CHAPTER 14

Mitigating Layer 2 Attacks

Unlike hubs, switches cannot regulate the flow of data between their ports by creating almost "instant" networks that contain only the two end devices communicating with each other. Data frames are sent by end systems, and their source and destination addresses are not changed throughout the switched domain. Switches maintain content-addressable memory (CAM) lookup tables to track the source addresses located on the switch ports. These lookup tables are populated by an address-learning process on the switch. If the destination address of a frame is not known or if the frame received by the switch is destined for a broadcast address, the switch forwards the frame out all ports. With their ability to isolate traffic and create the "instant" networks, switches can be used to divide a physical network into multiple logical or VLANs through the use of Layer 2 traffic segmentation.

VLANs enable network administrators to divide their physical networks into a set of smaller logical networks. Like their physical counterparts, each VLAN consists of a single broadcast domain isolated from other VLANs and work by tagging packets with an identification header and then restricting the ports that the tagged packets can be received on to those that are part of the VLAN. The two most prevalent VLAN tagging techniques are the IEEE 802.1q tag and the Cisco Inter-Switch Link (ISL) tag.

This chapter discusses Layer 2 attacks, mitigations, best practices, and functionality.

"Do I Know This Already?" Quiz

The purpose of the "Do I Know This Already?" quiz is to help you decide whether you really need to read the entire chapter. If you already intend to read the entire chapter, you do not necessarily need to answer these questions now.

The 11-question quiz, derived from the major sections in "Foundation Topics" section of the chapter, helps you determine how to spend your limited study time.

Table 14-1 outlines the major topics discussed in this chapter and the "Do I Know This Already?" quiz questions that correspond to those topics.

Table 14-1 *"Do I Know This Already?" Foundation Topics Section-to-Question Mapping*

Foundation Topics Section	Questions Covered in This Section
Types of Attacks	1–10
Factors Affecting Layer 2 Mitigation Techniques	11

CAUTION The goal of self-assessment is to gauge your mastery of the topics in this chapter. If you do not know the answer to a question or are only partially sure of the answer, you should mark this question wrong for purposes of the self-assessment. Giving yourself credit for an answer you correctly guess skews your self-assessment results and might provide you with a false sense of security.

1. What is the default inactivity expire time period on a Cisco Catalyst switch CAM table?
   a. 1 minute
   b. 5 minutes
   c. 10 minutes
   d. 50 minutes

2. What are three methods of implementing port security?
   a. Active secure MAC addresses, fixed secure MAC address, and closed secure MAC address
   b. Static secure MAC address, dynamic secure MAC addresses, and sticky secure MAC addresses
   c. Default secure MAC address, evasive secure MAC address, and evading secure MAC address
   d. Stinky secure MAC addresses, dynamite secure MAC, and clammy secure MAC addresses

3. Which command enables port security on an interface?
   a. **switchport mode port-security**
   b. **switchport mode interface-security**
   c. **switchport interface-security**
   d. **switchport port-security**

4. What is the default action mode for security violations?
   a. Protect
   b. Restrict
   c. Shutdown

5. The DTP state on a trunk port may be set to what?
   a. Auto, on, off, undesirable, or non-negotiate
   b. Auto, on, off, desirable, or non-negotiate
   c. Auto, on, off, desirable, or negotiate
   d. Auto, on, off, undesirable, or negotiate

6. What are the two different types of VLAN hopping attacks?
   a. Switch spoofing and double tagging
   b. Switch goofing and double teaming
   c. Switch impersonation and double grouping
   d. Switch imitation and double alliance

7. Which features of Cisco IOS Software enable you to mitigate STP manipulation? (Select two.)
   a. **spanning-tree portfast bpduguard**
   b. **spanning-tree guard rootguard**
   c. **set spantree global-default loopguard enable**
   d. **set udld enable**

8. What are the three types of private VLAN ports?
   a. Neighborhood, remote, and loose
   b. Community, isolated, and promiscuous
   c. Communal, remote, and licentious
   d. Area, secluded, and wanton

9. What common tool is used to launch MAC overflow attacks?
   a. Dsniff
   b. Macof
   c. Arpspoof
   d. Tablstuf

10. Protect mode security is recommended for trunk ports.

- **a.** True
- **b.** False

11. What are the three factors when designing a Layer 2 protected network?

- **a.** Number of users groups
- **b.** Number of DNS zones
- **c.** Number of switches
- **d.** Number of buildings
- **e.** Number of security zones

The answers to the "Do I Know This Already?" quiz are found in the appendix. The suggested choices for your next step are as follows:

- **9 or less overall score**—Read the entire chapter. This includes the "Foundation Topics" and "Foundation Summary" sections and the "Q&A" section.
- **10 or 11 overall score**—If you want more review on these topics, skip to the "Foundation Summary" section and then go to the "Q&A" section. Otherwise, move on to Chapter 15, "Context-Based Access Control."

Foundation Topics

The following sections provide an overview of the most common Layer 2 attacks and suggested mitigations. There is also a brief description of factors that you should consider when designing Layer 2 protected networks.

Types of Attacks

For years, the focus on security has been at the network edge or the IP level (Open System Interconnection [OSI] Layer 3). As the popularity of Ethernet switching and wireless LANs grow, however, the emphasis on Layer 2 security has become more important. Yet, less public information is available regarding security risks in a Layer 2 environment and mitigating strategies of these risks. In addition, switches and wireless access points are susceptible to many of the same Layer 3 attacks as routers.

The most common types of Layer 2 attacks are as follows:

- CAM table overflow
- VLAN hopping
- Spanning Tree Protocol (STP) manipulation
- MAC address spoofing
- Private VLAN
- DHCP "starvation"

The following sections discuss the most common Layer 2 attacks and recommended methods to reduce the effects of these attacks.

CAM Table Overflow Attacks

The content-addressable memory (CAM) table in a switch stores information, such as MAC addresses and associated VLAN parameters. It is similar to a router's routing table. CAM tables have a fixed size.

