Understanding Switch Security Issues

Overview

VLAN traffic, VLAN hopping, DHCP spoofing, Address Resolution Protocol (ARP) spoofing, at switch and its ports. You can take specific measures to guard against MAC flooding, which is a common Layer 2 malicious activity.

Objectives

Upon completing this lesson, you will be able to describe and implement security features in a switched network. This ability includes being able to meet this objectives:

- Describe switch and layer 2 security as a subject of an overall security plan
- Describe how a rouge device gains unauthorized access to a network
- Categorize switch attack types and list mitigation options
- Describe how a MAC flooding attack works to overflow a CAM Campus backbone Layer table
- Describe how port security is used to block input from devices based on Layer 2 restrictions
- Describe the procedure for configuring port security on a switch
- Describe the methods that can be used for authentication using AAA
- Describe port-based authentication using 802.1X

Overview of Switch Security issues

This topic describes switch and Layer 2 security as a subset of an overall network security plan

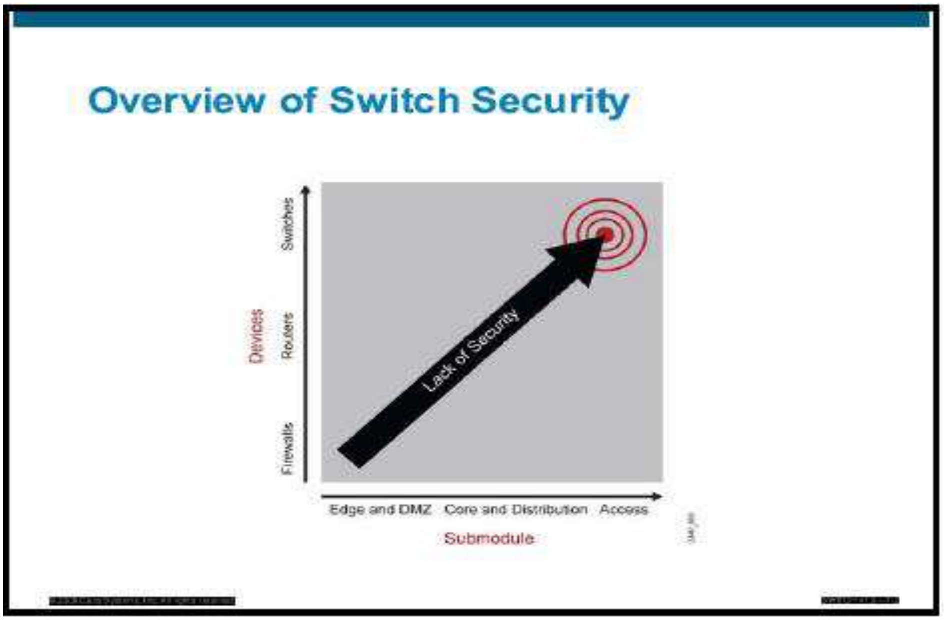

Overview of Switch Security

Much industry attention focuses on security attacks from outside the walls of an organization and at the upper Open System Interconnection (OSI) layers. Network security often focuses on edge routing devices and filtering of packets that are based on Layer 3 and Layer 4 headers, ports, stateful packet inspection, and so forth. This includes all issues related to layer 3and above, as traffic makes its way into the campus network from the internet. Campus access device and layer 2 communication are largely unconsidered in most security discussions.

The default state of networking equipment highlights this focus on external protection and internal open communication. Firewalls, placed at the organizational borders, arrive in a secure operational mode and allow no communication until they are configured to do so. Routers and switches that are internal to an organization and that are designed to accommodate communication, delivering needful campus traffic, have a default operational mode that forwards all traffic unless they are configured otherwise. Their function as devices that facilitate communication often result in minimal security configuration, and they become target for malicious attacks. If an attack is launched at layer 2 on an internal campus device, the rest of the network can be quickly compromised, often without detection.

Many security features are available for switches and routers, but they must be enabled to be effective. As with Layer 3, where security had to be tightened on devices within the campus as. Malicious activity that compromised this layer increased, now security measures must be taken to guard against malicious activity at Layer 2. A new security focus centers on attacks that are launched by maliciously using normal layer 2 switch operations. However, as with access control list (ACLs) for upper-layer security, a policy must be established and appropriate features configured to protect against potential malicious acts while maintaining daily network operations.

Security Infrastructure services

This topic describes the security design issues within an enterprise design network.

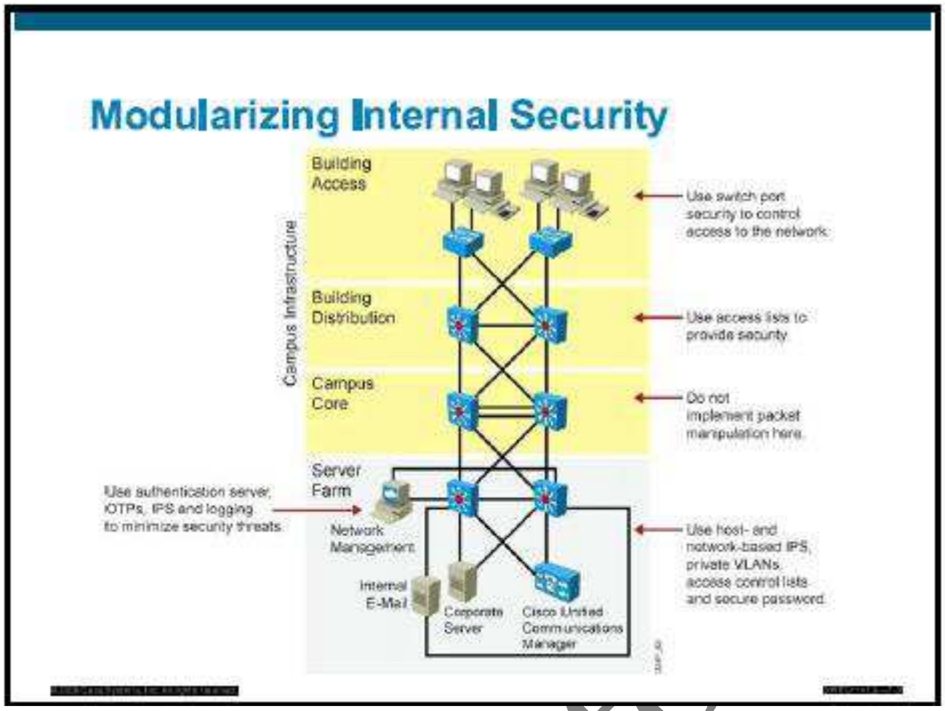

Modularizing Internal Security

Campus Infrastructure
Building Access
Building Distribution
Campus Core
Server Farm
Use switch port security to control access to the network.
Use access lists to provide security.
Do not implement packet manipulation here.
Use authentication server, OTPs, IPS and logging to minimize security threats.
Network Management
Use host- and network-based IPS, private VLANs, access control lists and secure password.
Internal E-Mail
Corporate Server
Cisco Unified Communications Manager

Security is an infrastructure service that increases the integrity of the networks by protecting network resources and users from internal and external threats. Without a full understanding of the threats that are involved, network security deployments tends to be incorrectly configured, too focused on security devices, or lacking in the appropriate threat-response options.

You can evaluate and apply security on a module-by-module basis within the Cisco Enterprise Architecture. The following are some recommended-practice security consideration for each module:

- The campus core layer in the campus infrastructure module switches packets as quickly as possible. It should not perform any security functions, because these would slow down packet switching.
- The building distribution layer performs packet filtering to keep unnecessary traffic from the campus core layer. Packet filtering at the building distribution layer is a security function because it prevents some undesired access to other modules. Given that switches in this layer are usually Layer 3 – aware multilayer switches, the building distribution layer is often the first location that can filter based on network layer information.
- At the building access layer access can be controlled at the port level with respect to the data link layer information (for example, MAC address).
- The server farm module provides application services to end users and devices. Given the high degree of access that most employees have to these servers, they often become the primary target of internally originated attacks. Use host-and network-based intrusion prevention system (IPSs), private VLANs, and access control to provide a much more comprehensive response to attacks. An comboardintrusion detection system (IDS) within multilayer switches can inspect traffic flows on the server farm modules.
- The server farm module typically includes a network management system that securely manages all devices and hosts within the enterprise architecture. Syslog provide important information regarding security violations and configuration changes by logging security related events (authentication and so on) other sever including an authentication, authorization, and accounting (AAA) security server can work in combination with the one-time password (OTP) server to provide a very high level of security to all local and remote users. AAA and OTP authentication reduce the likelihood of successful password attack.

