

Book III

Switching with Cisco Switches

The 5th Wave

By Rich Tennant



"My spam filter checks the recipient address, http links, and any writing that panders to postmodern English romanticism with conceits to 20th-century graphic narrative."

Contents at a Glance

Chapter 1: Introducing Layer 2 Switches	291
Layer 2 — Data Link Layer Review	291
Purpose of a Layer 2 Switch	292
Basic Switch Functions	296
Managing Port Security.....	306
Transmitting Unicast, Multicast, and Broadcast	307
Chapter 2: Managing a Switch Using Cisco IOS	313
Best Practice for Using Cisco Switches	313
Connecting to a Cisco Switch.....	315
Cisco Switch Startup Process	321
Configuring a Cisco Switch.....	324
Managing Cisco Switch Authentication	352
Chapter 3: Controlling Network Traffic with Cisco Switches	369
Sending to MAC Addresses in Remote Networks	369
Deciding the Fate of Frames	375
Switching in Half-Duplex and Full-Duplex Modes	378
Avoiding Loops with Spanning Tree Protocol (STP)	379
Chapter 4: Spanning Tree Protocol (STP)	385
Introducing the Spanning Tree Protocol (STP)	386
STP Operation Flow	389
Introducing Cisco Options for STP.....	401
Introducing Rapid Spanning Tree Protocol (RSTP)	405
EtherChannel.....	407
Monitoring STP	410
Chapter 5: Virtual Local Area Networks (VLANs)	415
Introducing Virtual Local Area Networks (VLANs).....	416
Benefits of VLANs	418
Managing VLANs	418
Identifying VLANs	421
VLAN Trunking.....	422
VLAN Trunking Protocol (VTP)	434
Routing Traffic from One VLAN to Another	438
Chapter 6: Voice over IP (VoIP)	445
Introducing Voice over IP (VoIP).....	446
VoIP Requires Quality of Service (QoS).....	446
Cisco IP Phone	447
Cisco Discovery Protocol (CDP).....	450
Configuring VoIP on Cisco Switches	451
Chapter 7: Troubleshooting a Switch Using Cisco IOS	455
Troubleshooting Cisco Switches	455

Chapter 1: Introducing Layer 2 Switches

Exam Objectives

- ✓ Describing data link OSI Layer 2 and how it relates to Layer 2 switches
- ✓ Describing the purpose of a Layer 2 switch
- ✓ Differentiating a Layer 2 switch from a hub or a bridge
- ✓ Describing the basic Layer 2 switch functions
- ✓ Managing Layer 2 switch port security
- ✓ Describing how a Layer 2 switch handles unicast, multicast, and broadcast transmissions
- ✓ Describing MAC address table thrashing and broadcast storms

Read this chapter to find out about Layer 2 switches.

Layer 2 — Data Link Layer Review

Layer 2 in the Open Systems Interconnection (OSI) and TCP/IP network models is the data link layer. The data link layer transmits data on the physical medium. The data link layer uses physical addresses assigned to each physical network device in the local network to route data from one physical device to another. These addresses are called Media Access Control (MAC) addresses in TCP/IP. MAC addresses uniquely identify a specific network device, such as a switch or a router, or a network interface card (NIC), in a host device.

The data link layer has the following features:

- ◆ Receives each packet from the network layer on the sending host
- ◆ Wraps up the packet in a *data frame* along with local routing data (the physical MAC address)

- ◆ Sends the data frame to the physical layer to code an electrical or optical signal

The physical layer transmits the data frame over a wire or over the air (wireless transmission).

On the receiving host, the data link layer does the following:

- ◆ Unwraps the data frame received
- ◆ Extracts the packet out of the data frame
- ◆ Sends the packet up to the network layer

The Ethernet TCP/IP protocol is used at Layer 2 for data link operations. Ethernet, defined by the IEEE 802.X standards, is the TCP/IP protocol that operates at Layer 2 to handle data link functions. Hence, whenever you consider Layer 2 in TCP/IP, you really need to think about Ethernet and physical MAC addresses. Layer 2, the data link layer, is only concerned with local-area networks (LANs). The mission of the data link layer is to handle data frame transmission locally between two devices connected on a LAN. The data link layer is concerned with LANs, implemented with Ethernet, and addressed using physical MAC addresses.

Layer 2 switches use the MAC address of each device that transmits data-link frames through the switch to decide whether to forward that frame on an outbound port. Layer 2 switches also remember MAC addresses of connected devices to decide on which outbound port to send data-link frames.

Purpose of a Layer 2 Switch

To understand the purpose of a Layer 2 switch, you need to know a little bit about older computer networks. First-generation LANs were connected using a single coaxial cable. Host devices in the LAN shared the bandwidth of that cable. Figure 1-1 illustrates such a network. Observe that hosts Alex, Claire, and John are connected to the same coaxial cable. The bandwidth is shared. Worse yet, everyone sees all frames sent on the cable: Alex sends a frame to John, but Claire also sees the frame and needs to discard it because it is not addressed to her. You can imagine that Claire needs to spend some time discarding frames, and so do Alex and John whenever they see a frame that is not addressed to them. This is not efficient.

Hubs

Later, hubs were introduced to interconnect several host devices using one cable for each device. This greatly simplified connections because each host was connected individually to the hub. It is very easy to connect or disconnect a host to a hub by simply plugging or unplugging the network cable linking the host to the hub. Hub networks were typically deployed using twisted-pair cable instead of coaxial cable. RJ-45 connectors were used, which greatly simplified connections.




Figure 1-1:
Flat bus
network
with single
collision
domain.

A hub works very much like a multiplexer or a multiple socket power bar: Data sent into a hub is sent out to all host devices connected to that hub except to the host that sent the data in. In other words, a hub forwards a data frame on all outbound ports, except on the port through which the frame came in.

Host devices share the bandwidth of the medium both when connected using coaxial cables and when connected using hubs. Because hosts share the same cable or the same hub, hosts can potentially send data frames at the same time on the medium. Hence, data frames can collide.

Ethernet uses carrier sense multiple access collision detect (CSMA/CD) to control data frame collisions. This works well to work around collisions, but it does consume bandwidth. The CSMA/CD collision-detection and back-off algorithms can consume a lot of bandwidth. This is why in older networks, only 50–60 percent of the bandwidth was really usable for data frame transmission. The rest was wasted dealing with data frame collisions. The solution to this problem is to limit the collision domain, to make it as small as possible. This decreases or eliminates the risk of collision. That's where Layer 2 bridges and Layer 2 switches come into play.

Figure 1-2 shows a hub-based network. This solution provides an easier way to connect the hosts, but the hosts share the hub and the bandwidth. The hub does not filter the frames.




Figure 1-2:
Hub-based
network
with single
collision
domain.

Bridges

Bridges, like hubs, are used to interconnect devices in a LAN using a point-to-point connection for each device. Unlike hubs, bridges were the first network device that was able to limit the collision domain.

Bridges, like switches, create one collision domain per port and can forward data frames only on the outbound port that reaches the destination of the frame, as opposed to hubs, which send the frame out on all ports.



Think about it this way:

- ◆ A hub is like phone conferencing. Everyone on the phone call hears what you say.
- ◆ A bridge or a switch is like a normal telephone call: Only the person you call hears what you say.

A bridge is slower than a switch because it uses software instead of hardware application-specific integrated circuits (ASICs) to learn about MAC addresses and to decide whether to forward a data frame. Otherwise, a bridge works like a switch.

Switches

A Layer 2 switch is a network device that creates one collision domain per port and forwards data frames only on the outbound port that reaches the destination of the frame. Switch characteristics are as follows:

- ◆ Switches are faster than bridges because they use hardware ASICs instead of software to perform their operations.
- ◆ Switches are typically faster than routers because they do not need to look at the network layer (Layer 3) IP packet header. They only inspect the data-link (Layer 2) frame to look at the source and destination MAC address of the frame. This is why they are called Layer 2 switches: They only operate on the data-link (Layer 2) frame.

Figure 1-3 shows a switched network. This solution has several advantages:

- ◆ It provides an easier way to connect the hosts.
- ◆ The hosts work in their own isolated collision domain (one per port). Hence, frames do not collide.
- ◆ The switch filters the frames and forwards them only to the outbound port connected to the destination host. This is much more efficient than flooding the frame out on all ports, as hubs do.




Figure 1-3: Switched network with one collision domain per port.

The features of hubs, bridges, and switches are outlined in Table 1-1.

Table 1-1 Hubs, Bridges, and Switches			
<i>Feature</i>	<i>Hubs</i>	<i>Bridges</i>	<i>Switches</i>
Technology	Port multiplexing	Software switching	ASIC switching
Duplex	Half	Half	Half/Full
Speed	Turtle	Bear	Leopard
VLAN support	No	No	Yes
Collision domain	Whole hub	1 per port	1 per VLAN
Broadcast domain	Whole hub	Whole bridge	1 per VLAN

Basic Switch Functions

A Layer 2 switch must accomplish three tasks:

- ◆ Learn about the MAC addresses of devices connected to the switch
- ◆ Decide whether to forward frames it receives from host devices or other switches
- ◆ Avoid creating any Layer 2 loops

The following sections explain each of these functions in detail.



Address learning

Layer 2 switches learn the MAC addresses of devices connected to the switch as follows:

- ◆ The switch inspects each data frame that enters the switch. It saves the port number where the frame entered along with the source MAC address of that frame. The MAC address and the corresponding port number are saved in a MAC address table.
- ◆ As devices send frames into the switch, the switch slowly builds a complete MAC address table that contains
 - The MAC address of each device connected
 - The port number through which that device is sending frames into the switch

Figure 1-4 illustrates hosts Alex, Claire, and John interconnected in a Layer 2 switch. The MAC address table is empty at this point because none of the hosts have sent frames into the switch.

Figure 1-5 illustrates hosts Alex, Claire, and John interconnected in a Layer 2 switch. Here, Alex sends a frame to Claire. The switch inspects the data-link (Ethernet) frame and registers the source MAC address in the MAC address table along with the switch port through which the frame came in. Hence, the switch knows about Alex at this point.



Next, suppose that Claire responds to Alex: Claire sends a frame back to Alex. The switch inspects the data-link (Ethernet) frame and registers the source MAC address in the MAC address table along with the switch port through which the frame came in. Hence, the switch also learns about Claire at this point. Figure 1-6 illustrates the state of the MAC address table when the switch knows about Alex and Claire.





Figure 1-6:
Learning
another
MAC
address.

Eventually, John will also send a frame into the switch. At that point, the switch would register John's MAC address along with the switch port where John sent the frame in.

Figure 1-7 shows the complete MAC address table that stores the MAC address of each host and the switch port where they connect. This is the state of the MAC address table after each host has sent at least one frame through the switch. The switch keeps the MAC address table in its flash memory as long as the switch is powered on. However, if a host sends no frames for a certain amount of time, its MAC address is removed from the MAC address table to keep the table clean and current.



Flooding, forwarding, and filtering frames

Layer 2 switches need to decide whether to forward frames they receive. They also need to figure out over which outbound port they forward the frame. To do this, switches use the MAC address table:

1. The switch inspects each frame that enters on an inbound port: It looks at the source and destination MAC address of the frame.
2. The switch searches for the destination MAC address of the frame in the MAC address table:
 - If the switch finds the destination, the switch forwards the frame on the outbound port saved in the MAC address table for that destination MAC address.
 - If the switch does not find the destination MAC address in its MAC address table, it forwards the frame out on all outbound ports except on the one through which that frame came in.

Flooding a frame

Recall from an earlier section that initially, the MAC address table is empty, as shown in Figure 1-4. When Alex sends a frame to Claire, the switch learns about Alex, but it still does not know about Claire at this point. How does the switch know where to send Alex's frame? It doesn't. So the switch *floods* the frame out on all ports except the port through which it came in, that is, the port where Alex connects. Figure 1-8 illustrates this.





Figure 1-8:
Flooding a
frame.

Forwarding a frame

Claire responds to Alex: Claire sends a frame back to Alex. In this case, the switch knows about Alex because it has just registered his MAC address in the MAC address table when Alex sent the initial frame. So, the switch *forwards* the frame *only* out on the port where Alex connects. Figure 1-9 illustrates this.

Filtering a frame

Assume that the network has been changed, as shown in Figure 1-10. Host Monica was added to the network. Observe that John and Monica are interconnected with a hub, and the hub connects to the switch. Assume that John sends a frame to Monica. First, John's frame enters the hub. The hub floods the frame out on all ports except the sending port, that is, except the port where John connects. So, the frame goes out to Monica and to the switch. At this point, Monica receives the frame.



Now, the switch looks at the destination MAC address and searches it in the MAC address table. It does not find it. Hence, the switch would normally flood the frame out on all ports except on the incoming port, that is, except on e0/2. However, the switch does have MAC addresses registered for ports e0/0 and e0/1 that are different from the destination MAC address (D-MAC) of John's frame. So, the switch *filters* the frame because it is not for any of the hosts plugged in to its outgoing ports. Filtering a frame basically means discarding it without sending it out on any ports.

Next, assume that Monica sends a frame back to John. First, Monica's frame enters the hub. The hub floods it out on all ports except on the sending port, that is, except on Monica's port. So, the frame goes out to John and to the switch. At this point, John receives the frame. Figure 1-11 illustrates this situation.

The switch uses the frame to learn about Monica's MAC address, and filters the frame. So here it goes.

The switch first registers Monica's MAC address in the MAC address table. Observe that it registers it to the same outgoing switch port as John's MAC address. This is normal, because from the switch's perspective, both John and Monica come in on the same port: e0/2. That's the port where the hub connects.




Figure 1-10:
Filtering a
frame.

Next, the switch looks at the source MAC address and at the destination MAC address: They are both connected to the same switch port. Hence, the switch simply discards the frame because there is no point in sending it back to where it came from. In this case, the switch also filters the frame: It does not forward it out on any ports.

Basically, the switch assumes that if the sender and the receiver are connected to the same switch port, this is either

- ◆ A host sending a frame to itself
- ◆ A host sending a frame to another host connected locally into a hub

In both cases, the destination host should have already received the frame, so there is no need to forward it back to the same port.




Figure 1-11:
Filtering
another
frame.

Avoiding loops

Data-link (Ethernet) frames do not expire. An Ethernet frame sent to an unknown MAC address can bounce forever in the network. This is not good because it wastes bandwidth. This problem is exacerbated when several switches are interconnected with redundant links. It is best practice to use redundant links to interconnect Layer 2 switches to mitigate interswitch link failure risks. However, this amplifies the risk to have frames bouncing back and forth between switches.

Figure 1-12 shows two switches interconnecting hosts Alex, Claire, John, and Monica. Suppose that the MAC address tables are empty on both switches because none of the hosts has sent any frames yet. Here's an example:

1. Assume that Alex sends a frame to Claire. Recall from Figure 1-8 that the first time Alex sends a frame to Claire, the switch learns about Alex, but it still does not know about Claire. So, switch S1 needs to flood the frame out on all ports except the port through which it came in, that is, the port where Alex connects to S1.




Figure 1-12:
Broadcast
loop.

2. Switch S1 floods the frame out on all ports except the port through which it came in. The frame reaches switch S2. Switch S2 does not know about Claire either, so it also needs to flood the frame out on all ports except the port through which it came in.
3. Switch S2 floods the frame out on all ports except the port through which it came in.

Because two links are connecting the S1 and S2 switches, the frame that came in through fa0/2 on S2 is flooded out on fa0/3, and the frame that came in through fa0/3 on S2 is flooded out on fa0/2. Hence, switch S1 gets the frame back, twice: on fa0/2 and on fa0/3.

4. Because S1 still does not know Claire's MAC address, it floods the frame again on both fa0/2 and fa0/3. S2 receives the frame again, twice: on fa0/2 and on fa0/3.
5. Because S2 still does not know Claire's MAC address, it floods the frame out again on all ports, including fa0/2 and fa0/3, and so on.

This is a typical example of a *broadcast storm*: a broadcast frame that bounces forever among switches interconnected with redundant links.

Observe another problem in Figure 1-12. Because switches S1 and S2 receive the same frame on two different ports, they register the source MAC address (S-MAC) of host Alex for both ports. This is a problem because the switch does not know over which port it needs to send frames out to Alex whenever D-MAC is set to Alex's MAC address in a frame. This is a typical example of *MAC address table thrashing*.

Broadcast storms and MAC address table thrashing are both caused by loops between Layer 2 switches interconnected with redundant links. To avoid loops and bouncing frames, the *Spanning Tree Protocol (STP)* is used on Layer 2 switches. STP is explained in detail in Book III, Chapter 4.



Address learning: The process by which a Layer 2 switch learns which MAC address is connected to each switch port. The process involves the switch saving the port number where each frame enters along with the source MAC address of that frame. The incoming port number and the MAC address are saved in a MAC address table.

Flooding: The process by which a Layer 2 switch forwards a frame out on all outbound ports except on the port through which that frame came in. A switch floods a frame whenever it does not find the frame's destination MAC address in the MAC address table. It also floods broadcast frames.

Forwarding: The process by which a Layer 2 switch forwards a frame on the outbound port saved in the MAC address table for that destination MAC address.

Filtering: The process by which a Layer 2 switch discards a frame without sending it out on any ports. A switch filters a frame whenever the frame's source MAC address and frame's destination MAC address are registered to the same switch port, according to the MAC address table.

Avoiding loops: The process by which Layer 2 switches eliminate transmission loops created by redundant interswitch links. The Spanning Tree Protocol (STP) handles this process. Spanning Tree Protocol (STP) is discussed in detail in Book III, Chapter 4.

Broadcast storm: A broadcast frame that bounces forever among switches interconnected with redundant links. Broadcast storms waste bandwidth and may thrash the MAC address table.

MAC address table thrashing: Multiple ports attached to the same MAC address. The switch does not know on which outgoing port it can reach that MAC address.

Managing Port Security

Figure 1-11 shows a situation that can trigger some security risks and performance issues: You added host Monica to a network by simply plugging both Monica and John into a hub-connected to the switch. Now, both Monica and John are using the switch port that was originally reserved only for John.

While this is a quick-and-easy way to expand the network, it may not be the best way to do it, especially in very large LANs. For example, assume that you have an Ethernet jack in each cubicle in a large company office. As you can see, you can plug a hub or another switch into any Ethernet jack and add as many hosts as you want in that cubicle. This can turn into a major security and performance issue.

It is not secure to add uncontrolled host devices to a company network. It can also be a performance issue because the access and distribution network layers are probably designed to handle only one host, or maybe two hosts per cubicle. If a user adds a hub or a switch and plugs several hosts into the same Ethernet jack through the hub or the switch, the network load originating from that cubicle could increase above planned levels.

So, how can you configure the port to limit it to just John or to a certain maximum number of hosts? You can set filters and limits on each switch port using the `switchport port-security` Cisco command.

Filter based on MAC address

You can set the port to only support a specific MAC address using the `switchport port-security mac-address <MAC>` command. For example, you can run `switchport port-security mac-address 07-a1-3f-ee-f5-c2` to set the port to only allow John to connect.



Setting a static filter on a specific MAC address may not be a flexible solution. If you need to change the network interface card (NIC) in the John host, for example, you would have to change the filter on your switch to allow the MAC address of the new NIC. The sticky option is probably a better option. The sticky option is explained shortly.

Filter based on number of devices connected

You can set the port to only support a certain maximum number of devices connected using the `switchport port-security maximum <max#>` command.

Filter based on sticky MAC address

The most interesting port security option is the *sticky* option. The sticky option basically allows you to specify that the MAC address of the first device to send through that port will stick to the port. You can even specify how many devices stick to the port. The command is `switchport port-security mac-address sticky`. Next, you specify how many MAC addresses you allow to stick to the port using the `switchport port-security maximum <max#>` command.

In other words, if you replace `<max#>` with 2 in this command, you have the first two MAC addresses stick to that port. This is a more flexible alternative to filtering on actual MAC addresses.

Action triggered by filter

You can choose what action the switch needs to take should the filter or the limit be broken. You can use the `switchport port-security violation <action>` command to do this. For example, to have the port shut down upon a policy violation, you can run the `switchport port-security violation shutdown` command.

Before you run any of the `switchport` commands, you first need to enter the Cisco switch configuration terminal using the `configure terminal` command. You can also use the short form of this command: `conf t`. Next, you need to select a port to work with using the `interface <interface#>` command. Again you can use the short form: `int <interface#>`. For example, to work with the first Fast Ethernet interface on your switch, you can run `int fastethernet 0/0` or the even shorter form `int fa 0/0`. Then, you can execute the previous commands.

Read more about Cisco commands in Book III, Chapter 2.

Transmitting Unicast, Multicast, and Broadcast

So far in this chapter, you've seen how Layer 2 switches behave during unicast transmissions. A *unicast* transmission involves one device sending frames to a single target device. Only two devices are involved in unicast transmissions: the sending device and the target device. The device can be a host or a network device, such as a switch or a router.

A *multicast* transmission involves a device sending frames to multiple target devices. Note that a single device can send multiple frames to multiple devices, but *multicast* really means the device sending the *same* frame to multiple target devices. Normally, a multicast group is created, and target devices are added to the group. By default, a Layer 2 switch that receives a frame addressed to a multicast MAC address floods that frame out on all ports except on the port through which the frame came in.



You can configure the switch to send the frame only to the ports of the devices in the multicast group. However, those configuration options are beyond the scope of the CCNA test.

Multicast transmissions are typically used to stream data out to a group of hosts, for example, in IP-based video or audioconferencing.

A *broadcast* transmission involves a device sending frames to all devices in its local network. Whereas multicast transmissions involve the sending device and a specific subset of target devices in the local network, broadcast transmissions involve the sending device and *all* devices in the local network. Whenever a frame contains a destination MAC address (D-MAC) of FF-FF-FF-FF-FF-FF, it is considered to be a broadcast frame. A Layer 2 switch always floods out broadcast frames on all ports except on the port in which the frame came in.

Broadcast transmissions are typically used whenever a host tries to find a target device. For instance, you saw previously in Figure 1-5 that Alex sends a frame to Claire. The figure shows that Alex fills the Ethernet frame with his MAC address in S-MAC and Claire's MAC address in D-MAC. Here's the question: How does Alex know Claire's MAC address? Well, before Alex can send that frame, it needs to find Claire's MAC address.

To determine Claire's MAC address, Alex first looks into his Address Resolution Protocol (ARP) table, searching for the MAC address corresponding to the IP address of Claire. The ARP table is a table where responses to previous ARP requests are cached. If Alex does not find an entry in its ARP table that corresponds to Claire's IP address, it broadcasts an ARP request. In other words, it sends an ARP request packaged in a frame with the destination MAC address (D-MAC) set to FF-FF-FF-FF-FF-FF. The switch automatically floods the frame out on all ports when D-MAC is set to the FF-FF-FF-FF-FF-FF broadcast MAC address. All devices in the local network see the ARP request. Claire eventually responds with her MAC address. Alex saves Claire's MAC address in his ARP table. The next time, Alex will find Claire's MAC address in his ARP table. He will fill the D-MAC field with Claire's MAC address without having to broadcast another ARP request. So, you just saw a typical use for broadcast transmissions. Each host builds its ARP table the same way: using ARP request broadcasts.



Unicast transmission: Involves a device sending a frame to a single target device.

Multicast transmission: Involves a device sending a frame to multiple target devices.

Broadcast transmission: Involves a device sending a frame to all devices in its local network.

Prep Test

- 1** The data link layer, Layer 2 in the OSI network reference model, is represented in TCP/IP by which of the following?
 - A ☐ IP addresses
 - B ☐ Ethernet
 - C ☐ EtherChannel
 - D ☐ All of the above

- 2** What is the main function of the data link layer?
 - A ☐ To handle data frame transmission locally between two devices connected on a local-area network (LAN)
 - B ☐ To handle data frame transmission between two devices connected on a wide-area network (WAN)
 - C ☐ To handle IP packet routing locally between two nodes connected on a local-area network (LAN)
 - D ☐ To handle IP packet routing between two nodes connected on a wide-area network (WAN)

- 3** Layer 2 switches rely on the _____ to determine whether they forward data-link frames on an outgoing port.
 - A ☐ Firmware version
 - B ☐ IP address
 - C ☐ MAC address
 - D ☐ All of the above

- 4** How does a hub forward a data frame?
 - A ☐ Only on the outbound port where the target firmware connects
 - B ☐ Only on the outbound port where the target IP address connects
 - C ☐ Only on the outbound port where the target MAC address connects
 - D ☐ On all outbound ports, except on the port through which the frame came in

5 Describe how a Layer 2 switch learns addresses.

- A ☐ The switch saves the port number where each data frame enters the switch along with the source IP address of that frame in the IP address table.
- B ☐ The switch saves the port number where each data frame enters the switch along with the source MAC address of that frame in the MAC address table.
- C ☐ The switch saves the port number where each data frame enters the switch along with the source firmware version of that frame in the MAC address table.
- D ☐ The switch saves the port number where each data frame enters the switch along with the source firmware version of that frame in the IP address table.

6 Describe how a Layer 2 switch floods a frame.

- A ☐ The switch eliminates redundant interswitch links.
- B ☐ The switch sends a data frame only on the outgoing port where the destination device can be reached.
- C ☐ The switch sends a data frame on all outgoing ports except on the port where it entered the switch.
- D ☐ The switch discards a data frame.

7 Describe how a Layer 2 switch forwards a frame.

- A ☐ The switch eliminates redundant interswitch links.
- B ☐ The switch sends a data frame only on the outgoing port where the destination device can be reached.
- C ☐ The switch sends a data frame on all outgoing ports except on the port where it entered the switch.
- D ☐ The switch discards a data frame.

8 Describe how a Layer 2 switch avoids loops.

- A ☐ The switch eliminates redundant interswitch links.
- B ☐ The switch sends a data frame only on the outgoing port where the destination device can be reached.
- C ☐ The switch sends a data frame on all outgoing ports except on the port where it entered the switch.
- D ☐ The switch discards a data frame.

9 Unicast transmission involves which of the following?

- A ☐ A device sending a frame to multiple target devices
- B ☐ A device sending a frame to all devices in its local network
- C ☐ A device sending a frame to a single target device
- D ☐ All of the above

Answers

- 1 B. Ethernet. The data link layer, Layer 2 in the OSI network reference model, is represented in TCP/IP by Ethernet. Refer to *"Layer 2 — Data Link Layer Review."*
- 2 A. The main function of the data link layer is to handle data frame transmission locally between two devices connected on a local-area network (LAN). Check out *"Layer 2 — Data Link Layer Review."*
- 3 C. MAC address. Layer 2 switches inspect the destination MAC address of each frame entering the switch to determine whether to forward the frame on an outgoing port. Review *"Basic Switch Functions."*
- 4 D. Hubs always forward frames on all outbound ports, except on the port through which the frame entered the hub. Hubs multiplex the data frames from one ingoing port to multiple outgoing ports. Peruse *"Purpose of a Layer 2 Switch."*
- 5 B. To learn addresses, a Layer 2 switch saves the port number where each data frame enters the switch along with the source MAC address of that frame in the MAC address table. Check out *"Address learning."*
- 6 C. A Layer 2 switch floods a frame by sending the frame on all outgoing ports except on the port where it entered the switch. Read *"Flooding, forwarding, and filtering frames."*
- 7 B. A Layer 2 switch floods a frame by sending the frame only on the outgoing port where the destination device can be reached. Review *"Flooding, forwarding and filtering frames."*
- 8 A. Eliminates redundant interswitch links. A Layer 2 switch avoids switching loops by eliminating redundant interswitch links. Check out *"Avoiding loops."*
- 9 C. A device sending a frame to a single target device is a unicast transmission. Read *"Transmitting Unicast, Multicast, and Broadcast."*

Chapter 2: Managing a Switch Using Cisco IOS

Exam Objectives

- ✓ Connecting to a Cisco switch
- ✓ Understanding the startup process of a Cisco switch
- ✓ Configuring a Cisco switch
- ✓ Managing Cisco switch configurations
- ✓ Managing Cisco switch authentication

Managing a Cisco switch is very similar to managing a Cisco router. Most IOS commands are the same for switches and routers, but the output differs in some cases. Most IOS commands and GUI tools are available for both switches and routers. Some GUI tools are only available for routers, such as the Cisco Router and Security Device Manager (SDM) and SDM Express. Other GUI tools are only available for switches, such as the Cisco Device Manager. You find out about switch management IOS commands and GUI tools in this chapter.

Best Practice for Using Cisco Switches

When designing networks, the key point to remember is that top-of-the-line switches are best suited for either the *core layer* or the *distribution layer* of the network. Entry-level and midrange switches are best suited for either the *access layer* or the *distribution layer*. Cisco defines three layers in a network: the core layer, distribution layer, and access layer.

Top-of-the-line switches manage specialized services in a network, such as STP root bridge role, VLAN Membership Policy Server (VMPS), VLAN Trunking Protocol (VTP) domain control, Inter-VLAN routing, and LAN gateway connectivity.

These services are used throughout the network. Hence, they need to run on a highly efficient and highly available switch.

Figures 2-1 and 2-2 show the front and rear panels, respectively, of a Cisco Catalyst 2960, 24-port entry-level switch.



Figure 2-1:
Cisco
Catalyst
2960
switch and
accessory
items —
Front panel.



Figure 2-2:
Cisco
Catalyst
2960
switch and
accessory
items —
Rear panel.

Connecting to a Cisco Switch

Unlike a computer host, Cisco switches do not have a keyboard, monitor, or mouse device to allow direct user interaction. Cisco switches can be managed by connecting either locally or remotely to the switch from a computer host. You basically leverage the computer host user interface to interact with the Cisco switch.



Connecting locally

Cisco switches have several ports. Think of switch ports as plugs where you connect either computer hosts or other network devices.

Some switch ports are reserved for other purposes than network connectivity. These are the console port and the auxiliary port on a Cisco switch.

Console port

The *console port* is used to connect a management computer host to the switch using a rollover cable. You use this port whenever you want to connect locally to the console of your switch. The console is the default monitoring and configuration input and output facility of the switch. All Cisco devices have a *console facility*, where the Cisco operating system (IOS) displays status messages, error messages, diagnostic messages, and user prompts.



Think of the console facility as a virtual computer monitor and keyboard:

- ◆ Switch output messages are going to the virtual computer monitor.
- ◆ Switch user input is gathered using the virtual keyboard.

To access the console facility, you need to connect to the console port on your switch. You use a rollover cable connected at one end to the console port of your switch and connected at the other end to the management computer host.

Figure 2-3 illustrates a Cisco Catalyst switch with a management computer host connected to the switch's console port using a rollover cable. This figure illustrates the rear panel of the switch.

After you physically connect the switch to your computer using a rollover cable, you need to use a terminal emulation application on your management computer to open a serial terminal connection to your switch. The following Telnet applications support terminal emulation:

- ◆ HyperTerminal
- ◆ TeraTerm
- ◆ SecureCRT

You need to configure a connection profile for your Cisco switch using specific RS232 serial communication parameters:

- ◆ Baud rate: 9600
- ◆ Data bits: 8
- ◆ Parity: none
- ◆ Stop bits: 1
- ◆ Flow control: none

Serial communications (RS232) are governed by these parameters. Older computer hosts used to have serial communications ports (RS232) named COM1, COM2, COM3, and so on. Since the invention of USB, serial ports are being phased out. Your laptop probably has several USB ports but no serial port.





Figure 2-3:
Cisco
Catalyst
switch
console
connection.

Serial/USB port adapters are on the market. If you do not have a serial port on your computer, you need one of these adapters. Serial/USB port adapters create a virtual serial port on your computer by emulating an RS232 serial port using the USB physical port. Both RS232 and USB are standards that define serial communication frameworks.

Figure 2-4 illustrates a DB9 serial-to-USB converter (on the right in the picture). The switch rollover cable (on the left in the picture) plugs in to the converter. The converter plugs in to one of the USB ports on the laptop.

Figure 2-4:
Serial-to-
USB port
adapter.



Auxiliary port

The *auxiliary* port is used to connect a management computer host to the switch using a rollover cable, similarly to the *console* port. However, you use this port whenever you want to connect a modem locally to the auxiliary port. Next, you configure the modem to receive calls from a management computer host over a telephone line. After the management computer host establishes a modem connection to the switch, you are connected to the *console facility* of the switch. This port is really used for remote connections to your switch, but you still need to connect an answering modem locally to the auxiliary port of the switch. Note that some switches do not have an auxiliary port.



Having a modem answering calls and allowing console connections to the switch from remote computer hosts are security risks. It is important to configure passwords to challenge console port and auxiliary port connections. You will see how to configure console, auxiliary, and telnet passwords in the “*Managing Cisco IOS Authentication*” section later in this chapter.

Figure 2-5 illustrates a Cisco Catalyst switch with a management computer host connected to the switch’s auxiliary port using two modems.




Figure 2-5:
Cisco
Catalyst
switch
auxiliary
modem
connection.

Connecting remotely

It's nice to be able to connect locally to a switch, but in a typical networking environment, you would have several switches. Some switches may not be located on the same floor or even in the same building as where you are located. So you can connect remotely to a Cisco switch. You can connect to both the switch console and its terminal window from a remote management computer host.

You can connect remotely to Cisco switches in several ways:

- ◆ Connect to the auxiliary port of the switch using a modem connection. This connects you to the console of the switch from a remote location.
- ◆ Connect to the network IP address of the switch using a Telnet or secure Telnet application. Each switch has a network IP address that identifies the switch as an IP node in the network. You can connect to this IP address using any Telnet or secure Telnet application, just as you connect to a computer host using Telnet or secure Telnet.
- ◆ Connect to the console IP address of the switch using a Console Terminal Server. You can configure a Console Terminal Server that makes the switch console facility available at a specific IP address and at a specific TCP/IP port number in the network. You can connect to this IP:port address using any Telnet application.

It is important to configure passwords to challenge console port, auxiliary port, and Telnet connections. You see how to configure console, auxiliary, and Telnet passwords in the “Configuring a Cisco Switch” and the “Managing Cisco Switch Authentication” sections, later in this chapter.

Figure 2-6 illustrates a Cisco Catalyst switch with a management computer host connected to the switch IP address using Telnet.




Figure 2-7 illustrates a Cisco Catalyst switch with a management computer host connected to the switch console via a Terminal Server.




Figure 2-7:
Cisco
Catalyst
switch
Console
Terminal
Server
connection.

Cisco Switch Startup Process

A Cisco switch always follows the same process when it is powered up:

1. Run the POST: The Cisco switch runs the power-on self test.

The POST is a microprogram, stored in ROM, that is used to verify basic functionality of the Cisco switch hardware.

2. Boot up: The Cisco switch runs the *bootstrap program*, also known as the *boot loader software*.

The boot loader is a microprogram, stored in ROM, that is used to bring up the switch and transition it to normal operation mode by loading the Cisco IOS from flash memory. If the boot loader does not find a valid Cisco IOS in flash memory, it tries to load the IOS from either

- A TFTP server
- ROM

On switches, the bootstrap program also initializes the central processing unit (CPU) registers. CPU registers control where physical memory is mapped, how much memory is mapped, memory speed, and a few parameters that are beyond the scope of CCNA.

3. Load the IOS: The Cisco IOS loads into RAM.

By default, the Cisco IOS is loaded from flash memory. However, if the bootstrap program does not find a valid IOS in flash memory, it tries to load it from a TFTP server, if one is configured.

If no TFTP server is configured, or if no valid IOS is on the TFTP server, the bootstrap program loads the Rx-boot image from ROM. (The Rx-boot image is a subset of the Cisco IOS operating system.) The Rx-boot image is used to manage the boot process. The Rx-boot image comes up with the `(boot)>` prompt, where you can run various commands to manage the boot process. For instance, you can manually specify a location from which to load the IOS image.

4. Load the startup configuration.

After the Cisco switch loads the IOS in RAM, the IOS loads the device *startup configuration* from NVRAM. The startup configuration is loaded into RAM and becomes the *running configuration*, the configuration that changes dynamically while the Cisco device is running.

At this point, the Cisco device is in normal operation mode, ready for business.

The following IOS CLI segment shows a Cisco Catalyst 2960 switch starting up for the first time.

```
Using driver version 1 for media type 1
Base ethernet MAC Address: [ ... output cut ... ]
Xmodem file system is available.
The password-recovery mechanism is enabled.
Initializing Flash...
mifs[2]: 0 files, 1 directories
mifs[2]: Total bytes      :   3870720
mifs[2]: Bytes used       :    1024
mifs[2]: Bytes available  :   3869696
mifs[2]: mifs fsck took 0 seconds.
mifs[3]: 516 files, 19 directories
mifs[3]: Total bytes      :  27998208
mifs[3]: Bytes used       :   8684544
mifs[3]: Bytes available  :  19313664
mifs[3]: mifs fsck took 6 seconds.
...done Initializing Flash.
done.
```

The first part of the preceding example shows output from the bootstrap program. It shows the media driver version being used on the switch and the MAC address of the switch. It also shows that Xmodem connectivity is enabled. The password recovery mechanism is enabled. Finally, the CPU registers are initialized to map flash memory structures.

The output of the preceding example starts at the second step in the process (bootstrap). You do not see the output from the first step of the process (POST) here. The first part of the POST only displays messages upon test failures.

```
Loading "flash:c2960-lanlite-mz.122-37.EY/c2960-lanlite-mz.122-37.EY.bin"...@@@@
@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@
@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@
File "flash:c2960-lanlite-mz.122-37.EY/c2960-lanlite-mz.122-37.EY.bin"
  uncompressed and installed, entry point: 0x3000
executing...
```

This part loads the Cisco IOS from flash memory and executes it in RAM. The bootstrap program loads the `c2960-lanlite-mz.122-37.EY/c2960-lanlite-mz.122-37.EY.bin` IOS image located in the flash: memory structure.

Restricted Rights Legend

Use, duplication, or disclosure by the Government is subject to restrictions as set forth in subparagraph (c) of the Commercial Computer Software - Restricted Rights clause at FAR sec. 52.227-19 and subparagraph (c) (1) (ii) of the Rights in Technical Data and Computer Software clause at DFARS sec. 252.227-7013.