A MAC address is a 48-bit hexadecimal number composed of two descriptive fields. The first 24 bits comprise the manufacturer code assigned by the IEEE. The second 24 bits comprise the specific interface number assigned by the hardware manufacturer. A MAC address of FF.FF.FF.FF.FF.FF

is a broadcast address. Each MAC address is a unique series of numbers, similar to serial numbers or LAN IP addresses. A manufacturer should not have two devices with the same MAC address.

When a Layer 2 switch receives a frame, the switch looks in the CAM table for the destination MAC address. If an entry exists for that MAC address, the switch forwards the frame to the port identified in the CAM table for that MAC address. If the MAC address is not in the CAM table, the switch forwards the frame out all ports on the switch. If the switch sees a response as a result of the forwarded frame, it updates the CAM table with the port on which the communication was received.

In a typical LAN environment where there are multiple switches connected on the network, all the switches receive the unknown destination frame. Figure 14-1 shows the CAM table operation.

Figure 14-1 *CAM Table Operation*

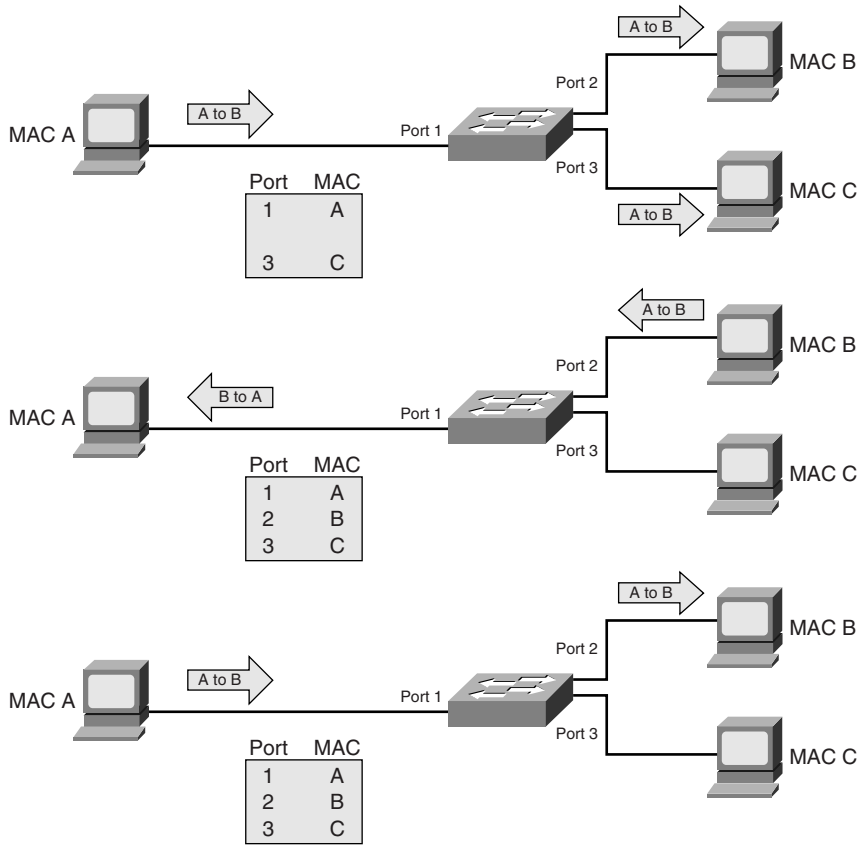

As previously mentioned, the CAM table has a limited size. Cisco Catalyst switches use the 63 bits of source (MAC, VLAN, and so on) and create a 14-bit hash value. If the value is the same, there are eight buckets in which to place CAM entries. These entries expire after a certain inactivity period. (The default on the Cisco Catalyst switch is 5 minutes.) If enough MAC addresses are flooded to a switch before existing entries expire, the CAM table fills up, and new entries are not accepted. When the CAM table is full, the switch starts flooding the packets out all ports. This scenario is called a *CAM table overflow*.

In a CAM table overflow attack, an attacker sends thousands of bogus MAC addresses from one port, which looks like valid hosts' communication, to the switch. One of the more popular tools used for launching this type of attack is called Macof, which was written using PERL code, ported to C language, and bundled into the Dsniff suite. Dsniff is a collection of tools for network auditing and penetration testing. Macof can generate 155,000 MAC entries on a switch per minute. The goal is to flood the switch with traffic by filling the CAM table with false entries. When flooded, the switch broadcasts traffic without a CAM entry out on its local VLAN, thus allowing the attacker to see other VLAN traffic that would not otherwise display.

Figure 14-2 shows a CAM table overflow attack.