Reason for internal security

This topic describes reasons for internal security

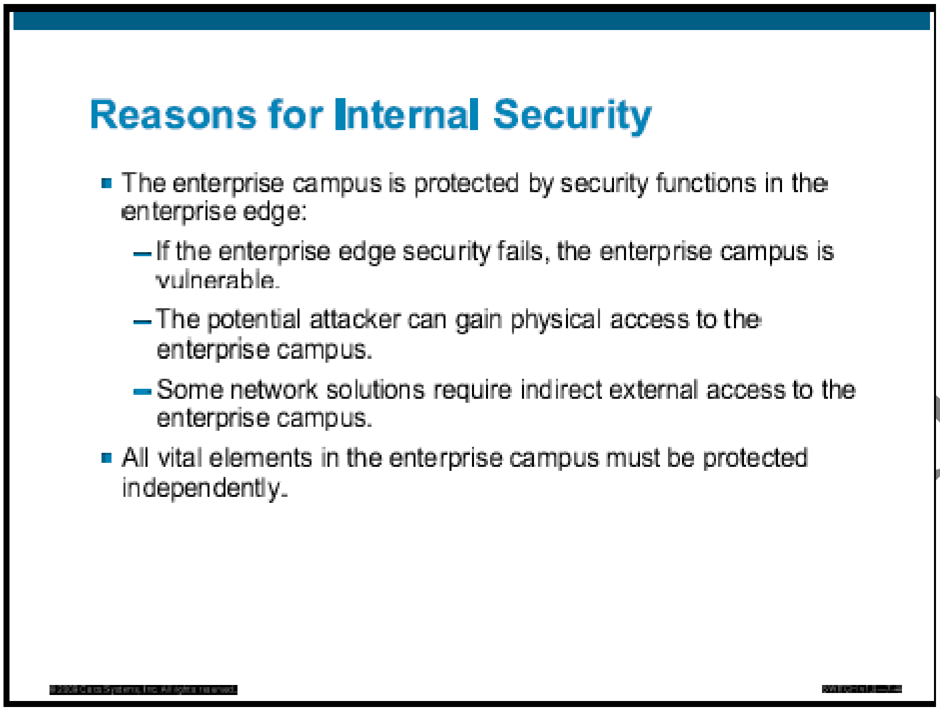

Reasons for Internal Security

- The enterprise campus is protected by security functions in the enterprise edge:
  - If the enterprise edge security fails, the enterprise campus is vulnerable.
  - The potential attacker can gain physical access to the enterprise campus.
  - Some network solutions require indirect external access to the enterprise campus.
- All vital elements in the enterprise campus must be protected independently.

- Several reasons exist for strong protection of the enterprise campus infrastructure, including security in each individual element of the enterprise campus, where usually the most strategic assets reside.
- Relying on the security that has been established at the enterprise edge fails as soon as security there is compromised. Having several layers of security increases the protection of the enerprise campus, where usually the most strategic assets reside.
- If the enterprise allows visitors into buildings, potentially an attacker gain physical access to devices in the enterprise campus. Relying on physical security is not enough.
- Very often external access does not stop at the enterprise edge. Application require at least an indirect access to the enterprise campus resources, which means that strong security is necessary.

Unauthorized Access by Rogue Devices

This topic describes how a rogue device gains unauthorized access to a network.

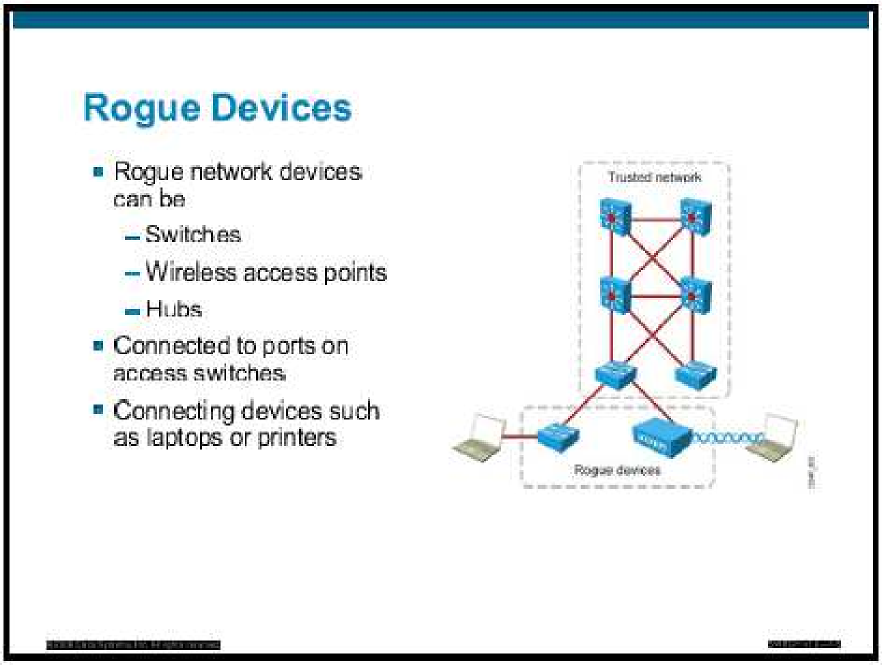

Rogue Devices

- Rogue network devices can be
  - Switches
  - Wireless access points
  - Hubs
- Connected to ports on access switches
- Connecting devices such as laptops or printers

Rogue access comes in several forms. For example, because unauthorized rogue access points are inexpensive and readily available, employees sometimes plug them into existing LANs and build ad hoc wireless networks without IT department knowledge or consent. These rogue access points can be a serious breach of network security because they can be plugged into a network port behind the corporate firewall. Because employees generally do not enable any security settings on the rogue access point, it is easy for unauthorized users to use the access point to intercept network traffic and hijack client sessions.

Malicious rogue access points, although much less common than employee-installed rogue access points, are also a security concern. These rogue access points create an unsecured wireless LAN connection that puts the entire wired network at risk. Malicious rogues present an even greater risk and challenge because they are intentionally hidden from physical and network view.

To mitigate Spanning Tree Protocol (STP) manipulation, use the **root guard** and the **BPDU guard** enhancement commands to enforce the placement of the root bridge in the network and to enforce the STP domain borders. The RootGuard feature is designed to provide a way to enforce the root bridge placement in the network. The STP BPDUGuard is designed to allow network designers to keep the active network topology predictable. Although BPDUGuard may seem unnecessary, given that the administrator can set the bridge priority to zero, there is still no guarantee that it will be elected as the root bridge, because there might be a bridge with priority zero and a lower bridge ID. BPDUGuard is best deployed toward user-facing ports to prevent rogue switch-network extensions by an attacker.

Switch Attack Categories

This topic categorizes switch attack types and lists mitigation options.

Switch Attack Categories

- MAC address-based attacks
- MAC address flooding
- VLAN attacks
- VLAN hopping
- Spoofing attacks
- Spoofing of DHCP, ARP, and MAC addressing
- Attacks on switch devices
- Cisco Discovery Protocol
- Management protocols

A device that is connected to the campus network typically launches Layer 2 malicious attacks. The attacks may originate from a physical rogue device that has been placed on the network for malicious purposes. The attack may also come from an external instruction that takes control of, and launches attacks from a trusted devices. In either case, the network sees all traffic as originating from a legitimate connected devices.

Attack that are launched against switches and Layer 2 can be grouped as follows:

- MAC layer attacks
- VLAN attacks
- Spoof attacks
- Attacks on switch devices

Significant attacks in these categories are discussed in more detail in subsequent section of the course. Each attack method is accompanied by a standard measure for mitigating the security compromise.

The table describes attack methods and the steps to mitigation.

Switch Security Concerns and Mitigation on Steps

Attack Method	**Description**	**Steps to Mitigation**
MAC Layer Attacks		
Mac address flooding	Frames with unique, invalid source MAC addresses flood the switch, exhausting content-addressable memory (CAM) table space, disallowing new entries from valid hosts. Traffic to valid hosts is subsequently flooded out all ports.	Port security, MAC address VLAN access map.
VLAN Attacks		
VLAN hopping	By altering the VLAN ID on packets that are encapsulated for trunking, an attacking devices can send or receive packets on various VLANs, bypassing Layer 3 security measures	Tighten up trunk: Configurations and the negotiation scale of unused ports. Place unused ports in a common VLAN.
Attacks between devices on a common VLAN	Devices may need protection from one another. Even though they are on a common VLAN. This is especially true on service provider segments that support devices from multiple customers.	Implement private VLANs (PVLANs).

Spoofing Attacks

DHCP starvation and DHCP spoofing	An attacking device can exhaust the address space available to the DHCP severs for a period of time or establish itself as a DHCP server in man-in-the-middle attacks.	Use DHCP Snooping
Spanning-tree compromises	Attacking devices spoofs the root bridge in the STP topology. If successful, the network attacker can see a variety of frames.	Proactively configure the primary and backup root devices Enable RootGuard.
MAC spoofing	Attacking device spoofs the MAC address of a valid host currently in the CAM table. Switch then forwards to an attacking devices any frames that are destined for the valid host.	Use DHCP snooping, port security.
Address Resolution Protocol (ARP) spoofing	Attacking device crafts ARP replies intended for valid hosts. The MAC address of the attacking device then becomes the destination address found in the layer 2 frames that were sent by the valid network device.	Use Dynamic ARP Inspection (DAI). DHCP snooping, port security.

Switch Device Attacks

Cisco Discovery Protocol manipulation	Information sent through Cisco Discovery Protocol is transmitted in clear text and unauthenticated, allowing it to be captured and to divulge network topology information.	Disable Cisco Discovery Protocol on all ports where it is not intentionally used.
Secure Shell (SSH) Protocol and Telnet attacks.	Telnet packets can be read in clear text. SSH is an option but has security issues in version 1.	Use SSH version 2. Use Telnet with vty ACLs.

MAC Flooding Attack

This topic describes how port security is used to block input from devices based on layer 2 restrictions.