Cisco Systems, Inc.
170 West Tasman Drive
San Jose, California 95134-1706

```
Cisco IOS Software, C2960 Software (C2960-LANLITE-M), Version 12.2(37)EY, RELEASE
SOFTWARE (fc2)
Copyright (c) 1986-2007 by Cisco Systems, Inc.
Compiled Thu 28-Jun-07 18:07 by antonino
Image text-base: 0x00003000, data-base: 0x00D00000

Initializing flashfs...
Using driver version 1 for media type 1
mifs[3]: 0 files, 1 directories
mifs[3]: Total bytes      : 3870720
mifs[3]: Bytes used      : 1024
mifs[3]: Bytes available : 3869696
mifs[3]: mifs fsck took 0 seconds.
mifs[3]: Initialization complete.

mifs[4]: 516 files, 19 directories
mifs[4]: Total bytes      : 27998208
mifs[4]: Bytes used      : 8684544
mifs[4]: Bytes available : 19313664
mifs[4]: mifs fsck took 1 seconds.
mifs[4]: Initialization complete.

...done Initializing flashfs.
```

This part shows the name and version of the Cisco IOS. The flash memory file system is initialized.

```
POST: CPU MIC register Tests : Begin
POST: CPU MIC register Tests : End, Status Passed

POST: PortASIC Memory Tests : Begin
POST: PortASIC Memory Tests : End, Status Passed

POST: CPU MIC interface Loopback Tests : Begin
POST: CPU MIC interface Loopback Tests : End, Status Passed

POST: PortASIC RingLoopback Tests : Begin
POST: PortASIC RingLoopback Tests : End, Status Passed

POST: PortASIC CAM Subsystem Tests : Begin
POST: PortASIC CAM Subsystem Tests : End, Status Passed

POST: PortASIC Port Loopback Tests : Begin
POST: PortASIC Port Loopback Tests : End, Status Passed

Waiting for Port download...Complete

cisco WS-C2960-24-S (PowerPC405) processor (revision C0) with 61440K/4088K bytes
of memory.

Last reset from power-on
1 Virtual Ethernet interface
24 FastEthernet interfaces
The password-recovery mechanism is enabled.

64K bytes of flash-simulated non-volatile configuration memory.

[ ... Some Serial, Assembly, Revision numbers cut ... ]

Model number: WS-C2960-24-S
```

```
[ ... More Serial, Assembly, and Revision numbers cut ... ]

Switch   Ports   Model           SW Version      SW Image
-----
*        1      24             WS-C2960-24-S   12.2(37)EY      C2960-LANLITE-M

Press RETURN to get started!
```

This part shows the results of the second part of the POST and lists information about hardware components of this switch. This is a 24-port switch. All ports are compliant to the Fast Ethernet standard. The maximum bandwidth supported by Fast Ethernet is 100 Mbps. This switch has 64K in nonvolatile RAM (NVRAM).

```
[ ... Some status messages cut ... ]

Cisco IOS Software, C2960 Software (C2960-LANLITE-M), Version 12.2(37)EY, RELEASE
SOFTWARE (fc2)
Copyright (c) 1986-2007 by Cisco Systems, Inc.
Compiled Thu 28-Jun-07 18:07 by antonino

Would you like to terminate autoinstall? [yes]: yes

      --- System Configuration Dialog ---

Would you like to enter the initial configuration dialog? [yes/no]: no
Switch>
```

This part shows the final stage of the boot process (load startup configuration). Because this is a new switch, no startup configuration exists.

The IOS asks whether you would like to enter the initial configuration dialog. This dialog allows you to quickly configure the switch as follows:

- ◆ IP address
- ◆ Default gateway (router)
- ◆ Subnet mask
- ◆ Console password
- ◆ Host name
- ◆ Telnet password

Configuring a Cisco Switch

Cisco switches ship with several items:

- ◆ Console rollover cable: This is the cable you need to connect to the console port on the switch.
- ◆ AC power cord.

- ◆ 19-inch rack mounting brackets.
- ◆ Getting Started documentation CD.




Figure 2-8:
Cisco
Catalyst
2960
switch and
accessory
items.

These items are needed to initially power up and configure the switch. First, use the mounting brackets to mount the switch in a standard EIA-310-D data center rack. Next, connect the power cord to the power feed of the EIA-310-D data center rack. Before you power up the switch, you may want to review the Getting Started documentation on the CD provided with the switch.

Initial switch configuration

Cisco switches are not configured when they are new. They do not have a startup configuration. Cisco switches can interconnect computer hosts locally in an isolated network comprised of the switch itself and the computer hosts, even without a startup configuration. However, if you need to interconnect the switch to another switch or if you need to connect the switch to a gateway to reach remote networks, you need to create a startup configuration.

Cisco switches revert to setup mode anytime no startup configuration is saved in NVRAM. This happens in two situations:

- ◆ The device is new: No startup configuration exists on new devices.
- ◆ The startup configuration has been deleted: You can delete the startup configuration using IOS commands or using the Mode button to reset the switch.

You should verify the switch and create a startup configuration before you start using the switch, even if you use the switch only to interconnect computer hosts locally in an isolated network comprised of the switch itself and some computer hosts. If you ever decide to expand this small network by adding another switch and more computer hosts, or if you decide to connect this small network to a gateway to reach the Internet, you definitely need a startup configuration in NVRAM.

You can create a startup configuration using one of the following methods:

- ◆ Filling in the Express Setup Web form
- ◆ Using the switch Autoinstall feature
- ◆ Using the initial configuration dialog
- ◆ Using the Cisco IOS setup mode commands

Express Setup mode

Cisco provides a utility, named *Express Setup*, that allows you to initially configure the switch. You run Express Setup when you power on the switch the first time to configure the following settings on the switch:

- ◆ IP address
- ◆ Default gateway (router)
- ◆ Subnet mask
- ◆ Console password
- ◆ Host name
- ◆ Telnet password

To set the switch in Express Setup mode, press and hold the Mode button for 3 seconds. Release the button when all LEDs are green.

Next, browse to the IP address of the switch. This loads the Express Setup form, as illustrated in Figure 2-9.

Initializing the switch using the initial configuration dialog

The CCNA test expects you to know how to initialize the switch using the Cisco IOS setup mode.

Figure 2-9:
Cisco
Catalyst
2960
Express
Setup Web
form.

The first time you power on a Cisco switch, it has no startup configuration in NVRAM. You can connect to the console of the switch to get access to the IOS prompt. The switch will be in setup mode, because no startup configuration exists. In setup mode, the IOS asks you to configure the following:

- ◆ IP address
- ◆ Default gateway (router)
- ◆ Subnet mask
- ◆ Enable and Enable secret password
- ◆ Host name
- ◆ Telnet password

If choose not to configure the switch and use it without a startup configuration, you can avoid or abort setup mode using one the following methods:

- ◆ Exit setup mode by pressing Ctrl+C.
- ◆ Answer no when setup mode asks whether you want to configure the switch.
- ◆ Answer no when setup mode asks whether you want to save the configuration at the end of the series of setup questions.

You should verify the switch and create a startup configuration before you start using the switch, even if you use the switch only to interconnect computer hosts locally in an isolated network comprised of the switch itself and some computer hosts.

Here is an example of the initial setup mode dialog:

```
Using driver version 1 for media type 1
Base ethernet MAC Address: 00:25:b4:10:58:80
Xmodem file system is available.
The password-recovery mechanism is enabled.

[ ... some boot messages cut ... ]

Loading "flash:c2960-lanlite-mz.122-37.EY/c2960-lanlite-mz.122-37.EY.bin"...@
@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@
@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@
lanlite-mz.122-37.EY.bin" uncompressed and installed, entry point: 0x3000
executing...

Restricted Rights Legend

Use, duplication, or disclosure by the Government is
subject to restrictions as set forth in subparagraph
(c) of the Commercial Computer Software - Restricted
Rights clause at FAR sec. 52.227-19 and subparagraph
(c) (1) (ii) of the Rights in Technical Data and Computer
Software clause at DFARS sec. 252.227-7013.

Cisco Systems, Inc.
170 West Tasman Drive
San Jose, California 95134-1706

Cisco IOS Software, C2960 Software (C2960-LANLITE-M), Version 12.2(37)EY, RELEASE
SOFTWARE (fc2)
Copyright (c) 1986-2007 by Cisco Systems, Inc.
Compiled Thu 28-Jun-07 18:07 by antonino
Image text-base: 0x00003000, data-base: 0x00D00000

[ ... some boot messages cut ... ]

Press RETURN to get started!

[ ... some boot messages cut ... ]
```

The first part of the output, shown in the preceding example, displays various boot messages. The switch is now up and running.

Because no startup configuration exists in NVRAM, the IOS first asks whether you would like to let the switch run the Autoinstall feature to automatically configure the switch with a baseline configuration:

- ◆ If you answer yes to this question, or if you just press Enter to keep the default answer, the switch automatically configures itself with a baseline configuration to allow minimum functionality. After the autoinstall process completes, your switch is operational.
- ◆ If you answer no to the autoinstall question, the IOS next asks whether you want to run the initial configuration dialog. You type **yes** as follows to run the initial configuration dialog.

```

Would you like to terminate autoinstall? [yes]: no

--- System Configuration Dialog ---

Would you like to enter the initial configuration dialog? [yes/no]: yes

At any point you may enter a question mark '?' for help.
Use ctrl-c to abort configuration dialog at any prompt.
Default settings are in square brackets '[]'.

Basic management setup configures only enough connectivity
for management of the system, extended setup will ask you
to configure each interface on the system

Would you like to enter basic management setup? [yes/no]: yes
Configuring global parameters:

Enter host name [Switch]: SW1

The enable secret is a password used to protect access to
privileged EXEC and configuration modes. This password, after entered, becomes
encrypted in the configuration.
Enter enable secret: my_priv_encrypt_password

The enable password is used when you do not specify an
enable secret password, with some older software versions, and some boot
images.
Enter enable password: my_priv_password

The virtual terminal password is used to protect
access to the router over a network interface.
Enter virtual terminal password: my_telnet_password
Configure SNMP Network Management? [no]:

```

The second part of the output, shown in the preceding example, prompts you to set a host name and passwords on your switch. In this example

- ◆ Name the switch SW1.
- ◆ Set the encrypted privileged mode password to `my_priv_encrypt_password`.
- ◆ Set the unencrypted privileged mode password to `my_priv_password`.
- ◆ Set the VTY access line password to `my_telnet_password`.

The next section of the initial configuration dialog, shown in the following example, displays a summary of the Ethernet interfaces (ports) available on your switch.

```

Current interface summary

Any interface listed with OK? value "NO" does not have a valid configuration

Interface      IP-Address  OK? Method Status  Protocol
Vlan1          unassigned YES unset  down   down
FastEthernet0/1 unassigned YES unset  down   down
FastEthernet0/2 unassigned YES unset  down   down

```

```
[... some output cut ...]

FastEthernet0/23 unassigned YES unset down down
FastEthernet0/24 unassigned YES unset down down

Enter interface name used to connect to the
management network from the above interface summary: vlan1

Configuring interface Vlan1:
Configure IP on this interface? [no]: yes
IP address for this interface: 192.168.75.10
Subnet mask for this interface [255.255.255.0] :
Class C network is 192.168.75.0, 24 subnet bits; mask is /24
Would you like to enable as a cluster command switch? [yes/no]: no
```

The third part of the output, shown in the preceding example, prompts you to set up a management interface for the switch. You choose the first VLAN interface: VLAN ID 1. You set an IP address and subnet mask on the management interface.



VLAN 1, the first VLAN, is reserved for management purposes. It is sometimes called the administrative VLAN.

The next section of the initial configuration dialog, shown in the following example, displays a summary of the configuration that will be created on your switch.



Only the secret password is encrypted by default: the encrypted privileged mode password.

You can also encrypt the other passwords for increased security.

The following configuration command script was created:

```
hostname SW1
enable secret 5 $1$3GJW$pVC5U4qVd1bzy5x8kDkwa.
enable password my_priv_password
line vty 0 15
password my_telnet_password
no snmp-server
!
!
interface Vlan1
no shutdown
ip address 192.168.75.10 255.255.255.0
!
interface FastEthernet0/1
!
interface FastEthernet0/2
!
interface FastEthernet0/3
!
interface FastEthernet0/4
!

[... some output cut ...]
```

```
interface FastEthernet0/22
!
interface FastEthernet0/23
!
interface FastEthernet0/24
!
end
```

The last section of the initial configuration dialog, shown in the following example, prompts you to do one of the following:

- ◆ Save the configuration you just created to NVRAM.
- ◆ Return to the setup without saving the configuration you just created.
- ◆ Go straight to the IOS prompt without saving the configuration you just created.

You choose “2” to save the configuration you just created to NVRAM. This configuration becomes the startup configuration.

```
[0] Go to the IOS command prompt without saving this config.
[1] Return back to the setup without saving this config.
[2] Save this configuration to nvram and exit.
```

Enter your selection [2]:

Building configuration...

[OK]

Use the enabled mode 'configure' command to modify this configuration.

SW1>

Initializing the switch using Cisco IOS setup mode commands

You can configure the switch settings manually using Cisco IOS setup mode commands at any time.

The following sections show you a few basic Cisco IOS configuration commands that you can use at any time to configure your Cisco device.

Naming the switch

You can name your switch using the `hostname` Cisco IOS command. You should name all your switches using meaningful names to ease identification and management of each switch.

To configure the host name of the switch, run the following commands:

```
Switch>enable (or en)
Switch#configure terminal (or config t)
Switch(config)#hostname SW1
SW1(config)#exit
SW1#disable
SW1>
```

Configuring management IP address for the switch

You can configure an IP address and an IP gateway for your switch using the `ip address` and `ip default-gateway` Cisco IOS commands. This allows you to connect to the switch from remote locations using either Telnet or HTTP.

To configure the management IP address and default gateway on your switch, run the following commands:

```
SW1>enable (or en)
SW1#configure terminal (or config t)
SW1(config)#interface vlan1 (or int vlan1)
SW1(config-if)#ip address 192.168.75.10 255.255.255.0
SW1(config-if)#no shutdown
SW1(config-if)#exit
SW1(config)#ip default-gateway 192.168.75.1
SW1(config)#exit
SW1#disable
SW1>
```

The first two commands (`enable` and `configure terminal`) set the IOS in privileged global configuration mode. You can now run commands that configure global switch settings, that is, settings that apply to the whole switch.

The `interface vlan1` command selects an interface to work with. In this example, you select `vlan1`. Cisco switches support several virtual LANs (VLANs). They are numbered from 1 to 4094. The first VLAN, VLAN 1, is reserved for switch management. It is called the management (or administrative) VLAN. This is the VLAN on which you want to set the management IP address and the management default IP gateway.

The `ip address 192.168.75.10 255.255.255.0` command sets the IP address and the subnet mask on the interface that you selected previously. Observe that the IOS prompt now shows `SW1(config-if)#`. `Config-if` means that you are configuring an interface (`if`) right now.

The `no shutdown` command starts up the VLAN 1 interface. Recall that you enable components and services on a Cisco switch using the component or service name. You disable components and services on a Cisco switch using the component or service name prefixed with the `no` keyword. Here, `shutdown` is actually a state, not a component or service. An interface is either shut down or not shut down (started up). To summarize:

- ◆ To shut down an interface, select it and execute `shutdown`.
- ◆ To start an interface, select it and execute `no shutdown`.

The `ip default-gateway 192.168.75.1` command sets the default gateway for the switch. Observe that you first exit from the interface configuration mode (`SW1(config-if)# exit`) because the default gate-

way applies to the whole switch, not just to an interface. Now, you are back in global configuration mode (`SW1 (config) #`).

The Cisco IOS prompt is designed to show you the configuration mode that you're in:

- ◆ `(config)`: You are in global configuration mode.

In this mode, you can execute commands that configure global switch settings, that is, settings that apply to the whole switch.

- ◆ `(config-if)`: You are in interface configuration mode.

In this mode, you can execute commands that configure switch interface settings, that is, settings that apply to just one interface of the switch. You select the interface to work with using the `interface` command.

- ◆ `(config-if-range)`: You are in interface range configuration mode.

In this mode, you can execute commands that configure a range of switch interfaces, that is, settings that apply to a range of interfaces on the switch. You select the range of interfaces to work with using the `interface range` command.

Configuring passwords

You can configure authentication passwords for your switch using the `password` and `login` Cisco IOS commands. By default, new Cisco devices do not have a password configured.

You can configure several passwords for different types of access:

- ◆ **Console password:** Used for console access through the console port or through a Console Terminal Server.
- ◆ **Auxiliary password:** Used for console access through the auxiliary port using a modem.
- ◆ **VTY lines password:** Virtual type terminal (VTY) lines are used for Telnet and Secure Shell (SSH) access. They are called virtual type terminal (VTY) lines because no physical terminal is connected to the switch. You connect a computer host to the network and access the switch remotely using the switch management IP address. A terminal emulation program on the computer host emulates a physical terminal (TTY).
- ◆ **Privileged password:** Used for privileged mode access. Privileged mode is an “expert” management operation mode used to run some Cisco IOS commands.

The console port and the auxiliary port are enabled by default, even if no password is specified for them. This is a security risk. It is best practice to specify at least a console password on new Cisco devices.

The VTY lines (Telnet and Secure Shell) are disabled by default. You need to specify a password for the VTY lines to enable them.

To configure passwords, you need to instruct the Cisco device to prompt for authentication. You use the `login` Cisco IOS command to do this.

Best practice

At a minimum, you should set passwords for console and VTY access to secure access through the console port and to enable and secure remote access through Telnet or SSH.

The following sections describe how to configure passwords for console and VTY access using Cisco IOS commands.

Console password

To configure the console password, run the following commands:

```
SW1>enable (or en)
SW1#configure terminal (or config t)
SW1(config)#line console 0 (or line con 0)
SW1(config-line)#password my_password
SW1(config-line)#login
SW1(config-line)#exit
SW1(config)#exit
SW1#disable
SW1>
```

The first two commands (`enable` and `configure terminal`) set the IOS in privileged global configuration mode. You can now run commands that configure global switch settings, that is, settings that apply to the whole switch.

The `line console 0` command selects the console line. Cisco devices have only one console line: console 0.

The `password my_password` command sets the `my_password` password on the console access line.

Finally, you instruct the Cisco device to prompt for authentication by entering the `login` Cisco IOS command. From now on, anytime you connect to the console line, you will be prompted to provide a password. The password you provide must be `my_password`. To disable the authentication prompt, issue `line console 0` again and enter the `no login` command.

You use the `exit` commands to exit the `config-line` mode and to exit the `(config)` global configuration mode.

Telnet password

To configure a password for the VTY lines, run the following commands:

```

SW1>enable (or en)
SW1#configure terminal (or config t)
SW1(config)#line vty 0 ?
    <0-15> last line number
SW1(config-line)#line vty 0-15
SW1(config-line)#password my_telnet_password
SW1(config-line)#login
SW1(config-line)#exit
SW1(config)#exit
SW1#disable
SW1>

```

The first two commands (`enable` and `configure terminal`) set the IOS in privileged global configuration mode. You can now run commands that configure global switch settings, that is, settings that apply to the whole switch.

The `line vty 0 ?` command asks the Cisco IOS how many VTY lines are available. The response (`<0-15> last line number`) shows that 16 VTY lines are available on this particular switch (you can have 16 simultaneous Telnet sessions opened to this switch). Now, you can configure a password either on one of the VTY lines (by selecting that particular line) or on all VTY lines (by selecting the whole 0–15 line range).

The `line vty 0-15` command selects the whole 0–15 VTY line range. Cisco devices have several VTY *access lines*. Older switches used to have only four VTY lines. Newer switches have as many as 1,180 VTY lines. It is best practice to query the Cisco IOS about how many VTY lines are available.

You have two main reasons to have several VTY access lines on a Cisco device:

- ◆ **Allowing you to connect to the switch and connect to another device from the switch:** Two VTY lines are needed in this case: one line to connect into the switch and another line to connect out of the switch to another device.
- ◆ **Allowing several administrators to work on the switch:** In large networks, more than one administrator may manage the network. More than one administrator may need to connect from a remote location to the same switch using Telnet or SSH. This is typical with large core switches.

The `password my_telnet_password` command sets the `my_telnet_password` password on the VTY access lines. You can set a different password on each line. However, this is not recommended. You don't know which VTY line you are on when you log in, so you wouldn't know which password to enter. Best practice is to select the whole range of VTY lines and set the same password on all of them.

Finally, you instruct the Cisco device to prompt for authentication by entering the `login` Cisco IOS command. From now on, anytime you connect to any of the VTY lines, you will be prompted to provide a password. The password you provide must be `my_telnet_password`. To disable the authentication prompt, issue `line vty 0-15` again and enter the `no login` command.

After you run the initial setup IOS commands, the switch has a startup configuration saved in NVRAM. The switch is reachable on the IP network. You can now connect to the switch either locally to its console or auxiliary port, or remotely to its IP address, as explained previously in the “Connecting to a Cisco Switch” section.

Configuring banners

You can configure a banner for your switch using the `banner` Cisco IOS command. The purpose of a banner is to display a brief message about the switch when you log in. This is useful when you have many switches spread out across multiple sites. It helps identifying the switch you log in to and its configuration and usage guidelines.

Banners are also useful for legal reasons: You can add a security warning in the banner message to warn users against unauthorized logins to the switch.

Although this does not prevent the user from logging in to the switch, it does provide grounds for legal action in the case of unauthorized logins. Always use passwords to secure your switches.

Four types of banners are available:

- ◆ **Message of the day (MOTD) banner:** This is the first banner to display when a user connects to a switch, regardless of the type of connection. This banner is used to display a message for all users connecting to the switch, *before* they log in. This is typically used for an unauthorized access warning.

- ◆ **Login banner:** This banner is displayed on TTY or VTY terminals after the MOTD banner.

This banner is typically used to display information about the switch and to provide guidelines on how to log in and use the switch.

- ◆ **Incoming terminal connection banner:** This banner is displayed on reverse TTY or VTY terminals after the MOTD banner.

This banner is used for the same reasons as the login banner. The only difference is the type of terminal connection: You can specify additional information in this banner for users connecting with reverse TTY or VTY terminals.

- ◆ **EXEC process creation banner:** This banner is displayed whenever an EXEC process is created, after the user is logged in. An EXEC process is created for each user EXEC prompt.

User EXEC prompts are established for each user connected to the switch, after they log in successfully. Hence, the EXEC process creation banner is displayed after the login/incoming banners, and after the MOTD banner.

To configure an MOTD banner on your switch, run the following commands:

```
SW1>enable (or en)
SW1#configure terminal (config t)
SW1(config)#banner motd -
Enter TEXT message. End with character '-'.
$This switch is property of silange.com networks. Unauthorized access is
  prohibited. Please disconnect if you are not a silange.com employee,
  customer, or business partner.
-
SW1(config)#
```

The `banner motd -` command starts the banner text editor using `' - '` as the delimiting character. The delimiting character is used by the IOS to determine when you are done typing the banner text. You enter the banner text of your choice and end by typing the delimiting character (`' - '` in this case) and pressing Enter. The delimiting character can be any character, but it is important to understand that you cannot use that character within the text of the banner. Cisco IOS would interpret that as the end of the banner text.

Resetting a Cisco switch

Express Setup and IOS setup mode run whenever no startup configuration exists in NVRAM. In other words, the switch reverts to setup mode anytime no startup configuration exists in NVRAM. Normally, this only happens when the switch is new and not yet configured.

However, you may want to clear the current configuration in NVRAM and start a new configuration from scratch. This is typically a last-resort troubleshooting method: When problems affect the switch and no troubleshooting method fixes the problems, your last resort is to reset the switch and reconfigure it.

If you need to run Express Setup or the IOS setup again, you need to reset the switch. Resetting the switch clears the configuration of the switch.

After it is reset, the switch is in the state it was when you purchased it:

- ◆ No IP address
- ◆ No host name

- ◆ No default gateway (router)
- ◆ No subnet mask
- ◆ No console password
- ◆ No Telnet password
- ◆ No startup configuration

You should back up the startup configuration of a switch to a backup computer host before resetting the switch. You learn how to backup the startup configuration file to a backup computer host in the “*Managing Cisco Switch Configuration*” section in this chapter.

To reset the switch and clear its startup configuration from NVRAM using the Mode button, follow these steps:

1. Press and hold the Mode button.

The switch LEDs begin blinking after a few seconds.

2. Continue to hold the Mode button.

The LEDs stop blinking after about 7 seconds. The switch reboots.

3. Release the Mode button after the LEDs stop blinking and the switch starts rebooting.

The switch comes back up in setup mode without a startup configuration in NVRAM.

Managing Cisco switch configuration

Ongoing, you can manage your switch using the following:

- ◆ Cisco IOS commands
- ◆ The Cisco Device Manager Web tool
- ◆ The Cisco Network Assistant (CNA) tool

The Cisco Device Manager and the Cisco Network Assistant are two graphical user interface (GUI)-based tools that allow you to manage your switch and its startup and running configurations from a remote location using a computer host. Their GUI is much easier to use and more intuitive than Cisco IOS commands.

You find out about Cisco IOS switch configuration management commands in the following sections.

Managing the switch boot process

Cisco switches always follow the same boot process when they are powered up. This process involves four stages:

1. **Run the POST:** The Cisco switch runs the power-on self test to verify basic functionality of the Cisco switch hardware.
2. **Boot up:** The Cisco switch runs the *bootstrap program* to bring up the switch and transition it to normal operation mode by loading up the first Cisco IOS image from flash memory.
3. **Load the IOS:** The Cisco IOS loads into RAM from flash memory.
4. **Load the startup configuration:** After the Cisco switch loads the IOS in RAM, the IOS loads the device *startup configuration* from NVRAM. The startup configuration is loaded into RAM and becomes the *running configuration*, the configuration that changes dynamically while the Cisco device is running. At this point, the Cisco device is in normal operation mode, ready for business.

This process is also called the *automatic boot process*. The switch automatically boots up with current options and loads the startup configuration.

Boot command

By default, the Cisco *bootstrap program* loads the first Cisco IOS image file from flash memory when the switch is powered up. In some cases, you may need to boot a specific IOS image file from flash memory. For example, suppose that you download a new Cisco IOS image file that contains additional features, but you would like to test it on your switch before you remove the original IOS image file. You can configure the switch to boot with the new IOS image file using the `boot` command.

The `boot` command allows you to configure the boot process of a Cisco switch. You can do the following:

- ◆ Control which Cisco IOS image file is loaded
- ◆ Control which startup configuration file is loaded
- ◆ Enable Ctrl+Break during booting
- ◆ Enable manual booting
- ◆ Change the size of the NVRAM area that is formatted in Cisco IFS format

You can use the Cisco IOS contextual help to see the options available with the `boot` command:

```
SW1>
SW1>enable (or en)
Password: my_priv_encrypt_password
SW1#configure terminal (or config t, or conf t)
SW1(config)#
SW1(config)#boot ?
    boothlpr      Boot Helper System Image
    buffersize     Specify the size for filesystem-simulated NVRAM
```

<code>config-file</code>	Configuration File
<code>enable-break</code>	Enable Break while booting
<code>helper</code>	Helper Image(s)
<code>helper-config-file</code>	Helper Configuration File
<code>manual</code>	Manual Boot
<code>private-config-file</code>	Private Configuration File
<code>system</code>	System Image

```
SW1(config)#exit
SW1#disable
SW1>
```

Observe that you need to be in privileged global configuration mode to execute the `boot` command.

The `bootlpr` option is beyond the scope of the CCNA test.

The `buffer-size` option allows you to specify the size of the NVRAM buffer (memory area) that is formatted in Cisco IOS File System (IFS). You may need to increase the size of the IFS buffer if you need to store additional Cisco IOS images in flash memory.

The `config-file` option allows you to specify which configuration file is loaded during the boot process. This is useful when you need to test alternate configurations stored in different startup configuration files.

The `enable-break` option enables the Ctrl+Break key-combination option during the boot process. This allows you to interrupt the startup of the switch by pressing Ctrl+Break.

The `helper` and `helper-config-file` options are beyond the scope of the CCNA test.

The `manual` option enables manual booting. This allows you to boot the switch manually: The automatic boot process interrupts and displays the `switch: manual boot` prompt. You can manually boot the switch from the `switch: manual boot` prompt, controlling which IOS image is loaded and which configuration file is loaded. You can also use the `switch: manual boot` prompt to troubleshoot startup problems on your switch.

The `private-config-file` option allows you to specify which private configuration file is loaded during the boot process. This is useful when you need to test alternate private configurations stored in different startup configuration files. Private configuration files are used to store secured configuration data such as cryptographic encryption keys used for SSH.

The `system` option allows you to boot a specific Cisco IOS image file. This is useful in cases when, for example, you download a new Cisco IOS image file, containing additional features, and you would like to test it on your switch.

Here is an example in which you set your switch to boot with a specific Cisco IOS image file stored in flash memory:

```
SW1>
SW1>enable (or en)
Password:my_priv_encrypt_password
SW1#configure terminal (or config t, or conf t)
SW1(config)#
SW1(config)#boot system flash:/c2960-lanlite-mz.122-37.EY
SW1(config)#exit
SW1#disable
SW1>
```

Observe that you need to be in privileged global configuration mode to execute the `boot` command.

You can use the `show boot` command to verify the current boot options on a switch:

```
SW1>
SW1>enable (or en)
Password: my_priv_encrypt_password
SW1#show boot
BOOT path-list      : flash:/c2960-lanlite-mz.122-37.EY
Config file         : flash:/config.text
Private Config file : flash:/private-config.text
Enable Break        : no
Manual Boot         : no
HELPER path-list    :
Auto upgrade        : yes
Auto upgrade path   :
NVRAM/Config file   :
                    buffer size: 65536
SW1#disable
SW1>
```

Observe that you need to be in privileged EXEC mode to execute the `show boot` command. This command shows the current boot options. All options can be changed with the corresponding `boot` command, as explained earlier. After you change a boot option, you need to reboot your switch using the `reload` command in privileged EXEC mode.

Interrupting the boot process

In some cases, you may need to interrupt the automatic boot process to control some of the boot options. For example, you may have lost or forgotten the passwords on your switch. In this case, you do not want to boot the switch automatically, because you will be locked out of the switch by a password you do not have or do not remember. You need to break out of the boot process to reset the password.

You have three methods of interrupting the boot process on a Cisco switch:

- ◆ Press the Mode button during the boot process. This is the simplest and default method to interrupt the boot process of a Cisco switch.
- ◆ Change the boot process from automatic to manual and reboot the switch.
- ◆ Enable the boot break option, reboot the switch, and press Ctrl+Break during the boot process.

Interrupting the boot process with the Mode button

Follow these steps to interrupt the boot process of a Cisco switch using the Mode button on the front panel of the switch:

- 1. Reboot the switch using the `reload` command in privileged EXEC mode.**
- 2. Press the Mode button while the System LED (SYST) is flashing green within the first 15 seconds of the switch boot process.**
- 3. Continue to press the Mode button until the System LED (SYST) flashes briefly amber and then turns solid green.**
- 4. Release the Mode button when the System LED (SYST) becomes solid green.**

Interrupting the boot process by enabling the manual boot option

Follow these steps to interrupt the boot process of a Cisco switch using the manual boot option:

- 1. Enable the manual boot option.**
- 2. Reboot the switch using the `reload` command in privileged EXEC mode.**

You have the following options to control the manual boot option on a Cisco switch:

- ◆ To enable the manual boot option, run `boot manual` at the privileged global configuration prompt.
- ◆ To disable the manual boot option, run `no boot manual` at the privileged global configuration prompt.

The manual boot option is off by default. The switch boots up automatically without showing the `switch: manual boot` prompt.

Here is an example of how you enable manual boot:

```
SW1>enable (or en)
SW1#configure terminal (or config t)
Enter configuration commands, one per line. End with CNTL/Z.
SW1(config)#boot manual
```

```

SW1(config)#
SW1(config)#do show boot (or do sh boot)
BOOT path-list      : flash:c2960-lanlite-mz.122-37.EY/c2960-lanlite-mz.122-37.
                     EY.bin
Config file         : flash:/config.text
Private Config file : flash:/private-config.text
Enable Break       : no
Manual Boot        : yes
HELPER path-list   :
Auto upgrade       : yes
Auto upgrade path  :
NVRAM/Config file  :
                   buffer size: 65536

SW1(config)#exit
SW1#reload
    System configuration has been modified. Save? [yes/no]: n
    Proceed with reload? [confirm] <press Enter>

00:20:51: %SYS-5-RELOAD: Reload requested by console. Reload Reason: Reload
command.

[... reboot output cut ...]

The system is not configured to boot automatically. The
following command will finish loading the operating system
software:

    boot

switch:

```

The first two commands (`enable` and `configure terminal`) set the IOS in privileged global configuration mode. You can now run commands that configure global router settings, that is, settings that apply to the whole router.

Next, enable the manual boot process using the `boot manual` IOS command. Then, look at the current boot options using the `show boot` command. Observe that you use the `do` prefix because the `show` command does not run in configuration mode; you need to run it at the privileged EXEC prompt. Recall that the “do” Cisco IOS prefix allows you to run non-configuration commands in configuration mode.

Think of the “do” Cisco IOS prefix as an override option that allows you to momentarily exit configuration mode, run the command that follows the “do” prefix and come back into configuration mode. Note that the manual boot option is enabled now (**Manual Boot: yes**).

The next few commands reboot the switch.

After the switch comes back up again, the bootstrap program interrupts the boot process and waits for commands at the `switch: manual boot` prompt.

Interrupting the boot process with Ctrl+Break

Follow these steps to interrupt the boot process of a Cisco switch using the Ctrl+Break key combination:

- 1. Enable the boot process break option.**
- 2. Reboot the switch using the `reload` command in privileged EXEC mode.**
- 3. Press Ctrl and Break simultaneously while the System LED (SYST) is flashing green within the first 15 seconds of the switch boot process.**

The following options allow you to control the boot process break option on a Cisco switch:

- ◆ To enable the boot process break option, run `boot enable-break` at the privileged global configuration prompt.
- ◆ To disable the boot process break option, run `no boot enable-break` at the privileged global configuration prompt.

The boot process break option is disabled by default.

Here is an example of how you enable the boot process break option:

```
SW1>enable (or en)
SW1#configure terminal (or config t)
Enter configuration commands, one per line. End with CNTL/Z.
SW1(config)#boot enable-break
SW1(config)#
SW1(config)#do show boot (or do sh boot)
BOOT path-list      : flash:c2960-lanlite-mz.122-37.EY/c2960-lanlite-mz.122-37.
                      EY.bin
Config file         : flash:/config.text
Private Config file : flash:/private-config.text
Enable Break       : yes
Manual Boot        : no
HELPER path-list   :
Auto upgrade       : yes
Auto upgrade path  :
NVRAM/Config file  :
                    buffer size: 65536

SW1(config)#exit
SW1#reload
    System configuration has been modified. Save? [yes/no]: n
    Proceed with reload? [confirm] <press Enter>

00:20:51: %SYS-5-RELOAD: Reload requested by console. Reload Reason: Reload
command.

[... reboot output cut ...]

[... Press CTRL-BREAK simultaneously ...]
```

The system boot process was interrupted. The following command will finish loading the operating system software:

```
boot
```

```
switch:
```

The first two commands (`enable` and `configure terminal`) set the IOS in privileged global configuration mode. You can now run commands that configure global router settings, that is, settings that apply to the whole router.

Next, enable the boot process break option using the `boot enable-break` IOS command. Then, look at the current boot options using the `show boot` command. Note that the enable break option is on now.

The next few commands reboot the switch.

After the switch comes back up again, press Ctrl and Break simultaneously to interrupt the boot process. After it is interrupted, the bootstrap program waits for commands at the `switch: manual boot` prompt.

Introducing the switch manual boot prompt

No matter how you interrupt the boot process, you end up at the `switch: manual boot` prompt. The manual boot prompt allows you to control the boot process of the switch.

The bootstrap program displays a message informing you that you can enter the `boot` command to continue booting up the switch. You can use the help command (?) at the `switch: boot` prompt to see the manual boot commands and options available.

The IOS has not yet loaded the startup configuration at this point. In fact, the bootstrap program has not even loaded the IOS at this point.

Managing configurations

Two configurations are available on any Cisco switch: the startup configuration and the running configuration. The following sections review the characteristics of startup and running configurations and describe how to manage them on a Cisco switch.

Startup configuration

The startup configuration is the configuration that the Cisco IOS loads when it boots up. The startup configuration is stored in NVRAM, which keeps its contents even when the Cisco device is powered down.

Whenever no startup configuration exists in NVRAM, the switch starts in setup mode to allow you to create a startup configuration. After you complete the setup mode, you are prompted to save the configuration to NVRAM. If you answer yes, the configuration you created is saved to NVRAM; it becomes the startup configuration. You can also manually save the current configuration to NVRAM by using the `copy running-config startup-config` Cisco IOS command.

Cisco devices use the startup configuration data to configure the device before normal operation starts. The Cisco IOS loads the startup configuration from NVRAM into RAM. At that point, the startup configuration becomes the running configuration. The switch is up and ready in normal operation mode.

Running configuration

The running configuration is the dynamic data that changes while the Cisco device is in normal operation mode. This includes the following:

- ◆ ARP cache (MAC address tables)
- ◆ Routing tables
- ◆ STP data
- ◆ VLAN data
- ◆ EtherChannel configuration data
- ◆ Temporary buffers

The Cisco IOS loads the startup configuration from NVRAM into RAM during the boot process. After it is in RAM, the startup configuration becomes the running configuration and can change dynamically.

You can save the running configuration to NVRAM to replace the startup configuration with updated data. To do this, use the `copy running-config startup-config` Cisco IOS command. You can also save the running configuration to a file on a computer host.

Saving the running configuration to NVRAM

To save the running configuration to NVRAM, run the following commands:

```
SW1>enable (or en)
SW1#copy running-config startup-config (or copy run start)
Destination filename [startup-config]? <press Enter>
Building configuration...
[OK]
SW1#
```

The first line shows the `copy` command. You can also use the short form of this command (`copy run start`).

The second line asks where you want to copy the running configuration to. You have several possible destinations:

- ◆ `startup-config`: Use this keyword to copy the current running configuration over the startup configuration in NVRAM.
- ◆ `flash`: Use this keyword to copy the current running configuration to a file in flash memory, alongside the Cisco IOS.
- ◆ `ftp`: Use this keyword to copy the current running configuration to an FTP server. This option is useful for backing up the switch configuration to a remote computer host.
- ◆ `archive`: This is similar to the `ftp` option, except that you archive to a tape- or disk-based backup system directly, instead of passing through a computer host.

Many more possible copy destinations exist. You can use the Cisco IOS contextual help system to find out more about the other possible copy destinations by running the following command:

```
SW1#copy running-config ? (or copy run ?)
```

The `?` sign instructs the Cisco IOS to show you all the possible values that you can provide for the second argument of the `copy` command.

The default destination is the `startup-config` file in NVRAM. To keep the default answer, just press Enter. Sometimes the Cisco IOS prompts you to enter a value, and it shows the default answer in between brackets []. To keep the default answer, just press Enter.

Monitoring the running configuration

To monitor the running configuration, run the following command:

```
SW1#show running-config (or sh run)
```

Monitoring the startup configuration

To monitor the startup configuration, run the following command:

```
SW1#show startup-config (or sh start)
```

Deleting the startup configuration

Resetting the switch erases the startup configuration from NVRAM and reboots the switch. Upon reboot, the switch starts in setup mode because no startup configuration exists in NVRAM. This allows you to clear the current configuration in NVRAM and build a new configuration from scratch.