Figure 14-2 *CAM Table Overflow Attack*

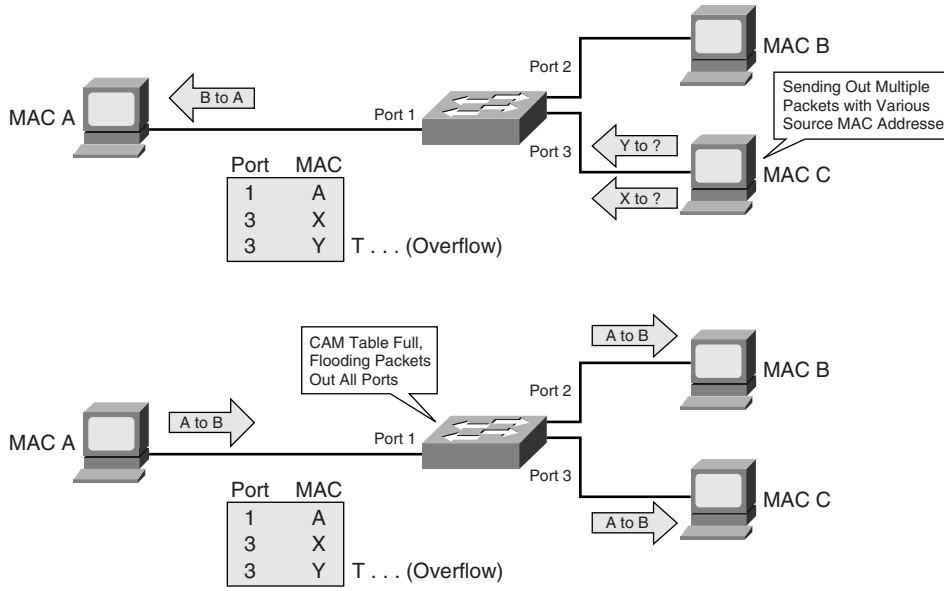

Mitigating CAM Table Overflow Attacks

You can mitigate CAM table overflow attacks in several ways. One of the primary ways is to configure port security on the switch. You can apply port security in three ways:

- **Static secure MAC addresses**—A switch port may be manually configured with the specific MAC address of the device that connects to it.
- **Dynamic secure MAC addresses**—The maximum number of MAC addresses that will be learned on a single switch port is specified. These MAC addresses are dynamically learned, stored only in the address table, and removed when the switch restarts.
- **Sticky secure MAC addresses**—The maximum number of MAC addresses on a given port may be dynamically learned or manually configured. The manual configuration is not a recommended method because of the high administrative overhead. The *sticky* addresses will be stored in the address table and added to the running configuration. If the addresses are saved in the configuration file, the interface does not need to dynamically relearn them when the switch restarts.

The type of action taken when a port security violation occurs falls into the following three categories:

- **Protect**—If the number of secure MAC addresses reaches the limit allowed on the port, packets with unknown source addresses are dropped until a number of MAC addresses are removed or the number of allowable addresses is increased. You receive no notification of the security violation in this type of instance.
- **Restrict**—If the number of secure MAC addresses reaches the limit allowed on the port, packets with unknown source addresses are dropped until some number of secure MAC addresses are removed or the maximum allowable addresses is increased. In this mode, a security notification is sent to the Simple Network Management Protocol (SNMP) server (if configured) and a syslog message is logged. The violation counter is also incremented.
- **Shutdown**—If a port security violation occurs, the interface changes to error-disabled and the LED is turned off. It sends an SNMP trap, logs to a syslog message, and increments the violation counter.

NOTE On some Cisco Catalyst switch platforms, such as 3550 and 6500, you can enable port security on trunk ports. On these platforms, Protect mode security is not recommended for a trunk port. Protect mode disables learning when any VLAN reaches its maximum limit, even if the port has not reached its maximum limit.

Table 14-2 shows the commands used to configure port security.

Table 14-2 *Configuring Port Security*

Command	Description
switchport mode {**access** | **trunk**}	Sets the interface mode as access or trunk. An interface in the default mode of dynamic auto cannot be configured as a secure port.
switchport port-security	Enables port security on the interface.
switchport port-security [**maximum** *value* [**vlan** {*vlan-list* | {**access** | **voice**}}]]	Sets the maximum number of secure MAC addresses for the interface. The active Switch Database Management (SDM) template determines the maximum number of available addresses. The default is 1.
switchport port-security violation {**protect** | **restrict** | **shutdown**}	Sets the action to be taken when a security violation is detected. The default mode for security violations is to shut down the interface.
switchport port-security [**mac-address** *mac-address* [**vlan** {*vlan-id* | {**access** | **voice**}}]	Sets a secure MAC address for the interface. This command may be used to enter the maximum number of secure MAC addresses. If fewer secure MAC addresses are configured than the maximum, the remaining MAC addresses are dynamically learned.
switchport port-security mac-address sticky	Enables sticky learning on the interface.
switchport port-security mac-address sticky [*mac-address* | **vlan** {*vlan-id* | {**access** | **voice**}}]	Sets a sticky secure MAC address. This command can be repeated as many times as necessary. If fewer secure MAC addresses are configured than the maximum, the remaining MAC addresses are dynamically learned, converted to sticky secure MAC addresses, and added to the running configuration.

NOTE When a secure port is in the error-disabled state, you can bring it out of this state by entering the **errdisable recovery cause psecure-violation** global configuration command, or you can manually re-enable it by entering the **shutdown** and **no shutdown** interface configuration commands.

The following example configures a switch port as an access port and sets dynamic port security with maximum number of addresses learned to 20. The violation mode is the default shutdown mode, sticky learning is enabled, and no static MAC addresses are configured. In the scenario where a twenty-first computer tries to connect, the port will be placed in error-disabled state and will send out an SNMP trap notification.

Example 14-1 *Configuring Dynamic Port Security*

```
Switch#configure terminal
Switch(config)#interface fastethernet0/0
Switch(config-if)#switchport mode access
Switch(config-if)#switchport port-security
Switch(config-if)#switchport port-security maximum 20
Switch(config-if)#switchport port-security mac-address sticky
```

VLAN Hopping Attacks

VLANs are a simple way to segment the network within an enterprise to improve performance and simplify maintenance. Each VLAN consists of a single broadcast domain. VLANs work by tagging packets with an identification header. Ports are restricted to receiving only packets that are part of the VLAN. The VLAN information may be carried between switches in a LAN using trunk ports. Trunk ports have access to all VLANs by default. They route traffic for multiple VLANs across the same physical link. Two types of trunks are used: 802.1q and ISL. The trunking mode on a switch port may be sensed using Dynamic Trunk Protocol (DTP), which automatically senses whether the adjacent device to the port may be capable of trunking. If so, it synchronizes the trunking mode on the two ends. The DTP state on a trunk port may be set to auto, on, off, desirable, or non-negotiate. The DTP default on most switches is auto.