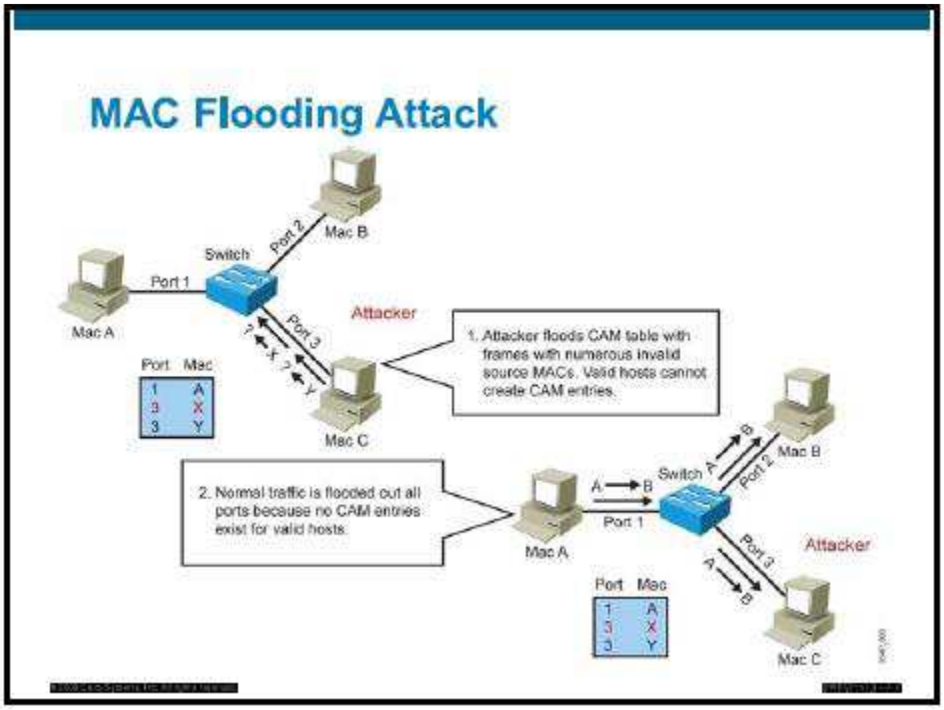

Port	Mac
1	A
3	X
3	Y

1. Attacker floods CAM table with frames with numerous invalid source MACs. Valid hosts cannot create CAM entries.

2. Normal traffic is flooded out all ports because no CAM entries exist for valid hosts.

Port	Mac
1	A
3	X
3	Y

A common Layer 2 or switch attack is MAC flooding, which results in an overflow of the CAM table of a switch. The overflow causes the flooding of regular data frames out all switch ports. This attack can be launched for the malicious purpose of collecting a broad sample of traffic or as a denial of service (DoS) attack.

The CAM tables of a switch are limited in size and therefore can contain only a limited number of entire at any one time. A network intruder can malicious flood a switch with a large number of frames from a range of invalid MAC address. If enough new entries are made before old ones expire, new valid entries will not be accepted. Then, when traffic arrives. At the switch for a legitimate device that is located on one of the switch ports that was not able to create a CAM table entry, the switch must flood the frames to that address out all ports. This has two adverse effects:

- The switch traffic forwarding is inefficient and voluminous.
- An intruding device can be connected to any switch port and capture traffic that is not normally detected on that port.

If the attack is launched before the beginning of the day, the CAM table would be full when the majority of the devices powered on. Then frames from those legitimate devices are unable to create CAM table entries as the power on. If this represents a large number of network devices, the number of MAC address that are flooded with traffic will be high, and any switch port will carry flooded frames from a large number of devices.

If the initial flood of invalid CAM table entries is a one-time event, the switch will eventually age out older, invalid CAM table entries, allowing new, legitimate devices to create entries.

Traffic flooding will cease and may never be detected, even though the intruder may have captured a significant amount of data from network.

As the figure shows, MAC flooding occurs in several steps. The table describes the progression of a MAC flooding attack.

MAC Flooding Attack Progression

step	Description
1.	Switch forwads traffic based on valid CAM table entries
2.	Attacker (MAC address C) sends out multiple packes with various source MAC addresses.
3.	Over a short time period, the CAM table in the switch fills up untill it cannot accept new entries. As long as the attack is running, the CAM table on the switch will remain full.
4.	Switch begins to flood all received packets out of every port so that from host A to host B are also flooded out of port 3 on the switch.

Suggested Mitigation for MAC Flooding Attacks

Configure port security to define the number of MAC address that are allowed on a given port. Port security can also specify which MAC address is allowed on a given port.

Port Security

This topic describes how port security is used to block input from devices based upon Layer 2 restrictions.

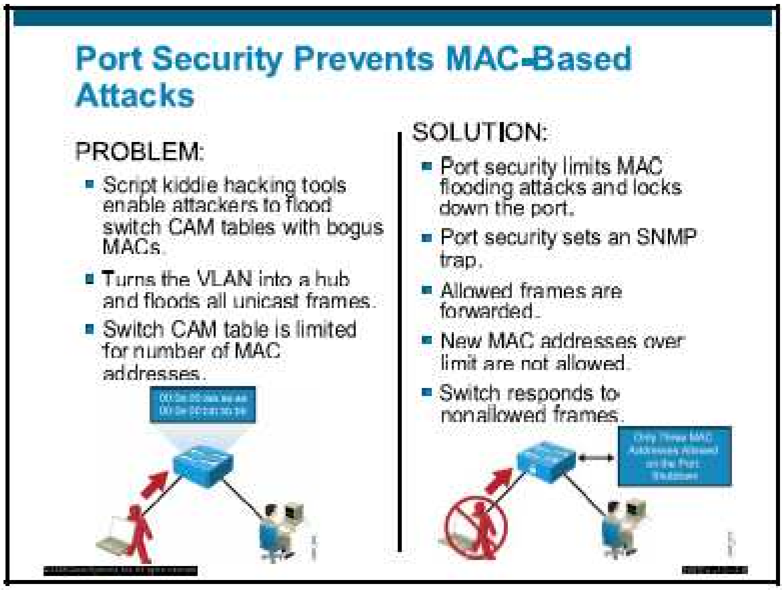

Port Security Prevents MAC-Based Attacks

PROBLEM:

- Script kiddie hacking tools enable attackers to flood switch CAM tables with bogus MACs.
- Turns the VLAN into a hub and floods all unicast frames.
- Switch CAM table is limited for number of MAC addresses.

SOLUTION:

- Port security limits MAC flooding attacks and locks down the port.
- Port security sets an SNMP trap.
- Allowed frames are forwarded.
- New MAC addresses over limit are not allowed.
- Switch responds to nonallowed frames.

Port security, a features that is supported on Cisco Catalyst switches, restricts a switch port to a specify set or number of MAC addresses. Those addresses can be learned dynamically or configured statically. The port will then provide access to frames from only those addresses. If, however, the number of addresses is limited to four but no specific MAC addresses are configured, the port will allow any four MAC addresses to be learned dynamically, and port access will be limited to those four dynamically learned addresses.

A port security features called “sticky learning,” available on some switch platforms, combines the features of dynamically learned and statically configured addresses. When this features is configured on an interface, the interface converts dynamically learned addresses to “sticky secure” addresses. This adds them to the running configuration as if they were configured with the switchport port-security mac-address command.

Scenario

Imagine five individuals whose laptops are allowed to connect to a specify port when they visit an area of the building. You want to restrict switch port access to the MAC addresses of those five laptops and allow no addresses to be learned dynamically on that port.

Process

The table describes the process that can achieve the desired result for this scenario.

Implementing Port Security

Step	Action	Notes
1.	Configure port security	Configure port security to allow only five connection on that port. Configure an entry for each of the five allowed MAC addresses. This configuration, in effect, populates the MAC address table with five entries for that port and allows no additional entries to be learned dynamically.
2.	Allowed frames are processed.	When frames arrive on the switch port, their source MAC address is checked against the MAC address table. If the frame source MAC address matches an entry in the table for that port, the frames are forwarded to the switch to be processed like any other frames on the switch.
3.	New addresses are not allowed to create new MAC address table entries.	When frames with a non-allowed MAC address arrive in the port, the switch determines that the address is not in the current MAC address table and does not create a dynamically entry for the new MAC address, because the number of allowed addresses has been limited.
4.	Switch takes action in response to non-allowed frames.	The switch will disallow access to the port and take one of these configuration – dependent actions: (a) the entries switch port can be shut down; (b) access can be denied for that MAC address only and a log error can be generated; (c) access can be denied for that MAC address but without generating a log message.

Note: Port security cannot be applied to trunk ports where addresses might change frequently. Implementation of the port security vary be Cisco Catalyst platform. Check documentation to determine whether particular hardware supports this features and how the hardware supports the features.

Configure Port Security

This topic explains the procedure for configuring port security on a switch.