This is typically a last-resort troubleshooting method: When problems affect the switch and no troubleshooting method fixes the problems, your last resort is to reset the switch and reconfigure it.

You can also reset the switch using the `erase startup-config` Cisco IOS command.

You should back up the startup configuration to a backup computer host before deleting it. You find out how to back up the configuration to a backup computer host later in this chapter.

To delete the startup configuration, run the following commands:

```
SW1>enable (or en)
SW1#erase startup-config
Erasing the nvram filesystem removes all configuration files!
Continue? [confirm] <press Enter to confirm>
[OK]
Erase of nvram: complete
SW1#
```

Observe that the `erase` command warns you that the NVRAM contents will be reinitialized and prompts you to confirm. The default answer is `confirm`. To confirm, just press Enter.

Switch configuration backup and recovery best practice

It is best practice to save the running configuration to both

- ◆ The startup configuration file in NVRAM on the switch itself
- ◆ A computer host

This protects you against a total failure of the switch. If the switch fails and you need to replace it with a new one, you don't want to have to re-create the startup configuration from scratch. If you have that configuration saved on a computer host, you can load the startup configuration from the computer host onto the new switch; you're now ready to go in production with the new switch.

Back up the switch running configuration to a computer host

To back up the running configuration to a computer host, run the following commands:

```
SW1>enable (or en)
SW1#copy running-config tftp (or copy run tftp)
Address or name of remote host []? 192.168.75.5
Destination filename [SW1-config]? <press Enter>
812 bytes copied in 0.910 secs (892 bytes/sec)
SW1#
```

The `copy running-config tftp` Cisco IOS command on the second line initiates copying the switch running configuration to a TFTP server.

The third line prompts you for the IP address or host name of the remote host destination. Observe that, in this example, the IOS does not suggest any destination IP address or hostname “[]”. You can configure a default TFTP server on your switch. The IP address or hostname of that default TFTP server will then show up suggested between the square brackets.

The fourth line prompts you for the filename of the configuration file to be saved on the remote host. The switch name followed by `-config` is suggested by default. Simply press Enter to accept the suggested filename.

Now your running configuration is saved in a file named `SW1-config` on the remote computer host with IP address 192.168.75.5.

Recover a switch configuration from a computer host

Suppose that you need to load a configuration from a remote host to your switch. The same `copy` Cisco IOS command allows you to recover a configuration from a remote host:

```
SW1>enable (or en)
SW1#copy tftp running-config (or copy tftp run)
Address or name of remote host []? 192.168.75.5
Source filename []? SW1-config
Destination filename [running-config]? <press Enter>
Accessing tftp://192.168.75.5/SW1-config...
Loading SW1-config from 192.168.75.5
[OK - 812 bytes]
812 bytes copied in 10.120 secs (80 bytes/sec)
SW1#
```

The `copy tftp running-config` Cisco IOS command on the second line copies the switch running configuration from a TFTP server.

The third line prompts you for the IP address or host name of the remote host destination. Observe that, in this example, the IOS does not suggest any source IP address or hostname “[]”. You can configure a default TFTP server on your switch. The IP address or hostname of that default TFTP server will then show up suggested between the square brackets.

The fourth line prompts you for the source filename of the configuration file that you want to load from the remote host. Enter the name of a configuration file that you saved previously on the TFTP server. In this example, you enter the name of the configuration that you just copied to the TFTP server in the previous example (“SW1-config”).

The fifth line prompts you for the destination filename. The `running-config` filename is suggested. This means that the configuration you load

from the TFTP server will replace the current running configuration. Press Enter to accept the suggested choice.

Next, the `copy` command displays a few status messages.

At this point, your running configuration has been loaded from a file named `SW1-config` from a remote TFTP server host with IP address 192.168.75.5.

Managing switch configurations using Cisco IOS File System (IFS)

You saw in the previous two sections how to back up a switch configuration to a file on a computer host and how to restore a switch configuration from a backup file on a computer host. You used the `copy` Cisco IOS command to do this. This section shows you a new trick: how to manage the memory of a Cisco switch using Cisco IOS File System (IFS) commands.

You can manage the memory of a Cisco switch like a file system on a computer host. The Cisco IFS presents the memory of a Cisco switch as directories (or folders) and files within those directories, just like Microsoft Windows, Apple Mac OS X, Linux, or any other computer operating system presents the hard drive of a computer host as directories (or folders) and files.

The Cisco IFS commands are only available in privileged mode. So, you need to run the `enable` command first (or `en` in its short form).

First, verify the current directory:

```
SW1>enable (or en)
SW1#pwd
flash:
```

You are currently in the flash memory main directory, as shown by the `pwd` command. Next, run the Cisco IFS command `dir` to see the contents of flash memory on your switch:

```
SW1#dir
Directory of flash:/
 1  -rw-  15572992 <no date>  c2960-ik9o3s3-mz.123-1a.bin
16777216 bytes total (1204160 bytes free)
```

This shows that the flash memory on your switch contains a single Cisco IOS image named `c2960-ik9o3s3-mz.123-1a.bin`.

So, what other directories and files exist on your switch? You can easily find that using the `dir all-filesystems` Cisco IFS command:

```
SW1#dir all-filesystems
Directory of nvram:/

 27  -rw-      0  <no date>  startup-config
 28  ----      0  <no date>  private-config
```

```

29688 bytes total (29636 bytes free)
Directory of system:/

 10  drwx   0          <no date>  its
   2  dr-x   0          <no date>  memory
   1  -rw-  582        <no date>  running-config
   9  dr-x   0          <no date>  vfiles

No space information available
Directory of flash:/

   1  -rw- 15572992 <no date> c2600-ik9o3s3-mz.123-1a.bin

16777216 bytes total (1204160 bytes free)

```

Observe that you have three main directories (file systems) on your Cisco switch:

- ◆ **nvr**am: Stores the startup-config and private-config files.
- ◆ **system**: This is the RAM, which contains the running-config file.
- ◆ **flash**: This is the flash memory, which contains the Cisco IOS operating system image. This is the Cisco IOS image that the bootstrap program loads in RAM when the switch boots up.

You can list the contents of any of these main directories (file systems) individually using the `dir` and `cd` Cisco IFS commands. You can change your working directory from `flash:` to any of the main directories using the `cd` Cisco IFS command.

For example, issue the following command to list the contents of NVRAM:

```

SW1#dir nvram:
Directory of nvram:/
 27  -rw-   0          <no date>  startup-config
 28  ----   0          <no date>  private-config
29688 bytes total (29636 bytes free)

```

Now, change the directory to the RAM directory (`system`) and look at the files and directories available there:

```

SW1#cd system:
SW1#dir
Directory of system:/
 10  drwx   0          <no date>  its
   2  dr-x   0          <no date>  memory
   1  -rw-  582        <no date>  running-config
   9  dr-x   0          <no date>  vfiles

```

Observe the `running-config` file. This is the current running configuration. You can copy a configuration file from a computer host to `system:/running-config` to load a configuration on the switch from a configuration backup file on a computer host.

To load a configuration on the switch from a configuration backup file on a computer host, run the following Cisco IFS commands:

```
SW1#copy tftp://192.168.75.5/SW1-config system:/running-config
Destination filename [running-config]? <press Enter>
Accessing tftp://192.168.75.5/SW1-config...
Loading SW1-config from 192.168.75.5
[OK - 812 bytes]
812 bytes copied in 10.120 secs (80 bytes/sec)
SW1#
```

The `copy tftp://192.168.75.5/SW1-config system:/running-config` command you use here is very similar to the `copy tftp running-config` command that you used in the previous section. The only difference is that you specify the path to the source and the path to the destination when you use the Cisco IFS version of the `copy` command.

You can also use this command to copy a configuration from a configuration backup file on a computer host to the `startup-config` file in NVRAM. Here is an example:

```
SW1#copy tftp://192.168.75.5/SW1-config nvram:/startup-config
Destination filename [startup-config]? <press Enter>
Accessing tftp://192.168.75.5/SW1-config...
Loading SW1-config from 192.168.75.5
[OK - 812 bytes]
812 bytes copied in 10.120 secs (80 bytes/sec)
```

Managing Cisco Switch Authentication

You can configure several passwords for different types of access to your switch:

- ◆ **Console password:** Used for console access through the console port or through a Console Terminal Server.
- ◆ **Auxiliary password:** Used for console access through the auxiliary port using a modem.
- ◆ **VTY lines password:** Virtual type terminal (VTY) lines are used for Telnet and Secure Shell (SSH) access. They are called virtual type terminal lines because no physical terminal is connected to the switch. Instead, you connect a computer host to the network and access the switch remotely using the switch management IP address. A terminal emulation program on the computer host emulates a physical terminal (TTY).
- ◆ **Privileged password:** Used for privileged mode access. Privileged mode is an “expert” management operation mode used to run some Cisco IOS commands.

The console port and the auxiliary port are enabled by default, even if no password is specified for them. This is a security risk. It is best practice to specify at least a console password on new Cisco devices.

Privileged EXEC mode is *not* secured with a password by default on Cisco switches. By default, you can access privileged mode by executing the `enable` IOS command at the user EXEC prompt. It is best practice to configure a privileged mode password to secure access to advanced IOS commands that can adversely affect the functionality of the switch if misused.

The VTY lines (Telnet and Secure Shell) are disabled by default. You need to specify a password for the VTY lines to enable them.

To configure passwords, you need to instruct the Cisco device to prompt for authentication. You use the `login` Cisco IOS command to do this.

Console password

This section reviews how to configure a console password. To configure the console password, run the following commands:

```
SW1>enable (or en)
SW1#configure terminal (or config t)
SW1(config)#line console 0 (or line con 0)
SW1(config-line)#password my_password
SW1(config-line)#login
SW1(config-line)#exit
SW1(config)#exit
SW1#disable
SW1>
```

The first two commands (`enable` and `configure terminal`) set the IOS in privileged global configuration mode. You can now run commands that configure global switch settings, that is, settings that apply to the whole switch.

The `line console 0` command selects the console line. Cisco devices have only one console line: console 0.

The `password my_password` command sets the `my_password` password on the console access line.

Finally, you instruct the Cisco device to prompt for authentication by entering the `login` Cisco IOS command. From now on, anytime you connect to the console line, you will be prompted to provide a password. The password you provide must be `my_password`. To disable the authentication prompt, issue `line console 0` again and enter the `no login` command.

You use the `exit` commands to exit the `config-line` mode and to exit the `(config)` global configuration mode.

Telnet password

This section reviews how to configure a password for the VTY access lines. To configure a password for the VTY lines, run the following commands:

```
SW1>enable (or en)
SW1#configure terminal (or config t)
SW1(config)# line vty 0 ?
    <0-15> last line number
SW1(config-line)# line vty 0-15
SW1(config-line)# password my_telnet_password
SW1(config-line)# login
SW1(config-line)# exit
SW1(config)# exit
SW1#disable
SW1>
```

The first two commands (`enable` and `configure terminal`) set the IOS in privileged global configuration mode. You can now run commands that configure global switch settings, that is, settings that apply to the whole switch.

The `line vty 0 ?` command asks the Cisco IOS how many VTY lines are available. The response (`<0-15> last line number`) shows that 16 VTY lines exist on this particular switch. Now, you can configure a password either on one of the VTY lines (by selecting that particular line) or on all VTY lines (by selecting the whole 0–15 line range).

The `line vty 0-15` command selects the whole 0–15 VTY line range. Cisco switches have several VTY *access lines*. Older switches used to have only four VTY lines. Newer switches have as many as 1,180 VTY lines. It is best practice to query the Cisco IOS about how many VTY lines are available.

Two main reasons exist to have several VTY access lines on a Cisco switch:

- ◆ **Allowing you to connect to the switch and connect to another device from the switch:** Two VTY lines are needed in this case: one to connect into the switch and another one to connect out of the switch to another device.
- ◆ **Allowing several administrators to work on the switch:** In large networks, more than one administrator may manage the network. More than one administrator may need to connect from a remote location to the same switch using Telnet or SSH. This is typical with large core switches.

The `password my_telnet_password` command sets the `my_telnet_password` password on the VTY access lines. You can set a different password on each line. However, this is not recommended because you don't

know which VTY line you are on when you log in, so you wouldn't know which password to enter. Best practice is to select the whole range of VTY lines and to set the same password on all of them.

Finally, instruct the Cisco switch to prompt for authentication by entering the `login` Cisco IOS command. From now on, anytime you connect to any of the VTY lines, you will be prompted to provide a password. The password you provide must be `my_telnet_password`. To disable the authentication prompt, issue `line vty 0-15` again and enter the `no login` command.

Auxiliary password

If your switch has an auxiliary port, it is best practice to secure access to the auxiliary port using a password. To configure the auxiliary password, run the following commands:

```
SW1>enable (or en)
SW1#configure terminal (or config t)
SW1(config)#line auxiliary 0 (or line aux 0)
SW1(config-line)#password my_aux_password
SW1(config-line)#login
SW1(config-line)#exit
SW1(config)#exit
SW1#disable
SW1>
```

The first two commands (`enable` and `configure terminal`) set the IOS in privileged global configuration mode. You can now run commands that configure global switch settings, that is, settings that apply to the whole switch.

The `line console 0` command selects the console line. Cisco switches have only one console line: console 0.

The `password my_password` command sets the `my_password` password on the console access line.

Finally, you instruct the Cisco switch to prompt for authentication by entering the `login` Cisco IOS command. From now on, anytime you connect to the console line, you will be prompted to provide a password. The password you provide must be `my_password`. To disable the authentication prompt, issue `line console 0` again and enter the `no login` command.

You use the `exit` commands to exit the `config-line` mode and to exit the (`config`) global configuration mode.

Privileged password

It is best practice to set up a privileged password to secure access to advanced IOS commands that can adversely affect the functionality of the switch if misused. You can use two different commands to configure a privileged password:

- ◆ `enable password my_priv_password`: This command sets `my_priv_password` as the privileged password. This password is not encrypted by default. This command is supported in all IOS versions.
- ◆ `enable secret my_priv_encrypt_password`: This command sets `my_priv_encrypt_password` as the privileged password. This password is encrypted. This command is supported in newer IOS versions.

Using the `enable secret` or the `enable password` commands, you configure a privileged password that is stored locally on your switch. The only difference between the two commands is that `enable secret` stores the password on the switch in encrypted form. This is more secure than `enable password`, which stores the password on the switch in unencrypted form.

Always use the Cisco IOS help to determine whether the `enable secret` command is available. If it is, use it to set the privileged password. It is best practice to encrypt the privileged password using the `enable secret` command whenever this command is available.

You can also configure your privileged password to be stored on a *Terminal Access Controller Access Control System (TACACS)* server. This is a good option in larger networks with many switches when it is better to configure the privileged password once, centrally, on a TACACS server, instead of configuring passwords on each of your switches.

To configure an encrypted privileged mode password, run the following commands:

```
SW1>enable (or en)
SW1#configure terminal (or config t)
SW1(config)#enable secret my_priv_encrypt_password
SW1(config)#exit
SW1#disable
SW1>
```

Encrypting passwords

By default, passwords are stored in plain text in the running configuration in RAM and in the startup configuration file in NVRAM. You can view the passwords on your switch in plain text, using either the `show running-config` or the `show startup-config` command. This is a security risk. It is best practice to encrypt all passwords on your switch.

The privileged password is encrypted if you use the `enable secret` command. It is best practice to set a privileged password using the `enable secret` command to make sure that it is encrypted. So, if you use `enable secret` instead of `enable password` to set a privileged password, your privileged password is encrypted.

How about other passwords? Console, auxiliary, and Telnet (VTY) passwords are *not* encrypted, even if you use the `enable secret` command for the privileged password. You need to encrypt them using the `service password-encryption` command.

To encrypt passwords on your switch, run the following commands:

```
SW1>enable (or en)
SW1#configure terminal (or config t)
SW1(config)#service password-encryption
SW1(config)#exit
SW1#disable
SW1>
```

The first two commands (`enable` and `configure terminal`) set the IOS in privileged global configuration mode. You can now run commands that configure global switch settings, that is, settings that apply to the whole switch.

The `service password-encryption` command asks the Cisco IOS to encrypt all passwords except the privileged password that is already encrypted if the `enable secret` command was used.

Enabling Secure Shell (SSH)

Secure Shell (SSH) is similar to Telnet: It allows remote users to connect to the switch by using the switch IP address. SSH is more secure than Telnet, because it encrypts data exchanged between the remote management computer and the switch.

You should encrypt data exchanged between the remote management computer and the switch.

Consider this example: You need to change the passwords on a switch located in a remote data center. The data center is too far to walk to, so you need to connect to the switch remotely from your computer. While changing the passwords, you enter the passwords on the keyboard of your computer. If you are connected to the switch using Telnet, the data exchanged between your computer and the switch is in clear text. That data includes the new passwords that you type at your computer keyboard. Hence, passwords are sent over the Internet in clear text. This is a security risk. SSH avoids this security risk by encrypting data exchanged between your computer and the switch.

SSH uses encryption keys to encrypt the data exchanged in an SSH session. Cisco supports both SSH version 1.0 (also known as SSH v1 or SSH.1) and SSH version 2.0 (also known as SSH v2 or SSH.2). Cisco started to support SSH v2 beginning with IOS Release 12.1(19)E on some switch and router models.

If you get a `%Invalid input` error when you execute the `crypto key generate rsa` command, the Cisco IOS version on your switch does not support SSH. You may need to download the cryptographic software image from Cisco. To download the cryptographic IOS software image for your switch, log in to www.cisco.com/public/sw-center/index.shtml. (You need to register at Cisco.com before downloading. Some software export limitations may apply.)

Four packages are available for each release:

- ◆ **LAN LITE w/o crypto:** This is the Cisco IOS image without cryptographic services. This IOS image does not support SSH.
- ◆ **LAN LITE w/o crypto with web-based DEV MGR:** This is the Cisco IOS image without cryptographic services, but it includes the Cisco Device Manager GUI. This IOS image does not support SSH.
- ◆ **LAN LITE:** This is the Cisco IOS image with cryptographic services. This IOS image supports SSH.
- ◆ **LAN LITE with web-based DEV MGR:** This is the Cisco IOS image with cryptographic services and with the Cisco Device Manager GUI. This IOS image supports SSH. *

To enable SSH, you need to do the following:

- ◆ Ensure that the Cisco IOS image on your switch supports SSH.
- ◆ Ensure that you assigned a host name to your switch: SSH needs the host name to generate the encryption keys.
- ◆ Ensure that you assigned an IP domain name to your switch: SSH needs the IP domain name to generate the encryption keys.
- ◆ Create a local or TACACS administrator user and assign it a password: SSH cannot use the Telnet password that you assigned to the VTY access lines. You need to create a specific user that will be used to authenticate SSH sessions.
- ◆ Create the encryption keys.
- ◆ Configure SSH options such as session timeout and number of login retries.
- ◆ Enable SSH on the VTY lines of your switch.

Consider the following example:

```
SW1>enable (or en)
SW1#configure terminal (or config t)
```

The `enable` and `configure terminal` commands set the IOS in privileged global configuration mode. You can now run commands that configure global switch settings, that is, settings that apply to the whole switch.

```
SW1(config)#ip domain-name silange.com
```

The `ip domain-name silange.com` command sets the IP domain name for this switch. SSH needs the host name and the IP domain name to generate the encryption keys. Observe that you did not use the `hostname SW1` command here because the switch already has a host name that is different from the default Cisco “Switch” host name. If your switch was still named “Switch” and you try to run the `crypto` command, the IOS would prompt you to assign a host name to your switch.

```
SW1(config)#username Admin007 password Admin007Passwrd
SW1(config)#line vty 0 15
SW1(config-line)#login local
SW1(config-line)#exit
```

The `line vty 0 15` and `login local` commands enable local authentication on VTY lines 0 to 15.



You can also use TACACS authentication by executing `login tacacs` instead of `login local`. However, you would need to have a TACACS server configured, which is beyond the scope of CCNA.

Next, you exit the line configuration mode by executing `exit` at the (config-line) prompt.

Now, you create the SSH encryption keys for your switch. You execute the `crypto key generate rsa` command. This command generates public and private encryption keys based on the RSA encryption algorithm. It uses the `hostname.domain-name` combination of the switch to generate the keys. The `crypto` command also uses a modulus number to generate the keys. You can choose the number of bits used for the modulus number: between 360 and 2048. The higher the number of bits in the modulus number, the harder it is for a hacker to break the keys. So, it is better to have more bits in the modulus number. However, the more bits in the modulus number, the longer it takes to generate it. Modulo 1024 is used in this example, a good compromise between security and performance:

```
SW1(config)#crypto key generate rsa
The name for the keys will be: SW1.silange.com
Choose the size of the key modulus in the range of 360 to 2048 for your General
Purpose Keys. Choosing a key modulus greater than 512 may take a few minutes.
```

```
How many bits in the modulus [512]: 1024
% Generating 1024 bit RSA keys ...[OK]
*Mar  1 00:14:40.071: %SSH-5-ENABLED: SSH 1.5 has been enabled
```

Next, you look at the current SSH configuration on the switch using the `show ip ssh` command. Observe that you prefix this command with the `do` keyword because you cannot run `show ip ssh` in configuration mode. Recall that the `do` prefix allows you to run a privileged EXEC command from the privileged EXEC configuration prompt:

```
SW1(config)#do show ip ssh
  SSH Enabled - version 2.0
  Authentication timeout: 120 secs; Authentication retries: 3

SW1(config)#ip ssh time-out 60
SW1(config)#ip ssh authentication-retries 2
```

The `ip ssh time-out 60` and `ip ssh authentication-retries 2` commands configure a 60-second SSH session timeout and configure SSH to allow two login retries. Observe that the default timeout is 120 seconds and the default number of retries is 2, as shown by the `show ip ssh` command in the previous step.

Finally, you need to enable SSH as an input transport protocol on the VTY access lines. The `line vty 0 15` and `transport input ssh telnet` commands accomplish this.



Note that you enabled both SSH and Telnet as input transport protocols on the VTY access lines. If you removed the `telnet` keyword from the `transport input` command, only SSH would be enabled:

```
SW1(config)#line vty 0 15
SW1(config-line)#transport input ssh telnet
SW1(config-line)#exit
```

The next two commands show the current SSH configuration on the switch (`show ip ssh`) and the status of current SSH connections if any (`show ssh`):

```
SW1(config)#do show ip ssh
  SSH Enabled - version 2.0
  Authentication timeout: 60 secs; Authentication retries: 2
SW1(config)#do show ssh
  %No SSH server connections running.
SW1(config)#exit
SW1#disable
SW1>
```

Recovering switch passwords

What happens if you forget your passwords? Suppose, for example, that you start a new job working at a company as the network administrator. The network consultant who set up the network infrastructure left without leaving

the passwords for some switches you have in your network. This happens quite often, unfortunately. Theoretically, you would have to wipe out the configuration of the switches and start from scratch. This is not productive.

The Cisco IOS provides a mechanism to recover passwords in case you lose them or you do not remember them.

You need to have physical access to the switch to connect to its console port. This is a good thing, because you don't want anyone else connecting from a remote location through Telnet to be able to reset the passwords on your switches.

The key point of the password recovery process is to boot up the Cisco switch ignoring its current startup configuration, which contains the current passwords. To do this, you need to interrupt the boot process of the switch before the startup configuration is loaded from NVRAM to RAM.

Reviewing the switch manual boot process

You have three ways to interrupt the boot process on a Cisco switch.

Interrupting the boot process with the Mode button

Follow these steps to interrupt the boot process of a Cisco switch using the Mode button on the front panel of the switch:

- 1. Reboot the switch using the `reload` command in privileged EXEC mode.**
- 2. Press the Mode button while the System LED (SYST) is flashing green within the first 15 seconds of the switch boot process.**
- 3. Continue to press the Mode button until the System LED (SYST) flashes briefly amber and then turns solid green.**
- 4. Release the Mode button when the System LED (SYST) becomes solid green.**

Enabling the manual boot option

Follow these steps to interrupt the boot process of a Cisco switch using the manual boot option:

- 1. Enable the manual boot option.**
- 2. Reboot the switch using the `reload` command in privileged EXEC mode.**

The following options allow you to control the manual boot option on a Cisco switch:

- ◆ To enable the manual boot option, run `boot manual` at the privileged global configuration prompt.

- ◆ To disable the manual boot option, run `no boot manual` at the privileged global configuration prompt.

The manual boot option is off by default: The switch boots up automatically without showing the `switch: manual boot` prompt.

Interrupting the boot process with Ctrl+Break

Follow these steps to interrupt the boot process of a Cisco switch using the Ctrl+Break key combination:

- 1. Enable the boot process break option.**
- 2. Reboot the switch using the `reload` command in privileged EXEC mode.**
- 3. Press Ctrl and Break simultaneously while the System LED (SYST) is flashing green within the first 15 seconds of the switch boot process.**

You have the following boot process break options on a Cisco switch:

- ◆ To enable the boot process break option, run `boot enable-break` at the privileged global configuration prompt.
- ◆ To disable the boot process break option, run `no boot enable-break` at the privileged global configuration prompt.

The boot process break option is disabled by default.

Reviewing the switch manual boot prompt

Regardless of how you interrupt the boot process, you end up at the `switch: manual boot` prompt. The manual boot prompt allows you to control the boot process of the switch.

The bootstrap program displays a message informing you that you can enter the `boot` command to continue booting up the switch.



You can use the `help (?)` command at the `switch: boot` prompt to see the manual boot commands and options available.

The IOS has not yet loaded the startup configuration at this point. In fact, the bootstrap program has not even loaded the IOS at this point.

You are now ready to start the password recovery process. This process varies depending on whether the password recovery feature is enabled on the switch.

Introducing the password recovery option

A switch administrator can disable the password recovery feature for security reasons. If the password recovery feature is disabled, the password can

still be reset, but the switch configuration is reset as well. This prevents an unauthorized user from resetting the password and gaining access to a configured switch.

You control the password recovery feature availability using the `service password-recovery` IOS command. The `service password-recovery` IOS command needs to be run in privileged global configuration mode. You have the following options to control the password recovery feature on a Cisco switch:

- ◆ To enable password recovery, run `service password-recovery` at the privileged global configuration prompt.
- ◆ To disable password recovery, run `no service password-recovery` at the privileged global configuration prompt.

Password recovery process with password recovery enabled

The password recovery feature is enabled if you see the message `The password-recovery mechanism is enabled during the boot process`. Follow these steps to recover the password on a switch with the password recovery option enabled:

- 1. Interrupt the normal boot process using one of the methods previously described.**
- 2. Manually initialize the flash file system.**
- 3. Hide the startup configuration file to prevent the IOS from loading it.**
- 4. Boot the switch manually and wait for the IOS to finish loading in RAM.**
- 5. Unhide the startup configuration file.**
- 6. Manually load the startup configuration file from NVRAM to RAM.**
- 7. Reset the password.**
- 8. Save the running configuration over the startup configuration.**
- 9. Revert the boot process to its default options.**
- 10. Reboot the switch.**

Password recovery process with password recovery disabled

The password recovery feature is disabled if you see the message `The password-recovery mechanism is disabled during the boot process`. In this case, you need to reset the whole configuration of the switch to be able to reset the password. This basically resets the switch to its default out-of-factory configuration. You need to ensure that you do not lose any valuable information about your network configuration when you do this. It is best practice to save the configuration of the switch before resetting it.

Follow these steps to recover the password on a switch with the password recovery option disabled:

- 1. Interrupt the normal boot process using one of the methods described previously.**
- 2. Accept to reset the switch to its default configuration.**
- 3. Boot the switch manually and wait for the IOS to finish loading in RAM.**
- 4. Reset the password.**
- 5. Save the running configuration over the startup configuration.**
- 6. Revert the boot process to its default options.**
- 7. Reboot the switch.**

Here is an example:

```
!Lines starting with "!" in this example are comments that
!explain certain commands within the example. Cisco IOS
!does not interpret lines starting with the "!" character.
!The IOS interprets these lines as comments, not as commands.
!The first few lines below restart the switch using the
!reload IOS command after the manual boot option was enabled.

SW1>en
SW1#reload
Proceed with reload? [confirm]

00:01:35: %SYS-5-RELOAD: Reload requested by console. Reload Reason: Reload
      command.
!The switch is now starting up. Observe the message in bold
!below, showing that password recovery is enabled.

Using driver version 1 for media type 1
Base ethernet MAC Address: 00:25:b4:10:58:80
Xmodem file system is available.
The password-recovery mechanism is disabled.

The password-recovery mechanism has been triggered, but
is currently disabled. Access to the boot loader prompt
through the password-recovery mechanism is disallowed at
this point. However, if you agree to let the system be
reset back to the default system configuration, access
to the boot loader prompt can still be allowed.

!here the IOS prompts you to reset the configuration of the
!switch to be able to reset its password.

Would you like to reset the system back to the default configuration (y/n)? y

!next line manually boots the switch
switch:boot
```

```

Loading "flash:c2960-lanlite-mz.122-37.EY/c2960-lanlite-mz.122-37.EY.bin"...@@@
@@@@@@@@@@@@@@@@@@@@

[... boot output cut ...]

SW1>en
SW1#configure terminal (or config t)

!next line sets the new password
SW1(config)#enable secret my_priv_encrypt_password

SW1(config)#exit

!next line replaces the startup-config file
!   in NVRAM (flash:config.text) with the current
!   running-config file in RAM (system:running-config)

SW1#copy running-config startup-config

!next line reboots the switch
SW1#reload

System configuration has been modified. Save? [yes/no]: yes
Proceed with reload? [confirm] <press Enter>

00:15:05: %SYS-5-RELOAD: Reload requested by console. Reload Reason: Reload
command.

[... reboot output cut ...]

```

Observe that the password recovery feature is disabled on this switch. You are prompted to reset the switch to its default out-of-factory configuration to reset the password. You answer *y* to the prompt.

Next, manually boot the switch and wait for the IOS to finish loading in RAM. You know that the IOS is finished loading when it displays the *SW1>* user EXEC prompt.

Now it's time to set the new password. This basically resets the password to a new one.

Finally, replace the startup-config file in NVRAM (*flash:config.text*) with the current running-config file in RAM (*system:running-config*) and reboot the switch.

Upon reboot, the switch is reset to its default out-of-factory configuration. The password of the switch is set to *my_priv_encrypt_password*. This is the password that you just set with the *enable secret* command during the password recovery procedure.

Prep Test

- 1 Describe the Cisco IOS (Internetwork Operating System).**
 - A ☐ Cisco's proprietary switch and router graphical user interface
 - B ☐ Cisco's proprietary switch and router operating system
 - C ☐ Cisco's proprietary switch and router console facility
 - D ☐ All of the above
- 2 What are some of the graphical user interfaces that you can use to manage a Cisco switch?**
 - A ☐ Cisco Network Assistant (CNA) and Cisco Device Manager
 - B ☐ Cisco Switch Windows Control Panel
 - C ☐ Cisco Router and Security Device Manager (SDM)
 - D ☐ All of the above
- 3 What is the purpose of the boot image (Rx-boot) microprogram, accessible using the switch: manual boot prompt?**
 - A ☐ Control the log level during the boot process
 - B ☐ Control the power supply level during the boot process
 - C ☐ Control the boot process
 - D ☐ All of the above
- 4 When does the POST (power-on self test) run?**
 - A ☐ Immediately after the Cisco IOS loads on a switch or router
 - B ☐ Immediately after the startup configuration loads on a switch or router
 - C ☐ Immediately after the startup configuration loads on a switch or router
 - D ☐ Immediately after a Cisco switch or router is powered up
- 5 What is the main purpose of the bootstrap program, also known as the boot loader?**
 - A ☐ To bring up a Cisco switch or router by loading the startup configuration from flash memory to RAM
 - B ☐ To bring up a Cisco switch or router by loading the Cisco IOS from flash memory to RAM
 - C ☐ To bring up a Cisco switch or router by loading the running configuration from flash memory to RAM
 - D ☐ To bring up a Cisco switch or router by loading the Cisco Device Manager from flash memory to RAM

- 6** At what layer is it best practice to use entry-level switches, such as the Cisco Catalyst 2960 or the Cisco Catalyst 3560?
- A ☐ Access and distribution layer
 - B ☐ Distribution and core layer
 - C ☐ Core layer
 - D ☐ Access and core layer
- 7** What does the flash memory on a Cisco switch store?
- A ☐ The startup configuration of a Cisco switch or router
 - B ☐ The image file of the Cisco IOS operating system
 - C ☐ The Cisco Device Manager software application program
 - D ☐ All of the above
- 8** Where does a Cisco switch load its running configuration?
- A ☐ From the running configuration file during startup
 - B ☐ From Cisco's Web site during startup
 - C ☐ From the startup configuration file during startup
 - D ☐ All of the above
- 9** When does the Cisco IOS command-line interface operate in setup mode?
- A ☐ To initially configure the ROM Monitor (ROMMON) microprogram
 - B ☐ To initially configure the Device Manager software application
 - C ☐ To initially configure the IOS
 - D ☐ To initially configure the switch or router
- 10** When does the Cisco IOS command-line interface operate in privileged configuration mode?
- A ☐ To configure the ROM Monitor (ROMMON) microprogram
 - B ☐ To configure the Device Manager software application
 - C ☐ To configure the IOS
 - D ☐ To configure the switch or router

Answers

- 1 B. Cisco's proprietary switch and router operating system. The Cisco IOS (Internetworking Operating System) is Cisco's proprietary switch and router operating system. Review the chapter introduction.
- 2 A. Cisco Network Assistant (CNA) and Cisco Device Manager. Some of the Graphical User Interfaces that you can use to manage a Cisco switch are the Cisco Network Assistant (CNA) and the Cisco Device Manager. Review the chapter introduction.
- 3 C. Maintain, test and troubleshoot the configuration stored in ROM and in the flash memory of a Cisco device. The purpose of the Boot image (Rx-boot) micro-program, accessible using the switch: manual boot prompt, is to maintain, test and troubleshoot the configuration stored in ROM and in the Flash memory of a Cisco device. Check "*Cisco switch startup process and Managing Cisco Switch Configuration.*"
- 4 D. Immediately after a Cisco switch or router is powered up. The Power-On-Self-Test runs immediately after a Cisco switch is powered up. Check "*Cisco switch startup process.*"
- 5 B. Bring up a Cisco switch or router by loading the Cisco IOS from flash memory to RAM. The main purpose of the bootstrap program, also known as the boot loader, is to bring up a Cisco switch by loading the Cisco IOS from Flash memory to RAM. See "*Cisco switch startup process.*"
- 6 A. Access and distribution layer. It is best practice to use entry-level switches, such as the Cisco Catalyst 2960 or the Cisco Catalyst 3560, at the access and distribution layers. Review "*Best Practice For Using Cisco Switches.*"
- 7 D. All of the above. The flash memory on a Cisco switch stores the startup configuration of a Cisco switch or router, the image file of the Cisco IOS operating system, and the Cisco Device Manager software application program. See "*Cisco switch startup process.*"
- 8 C. Is loaded from the startup configuration file during startup. The running configuration of a Cisco switch is loaded from the startup configuration file during startup. Review "*Cisco switch startup process.*"
- 9 D. Initially configure the switch or router. The Cisco IOS command line interface operates in setup mode to initially configure the switch or router. Review "*Initial switch configuration.*"
- 10 D. Configure the switch or router. The Cisco IOS command line interface operates in privileged configuration mode to configure the switch. Review "*Managing Cisco Switch Authentication.*"

Chapter 3: Controlling Network Traffic with Cisco Switches

Exam Objectives

- ✓ Describing how MAC addresses are handled in local and remote transmissions
- ✓ Describing Layer 2 switch modes
- ✓ Setting the duplex mode of a Layer 2 switch
- ✓ Reviewing MAC address table thrashing and broadcast storms and seeing how STP fixes these two issues

Read this chapter to find out how Cisco Layer 2 switches control and optimize traffic in a local-area network (LAN).

Sending to MAC Addresses in Remote Networks

In the following sections, you discover how a sending host device interacts with a Layer 2 switch and with a gateway (a router) to send frames to a target device located in a remote network. Most data transfers involve sending and receiving devices that are not located in the same LAN. Data frames need to be sent to the LAN gateway first. Next, the LAN gateway routes the data frames wrapped in packets through a wide-area network (WAN) to the destination LAN.

Sending frames within the LAN

A host determines the MAC address of a target device in the local network by first looking into its Address Resolution Protocol (ARP) table, searching for the MAC address that corresponds to the IP address of the target device:

- ◆ If the sending host finds an entry in its ARP table corresponding to the IP address of the target device, it simply writes that MAC address in the destination MAC address field in the Ethernet frame.
- ◆ If the sending host does not find an entry in its ARP table corresponding to the IP address of the target device, it broadcasts an ARP request. In other words, it sends an ARP request packaged in a frame with the destination MAC address set to the broadcast address: FF-FF-FF-FF-FF-FF. All devices in the LAN see the ARP request.

The target device eventually responds with its MAC address. The sending host saves the target's MAC address in its ARP table. The next time the sending host will find the target's MAC address in its ARP table.

Figure 3-1 illustrates how Alex, the sending host, interacts with the Layer 2 switch and with Claire, the target host, to find Claire's MAC address and to send her frames:

1. Alex looks up his ARP table to see whether he already has Claire's MAC address. He doesn't, so he needs to broadcast an ARP request.
2. Alex broadcasts an ARP request.

The Ethernet frame contains Alex's MAC address in the S-MAC (source) field and the standard broadcast MAC address (FF-FF-FF-FF-FF-FF) in the D-MAC (destination) field.

3. The switch receives this frame, and upon inspection of the D-MAC field, it floods the frame out on all ports except the one through which it came in.
4. The frame reaches both John and Claire. John discards it, because it's not for him. Claire recognizes her own IP address in the ARP request, so she responds with her MAC address.

The Ethernet frame contains Claire's MAC address in the S-MAC field and Alex's MAC address in the D-MAC field.

5. Alex receives Claire's frame that basically answers his ARP request.
Alex now has Claire's IP address and her MAC address. He saves them both in his ARP table. So, the next time he needs to send something to Claire, he doesn't need to broadcast an ARP request again.
6. Now, Alex can send frames directly to Claire because he knows her MAC address. He reads Claire's MAC address from his ARP table and builds the Ethernet frame, putting his own MAC address in the S-MAC field and Claire's MAC address in the D-MAC field.

Sending frames to a remote network

What happens when the target device is *not* in the local network? It starts the same way as when the device is in the local network. The sending host first looks up its Address Resolution Protocol (ARP) table, searching for the MAC address that corresponds to the IP address of the target device:

- ◆ If the sending host finds an entry in its ARP table corresponding to the IP address of the target device, it simply writes that MAC address in the destination MAC address field in the Ethernet frame.




Figure 3-1:
Sending frames within the LAN.