One of the areas of concern with Layer 2 security is the variety of mechanisms by which packets that are sent from one VLAN may be intercepted or redirected to another VLAN, which is called *VLAN hopping*. VLAN hopping attacks are designed to allow attackers to bypass a Layer 3 device when communicating from one VLAN to another. The attack works by taking advantage of an incorrectly configured trunk port.

It is important to note that this type of attack does not work on a single switch because the frame will never be forwarded to the destination. But in a multiswitch environment, a trunk link could be exploited to transmit the packet. There are two different types of VLAN hopping attacks:

- **Switch spoofing**—The network attacker configures a system to spoof itself as a switch by emulating either ISL or 802.1q, and DTP signaling. This makes the attacker appear to be a switch with a trunk port and therefore a member of all VLANs.

- **Double tagging**—Another variation of the VLAN hopping attack involves tagging the transmitted frames with two 802.1q headers. Most switches today perform only one level of decapsulation. So when the first switch sees the double-tagged frame, it strips the first tag off the frame and then forwards with the inner 802.1q tag to all switch ports in the attacker's VLAN as well as to all trunk ports. The second switch forwards the packet based on the VLAN ID in the second 802.1q header. This type of attack works even if the trunk ports are set to off.

Figure 14-3 shows VLAN hopping with a double-tagging scenario.

Figure 14-3 *Double-Tagging VLAN Hopping Attack*

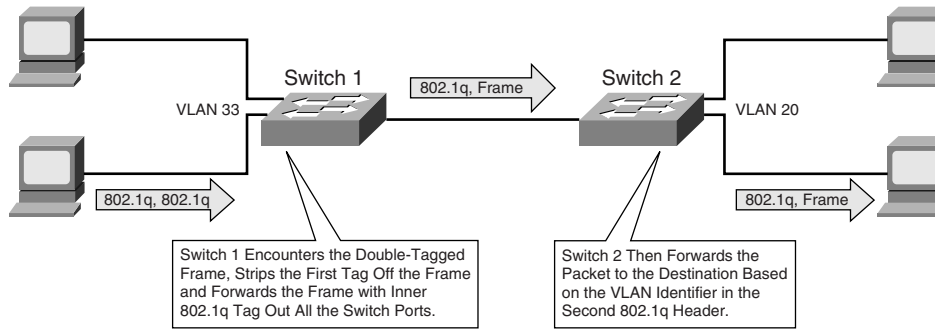

Mitigating VLAN Hopping Attacks

Mitigating VLAN hopping attacks requires the following configuration modifications:

- Always use dedicated VLAN IDs for all trunk ports.
- Disable all unused ports and place them in an unused VLAN.
- Set all user ports to nontrunking mode by disabling DTP. Use the **switchport mode access** command in the interface configuration mode.
- For backbone switch-to-switch connections, explicitly configure trunking.
- Do not use the user native VLAN as the trunk port native VLAN.
- Do not use VLAN 1 as the switch management VLAN.

STP Manipulation Attacks

STP prevents bridging loops in a redundant switched network environment. By avoiding loops, you can ensure that broadcast traffic does not become a traffic storm.

STP is a hierarchical tree-like topology with a "root" switch at the top. A switch is elected as root based on the lowest configured priority of any switch (0 through 65,535). When a switch boots up, it begins a process of identifying other switches and determining the root bridge. After a root bridge is elected, the topology is established from its perspective of the connectivity. The switches determine the path to the root bridge, and all redundant paths are blocked. STP sends configuration and topology change notifications and acknowledgments (TCN/TCA) using bridge protocol data units (BPDU).

An STP attack involves an attacker spoofing the root bridge in the topology. The attacker broadcasts out an STP configuration/topology change BPDU in an attempt to force an STP recalculation. The BPDU sent out announces that the attacker's system has a lower bridge priority. The attacker can then see a variety of frames forwarded from other switches to it. STP recalculation may also cause a denial-of-service (DoS) condition on the network by causing an interruption of 30 to 45 seconds each time the root bridge changes. Figure 14-4 shows an attacker using STP network topology changes to force its host to be elected as the root bridge.

Figure 14-4 *STP Attack*

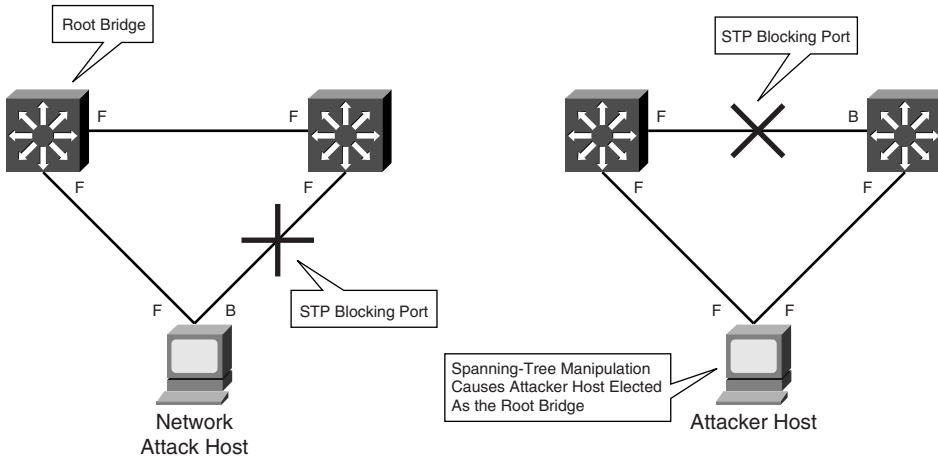

Preventing STP Manipulation Attacks

To mitigate STP manipulation, use the root guard and BPDU guard features in the Cisco IOS Software. These commands enforce the placement of the root bridge and the STP domain borders. The STP root guard feature is designed to allow the placement of the root bridge in the network. The STP BPDU guard is used to keep all active network topology predictable.