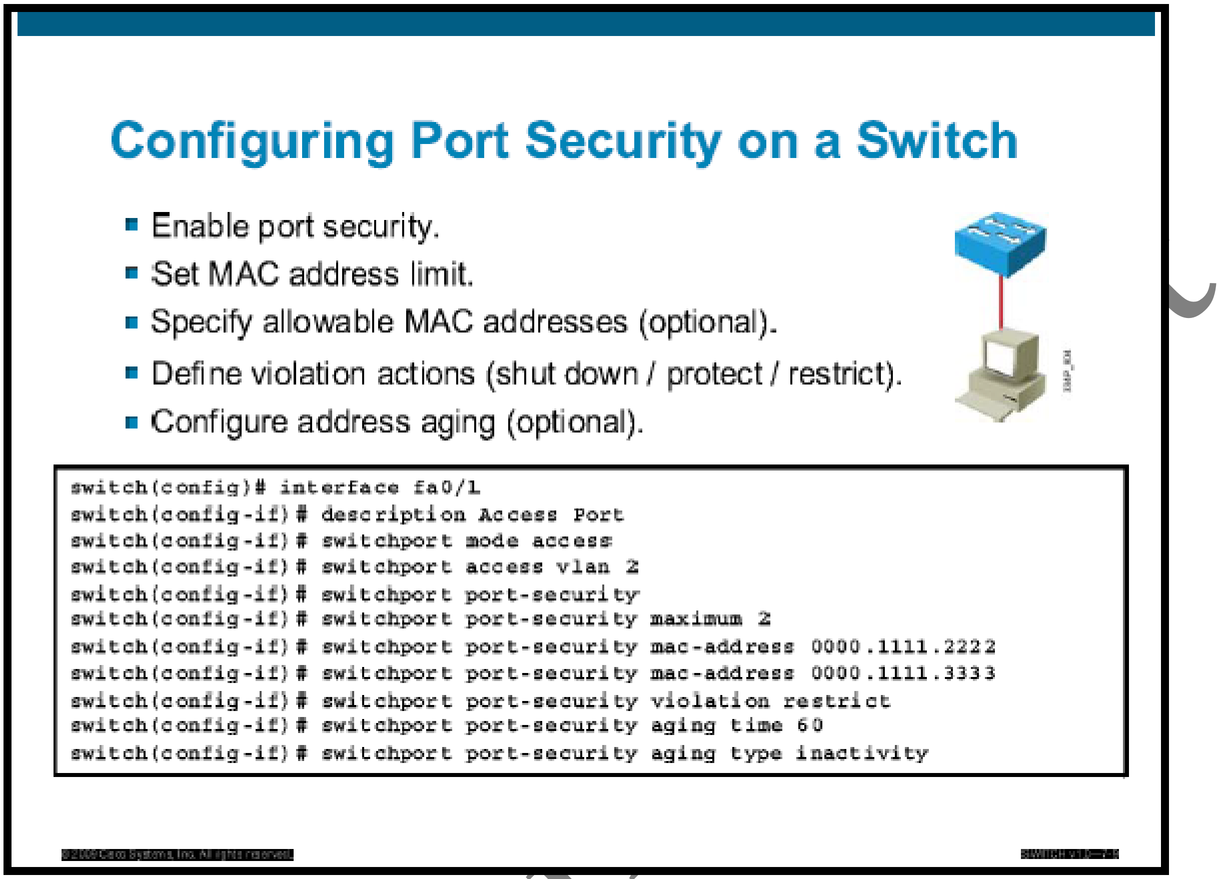

Configuring Port Security on a Switch

- Enable port security.
- Set MAC address limit.
- Specify allowable MAC addresses (optional).
- Define violation actions (shut down / protect / restrict).
- Configure address aging (optional).

```
switch(config)# interface fa0/1
switch(config-if)# description Access Port
switch(config-if)# switchport mode access
switch(config-if)# switchport access vlan 2
switch(config-if)# switchport port-security
switch(config-if)# switchport port-security maximum 2
switch(config-if)# switchport port-security mac-address 0000.1111.2222
switch(config-if)# switchport port-security mac-address 0000.1111.3333
switch(config-if)# switchport port-security violation restrict
switch(config-if)# switchport port-security aging time 60
switch(config-if)# switchport port-security aging type inactivity
```

SWITCH v1.0—7-8

To configure port security so that switch port access is limited to a finite number and a specific set of end-device MAC addresses, follow the steps in the table.

To configure port security, follow the steps listed in the table.

Port Security Configuration Steps

Step	Description
1.	Enables port security. Switch (config-if) # **switchport port-security**
2.	Sets a maximum number of MAC addresses that will be allowed on this port. The default is one. Switch(config-if) # **switchport port-security maximum** value
3.	(optional) Specifies which MAC addresses will be allowed on this port Switch(config-if) # **switchport port-security mac-address** mac-address Switch(config-if) # **switchport port-security mac-address** mac-address
4.	Defines what action an interface will take if a nonallowed MAC address attempts access. Switch(config-if) # **switchport port-security violation** { shutdown | restrict | protect }

Caveats to Port Security Configuration Steps

There are some caveats to bear in mind:

Step 1. Port security is enabled on a port-by-port basis.

Step 2. By default, only one MAC address is allowed access through a given switch port when port security is enabled. This parameter increases that number. It implies no restriction on specific MAC addresses, but only on the total number of addresses that can be learned by the port. Learned addresses are not aged out by default but can be configured to do so after a specified time when you use the **switchport port-security aging** command. The value parameter can be any number from 1 to 1024, with some restrictions related to the number of ports on a given switch with port security enabled.

Note: - Be sure to set the value parameter to a value of 2 when you are configuring a port to support VOIP with a phone and computer that are accessible on the port. If the default value is used, a port-security violation will result.

Step 3. Access to the switch port can be restricted to one or more specific MAC addresses. If the number of specific MAC addresses that are assigned when you use this command is lower than the value parameter that you set in Step 2, then the remaining allowed addresses can be learned can be learned dynamically. If you specify a set of MAC addresses that is equal to the maximum number allowed, access is limited to that set of MAC addresses.

Step 4. By default, if the maximum number of connections is achieved and a new MAC address attempts to access the port, the switch must take one of these actions:

- **Protect:** Frames from the nonallowed address are dropped, but there is no log of the violation.

Note: - The protect argument is platform or version dependent.

- **Restrict:** Frames from the nonallowed address are dropped, a log message is created, and a Simple Network Management Protocol (SNMP) trap is sent.
- **Shut down:** If any frames are detected from a nonallowed address, the interface is errdisabled, a log entry is made, an SNMP trap is sent, and manual intervention or errdisable recovery must be used to make the interface usable.

Verifying Port Security

This subtopic describe how to verify port security.

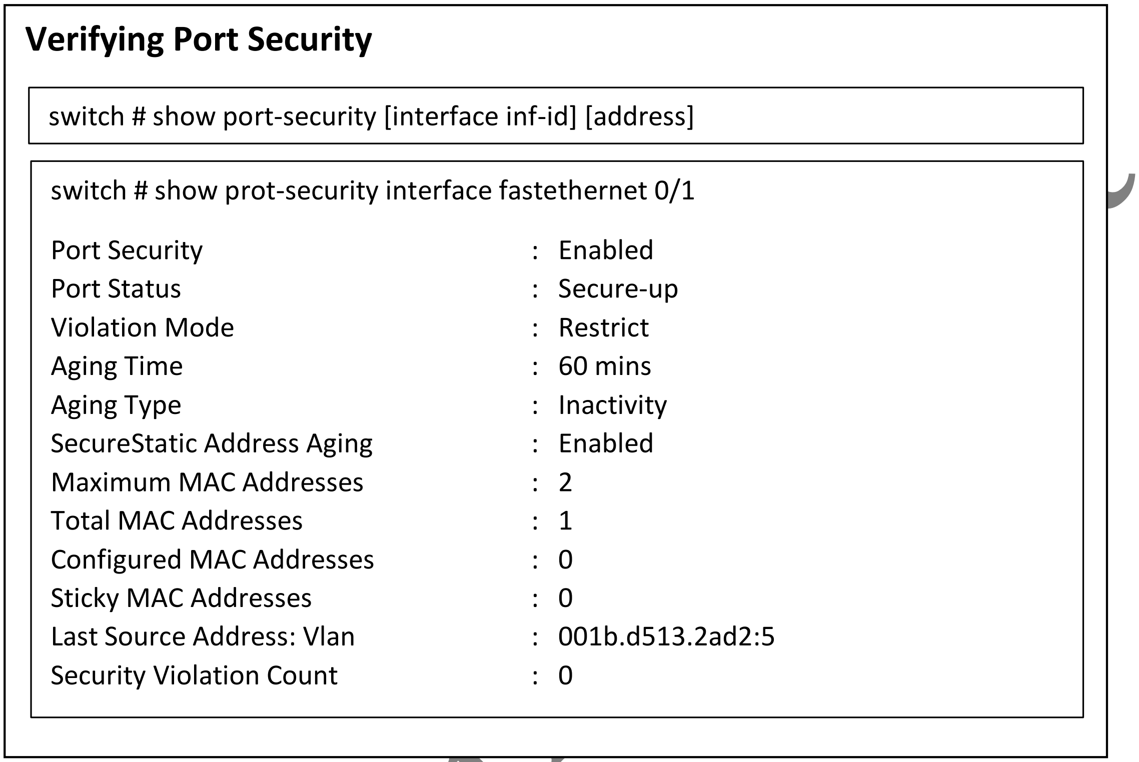

Verifying Port Security

```
switch # show port-security [interface inf-id] [address]
```

```
switch # show prot-security interface fastethernet 0/1

Port Security                       :  Enabled
Port Status                         :  Secure-up
Violation Mode                      :  Restrict
Aging Time                          :  60 mins
Aging Type                          :  Inactivity
SecureStatic Address Aging          :  Enabled
Maximum MAC Addresses               :  2
Total MAC Addresses                 :  1
Configured MAC Addresses            :  0
Sticky MAC Addresses                :  0
Last Source Address: Vlan           :  001b.d513.2ad2:5
Security Violation Count            :  0
```

You can use the **show port-security** command to verify the ports on which port security has been enabled. It also displays count information and security actions to be taken per interface.

The full command syntax is as follows:

Switch # **show port-security** [interface inif_id] **address**

Arguments are provided to view the port security status by interface or to view the addresses that are associated with port security on all interfaces.

Use the interface argument to provide output for a specific interface.