- ◆ If the destination device is located in a remote network, the host will never find the MAC address of the target host in its ARP table, but rather, it will find the MAC address of the LAN gateway.

The LAN gateway is the router device that routes data frames wrapped in IP packets from the LAN to a WAN and from the WAN back into the LAN. The sending host will never find the MAC address of the target host in its ARP table, but rather, it will eventually find the MAC address of the LAN gateway.

If the sending host does not find an entry in its ARP table corresponding to the IP address of the target device, it broadcasts an ARP request. In other words, it sends an ARP request packaged in a frame with the destination MAC address set to the broadcast address: FF-FF-FF-FF-FF-FF. All devices in the LAN see the ARP request, but the target device is not in the local network. So, it does not see the ARP request. However, the LAN gateway (that is, the router), like all devices in the LAN, does see the ARP request. It recognizes that the ARP request is for an IP address that is located in a remote network.

The router responds to the sending host with its own MAC address. In other words, the router tells the sending host: Anytime you need to send frames to this target device, send the frames to me, and I'll route the frames to the remote LAN where this target device is located. The sending host saves the gateway's MAC address in its ARP table. The next time the sending host needs to send to this same target device in a remote network, it will find the MAC address of the gateway in its ARP table. Hence, the sending host will send the frame to the router, directly, without broadcasting an ARP request.

Figure 3-2 illustrates how Alex, the sending host, interacts with the Layer 2 switch and with the gateway router, to find the gateway's MAC address and to send frames to Mitch, the target host, in a remote network:

1. Alex looks up his ARP table to see whether he already has Monica's MAC address. He doesn't, so he needs to broadcast an ARP request.
2. Alex broadcasts an ARP request.

The Ethernet frame contains Alex's MAC address in the S-MAC (source) field and the standard broadcast MAC address (FF-FF-FF-FF-FF-FF) in the D-MAC (destination) field.

3. The switch receives this frame, and upon inspection of the D-MAC field, it floods the frame out on all ports except the one through which it came in.




Figure 3-2:
Sending frames to a remote network.

4. The frame reaches both Claire and the router that links this LAN to a WAN:
 - Claire discards the frame, because it's not for her.
 - The router recognizes that the IP address requested in the ARP is in a different network, so it sends back its MAC address. This enables Alex to send the frame to the router, and the router handles it from there.

The Ethernet frame contains the router's MAC address in the S-MAC field and Alex's MAC address in the D-MAC field.

5. Alex receives the router's frame that basically answers his ARP request. Alex now saves Monica's IP address and the router's MAC address in his ARP table.

The next time he needs to send something to Monica, he doesn't need to broadcast an ARP request again. He simply sends the frame to the router, and the router handles it from there.

6. Now, Alex can send frames to Monica without broadcasting an ARP request. He looks up Monica's IP address in his ARP table and builds the Ethernet frame, putting his own MAC address in the S-MAC field and the router's MAC address in the D-MAC field. Thus, the frame bound for Monica goes directly to the router now. The router sends the frame over the WAN to the gateway of the remote network where Monica is located.
7. The remote router encapsulates the IP packet in a new Ethernet frame that contains his MAC address in the S-MAC field and Monica's MAC address in the D-MAC field.

Ethernet frames are always bound to a given LAN. The Ethernet frame in Alex's network involved Alex and his local gateway. The Ethernet frame in Monica's network involved Monica and the remote router (Monica's gateway). Ethernet frames always identify the next hop in a LAN: The source and destination MAC addresses change at each hop. In contrast, IP packets identify the sender and the receiver of the data payload. The source and destination IP addresses remain the same throughout the journey of a packet through networks linking the sender and the receiver.

Sending frames within a LAN: The destination MAC address of the frame is set to the MAC address of the target host device.

Sending frames to a remote network: The destination MAC address of the frame is set to the MAC address of the LAN gateway. The LAN gateway is a router that links the LAN to a WAN.

The destination MAC address of a Layer 2 data-link (Ethernet) frame is never set to the MAC address of a Layer 2 switch. A Layer 2 switch is merely a conduit that provides an optimized point-to-point link between the sending device and the target device.

The sending device is usually a host in the LAN. Hence, the source MAC address in the Ethernet frame is the MAC address of sending host device.

The target device is either another host in the LAN or the LAN gateway. Hence, the destination MAC address in the Ethernet frame is either the MAC address of a target host device within the LAN or the MAC address of the LAN gateway.

Deciding the Fate of Frames

In the following sections, you find out about switching modes and how they control the fate of each frame: filtering, forwarding, or flooding the frame. One of the most important functions of a Layer 2 switch is to decide what to do with a frame it receives. The switch either

- ◆ Filters (that is, discards) the frame
- ◆ Forwards the frame out on a specific port
- ◆ Sends the frame out on all ports except the port through which it came in

The following sections describe how the switch decides the fate of a frame.

Switching modes

Layer 2 switches support three switching modes: Store-and-forward, Cut-through, and Fragment-free

Store-and-forward

The *store-and-forward* switching mode follows this process:

1. The switch stores the entire frame in the switch buffer (temporary) memory.
2. FCS (frame check sequence) is run to ensure that the frame is valid.
3. The switch inspects the source MAC address and destination MAC address of the frame.
4. The switch saves the source MAC address along with the incoming port in the MAC address table.
5. The switch looks up the MAC address table for the destination MAC address:
 - If it finds the destination MAC address in the MAC address table, it *forwards* the frame only on the outgoing port registered in the MAC address table.

- If it does not find the destination MAC address in the MAC address table, it *floods* the frame on all ports, except the port where the frame came in.

Cut-through

The *cut-through* switching mode is faster than the store-and-forward mode. In cut-through mode, the switch does not store the whole frame in its buffer memory. Instead, it processes the frame as soon as it receives the first 6 bytes of the frame:

1. The switch inspects the first 6 bytes of the Ethernet frame, looking at the destination MAC address.
2. The switch searches an entry for the destination MAC address in its MAC address table.
3. Assuming that the destination MAC address has already been catalogued in the MAC address table, the switch forwards the bytes of the frame right away on the corresponding outgoing port, even before the rest of the frame has been completely received.

The switch does *not* calculate the FCS before it forwards the frame. The cut-through mode speeds frame forwarding considerably, but at a price: Because the FCS is not calculated, both valid and invalid frames are forwarded. Normally, the receiving device can handle this. However, in networks with high traffic volume, by forwarding both valid and invalid frames, the switch unnecessarily overloads the network segment between the switch and the receiving device, which is not good. This is why some switch models still calculate the FCS and keep track of invalid frames. When the number of invalid frames passes a certain threshold, the switch reverts to store-and-forward mode.

Fragment-free

The *fragment-free* switching mode combines the advantages of both the store-and-forward mode and the cut-through mode. Valid Ethernet frames are normally at least 64 bytes in length. Fragment-free considers any frame that is at least 64 bytes to be valid.

A switch in fragment-free mode acts as follows:

1. The switch stores the first 64 bytes of a frame in the switch buffer.
2. It inspects the source MAC address and destination MAC address of the frame.
3. It saves the source MAC address along with the incoming port in the MAC address table.

4. The switch looks up the MAC address table for the destination MAC address:
 - If it finds the destination MAC address in the MAC address table, it *forwards* the frame only on the outgoing port registered in the MAC address table.
 - If it does not find the destination MAC address in the MAC address table, it *floods* the frame on all ports, except the port where the frame came in.

Fragment-free starts forwarding the frame as soon as its first 64 bytes have been received, as opposed to store-and-forward, which stores the whole frame before forwarding it. Hence, fragment-free is faster than store-and-forward. However, fragment-free is slower than cut-through, because cut-through forwards the frame as soon as the first 6 bytes have been received.

A switch in fragment-free mode discards any frame that is shorter than 64 bytes. Because frames shorter than 64 bytes are typically invalid frames, this lowers the risk of forwarding invalid frames. This is similar to the store-and-forward mode, which does not forward invalid frames.

Fragment-free mode does not calculate the FCS. This is similar to the cut-through mode. As in cut-through mode, some switches calculate the FCS anyway and revert to store-and-forward when the number of invalid frames passes a certain threshold.

Fragment-free is a good compromise between store-and-forward and cut-through switching modes:

- ◆ Store-and-forward is reliable, forwarding only valid frames, but slower because it always calculates the FCS.
- ◆ Cut-through is faster because it doesn't store the whole Ethernet frame before forwarding it, but is potentially unreliable because it may forward invalid frames.
- ◆ Fragment-free relies on the minimum valid Ethernet frame size to determine whether a frame is valid without calculating the FCS:
 - Because most invalid frames are smaller than 64 bytes, fragment-free is quite reliable.
 - Because fragment-free does not store more than 64 bytes of the frame in the switch buffer, it is quite efficient as well.

Switching in Half-Duplex and Full-Duplex Modes

The first few iterations of the Ethernet standard supported only half-duplex transmissions. Current Ethernet standards support both half-duplex and full-duplex transmission modes.

Reviewing half-duplex Ethernet

Half-duplex Ethernet uses a single pair of wires for both sending and receiving frames. Frame collisions occur and need to be mitigated with carrier sense multiple access collision detect (CSMA/CD), so half-duplex throughput is usually 30–40 percent of theoretical throughput. This means that in a typical 100BASE-T Ethernet network, although theoretical throughput is 100 Mbps, half-duplex only provides 40 to 50 Mbps. All hubs, bridges, and switches support half-duplex transmission mode. However, half-duplex Ethernet should only be used with older devices that do not support full-duplex Ethernet. In all other cases, it is best to use full-duplex Ethernet.

Reviewing full-duplex Ethernet

Full-duplex Ethernet uses two pairs of wires, or even four pairs of wires, to send and receive frames at the same time on several connection paths using one connection path per pair. Full-duplex twisted-pair connections provide a point-to-point connection between the transmitter on the sending host and the receiver on the receiving host. This eliminates frame collisions and speeds data transfer. Using full-duplex twisted-pair connections, you can theoretically achieve full throughput in both directions. Considering the 100BASE-T example, using full-duplex twisted-pair, you should be able to get 100 Mbps sending and 100 Mbps receiving, for a total bandwidth of 200 Mbps. However, in reality, you will likely get less than 100 Mbps due to potential electrical noise, crosstalk, and other factors that may distort the signal. Most Layer 2 switches support full-duplex transmission mode.

Duplex mode best practice

It is best practice to use full-duplex transmission mode because it improves network throughput and reliability by eliminating frame collisions. Particularly, always use full-duplex transmission mode to connect hosts to switches, and interconnect switches using full-duplex twisted-pair.

Configuring port duplex mode on a Cisco switch

You configure the transmission mode on a Cisco switch using the `duplex` IOS command. For example, to set the transmission mode of a switch interface in full-duplex mode, execute the following command:

```
SW1(config-if)#duplex full
```

To set the transmission mode of a switch interface in half-duplex mode, execute the following command:

```
SW1(config-if)#duplex half
```

By default, all switch ports are set to autonegotiate the duplex mode with the device that connects to the port. To manually set the duplex mode to autonegotiate, execute the following command:

```
SW1(config-if)#duplex auto
```

Configuring port speed on a Cisco switch

You can follow the `duplex` IOS command with the `speed` IOS command to set the speed of the interface. The `speed` command can configure the interface to run at various speeds: 10, 100, 1000, and auto.

Some interfaces are limited in the speeds they support. For example, a Fast Ethernet switch interface supports speeds only up to 100 Mbps. In this case, you can only set the Fast Ethernet interface speed to 10, 100, or auto.

Selecting a switch port

Before you run any of the `duplex` and `speed` commands, you need to select the port you want to configure:

1. Transition the IOS to privileged EXEC mode:

```
SW1>enable (or en)
```

2. Enable the Cisco switch configuration terminal:

```
SW1#configure terminal (or config t or conf t)
```

3. Select a port to work with using the `interface <interface#>` command.

Again, you can use the short form: `int <interface#>`. For example, to work with the first Fast Ethernet interface on your switch (0/0), you can run the following:

```
SW1(config)#int fastethernet 0/0 (or int fa 0/0)
SW1(config-if)#
```

Avoiding Loops with Spanning Tree Protocol (STP)

This section introduces the Spanning Tree Protocol (STP) that is used by Layer 2 switches to avoid loops in topologies with redundant interswitch links. The next chapter is devoted to a more detailed discussion of STP.

Whenever two Layer 2 switches are interconnected with redundant links, transmission loops may occur. Frames are continuously retransmitted on alternate links between switches, effectively causing frames to loop, or bounce, between switches. This can cause broadcast storms and MAC address table thrashing.

A broadcast frame that bounces forever between switches interconnected with redundant links causes a broadcast storm. Because each switch floods broadcast frames out on all ports, they continuously send the frame to each other on alternate links. Data-link (Ethernet) frames do not expire, so they can bounce forever between Layer 2 switches. The Spanning Tree Protocol helps avoiding these broadcast storms.

MAC address tables can be thrashed when a MAC address is registered for more than one switch port. This may happen when the same Ethernet frame is received through more than one port on a given switch. It is common to get thrashed MAC address tables when loops exist in a network because frames are sent between switches on alternate ports. Hence, the same source MAC address in the frame is registered for each alternate port.

So, how exactly does STP eliminate loops in LANs? STP basically monitors the network and catalogs each link, particularly redundant links. Next, STP disables redundant links, setting up preferred, optimized links between switches. The preferred, optimized links are used under normal circumstances. Should any of the preferred links fail, one of the nonpreferred redundant links is enabled and used instead.

STP assigns a root bridge that acts sort of like the decision maker in the network. The term *root bridge* typically refers to a Layer 2 switch. It can also really be a bridge, but more commonly it is a switch. The root bridge decides which routes are preferred and which routes are nonpreferred.

The root bridge interacts with *nonroot bridges*: other switches in the LAN. STP is enabled on all switches in the LAN. The root bridge and the nonroot bridges have specific roles within STP.

Switch ports are categorized by whether they forward frames and whether they are the endpoints of a preferred, optimized link within the LAN.

STP operation, management, and optimization are covered in Book III, Chapter 4.

Prep Test

- 1** When sending data frames to a host in a remote network, the destination MAC address in the data frame is set to what?
 - A ☐ The IP address of the destination host
 - B ☐ The IP address of the LAN gateway
 - C ☐ The MAC address of the destination host
 - D ☐ The MAC address of the LAN gateway
- 2** When sending data frames to a host within the LAN, the destination MAC address in the data frame is set to what?
 - A ☐ The IP address of the destination host
 - B ☐ The IP address of the LAN gateway
 - C ☐ The MAC address of the destination host
 - D ☐ The MAC address of the LAN gateway
- 3** Hosts use the _____ to determine the MAC address of a target device.
 - A ☐ UDP (User Datagram Protocol)
 - B ☐ ARP (Address Resolution Protocol)
 - C ☐ DHCP (Dynamic Host Configuration Protocol)
 - D ☐ TCP (Transport Control Protocol)
- 4** Hosts need to broadcast _____ whenever their _____ does not contain the MAC address of a target device.
 - A ☐ an ARP request, ARP table
 - B ☐ a UDP request, UDP table
 - C ☐ a DHCP request, DHCP table
 - D ☐ a TCP request, TCP table
- 5** Which of the following describes the store-and-forward switching mode?
 - A ☐ Stores only the first 64 bytes of the data frame before processing it
 - B ☐ Does not store any bytes in the switch buffer before processing the data frame
 - C ☐ Stores only the first 6 bytes of the data frame before processing it
 - D ☐ Stores the entire data frame in the switch buffer before processing it

- 6 Which of the following describes the cut-through switching mode?**
- A ☐ Stores only the first 64 bytes of the data frame before processing it
 - B ☐ Does not store any bytes in the switch buffer before processing the data frame
 - C ☐ Stores only the first 6 bytes of the data frame before processing it
 - D ☐ Stores the entire data frame in the switch buffer before processing it
- 7 Which of the following describes the fragment-free switching mode?**
- A ☐ Stores only the first 64 bytes of the data frame before processing it
 - B ☐ Does not store any bytes in the switch buffer before processing the data frame
 - C ☐ Stores only the first 6 bytes of the data frame before processing it
 - D ☐ Stores the entire data frame in the switch buffer before processing it
- 8 How does the store-and-forward switching mode ensure that the data frame is valid?**
- A ☐ By reading the FCS from the frame
 - B ☐ By calculating the FCS and verifying it against the FCS stored in the frame
 - C ☐ By discarding any data frame that is smaller than 64 bytes
 - D ☐ By discarding any data frame that is larger than 64 bytes
- 9 How does the fragment-free switching mode ensure that the data frame is valid?**
- A ☐ By reading the FCS from the frame
 - B ☐ By calculating the FCS and verifying it against the FCS stored in the frame
 - C ☐ By discarding any data frame that is smaller than 64 bytes
 - D ☐ By discarding any data frame that is larger than 64 bytes
- 10 Although you can manually set the duplex mode and the speed of switch ports on a Cisco switch, by default, how are both the duplex mode and the speed set?**
- A ☐ To the lowest value
 - B ☐ To the highest value
 - C ☐ Autonegotiated based on MAC address table entries
 - D ☐ Autonegotiated with the device connecting in the switch port

Answers

- 1** D. When sending data frames to a host in a remote network, the destination MAC address in the data frame is set to the MAC address of the LAN gateway. Refer to *"Sending frames to a remote network."*
- 2** C. When sending data frames to a host in a remote network, the destination MAC address in the data frame is set to the MAC address of the destination host. Review *"Sending frames within the LAN."*
- 3** B. ARP (Address Resolution Protocol). Hosts use ARP to determine the MAC address of a target device. Check out *"Sending frames within the LAN."*
- 4** A. Hosts need to broadcast an ARP request whenever their ARP table does not contain the MAC address of a target device. If the ARP table does not contain the MAC address of the destination device, the sending host needs to broadcast an ARP request to find out the MAC address of the target device. Read *"Sending frames within the LAN."*
- 5** D. The store-and-forward switching mode stores the entire data frame in the switch buffer before processing it. Peruse *"Switching modes."*
- 6** C. The cut-through switching mode stores only the first 6 bytes of the data frame before processing it. Review *"Switching modes."*
- 7** A. The fragment-free switching mode stores only the first 64 bytes of the data frame before processing it. Check out *"Switching modes."*
- 8** B. The store-and-forward switching mode ensures that the data frame is valid by calculating the FCS and verifying against the FCS stored in the frame. Read *"Switching modes."*
- 9** C. The fragment-free switching mode ensures that the data frame is valid by discarding any data frame that is smaller than 64 bytes. See *"Switching modes."*
- 10** D. The duplex mode and the speed of switch ports on a Cisco switch are by default set to autonegotiate with the device connecting in the switch port. Review *"Switching in Half-Duplex and Full-Duplex Modes."*

Chapter 4: Spanning Tree Protocol (STP)

Exam Objectives

- ✓ Describing problems related to switching loops and seeing how STP fixes them
- ✓ Describing STP operation flow, root bridge, and nonroot bridge functions
- ✓ Managing STP root bridge selection
- ✓ Managing STP port type assignment
- ✓ Describing the STP convergence process
- ✓ Describing the STP port states
- ✓ Understanding best practices to decrease STP convergence duration
- ✓ Managing Cisco switch and port configuration options for STP
- ✓ Describing Rapid Spanning Tree Protocol (RSTP) improvements over STP
- ✓ Describing and managing EtherChannel

You found out in previous chapters that switches are very often interconnected using redundant links to improve reliability of interswitch links. You also discovered that redundant links can create transmission loops, causing broadcast storms and MAC address table thrashing. Switches need to avoid transmissions loops when they are interconnected using redundant links.

Read this chapter to see how the Spanning Tree Protocol (STP) (IEEE 802.1d) can be used to help Layer 2 switches avoid loops. You see how to configure and manage STP. You are also introduced to STP port types. The convergence process flow is explained, along with some parameters you can set to control STP convergence. Finally, you see how the Rapid Spanning Tree Protocol (RSTP) (IEEE 802.1w) improves STP convergence speed.

Introducing the Spanning Tree Protocol (STP)

Layer 2 switches need to avoid switching loops. Here you see how switches can avoid switching loops using STP. The STP protocol is defined in the IEEE 802.1d standard.

Before I look at STP, I review the problem that STP resolves: data-link frames being continuously sent back and forth between Layer 2 switches interconnected with redundant links.

Data-link (Ethernet) frames do not expire. An Ethernet frame sent to an unknown MAC address can bounce around forever in the network. This is not good because it wastes bandwidth. This problem is exacerbated when several switches are interconnected with redundant links. It is best practice to use redundant links to interconnect Layer 2 switches to mitigate inter-switch link failure risks. However, this amplifies the risk of having frames bouncing back and forth between switches.

Figure 4-1 shows a typical dual-switch LAN with two interswitch links. Hosts Alex, Claire, John, and Monica are interconnected in a LAN using two switches connected with redundant interswitch links. Suppose that the MAC address tables are empty on both switches because none of the hosts have sent any frames yet.

Assume that Alex sends a frame to Claire:

1. The first time Alex sends a frame to Claire, the switch learns about Alex, but it still does not know about Claire. So, switch SW1 needs to flood the frame out on all ports except the port through which it came in, that is, the port where Alex connects to SW1.
2. Switch SW1 floods the frame out on all ports except the port through which it came in.

The frame reaches switch SW2. Switch SW2 does not know about Claire either, so it also needs to flood the frame out on all ports except the port through which it came in.

3. Switch SW2 floods the frame out on all ports except the port through which it came in.

Because you have two links connecting the SW1 and SW2 switches, the frame that came in through fa0/2 on SW2 is flooded out on fa0/3, and the frame that came in through fa0/3 on SW2 is flooded out on fa0/2. Hence, switch SW1 gets the frame back, twice: on fa0/2 and on fa0/3.

4. Because SW1 still does not know Claire's MAC address, it floods the frame again on both fa0/2 and fa0/3. SW2 receives the frame again, twice, on fa0/2 and on fa0/3.

5. Because SW2 still does not know Claire's MAC address, it floods the frame out again on all ports, including fa0/2 and fa0/3, and so on.

This is a typical example of a *broadcast storm*: a broadcast frame that bounces forever between switches interconnected with redundant links.

Observe another problem in Figure 4-1. Because switches SW1 and SW2 receive the same frame on two different ports, they register the source MAC address (S-MAC) of host Alex for both ports. This is a problem because the switch will not know over which port it needs to send frames out to Alex whenever D-MAC is set to Alex's MAC address in a frame. This is a typical example of *MAC address table thrashing*.




Figure 4-1:
Layer 2
transmission
loop.

Layer 2 transmission loops cause the following:

- ◆ **Broadcast storms:** A broadcast frame that bounces forever between switches interconnected with redundant links is called a broadcast storm. Broadcast storms waste bandwidth and may thrash the MAC address table.
- ◆ **MAC address table thrashing:** The MAC address table gets thrashed when multiple ports are attached to the same MAC address. The switch does not know on which outgoing port it can reach that MAC address.

To avoid loops and bouncing frames, enable the *Spanning Tree Protocol (STP)* on Layer 2 switches.

So, how exactly does STP eliminate Layer 2 loops in LANs? STP basically monitors the network and catalogs every link, particularly redundant links. Next, STP disables redundant links, setting up preferred, optimized links between switches:

- ◆ The preferred, optimized links are used under normal circumstances.
- ◆ Should any of the preferred links fail, one of the nonpreferred redundant links is enabled and used instead.

STP assigns a root bridge that acts sort of like the decision maker in the network. The term *root bridge* refers typically to a Layer 2 switch. It can also really be a bridge, but more commonly it is a switch. The root bridge decides which routes are preferred and which routes are nonpreferred.

The root bridge interacts with *nonroot bridges*: other switches in the LAN. STP is enabled on all switches in the LAN. The root bridge and the nonroot bridges have specific roles within STP.

Switch ports are categorized depending on whether they forward frames and depending on whether they are the endpoints of a preferred, optimized link within the LAN.

Root bridge: A Layer 2 switch (or bridge) that decides which routes are preferred and which routes are nonpreferred.

Nonroot bridge: A Layer 2 switch (or bridge) that is not a root bridge. Nonroot bridges coordinate with the root bridge to assist in the determination of preferred and nonpreferred routes from each nonroot bridge to the root bridge in the LAN.

STP Operation Flow

The Spanning Tree Protocol executes three operations to achieve a stable, loop-free LAN:

- ◆ **Electing STP root bridge:** The switch with the lowest *bridge ID* is elected the root bridge.
- ◆ **Assigning STP port types:** Each port on every switch in the LAN is assigned a type that defines its behavior. Switches communicate using STP to assign a type to their ports. This determines the ports' behavior to be either *forwarding* (that is, the port forwards data-link frames) or *blocking* (that is, the port sends no data-link frames).
- ◆ **Achieving STP convergence:** Convergence is the result of assigning port types. After switches have assigned a forwarding type or a blocking type to all ports that interconnect them, STP achieves a stable, loop-free LAN.

I now look at each of these steps in detail.

Electing a root bridge

STP's first step is to choose a root bridge. STP relies on the bridge ID to decide which switch should become the root bridge. STP chooses the switch with the lowest bridge ID to be the root bridge.

Bridge ID

The bridge ID is a number composed of the switch's MAC address and its STP priority. The STP priority on any given Layer 2 switch is by default 32768. So, because the priority of all switches is by default the same (32768), STP really chooses the switch with the lowest MAC address to be the root bridge.

Figure 4-2 shows an example with four switches having the same STP priority.


The switch with the lowest MAC address is chosen as the root bridge.

STP priority

What if the switch with the lowest MAC address is actually an old, slow switch? You wouldn't want the slowest switch in your LAN to become the root bridge because it will slow every other switch in your network. You can use the STP priority setting to control which switch will have the lowest bridge ID.



Best practice is to set the STP priority to a low value (smaller than the default 32768 STP priority) on the fastest switch in your LAN. This causes STP to choose the fastest switch in your LAN as the root bridge.



You can set the STP priority to any number between 0 and 61440. However, the STP priority value has to be a multiple of 4096. So, the STP priority values you can choose are 0, 4096, 8192, 12288, . . . 61440. If you set the STP priority to 0, that switch will be the root bridge.



It is not best practice to set the STP priority to 0 on more than one switch. Setting the STP priority to 0 on more than one switch would basically tell STP to use those switches as the root bridge. However, you can only have one root bridge in any given Layer 2 spanning tree, so STP will simply choose the switch with the lowest MAC address among the switches with STP priority set to 0.

Figure 4-3 shows the four switches example again, this time with a lower STP priority value set on the faster Cisco Catalyst 3560 SW2 switch.

Observe that the faster Cisco Catalyst 3560 switch with the lower STP priority value is chosen as the root bridge.

STP priority is VLAN-bound. In other words, each virtual LAN can have a different root bridge. However, this is not a common setup. Usually, all VLANs have the same root bridge. You can subdivide a LAN into several logical (virtual) LANs (VLANs) using the same physical equipment. In other words, you can have some devices communicate within one virtual LAN and other devices communicate within another virtual LAN on the same cables, switches, and routers. This allows you to identify and separate network traffic for different purposes on any given network.




Figure 4-3: Electing the STP root bridge with STP priority control.

STP priority related to the Cisco hierarchical network architecture

Cisco networks are designed using a hierarchical architecture. Three layers make up the Cisco hierarchical network architecture:

- ◆ Core layer
- ◆ Distribution layer
- ◆ Access layer

The *core layer* is the layer that sits at the center of the network. This layer is also called the *backbone*.

Ultimately, traffic from all devices in the network may end up being routed to the core of the network. The core layer is “where networks meet.”

The *distribution layer* links the core layer to the access layer. This layer is also called the *workgroup layer*. The distribution layer is involved in routing packets between nodes connected at the access layer. Packets may need to be routed to a different network, which may be located within the same distribution layer network or up through the core layer in a different distribution layer network.

The *access layer* is the layer that interconnects host devices to LANs. The access layer is where workgroup LANs are defined. This layer is also called the *desktop layer*, because it usually involves interconnecting desktop computers and concentrating their traffic into a distribution layer switch or router. Switches, small routers, and computer host devices typically interconnect at the access layer.

Layer 2 switches are typically used at the access layer and at the distribution layer. It is important to understand that the larger the LAN, the more choices are available for a root bridge. As mentioned earlier, it is best practice to assign STP priorities to ensure that the fastest switch in the LAN becomes the root bridge.

Figure 4-4 shows a very large Cisco LAN that spans both the distribution and the access layers.




Figure 4-4:
STP priority
in a large
LAN.

Observe that you set the STP priority values in a hierarchical way to ensure that the fastest Cisco Catalyst 6500 series switch is elected the root bridge. One of the Cisco Catalyst 4500 series would be next in line for becoming the root bridge, should the 6500 series ever fail. You set the STP priority in a hierarchical way to ensure that one of the next-fastest switches is elected, if the root bridge ever fails or is unavailable. If you set the STP priorities all equal below the 6500 series, any of those switches could become the root bridge. Do you remember which one would, in that case? If you answered the one with the lowest MAC address because all STP priorities would be equal, you are right. Note also that within the same level, faster switches have lower STP priorities to ensure that they would be the first ones to become the root bridge, if necessary.

Assigning STP port types

The second step in the STP operation flow is assigning port types to every port in the network. STP assigns port types according to the cost of each path between the root switch and a nonroot switch.

STP path cost

Switches are typically interconnected using redundant links to increase resilience, should one of the links fail. Each of these links provides a connection path. Each path has a certain cost associated with it. The cost is defined according to the bandwidth of the link. Table 4-1 lists the costs associated with each Ethernet bandwidth.

Table 4-1 STP Path Cost for Each Ethernet Bandwidth	
<i>Bandwidth</i>	<i>STP Cost</i>
10 Mbps	100
100 Mbps	19
1 Gbps	4
10 Gbps	2

STP prefers paths with lower costs. Looking at Table 4-1, you can see that STP prefers faster bandwidths. Consider a switch that has a 100-Mbps link and a 1-Gbps link connecting it to the root bridge. STP will designate the port that connects to the 1-Gbps link to be *root port*, the port that connects to the root bridge. The port connecting to the slower 100-Mbps link will be either a *designated port* or a *blocking port*. A designated port is a port that forwards data-link frames. A blocking port is a port that does not forward data-link frames. You read more about designated and blocking ports a bit later in this chapter.

Setting root ports

After switches have elected a root bridge, nonroot switches assign one of their ports to be their root port. The STP root port is the port that connects a switch to the STP root bridge. The port that is elected to be the root port is either

- ◆ A port that connects the nonroot switch directly to the root bridge
- ◆ A port that connects to the least expensive path (STP cost) to the root bridge

Switches communicate between themselves using STP to calculate the cost of each path to the root bridge. Each switch adds the cost of its own path to the cost received from the neighboring switches to determine the total cost of a given path to the root. After the costs of all paths to the root have been calculated, each switch assigns the root port to the port connecting to the path with the lowest cost.

Figure 4-5 illustrates a simple LAN with a root bridge switch and three non-root switches. Each nonroot switch sets one of its ports to be the root port.




Figure 4-5:
Setting the
root port.

Observe that each of the three nonroot switches chooses a root port that connects directly to the root bridge.

Note also that switch SW4 only has one link connecting it to the SW1 root bridge. In this case, there isn't a choice for the root port: It is the one port that connects SW4 to SW1, the root bridge.

The SW2 switch chooses the direct path to the root bridge because it is less costly than the path going through the SW3 switch. Similarly, the same is true for the SW3 switch.

The SW1 switch, the root bridge, has no root ports because it is the root bridge itself. SW1 only has designated ports.

Setting designated ports


After each switch assigned one of its ports to be the root port, connecting to the root bridge, the remaining ports are either *designated* or *blocking*:

- ◆ An STP designated port is a port that forwards data-link frames.
Each LAN segment needs to have a designated port, a port that forwards traffic in and out of that LAN segment.
- ◆ An STP blocking port is a port that does not forward data-link frames.

Figure 4-6 shows the same four switches, this time with root ports, with designated ports, and with blocking ports.

Observe the link between switch SW2 and SW3. Switches SW2 and SW3 need to decide which port is designated and which port is blocking on the link that interconnects them. Because both SW2 and SW3 connect to SW1, the root bridge, through a path with an STP cost of 19 (100 Mbps), the total STP cost to the root is the same in both directions on the SW2/SW3 link: The SW2 to SW3 to SW1 STP cost is $19 + 19$, and SW3 to SW2 to SW1 is also $19 + 19$. So, SW2 and SW3 cannot set their designated and blocking ports relying solely on the STP path cost. SW2 and SW3 need to look at other criteria to decide which port is designated and which port is blocking on the link that interconnects them. In this case, SW2 and SW3 look at their bridge ID: The switch with the lower bridge ID sets its port to be designated, while the switch with the higher bridge ID sets its port to be blocking. Recall that the bridge ID is a number composed of the switch's MAC address and its STP priority. Both SW2 and SW3 have the same STP priority, so SW2 wins because its MAC address is lower than SW3's MAC address.

Switches look at other criteria anytime a tie exists in the STP path cost. The following sections describe the criteria that switches analyze to decide which of their ports is the root port, which of their ports are designated ports, and which of their ports are blocking ports.



Choosing ports according to STP path cost

The port that connects to the path reaching the root bridge with the lowest cumulative STP cost becomes either the root port or a designated port.

For example, SW2 connects to SW1, the root bridge, through port fa0/1. The STP cost of that path is 19. SW2 also connects to SW1, the root bridge, through port fa0/2, via SW3. The total STP cost of that path is $19 + 19 = 38$. Because the fa0/2 STP cost of 38 is greater than the fa0/1 STP cost of 19, SW2 will set fa0/1 to be its root port.

Choosing ports according to bridge ID

If the switch has multiple ports and the STP path cost is the same for all ports, the port that connects to the path reaching the switch with the lowest bridge ID becomes either the root port or a designated port.

For example, Figure 4-7 adds a new switch, SW5 to the network. SW5 is connected to both SW2 and SW3, but it is not directly connected to SW1, the

root bridge. Several routes exist from SW5 to SW1. One route is SW5 to SW2 to SW1. Similarly, another route is SW5 to SW3 to SW1. You can see that the STP cost is equal for these two routes: $19 + 19 = 38$. In this case, SW5 will choose port fa0/1 that connects to SW2 to be its root port because SW2 has a lower bridge ID than SW3.




Figure 4-7: Setting the designated and blocking ports based on bridge ID.

Observe that the second port on SW5 connecting to SW3 ends up blocking. Here is why: After SW5 has chosen its root port, SW5 and SW3 communicate to decide how they assign the ports that interconnect them. Which port will be designated on that segment, and which port will be blocking? As usual, the STP path cost is analyzed first: The STP path cost from SW3 to the root bridge is 19. The STP path cost from SW5 to the root bridge is $19 + 19 = 38$, because the root port on SW5 connects to SW2 and then SW2 connects to SW1. The STP path cost is greater on SW5 than on SW3, so SW3 wins. The port on SW3 becomes designated for the SW5-to-SW3 link. The port on SW5 becomes blocking for the SW5-to-SW3 link.

Choosing ports according to port number

If the switch has multiple ports, the STP path cost is the same for all ports, and the bridge ID is the same (that is, two links interconnecting the same two switches), the port with the lowest port number becomes either the root port or a designated port.

Setting blocking ports

Earlier you read that each LAN segment needs to have a designated port, a port that forwards traffic in and out of that LAN segment. The port at the other end of the link becomes a blocking port. An STP blocking port is a port that does not forward data-link frames, thereby closing the loop. A blocking port still receives data frames, but it sends no data frames out on the link. In other words, the link becomes one-way, from the designated port to the blocking port.

Achieving STP convergence

After switches have assigned a forwarding type or a blocking type to all ports that interconnect them, STP achieves a stable, loop-free LAN. This is a *converged network*. Convergence is the result of assigning port types to eliminate loops in the LAN.

Introducing bridge protocol data units (BPDUs)

Switches communicate by sending each other bridge protocol data units (BPDUs).

A BPDU is a special-purpose data-link frame that is multicast every 2 seconds. BPDUs contain information about STP path costs, bridge IDs, port IDs, and some other parameters that help switches to elect the root bridge and to decide how to assign port types to their respective ports.

Figure 4-8 illustrates the structure of IEEE 802.1d BPDU data-link frames.

Figure 4-8:
Bridge
protocol
data unit
(BPDU)
structure.

Protocol ID
Version
BPDU Type
Flags (TCN, ACK)
Root Priority
Root ID
STP Root Path Cost
Bridge Priority
Bridge ID
Port ID
Message Age
Max Age
Hello Time
Forward Delay

Topology change notification (TCN) BPDUs

What if the topology of the LAN changes? What if you add or remove a link between two switches, or if you add or remove a switch? In that case, the switches involved send a topology change notification (TCN) BPDU.

A TCN is a specialized BPDU that informs every switch in the LAN that the topology has changed. At that point, switches need to go through the convergence process again: They need to reelect a root if the root bridge was affected by the topology change. Next, they need to decide how they set up their ports. Finally, they achieve a converged state.

STP port states

Ports on switches running STP are in one of the following states:

- ◆ **Blocking:** Ports start in the blocking state. A blocking port forwards no data-link frames. It listens to what's happening in the LAN: It receives and analyzes BPDUs. Having all ports blocking when a switch is powered up prevents the switch from creating and using transmission loops while STP is converging. A port that is an endpoint to an interswitch link stays in the blocking state unless it becomes a root port or a designated port during the STP convergence process.

A port can also transition to the blocking state if the switch receives a topology change notification (TCN) BPDU. Whenever the network topology changes, STP blocks all ports until the STP convergence process is restarted. A delay occurs before the STP convergence is restarted. Whenever a switch or a link fails or becomes unavailable, switches connected to the failing switch or link wait for 20 seconds before they start the STP convergence process. This 20-second wait is called the *max age timer*. You can view and change the max age timer using Cisco IOS commands.

Blocking state duration: 20 seconds (*max age timer*)

- ◆ **Listening:** Next, ports transition to the listening state. A listening port listens to BPDUs to prepare to transmit frames. A port in the listening state would only send and receive BPDUs without populating the MAC address table of the switch. This is the state in which switches communicate using BPDUs to assign a port type to the listening port. A port stays in the listening state for 15 seconds by default. The listening time and the learning time make up the *forward delay timer*. You can view and change the forward delay timer using Cisco IOS commands.

Listening state duration: 15 seconds (*forward delay timer — part 1*)

- ◆ **Learning:** Ports transition to the learning state. A learning port listens to BPDUs and populates its MAC address table. The purpose of the learning state is to allow the switch to gather information about the MAC address reachable on each port. A port in learning mode sends no data frames. A port stays in the learning state for 15 seconds by default. The listening time and the learning time make up the *forward delay timer*. You can view and change the forward delay timer using Cisco IOS commands.