Example 14-2 shows an example of enabling BPDU guard, using **portfast**, to disable ports upon detection of a BPDU message and to disable ports that would become the root bridge because of their BPDU advertisement.

Example 14-2 *Enabling BPDU and Root Guard*

```
Switch#configure terminal
Switch(config)#spanning-tree portfast bpduguard
Switch(config)#interface fa0/10
Switch(config)#spanning-tree guard root
```

MAC Address Spoofing—Man-in-the-Middle Attacks

MAC spoofing involves the use of a known MAC address of another host that is authorized to access the network. The attacker attempts to make the target switch forward frames destined for the actual host to the attacker device instead. This is done by sending a frame with the other host's source Ethernet address with the objective to overwrite the CAM table entry. After the CAM is overwritten, all the packets destined for the actual host will be diverted to the attacker. If the original host sends out traffic, the CAM table will be rewritten again, moving the traffic back to the original host port. Figure 14-5 shows how MAC spoofing works.

Figure 14-5 *MAC Spoofing Attack*

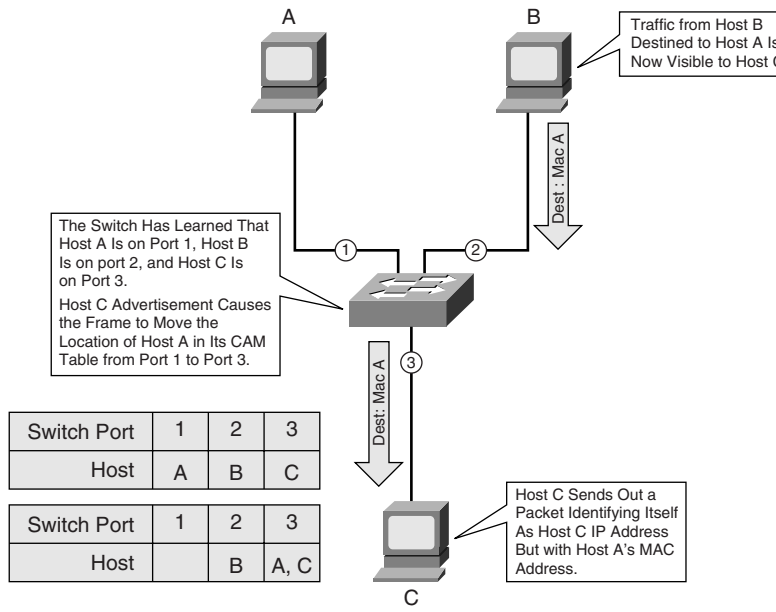

Another method of spoofing MAC addresses is to use Address Resolution Protocol (ARP), which is used to map IP addressing to MAC addresses residing on one LAN segment. When a host sends out a broadcast ARP request to find a MAC address of a particular host, an ARP response comes from the host whose address matches the request. The ARP response is cached by the requesting host. ARP protocol also has another method of identifying host IP-to-MAC associations, which is called Gratuitous ARP (GARP), which is a broadcast packet used by hosts to announce their IP address to the LAN to avoid duplicate IP addresses on the network. GARP can be exploited maliciously by an attacker to spoof the identity of an IP address on a LAN segment. This is typically used to spoof the identity between two hosts or all traffic to and from the default gateway.

One of the tools used to spoof ARP entries is called Arpspoof and is part of a collection of tools known as *Dsniff*.

Mitigating MAC Address Spoofing Attacks

Use the **port-security** command described in the "Mitigating CAM Table Overflow Attacks" section to specify MAC addresses connected to particular ports; however, this type of configuration has a high administrative overhead and is prone to mistakes. There are other mechanisms, such as hold-down timers, that you can use to mitigate ARP spoofing attacks by setting the length of time an entry will stay in the ARP cache. Hold-down timers by themselves are insufficient to mitigate attacks. It is possible to combine this with the modifications to the ARP cache expiration time for all the hosts, but this is also unmanageable. One recommended alternative is to use private VLANs to mitigate these types of network attacks. A couple of other features of Cisco IOS Software provide protection from this type of attack: dynamic ARP inspection (DAI) and DHCP snooping (as described in the following sections).

Using DHCP Snooping

DHCP snooping is a DHCP security feature that provides network security by filtering untrusted DHCP messages using a DHCP snooping binding database it builds and maintains (referred to as the *DHCP snooping binding table*). An untrusted message is a message received from outside the network or firewall. When a switch receives a packet on an untrusted interface and DHCP snooping is enabled on that interface or VLAN, the switch compares the source MAC address and the DHCP client requester hardware address. If the addresses match, the switch forwards the packet. If the addresses do not match, the switch drops the packet. DHCP snooping considers DHCP messages originating from any user port to a DHCP server as untrusted. An untrusted port should not send DHCP server type responses, such as DHCPOffer, DHCPAck, and so on.

A trusted interface is an interface configured to receive messages only from within the network.

The DHCP snooping binding table contains information such as the host MAC address (dynamic and static), IP address, lease time, binding type, and VLAN number. The database can have up to 8192 bindings.

To enable DHCP snooping on a switch, use the **ip dhcp snooping** command in global configuration mode. To enable the switch to insert and remove DHCP relay information (option 82 field) in forwarded DHCP request messages to the DHCP server, use the **ip dhcp snooping information option** global configuration command.

Example 14-3 shows the configuration for enabling DHCP snooping for VLAN 34 and how to configure a rate limit of 70 packets per second on FastEthernet port 0/0.