Verifying Port Security (cont.)

```
Switch# show port-security
Secure  Port   MaxSecureAddr    CurrentAddr    SecurityViolation     Security Action
                  ( Count )          ( Count )          ( Count )
----------------------------------------------------------------------------------------------------------------
   Fa 0/1               2                    1                   0                     Restrict
----------------------------------------------------------------------------------------------------------------
Total Addresses in System ( excluding on mac per port )              :  0
Max Addresses limit in System ( excluding one mac per port )      :  6144
```

```
Switch # show port-security address
                     Secure Mac Address Table
-----------------------------------------------------------------------------------------------------------------------------
Vlan      Mac Address                  Type                                         Ports     Remaining Age
                                                                                                         (mins)
-------    -----------------           --------                                      --------   --------------------
   2       001b.d513.2ad2          SecureDynamic                            Fa0/1          60 (I)
----------------------------------------------------------------------------------------------------------------------------
Total Addresses in System (excluding one mac per port)              :  0
Max Addresses limit in System (excluding one mac per port)       :  6144
```

You can use the show port-security command to verify the ports on which port security has been enabled. It also displays count information and security actions to be taken per interface.

The full command syntax is as follows:

Switch # **show port-security** [interface inif_id] **address**

Arguments are provided to view the port security status by interface or to view the addresses that are associated with port security on all interfaces.

Use the address argument to display MAC address table security information. The remaining age column is populated only when it is specifically configured for a given interface.\

The example displays output from the **show port-security address** privileged EXEC command.

Port security with Sticky Mac Addresses

This subtopic describes the sticky MAC option with port security.

Configuring Sticky MAC Addresses

```
Switch (config) # interface fa0/1
Switch (config-if) # switchport prot-security mac-address sticky
```

```
Switch # show port-security address
Secure MAC Address Table
-----------------------------------------------------------------------------------------------------------------------------
Vlan    Mac Address          Type                        Ports      Remaining Age
                                                                        (mins)
------  -----------------    -------                     -------    --------------------
  2     001b.d513.2ad2       SecureSticky                Fa0/1            -
```

```
Switch # show running-config fastethernet 0/1
Interface FastEthernet0/1
          Switchport access vlan 2
          Switchport mode access
          Switchport port-security maximum 2
          Switchport port-security
          Switchport port-security violation restrict
          Switchport port-security mac-address sticky
          Switchport port-security mac-address sticky 001b.d513.2ad2
```

Port security can be used to mitigate spoof attacks by limiting access through each switch port to a single MAC address. This prevents intruders from using multiple MAC addresses over a short time period but does not limit port access to a specific MAC address. The most restrictive port security implementation would specify the exact MAC address of the single device that is to gain access through each port. Implementing this level of security, however, requires considerable administrative overhead.

Port security has a feature called sticky MAC addresses that can limit switch port access to a single, specific MAC address without the network administrator having to gather the MAC address of every legitimate device and manually associate it with a particular switch port.

When sticky MAC addresses are used, the switch port will convert dynamically learned MAC addresses to sticky MAC addresses and subsequently add them to the running configuration as if they were static entries for a single MAC address to be allowed by port security. Sticky secure MAC addresses will be added to the running configuration is copied to the startup configuration after addresses have been learned. If they are saved in the startup configuration, they will not have to be relearned upon switch reboot, and this provides a higher level of network security.

The command that follows will convert all dynamic port-security learned MAC addresses to sticky secure MAC address: **Switchport port-security mac-address sticky**

This command cannot be used on ports where voice VLANs are configured.

Authentication and Authorization Methods

This topic describes security in a multilayer switched network.

AAA Network Configuration

- Authentication
  - Verifies a user identify
- Authorization
  - Specifies the permitted tasks for the user
- Accounting
  - Provides billing, auditing, and monitoring

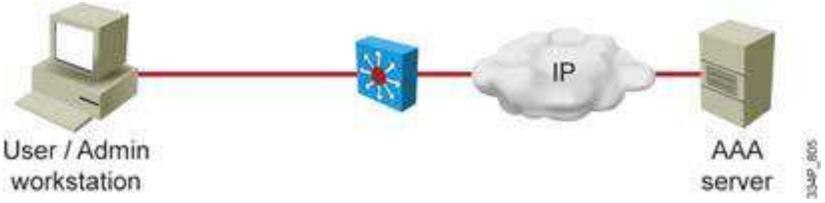

Authentication, authorization, and accounting (AAA) network security services provide the primary framework through which access control is set up on a switch. AAA is an architectural framework for configuring a set of three independent security functions in a consistent manner. AAA provides a modular way of performing these services. For purposes of this course, only authentication will be discussed.

Authentication is the way that a user is identified before being allowed access to the network and network services. You configure AAA authentication by defining a list of named authentication methods and then applying that list to various interfaces. The method list defines the types of authentication to be performed and the sequence in which they will be performed; it must be applied to a specific interface before any of the defined authentication methods will be performed.

The only exception is the default method list (which is named “default”). The default method list is automatically applied to all interfaces if no other method list is defined. A defined method list overrides the default method list.

In many circumstances, AAA uses protocols such as RADIUS, TACACS+, or IEEE 802.1X to administer its security functions. If the switch is acting as a network access server, AAA is the means through which the switch establishes communication between the network access server and the RADIUS, TACACS+, or 802.1X security server.

Local User Authentication with AAA

This subtopic describes how to configure user authentication with AAA.

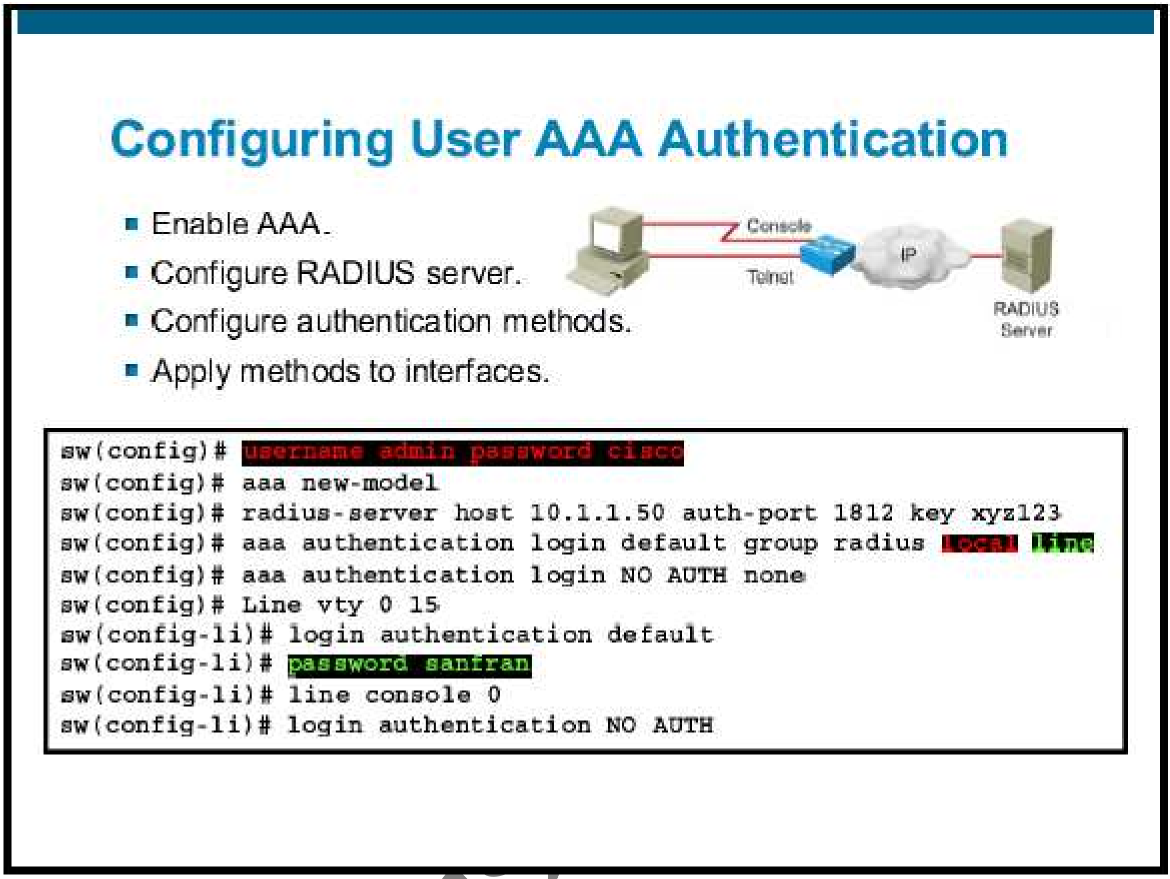

Configuring User AAA Authentication

- Enable AAA.
- Configure RADIUS server.
- Configure authentication methods.
- Apply methods to interfaces.

Console
Telnet
IP
RADIUS Server

```
sw(config)# username admin password cisco
sw(config)# aaa new-model
sw(config)# radius-server host 10.1.1.50 auth-port 1812 key xyz123
sw(config)# aaa authentication login default group radius local line
sw(config)# aaa authentication login NO_AUTH none
sw(config)# Line vty 0 15
sw(config-li)# login authentication default
sw(config-li)# password sanfran
sw(config-li)# line console 0
sw(config-li)# login authentication NO_AUTH
```

To configure user authentication, start by using the **aaa new-model** command. As soon as you enter this command, the default authentication methods do not apply any more, and you need to define how access is granted to the console or vty lines.

In this example, several methods are used.

