarts.
- **Sticky secure MAC addresses** - MAC addresses that can be dynamically learned or manually configured, then stored in the address table and added to the running configuration.

Sticky Secure MAC addresses

To configure an interface to convert dynamically learned MAC addresses to sticky secure MAC addresses and add them to the running configuration, you must enable sticky learning. Sticky learning is enabled on an interface by using the `switchport port-security mac-address sticky` interface configuration mode command.

When this command is entered, the switch converts all dynamically learned MAC addresses, including those that were dynamically learned before sticky learning was enabled, to sticky secure MAC addresses. All sticky secure MAC addresses are added to the address table and to the running configuration.

Sticky secure MAC addresses can also be manually defined. When sticky secure MAC addresses are configured by using the `switchport port-security mac-address sticky mac-address` interface configuration mode command, all specified addresses are added to the address table and the running configuration.

If the sticky secure MAC addresses are saved to the startup configuration file, then when the switch restarts or the interface shuts down, the interface does not need to relearn the addresses. If the sticky secure addresses are not saved, they will be lost.

If sticky learning is disabled by using the `no switchport port-security mac-address sticky` interface configuration mode command, the sticky secure MAC addresses remain part of the address table, but are removed from the running configuration.

Figure 2 shows the characteristics of stick secure MAC addresses.

Note that `switchport port-security` commands will not function until port security is enabled.

Refer to
Online Course
for Illustration

2.2.4.4 Port Security: Violation Modes

It is a security violation when either of these situations occurs:

- The maximum number of secure MAC addresses have been added to the address table for that interface, and a station whose MAC address is not in the address table attempts to access the interface.
- An address learned or configured on one secure interface is seen on another secure interface in the same VLAN.

An interface can be configured for one of three violation modes, specifying the action to be taken if a violation occurs. The figure presents which kinds of data traffic are forwarded when one of the following security violation modes are configured on a port:

- **Protect** - When the number of secure MAC addresses reaches the limit allowed on the port, packets with unknown source addresses are dropped until a sufficient number of secure MAC addresses are removed, or the number of maximum allowable addresses is increased. There is no notification that a security violation has occurred.
- **Restrict** - When the number of secure MAC addresses reaches the limit allowed on the port, packets with unknown source addresses are dropped until a sufficient number of secure MAC addresses are removed, or the number of maximum allowable addresses is increased. In this mode, there is a notification that a security violation has occurred.

- **Shutdown** - In this (default) violation mode, a port security violation causes the interface to immediately become error-disabled and turns off the port LED. It increments the violation counter. When a secure port is in the error-disabled state, it can be brought out of this state by entering the `shutdown` and `no shutdown` interface configuration mode commands.

To change the violation mode on a switch port, use the `switchport port-security violation {protect | restrict | shutdown}` interface configuration mode command.

Refer to
Online Course
for Illustration

2.2.4.5 Port Security: Configuring

Figure 1 summarizes the default port security configuration on a Cisco Catalyst switch.

Figure 2 shows the Cisco IOS CLI commands needed to configure port security on the Fast Ethernet F0/18 port on the S1 switch. Notice that the example does not specify a violation mode. In this example, the violation mode is shutdown (the default mode).

Figure 3 shows how to enable sticky secure MAC addresses for port security on Fast Ethernet port 0/19 of switch S1. As stated earlier, the maximum number of secure MAC addresses can be manually configured. In this example, the Cisco IOS command syntax is used to set the maximum number of MAC addresses to 50 for port 0/19. The violation mode is set to shutdown, by default.

Refer to
Online Course
for Illustration

2.2.4.6 Port Security: Verifying

Verify Port Security

After configuring port security on a switch, check each interface to verify that the port security is set correctly, and check to ensure that the static MAC addresses have been configured correctly.

Verify Port Security Settings

To display port security settings for the switch or for the specified interface, use the `show port-security [interface interface-id]` command. The output for the dynamic port security configuration is shown in Figure 1. By default, there is one MAC address allowed on this port.

The output shown in Figure 2 shows the values for the sticky port security settings. The maximum number of addresses is set to 50, as configured.

Note The MAC address is identified as a sticky MAC.

Sticky MAC addresses are added to the MAC address table and to the running configuration. As shown in Figure 3, the sticky MAC for PC2 has been added to the running configuration for S1.