Example 14-3 *Configuring DHCP Snooping*

```
Switch#configure terminal
Switch(config)#ip dhcp snooping
Switch(config)#ip dhcp snooping vlan 34
Switch(config)#ip dhcp snooping information option
Switch(config)#interface fa0/0
Switch(config-if)#ip dhcp snooping limit rate 70
```

Using DAI

DAI is a security feature that intercepts and verifies IP-to-MAC address bindings and discards invalid ARP packets. DAI uses the DHCP snooping database to validate bindings. It associates a trust state with each interface on the switch. Packets arriving on trusted interfaces bypass all DAI validation checks, and those arriving on untrusted interfaces undergo the DAI validation process. In a typical network, all ports on the switch connected to host are configured as untrusted, and switch ports are considered trusted. Use the **ip arp inspection trust** interface command to configure the trust settings. When the switch is configured for DAI, it rate-limits incoming ARP packets to prevent DoS attacks. The default rate for an untrusted interface is 14 packets per second. Trusted interfaces are not rate-limited. You can change the rate limit by using the **ip arp inspection limit** interface configuration command. DAI uses the ARP access control lists (ACLs) and DHCP snooping database for the list of valid IP-to-MAC address bindings. ARP ACLs take precedence over entries in the DHCP snooping bindings database but have to be manually configured. Use the **ip arp inspection filter** global configuration command to configure ARP ACLs. The switch will drop a packet that is denied in the ARP ACLs even if the DHCP snooping database has a valid binding for it. When a switch drops a packet, it is logged in the buffer and generates a system message. Use the

ip arp inspection log-buffer to configure the number of buffers and the number of entries needed in the specified interval to generate system messages. Example 14-4 shows a sample DAI configuration for VLAN 34, setting fa0/0 to a trusted port with a rate limit of 20 packets per second and a logging buffer of 64 messages.

Example 14-4 *DAI Configuration*

```
Switch#configure terminal
Switch(config)#ip arp inspection vlan 34
Switch(config)#interface fa0/0
Switch(config-if)#ip arp inspection trust
Switch(config-if)#ip arp inspection limit rate 20 burst interval 2
Switch(config-if)#exit
Switch(config)#ip arp inspection log-buffer entries 64
```

Private VLAN Vulnerabilities

Private VLANs isolate ports within a VLAN to communicate only with other ports in the same VLAN. There are three types of private VLAN ports:

- **Community**—Communicate among themselves and with other promiscuous ports. These interfaces are isolated at Layer 2 from all other interfaces in other communities or isolated ports within their private VLAN.
- **Isolated**—Has complete Layer 2 separation from other ports within the same private VLAN except for the promiscuous port. Private VLANs block all traffic to isolated ports except traffic from promiscuous ports. Traffic received from an isolated port is forwarded only to promiscuous ports.
- **Promiscuous**—Communicates with all interfaces, including community and isolated ports within a private VLAN.

NOTE To configure a private VLAN, the switch must be in Virtual Terminal Protocol (VTP) transparent mode.

A network vulnerability of private VLANs involves the use of a proxy to bypass access restrictions of the private VLAN. In a proxy attack, frames are forwarded to a host on the network connected to a promiscuous port, such as a router. The network attacker sends a packet with its source IP and MAC address and a destination IP address of the target system but a destination MAC address of the router. The switch forwards the frame to the router. The router routes the traffic, rewrites the destination MAC address as that of the target, and sends the packet out. Because the router is

authorized to communicate with the private VLANs, the packet is forwarded to the target system. This type of attack allows for unidirectional traffic only because the private VLAN filter blocks the target's attempts to respond. This vulnerability is not a private VLAN vulnerability per se because all the rules of that VLAN were enforced.

Figure 14-6 shows how the private VLAN proxy vulnerability works.

Figure 14-6 *Private VLAN Proxy Attack*

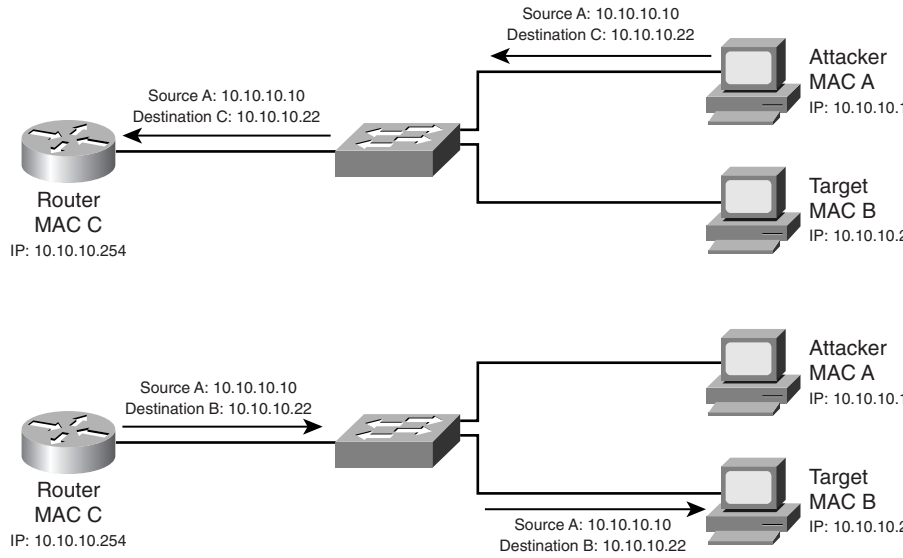

Defending Private VLANs

You can configure ACLs on the router port to mitigate private VLAN attacks. You can also use virtual ACLs on the Cisco Catalyst Layer 3 switch platforms to help mitigate the effects of private VLAN attacks. Example 14-6 shows the ACLs configured on a router port to segment the private VLAN subnet 192.168.200.0/24.

Example 14-5 *Mitigating Private VLAN Proxy Attack*

```
Router1#configure terminal
Router1(config)#access-list 150 deny ip 192.168.200.0 0.0.0.255 192.168.200.0 0.0.0.255 log
Router1(config)#access-list 150 permit ip any any
Router1(config)#interface fa0/0
Router1(config-if)#ip access-group 150 in
```

DHCP Starvation Attacks

A DHCP server dynamically assigns IP addresses to hosts on a network. The administrator creates pools of addresses available for assignment. A lease time is associated with the addresses. DHCP is a standard defined in RFC 2131.