Learning state duration: 15 seconds (*forward delay timer — part 2*)

- ◆ **Forwarding:** At this point, the port has become a root port, a designated port, or a blocking port. If the port is either root or designated, it transitions to the forwarding state. A port in the forwarding state can send and receive data frames.
- ◆ **Disabled:** A disabled port is basically a port that has been shut down manually by the switch administrator. Disabled ports do not participate in the network: They do not go through the convergence process, and they do not send or receive frames. They are basically turned off.

STP convergence duration

The STP convergence process takes some time to complete. During this time, no data communication passes through the switches involved. In other words, while switches go through the STP convergence process, host devices connected to the LAN do not have access to the network.



How long must the hosts wait before they get network access? Look at port-state default durations listed previously:

- ◆ 20 seconds in the blocking state (max age timer)
- ◆ 15 seconds in listening state (forward delay — part 1)
- ◆ 15 seconds in the learning state (forward delay — part 2)

The total time is 50 seconds. Hence, hosts need to wait up to 50 seconds before they get network access anytime a switch is rebooted, or a link becomes unavailable. In other words, the network becomes unavailable for up to 50 seconds every time a change in topology occurs. This is not good. This is why Cisco and the IEEE improved STP to decrease STP convergence delays. These improvements are explained in the following sections.

Best practices to decrease STP convergence duration

It is best practice to set the STP priority to a low value (smaller than the default 32768 STP priority) on the fastest switch in your LAN. This causes STP to choose the fastest switch in your LAN as the root bridge.

Also, in larger LANs, when using several switches with various performance levels, it is best practice to set STP priority values in a hierarchical way according to the performance level of each group of switches to ensure that one of the next-fastest switches is elected the root bridge, if the root bridge ever fails or becomes unavailable.

Introducing Cisco Options for STP

Even when best practices are followed, the STP convergence delays may still be too long. You can use the following Cisco configuration options for STP to decrease STP convergence delays.

PortFast

The Cisco PortFast port option is used on ports that do not need to participate in STP. These are typically ports that do not interconnect switches, bridges, or hubs. For example, a switch port that connects directly to a host device using a single link does not need to participate in STP because there are no chances that that port will ever create a switching loop.

Enabling PortFast on a switch port basically turns off STP for that particular port. The advantage is that the port does not need to wait for the STP convergence to complete (up to 50 seconds) before it is active.

To configure PortFast on interface fa0/1, run the following commands:

```
SW1>en
SW1>configure terminal (or config t)
SW1(config)>interface fastethernet 0/1 (or int fa0/1)
SW1(config-if)>spanning-tree portfast
```

To configure PortFast on interfaces fa0/1 to fa0/4, run the following commands:

```
SW1>en
SW1>configure terminal (or config t)
SW1(config)>interface range fastethernet 0/1 - 4
SW1(config-if-range)>spanning-tree portfast
```

BPDUGuard

The BPDUGuard Cisco option is used in conjunction with the PortFast option on access layer switches. The PortFast option is dangerous if it is enabled by mistake on a port that interconnects switches: Enabling PortFast turns off STP on that port. This can potentially create a switching loop. The BPDUGuard option monitors the frames received on the PortFast-enabled port.

If the port ever receives a BPDU, it means that it is connected to a switch. In this case, BPDUGuard sets the port into an error-disabled state. This generates an error message on the switch console to inform the switch administrator that the port should not be PortFast-enabled. This situation can happen, for example, when an administrator mistakenly connects a switch to a port that was previously connected to a host. The port was PortFast-enabled because a host cannot create a switching loop. But now a switch connects to the port, so PortFast needs to be disabled. The BPDUGuard option ensures that an error is raised and that the port is disabled whenever this happens.

To configure BPDUGuard on interface fa0/1, run the following commands:

```
SW1>en
SW1>configure terminal (or config t)
SW1(config)>interface fastethernet 0/1 (or int fa0/1)
SW1(config-if)>spanning-tree bpduguard enable
```

To configure BPDUGuard on interfaces fa0/1 to fa0/4, run the following commands:

```
SW1>en
SW1>configure terminal (or config t)
SW1(config)>interface range fastethernet 0/1 - 4
SW1(config-if-range)>spanning-tree bpduguard enable
```

BPDUFILTER

The BPDUFILTER option does not allow BPDU frames to enter or exit a PortFast-enabled port. Without BPDUFILTER enabled, a port that is PortFast enabled still receives BPDU frames. BPDUs are only useful in the context of STP. So, it makes sense to enable BPDUFILTER to fence off the BPDUs on a port that is PortFast enabled because that port does not participate in STP.

If a BPDU reaches the PortFast-enabled port with BPDUFILTER option set, the PortFast option is disabled and the port starts to participate in STP. This is different from BPDUGuard: When a PortFast-enabled port with BPDUGuard option set receives a BPDU, the port is turned off and an error is logged. BPDUFILTER leaves the port on but turns off the PortFast option, thereby enabling STP on the port. You can use either the BPDUGuard option, the BPDUFILTER option, or both on a PortFast-enabled port. Best practice is to use at least one of these options along with PortFast for ports that do not need to participate in STP.

To configure BPDUFILTER on interface fa0/1, run the following commands:

```
SW1>en
SW1>configure terminal (or config t)
SW1(config)>interface fastethernet 0/1 (or int fa0/1)
SW1(config-if)>spanning-tree bpdufilter enable
```

To configure BPDUFILTER on interfaces fa0/1 to fa0/4, run the following commands:

```
SW1>en
SW1>configure terminal (or config t)
SW1(config)>interface range fastethernet 0/1 - 4
SW1(config-if-range)>spanning-tree bpdufilter enable
```

UplinkFast

The UplinkFast Cisco option speeds STP convergence when at least one alternate or backup root port exists on the switch.

This option should only be enabled if an alternate root port exists. The alternate root port would be blocked in normal operation. If the root port fails or the link starting at the root port fails, UplinkFast switches right away to the use the alternate root port without waiting for the STP convergence process to complete. This results in quicker failover time whenever the root port/link fails. This option is typically enabled on access layer switches that have redundant links to distribution layer switches.

Figure 4-9 illustrates the network shown in Figure 4-7 with an additional link between SW2 and SW5.




Figure 4-9:
STP
UplinkFast
topology.

In this setup, you can enable STP UplinkFast on switch SW5 because it has an alternate root port connecting it to switch SW2. Port fa0/3 on SW5 is now blocked. By enabling STP UplinkFast on SW5, port fa0/3, the alternate root port on SW5, would become root if the current root port, fa0/1, ever becomes unavailable, or if the link starting at fa0/1 on SW5 ever becomes unavailable.

To configure UplinkFast on switch SW1, run the following commands:

```
SW1>en
SW1>configure terminal (or config t)
SW1(config)>spanning-tree uplinkfast
```

BackboneFast

The BackboneFast Cisco option allows a switch to detect failures on links that are not connected directly to the switch. This option can be set on any switch including the root switch to allow all switches in the LAN to detect indirect link failures before TCN BPDUs are sent. This speeds STP convergence.

To configure BackboneFast on switch SW1, run the following commands:

```
SW1>en
SW1>configure terminal (or config t)
SW1(config)>spanning-tree backbonefast
```

Introducing Rapid Spanning Tree Protocol (RSTP)

Layer 2 switches need to avoid switching loops. However, STP convergence delays are a problem. Following STP best practices increases the chances to have a quick STP convergence after a topology change, but you have no guarantee that it will be any quicker than 50 seconds. Cisco-specific STP options also increase the chances to have a quick STP convergence after a topology change, but still, there is no guarantee that it will be any quicker than 50 seconds.

The Rapid Spanning Tree Protocol (RSTP) resolves this problem. The RSTP protocol is defined in the IEEE 802.1w standard. RSTP integrates the STP Cisco options described previously and adds more port types to decrease the time required to achieve STP convergence.

Shorter delay before STP recalculation (max age timer)

Remember that ports start in the blocking state whenever a switch is powered up. Switch ports also transition to the blocking state whenever a change in topology occurs (that is, whenever they receive a TCN BPDU). Switches wait for 20 seconds before they start the STP recalculation process. This 20-second wait time is called the *max age timer*.

Considering that BPDUs are sent every 2 seconds, STP switches wait for 10 BPDUs before they start the STP recalculation process.

RSTP reduces this 20-second delay. RSTP waits for just three BPDUs before it starts the STP convergence process. This considerably speeds the STP convergence process: a 6 max age timer delay in RSTP instead of the 20-second max age timer delay in STP.

Alternate port and backup port

RSTP adds two port types to the mix: *alternate ports* and *backup ports*. The purpose of these additional ports is to reduce STP convergence time by allowing switches in the LAN to have failover plans in the event that their root port or designated port fails.

The BPDUs in RSTP are extended to contain details not just about root, designated, and blocking ports, but also information about alternate and backup ports.

Alternate port

An *alternate port* is a port that becomes the root port in the event that the root port or the link starting at the root port fails or becomes unavailable.

The idea is to allow the switch to quickly use the alternate port as a root failover port without the need to go through the lengthy STP convergence process to find an alternate root port. The switch basically keeps in mind a “plan B” (alternate port) for the root port in the event that the current root port fails.

Backup port

A *backup port* is a port that becomes the designated port in the event that the designated port or the link starting at the designated port fails or becomes unavailable.

The idea is to allow the switch to quickly use the backup port as a designated failover port without the need to go through the lengthy STP convergence process to find an alternate designated port. The switch basically keeps in mind a “plan B” (backup port) for the designated port in the event that the current designated port fails.

Rapid Transitioning to Forwarding (RTF)

The IEEE 802.1w standard that defines RSTP adopted the Cisco PortFast port configuration option and named it Rapid Transitioning to Forwarding (RTF).

The same as the Cisco PortFast port configuration option, RTF speeds the STP convergence by transitioning *edge ports* very quickly to a forwarding state. A switch edge port is a port that connects to the edge of the network: a port that connects to an end device, such as a computer host, instead of connecting to another switch. These ports do not need to participate in

STP, so they can transition to forwarding mode before STP convergence is achieved.

Enabling RSTP on a Cisco switch

RTF (PortFast) and other Cisco switch and port options for STP need to be enabled manually for RSTP as for STP. After Cisco switch and port options for STP have been set, you can turn on RSTP.

To enable RSTP on switch SW1, run the following commands:

```
SW1>en
SW1>configure terminal (or config t)
SW1(config)>spanning-tree mode rapid-pvst
```

EtherChannel

The STP convergence process assigns STP port types to all switches in the LAN, setting the ports in either the forwarding mode or blocking mode. On any given LAN segment, you only have one forwarding port. All other ports are blocking. The blocking ports would only be used if the current forwarding port fails.

Why would you have redundant links between switches if STP effectively shuts them down? Why would you spend extra money to have redundant links if you cannot use them? The additional links are only used for failover. Wouldn't it be nice if you could use these extra links to send data, thereby increasing the throughput of your LAN? That's exactly what EtherChannel does.

EtherChannel allows you to group several physical links into a single logical link. This is also called *port trunking* because you put several ports in a logical trunk. Cisco calls it *link aggregation* because you group several physical links together to create a single logical (aggregated) link. Both terms are valid because you group several ports in a trunk, thereby grouping several physical links in a single logical link.

You can group up to eight ports in a same trunk using EtherChannel.

EtherChannel and STP are friends

It is best practice to enable EtherChannel on redundant interswitch links when using STP. This groups redundant physical ports into a single logical port trunk. STP considers the whole port trunk to be a single link, so it will configure it to be forwarding (either a root port or designated port).

Although STP sees the port trunk as a single port, you still have several physical links between your switches. So, STP is happy because only one (logical trunk) port exists between the switches: You have no danger of creating a switching loop, so you don't need to block any ports. On the other hand, you still have several physical links between your switches to provide the desired redundancy.

Enabling EtherChannel on your redundant links provides three main advantages:

- ◆ By grouping up to eight physical ports into a single logical port trunk, you add the bandwidth of each port to provide an *increased bandwidth* in the port trunk. For example, if you group eight Fast Ethernet ports into a port trunk, your total bandwidth in the port trunk would be $8 \times 100 \text{ Mbps} = 800 \text{ Mbps}$.
- ◆ EtherChannel uses *load-balancing* algorithms to spread the network traffic across all ports in the port trunk. If one of the ports becomes overloaded, traffic is distributed across the remaining ports.
- ◆ EtherChannel has built-in *fault tolerance*: If one port or link fails, EtherChannel sends traffic across the remaining ports in the trunk.

Network applications consider the whole EtherChannel port trunk as a single network link, so they are not affected by a single port failure within the EtherChannel trunk. Better yet, STP also considers the whole EtherChannel port trunk as a single link, so STP recalculation is not necessary when a port fails within the EtherChannel trunk. STP recalculation would only be necessary if the whole EtherChannel port trunk fails. The EtherChannel port trunk would only fail when all the ports in the EtherChannel trunk fail.

EtherChannel versions

EtherChannel was initially developed by Cisco — actually, a company that Cisco acquired. Later, IEEE released the 8023.3ad standard, which defines an open-standard version of EtherChannel. You can use either version on Cisco switches.

Port Aggregation Protocol (PAgP)

The Port Aggregation Protocol (PAgP) is the Cisco-proprietary protocol that can only be used on Cisco switches to manage EtherChannel.

Link Aggregation Control Protocol (LACP)

The Link Aggregation Control Protocol (LACP) is a protocol defined by the IEEE 802.3ad standard to manage EtherChannel. LACP can be used on any switch brands.

Enabling EtherChannel on SW2 and SW5

To use link aggregation on a redundant link between two switches, you need to enable EtherChannel on both switches. You can enable EtherChannel using the `interface port-channel` Cisco IOS command.

To configure EtherChannel on switch SW2 and assign interfaces fa0/3 to fa0/4 to the EtherChannel, run the following commands:

```
SW2>en
SW2>configure terminal (or config t)
SW2(config)>interface port-channel 1 (or int port-channel 1)
SW2(config)>interface range fastethernet 0/3 - 4
SW2(config-if-range)>channel-group 1 mode on
SW2(config-if-range)>exit
SW2(config)>
```

You need to run the previous commands on SW5 as well.

The `interface port-channel 1` command creates the port channel no. 1, also known as EtherChannel 1.

The `interface range fastethernet 0/3 - 4` command selects a range of physical interfaces to work with.

The `channel-group 1 mode on` command assigns the selected interface, or the selected interface range, to port channel group (EtherChannel) no. 1.

The `channel-group mode` command has a few options that control which port aggregation protocol is used to aggregate the physical ports in the EtherChannel you create. Recall that you can choose between two port (or link) aggregation protocols: Cisco's proprietary Port Aggregation Protocol (PAgP) and IEEE's 802.3ad standard named Link Aggregation Control Protocol (LACP). Table 4-2 summarizes the options available with the `channel-group mode` Cisco IOS command.

Table 4-2 Choosing a Port (Link) Aggregation Protocol			
Channel-group command	Protocol	Behavior	
channel-group <channel #> mode on	LACP	Sends LACP frames to neighbor switch	
channel-group <channel #> mode off	None	Port aggregation is off	
channel-group <channel #> mode auto	PAgP passive	Waits for neighbor switch to send PAgP frames	

(continued)

Table 4-2 (continued)

<i>Channel-group command</i>	<i>Protocol</i>	<i>Behavior</i>
channel-group <channel #> mode passive	LACP passive	Waits for neighbor switch to send LACP frame
channel-group <channel #> mode desirable	PAgP active	Sends PAgP frames to neighbor switch
channel-group <channel #> mode active	LACP active	Sends LACP frames to neighbor switch

Monitoring STP

By default, STP is enabled on Cisco switches. However, you need to set port options and switch-wide STP options manually. Also, RSTP needs to be enabled manually. The following sections show how to look at the current STP configuration on a switch.

To monitor STP configuration on a switch, run the ubiquitous Cisco IOS show command with the spanning-tree option.

Monitoring switch STP configuration

You need to be in the configuration terminal mode (using the configure terminal IOS command) for switch-bound configuration monitoring:

- ◆ Displaying general STP configuration:

```
SW1>show spanning-tree
```

- ◆ Displaying UplinkFast configuration:

```
SW1>show spanning-tree uplinkfast
```

- ◆ Displaying BackboneFast configuration:

```
SW1>show spanning-tree backbonefast
```

Monitoring port STP configuration

You need to be in the interface configuration terminal mode (configure terminal IOS command and interface selection IOS command) for switch-bound configuration monitoring:

- ◆ To display PortFast on interface fa0/1, run the following command:

```
SW1>show spanning-tree interface fastethernet 0/1  
portfast
```

- ◆ To display BPDUGuard on interface fa0/1, run the following command:

```
SW1>show spanning-tree interface fastethernet 0/1  
bpduguard
```

- ◆ To display BPDUFilter on interface fa0/1, run the following command:

```
SW1>show spanning-tree interface fastethernet 0/1  
bpdufilter
```

Prep Test

- 1 What is the main purpose of STP?**
 - A ☐ Routing IP packets
 - B ☐ Avoiding switching loops
 - C ☐ Managing network segmentation in LANs using UplinkFast Cisco technology
 - D ☐ All of the above
- 2 What is the main purpose of a root bridge?**
 - A ☐ Routing IP packets
 - B ☐ Managing port security in LANs using PortFast Cisco technology
 - C ☐ Helping decide which ports and links are preferred and which ports and links are nonpreferred
 - D ☐ None of the above
- 3 What criteria are analyzed to select the root bridge?**
 - A ☐ STP priority and MAC address of the switch
 - B ☐ Port type and port state of the root port
 - C ☐ MAC address table and port access credentials on the switch
 - D ☐ All of the above
- 4 What is best practice when setting STP priority?**
 - A ☐ Set it equal on all switches
 - B ☐ Set it in a hierarchical way based the number of licensed switch ports
 - C ☐ Set it in a hierarchical way based on switch capabilities
 - D ☐ Set it to the default value of 32768
- 5 What determines the STP path cost?**
 - A ☐ The number of ports (hops) of a link to the root bridge
 - B ☐ The sum of the bandwidth of each segment of a link to the root bridge
 - C ☐ The number of segments of a link to the root bridge
 - D ☐ The duplex mode of ports (hops) of a link to the root bridge
- 6 How does STP assign port types?**
 - A ☐ According to port MAC address, bridge ID, and port number
 - B ☐ According to port bootup order, port MAC address, and bridge ID
 - C ☐ According to STP path cost, bridge ID, and port number
 - D ☐ According to STP path cost, port number, and MAC address

- 7 What are the STP port types?**
- A ☐ Root port, or designated port or blocking port
 - B ☐ Root port, or alternate or backup port
 - C ☐ Root port, fabric port, failover port
 - D ☐ All of the above
- 8 The main advantage of RSTP over STP is that it allows switches to do which of the following?**
- A ☐ Avoid loops
 - B ☐ Wake up on LAN activity
 - C ☐ Log STP status messages
 - D ☐ Converge faster
- 9 Describe the Cisco options for STP.**
- A ☐ They are useless when more than three switches are in the LAN.
 - B ☐ They are enabled manually.
 - C ☐ They are enabled by default when STP is started.
 - D ☐ They are not available if RSTP is used.
- 10 What are the main advantages of using EtherChannel to logically group redundant interswitch links?**
- A ☐ Fewer cables to run between switches, faster switch bootup time, cost
 - B ☐ Load balancing, fewer link failures, increased bandwidth per physical link
 - C ☐ Load balancing, fault tolerance, link aggregation
 - D ☐ Load balancing, fewer link failures, link aggregation
- 11 What is the Cisco IOS command used to monitor most STP?**
- A ☐ show statistics
 - B ☐ show spanning-tree
 - C ☐ show STP
 - D ☐ monitor STP

Answers

- 1 B. Avoiding switching loops. The main purpose of STP is to help switches communicate to set their ports in root, designated, or blocking state to avoid switching loops. Refer to *"Introducing the Spanning Tree Protocol (STP)."*
- 2 C. Helping decide which ports and links are preferred and which ports and links are nonpreferred. STP calculates optimal routes to the root bridge to determine which ports are forwarding and which ports are blocking on each switch. See *"Electing a root bridge."*
- 3 A. STP priority and MAC address of the switch. Switches elect a root bridge according to the bridge ID of each switch. The bridge ID of each switch is a combination of the MAC address of the switch and its STP priority. Check out *"Electing a root bridge."*
- 4 C. Set STP priority in a hierarchical way based on switch capabilities to ensure that most powerful switches are always first in line to be elected root bridge. Review *"STP priority related to the Cisco hierarchical network architecture."*
- 5 B. The sum of the bandwidth of each segment of a link to the root bridge. Switches add the bandwidth of each link to the root bridge to determine the STP path cost. See *"Setting designated ports."*
- 6 C. According to STP path cost, bridge ID, and port number. These three criteria are analyzed to determine the STP type of each port. Refer to *"Setting designated ports."*
- 7 A. Root port, or designated port or blocking port. RSTP adds the alternate port and backup port. Check out *"Assigning STP port types."*
- 8 D. RSTP helps switches converge faster. Read *"Introducing Rapid Spanning Tree Protocol (RSTP)."*
- 9 B. They are enabled manually. Cisco options for STP are enabled manually either on the whole switch or on a particular port. See *"Introducing Cisco Options for STP."*
- 10 C. Load balancing, fault tolerance, link aggregation. EtherChannel groups several physical links into a single logical link, thereby balancing the load across all physical links in the trunk, tolerating one or a few link failures in the trunk and adding up the bandwidth of the physical links to use for the logical link. Refer to *"EtherChannel."*
- 11 B. The `show spanning-tree` command is the Cisco IOS command used to monitor most STP. Check out *"Monitoring STP."*

Chapter 5: Virtual Local Area Networks (VLANs)

Exam Objectives

- ✓ Describing the purpose and implications of VLANs
- ✓ Describing the benefits of using VLANs
- ✓ Creating and managing VLANs
- ✓ Identifying VLANs using Cisco ISL and IEEE 802.1q frame tagging
- ✓ Creating and managing VLAN port trunks
- ✓ Creating and managing EtherChannel logical ports
- ✓ Understanding the difference between Port Aggregation Protocol (PAgP) and Link Aggregation Control Protocol (LACP)
- ✓ Understanding the difference between a VLAN port trunk and an EtherChannel logical port
- ✓ Creating and managing VLAN port trunk over EtherChannel logical port
- ✓ Managing switch port types: access ports and trunk ports
- ✓ Managing VLAN port trunks using the Dynamic Trunking Protocol (DTP)
- ✓ Managing the VLAN Trunking Protocol (VTP)
- ✓ Managing inter-VLAN routing

VLANs are an important feature that allows you to subdivide a LAN into several logical (virtual) LANs (VLANs) using the same physical equipment. In other words, you could have some devices communicate within one virtual LAN and other devices communicate within another virtual LAN on the same cables, switches, and routers. You can identify and separate network traffic for different purposes on any given network using VLANs. The nice thing about VLANs is that you don't need to have separate physical networks to separate traffic within the LAN.

VLANs also limit the broadcast domain. Switches do not limit broadcast domains. A broadcast frame sent on a Layer 2 switch will reach every host and network device within that LAN. This may be a problem in a large LAN. VLANs fix this problem: Each VLAN is a broadcast domain. In other words, broadcast frames only reach host and network devices within the VLAN where the broadcast frame is sent. This dramatically improves the performance and reliability of LANs.

Introducing Virtual Local Area Networks (VLANs)

A VLAN is a logical grouping of network resources and host devices based on the ports where those hosts connect to the switch, or based on the MAC address of those hosts. A VLAN is basically a logical grouping of switch ports or MAC addresses. A VLAN can span more than one physical switch. This is interesting because you could have a VLAN that groups network and host devices across many switches in a network. Data-link frames can be tagged with a VLAN identifier to indicate that they are part of a certain VLAN.

Figure 5-1 illustrates three VLANs that group the computer hosts by business department.




Figure 5-1:
VLAN
segmenting
a LAN by
business
department.

VLANs keep things tidy

The most apparent benefit of using VLANs is keeping things well organized by logically grouping your network along business departments and functions. You can even group network and host devices in VLANs by project. For example, if Engineering starts a new project and it needs to experiment with network connectivity, new hardware, and new software, it is a good idea to group the resources needed for that project into a specific VLAN.

This avoids running into a situation where the engineers run a test that would bring down the whole network.

VLANs subdivide the broadcast domain

Another benefit is not as apparent as keeping traffic separate along business functions. VLANs break up the broadcast domain:

- ◆ A *collision domain* is a logical network space where frames can collide because several hosts are sharing the bandwidth of the network medium and they can potentially send frames on the wire at the same time. It is best to *segment networks* into several smaller collision domains to lower the chances of having frame collisions. Layer 2 switches break up the collision domain by providing point-to-point links between devices connected to the switch.
- ◆ A *broadcast domain* is a logical network space where broadcast frames are sent. A broadcast frame is a data-link frame that is sent to all devices in the network. Broadcast frames are typically used to locate a device in the network. Layer 2 switches forward a broadcast frame out on all their ports except on the port where the broadcast frame entered the switch. This may become a problem in large networks. That's where VLANs enter the scene. VLANs basically subdivide the LAN into several virtual LANs. Broadcast frames are sent within a LAN. So, if you subdivide the LAN into several smaller VLANs, broadcast frames are only sent within each VLAN. That's the beauty of VLANs: They break up the broadcast domain.

This works quite nicely because most of the time, hosts need to communicate within their own VLAN. They need to communicate with hosts outside of their own VLAN more rarely. That's why it is best practice to organize VLANs along business departments and functions and/or along projects.

Referring to Figure 5-1, a computer in the Marketing group VLAN needs to access Marketing servers and computers most of the time. Only rarely does it need to communicate with a computer in the Engineering or Sales department.

Subdividing the network in the Marketing, Engineering, and Sales VLANs breaks up the broadcast domain into three subdomains: Marketing, Engineering, and Sales. So, a computer in Marketing broadcasts a frame that only reaches computers in Marketing whenever it needs to find the IP and MAC address of a server in Marketing to connect to it.

What if a computer in Marketing needs to find the IP and MAC address of a server in Sales or Engineering? In that case, routing needs to be enabled among those three VLANs. By default, data-link frames do not cross over from one VLAN to another. You need to specifically configure routing between VLANs to allow that. You see how to enable inter-VLAN routing later in this chapter.

There are many more reasons to use VLANs in a network. The following section reviews some of these benefits of VLANs and explains why you should use VLANs in your networks.

Benefits of VLANs

Here are the main benefits of using VLANs in your network:

- ◆ **Limit the size of each broadcast domain:** This is probably the most important benefit of VLANs. VLANs increase the number of broadcast domains and reduce the size of each broadcast domain. VLANs subdivide a LAN into several logical LANs. Broadcast frames are sent within a LAN. So, if you subdivide the LAN into several smaller VLANs, broadcast frames are only sent within each VLAN.
- ◆ **Improve security:** VLANs allow network administrators to logically group switch ports for specific purposes. They can assign users (computer hosts) to those VLANs by controlling the port where the computer host connects to the switch. Ports are assigned to the port group (VLAN) where that computer host needs to be. Those ports can be located on any switch in the network: A VLAN can span more than just one switch. This is a very efficient way to control access to network resources.
- ◆ **Improve network management and flexibility:** Previously you read that VLAN membership allows network administrators to control access to network resources. The nice thing is that they can manage this membership from a single location for all switches in the network. Better yet, by using dynamic VLAN membership, VLAN Trunking Protocol (VTP), and inter-VLAN routing, they can control access to network resources in a very large network with minimum effort.
- ◆ **Improve network usage and efficiency:** Network administrators assign network resources to a certain VLAN. Only computer hosts and network devices in that VLAN have access to those resources. Referring to Figure 5-1, a Marketing server should not be disturbed by requests from Engineering computers. By subdividing the network in three VLANs, you ensure that Engineering computers will not disturb Marketing servers and vice versa.



You can configure inter-VLAN routing to allow certain requests to cross over from one VLAN to another, but using VLANs, you have control; it is not all open access.

Managing VLANs

I now discuss how to create and manage VLANs. You discover VLAN 1, a special-purpose VLAN. You also find out about the difference between static and dynamic VLAN membership.

Create VLANs

To create a VLAN on a Cisco switch, you first use the `vlan` Cisco IOS command with the number that you want to assign to the new VLAN. You can assign a VLAN number between 2 and 4094. Then, you can name that VLAN using a name that identifies the purpose of the VLAN.

Create VLAN example

To create the Sales VLAN on switch SW1, run the following commands:

```
SW1>en
SW1>configure terminal (or config t)
SW1(config)>vlan 2
SW1(config-vlan)>name Sales
SW1(config-vlan)>exit
SW1(config)>
```

To create the other two VLANs shown in Figure 5-1, run the following commands:

```
SW1(config)>vlan 3
SW1(config-vlan)>name Engineering
SW1(config)>vlan 4
SW1(config-vlan)>name Marketing
SW1(config-vlan)>exit
SW1(config)>
```

Special-purpose VLANs

Looking at the previous example, you can see that you chose numbers 2, 3, and 4, respectively, for the Sales, Engineering, and Marketing VLANs.

Why did you not choose VLAN 1 for the Sales VLAN? VLAN 1 is a special-purpose VLAN. VLAN 1 is the administrative VLAN. Although you can assign ports to VLAN 1, Cisco recommends reserving that VLAN for switch and network management purposes only. VLAN 1 is the only VLAN that is precreated on a Cisco switch, and you cannot delete it.

By default, all ports on the switch are members of VLAN 1. Whenever you assign a port to a VLAN, you basically move the port out of VLAN 1 and into a new VLAN that you created.

Static and dynamic VLAN membership

How you assign ports to a VLAN depends on whether you're using static or dynamic VLAN membership.

A switch can have some ports configured with static VLAN membership and some ports configured with dynamic VLAN membership. Static and dynamic VLAN membership modes can coexist on the same switch. However, you should ensure that ports configured with dynamic VLAN membership connect to end devices, not to another switch.

Static VLAN membership

With *static VLAN membership*, you assign a specific port to a specific VLAN using the `switchport access vlan` Cisco IOS command.

Static VLAN membership is also called port-based VLAN membership because you control VLAN membership based on the port where the computer host connects to the switch.

This works quite well in small networks. However, if you need to move the computer host to a different port, you need to make sure that the new port is also configured in the correct VLAN. You may also need to change the VLAN membership of the port where the computer host used to connect, to disable access to the VLAN from that port, if that port remains unoccupied, for instance.

Port management and port-based (static) VLAN membership can become a daunting management task in larger networks. In that case, you may want to use dynamic VLAN membership instead.

Static VLAN membership port assignment example

To assign port 2 on switch SW1 to VLAN 2 (Sales), run the following commands:

```
SW1>en
SW1>configure terminal (or config t)
SW1(config)>interface fa0/2 (int fa0/2)
SW1(config-if)>switchport access vlan 2
SW1(config-if)>exit
SW1(config)>
```

Dynamic VLAN membership

With *dynamic VLAN membership*, you enable a service called VLAN Membership Policy Server (VMPS) on your switches. The VMPS keeps a table of the MAC addresses of all devices connected to your switch. You can assign each MAC address to a specific VLAN. This is great because, no matter where that device connects to your network, it will always be assigned to the correct VLAN. The MAC address of the device does not change, even if you connect the device to a different port on the switch, or even to a different switch. That MAC address is assigned to a specific VLAN on the VMPS. So, the device will always operate in the correct VLAN, no matter where it connects from.



Certain Cisco switch models, such as the Cisco Catalyst 2960 series, can only be configured as VMPS clients. You would have to configure VMPS on a Cisco Catalyst 3560 series or better.

Identifying VLANs

You can identify VLANs by their number and by their name. You found out previously that you assign a number between 2 and 4094 to VLANs that you create. VLAN 1 is reserved for administrative purposes only; it is best practice to assign your ports to other VLANs than VLAN 1. The number of a VLAN is also called the *VLAN ID*.

The following sections describe how a switch identifies a VLAN.

Tagging data-link frames with a VLAN ID

Data-link frames can be tagged with the VLAN ID. This allows any switch in the network to identify the VLAN that the frame belongs to. The VLAN ID tag changes the data-link frame: It adds a header and a trailer, in some cases. There are two methods for VLAN frame tagging:

- ◆ The Cisco Inter-Switch Link (ISL) proprietary method
- ◆ The IEEE 802.1q standard

Cisco ISL frame tagging

The Cisco ISL frame-tagging method encapsulates data-link frames in Cisco ISL frames by adding a 26-byte header and a 4-byte trailer to the original data-link frame. The header contains the VLAN ID, identifying the VLAN to which the encapsulated data-link frame belongs. The trailer includes a cyclic redundancy check (CRC) to ensure the validity of the ISL frame.

The Cisco ISL frame-tagging method is Cisco proprietary: It works only between Cisco switches. So, if the network contains switches from other vendors, using Cisco ISL may not be the best choice. In fact, Cisco is moving away from its proprietary Cisco ISL frame-tagging method to use the IEEE 802.1q standard instead.

The Cisco ISL frame-tagging method changes the size of the data-link (Ethernet) frame. This can be a problem if a switch or a computer host that does not support Cisco ISL ever receives the frame. Basically, this is a problem if anything else but a Cisco switch receives the frame. The modified data-link frame, encapsulated in a Cisco ISL tagged frame, is potentially larger than the maximum size of an Ethernet frame. Devices that do not support Cisco ISL would drop a Cisco ISL frame.

Some Cisco switch models do not support the Cisco ISL frame-tagging method anymore. For example, the Cisco Catalyst 2960 series switches support only the IEEE 802.1q frame-tagging standard. The Cisco Catalyst 3560 series switches, on the other hand, support both the Cisco ISL frame-tagging method and the IEEE 802.1q frame-tagging standard.

IEEE 802.1q frame tagging (a.k.a. dot1q)

The IEEE 802.1q frame-tagging standard changes the data-link frame header: It adds an “802.1q tag” that identifies the VLAN to which the modified data-link frame belongs. The 802.1q tag is 4 bytes long. Hence, the data-link frame size may be a bit larger than a normal Ethernet frame. However, this is not a problem, because all network devices, even computer host network interface cards (NICs), understand the IEEE 802.1q standard. Whenever they see an 802.1q tag in the Ethernet frame, they can simply extract it and process the rest of the data-link (Ethernet) frame.

The IEEE 802.1q standard is also known as the *dot1q* encapsulation method. In fact, you use the `switchport trunk encap dot1q` Cisco IOS command to set a trunk port to use the IEEE 802.1q VLAN ID frame-tagging standard.

VLAN Trunking

Connecting group hosts and network resources by using a single switch may be an efficient setup in a small network, but it does not scale very well.

What happens if you need to add more computer hosts into that network but no more ports are available on the SW1 switch? What if one of the computer hosts needs to be moved to a different location? In both cases, you would need to add a second switch, SW2, to the network. But then, how do you interconnect the two switches, SW1 and SW2, and more importantly, how do you manage the three VLANs across two switches? This is when you need VLAN trunking.

VLAN trunking allows switches to send VLAN information across interswitch links that are configured as a trunk port. By using trunk ports, VLANs can span more than one switch. You can now add new computer hosts to the second switch, SW2, and have them operate in their corresponding VLAN that now spans SW1 and SW2.

Figure 5-2 illustrates an expanded network: You now have two switches, SW1 and SW2, that are interconnected by a trunk port. The three VLANs, Sales, Engineering, and Marketing, now span both switches SW1 and SW2. You can add computer hosts to either SW1 or SW2 and assign them to one of the three VLANs that span the whole network.




Figure 5-2:
VLANs
spanning
two switches
inter-
connected
with a VLAN
trunk port.

If you also use dynamic VLAN membership, you can move a computer host from SW1 to SW2, because that computer VLAN membership depends on the MAC address of the computer host, not on the SW1 switch port where it used to connect.

EtherChannel and VLANs are friends

It is best practice to enable EtherChannel on redundant interswitch links when using STP. It is also best practice to enable EtherChannel on redundant interswitch links when using VLANs. This groups redundant physical ports into a single logical port trunk. This is also called *port trunking* because you put several ports in a logical trunk. Cisco calls it *link aggregation* because you group several physical links together to create a single logical (aggregated) link. Both terms are valid because you group several ports in a trunk, thereby grouping several physical links in a single logical link.

You can group up to eight ports in a same trunk using EtherChannel.

Although switches see the EtherChannel port trunk as a single port, you still have several physical links between your switches. So, only one logical EtherChannel trunk port exists between the switches backed by several physical interswitch links to provide the desired redundancy. Using EtherChannel with VLANs provides the same advantages as using EtherChannel with STP: increased bandwidth, load balancing, and fault tolerance.

Increasing bandwidth with EtherChannel

By grouping up to eight physical ports into a single logical port trunk, you add the bandwidth of each port to provide an *increased bandwidth* in the port trunk. For example, if you group eight Fast Ethernet ports into a port trunk, your total bandwidth in the port trunk would be $8 \times 100 \text{ Mbps} = 800 \text{ Mbps}$.

Balancing the load with EtherChannel

EtherChannel uses *load-balancing* algorithms to spread the network traffic across all ports in the port trunk. If one of the ports becomes overloaded, traffic is distributed across the remaining ports.

Tolerating faults with EtherChannel

EtherChannel has built-in *fault-tolerance*: If one port or link fails, EtherChannel sends traffic across the remaining ports in the trunk. Network applications consider the whole port trunk as a single network link, so they are not affected by a single port failure within the trunk. Better yet, STP also considers the whole port trunk as a single link, so STP recalculation is not necessary when a port fails within the trunk. STP recalculation would only be necessary if the whole port trunk fails. The port trunk would only fail when all the ports in the trunk failed.

EtherChannel versions

EtherChannel was initially developed by Cisco — actually, a company that Cisco acquired. Later, IEEE released the 802.3ad standard, which defines an open version of EtherChannel. You can use either version on Cisco switches.

Port Aggregation Protocol (PAgP)

The Port Aggregation Protocol (PAgP) is the Cisco-proprietary protocol that can only be used on Cisco switches to manage EtherChannel.

Link Aggregation Control Protocol (LACP)

The Link Aggregation Control Protocol (LACP) is a protocol defined by the IEEE 802.3ad standard to manage EtherChannel. LACP can be used on any switch brands.

VLAN or EtherChannel trunking? Both?

VLAN port trunks and EtherChannel port trunks are different. However, switch ports can be configured to be both an EtherChannel port trunk and a VLAN port trunk.

It is best practice to configure VLAN port trunks over EtherChannel port trunks whenever redundant interswitch links are available.

VLAN port trunk

You create a VLAN port trunk by configuring a switch port to be a trunk port interconnecting two switches.

The key point here is that this port may be just a single physical connection, a single link, interconnecting the switches. It is still called a VLAN port trunk because that port carries data for all VLANs in the network. So, it is a port trunk in the sense that it carries data for more than one VLAN. It is not a port trunk in the sense that you group several physical links in a single logical link. So, by default, if you set a *single* physical or logical port to be a port trunk interconnecting two switches, that port is a VLAN port trunk.