Verify Secure MAC Addresses

To display all secure MAC addresses configured on all switch interfaces, or on a specified interface with aging information for each, use the `show port-security address` command. As shown in Figure 4, the secure MAC addresses are listed along with the types.

Refer to
Online Course
for Illustration

2.2.4.7 Ports in Error Disabled State

When a port is configured with port security, a violation can cause the port to become error disabled. When a port is error disabled, it is effectively shut down and no traffic is sent or received on that port. A series of port security related messages display on the console (Figure 1).

Note The port protocol and link status is changed to down.

The port LED will change to orange. The `show interface` command identifies the port status as `err-disabled` (Figure 2). The output of the `show port-security interface` command now shows the port status as `secure-shutdown`. Because the port security violation mode is set to `shutdown`, the port with the security violation goes to the error disabled state.

The administrator should determine what caused the security violation before re-enabling the port. If an unauthorized device is connected to a secure port, the port should not be re-enabled until the security threat is eliminated. To re-enable the port, use the `shutdown` interface configuration mode command (Figure 3). Then, use the `no shutdown` interface configuration command to make the port operational.

Refer to
Online Course
for Illustration

2.2.4.8 Network Time Protocol (NTP)

Having the correct time within networks is important. Correct time stamps are required to accurately track network events such as security violations. Additionally, clock synchronization is critical for the correct interpretation of events within syslog data files as well as for digital certificates.

Network Time Protocol (NTP) is a protocol that is used to synchronize the clocks of computer systems over packet-switched, variable-latency data networks. NTP allows network devices to synchronize their time settings with an NTP server. A group of NTP clients that obtain time and date information from a single source will have more consistent time settings.

A secure method of providing clocking for the network is for network administrators to implement their own private network master clocks, synchronized to UTC, using satellite or radio. However, if network administrators do not wish to implement their own master clocks because of cost or other reasons, other clock sources are available on the Internet. NTP can get the correct time from an internal or external time source including the following:

- Local master clock
- Master clock on the Internet
- GPS or atomic clock

A network device can be configured as either an NTP server or an NTP client. To allow the software clock to be synchronized by an NTP time server, use the `ntp server ip-address` command in global configuration mode. A sample configuration is shown in the Figure 1. Router R2 is configured as an NTP client, while router R1 serves as an authoritative NTP server.

To configure a device as having an NTP master clock to which peers can synchronize themselves, use the `ntp master [stratum]` command in global configuration mode. The stratum value is a number from 1 to 15 and indicates the NTP stratum number that the system will claim. If the system is configured as an NTP master and no stratum number is specified, it will default to stratum 8. If the NTP master cannot reach any clock with a lower stratum number, the system will claim to be synchronized at the configured stratum number, and other systems will be willing to synchronize to it using NTP.

Figure 2 displays the verification of NTP. To display the status of NTP associations, use the `show ntp associations` command in privileged EXEC mode. This command will indicate the IP address of any peer devices that are synchronized to this peer, statically configured peers, and stratum number. The `show ntp status` user EXEC command can be used to display such information as the NTP synchronization status, the peer that the device is synchronized to, and in which NTP strata the device is functioning.

Refer to **Packet Tracer Activity** for this chapter

2.2.4.9 Packet Tracer - Configuring Switch Port Security

In this activity, you will configure and verify port security on a switch. Port security allows you to restrict a port's ingress traffic by limiting the MAC addresses that are allowed to send traffic into the port.

Refer to **Packet Tracer Activity** for this chapter

2.2.4.10 Packet Tracer - Troubleshooting Switch Port Security

The employee who normally uses PC1 brought his laptop from home, disconnected PC1 and connected the laptop to the telecommunication outlet. After reminding him of the security policy that does not allow personal devices on the network, you now must reconnect PC1 and re-enable the port

Refer to **Lab Activity** for this chapter

2.2.4.11 Lab - Configuring Switch Security Features

In this lab, you will complete the following objectives:

- Part 1: Set Up the Topology and Initialize Devices
- Part 2: Configure Basic Device Settings and Verify Connectivity
- Part 3: Configure and Verify SSH Access on S1
- Part 4: Configure and Verify Security Features on S1

Refer to **Online Course** for illustration

2.3 Summary

Refer to **Lab Activity** for this chapter

2.3.1.1 Class Activity – Switch Trio

Switch Trio

Scenario

You are the network administrator for a small- to medium-sized business. Corporate headquarters for your business has mandated that on all switches in all offices, security must be implemented. The memorandum delivered to you this morning states:

“By Monday, April 18, 20xx, the first three ports of all configurable switches located in all offices must be secured with MAC addresses – one address will be reserved for the PC, one address will be reserved for the laptop in the office, and one address will be reserved for the office server.