A DHCP starvation attack works by broadcasting DHCP requests with spoofed MAC addresses. This scenario is achieved with attack tools such as gobbler, which looks at the entire DHCP scope and tries to lease all the DHCP addresses available in the DHCP scope. This is a simple resource starvation attack, similar to a SYN flood attack. The attacker can then set up a rogue DHCP server and respond to new DHCP requests from clients on the network. This might result in a "man-in-the-middle" attack.

Mitigating DHCP Starvation Attacks

The methods used to mitigate a MAC address spoofing attack may also prevent DHCP starvation by using the DHCP snooping feature. Implementation of RFC 3118, *Authentication for DHCP Message*, will also assist in mitigating this type of attack.

You can also limit the number of MAC addresses on a switch port, a mitigation strategy for CAM table flooding, to mitigate DHCP starvation attacks.

Other features on the Cisco Catalyst switch, such as IP source guard, may also provide additional defense against attacks. IP source guard initially blocks all IP traffic except for DHCP packets captured by DHCP snooping process. When a client receives a valid IP address from the DHCP server, an ACL is applied to the port. This ACL restricts the traffic from the client to those source IP addresses configured in the binding.

One way to prevent a rogue DHCP server from responding to DHCP requests is to use VLAN ACLs (VACL). You can use VACLs to limit DHCP replies to legitimate DHCP servers and deny them to all others. You should use this type of configuration if the network does not support DHCP snooping.

IEEE 802.1x EAP Attacks

IEEE 802.1x is an IEEE standard link layer (Layer 2) protocol designed to provide port-based network access control using authentication unique to a device or user. Extensible Authentication Protocol (EAP) is the transport mechanism used in 802.1x to authenticate supplicants/clients against

a back-end data store, which is typically a RADIUS server. Figure 14-7 shows client authentication using 802.1x and EAP.

Figure 14-7 *Client Authentication Using 802.1x and EAP*

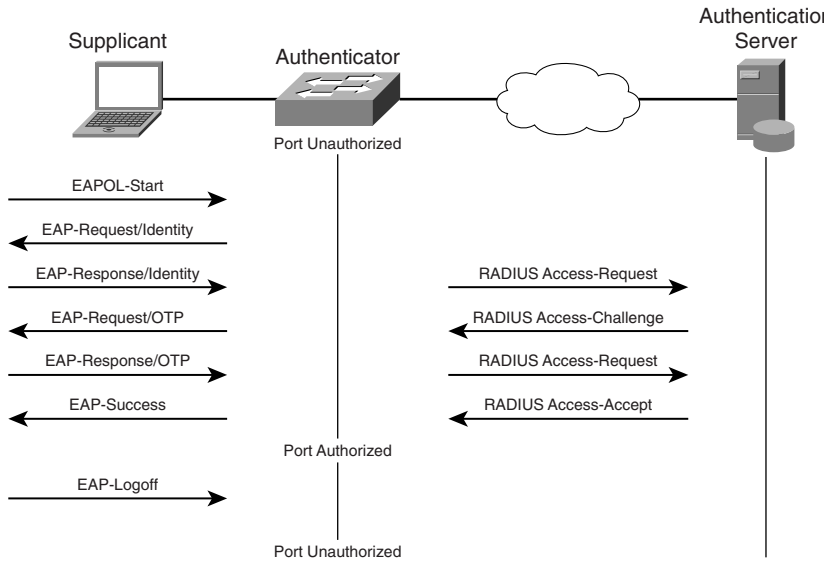

Two types of vulnerabilities are associated with EAP:

- **Man-in-the-middle attack**—At the end of the EAP Over LAN (EAPOL) authentication, the attacker sends the client an EAP-Success message that identifies the attacker as the authenticator. When this action is successful, the attacker is in the path between the client and the authenticator.
- **Session-hijacking attack**—This attack occurs after the authentication process between the client and the authentication server is complete. If the attacker sends a disassociate management frame with the authenticator's MAC address to the client, it will force the client to disconnect from the network; however, the authenticator state is still in an authenticated and associated state, which allows the attacker to access the network.

Mitigating IEEE 802.1x EAP Attacks

Cisco recommends deploying Protected EAP (PEAP) for use in a wireless LAN environment and deploying 802.1x on all access switches to limit physical access to the network. Example 14-6 shows a sample configuration enabling 802.1x authentication on a Cisco router.

Example 14-6 *Enabling 802.1x on a Cisco Router*

```
Router#configure terminal
Router(config)#aaa new-model
Router(config)#aaa authentication dot1x default group radius
Router(config)#dot1x system-auth-control
Router(config)#interface fa0/0
Router(config-if)#dot1x port-control auto
```

For more information on functionality and configuration of IEEE 802.1x and EAP, see Chapters 17, "Identity-Based Networking Services," and 18, "Configuring 802.1x Port-Based Authentication," of this book.

Factors Affecting Layer 2 Mitigation Techniques

You must consider several factors when designing a protected Layer 2 network, as divided into the following three categories:

- **The number of user groups**—Depending on the size of the network, users can be grouped by function, location, or access level. For example, users in a company's financial division can be grouped together to provide restricted access to the financial resources.
- **The number of switches**—The number of switches depends on the scale of the network. As the number of switches in the network grows, the manageability and security of the devices becomes more difficult, and the security vulnerabilities increase.
- **The number of security zones**—A security zone is a logical and physical environment under the common set of policies and procedure with visible management support. In this environment, access privileges to all information assets are defined and followed. For example, an organization's internal LAN could be considered on a security zone or it could be broken up into several depending on the types of users. If the organization has external partners it exchanges information with, it will typically have a separate security zone for this communication, sometimes referred to as an *extranet*. If the network is connected to the Internet, there is typically a demilitarized zone (DMZ), separating the internal network from the Internet. A DMZ is a security zone that might contain resources (such as web servers, mail, and so on) accessible via the public Internet without compromising the internal network security.