The main purpose of a VLAN port trunk is to carry data for several VLANs over a *single* physical or logical link interconnecting two switches.

EtherChannel port trunk

You create an EtherChannel port trunk by grouping together *several* physical ports interconnecting two switches in a single logical port.

The key point here is that you group several physical links into a single logical link. So, an EtherChannel port trunk is a port trunk in the sense that you group several physical links into a single logical link.

The main purpose of an EtherChannel port trunk is to provide a single logical interswitch link backed by *several* physical interswitch links to provide increased bandwidth, load balancing, and fault tolerance on that interswitch link.

Configuring EtherChannel and VLAN trunking

You can configure ports to be both an EtherChannel port trunk and a VLAN port trunk. Follow these steps:

- ◆ Create an EtherChannel port trunk:
 1. Use the `interface port-channel` Cisco IOS command to create the EtherChannel port trunk.

2. Use the `interface range` Cisco IOS command to select the physical interfaces to assign to the EtherChannel port trunk.
 3. Use the `channel-group <#> mode` command to assign the selected interfaces to the EtherChannel port trunk.
- ◆ Configure a VLAN trunk over an EtherChannel logical port trunk:
 1. Use the `interface port-channel` Cisco IOS command to select the EtherChannel port.
 2. Use the `switchport mode trunk` Cisco IOS command to set the EtherChannel port as a trunk port.

It is best practice to configure VLAN port trunks over EtherChannel port trunks whenever redundant interswitch links are available.

Figure 5-3 illustrates two switches that are interconnected with two physical links. Those links are grouped into an EtherChannel. The EtherChannel is configured as a VLAN trunk. The VLAN trunk carries frames for all VLANs in the network.

A VLAN trunk over EtherChannel is a bit like a freeway overpass:

- ◆ The physical Ethernet connections are like the support beams of the overpass.
- ◆ The EtherChannel logical port supported by the physical Ethernet connections is like the overpass flatbed held up by the support beams.
- ◆ The VLAN trunk carrying VLAN data frames is like the smooth asphalt surface that carries vehicles on top of the overpass flatbed.

The section “Managing VLAN trunk ports,” later in this chapter, shows an example of how to configure a few physical ports in an EtherChannel port trunk. The example then shows how to configure a VLAN port trunk over the EtherChannel port trunk.

Introducing switch port types

Switch ports can be used either as access ports or as trunk ports:

- ◆ You can configure a switch port to be an access port or to be a trunk port manually, using the `switchport mode` Cisco IOS command.
- ◆ You can also use the *Dynamic Trunking Protocol (DTP)* to set the switch port operation mode automatically:
 - DTP sets a switch port to be a trunk port if the port connects to another switch.
 - DTP sets the port to be an access port if the port connects to an end device, such as a computer host or an IP telephone.




Figure 5-3:
VLAN trunk
created over
Ether-
Channel port
trunk.

Access ports

Access ports are ports that connect a switch to end devices, such as computer hosts. By default, switch ports are configured as access ports and they are all part of VLAN 1, the default administrative VLAN. It is best practice to create new VLANs and assign access ports to those new VLANs, avoiding using VLAN 1 for access ports.

Access ports can only be members of one VLAN, and they only carry traffic for that VLAN. Access ports always carry traffic in native Ethernet format. The switch removes the VLAN ID tag from the Ethernet frame before it hands off the frame out on an access port. Similarly, the switch adds the VLAN ID tag to Ethernet frames it receives through an access port from an end device. If the switch receives an Ethernet frame tagged with a VLAN ID from an access port, it drops the frame because it considers it invalid. This is a “free” security feature: It prevents connecting a switch to a port that should be used by a single end device. It is “free” because the switch does not have to do anything special to provide this security feature: Switches simply do not look at the source field of an Ethernet frame when that frame comes in on an access port. Recall that tagging an Ethernet frame with a VLAN ID increases the size of the frame over the normal Ethernet frame size. So, the switch sees an oversized frame coming in on an access port: It drops it as invalid because it is oversized. The VLAN assigned to an access port is also called the *configured VLAN*.

There is an exception to the rule that an access port can only be a member of one VLAN and that it can only carry traffic for that VLAN. Access ports can be part of a single *data* VLAN and carry *data* frames only for that VLAN. However, access ports can simultaneously be part of a VLAN reserved for Voice over IP (VoIP) and carry VoIP frames for that VLAN. In other words, access ports on newer switches can be configured with two VLANs: one VLAN for data and one VLAN for VoIP. This allows you to connect both a computer host and an IP telephone to the same access port.

Trunk ports

Trunk ports are ports that connect a switch to another switch. By default, switch ports are not configured as trunk ports. You need to configure trunk ports either manually, using the `switchport mode Cisco IOS` command, or automatically, using *Dynamic Trunking Protocol (DTP)*.

Before you configure trunk ports you need to ensure the following:

- ◆ You have Fast Ethernet (100-Mbps) or Gigabit Ethernet (1-Gbps) ports on your switch. Trunk ports cannot be configured on normal Ethernet ports (10-Mbps) because they are not fast enough to carry the additional traffic load generated by several VLANs.
- ◆ You interconnect your switches using crossover cables. Switches are always interconnected using crossover cables, not straight-through cables. Newer Cisco switches have a feature called SmartPort. SmartPort adapts connection parameters according to the type of cable that is plugged in to the port and according to the device at the other end of the cable. Hence, using the Cisco SmartPort feature, you could theoretically use any cable. However, for the CCNA test purposes, switches are always connected using crossover cables.

To configure trunk ports, you can do either of the following:

- ◆ Use the `switchport mode Cisco IOS` command.
- ◆ Use the Dynamic Trunking Protocol (DTP). DTP sets a switch port to be a trunk port if the port connects to another switch. DTP sets the port to be an access port if the port connects to an end device, such as a computer host or an IP telephone.

Trunk ports are members of all VLANs, and they carry traffic for all VLANs. Trunk ports carry traffic in Ethernet frames that are tagged with VLAN ID information. When you create a trunk port, you need to choose one of the frame-tagging methods discussed earlier: the Cisco ISL proprietary tagging method or the IEEE 802.1q tagging standard. You can use the `switchport trunk encapsulation Cisco IOS` command to choose a frame-tagging method.

Switches look at the source field of an Ethernet frame when the frame comes in on a trunk port, and they process the frame. The switch extracts the VLAN ID from the Cisco ISL tag or from the IEEE 802.1q tag:

- ◆ If the frame destination is an end device on the switch, the switch rebuilds the Ethernet frame without the VLAN ID tag and forwards the frame out on the access port connecting to the end device.
- ◆ If the frame destination is an end device on another switch, the switch forwards the tagged Ethernet frame on an outgoing trunk port that reaches that destination switch.

Managing VLAN trunk ports

Read the following sections to find out more about DTP and to see a complete example of how to configure a few physical ports in an EtherChannel port trunk. The example then shows how to configure a VLAN port trunk over the EtherChannel port trunk.

Dynamic Trunking Protocol (DTP)

The Dynamic Trunking Protocol (DTP) is used to configure VLAN port trunks automatically. DTP sets a switch port to be a trunk port if the port connects to another switch. It sets the port to be an access port if the port connects to an end device, such as a computer host or an IP telephone.

DTP can also set the VLAN ID tag encapsulation method. Two VLAN ID tag encapsulation methods are available on Cisco switches:

- ◆ The Cisco ISL proprietary VLAN ID tag encapsulation method
- ◆ The IEEE 802.1q VLAN ID tagging standard, also known as the *dot1q* encapsulation method

DTP helps two switches negotiate the choice of one of these VLAN ID tagging methods.

Whenever DTP determines that a switch port connects to another switch, it communicates with DTP on the remote switch to set up a VLAN port trunk according to the mode of the ports on each switch.

Configuring switch port DTP operation mode

Switch ports are either ready to set up a trunk with another switch port or not ready to set up a trunk with another switch port. Five operation modes determine whether a switch port becomes a trunk port. You can configure these five modes on each physical port on a switch.

Switch ports can be operating in any of the following five modes.

Access

A port in access mode is a port that is not available to set up a VLAN port trunk, even if DTP on the neighbor switch port tries to set up a VLAN port trunk. You set a port in access operating mode if it connects to end devices, such as computer hosts. A port in access operating mode is a member of a single VLAN, and it carries data only for that VLAN. Access ports do not send or receive any DTP frames: DTP is disabled on these ports.

You can use the `switchport mode access` Cisco IOS command to set up a port in access operating mode.

Trunk

A port in trunk mode is a port that is available to set up a VLAN port trunk, even if DTP on the neighbor switch port does not try to set up a VLAN port trunk. In other words, no matter what the neighbor switch is doing, this port, in trunk operating mode, will try to set up a port trunk. It will succeed if the port on the remote switch is in any operating mode other than access.

You can use the `switchport mode trunk` Cisco IOS command to set up a port in trunk operating mode.

Nonegotiate

A port in nonegotiate mode is a port that does not send or receive any DTP frames: DTP is disabled on these ports. You can set a trunk port or an access port in nonegotiate operating mode. Anytime you set a port in nonegotiate mode, you are expected to set the operating mode of the port manually, instead of relying on DTP to do that.

You can use the `switchport nonegotiate` Cisco IOS command to set up a port in nonegotiate operating mode.

Dynamic desirable

A port in dynamic desirable mode is actively ready to become a trunk port. It sends and receives DTP frames. It actively seeks to set up a trunk with its neighbor port. It sets up a trunk port with its neighbor if the port at the other end of the link is a trunk port, a dynamic desirable port, or a dynamic auto port.

You can use the `switchport mode dynamic desirable` Cisco IOS command to set up a port in dynamic desirable operating mode.

Dynamic auto

A port in dynamic auto mode is passively ready to become a trunk port. It receives DTP frames but does not send any. It sets up a trunk port with its neighbor port if the port at the other end of the link is either a trunk port or a dynamic desirable port. Cisco switch ports operate by default in dynamic auto mode.

You can use the `switchport mode dynamic auto` Cisco IOS command to set up a port in dynamic auto operating mode.

Enabling EtherChannel and port trunking

In the previous chapter, you discovered how to set up an EtherChannel. Here you review how to set up an EtherChannel and you see how to create a VLAN port trunk over that EtherChannel.

Creating port channel (EtherChannel) 1

To use link aggregation on a redundant link between two switches, you need to enable EtherChannel on both switches. You can enable EtherChannel using the `interface port-channel` Cisco IOS command.

To configure EtherChannel on switch SW2 and on switch SW5 and to assign interfaces fa0/3 to fa0/4 to the EtherChannel, run the following commands on SW2 and SW5:

```
SW2>en
SW2>configure terminal (or config t)
SW2(config)>interface port-channel 1 (or int port-channel 1)
SW2(config)>interface range fastethernet 0/3 - 4
SW2 (config-if-range)>channel-group 1 mode on
SW2 (config-if-range)>exit
SW2 (config)>
```

Here are the details of the preceding commands:

- ◆ The `interface port-channel 1` command creates the port channel no. 1, also known as EtherChannel 1.

- ◆ The interface range `fastethernet 0/3 - 4` command selects a range of physical interfaces to work with.
- ◆ The `channel-group 1 mode on` command assigns the selected interface, or the selected interface range, to port channel group (EtherChannel) 1.
- ◆ The `channel-group mode` command has a few options that control which port aggregation protocol is used to aggregate the physical ports in the EtherChannel you create. You can choose between two-port (or link) aggregation protocols:
 - Cisco's proprietary Port Aggregation Protocol (PAgP)
 - IEEE's 802.3ad Link Aggregation Control Protocol (LACP) standard

Table 5-1 summarizes the options available with the `channel-group mode` Cisco IOS command.

Table 5-1 Choosing a Port (Link) Aggregation Protocol		
<i>Channel-group command</i>	<i>Protocol</i>	<i>Behavior</i>
<code>channel-group <channel #> mode on</code>	LACP	Sends LACP frames to the neighbor switch
<code>channel-group <channel #> mode off</code>	None	Port (link) aggregation is off
<code>channel-group <channel #> mode auto</code>	PAgP passive	Waits for the neighbor switch to send PAgP frames
<code>channel-group <channel #> mode passive</code>	LACP passive	Waits for the neighbor switch to send LACP frames
<code>channel-group <channel #> mode desirable</code>	PAgP active	Sends PAgP frames to the neighbor switch
<code>channel-group <channel #> mode active</code>	LACP active	Sends LACP frames to the neighbor switch

Configuring a VLAN port trunk over port channel (EtherChannel) 1

To manually configure port trunking over EtherChannel 1, run the following commands on SW2 and SW5:

```
SW2>en
SW2>configure terminal (or config t)
SW2(config)>interface port-channel 1 (or int port-channel 1)
SW2(config-if)>switchport trunk encap dot1q
```

```
SW2(config-if)>switchport mode trunk
SW2(config-if)>switchport nonegotiate
SW2(config-if)>exit
SW2(config)>
```

Here are the details of the preceding commands:

- ◆ The `switchport trunk encap dot1q` command sets the IEEE 802.1q VLAN ID tagging standard on the VLAN port trunk.
- ◆ The `switchport mode trunk` command enables port trunking over the selected port channel group.
- ◆ The `switchport nonegotiate` command disables automatic port negotiation on the trunk port that you just created over the port channel group (EtherChannel) 1.

You did everything manually here:

- ◆ You used the `switchport nonegotiate` command to disable DTP negotiation.
- ◆ You set the operating mode of the EtherChannel port manually, using the `switchport mode trunk` command.

Alternately, you can use DTP to automatically create the VLAN port trunk over the EtherChannel and to automatically negotiate the VLAN ID tagging method used by SW2 and SW5. To do this, you would need to

- ◆ Set at least one of the EtherChannel ports to dynamic desirable
- ◆ Set the other port to either dynamic auto or dynamic desirable

To let DTP automatically configure port trunking over EtherChannel 1, run the following commands:

```
SW2>en
SW2>configure terminal (or config t)
SW2(config)>interface port-channel 1 (or int port-channel 1)
SW2(config-if)>switchport mode dynamic desirable
SW2(config-if)>exit
SW2(config)>
```

You need to run the previous commands on SW5 as well. However, you can leave the SW5 port in dynamic auto operating mode because the SW2 port is dynamic desirable: SW2 initiates the DTP negotiation (dynamic desirable) while SW5 can just wait (dynamic auto).

VLAN Trunking Protocol (VTP)

You previously discovered how to create VLANs manually, using the `vlan` Cisco IOS command. Read the following sections to see how you can use the VLAN Trunking Protocol (VTP) to create and manage VLANs automatically.

VTP creates and manages VLANs

VTP is a data-link (Layer 2) protocol that facilitates the management of VLANs across several switches in a network. Using VTP, you do not need to log in to each switch to create and name each VLAN manually. Managing VLANs manually on each switch in your network works well for a few switches. VTP is a better solution in large enterprise networks.

VTP does not manage VLAN port membership

You still need to assign ports to each VLAN either manually or automatically:

- ◆ **Manually:** You can assign ports to each VLAN on each switch, using the `switchport access vlan <vlan #>` Cisco IOS command.
- ◆ **Automatically:** You can let VMPS (VLAN Membership Policy Server) assign ports to each VLAN across your network according to the MAC address of each device connecting to the LAN.

VTP benefits

Here are the advantages of using VTP to manage VLANs in your network:

- ◆ **Simplified VLAN management:** Managing VLANs manually on each switch in a large network can become a management nightmare. VTP solves this problem by centralizing VLAN management on a single switch, or a few switches, and keeping all other switches in sync, automatically.
- ◆ **Flexible VLAN management:** VTP keeps track of any VLAN change. Whenever a VLAN is created or removed, VTP notifies all switches in the network about the change with no intervention of the network administrator.

VTP domain

VTP works with VTP domains: group of switches that have the same VTP domain name. VTP domains allow a network administrator to have different VLAN management policies for different parts of the network. VTP manages switches within each VTP domain. So, you can set certain management rules for one VTP domain that includes switches at the head office, for example. You can set different management rules for VTP domains that include

switches at branch offices. A switch can only be in one VTP domain at a time.

VTP server

VTP manages VLANs using a switch that acts as the VLAN manager in the VTP domain. This switch is called the *VTP server* or *VTP domain controller*.

VTP switch operating mode

A switch can operate in three modes in VTP: server, client, and transparent.

Server mode

A switch in server mode is the VLAN manager in the domain, the VTP domain controller. The VTP server can create, change, and delete VLANs in the VTP domain and propagate VLAN changes to the other switches in the VTP domain.

You can use the `ntp mode server` Cisco IOS command to configure a switch to be a VTP server.

VLAN configurations are saved in NVRAM. This is the default mode for all Cisco switches.

Client mode

A switch in client mode gets VLAN configuration information from the VTP server and sends and receives updates to/from other switches. A VTP client cannot create, change, or delete VLANs.

You can use the `ntp mode client` Cisco IOS command to configure a switch to be a VTP server.

VLAN configuration is *not* saved in NVRAM.

Transparent mode

A switch in transparent mode does not take into account VLAN configuration information pushed by the VTP server. It maintains its own isolated list of VLANs. A VTP transparent switch sends updates to other switches about its own VLANs. However, it does not accept VLANs pushed by the VTP server. A VTP transparent switch can create, change, or delete its own local VLANs.

You can use the `ntp mode transparent` Cisco IOS command to configure a switch to be a VTP server.

VLAN configuration is saved in NVRAM.

Best practice for VTP switch operating mode

It is best practice to configure a core switch as a VTP server and configure all other switches as VTP clients. Core switches are faster, more resilient, and more easily accessed by all switches in the network. Always configure the VTP server on a core switch or on the fastest switch in your network.

Cisco switches are by default configured to be a VTP server. Anytime a new switch is added to the network, configure it as a VTP client before connecting it to the network. This is to avoid the new switch pushing out its empty VLAN configuration before it gets the current VLAN configuration from the VTP server in the network. If the new switch pushes out its empty VLAN configuration before it gets the current VLAN configuration from the VTP server, it can erase all VLAN information in the network.

VTP updates

The VTP server distributes information about new VLANs defined in the VTP domain using a VTP notification. VTP notifications are sent to all switches in the domain. VTP client switches forward the VTP notification on their downstream trunk ports. Each VTP client switch updates its VLAN database when it receives a VTP update notification.

Each VTP notification frame has a *VTP revision number*. Every time the VTP server sends new information, it increases the VTP revision number. VTP client switches keep track of the VTP revision number to make sure that they update their VLAN information with each VTP revision.

VTP pruning

VTP pruning saves some bandwidth on trunk ports and on switches by limiting the number of VTP-update transmissions. When the VTP pruning option is enabled in a VTP domain, VTP client switches only receive VTP-update frames for VLANs that are enabled on each switch.

In other words, if a VTP client switch has ports in the Sales VLAN and in the Marketing VLAN, but has no ports in the Engineering VLAN, that switch will only receive VTP updates for the Sales and Marketing VLANs.

The VTP pruning option is disabled by default, but it's a very useful option, so it is best practice to enable it.

VLAN ID range

Two ranges of VLAN IDs are available:

- ◆ The normal VLAN range is comprised of VLANs with IDs of 1 to 1005.
- ◆ The extended VLAN range is comprised of VLANs with IDs of 1006 to 4094.

VTP only manages VLANs in the normal range, with VLAN ID 1–1005. Because VLAN 1 is reserved for administrative purposes, you can only use VLANs 2–1005 with VTP. If you configure a switch in transparent VTP operating mode, you can use the additional VLAN IDs 1006–4094 locally on that switch.

VTP requirements

Before you set up VTP in your network, your switches need to satisfy a few requirements:

- ◆ **Have a few switches in your network:** This is not really a requirement, but rather a recommendation. VTP works best when you have at least two or three switches in your network. Actually, with two switches, you could probably get away without VTP, but if you plan to scale up your network, you might as well configure VTP from the start. It doesn't make sense to configure VTP when you have just one switch in your network.

- ◆ **Interconnect switches with crossover cables:** You always connect switches using crossover cables.

The Cisco SmartPort feature in newer switches enables interswitch connectivity even if you don't use crossover cables. However, the CCNA test expects you to know that switches are always interconnected using crossover cables.

- ◆ **Interconnect switches with VLAN trunk and EtherChannel** over redundant links.

- ◆ **Configure one VTP server per VTP domain:** Configure your fastest switch as the VTP server, and configure all other switches in your network as VTP clients. You may also configure VTP transparent switches if you need to have a specific isolated list of VLANs, such as in a branch office, for example.

- ◆ **Ensure that the VTP domain name is identical:** Configure the same VTP domain name on the VTP server and on all VTP client/transparent switches if you want them to have a uniform VLAN database.

- ◆ **Ensure that the VTP password is identical:** You can optionally configure a password for VTP communication in the VTP domain. All switches in the VTP domain must have the same password to be able to send and receive VTP updates.



Enabling VTP

After you have verified that VTP requirements are satisfied, you can enable VTP by running the following Cisco IOS commands on switches in your network.

VTP server switch

```
SW1>en
SW1>configure terminal (or config t)
SW1(config)>vtp mode server
SW1(config)>vtp domain my_vtp_domain
SW1(config)>vtp password my_password
```

VTP client switches

```
SW1>en
SW1>configure terminal (or config t)
SW1(config)>vtp mode client
SW1(config)>vtp domain my_vtp_domain
SW1(config)>vtp password my_password
```

Monitoring and troubleshooting VTP

To monitor your VTP configuration, you can use the `show vtp` Cisco IOS commands on switches in your network. Here is an example:

```
SW1>show vtp status
```

The most common source of problems with VTP is having different VTP domain names or different VTP passwords on your switches. Make sure that the `show vtp status` command reports the same VTP domain name and the same VTP password on all switches in your network.

Another source of problems with VTP is having mismatches between the VTP switch operating modes. Make sure that only one of your switches is set to operate in VTP server mode. Set the other switches to operate in VTP client mode or in VTP transparent mode. Remember that Cisco switches are set to operate in VTP server mode by default, so you need to keep your fastest switch in the default VTP server operating mode and set the other switches to operate in VTP client mode or in VTP transparent mode.

Routing Traffic from One VLAN to Another

At the beginning of this chapter, you read that data-link frames do not cross over from one VLAN to another, by default. VLANs break up broadcast domains at the data-link layer (Layer 2).

You need to specifically configure routing between VLANs to allow data-link frames to travel from one VLAN to another.

This section describes how to configure inter-VLAN routing and discusses the various options available to implement inter-VLAN routing using routers, network (Layer 3) switches, or both.

Each VLAN represents its own isolated LAN (virtual LAN). Imagine that you have 15 computers in three different business departments, five computers in each department. You create three networks of five computers. The five computers are interconnected with each other in each network. But, the three networks are not connected. To interconnect the networks, you need a router or a network (Layer 3) switch.

It's the same with VLANs: You have three virtual LANs, isolated logically. They are interconnected physically, but the network is logically divided along three departments just like in the previous example. As in this example, each five computer sets can communicate within their department VLAN, but they cannot communicate with a computer in another department VLAN. To enable inter-VLAN communication, you need to interconnect the VLANs with a router or a network (Layer 3) switch.

One router per VLAN

This method requires one router for each VLAN. You interconnect the routers. You can choose inexpensive routers because you only need one port to connect to the VLAN and one port to connect to another router. However, routers, especially inexpensive ones, are typically slower than switches. So, this solution does not offer the best performance.

- ◆ Advantages: Easy to set up and can use inexpensive routers.
- ◆ Disadvantages:
 - *Cost*: Even if you use inexpensive routers, the cost still adds up pretty quickly.
 - *Speed*: Routers are slower than switches. It is best to do as much as possible with switches.
 - *Latency*: Routers have much higher latency than switches. Also, interconnecting several routers, each handling just one VLAN, is a waste of routing capability.

One large router with one port per VLAN

This method requires one large router that has enough ports to connect to each VLAN. Router prices tend to increase proportionately with the number of ports. So, although this solution may be fairly easy to set up and administer, it is more expensive.

- ♦ Advantages: Easy to set up, easy to maintain, and good performance.
- ♦ Disadvantages:
 - *Cost*: More router ports means higher cost.
 - *Speed*: Although performance is good, routers are slower than switches. You may find that you can get a network (Layer 3) switch for less money than a large router and have the same or better performance.

One subinterface per VLAN (router-on-a-stick)

This method has a funny name, but it is an interesting inter-VLAN routing option. *Router-on-a-stick* leverages the capability to configure trunk ports on a router. It requires at minimum one router with one Fast Ethernet physical interface that can be trunked. You configure subinterfaces on one physical router interface and have each subinterface connect to each VLAN. You use the router to route between subinterfaces.

This method provides better performance, and it's cheaper than having one router per VLAN. It's also cheaper than having one large router with enough ports to connect to each VLAN.

- ♦ Advantages:
 - *Cost*: Cheaper than one router per VLAN or one large router.
 - *Good performance*: Much better latency than one router per VLAN.
- ♦ Disadvantages:
 - Need to configure subinterfaces on the router.
 - Physical interface is a single point of failure.

Network (Layer 3) switch

This method is similar to router-on-a-stick, except that the virtual interfaces connecting to each VLAN are created internally on a network (Layer 3) switch, not on an external router. This provides better performance than router-on-a-stick. The network (Layer 3) switch must also support routing traffic between the virtual interfaces. Some Cisco switches do not support this feature. For example, Cisco Catalyst 2960 series do not support routing traffic between the virtual interfaces. However, Cisco Catalyst 3550 and 3560 series do. It is best practice to use a core switch such as a Cisco Catalyst 6500 series to handle inter-VLAN routing.

This method provides the best performance, and it may be cheaper than using routers for inter-VLAN data frame routing. The choice depends on The number of VLANs in the network, the budget available, and performance expectations

In most combinations of these criteria, network (Layer 3) switches are the best choice. The advantages and disadvantages of using a network switch are as follows:

- ◆ Advantages:
 - *Cost*: Cheaper than routing inter-VLAN data frames with routers.
 - *Best performance*: Switches are faster than routers.
- ◆ Disadvantages:
 - Not supported on all Cisco switches.
 - Need to configure virtual interfaces and routing between virtual interfaces.

Prep Test

- 1** Virtual local-area networks (VLANs) allow you to do which of the following?
 - A ☐ Subdivide a LAN into several logical LANs
 - B ☐ Subdivide a MAC address into several logical MAC addresses
 - C ☐ Subdivide an IP address into several logical IP addresses
 - D ☐ All of the above
- 2** Virtual local-area networks (VLANs) limit the _____ domain.
 - A ☐ collision
 - B ☐ search
 - C ☐ broadcast
 - D ☐ All of the above
- 3** What is a benefit of using VLANs?
 - A ☐ Keeps things well organized by logically grouping your network based on MAC addresses
 - B ☐ Keeps things well organized by logically grouping your network based on business departments and functions
 - C ☐ Keeps things well organized by logically grouping your network based on Web sites names
 - D ☐ All of the above
- 4** You create a VLAN using the _____ IOS command.
 - A ☐ virtual lan create
 - B ☐ vlan create
 - C ☐ vlan
 - D ☐ virtual lan enable
- 5** Static VLAN membership involves which of the following?
 - A ☐ Manually assigning a specific port to a specific VLAN
 - B ☐ Automatically assigning a specific port to a specific VLAN using VMPS
 - C ☐ Manually enabling a specific port with the `virtual lan` command
 - D ☐ Automatically enabling a specific port with the `virtual lan` command

- 6** How is data-link traffic identified to belong to a specific VLAN?
- A ☐ By tagging each data-link frame with the LAN ID
 - B ☐ By tagging each data-link frame with the IP address
 - C ☐ By tagging each data-link frame with the MAC address
 - D ☐ By tagging each data-link frame with the VLAN ID
- 7** What do you need to do to create a VLAN trunk over an EtherChannel logical port?
- A ☐ Create the EtherChannel logical port first and set it up as a trunk port
 - B ☐ Create the trunk port first and set it up as an EtherChannel logical port
 - C ☐ Create a VLAN first and set it up as a trunk port
 - D ☐ Create a VLAN first and set it up as an EtherChannel logical port
- 8** What is the main purpose of the Dynamic Trunking Protocol (DTP)?
- A ☐ To configure EtherChannel logical ports automatically
 - B ☐ To configure EtherChannel logical ports and VLAN port trunks automatically
 - C ☐ To configure VLAN port trunks automatically
 - D ☐ All of the above
- 9** VTP updates represent the process by which a VTP server does which of the following?
- A ☐ Distributes information about new MAC addresses added to MAC address tables in the VTP domain using a VTP notification
 - B ☐ Distributes information about new EtherChannel logical ports defined in the VTP domain using a VTP notification
 - C ☐ Distributes information about new VLAN trunk ports defined in the VTP domain using a VTP notification
 - D ☐ Distributes information about new VLANs defined in the VTP domain using a VTP notification
- 10** VTP pruning represents the process by which a VTP server or client does which of the following?
- A ☐ Refuses VTP updates for VLANs that are not active on its ports
 - B ☐ Refuses VTP updates for VLAN trunk ports that are not active on its ports
 - C ☐ Refuses VTP updates for EtherChannel logical ports that are not active on its ports
 - D ☐ Refuses VTP updates for MAC addresses that are not active on its ports

Answers

- 1 A. Virtual local-area networks (VLANs) allow you to subdivide a LAN into several logical (virtual) LANs. Refer to *"Benefits of VLANs."*
- 2 C. Virtual local-area networks (VLANs) limit the broadcast domain. Review *"VLANs subdivide the broadcast domain."*
- 3 B. VLANs keep things well organized by logically grouping your network along business departments and functions. Check out *"VLANs keep things tidy."*
- 4 C. You create a VLAN using the `vlan` IOS command. Check out *"Managing VLANs."*
- 5 A. Static VLAN membership involves manually assigning a specific port to a specific VLAN. See *"Static and dynamic VLAN membership."*
- 6 D. Data-link traffic is identified to belong to a specific VLAN by tagging each data-link frame with the VLAN ID. See *"Identifying VLANs."*
- 7 A. To create a VLAN trunk over an EtherChannel logical port, you need to create the EtherChannel logical port first and set it up as a trunk port. Review *"EtherChannel and VLANs are friends."*
- 8 C. The main purpose of the Dynamic Trunking Protocol (DTP) is to configure VLAN port trunks automatically. Check out *"Dynamic Trunking Protocol (DTP)."*
- 9 D. VTP updates represent the process by which a VTP server distributes information about new VLANs defined in the VTP domain using a VTP notification. Look over *"VTP updates."*
- 10 A. VTP pruning represents the process by which a VTP server or client refuses VTP updates for VLANs that are not active on its ports. Review *"VTP pruning."*

Chapter 6: Voice over IP (VoIP)

Exam Objectives

- ✓ Describing the purpose and implications of VoIP
- ✓ Describing the quality of service (QoS) used with VoIP
- ✓ Differentiating between IP priority at network Layer 3 and class of service (CoS) at data link Layer 2
- ✓ Describing how a Cisco IP phone produces VoIP packets
- ✓ Describing how a Cisco IP phone interacts with a Cisco switch
- ✓ Describing how a Cisco IP phone interacts with a computer host connected to its PC port
- ✓ Describing Cisco Discovery Protocol (CDP)
- ✓ Differentiating between trusting and nontrusting switch access ports
- ✓ Configure switch access ports for VoIP traffic

Read this chapter to find out about Voice over IP (VoIP). VoIP defines a group of network applications, network protocols, and network devices that carry voice signals over the Internet Protocol (IP). More organizations are choosing to use IP telephony to save costs by concentrating their phone and data traffic over the same IP infrastructure.

Using industry standards, Cisco IP phones, Cisco VoIP gateways, Cisco switches, and Cisco routers can provide IP telephony over the same IP network that is used to provide data connectivity. Cisco IP telephony solution is now part of the Cisco Unified Communications solution framework that provides a very large array of data, storage, and telephony networking solutions. Cisco IP telephony provides advanced features such as voice mail, contact centers, fax services, advanced call routing and call forwarding, caller ID, and global corporate calling extension numbers.

Cisco IP telephony offers all the communications features one has come to expect from a major telephony solutions provider, yet without the cost of a dedicated corporate phone line network.

Introducing Voice over IP (VoIP)

VoIP protocols and network applications digitize the audio signal received from an IP phone handset microphone. Next, VoIP protocols and network applications cut the digital signal into small pieces and wrap those signal bits in IP packets. The IP packets are sent over the network to an IP telephony gateway that forwards those packets to the destination IP telephone. The destination IP telephone unwraps the IP packets, extracts the digital audio bits, reassembles them in order, converts them to analog sound, and sends the analog sound signal out on the handset speaker.

Several networking products and networking protocols work together to provide the link between two IP phones. However, those networking products, networking standards, and networking protocols are beyond the scope of the CCNA test. You need to know though that an IP phone typically connects to an access port on a Layer 2 switch. That access port needs to be configured for VoIP by enabling the VoIP VLAN and configuring quality of service on the switch and on the access port itself.

VoIP Requires Quality of Service (QoS)

Have you ever noticed sound breaking off at times, or poor sound quality, when using audio instant-messaging programs or IP telephony products and services? This is typically due to VoIP without QoS or VoIP with poor QoS configuration.

Without proper QoS configuration, IP packets carrying sound bits are not sent with high priority. They may get sent behind other packets. This causes the sound to break off or be delayed, or be “scratchy.” Voice traffic sent over IP networks requires quality of service (QoS) because sound deteriorates if VoIP packets are not transmitted in an orderly and efficient manner.

Consider the following example. When you send an e-mail, the text is cut into smaller pieces and sent over IP. The receiving e-mail application can simply reorder the text before it displays it, if the text arrives out of order. Also, the receiving e-mail application can request a retransmit and wait to receive the whole text before it displays it, if an IP packet got lost in the transmission. IP phones are more sensitive to IP packet transmission errors and delays. If an IP packet carrying part of a sound gets lost, the IP phone can either render the conversation without that sound, in which case the sound breaks off, or the IP phone can request a retransmit of that packet, in which case you hear a pause in the conversation. In both cases, the sound quality is poor. This is why IP telephony packets are always sent with the highest priority. This is also why you need to configure your switch and your access port to support the VoIP VLAN and to enable QoS on the switch and on the access port.

VoIP uses QoS at two OSI layers:

- ◆ IP priority at the network layer (Layer 3)
- ◆ Class of service (CoS) at the data link layer (Layer 2)

Both values are set by default to 0. VoIP sets them to 5: higher priority.

Class of service (CoS) (IEEE 802.1p)

Data-link (Layer 2) frames can be configured with a certain class of service (CoS). By default, CoS is set to 0. Data-link frames carrying VoIP traffic are usually configured with CoS 5 (higher priority). A switch always processes and sends a VoIP data-link frame with CoS 5 before a regular data-link frame with CoS 0.

The class of service (CoS) data-link (Layer 2) option is defined in the IEEE 802.1p standard. Data-link frames are tagged with the VLAN ID, either by the Cisco ISL VLAN ID tagging method or by the IEEE 802.1q (dot1q) VLAN ID tagging standard to identify the VLAN each data link belongs to. The CoS value is the priority field in the 802.1q (dot1q) VLAN tag field. Hence, every data-link frame carries a VLAN ID and a CoS value in the VLAN ID 802.1q tag.

Cisco IP Phone

The Cisco IP phone is an end device that connects to a switch access port configured for VoIP. You see how to configure the switch access port for VoIP later in this chapter. Here you discover a bit about the Cisco IP phone device. Figure 6-1 illustrates a Cisco IP phone.

The Cisco IP phone looks and behaves like a normal phone. It has all the features of a typical business-class phone set, such as a large display, hands-free communication, dual or multiple lines, call waiting, call forwarding, caller ID display, and an illuminated keyboard.

The Cisco IP phone is also a three-port Layer 2 switch. Here are the three ports and their usage:

- ◆ **Uplink (10/100 SW) connection:** This port is reserved to connect upstream to the network switch access port.
- ◆ **PC (10/100 PC) connection:** This port can be used to connect a computer host to the phone.
- ◆ **Internal connection:** The third port is an internal port that connects to the IP phone's central processing unit (CPU).

Figure 6-2 illustrates the ports available on a Cisco IP phone.

Figure 6-1:
Cisco IP
phone.



Figure 6-2:
Cisco IP
phone ports.




You connect the uplink port on the Cisco IP phone to the upstream switch access port. The Cisco IP phone operates in the VoIP VLAN.

You can connect a computer host to the PC port on the Cisco IP phone. The computer host operates in the data VLAN configured on the upstream switch access port.

Now, you realize that the uplink port that connects the Cisco IP phone to the upstream switch access port operates like a VLAN trunk port: It interconnects two switches, the upstream network switch and the Cisco IP phone mini-switch, and carries data for two VLANs — the data VLAN and the VoIP VLAN. In the previous chapter, you read that a Cisco switch access port can carry data for two VLANs: a data VLAN and a VoIP VLAN.

Figure 6-3 illustrates a configuration where you have

- ◆ A Cisco IP phone connected to an upstream switch access port
- ◆ A computer host connected to the Cisco IP phone



The Cisco IP phone builds data-link frames carrying VoIP and sets the CoS at 5 (high priority). It tags the VoIP data-link frames with the VoIP VLAN ID. Next, the Cisco IP phone sends the VoIP frames out on the uplink port to the upstream switch access port.

You would think that the upstream switch processes the VoIP frames with high priority and forwards them to the Cisco VoIP gateway as quickly as possible. Unfortunately, by default, this is not the case. You must specifically configure the upstream switch to trust the IP priority (5) and CoS level (5) set by the Cisco IP phone in the VoIP packets. By default, in untrusted mode, switches override IP priority and CoS values they receive, with the default low priority value (0). You see how to configure the upstream switch in the section “Configuring VoIP on Cisco Switches,” later in this chapter.

The Cisco IP phone may receive data frames from the PC port if a computer host is connected. The Cisco IP phone leaves the CoS at 0, the default low-priority value. The Cisco IP phone tags these frames with the data VLAN ID. Next, it sends the data frames on the uplink port to the upstream switch access port. The upstream switch processes the data frames with normal (low) priority and forwards them to the appropriate outgoing switch port.

This is the typical operation flow. You can change the typical operation flow by changing the following:

- ◆ The default configuration of the upstream switch
- ◆ The default configuration of the Cisco IP phone

These configuration options are beyond the scope of the CCNA test, but you do need to know how to configure the upstream switch to trust the IP priority and CoS level set by the Cisco IP phone in the VoIP packets.

Cisco Discovery Protocol (CDP)

Cisco created the Cisco Discovery Protocol (CDP), which allows a Cisco switch to discover the devices connected to its ports. CDP is enabled by default on Cisco switches. CDP is also enabled by default on Cisco IP phones. This protocol is useful in VoIP environments. CDP allows the upstream switch to discover the Cisco IP phone and to negotiate interconnection parameters that are optimum for VoIP.