The first method is the default method, specified by the line **aaa authentication login default**. This line defines that, by default, the RADIUS group should be used to authenticate users who are trying to log in to the local device. If the servers that are defined in the RADIUS group do not answer, local authentication is possible. If local authentication is not possible (because no local user was defined), line authentication is used. Several methods are allowed, but the main method here is RADIUS. Local authentication will be used only if the RADIUS servers are unreachable. Similarly, line authentication will be used only if local authentication is not possible. If RADIUS servers respond and authentication fails, the local or line method will not be used and the user will be denied access. If the RADIUS servers do not answer and a local user is defined, local authentication will be used (it may succeed or fail) and line authentication will not be used.

Because this default method defines “group RADIUS” as the first default method, one or several RADIUS servers must be defined. This example defines 10.1.1.50 as the RADIUS server, listening on UDP port 1812. The shared secret used to be allowed to query the RADIUS is xyz123.

Because the default method also allows local authentication, a local username and password pair are defined. You can create several users.

Because the default method also allows line authentication, a password is defined on the vty lines.

In a real implementation, you would probably use two of these methods but not all three. If local authentication is allowed, there is little purpose in defining a possible line authentication. All three methods are mentioned here to show that several backup methods can be used in combination with a primary authentication method. Each method requires specific configuration.

In the example, a second **aaa authentication login** mechanism is used. It is called NO_AUTH; and is defined as needing no authentication when you use this method.

On vty lines, the default method is used, thus requiring RADIUS authentication, and then local or line authentication as secondary or tertiary methods.

On the console line, the NO_AUTH method is called, which implies that no authentication is needed when you are connecting through the console.

You can see in this example that “default” is the name of an authentication method. This name is special in the sense that all lines requiring authentication will use the default method if no other specific method is called on the corresponding line. This means, for example, that if you did not specify a method on the console line, it would have used the default method. You still have to specify login requirements. In the previous example, if you remove the line login authentication NO_AUTH completely, no login requirement is specified for the console line, and no console connection is possible anymore. If you specify “login” without defining the method to use, the “default” method is called by default.

802.1 X Port-Based Authentication

This topic describes IEEE 802.1X port-based authentication.

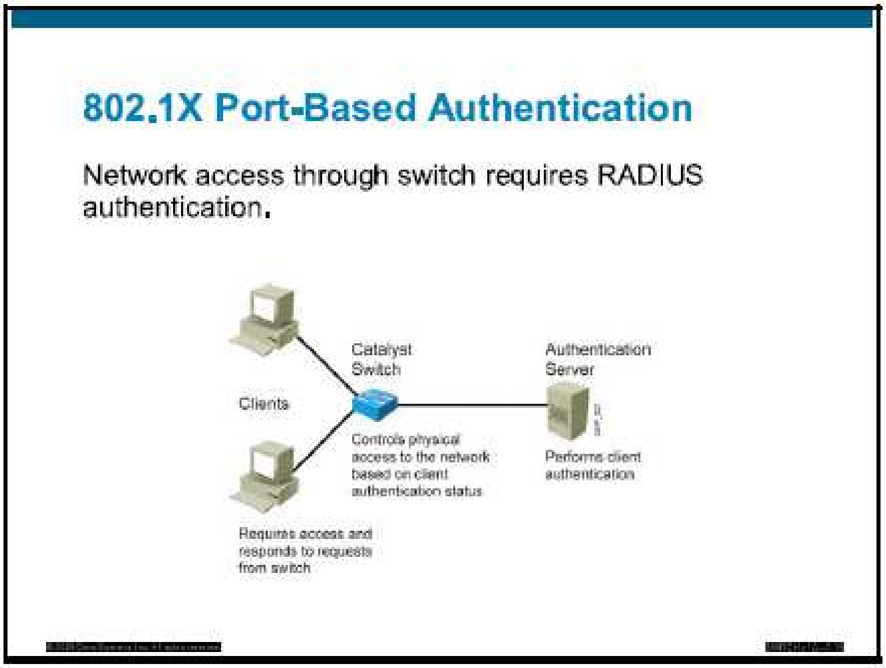

802.1X Port-Based Authentication

Network access through switch requires RADIUS authentication.

Clients

Requires access and responds to requests from switch

Catalyst Switch

Controls physical access to the network based on client authentication status

Authentication Server

Performs client authentication

The 802.1X standard defines a port-based access control and authentication protocol that restricts unauthorized workstation from connecting to a LAN through publicly accessible switch ports. The authentication server authenticates each workstation that is connected to a switch ports before making available any services that are offered by the switch or the LAN.

Until the workstation is authentication, 802.1X access control allows only Extensible Authentication Protocol over LAN (EAPOL) traffic through the port to which the workstation is connected. After authentication succeeds, normal traffic can pass through the port.

With 802.1X port-based authentication, the devices in the network have specific roles, as follows:

- **Client:** The devices (workstation) that requests access to LAN and switch services, and then respond to request from the switch. The workstation must be running 802.1X – complaint client software, such as what is offered in the Microsoft Windows XP operating system. (The port that the client is attached to is the supplicant [client] in the 802.1X specification)
- **Authentication server:** Perform the actual authentication of the client. The authentication server validates the identity of the client and notifies the switch whether or not the client is authorized to access the LAN and switch services. Because the switch acts as the proxy, the authentication service is transparent to the client. The RADIUS security system with Extensible Authentication Protocol (EAP) extension is the only supported authentication server.

Switch (also called the authenticator): Controls physical access to the network based on the authentication status of the client. The switch acts as an intermediary (proxy) between the client (supplicant) and the authentication server, requesting identifying information from the client, verifying that information with the authentication server, and relaying a response to the client. The switch uses a RADUIS software agent, which is responsible for encapsulation and decapsulating the EAP frames and interacting with the authentication server.

The switch port state determines whether the client is granted access to the network. The port starts in the unauthorized state. While in this state, the port disallow ingress and egress traffic except for 802.1X protocol packets. When a client is successfully authentication, the port transition to the authorized state, allowing all traffic for the client to flow normally.

If the switch request the client identity (authenticator initiation) and the client does not support 802.1X, the port remains in the unauthorized state, and the client is not granted access to the network.

In contrast, when an 802.1X – enabled client connects to a port and the client initiates the authentication process (supplicant initiation) by sending the EAPOL start frame to a switch that is not running 802.1X protocol, no is received, and the client begins sending frames as if the port is in the authorized state.

You control the port authorized state by using the **dot1x port-channel** interface configuration command and these keywords:

- **Forced-authorized:** Disable 802.1X port-based authentication and causes the port to transition to the authorized state without any authentication exchange required. The port transmits and receive normal traffic without 802.1X – based authentication of the client. This is default setting.
- **Forced-unauthorized:** Causes the port to remain in the unauthorized state, ignoring all attempts by the client to authenticate. The switch cannot provide authentication services to the client through the interface.
- **Auto:** Enables 802.1X port-based authentication and causes the port to begin in the unauthorized state, allowing only EAPOL frames to be sent and receive through the port. The authentication process begins when the link state of the port transition from down to up (authenticator initiation) or when an EAPOL-start frame is receive (supplicant initiation). The switch request the identity of the client and begins relaying authentication messages between the client and the authentication server. The switch uniquely identifies each client that is attempting to access the network by using the client MAC address.

If the client is successfully authenticated (that is, if it receives an “accept” frame from the authentication server), the port state changes to authorized, and all frames from the authenticated client are allowed through the port.

If the authentication fails, the port remains in the unauthorized state, but authentication can be retried. If the authentication server cannot be reached, the switch can retransmit the request. If no response is received from the server after the specific number of attempts, authentication fails, and network access in not granted.

When a client logs out, it send an EAPOL-logout message, causing the switch port to transition to the unauthorized state.

Configuring 802.1X

Implementing 802.1X Port-Based Authentication

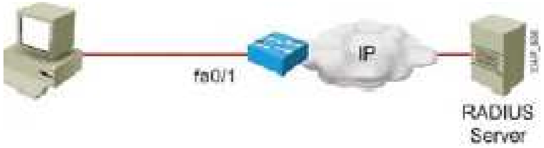

- Enable AAA
- Configure RADIUS server.
- Enable 802.1X globally.
- Configure interface for 802.1X

```
sw (config) # aaa new-model
sw (config) # radius-server host 10.1.1.50 auth-port 1812 key xyz123
sw (config) # aaa authentication dot1x default group radius
sw (config) # dot1X system-auth-control
sw (config) # interface fa0/1
sw (config-if) # description Access Port
sw (config-if) # switchport mode access
sw (config-if) # dot1X port-control auto
```

Steps	Description
1.	Enable AAA sw (config) # aaa new-model
2.	Create an 802.1X port-based authentication method list. sw (config) # aaa authentication dot1x {default} mehtod1 [method 2. . .]
3.	Globally enable 802.1X port-based authentication sw (config) # dot1X system-auth-control
4.	Enter interface configuration mode and specify the interface to be enabled for 802.1X port-based authentication. sw (config) # interface type slot/port
5.	Enable 802.1X port-based authentication on the interface. Switch (config-if) # dot1X port-control auto
6.	Return to priviledge EXEC mode sw (config) # end

Be aware the entering the **aaa new-model** command disables the standard authentication process on the switch. You also need to redefine user login policies as specified in the earlier subtopic “Local User Authentication with AAA.”

Summary

Summary

This topic summarizes the key point that were discussed in this lesson.