Based on the preceding categories, a combination of eight scenarios could be developed to address security vulnerabilities in typical networking environments. Table 14-3 briefly describes the vulnerabilities associated with these combinations.

Table 14-3 *Layer 2 Security Vulnerabilities in Different Network Combinations*

Combination #	Security Zones	Number of User Groups	Number of Switch Devices	Layer 2 Vulnerabilities
1	Single	Single	Single	MAC spoofing CAM table overflow
2	Single	Single	Multiple	MAC spoofing CAM table overflow VLAN hopping STP
3	Single	Multiple	Single	MAC spoofing CAM table overflow VLAN hopping
4	Single	Multiple	Multiple	MAC spoofing CAM table overflow VLAN hopping STP
5	Multiple	Single	Single	MAC spoofing CAM table overflow VLAN hopping
6	Multiple	Single	Multiple	MAC spoofing CAM table overflow VLAN hopping STP

continues

Table 14-3 *Layer 2 Security Vulnerabilities in Different Network Combinations (Continued)*

Combination #	Security Zones	Number of User Groups	Number of Switch Devices	Layer 2 Vulnerabilities
7	Multiple	Multiple	Single	MAC spoofing CAM table overflow VLAN hopping Private VLAN attacks
8	Multiple	Multiple	Multiple	MAC spoofing CAM table overflow VLAN hopping STP VTP attacks

Part VII of this book, "Scenarios," outlines specific examples of these combinations.

Foundation Summary

The "Foundation Summary" section of each chapter lists the most important facts from the chapter. Although this section does not list every fact from the chapter that will be on your SNRS exam, a well-prepared candidate should at a minimum know all the details in each "Foundation Summary" before going to take the exam.

The most common types of Layer 2 attacks and mitigation strategies are as follows:

CAM table overflow—In a CAM table overflow attack, an attacker sends thousands of bogus MAC addresses from one port, which looks like valid hosts' communication to the switch. You can mitigate CAM table overflow attacks in several ways. One of the primary ways is to configure port security on the switch. You can apply port security in three ways: static secure MAC addresses, dynamic secure MAC addresses, and sticky secure MAC addresses.

VLAN hopping—There are two different types of VLAN hopping attacks: switch spoofing and double tagging. Mitigating VLAN hopping attacks requires the following configuration modifications:

- Always use dedicated VLAN IDs for all trunk ports.
- Disable all unused ports and place them in an unused VLAN.
- Set all user ports to nontrunking mode by disabling DTP. Use the **switchport mode access** command in the interface configuration mode.
- For backbone switch-to-switch connections, explicitly configure trunking.
- Do not use the user native VLAN as the trunk port native VLAN.
- Do not use VLAN 1 as the switch management VLAN.

STP manipulation—An STP attack involves an attacker spoofing the root bridge in the topology. The attacker broadcasts out an STP configuration/topology change BPDU in an attempt to force an STP recalculation. The BPDU sent out announces that the attacker's system has a lower bridge priority. The attacker can then see a variety of frames forwarded from other switches to it.

To mitigate STP manipulation, use the root guard and BPDU guard features in the Cisco IOS Software.

MAC address spoofing—MAC address spoofing involves the use of a known MAC address of another host authorized to access the network. The attacker attempts to make the target switch forward frames destined for the actual host to the attacker device instead. Another way to spoof MAC addresses is by using ARP.

Use the **port-security** command described in the "Mitigating CAM Table Overflow Attacks" section to specify MAC addresses connected to particular ports. DHCP snooping could be used as a method to mitigate MAC address spoofing. Another method of mitigating MAC address spoofing is DAI.

Private VLAN—Private VLANs isolate ports within a VLAN to communicate only with other ports in the same VLAN. The three types of private VLAN ports are community, isolated, and promiscuous.

You can configure ACLs on the router port to mitigate private VLAN attacks. You can also use virtual ACLs on the Cisco Catalyst Layer 3 switch platforms to help mitigate the effects of private VLAN attacks.

DHCP "starvation"—A DHCP starvation attack works by broadcasting DHCP requests with spoofed MAC addresses.

The methods used to mitigate MAC address spoofing attack may also prevent DHCP starvation by using the DHCP snooping feature. You can limit the number of MAC addresses on a switch port, a mitigation strategy for CAM table flooding, to mitigate DHCP starvation attack. Other features on the Cisco Catalyst switch, such as IP source guard, may also provide additional defense against attacks.

IEEE 802.1x attack—IEEE 802.1x is an IEEE standard link layer (Layer 2) protocol designed to provide port-based network access control using authentication unique to a device or user.

Two types of vulnerabilities are associated with EAP: man-in-the-middle attacks and session-hijacking attacks. Cisco recommends deploying PEAP for use in a wireless LAN environment and deploying 802.1x on all access switches to limit physical access to the network.

Q&A

As mentioned in the section "How to Use This Book" in the Introduction, you have two choices for review questions. The questions that follow present a bigger challenge than the exam itself because they use an open-ended question format. By reviewing using this more difficult format, you can exercise your memory better and prove your conceptual and factual knowledge of this chapter. You can find the answers to these questions in the appendix.

For more practice with exam-like question formats, including questions using a router simulator and multiple-choice questions, use the exam engine on the CD-ROM.

1. What are the most common types of Layer 2 attacks?
2. Describe the CAM table overflow attack.
3. Explain the three categories of action that can be taken when a port security violation occurs.
4. When a secure port is in the error-disabled state, how can it be brought out of this state?
5. How can you mitigate VLAN hopping attacks?
6. What is involved in an STP attack?
7. How does MAC spoofing–man-in-the-middle attacks work?
8. How can you mitigate MAC spoofing attacks?
9. Describe how a proxy attack bypasses access restrictions of private VLANs.
10. Explain how a DHCP starvation attack is performed.