Negotiating VLAN

The upstream switch communicates with the Cisco IP phone using CDP to set up an interconnection link that allows the Cisco IP phone to send VoIP packets on its uplink port back to the upstream switch, either in the VoIP VLAN or in the data VLAN.

Negotiating CoS

The upstream switch also communicates with the Cisco IP phone using CDP to set up an interconnection link that allows the Cisco IP phone to send VoIP packets on its uplink port back to the upstream switch, either with default CoS level 0 or with high-priority CoS level 5.

Negotiating Cisco IP phone PC port

You can connect a computer host to the PC port on the Cisco IP phone. The computer host operates in the data VLAN configured on the upstream switch access port. By default, the Cisco IP phone leaves the CoS at 0 on data frames received from the PC port. This default option can be changed

on the Cisco IP phone. It can also be changed by the upstream switch. The upstream switch can communicate with the Cisco IP phone using CDP to set up the PC port to be trusting or nontrusting:

- ◆ A *trusting* PC port on the Cisco IP phone trusts the IP priority and CoS level set on incoming IP packets by the computer host connected to the PC port.

If, for example, the computer host connected in the Cisco IP phone PC port sets the IP priority and the CoS level at 3, and the Cisco IP phone PC port is trusting, it will keep the IP priority and the CoS level at 3.

- ◆ A *nontrusting* PC port on the Cisco IP phone does not trust the IP priority and CoS level set on incoming IP packets by the computer host connected to the PC port.

If, for example, the computer host connected in the Cisco IP phone PC port sets the IP priority and the CoS level at 3, and the Cisco IP phone PC port is nontrusting, the Cisco IP phone will reset the IP priority and the CoS level at 0, the default value for IP data packets.

Configuring VoIP on Cisco Switches

Switch access ports can operate in two VLANs: a data VLAN and a VoIP VLAN. This allows you to connect both a computer host and an IP telephone to the same upstream access port. This setup is illustrated in Figure 6-2.

The following sections describe how you configure an access port to support both data and voice VLANs. You also find out how to enable QoS on the switch and on the access port to support high IP priority and high CoS required by VoIP traffic.

Enabling QoS on the upstream switch

To configure VoIP support on the upstream switch, you first need to enable quality of service (QoS) on the switch. To do this, you use the `mls` Cisco IOS command.

For example, to enable QoS on the switch, run the following commands:

```
SW3>en
SW3>configure terminal (or config t)
SW3>mls qos
```

Configuring switch access port to trust CoS

Next, you need to configure the upstream switch access port to trust the IP priority and class of service (CoS) settings of incoming packets from the Cisco IP phone.

For example, to configure the upstream switch access port to trust the CoS level set by the Cisco IP phone on its IP packets, run the following commands:

```
SW3>en
SW3>configure terminal (or config t)
SW3(config)>interface f0/24
SW3(config-if)>mls qos trust cos
SW3(config-if)>switchport priority extend trust
```

Enabling VoIP VLAN on the switch access port

To complete the configuration of VoIP support on the upstream switch, you need to enable the VoIP VLAN on the upstream switch access port. You also configure the VoIP VLAN to use the IEEE 802.1p (CoS) class of service setting to decide the priority of IP packets coming in through the port.

For example, to configure the upstream switch access port to trust the CoS level set by the Cisco IP phone on its IP packets, run the following commands:

```
SW3>en
SW3>configure terminal (or config t)
SW3(config)>interface f0/24
SW3(config-if)>switchport voice vlan dot1p
SW3(config-if)>switchport mode access
SW3(config-if)>switchport access vlan 7
SW3(config-if)>switchport voice vlan 5
```

Here, you configure the upstream switch access port to use CoS to determine the priority of incoming IP packets.

You also set the port to be an access port, and you enable two VLANs on the access port: a VoIP VLAN with VLAN ID 5 and a data VLAN with VLAN ID 7.

Prep Test

1 What is VoIP?

- A ☐ A group of network applications and protocols that carries voice over IP networks
- B ☐ A group of network applications and protocols that allows a Cisco switch port to connect to a Cisco IP phone
- C ☐ A group of network applications and protocols that allows cell phone service providers to use Cisco IP phones for in-house cell phone network communications
- D ☐ All of the above

2 Why does VoIP traffic require quality of service (QoS)?

- A ☐ Because VoIP is unreliable
- B ☐ Because it is best practice to use QoS, generally, in IP networks
- C ☐ Because VoIP is not fast enough
- D ☐ Because sound deteriorates if VoIP packets are not transmitted orderly and efficiently

3 How is VoIP packet priority determined?

- A ☐ By VLAN ID tagging
- B ☐ By STP priority and MAC address of the switch
- C ☐ By IP priority setting on IP packets and CoS setting on data frames
- D ☐ By IP priority setting on data frames and CoS setting on IP packets

4 The Cisco IP phone sends which of the following?

- A ☐ VoIP packets with high priority and data packets with default priority
- B ☐ VoIP packets with high priority and data packets with either default priority or priority set by the computer host connected to the PC port
- C ☐ VoIP packets with high priority and data packets with priority set by the computer host connected to the PC port
- D ☐ VoIP packets with default priority and data packets with high priority

5 How is the Cisco Discovery Protocol (CDP) used?

- A ☐ To help a switch update its MAC address table
- B ☐ To allow a switch to discover the devices connected to its ports
- C ☐ To help a switch negotiate connection settings with devices connected to its ports
- D ☐ All of the above

Answers

- 1** A. VoIP is a group of network applications and protocols that carries voice over IP networks. Refer to *"Introducing Voice over IP (VoIP)."*
- 2** D. VoIP traffic requires quality of service (QoS) to ensure that VoIP packets are transmitted orderly and efficiently to avoid sound deterioration. Review *"VoIP Requires Quality of Service (QoS)."*
- 3** C. The priority of VoIP packets is determined by the combination of IP priority setting on IP packets and CoS setting on data frames. Check out *"VoIP Requires Quality of Service (QoS)."*
- 4** B. The Cisco IP phone sends VoIP packets with high priority and data packets with either default priority or priority set by the computer host connected to the PC port. Read *"Cisco IP Phone."*
- 5** D. All of the above. CDP allows the switch to discover devices connected to its ports. After a device is discovered, CDP helps the switch to get the MAC address of this device and update its MAC address table. The switch also uses CDP to communicate with devices connected to its ports, to negotiate connection parameters and other settings. Look over *"Cisco Discovery Protocol (CDP)."*

Chapter 7: Troubleshooting a Switch Using Cisco IOS

Exam Objectives

- ✓ Describing Cisco hardware and software products and their purposes
- ✓ Gathering troubleshooting information about a Cisco switch
- ✓ Inspecting Cisco switch memory contents
- ✓ Using the show tech-support and the debug commands to inspect a Cisco switch
- ✓ Troubleshooting switch connectivity
- ✓ Gathering troubleshooting information about your network
- ✓ Discovering your network using Cisco Discovery Protocol (CDP)
- ✓ Troubleshooting the startup configuration
- ✓ Troubleshooting the running configuration

Troubleshooting a Cisco switch is very similar to troubleshooting a Cisco router. Most IOS troubleshooting commands, such as the debug command, are the same for switches and routers, but the output differs in some cases. Most troubleshooting tools are available for both switches and routers, such as the ping and trace route tools, or the Cisco Network Assistant (CNA). Some troubleshooting tools are only available for routers, such as the Cisco Router and Security Device Manager (SDM) and SDM Express. Other troubleshooting tools are only available for switches, such as the Cisco Device Manager. You learn about switch troubleshooting tools and IOS commands in this chapter.

Troubleshooting Cisco Switches

Read this section to find out how to troubleshoot your Cisco switch using IOS commands and graphical user interface tools, such as the Cisco Device Manager and the Cisco Network Assistant (CNA).

Gathering information about the switch

There are several commands that allow you to gather information about your switch. You review information gathering and troubleshooting commands in this section and find out about some additional commands. You can also use the IOS graphical user interfaces (GUI tools) to gather information about your switch.

Obtaining the IOS version

The first step of the troubleshooting process is finding the particular version of the Cisco IOS that operates your switch. There are several ways to find out the IOS version on your switch:

- ◆ The boot process output messages display the IOS version.
- ◆ The IOS image file name contains the IOS version.
- ◆ The `show version` IOS command displays the IOS version.
- ◆ GUI tools display the IOS version.

Obtaining IOS version from boot output

The Cisco IOS version and release numbers are displayed during the startup process. Here is an example of the boot loader output displayed on the console of a Cisco Catalyst 2960 switch. The lines in bold font show the version and release numbers of the IOS installed on this switch.

```
Loading "flash:c2960-lanlite-mz.122-37.EY/c2960-lanlite-mz.122-37.EY.bin"...@@@
@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@
@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@
File "flash:c2960-lanlite-mz.122-37.EY/c2960-lanlite-mz.122-37.EY.bin"
uncompressed and installed, entry point: 0x3000
executing...
```

```
Restricted Rights Legend
Use, duplication, or disclosure by the Government is
subject to restrictions as set forth in subparagraph
(c) of the Commercial Computer Software - Restricted
Rights clause at FAR sec. 52.227-19 and subparagraph
(c) (1) (ii) of the Rights in Technical Data and Computer
Software clause at DFARS sec. 252.227-7013.
```

```
cisco Systems, Inc.
170 West Tasman Drive
San Jose, California 95134-1706
```

```
Cisco IOS Software, C2960 Software (C2960-LANLITE-M), Version 12.2(37)EY,
RELEASE SOFTWARE (fc2)
Copyright (c) 1986-2007 by Cisco Systems, Inc.
Compiled Thu 28-Jun-07 18:07 by antonino
Image text-base: 0x00003000, data-base: 0x00D00000
```

Obtaining IOS version from the IOS image file name

Observe that the Cisco IOS c2960-lanlite-mz.122-37.EY.bin image file is loaded in the preceding example. You can see the Cisco IOS version and release numbers in the filename of the Cisco IOS image file. The image file is named according to the:

- ◆ Model of the Cisco switch: Catalyst 2960 (c2960)
- ◆ Edition of the IOS (lanlite)
- ◆ Version, release of the IOS: Version 12.2 release 37.EY (122-37.EY)

Obtaining IOS version using the Show version command

You can use the show version IOS command to verify the IOS version and release numbers of the IOS running on a Cisco switch. You run the show version IOS command in privileged EXEC mode, as shown in the following example. The lines in bold show:

- ◆ The version of the Cisco IOS
- ◆ The version of the Boot loader (bootstrap) program
- ◆ The Cisco IOS image file that was loaded

```
SW1>en
Password:
SW1#sh version
Cisco IOS Software, C2960 Software (C2960-LANLITE-M), Version 12.2(37)EY, RELEASE
SOFTWARE (fc2)
Copyright (c) 1986-2007 by Cisco Systems, Inc.
Compiled Thu 28-Jun-07 18:07 by antonino
Image text-base: 0x00003000, data-base: 0x00D00000

ROM: Bootstrap program is C2960 boot loader
BOOTLDR: C2960 Boot Loader (C2960-HBOOT-M) Version 12.2(44r)SE1, RELEASE SOFTWARE
(fc2)

SW1 uptime is 13 minutes
System returned to ROM by power-on
System image file is „flash:c2960-lanlite-mz.122-37.EY/c2960-lanlite-mz.122-37.
EY.bin“

cisco WS-C2960-24-S (PowerPC405) processor (revision C0) with 61440K/4088K bytes
of memory.

[... some output cut ...]

1 Virtual Ethernet interface
24 FastEthernet interfaces
The password-recovery mechanism is enabled.

64K bytes of flash-simulated non-volatile configuration memory.

[... some output cut ...]
```

Switch	Ports	Model	SW Version	SW Image
-----	-----	-----	-----	-----
* 1	24	WS-C2960-24-S	12.2 (37) EY	C2960-LANLITE-M

Configuration register is 0xF

Obtaining IOS version using the GUI tools

Cisco GUI tools also display the IOS version. The Cisco Device Manager and the Cisco Network Assistant (CNA) show the version of the IOS operating system running on your switch.

Figure 7-1 illustrates the Cisco Device Manager Dashboard screen showing the version of the IOS running on a Cisco Catalyst 2960 switch. The IOS version shows in the *Switch Information* frame.




Figure 7-1:
Cisco
Device
Manager
Dashboard.

Figure 7-2 illustrates the Cisco Device Manager Software Upgrade screen showing the version of the IOS running on a Cisco Catalyst 2960 switch. The Software Upgrade Web form allows you to upgrade the Cisco Internetworking Operating System (IOS) installed on your switch. Here is the procedure you need to follow to upgrade the IOS on your switch:

- ◆ Click on the www.cisco.com/public/sw-center/ URL link to access the latest version of the Cisco IOS on Cisco's Web site.
- ◆ Download the Cisco IOS image (tar file) to your computer hard drive.
- ◆ Click on *Choose File* and browse to the location where you saved the Cisco IOS image (tar file) on your computer hard drive.
- ◆ Select the Cisco IOS image (tar file) on your computer hard drive.
- ◆ Click *Upgrade* to upgrade the Cisco IOS on your switch with the Cisco IOS image file you just downloaded to your computer hard drive.




Figure 7-2:
Cisco
Device
Manager
Software
Upgrade.

Figure 7-3 illustrates a screenshot of the Cisco Network Assistant (CNA) Device Properties showing the version of the IOS running on a Cisco Catalyst 2960 switch. You access the *Device Properties* in CNA by right-clicking the device in the *Topology View* or in *Front Panel View* and selecting the *Properties...* menu item.




Figure 7-3:
Cisco
Network
Assistant
Device
Properties.

Inspect Cisco switch configuration and memory contents

During the troubleshooting process you need to know the particular configuration of your switch. You also need to know the configuration data stored in the memory of the switch. There are several methods to inspect the configuration of your switch and the data stored in the switch memory:

- ◆ Use the `show flash` command to inspect flash memory contents.
- ◆ Use the `show running-config` command to inspect the current running configuration in RAM.
- ◆ Use the `show startup-config` command to inspect the startup configurations in NVRAM.
- ◆ Use the `show tech-support` command to inspect all technical parameters on your switch.
- ◆ Use the Cisco IOS File System (IFS) commands to inspect RAM, NVRAM and flash contents.

Let's review each of these tools.

Inspecting Cisco switch memory contents using the show flash command

You can verify the following technical aspects of your switch using the show flash IOS command:

- ◆ **Presence of a Cisco IOS image file:** You should have at least one Cisco IOS image file in flash memory on your switch.
- ◆ **Presence of config.text:** This is the startup-configuration file. When the switch is started up, the flash memory area that stores config.text is mapped as NVRAM and the config.text file is mapped as the nvram:startup-config file by the Cisco IOS File System (IFS). Hence, if you do not have a config.text file in flash memory, the switch does not have a startup configuration: It starts in Setup mode.
- ◆ **Presence of private-config.text:** This file is used to store secured configuration data such as cryptographic encryption keys used for SSH. This file should be there in flash memory when config.text is there.
- ◆ **Presence of config.text.renamed:** This file is created whenever you, or the switch IOS, erase the startup configuration file. Recall that you can erase the startup configuration file using the erase startup-config IOS command. The switch erases the startup configuration file whenever you reset the switch using the Mode button. In both cases, the startup configuration file is not really removed from flash memory, but rather, renamed to config.text.renamed. This ensures that the switch IOS will not find config.text when it reboots, but it provides a safety net in case you erase the startup configuration file by mistake.
- ◆ **Presence of private-config.text.renamed:** This file is a hidden version of the private-config.text file. The file private-config.text is renamed to private-config.text.renamed whenever you, or the switch, remove the startup configuration file.
- ◆ **Presence of vlan.dat:** This file stores VLAN data on your switch. The IOS creates some default VLANs that do not require the presence of the vlan.dat file. However, if you create additional VLANs, the vlan.dat file should be there.
- ◆ **Free space available in flash memory:** The show flash IOS command displays the free space available in flash memory. It is useful to know how much free space you have in flash memory whenever you want to upload several Cisco IOS image files to the flash memory of your switch. For example, you may want to upload a Cisco IOS image that supports encryption to enable SSH access on your switch. Before you do that, you need to verify that you have enough room in flash memory for both the current Cisco IOS image file and the new one.

You can run the `show flash` IOS command both in user EXEC mode and in privileged EXEC mode. The following example shows the contents in flash memory of a Cisco Catalyst 2960 switch.

```
SW1>show flash
Directory of flash:/
 2 -rwx 1993 Mar 1 1993 00:30:29 +00:00 private-config.text.renamed
 3 -rwx 1398 Mar 1 1993 00:30:29 +00:00 config.text.renamed
 4 -rwx 1398 Mar 1 1993 00:30:29 +00:00 config.text
 5 drwx 512 Mar 1 1993 00:11:25 +00:00 c2960-lanlite-mz.122-37.EY
539 -rwx Mar 1 1993 00:30:29 +00:00 private-config.text

27998208 bytes total (19307520 bytes free)
```

Observe the Cisco IOS image file named `c2960-lanlite-mz.122-37.EY`. This is the IOS image file that is loaded by default when the switch is started up.

Note that you have both a `config.text` file and a `config.text.private` file in flash memory on this switch. This shows that the switch was reset at some point: This is when the `.renamed` files were created. At that point there was no `config.text` file and no `private-config.text` file in flash memory, just the `.renamed` files. Then, a new startup configuration was created. At that point the `config.text` and the `private-config.text` files were created. However, the `.renamed` files were not removed. This is why you see both the `config.text`, `private-config.text` files and the corresponding `.renamed` files in flash memory.

Inspecting Cisco switch current configuration using show running-config

You can use the `show running-config` IOS command to verify the current running configuration on your switch. You run the `show running-config` command in privileged EXEC mode. The example shows the running configuration of a Cisco Catalyst 2960 switch.

```
SW1>enable
Password:
SW1#show running-config
Building configuration...

Current configuration : 1398 bytes
!
version 12.2
no service pad
service timestamps debug uptime
service timestamps log datetime
no service password-encryption
service sequence-numbers
!
hostname SW1
!
enable secret 5 $1$/nQ0$AJ.w7nP4fVgH44NU4gzkZ1
```

```

enable password my_priv_password
!
no aaa new-model
system mtu routing 1500
ip subnet-zero
!
!
!
no file verify auto
spanning-tree mode pvst
spanning-tree extend system-id
!
vlan internal allocation policy ascending
!
interface FastEthernet0/1
!
[... some output cut ...]

interface FastEthernet0/24
!
interface Vlan1
 ip address 192.168.75.10 255.255.255.0
 no ip route-cache
!
ip default-gateway 192.168.75.1
ip http server
!
control-plane
!
!
line con 0
 exec-timeout 0 0
line vty 0 4
 password my_telnet_password
 login
line vty 5 15
 password my_telnet_password
 login
!
end
SW1#disable
SW1>

```

Inspecting Cisco switch startup configuration using show startup-config

You can use the `show startup-config` IOS command to verify the startup configuration on your switch. You run the `show startup-config` command in privileged EXEC mode. The following example shows the startup configuration of a Cisco Catalyst 2960 switch.

```

SW1>enable
Password:
SW1#show startup-config
Using 1398 out of 65536 bytes
!
version 12.2
no service pad
service timestamps debug uptime

```

```
service timestamps log datetime
no service password-encryption
service sequence-numbers
!
hostname SW1
!
enable secret 5 $1$/nQ0$AJ.w7nP4fVgH44NU4gzkZ1
enable password my_priv_password
!
no aaa new-model
system mtu routing 1500
ip subnet-zero
!
!
!
no file verify auto
spanning-tree mode pvst
spanning-tree extend system-id
!
vlan internal allocation policy ascending
!
interface FastEthernet0/1
!
[... some output cut ...]

interface FastEthernet0/24

!
interface Vlan1
 ip address 192.168.75.10 255.255.255.0
 no ip route-cache
!
 ip default-gateway 192.168.75.1
 ip http server
!
 control-plane
!
!
 line con 0
   exec-timeout 0 0
 line vty 0 4
   password my_telnet_password
   login
 line vty 5 15
   password my_telnet_password
   login
!
end
SW1#disable
SW1>
```

Inspecting RAM, NVRAM and flash memory contents using the Cisco IFS

You can use Cisco IOS File System (IFS) commands to verify the contents of the RAM, NVRAM, and flash memory areas on your switch. You run Cisco IFS command in privileged EXEC mode.

To see the Cisco IFS directories available on your switch, execute the `dir all file-systems` Cisco IFS command, as shown in the following example:

```
SW1>enable
Password:
SW1# dir all-file-systems
Directory of nvram:/

 27  -rw-   0   <no date>  startup-config
 28  ----   0   <no date>  private-config

29688 bytes total (29636 bytes free)
Directory of system:/

 10  drwx   0               <no date>  its
  2  dr-x   0               <no date>  memory
  1  -rw-  582             <no date>  running-config
  9  dr-x   0               <no date>  vfiles

No space information available
Directory of flash:/

 1  -rw- 15572992 <no date> c2600-ik9o3s3-mz.123-1a.bin

16777216 bytes total (1204160 bytes free)
SW1#disable
SW1>
```

Observe the three main directories on your Cisco switch:

- ◆ **nvram:** This is the Non-Volatile Random Access Memory (NVRAM), storing the `startup-config` file, and the `private-config` file.
- ◆ **system:** This is the Random Access Memory (RAM), containing the `running-config` file.
- ◆ **flash:** This is the flash memory containing the Cisco IOS operating system image. This is the Cisco IOS image that the bootstrap program loads in RAM when the switch boots up.

Observe how the Cisco IFS maps a portion of flash memory as NVRAM during the switch boot process. When you run the `show flash` command, you see the `config.text` and the `private-config.text` files in the listing. These files correspond to the `nvram:startup-config` file and to the `nvram:private-config`, respectively. The `dir all file-systems` IFS command only lists the Cisco IOS image file in flash, because the other files stored in flash memory are mapped to the `nvram:` directory.

You can list the contents of any of these main directories (file systems) individually using the `dir` Cisco IFS command. You can change your current working directory to any of the main directories using the `cd` Cisco IFS command.

For example, to list the contents of NVRAM specifically, execute the following command:

```
SW1# dir nvram:
Directory of nvram:/
 27  -rw-   0          <no date>  startup-config
 28  ----   0          <no date>  private-config
29688 bytes total (29636 bytes free)
```

Now, change directory to the RAM directory (system) and look at the files and directories available there:

```
SW1# cd system:
SW1# dir
Directory of system:/
 10  drwx   0  <no date>  its
  2  dr-x   0  <no date>  memory
  1  -rw- 582  <no date>  running-config
  9  dr-x  0          <no date>  vfiles
```

Finally, to list the contents of flash memory, as seen by IFS, execute the following command:

```
SW1# dir flash:
Directory of flash:/
 1  -rw- 15572992 <no date>  c2600-ik9o3s3-mz.123-1a.bin
16777216 bytes total (1204160 bytes free)
```

Inspecting all technical parameters on your switch using show tech-support

The `show tech-support` IOS command collects and displays information about all technical parameters on a switch. You run the `show tech-support` command in privileged EXEC mode.

The `show tech-support` command executes IOS commands to collect and display as much troubleshooting data as possible on your switch, such as these:

- ◆ `show clock`: Displays the switch system date and time.
- ◆ `show interfaces`: Displays configuration data about interfaces on your switch. This command displays IP configuration data, duplex mode, speed, media type, queue configuration and input and output rates. This data is different from the data returned by the next command (`show controllers`). Here is the difference: The `show interfaces` command shows configuration data, whereas the `show controllers` command shows real time operating statistics collected and summarized by the IOS. Also, note that `show interfaces` shows configuration data about both the VLAN (port trunk) logical interfaces and the underlying physical interfaces (ports on the switch). The `show controllers` command only shows statistics about physical interfaces (ports on the switch).

- ◆ `show logging`: Displays messages logged in the switch system log file. This command is very useful, because it shows what happened since the switch powered up.
- ◆ `show env all`: Displays environment data about your switch, such as temperature, fan status, and power supply status.
- ◆ `show interfaces status`: Displays the state of interfaces on your switch.
- ◆ `show interfaces status err-disabled`: Displays the state of interfaces that transitioned to err-disabled state. An interface transitions to err-disabled state if it was configured to be up and connected, but something prevented it from transitioning to the “up” state.
- ◆ `show interfaces switchport`: Displays VLAN and port trunk data about the interfaces on your switch.
- ◆ `show interfaces trunk`: Displays data about port trunks on your switch, if you created any.
- ◆ `show vlan`: Displays data about VLANs on your switch.
- ◆ `show mac-address-table`: Displays data about the MAC address table on your switch. This is a good place to look if you see an excessive number of Multicast and Broadcast frames reported by the `show controllers` command. The MAC address table should contain the MAC addresses of all computer hosts in your LAN. If it doesn’t, then, nodes that need to send frames to computers that are not known by your switch, need to send multicast or broadcast frames in the network to find the MAC address of the target. You can ping target hosts from the switch to manually help the switch complete its MAC address table. This would cause the switch to register in its MAC address table the MAC address of the hosts you ping.
- ◆ `show spanning-tree`: Displays data about STP (Spanning Tree Protocol) on your switch.
- ◆ `show etherchannel summary`: Displays data about EtherChannels on your switch.

Some commands do not return any output, because the technical aspect they would normally show does not apply to the switch being inspected.

Inspect Cisco switch logs and system messages

System messages and log files are the most useful troubleshooting tools. Cisco switches register events and errors in the system log files as they happen. Each event or error message is logged with a timestamp. Hence, you can look at the log files to know what happened since the switch started. Cisco provides several methods to inspect system event messages and alerts in log files:

- ◆ Using the `show logging` IOS command
- ◆ Using the Cisco Network Assistant (CNA)

Let's review each of these tools.

Inspecting the switch logs and system messages with show logging

You can run the `show logging` IOS command both in user EXEC mode and in privileged EXEC mode. The following example shows the results returned by `show logging` on a Cisco Catalyst 2960 switch.

```
SW1>show logging
Syslog logging: enabled (0 messages dropped, 0 messages rate-limited, 0 flushes,
 0 overruns, xml disabled, filtering disabled)
Console logging: level debugging, 7 messages logged, xml disabled, filtering
disabled
Monitor logging: level debugging, 0 messages logged, xml disabled, filtering
disabled
Buffer logging: level debugging, 7 messages logged, xml disabled, filtering
disabled
Exception Logging: size (4096 bytes)
Count and timestamp logging messages: disabled
File logging: disabled
Trap logging: level informational, 10 message lines logged

Log Buffer (4096 bytes):

00:00:29: %LINEPROTO-5-UPDOWN: Line protocol on Interface Vlan1, changed state to
down
00:00:29: %SPANTREE-5-EXTENDED_SYSID: Extended SysId enabled for type vlan
000003: *Mar  1 00:00:31: %SYS-5-CONFIG_I: Configured from memory by console
000004: *Mar  1 00:00:31: %SYS-5-RESTART: System restarted --
Cisco IOS Software, C2960 Software (C2960-LANLITE-M), Version 12.2(37)EY, RELEASE
SOFTWARE (fc2)
Copyright (c) 1986-2007 by Cisco Systems, Inc.
Compiled Thu 28-Jun-07 18:07 by antonino
000005: *Mar  1 00:00:33: %LINK-3-UPDOWN: Interface FastEthernet0/1, changed
state to up
000006: *Mar  1 00:00:34: %LINEPROTO-5-UPDOWN: Line protocol on Interface
FastEthernet0/1, changed state to up
000007: *Mar  1 00:01:04: %LINEPROTO-5-UPDOWN: Line protocol on Interface Vlan1,
changed state to up
```

Observe log messages 5 and 6 in bold font type. These messages show that the first physical interface (FastEthernet0/1) was started up (connected). Message 7 shows that the VLAN1 logical interface also started up.

Managing the logging system

Cisco switches enable logging by default. All log messages are sent to the console and to the internal log buffer by default. You can use the `logging` IOS command in privileged EXEC configuration mode to control logging settings on a Cisco switch. For example:

- ◆ `logging on`: Enables logging
- ◆ `no logging on`: Disables logging

The `logging` command also allows you to control where system event messages and alerts are logged. Log messages can be sent to the:

- ◆ **Console:** All system event messages and alerts are displayed on the switch console by default. The `logging console [<level>]` command enables logging to the console and sets the logging level. The *logging level* defines which messages are logged. Logging levels are explained further in this section. The `no logging console` command disables logging to the console.
- ◆ **Internal log buffer:** All system event messages and alerts are written to the internal log buffer by default. The `logging buffered [<level>]` command enables logging to the internal log buffer. The `no logging buffered` command disables logging to the internal log buffer.
- ◆ **VTY session (Telnet/SSH):** System event messages and alerts are not displayed in virtual terminal (VTY) sessions by default. The `logging monitor [<level>]` command enables logging to Telnet/SSH sessions. The `no logging monitor` command disables logging to Telnet/SSH sessions.
- ◆ **File in flash memory:** As mentioned earlier, system event messages and alerts are written to the internal log buffer by default. You can configure the logging system to write system event messages and alerts to a specific file in flash memory using the `logging file [flash:<filename>]` command. The `no logging file` command disables logging to a specific file in flash memory.
- ◆ **Syslog server:** You can configure the SYSLOG protocol on Cisco switches to log system events and alerts to an external computer system. As mentioned earlier, system event messages and alerts are written to the internal log buffer by default. However, this buffer can contain only a limited number of messages, due to the limited amount of memory on a Cisco switch. Configuring a Syslog server allows you to leverage larger hard drive storage available on a management computer system and to centralize system events and alerts logging from several switches onto that single computer host. The `logging trap [<level>]` command enables logging to a Syslog server. The `no logging trap` command disables logging to a Syslog server. The `logging [<IP address>]` command identifies the IP address of the destination Syslog server where system events and alerts are logged. You also need to configure the SYSLOG protocol on the Syslog server computer. You can log system events and alerts to more than one Syslog server by executing the `logging [<IP address>]` command with the IP address of each destination Syslog server. To remove a Syslog server from the logging destination list, execute `no logging [<IP address>]` with the IP address of the Syslog server to be removed from the list.

Cisco logging levels are designed to control the verbosity of the logs. Table 7-1 illustrates the seven log levels.

Table 7-1		Logging Levels	
<i>Level</i>	<i>Severity</i>	<i>Description</i>	<i>Syslog Type</i>
Emergencies	0	Switch is unusable	LOG_EMERG
Alerts	1	Switch requires immediate attention	LOG_ALERT
Critical	2	Switch experienced critical condition	LOG_CRITICAL
Errors	3	Switch experienced error condition	LOG_ERROR
Warnings	4	Switch experienced warning condition	LOG_WARNING
Notifications	5	Switch experienced normal significant condition	LOG_NOTICE
Informational	6	Information message	LOG_INFO
Debugging	7	Debugging message	LOG_DEBUG

Lower severity number means higher severity. This is a bit confusing. Think of it this way:

- ◆ Severity 0: You have 0 functionality: Switch is unusable due to fatal condition.
- ◆ Severity 7: You have maximum functionality and log verbosity.

So, for example, if you set logging level 3 (Errors) on the virtual terminal lines (Telnet/SSH sessions) using the command, `logging monitor 3` or `logging monitor errors` the switch only logs messages with severity 3, 2, 1 or 0 out on the virtual terminal lines (Telnet/SSH sessions). It only logs errors, critical events, alerts, and emergency events.

Understanding the output of the show logging command

Observe the first lines of output of the `show logging` command:

```
SW1>show logging
Syslog logging: enabled (0 messages dropped, 0 messages rate-limited, 0 flushes,
0 overruns, xml disabled, filtering disabled)
Console logging: level debugging, 7 messages logged, xml disabled, filtering
disabled
Monitor logging: level debugging, 0 messages logged, xml disabled, filtering
disabled
Buffer logging: level debugging, 7 messages logged, xml disabled, filtering
disabled
Exception Logging: size (4096 bytes)
Count and timestamp logging messages: disabled
File logging: disabled
Trap logging: level informational, 10 message lines logged
```

```
Log Buffer (4096 bytes):
[... some output cut ...]
```

These lines provide important information about the configuration of the logging system on your switch.

- ◆ **Syslog logging:** Shows the configuration of the Cisco *system logging process*. This refers to the logging system as a whole on a Cisco switch, not to the SYSLOG protocol and destination Syslog servers. The `Trap logging` field in the `show logging` output refers to the SYSLOG protocol and Syslog server configuration.
- ◆ **Console logging:** Shows configuration of logging to console.
- ◆ **Monitor logging:** Shows configuration of logging to virtual terminal lines (Telnet/SSH session).
- ◆ **Buffer logging:** Shows the configuration of logging to internal log buffer.
- ◆ **Exception logging:** Shows size of exception logging buffer. This is the size of the internal log buffer.
- ◆ **Count and timestamp logging:** When enabled this option counts every log message and timestamps last occurrence of each message. This option is disabled by default: Messages are not counted, they are logged and time stamped as events occur. If an event occurs more than once, you see one message logged for each occurrence of the event.
- ◆ **File logging:** Shows configuration of logging to specific file in flash memory.
- ◆ **Trap logging:** Shows configuration of logging to Syslog server(s).
- ◆ **Log buffer:** Shows size of internal log buffer.



The term *syslog* refers to two different concepts on Cisco switches:

- ◆ The *Cisco System Logging Process* is the process that manages logging on Cisco switches. Cisco sometimes abbreviates *System Logging Process* to *syslog*. You may also see the term SEMs (system error messages). The term SEMs refers to system error messages generated by the Cisco *System Logging Process*. The `Syslog logging` field in the output of `show logging` refers to the Cisco *System Logging Process*.
- ◆ SYSLOG is a protocol that allows computer systems to exchange log messages. Cisco supports the SYSLOG protocol to send log messages to a Syslog server to leverage larger hard drive storage available on a management computer system and to centralize system events and alerts from several switches onto that single computer host. The `Trap logging` field in the output of `show logging` refers to the SYSLOG protocol configuration on the Cisco switch.


You can distinguish between the two meanings of *syslog* based on the context of the discussion.

Inspecting the switch logs and system messages with CNA

You learned so far how to inspect system messages in log files using the `show logging` Cisco IOS command. Read this section to find out how to inspect system messages in log files using the Cisco Network Assistant (CNA).

Figure 7-4 illustrates the System Messages Monitoring tool in the Cisco Network Assistant (CNA) GUI.

Figure 7-4:
Cisco
Network
Assistant
(CNA) –
System
Messages.



To access the System Messages of a particular switch you need to:

- ◆ Select the switch you want to inspect in *Topology* or *Front Panel* view.
- ◆ Select the *Monitoring* tab in the menu.
- ◆ Click *System Messages...*

Observe that each message line displays:

- ◆ **Severity of the message:** One of the 7 severity levels indicating whether this is warning, an error, a critical condition, or just an informational message, or a debugging message.
- ◆ **Timestamp of the event:** Data and time when the event occurred.
- ◆ **Event description:** Short description of what happened.
- ◆ **Device:** Host name of the device where the event occurred.

If you click any message line, CNA displays event details in the message details text box, under the message list:

- ◆ **Description:** This is the event description displayed on each message line.
- ◆ **Explanation:** This explains the event providing more details about what happened.
- ◆ **Recommended Action:** This displays a recommended action for events that require the network administrator to remediate a situation that caused the event. Note that in Figure 7-4, the recommended action for the selected message line is LOG_STD_NO_ACTION. This is normal, because the message selected is simply a notification.

Troubleshooting switch connectivity

To troubleshoot connectivity problems on a Cisco switch, you need to:

- ◆ Verify cables and patch panels involved in the link.
- ◆ Verify the status of ports on your switch.
- ◆ Test connectivity using `ping` and `tracert` tools.

Verify cables and patch panels

To verify connectivity between your switch and another device in the network, you first need to verify that you have a valid physical connection between the switch and that device. You need to:

- ◆ Ensure cable connectors are not physically damaged. Pay close attention to the clip securing the connector in the switch port. Look also at the contact lids of the connector to ensure they are not physically damaged and to ensure that they are clean.
- ◆ Ensure the contact lids of the switch port are not physically damaged and that they are clean.

- ◆ Ensure cable connectors are properly inserted in the switch port and in the device at the other end of the cable.
- ◆ Ensure cables are not physically damaged. Pay close attention to the cable section exiting the connector; make sure the cable sleeve is not damaged in any way at the junction point in the connector.
- ◆ Use valid Ethernet cabling: Recall that Ethernet comes in several versions. Each version recommends a specific Unshielded Twisted Pair (UTP) cable. Most network installations today use at least UTP Category 5e (shortly known as Cat5e) cabling. Cat5e cables are well suited for most Ethernet standards. There are several tools that allow you to verify connectivity from one end of the cable to the other. Those tools are beyond CCNA.



When you suspect a broken cable segment (that is, there's no connectivity within that segment), you can replace the broken cable with a cable you know is healthy. This helps you eliminate a potential connectivity trouble. If the switch still experiences connectivity problems, you need to continue the troubleshooting process.

- ◆ Ensure the link does not extend beyond the distance supported by the Ethernet standard and type of cabling used to connect your switch.
- ◆ Ensure patch panels are operating within specification.
- ◆ Ensure distance added by the patch panels does not extend the link beyond the distance supported by the Ethernet standard and type of cabling used to connect your switch.

Verify the ports status on your switch

Once you eliminated cabling from the list of culprits, you verify the status of the ports on your switch. You can either use IOS commands or use Cisco GUI tools to verify port status on your switch.

Verify the ports status on your switch with IOS commands

You can use the following IOS commands to verify port status on your switch:

- ◆ `show interfaces status`: Displays the state of interfaces on your switch.
- ◆ `show interfaces status err-disabled`: Displays the state of interfaces that transitioned to err-disabled state. An interface transitions to err-disabled state if it was configured to be up and connected, but something prevented it from transitioning to the "up" state.
- ◆ `show interfaces switchport`: Displays VLAN and port trunk data about the interfaces on your switch.

- ◆ `show interfaces trunk`: Displays data about port trunks on your switch, if you created any.
- ◆ `show vlan`: Displays data about VLANs on your switch.
- ◆ `show ip interface`: Displays IP information about interfaces on your switch.
- ◆ `show logging`: Displays system messages logged in the system log since the switch started up. These system messages include port transitions from down state to up state and vice versa. In some cases, you can also see a port transition to `err-disabled` state if there were any errors preventing the port to open up.