- Layer 2 security measures must be taken as a subset of the overall network security plan.
- Rouge devices can allow access to the network and undermine the security.
- Switch attacks fall into four main categories.
- MAC flooding attacks are launched against Layer 2 access switches and can cause the CAM table to overflow.
- Port security can be configured at Layer 2 to block input from devices.
- Sticky MAC addresses allow port security to limit access to a specify, dynamically learned MAC address.
- AAA can be used for authentication on a multilayer switch.
- 802.1X port-based authentication can mitigate risk of rouge devices gaining unauthorized access.

Protect against VLAN Attacks

Overview

On networks using trunking protocols, there is a possibility of rouge traffic “hopping” from one VLAN to another, thereby creating security vulnerabilities. These VLAN hopping attacks are best mitigate by close control of trunk links.

You can configure close control of trunk links to mitigate VLAN hopping attacks and configure VLAN access control lists (VACLs) to filter traffic within a VLAN.

Private VLANs (PVLANs) can be configured to establish security region within a single VLAN without subnetting, and VACLs can be used to filter traffic within a VLAN.

Objectives

Upon completing this lesson, you will be able to confirm various features to prevent VLAN hopping and to address VLAN security issues. This ability includes being able to meet these objectives:

- Describe how VLAN hopping occurs and why it is a security vulnerability
- Explain the procedure for configuring a switch to mitigate VLAN hopping attacks
- Describe VACLs and their purpose a part of VLAN security
- Explain the procedure for configuring VACLs

Explaining VLAN Hopping

This topic describes how VLAN hopping occurs and why it is a security vulnerability.

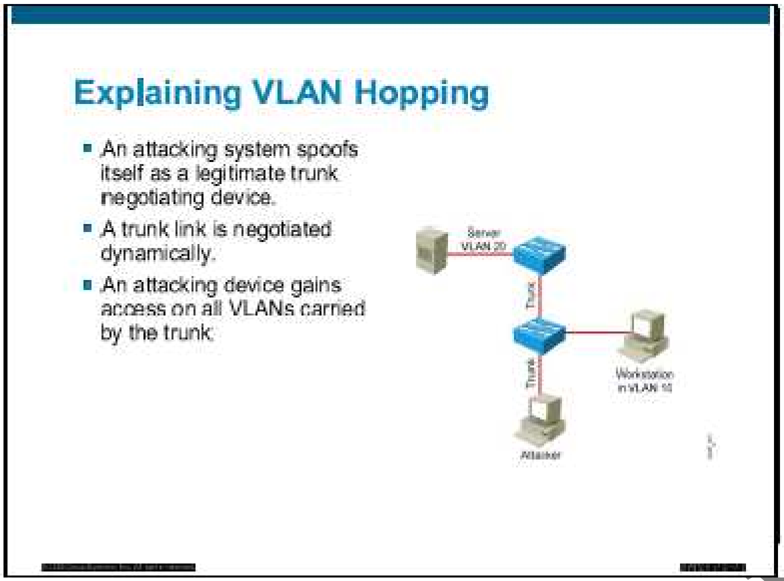

Explaining VLAN Hopping

- An attacking system spoofs itself as a legitimate trunk negotiating device.
- A trunk link is negotiated dynamically.
- An attacking device gains access on all VLANs carried by the trunk.

VLAN hopping is a network attack whereby an end system send packets to, or collects packets from, a VLAN that should not be accessible to that end system. This is accomplished by tagging the invasive traffic with a specific VLAN ID (VID) or by negotiating a trunk link to send or receiving traffic on penetrated VLANs. VLAN hopping can be accomplished by switch spoofing or double tagging.

Spoofing

In a switch spoofing attack, the network attacker configure a system to spoof itself as a switch.
The attack emulates Inter-Switch Link (ISL) or IEEE 802.1Q signaling along with Dynamic Trunking Protocol (DTP). This is signaling in an attempt to establish a trunk connection to a switch.

Any switch port configured as DTP auto, upon receipt of a DTP packet generated by the attacking device, may become a trunk port and thereby accept traffic that is destined for any VLAN supported in the trunk. The malicious device can then send packets to, or collect packets from, any VLAN that is carried on the negotiated trunk.

The table describes the switch spoofing sequence of events.

1.	Attacker gains access to a switch port and sends DTP negotiating frames towards a switch with DTP running and auto negotiation turned on (often, the default settings).
2.	Attacker and switch negotiate trunking over the port.
3.	Switch allows all VLANs (default) to transverse the trunk link.
4.	Attacker sends data to, or collects data from, all VLANs carried on that trunk.

VLAN Hopping with Double Tagging

This subtopic describes double tagging as a means of VLAN hopping.

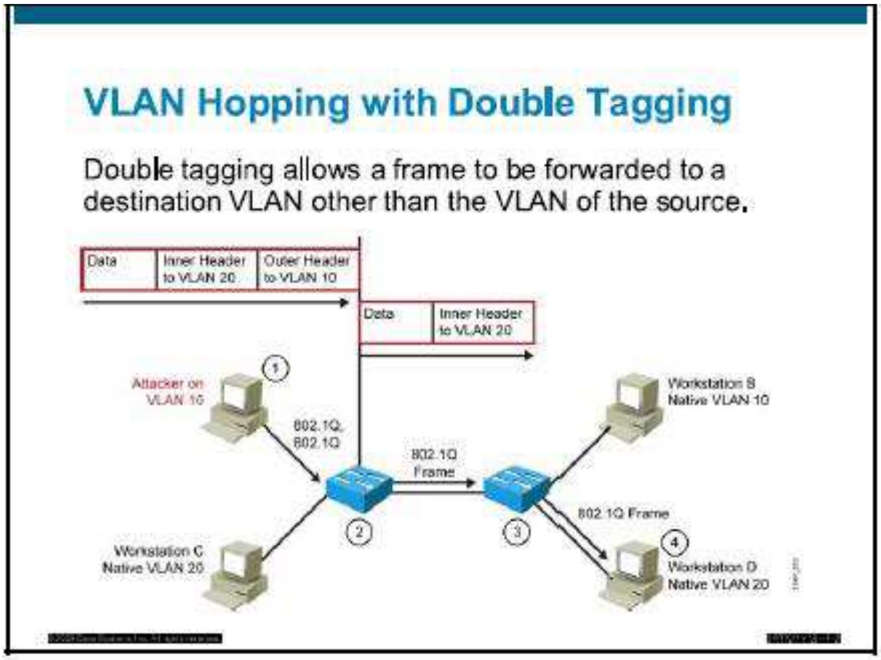

VLAN Hopping with Double Tagging

Double tagging allows a frame to be forwarded to a destination VLAN other than the VLAN of the source.

Another method of VLAN hopping is for any workstation to generate frames with two 802.1Q headers to cause the switch to forward the frames onto a VLAN that would be inaccessible to the attacker through legitimate means.

The first switch to encounter the double-tagged frame strips the first tag off the frame, because the first tag (VLAN 10) matches the trunk port native VLAN, and then forwards the frame out.

The result is that the frame is forwarded, with the inner 802.1Q tag, out all switch ports, including trunk ports that are configured with the native VLAN of the network attacker. The second switch than forwards the packets to the destination based on the VLAN ID in the second 802.1Q header. If the trunk does not match the native VLAN of the attacker, the frame is untagged and is flooded to only the original VLAN.

The table describes the double-tagging method of VLAN hopping.

Double-Tagging Method of VLAN Hopping

Steps	Description
1.	Workstation A (native VLAN 10) sends a frame with two 802.1Q headers to switch 1.
2.	Switch 1 strips the outer tag and forwards the frame to all ports within same native VLAN.
3.	Switch 2 interprets frame according to information in the inner tag marked with VLAN ID 20.
4.	Switch 2 forwards the frame out all ports associated with VLAN 20, including trunk ports.

Mitigating VLAN Hopping

This topic describes how to mitigate VLAN hopping attacks.

Mitigate VLAN Hopping
Unused ports

- Shut down all unused ports.
- Configure all unused port to access mode.
- Configure an access VLAN on all unused ports to an unused VLAN.
- Configure a native trunk VLAN on all unused ports to an unused VLAN.

Trunk Ports

- Configure a trunk port with trunk mode on, and disable trunk negotiation.
- Configure a native trunk VLAN on trunk ports to an unused VLAN.
- Configure the allowed VLANs on the trunk ports, and do not allow a native VLAN.

The measures for defending the network from VLAN hopping are a series of best practice configuring all switch ports and parameter when establishing a trunk port.

- Configure all unused ports as the access ports so that trunking cannot be negotiated across those links.
- Configure all unused port in the shutdown state and associate them with a VLAN that is designed for only unused ports, carrying no user traffic.
- When establishing a trunk link, purposefully configure arguments to achieve the following results:
- The native VLAN is different from any data VLANs.
- Trunking is set up as "on" rather than as "negotiated".
- The specific VLAN range is carried on the trunk.

VLAN Access Control Lists

Access control lists (ACLs) are useful for controlling access in a multilayer switched network. The topic describes VACLs and their purpose as part of VLAN security.