If a port’s security is breached, we ask you to shut it down until the reason for the breach can be certified.

Please implement this policy no later than the date stated in this memorandum. For questions, call 1.800.555.1212. Thank you. The Network Management Team”

Work with a partner in the class and create a Packet Tracer example to test this new security policy. After you have created your file, test it with, at least, one device to ensure it is operational or validated.

Save your work and be prepared to share it with the entire class.

Refer to **Packet Tracer Activity** for this chapter

2.3.1.2 Packet Tracer - Skills Integration Challenge

The network administrator asked you to configure a new switch. In this activity, you will use a list of requirements to configure the new switch with initial settings, SSH, and port security.

Refer to **Online Course** for illustration

2.3.1.3 Summary

When a Cisco LAN switch is first powered on it goes through the following boot sequence:

1. First, the switch loads a power-on self-test (POST) program stored in ROM. POST checks the CPU subsystem. It tests the CPU, DRAM, and the portion of the flash device that makes up the flash file system.
2. Next, the switch loads the boot loader software. The boot loader is a small program stored in ROM and is run immediately after POST successfully completes.
3. The boot loader performs low-level CPU initialization. It initializes the CPU registers, which control where physical memory is mapped, the quantity of memory, and its speed.
4. The boot loader initializes the flash file system on the system board.
5. Finally, the boot loader locates and loads a default IOS operating system software image into memory and hands control of the switch over to the IOS.

The specific Cisco IOS file that is loaded is specified by the BOOT environmental variable. After the Cisco IOS is loaded it uses the commands found in the startup-config file to initialize and configure the interfaces. If the Cisco IOS files are missing or damaged, the boot loader program can be used to reload or recover from the problem.

The operational status of the switch is displayed by a series of LEDs on the front panel. These LEDs display such things as port status, duplex, and speed.

An IP address is configured on the SVI of the management VLAN to allow for remote configuration of the device. A default gateway belonging to the management VLAN must be configured on the switch using the `ip default-gateway` command. If the default gateway is not properly configured, remote management is not possible. It is recommended that Secure Shell (SSH) be used to provide a secure (encrypted) management connection

to a remote device to prevent the sniffing of unencrypted user names and passwords which is possible when using protocols such as Telnet.

One of the advantages of a switch is that it allows full-duplex communication between devices effectively doubling the communication rate. Although it is possible to specify the speed and duplex settings of a switch interface, it is recommended that the switch be allowed to set these parameters automatically to avoid errors.

Switch port security is a requirement to prevent such attacks as MAC Address Flooding and DHCP Spoofing. Switch ports should be configured to allow only frames with specific source MAC addresses to enter. Frames from unknown source MAC addresses should be denied and cause the port to shut down to prevent further attacks.

Port security is only one defense against network compromise. There are 10 best practices that represent the best insurance for a network:

- Develop a written security policy for the organization.
- Shut down unused services and ports.
- Use strong passwords and change them often.
- Control physical access to devices.
- Avoid using standard insecure HTTP websites, especially for login screens. Instead use the more secure HTTPS.
- Perform backups and test the backed up files on a regular basis.
- Educate employees about social engineering attacks, and develop policies to validate identities over the phone, via email, and in person.
- Encrypt sensitive data and protect it with a strong password.
- Implement security hardware and software, such as firewalls.
- Keep IOS software up-to-date by installing security patches weekly or daily, if possible.

These methods are only a starting point for security management. Organizations must remain vigilant at all times to defend against continually evolving threats.

Go to the online course to take the quiz and exam.

Chapter 2 Quiz

This quiz is designed to provide an additional opportunity to practice the skills and knowledge presented in the chapter and to prepare for the chapter exam. You will be allowed multiple attempts and the grade does not appear in the gradebook.

Chapter 2 Exam

The chapter exam assesses your knowledge of the chapter content.

Your Chapter Notes