Here is an example of `show logging` command:

```
SW1>enable
Password:
SW1#show logging
Syslog logging: enabled (0 messages dropped, 0 messages rate-limited, 0 flushes,
0 overruns, xml disabled, filtering disabled)
Console logging: level debugging, 7 messages logged, xml disabled, filtering
disabled
Monitor logging: level debugging, 0 messages logged, xml disabled, filtering
disabled
Buffer logging: level debugging, 7 messages logged, xml disabled, filtering
disabled
Exception Logging: size (4096 bytes)
Count and timestamp logging messages: disabled
File logging: disabled
Trap logging: level informational, 10 message lines logged

Log Buffer (4096 bytes):

00:00:29: %LINEPROTO-5-UPDOWN: Line protocol on Interface Vlan1, changed state to
down
00:00:29: %SPANTREE-5-EXTENDED_SYSID: Extended SysId enabled for type vlan
000003: *Mar 1 00:00:31: %SYS-5-CONFIG_I: Configured from memory by console
000004: *Mar 1 00:00:31: %SYS-5-RESTART: System restarted --
Cisco IOS Software, C2960 Software (C2960-LANLITE-M), Version 12.2(37)EY, RELEASE
SOFTWARE (fc2)
Copyright (c) 1986-2007 by Cisco Systems, Inc.
Compiled Thu 28-Jun-07 18:07 by antonino
000005: *Mar 1 00:00:33: %LINK-3-UPDOWN: Interface FastEthernet0/1, changed
state to up
000006: *Mar 1 00:00:34: %LINEPROTO-5-UPDOWN: Line protocol on Interface
FastEthernet0/1, changed state to up
000007: *Mar 1 00:01:04: %LINEPROTO-5-UPDOWN: Line protocol on Interface Vlan1,
changed state to up
SW1#disable
SW1>
```

Observe log messages 5 and 6 in bold font type. These messages show that the first physical interface (FastEthernet0/1) was started up (connected). Message 7 shows that the VLAN1 logical interface also started up.

Since physical interface FastEthernet0/1 (fa0/1) is “powering” logical interface VLAN 1, as soon as physical interface fa0/1 starts up, logical interface VLAN 1 also automatically starts. If physical interface fa0/1 was shut down, logical interface VLAN 1 also would shut down, unless there are other physical interfaces “powering” it up.

Here is an example of the `show interface status` command:

```
SW1>enable
Password:
SW1#show interfaces status
Port Name   Status      Vlan    Duplex  Speed Type
Fa0/1       connected   1       a-full  a-100  10/100BaseTX
Fa0/2       notconnect  1       auto    auto   10/100BaseTX
Fa0/3       notconnect  1       auto    auto   10/100BaseTX

[... some output cut ...]

Fa0/24      notconnect  1       auto    auto   10/100BaseTX
SW1#disable
SW1>
```

Observe physical interface (port) FastEthernet0/1 (fa0/1) is:

- ◆ Connected
- ◆ Member of VLAN 1
- ◆ Running in Full Duplex Mode
- ◆ Running at 100Mbps (maximum bandwidth for FastEthernet)
- ◆ Type: 10/100BaseTx (compliant to FastEthernet standard)

All other ports are currently in `notconnect` state: Ports 2 to 24 are not connected right now.

Here is an example of the `show interface switchport` command:

```
SW1>enable
Password:
SW1#show interfaces switchport
Name: Fa0/1
Switchport: Enabled
Administrative Mode: dynamic auto
Operational Mode: static access
Administrative Trunking Encapsulation: dot1q
Operational Trunking Encapsulation: native
Negotiation of Trunking: On
Access Mode VLAN: 1 (default)
Trunking Native Mode VLAN: 1 (default)
Administrative Native VLAN tagging: enabled
Voice VLAN: none
Administrative private-vlan host-association: none
```

```
Administrative private-vlan mapping: none
Administrative private-vlan trunk native VLAN: none
```

```
[... some output cut ...]
```

```
Name: Fa0/24
Switchport: Enabled
Administrative Mode: dynamic auto
Operational Mode: down
Administrative Trunking Encapsulation: dot1q
Negotiation of Trunking: On
Access Mode VLAN: 1 (default)
Trunking Native Mode VLAN: 1 (default)
Administrative Native VLAN tagging: enabled
Voice VLAN: none
Administrative private-vlan host-association: none
Administrative private-vlan mapping: none
Administrative private-vlan trunk native VLAN: none
[... some output cut ...]
```

```
SW1#disable
SW1>
```

Compare the text in bold font type between physical interface FastEthernet0/1 (fa0/1) and physical interface FastEthernet0/24 (fa0/24):

Both physical interfaces are enabled: If you plug a network cable in the corresponding fa0/1 or fa0/24 port, you get connectivity through that interface.

Both physical interfaces are member of VLAN 1 (Access Mode VLAN). Out of factory, all Cisco switches have a few VLANs pre-created. VLAN 1 is one of them. These pre-created VLANs are reserved for special purposes. VLAN 1 contains all switch ports by default.

The Operational mode of physical interface fa0/1 is static access. This means that the physical interface (port) fa0/1 is currently connected with a network cable to another device. You see static (or dynamic) access operational mode, as soon as you connect a cable into a physical interface (port) that is enabled and member of a VLAN that is configured for IP networks.

The Operational mode of physical interface fa0/24 is down. This means that the physical interface (port) fa0/24 is currently not connected with a network cable to another device.

When troubleshooting connectivity, the interface (port) should show up as fa0/1, not as fa0/24. If the operational mode is down, as with fa0/24, either there is no cable connected, or there is a connectivity problem with the cable or patch panels. If you know for sure that a cable is connected in the fa0/24 port and that there are no connectivity problems within the cable or the patch panels, you need to verify the IP configuration of the VLAN where

physical interface (port) fa0/24 is assigned. In this case, fa0/24 is assigned to VLAN 1. Note that, in this case, you know for sure that VLAN 1 is configured properly for IP, because physical interface (port) fa0/1 is also assigned to VLAN 1: If there was a problem with VLAN 1, the operational state of fa0/1 would also be down, and you wouldn't have any connectivity over fa0/1, either. This tells you for sure that fa0/24 is down, because there is no cable connected into that port.

Here is an example of the `show vlan` command:

```
SW1>enable
Password:
SW1#show vlan
VLAN Name                Status    Ports
-----
1    default                active    Fa0/1, Fa0/2, Fa0/3, Fa0/4
                                           Fa0/5, Fa0/6, Fa0/7, Fa0/8
                                           Fa0/9, Fa0/10, Fa0/11, Fa0/12
                                           Fa0/13, Fa0/14, Fa0/15, Fa0/16
                                           Fa0/17, Fa0/18, Fa0/19, Fa0/20
                                           Fa0/21, Fa0/22, Fa0/23, Fa0/24
1002 fddi-default          act/unsup
1003 token-ring-default    act/unsup
1004 fddinet-default       act/unsup
1005 trnet-default         act/unsup
[... some output cut ...]

SW1#disable
SW1>
```

This switch has 24 ports, all members of VLAN 1, the default VLAN on Cisco switches. Observe that VLAN 1 status is `active`. This command helps you verify that:

- ◆ The physical interface (Fa0/1) was assigned to a VLAN on the switch.
- ◆ The VLAN status is active for the VLAN powered by the interface.

If you created additional VLANs, they would appear in this list.

Now, you need to verify that the VLAN powered by this interface has a valid IP address. You can do this by using the `show ip interface` IOS command:

```
SW1>show ip interface
Vlan1 is up, line protocol is up
  Internet address is 192.168.75.10/24
  Broadcast address is 255.255.255.255
  Address determined by non-volatile memory
  MTU is 1500 bytes
[... some output cut ...]
FastEthernet0/1 is up, line protocol is up
  Inbound access list is not set
```

```
FastEthernet0/2 is down, line protocol is down
  Inbound access list is not set
[... some output cut ...]
FastEthernet0/24 is down, line protocol is down
  Inbound access list is not set
```

The lines in bold font type show that VLAN 1 is assigned a valid IP address (192.168.75.10) and that the underlying physical interface (port Fa0/1) is up.

Verify the ports status on your switch with Cisco Device Manager

You can use the Cisco Device Manager Graphical User Interface to verify port status on your switch. You need to:

- ◆ Browse to the IP address of the switch.
- ◆ Expand the *Monitor* tab in the *Contents* section.
- ◆ Select *Port Status* in the *Monitor* tab.

Figure 7-5 illustrates the *Port Status* screen in the Cisco Device Manager.




Figure 7-5:
Cisco
Device
Manager –
Port Status.

Verify the ports status on your switch with Cisco Network Assistant (CNA)


You can use the Cisco Network Assistant (CNA) Graphical User Interface to verify port status on your switch.

First, you can verify the port is enabled and that its speed and media type are set to auto-negotiate:

- ◆ Start CNA and Logon to the switch using `level_15_access` (privileged EXEC mode).
- ◆ Select the *Configure* tab.
- ◆ Expand the *Ports* tab.
- ◆ Select *Port Settings...*

Figure 7-6 illustrates the *Port Settings* screen in CNA.

Figure 7-6:
Cisco
Network
Assistant –
Port
Settings.




Next, you verify that the physical interface was assigned to a VLAN on the switch and that the status of that VLAN is active.

- ◆ Start CNA and Logon to the switch using `level_15_access` (privileged EXEC mode).
- ◆ Select the *Configure* tab.
- ◆ Expand the *Switching* tab.
- ◆ Select *VLANs*...

Figure 7-7 illustrates the VLANs screen in CNA.

Figure 7-7:
Cisco
Network
Assistant –
VLANs.



Book III
Chapter 7


Troubleshooting
a Switch Using
Cisco IOS

Now, you verify that the VLAN powered by the physical interface has a valid IP address.

- ◆ Start CNA and Logon to the switch using `level_15_access` (privileged EXEC mode)
- ◆ Select the *Configure* tab.
- ◆ Expand the *Device Properties* tab.
- ◆ Select *IP Addresses*...

Figure 7-8 illustrates the IP Addresses screen in CNA.

Figure 7-8:
Cisco
Network
Assistant –
IP
Addresses.



You can look at Port Statistics to get some performance data about your physical interface. Note that this not only tells you whether you have connectivity over a particular physical interface, it also tells you how fast that interface is.

- ◆ Start CNA and Logon to the switch using `level_15_access` (privileged EXEC mode).
- ◆ Select the *Monitor* tab.
- ◆ Expand the *Reports* tab.
- ◆ Select *Port Statistics...*


Figure 7-9 illustrates the *Port Statistics* screen in CNA.

Testing connectivity with ping and trace route tools

Finally, you verify that the switch can reach a particular device over the TCP/IP network using the `ping` and `traceroute` IOS commands.

Let's look at each of these tools.

Figure 7-9:
Cisco
Network
Assistant –
Port
Statistics.



Ping tool

The *Ping* tool is an IOS command and CNA tool allowing you to send IP packets to an IP address to test end-to-end connectivity. This is very often the first command that you use to test connectivity between two nodes on a network when you suspect connectivity problems. The nodes can be network devices (switches, routers), computer hosts, IP phones, or any other device that supports TCP/IP network connectivity. Note that the target device must have a valid IP address to use the *ping* tool. You use the *ping* in two situations:

- ◆ Before the troubleshooting process to confirm that there is a connectivity problem. If the ping tool reports failure to send and receive IP packets, you need to troubleshoot connectivity as explained in the *Troubleshooting Switch Connectivity* section.
- ◆ After the troubleshooting process to confirm that connectivity has been restored and that you can send and receive IP packets between your switch and the remote node. If the ping tool successfully sends and receives IP packets between your switch and the remote node, you successfully restored connectivity.

Here is an example of the ping IOS command:

```
SW1>ping 192.168.75.100
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.168.75.100, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/8 ms
```

To use the ping tool in CNA:

- ◆ Start CNA and Logon to the switch using `level_15_access` (privileged EXEC mode).
- ◆ Select the *Troubleshooting* tab.
- ◆ Select *Ping and Trace...*
- ◆ Select *Ping* tool in the *Ping and Trace* screen.
- ◆ Enter the *Destination (IP Address/Hostname)*.
- ◆ Click *Start*.

Figure 7-10 illustrates the *Ping* screen in CNA.




Figure 7-10:
Cisco
Network
Assistant –
Ping.

Trace route tool

The *Trace Route* tool is an IOS command (`tracert`) allowing you to trace the IP route between two nodes on a TCP/IP network. The nodes can be network devices (switches, routers), computer hosts, IP phones, or any other device that supports TCP/IP network connectivity. Note that the target device must have a valid IP address to use the Layer 3 (Network) `tracert` tool.

You can also use the Layer 2 (data link) `tracert mac` tool to trace the Layer 2 (data link) route between two nodes on an Ethernet network. The Layer 2 (data link) `tracert mac` tool requires that:

- ◆ Both nodes have valid MAC addresses.
- ◆ CDP is configured properly and running on your switch.
- ◆ You execute the `tracert mac` IOS command in privileged EXEC mode.

The *Trace Route* tool is exposed in the CNA Graphical User Interface (GUI). To use the Layer 3 (network) `tracert` tool in CNA:

- ◆ Start CNA and Logon to the switch using `level_15_access` (privileged EXEC mode).
- ◆ Select the *Troubleshooting* tab.
- ◆ Select *Ping and Trace...*
- ◆ Select *Layer 3 Trace* tool in the *Ping and Trace* screen.
- ◆ Enter the *Destination (IP Address/Hostname)*
- ◆ Click *Start*.


Figure 7-11 illustrates the *Layer 3 Trace Route* screen in CNA.

Gather information about your network

So far, you learned how to gather information about your switch and how to troubleshoot connectivity on a Cisco switch. In this section, you find out how to gather information about the network where your switch connects.

There are several methods allowing you to gather information about your network. You have already seen some of them in previous sections in this chapter. You review information gathering and troubleshooting commands in this section and find out about some additional commands. You can use both IOS commands and graphical user interfaces (GUI tools) to gather information about your network.

Figure 7-11:
Cisco
Network
Assistant –
Layer 3
Trace
Route.



Using ICMP

The Internet Control Message Protocol (ICMP) is designed to send control and test IP packets on IP networks. ICMP is a Layer 3 (Network) TCP/IP protocol. Think of an ICMP packet as a specialized IP packet used to discover and control IP network services. Any device that supports IP connectivity also supports ICMP. However, ICMP, or parts of ICMP, can be disabled on a network for security reasons.

Both the *Ping* tool and the *Trace Route* tool are using the Internet Control Message Protocol (ICMP). If ICMP is disabled, the *Ping* tool and the *Trace Route* tool fail. This doesn't mean that there is no connectivity. It is just not possible to verify connectivity using *ping* and *tracert* on a network where ICMP is disabled.

Typically, ICMP is:

- ◆ Enabled inside the firewall, within the internal network: ICMP is useful to troubleshoot IP connectivity and performance problems using the *Ping* and *Trace Route* tools. So, it makes sense to enable these tools within a local, firewall-protected network.

- ◆ Disabled for incoming requests through the firewall: The *Ping* and *Trace Route* tools can be used to find out about the nodes and configuration of a network. Hence, hackers can use these tools to build a “map” of your network. If the firewall allows outside ping and traceroute requests to flow into your network and if the firewall allows ping and traceroute results to be returned out of your network, hackers can use the *Ping* and *Trace Route* tools to discover your network.

The *Ping* and *Trace Route* tools can be used on most devices that support TCP/IP connectivity. All UNIX, Linux, Mac OS X, and Windows-based computers provide the *Ping* and *Trace Route* tools. Hence, you can test connectivity and discover your network using these tools both ways between a Cisco switch and a computer host.

Here is an example of the ping command executed on a Mac OS X computer host to test connectivity to the SW1 (192.168.75.10) switch.

```
Macintosh-8:~ silange$ ping 192.168.75.10
PING 192.168.75.10 (192.168.75.10): 56 data bytes
64 bytes from 192.168.75.10: icmp_seq=0 ttl=255 time=0.546 ms
64 bytes from 192.168.75.10: icmp_seq=1 ttl=255 time=0.532 ms
[... some output cut ...]
```

The ping output shows computer host Macintosh-8 can reach switch SW1 within approximately 0.5 milliseconds (time=0.546 ms, time=0.532 ms). Observe that ping sends each ICMP packet from the Macintosh-8 computer host to the SW1 switch with a sequence number (icmp_seq=0, icmp_seq=1, ...)

Here is another example of the ping command this time executed on a Cisco switch to test connectivity to the Dummies Web site (www.dummies.com):

```
SW1>ping www.dummies.com
Translating "www.dummies.com"...domain server (255.255.255.255) [OK]

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 208.215.179.139, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
```

The ping output shows that switch SW1 cannot reach the Dummies Web site (Success rate is 0 percent (0/5)). This is bizarre, because I can browse to www.dummies.com from SW1's network. Any ideas why ping reports no connectivity? If you answer because incoming ping (ICMP echo) requests are disabled through the firewall of the www.dummies.com network, you are right.

Here is an example of the `traceroute` command executed on a Cisco switch to trace the IP route to the *Dummies* Web site (`www.dummies.com`):

```
SW1>traceroute www.dummies.com
Translating "www.dummies.com"...domain server (255.255.255.255) [OK]

Type escape sequence to abort.
Tracing the route to dummies.com (208.215.179.139)

 0  192.168.75.1 1007 msec 0 msec 0 msec
 1  cpe-174-099-120-001.nc.res.rr.com (174.99.120.1) 9 msec 8 msec 8 msec
 2  ten13-0-0-202.rlghnca-rtr1.nc.rr.com (66.26.32.65) 17 msec 8 msec 17 msec

[... some output cut ...]

 9  ae-2-79.edge1.Washington3.Level3.net (4.68.17.75) 25 msec
    ae-4-99.edge1.Washington3.Level3.net (4.68.17.203) 25 msec
    ae-1-69.edge1.Washington3.Level3.net (4.68.17.11) 42 msec

[... some output cut ...]

13  gar3.nw2nj.ip.att.net (12.122.130.109) 25 msec 34 msec 50 msec
14  12.88.61.178 34 msec 34 msec 33 msec
15  * * *
```

The results returned by the *Trace Tool* show that the SW1 switch reached the `www.dummies.com` site server through several network routers through the `rr.com`, `Level3.net`, and `att.net` networks. This helps you define where your network sits within the larger global networks. You can see that:

- ◆ Switch SW1 connects into Local Area Network (LAN) 192.168.75.0.
- ◆ The LAN 192.168.75.0 connects to the `rr.com` network through gateway (router) 192.168.75.1.
- ◆ The `rr.com` network connects to the `Level3.net` global network.
- ◆ The `Level3.net` global network is interconnected with the `att.net` global network.
- ◆ The `www.dummies.com` site server connects into the `att.net` global network.

Here is an example of the `traceroute` command executed on a Mac OS X computer host to trace the IP route from computer host Macintosh-8 to router RT1 (192.168.75.40):

```
Macintosh-8:~ silange$ traceroute 192.168.75.40
traceroute to 192.168.75.40 (192.168.75.40), 64 hops max, 40 byte packets
 1  192.168.75.40 (192.168.75.40)  1.661 ms * 1.227 ms
```

Here is an example of the `traceroute` command executed on a Mac OS X computer host to trace the IP route from computer host Macintosh-8 to switch SW1:

```
Macintosh-8:~ silange$ traceroute 192.168.75.10
traceroute to 192.168.75.10 (192.168.75.10), 64 hops max, 40 byte packets
 1  192.168.75.10 (192.168.75.10)  0.990 ms * 0.826 ms
```

Observe the preceding performance figures in bold font type. Why are switch SW1 times faster than router RT1 times? If you answer, because a switch is typically faster than a router, because it doesn't need to process IP packets, just data-link frames, you are right.

Here is an example of the `traceroute` command executed on a Mac OS X computer host to trace the IP route to the Dummies Web site (www.dummies.com):

```
Macintosh-8:~ silange$ traceroute dummies.com
traceroute to dummies.com (208.215.179.139), 64 hops max, 40 byte packets
 1  192.168.75.1 (192.168.75.1)  2.454 ms  1.891 ms  1.686 ms
 2  cpe-174-099-120-001.nc.res.rr.com (174.99.120.1)  9.994 ms  11.131 ms  10.260 ms
 3  ten13-0-0-202.rlghnca-rtr1.nc.rr.com (66.26.32.65)  26.870 ms  13.329 ms  13.247 ms

[... some output cut ...]

 9  ae-2-79.edge1.Washington3.Level3.net (4.68.17.75)  25.284 ms ae-4-99.edge1.
    Washington3.Level3.net (4.68.17.203)  26.751 ms ae-3-89.edge1.Washington3.
    Level3.net (4.68.17.139)  37.998 ms

[... some output cut ...]

13  gar3.nw2nj.ip.att.net (12.122.130.109)  32.668 ms  31.946 ms  31.588 ms
14  12.88.61.178 (12.88.61.178)  34.915 ms  32.394 ms  31.880 ms
15  * * *
```

The output of `traceroute dummies.com` on the Mac OS X computer host is very similar to the output of `traceroute dummies.com` on the SW1 switch.

The `traceroute` command on Windows is actually named `tracert`. So, you execute `tracert` instead of `traceroute` on Windows-based computer hosts.

Here is an example of the `tracert` command executed on a Windows computer host to trace the IP route to the *eTips Dummies* Web site (etips.dummies.com):

```
C:\> tracert etips.dummies.com

Tracing route to etips.dummies.com [12.165.240.180] over a maximum of 30 hops:

 1  2.200 ms  1.649 ms  1.726 ms 192.168.75.1 [192.168.75.1]
```

```

2  9.315 ms  8.800 ms  9.282 ms cpe-174-099-120-001.nc.res.rr.com [174.99.120.1]
3  13.422 ms 15.031 ms 13.974 ms ten13-0-0-202.rlghnca-rtr1.nc.rr.com
   [66.26.32.65]

[... some output cut ...]

9  32.835 ms 31.070 ms 36.398 ms ae-61-61.ebr1.Washington1.Level3.net
   (4.69.134.129)

[... some output cut ...]

16 72.789 ms 59.959 ms 60.742 ms gar1.ipsin.ip.att.net [12.122.133.109]
17 62.420 ms 59.553 ms 61.310 ms 12.86.112.170 [12.86.112.170]
18 * * *
```

Observe that Windows writes the performance statistics before the hop host name and IP address. Note also that Windows writes the IP addresses between square brackets instead of curved brackets. Otherwise, the *Trace Route* tool results are similar on all platforms.

Using CNA

You can use the Cisco Network Assistant to discover the network surrounding your switch. CNA provides a Graphical User Interface (GUI) allowing the network administrator to have both a global and detailed views of the Cisco networking environment.


Although CNA monitors and manages both Cisco switches and Cisco routers, it is intended more for switches than for routers. Routers are typically monitored and managed using either the Cisco SDM (Security and Device Manager) tool, or the Cisco SDM Express tool.

You download the CNA software package from Cisco Web site (cisco.com). CNA is provided for free by Cisco.

Here are the features of CNA that allow you to discover the network surrounding your switch:

- ◆ ***Topology View:*** The topology view provides a global view of your Cisco network. You see both the switch you connect to and network devices that are connected to the switch. Figure 7-12 shows the SW1 switch and its neighbor router RT1. Observe that router RT1 is not supported by CNA. RT1 is an older Cisco 2621 series router that is unsupported by CNA version 2.5.
- ◆ ***Ping and Trace tools:*** The *Ping and Trace* tools are used to verify connectivity to neighboring devices from your switch. The topology view provides you a quick overview of your network. The ping and trace route tools allow you to verify connectivity to those devices. You learned previously how to use the Ping and Trace tools in the section “Using ICMP to Discover Network,” and in the section “Testing connectivity using ping and traceroute tools.”

Figure 7-12:
Cisco
Network
Assistant –
Topology
View.



Using CDP

The *Cisco Discovery Protocol (CDP)* is a proprietary Cisco Layer 2 (data link) protocol. CDP is enabled by default on most Cisco switches and routers.

CDP is designed to enable Cisco networking devices to discover each other. The CDP protocol is used, for example, to build the topology view in CNA. Recall that the CNA topology view shows the neighbors of a Cisco switch or router.

CDP is very useful to discover the landscape of your network. Assume that you know the IP address of only one switch or only one router in your network. To discover other devices in your network, you can:

- ◆ Connect to the IP address of the switch or router you know.
- ◆ Use CDP IOS commands to discover:
 - CDP configuration
 - Devices in the network

Configure CDP

For example, to view the CDP configuration on your switch, execute the `show cdp` IOS command:

```
SW1>show cdp
Global CDP information:
  Sending CDP packets every 60 seconds
  Sending a holdtime value of 180 seconds
  Sending CDPv2 advertisements is enabled
```

If the CDP protocol is not enabled on your switch or router, the `show cdp` command reports "% CDP is not enabled". You can enable CDP using the `cdp run` IOS command like this:

```
SW1> enable (or en)
Password: my_priv_encrypt_password
SW1# configure terminal (or config t, or conf t)
SW1(config)# cdp run
SW1(config)# exit
SW1# copy running-configuration startup-configuration
Destination filename [startup-config]? <press Enter>
Building configuration...
[OK]
SW1# disable
SW1>
```

To see available CDP IOS commands, use the contextual help system:

```
SW1> enable (or en)
Password: my_priv_encrypt_password
SW1# configure terminal (or config t, or conf t)
SW1(config)# cdp ?
  advertise-v2 CDP sends version-2 advertisements
  holdtime     Specify the holdtime to be sent in packets
  run          Enable CDP
  timer        Specify the rate at which CDP packets are sent
```

Observe that you need to be in privileged EXEC global configuration mode to execute CDP configuration commands.

There are two CDP parameters that you can control with the `cdp` command:

- ◆ **holdtime:** CDP packets have a hold time value (in seconds) that determines how long the receiver keeps packets. This helps receivers that keep CDP information about their neighbors, even if CDP momentarily stops sending updates.
- ◆ **timer:** CDP packets are constantly exchanged between CDP-enabled switches and routers to keep track of changes in the network. The *timer* parameter determines the frequency (in seconds) at which these packets are sent.

You can enable or disable CDP individually on each interface. To do so, you:

- ◆ Select the interface on which you want to enable or disable CDP.
- ◆ Execute `cdp enable` to enable CDP on that interface.
- ◆ Execute `no cdp enable` to disable CDP on that interface.

By default, CDP is enabled globally on Cisco devices: CDP is enabled on all interfaces on Cisco devices. CDP consumes some bandwidth sending packets about networking devices in the LAN. You may choose to disable CDP on certain interfaces if, for example, you need those interfaces to perform at top speed.

Discover your network using CDP

When CDP is enabled and configured, you can use it to discover the landscape of your network. There are several discovery options for the `show cdp` command. Read on to find out about each of these CDP network discovery commands.

Discover neighbor devices

To discover the neighbors of switch SW1, use the `show cdp neighbors` IOS command:

```
SW1>sh cdp neighbors
Capability Codes:
R-Router, T-Trans Bridge, B-Source Route Bridge
S-Switch, H-Host, I-IGMP, r - Repeater, P - Phone

Device ID  Local Intrfce  Holdtme  Capability  Platform  Port ID
RT2.silange.com  Fas 0/7  175      R S I      1801      Fas 7
RT2.silange.com  Fas 0/5  175      R S I      1801      Fas 1
RT1           Fas 0/1  146      R I        2621      Fas 0/0
```

The output shows that switch SW1 is connected to:

- ◆ Router RT2 using two interfaces:
 - SW1 interface 7 (Fas 0/7) connects to RT2 interface 7 (Fas7)
 - SW1 interface 5 (Fas 0/5) connects to RT2 interface 1 (Fas1)
- ◆ Router RT1 using one interface:
 - SW1 interface 1 (Fas 0/1) connects to RT2 interface 0 (Fas 0/0)

Observe that you see both the local interface on SW1 (Local Intrfce) and the remote interface on RT2 (Port ID) in the output of `show cdp neighbors`.

The `show cdp neighbors` IOS command reports the platform and capabilities of each remote device:

- ◆ RT2 is reported to be an 1801 series router with router (R), switch (S) and ICMP (I) capabilities.
- ◆ RT1 is reported to be a 2621 series router with router (R) and ICMP (I) capabilities.

To get more details about the remote RT2 router, use the `detail` argument with the `show cdp neighbors` command.



You can also execute the `show cdp entry *` command on SW1 to view details about devices in the vicinity of switch SW1. The `show cdp entry *` command and the `show cdp neighbors detail` command display the same output.

Here are the items that are reported by the `show cdp neighbors detail` command, in addition to the summary command:

- ◆ IP address of remote device
- ◆ Cisco IOS software version of remote device
- ◆ CDP advertisement version
- ◆ VTP (VLAN Trunk Protocol) Management Domain of remote device
- ◆ Native VLAN of remote device
- ◆ Duplex mode of the link to the remote device

You can use either the `show cdp neighbors detail` command or the `show cdp entry *` command to view details about devices in the vicinity of the switch or router where you run the `show cdp` command.

The `show cdp entry *` command can filter the output if you use one of the following arguments:

- ◆ `protocol`: Shows only the IP address of each remote node.
- ◆ `version`: Shows only the Cisco IOS software version of each remote node.

Viewing interface status and configuration

To discover the status and configuration of interfaces on switch SW1, use the `show cdp interface` IOS command.

Troubleshooting the startup configuration

If your switch has trouble starting up, there may be problems with the startup configuration. To verify whether the startup configuration is causing problems during startup, you can:

- ◆ Boot the switch loading a known good startup configuration file, if you have one. You should have one if you followed the best practice to backup your startup configuration file to a backup computer host. If the switch starts up fine when you load the known good startup configuration file, you can simply replace the faulty startup configuration file with the known good one.

- ◆ Boot the switch without loading the startup configuration file. If the switch starts up fine when you do not load the startup configuration, chances are there is something wrong in the startup configuration, or the startup configuration file is corrupted.

To boot the switch without loading the startup configuration file, you need to interrupt the automatic boot process using one of these methods:

- ◆ Use the `boot` command to enable manual booting.
- ◆ Press the Mode button during the boot process.
- ◆ Press Ctrl-Break during the boot process.

Either of these methods brings you at the `switch:` manual boot prompt. The manual boot prompt allows you to control the boot process of the switch. Here are the steps to boot up the switch manually, without loading the startup configuration file from NVRAM:

- ◆ Manually initialize the flash file system:

```
switch: flash_init

Initializing Flash...
mifs[2]: 0 files, 1 directories
mifs[2]: Total bytes      : 3870720
mifs[2]: Bytes used      : 1024
mifs[2]: Bytes available : 3869696
mifs[2]: mifs fsck took 0 seconds.
mifs[3]: 518 files, 19 directories
mifs[3]: Total bytes      : 27998208
mifs[3]: Bytes used      : 8687616
mifs[3]: Bytes available : 19310592
mifs[3]: mifs fsck took 5 seconds.
...done Initializing Flash.
```

- ◆ Hide the startup configuration file to prevent the IOS from loading it.

```
switch: rename flash:config.text flash:config.text.old
```

- ◆ Boot the switch manually and wait for IOS to finish loading in RAM.

```
switch: boot
```

If the switch starts up fine, you need to rebuild a new, valid startup configuration file, using the setup utility:

- ◆ Execute the setup utility to setup your switch.

```
SW1>enable (or en)
SW1# setup
--- System Configuration Dialog ---
Continue with configuration dialog? [yes/no]: yes

[... some output cut ...]
```

- ◆ Save the current running-configuration over the startup configuration file in NVRAM.

```
SW1#copy running-config startup-config
Destination filename [startup-config]?
Building configuration...
[OK]
```

- ◆ Reboot your switch.

```
SW1# reload
```

Troubleshooting the running configuration

If your switch reports errors, alerts, or warnings during normal operation, there may be something wrong with the running configuration. To verify whether the running configuration is causing problems, you can:

- ◆ Boot the switch loading a known good startup configuration file, if you have one. You should have one if you followed the best practice to backup your startup configuration file to a backup computer host.
 - If the switch operates fine when you load the known good startup configuration file, you can simply replace the faulty startup configuration file with the known good one.
 - If the switch still reports errors, alerts, or warnings when you load the known good startup configuration file, the problem may not be related to the running configuration. Read on.
- ◆ If problems persist even after you start the switch with a known good startup configuration, boot the switch without loading a startup configuration at all.
 - If the switch starts up and operates fine without a startup configuration, chances are there is something wrong in the startup configuration, or the startup configuration file is corrupted. You need to fix the startup configuration. Refer to the section “Troubleshooting the startup configuration” to see how to fix the startup configuration.
 - If the switch still reports errors, alerts, or warnings, even without a startup configuration, the problem may not be related to the startup configuration. Read on.

If problems persist when you start the switch without a startup configuration, the startup and the running configurations are not the cause of the problems. You need to use other troubleshooting tools to find out the culprit:

- ◆ **Inspect the logs and system messages:** First, look at the logs and system messages, searching for any error messages, fatal exceptions, alerts, and warnings. Refer to the section “Inspect Cisco switch logs and system messages” earlier in this chapter to see how to manage and inspect the logs and system messages.

- If you notice any error message, fatal exception, alert, or warning in the logs or in the system messages, you need to correct that issue.
- If you do not notice error messages, fatal exceptions, alerts, and warnings in the logs or in the system messages, read on.
- ◆ **Use the `show tech-support` command:** You can use this command to inspect many technical aspects on your switch. Look out for any error messages, fatal exceptions, alerts, or warnings in the output of the `show tech-support` command. Refer to the section “Inspect Cisco Switch Configuration and Memory Contents” earlier in this chapter to see how to use the `show tech-support` command and how to interpret its output.
- ◆ **Use the `debug` command:** This command provides a wealth of technical information about the operation of your switch.

Using the Debug command

The `debug` command is a very useful troubleshooting tool. This command enables *debug mode* on the switch. In debug mode, the switch displays and logs detailed information about ongoing operations, including any error messages, fatal exceptions, alerts, and warnings.

Best practice

You can enable debug mode using the `debug` command on specific components and protocols. It is best practice to enable debug mode selectively on certain components or protocols that you suspect are troublesome. It is not recommended to enable debug mode on all components, features and protocols, because the debug process has priority over all other processes on your switch. This may severely impact the performance of your switch, to the point of rendering it unusable.

You execute the `debug` command in privileged EXEC mode.

To enable debug mode on a specific component, feature, or protocol, execute the `debug` command suffixed by the component, feature, or protocol name. Many components, features and protocols have various items that can be monitored in debug mode. For example, there are several aspects of STP (Spanning Tree Protocol) that you can debug:

```
SW1#debug spanning-tree ?
all                All Spanning Tree debugging messages
backbonefast       BackboneFast events
bpdu               Spanning tree BPDU
bpdu-opt           Optimized BPDU handling
config             Spanning tree config changes
csuf/csrt          STP CSUF/CSRT
etherchannel       EtherChannel support
events            Spanning tree topology events
exceptions         Spanning tree exceptions
general           Spanning tree general
```

```
mstp           MSTP debug commands
pvst+          PVST+ events
root           Spanning tree root events
snmp           Spanning Tree SNMP handling
switch         Switch Shim debug commands
synchronization STP state sync events
uplinkfast     UplinkFast events
SW1#
SW1# debug spanning-tree general
Spanning Tree general debugging is on
SW1#
```

You just enabled general debugging for STP. Now, you will see the following messages displayed on the switch console and logged in the system logs whenever connected switch ports return STP statistics:

```
SW1#
000012: 00:18:35: Returning spanning tree port stats: FastEthernet0/1 (1A32264)
000013: 00:18:35: Returning spanning tree port stats: FastEthernet0/5 (1A2E790)
000014: 00:18:35: Returning spanning tree port stats: FastEthernet0/7 (1A330B0)
000015: 00:19:35: Returning spanning tree port stats: FastEthernet0/1 (1A32264)
000016: 00:19:35: Returning spanning tree port stats: FastEthernet0/5 (1A2E790)
000017: 00:19:35: Returning spanning tree port stats: FastEthernet0/7 (1A330B0)
```

To disable debugging of a specific component, feature, or protocol, execute the `no debug` command followed by the component, feature, or protocol. For example, to disable general STP debugging, execute:

```
SW1# no debug spanning-tree general
Spanning Tree general debugging is off
SW1#
```

To disable all debugging, execute:

```
SW1# no debug all
All possible debugging has been turned off
SW1#
```

Prep Test

- 1** You can obtain the version of the IOS on your switch by looking at
 - A ☐ The output messages displayed during the boot process
 - B ☐ The IOS image file name
 - C ☐ The show version command output
 - D ☐ All of the above
- 2** You can inspect the contents stored in the flash memory of a Cisco switch by using
 - A ☐ The show flash and show tech-support IOS commands
 - B ☐ Cisco IFS (IOS file System) commands
 - C ☐ Cisco Network Assistant (CNA)
 - D ☐ All of the above
- 3** You can inspect the inspect the configuration of a Cisco switch by using
 - A ☐ The show flash IOS command
 - B ☐ The show startup-config command and the show running-config command
 - C ☐ The Cisco Router and Security Device Manager (SDM)
 - D ☐ All of the above
- 4** You can inspect the log files and system messages on a Cisco switch by using
 - A ☐ The show logging IOS command
 - B ☐ The show log-config IOS command
 - C ☐ The Cisco Router and Security Device Manager (SDM)
 - D ☐ All of the above
- 5** You can configure the logging system on a Cisco switch by using
 - A ☐ The log-config IOS command
 - B ☐ The logging IOS command
 - C ☐ The Cisco Router and Security Device Manager (SDM)
 - D ☐ All of the above
- 6** To troubleshoot connectivity on a Cisco switch, you need to verify
 - A ☐ Continuity of signal through cables and patch panels involved in the link
 - B ☐ Connectors involved in the link are properly making contact
 - C ☐ Switch ports involved in the link are up, and properly configured
 - D ☐ All of the above

7 The ICMP (Internet Control Message Protocol) is designed to send

- A ☐ HTTP text over IP networks
- B ☐ E-mail messages over IP networks
- C ☐ Control and test IP packets over IP networks
- D ☐ All of the above

8 The ping and trace route tools use

- A ☐ The ICMP protocol
- B ☐ The DHCP protocol
- C ☐ CHAP authentication
- D ☐ All of the above

9 The ping tool allows you to

- A ☐ Discover the network hops in a WAN
- B ☐ Send IP packets to an IP address to test end-to-end connectivity
- C ☐ Trace the IP route between two nodes on a TCP/IP network
- D ☐ All of the above

10 The trace route tool allows you to

- A ☐ Discover the network hops in a WAN
- B ☐ Send IP packets to an IP address to test end-to-end connectivity
- C ☐ Trace the IP route between two nodes on a TCP/IP network
- D ☐ All of the above

Answers

- 1** D. All of the above
- 2** D. All of the above
- 3** B. Using show startup-config command and the show running-config command
- 4** A. The show logging IOS command
- 5** B. The logging IOS command
- 6** D. All of the above
- 7** C. Send control and test IP packets over IP networks
- 8** A. Use the ICMP protocol
- 9** B. An IOS command and CNA tool allowing you to send IP packets to an IP address to test end-to-end connectivity
- 10** C. An IOS command and CNA tool allowing you to trace the IP route between two nodes on a TCP/IP network