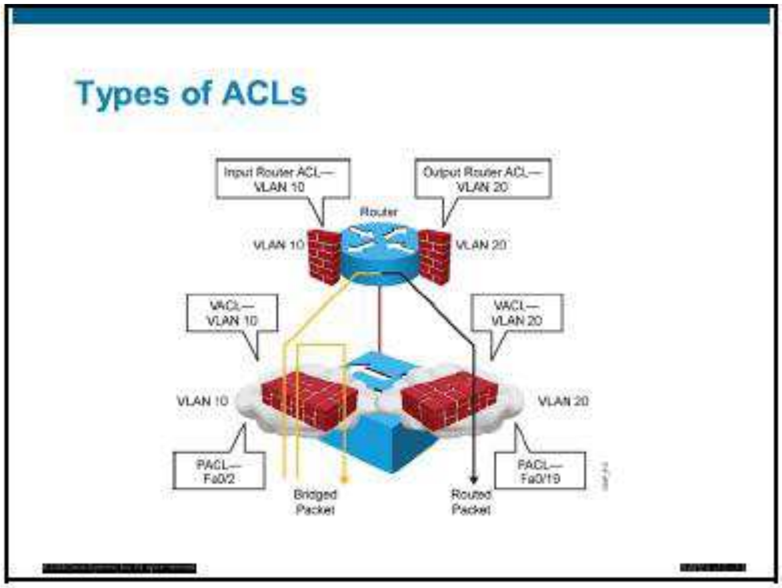

Types of ACLs

Cisco multiplayer switches support three types of ACLs:

- **Router ACL:** Support in the ternary content addressable memory (TCAM) hardware on Cisco multiplayer switches. In Cisco Catalyst switches, a router ACL can be applied to any routed interface, such as a switch virtual interface (SVI) or a Layer 3 routed port.
- **Port access control list (PACL):** Filters traffic at the port level. PACLs can be applied on a Layer 2 switch port, or EhterChannel port.
- **VACL:** Supported in software on Cisco multiplayer switches.

Catalyst switches support four ACL lookups per packets; input and output security ACL, and input and output quality of service (QoS) ACL.

Catalyst switches use two methods of performing a merge, ACLs are transformed from a series of order-dependent actions to a set of order-independent masks and patterns. The resulting access control entry (ACE) can be very large. The merge is processor and memory intensive.

An order-dependent merge is a recent improvement on some Catalyst switches in which ACLs retain their order-dependent aspects. The computation is much faster and is less processor-intensive.

Router ACLs are supported is hardware through IP standard ACLs and IP extended ACLs, with permit and deny actions. ACL processing is an intrinsic part of the packet-forwarding process. ACL entries are programmed in hardware. Lookups occur in the pipeline, whether ACLs are configured or not. With router ACLs, access list statistics and logging are not supported.

Configuring VACLs

This topic describes how to configure VACLs.

Configuring VACLs

- Create an access list.
- Configure an access map.
- Create a VLAN filter.
- Example: Drop all traffic from network 10.1.9.0/24 on VLAN 10 and 20, and drop all traffic to backup server 0000.1111.4444.

```
Switch (config)# access-list 100 permit ip 10.1.9.0  0.0.0.255 any
Switch (config)# mac access-list extended BACKUP SERVER
Switch (config-ext-mac)# permit any host 0000.1111.4444
Switch (config)# vlan access-map XYZ 10
Switch (config-map)# match ip address 100
Switch (config-map)# action drop
Switch (config-map)# vlan access-map XYZ 20
Switch (config-map)# match mac address BACKUP SERVER
Switch (config-map)# action drop
Switch (config-map)# vlan access-map XYZ 30
Switch (config-map)# action forward
Switch (config)# vlan filter XYZ vlan-list 10,20
```

VACLs (also called VLAN access maps in Cisco IOS Software) apply to all traffic on the VLAN. You can configure VACLs for IP and for MAC-layer traffic.

VACLs follow route-map conventions, in which map sequences are checked in order.

When a matching permit ACE is encountered, the switch takes the action. When a matching deny ACE is encountered, the switch checks the next ACL in the sequence or checks the next sequence.

Three VACL actions are permitted:

- Permit (with capture, Cisco Catalyst 6500 Series Switches only)
- Redirect (Catalyst 6500 Series Switches only)
- Deny (with logging, Catalyst 6500 Series Switches only)

The VACL capture option copies traffic to specify capture ports. VACL ACEs that are installed in hardware are merged with Router ACLs and other features.

Two features are supported on only the Catalyst 6500 Series Switches:

- **VACL capture:** Forwarded packets are captured on capture ports. The capture option is on only permits ACEs. The capture port can be an intrusion detection system (IDS) monitor port or any Ethernet port. The capture port must be in an output VLAN for Layer 3 switched traffic.
- **VACL redirect:** Matching packets are redirected to specify ports. You can configure up to five redirect ports. Redirect port must be in a VLAN where a VACL is applied.

To configure VACLs, complete the following steps.

Configure VACLs

<table>
<tr><th>Step</th><th>Description</th></tr>
<tr><td>1.</td><td>Define a VLAN access map.
Switch (config)# vlan access-map map name [seq#]</td></tr>
<tr><td>2.</td><td>Configure a match clause.
Switch (config-access-map)# action {drop [log]} | {forward [capture]}
{redirect {{ fastethernet | gigabitehternet | tengigabitethernet}
Slot/port | [port-channel channel id]</td></tr>
<tr><td>3.</td><td>Configure a match clause.
Switch {config-access-map}# action {drop [log]} {[forward [capture]]} {redirect {fastethernet} | gigabitehternet | tengigabitthernet} slotport} | {port-channel channel id}</td></tr>
<tr><td>4.</td><td>Apply a map to VLANs.
Switch {config} # vlan filter map name vlan list list</td></tr>
<tr><td>5.</td><td>Verify the VACL configuration.
Switch# show vlan access-map map name
Switch# show vlan filter [access-map] map name | vlan id]</td></tr>
</table>

Summary

This topic summarizes the key points discussed in this lesson.

Summary

- VLAN hopping can allow Layer 2 unauthorized access to another VLAN.
- VLAN hopping can be mitigate by:
- Properly configuration 802.1Q trunks
- Turning off trunk negotiation
- Access list can be applied to VLANs to limit Layer 2 access.
- VACLs can be configured on Cisco Catalyst switches.

Protecting Against Spoofing Attacks

Overview

Spoofing attacks can occur because several protocols allow a reply from a host even if a request was not received. By spoofing, or pretending to be another machine, the attacker can redirect part or all the traffic coming from, or going to, a predefined target. After the attack, all traffic from the device under attack flows through the computer of the attacker and then to the router, switch, or host.

A spoofing attack can affect hosts, switches, and routers that are connected to your Layer 2 network by sending false information to the devices that are connected to the subnet. Spoofing attacks can also intercept traffic that is intended for other hosts on the subnet. This lesson describes how to mitigate these attacks, and how to configure switches to guard against DHCP, MAC, and Address Resolution Protocol (ARP) threats.

Objectives

Upon completing this lesson, you will be able to protect your network against spoofing attacks. This ability includes being able to meet these objectives:

- Identify DHCP spoofing attacks
- Prevent attacks by using DHCP snooping
- Configure DHCP snooping
- Describe ARP poisoning
- Protect against ARP spoofing attacks with DAI

DHCP Spoofing Attacks

This topic describes DHCP spoofing attacks. Protection against them is part of switch integrated security.

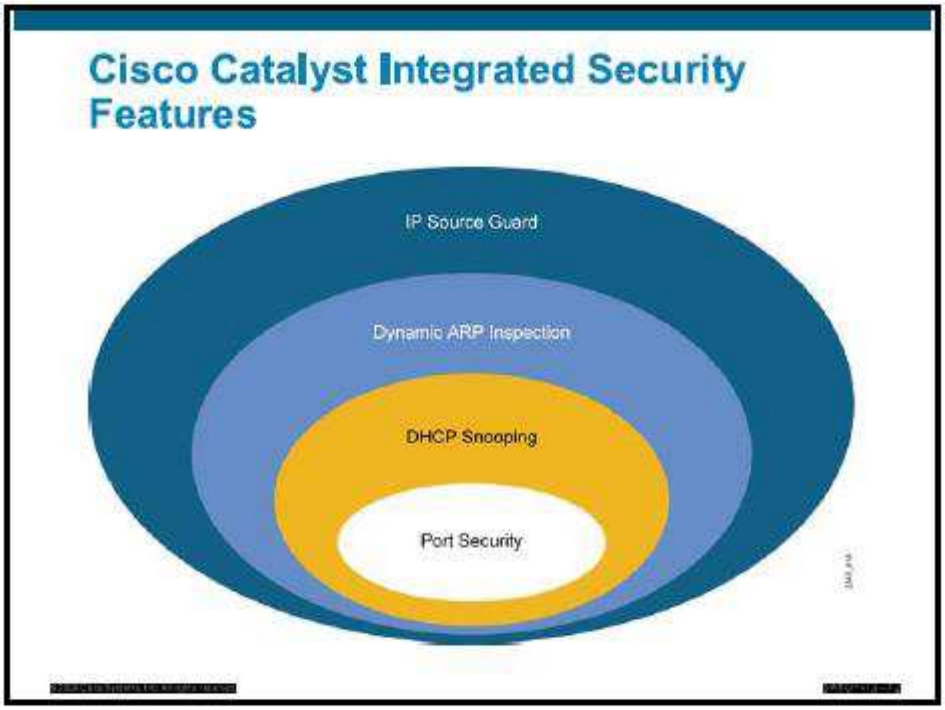

Cisco Catalyst Integrated Security Features

IP Source Guard

Dynamic ARP Inspection

DHCP Snooping

Port Security

Cisco Catalyst integrated security capabilities provide campus security on the Cisco Catalyst switches using integrated tools:

- Port security prevents MAC flooding attacks.
- DHCP snooping prevents client attacks on the DHCP server and switch.
- Dynamic ARP Inspection (DAI) adds security to ARP by using the DHCP snooping table to minimize the impact of ARP poisoning and spoofing attacks.
- IP Source Guard prevents IP spoofing addresses by using the DHCP snooping table.

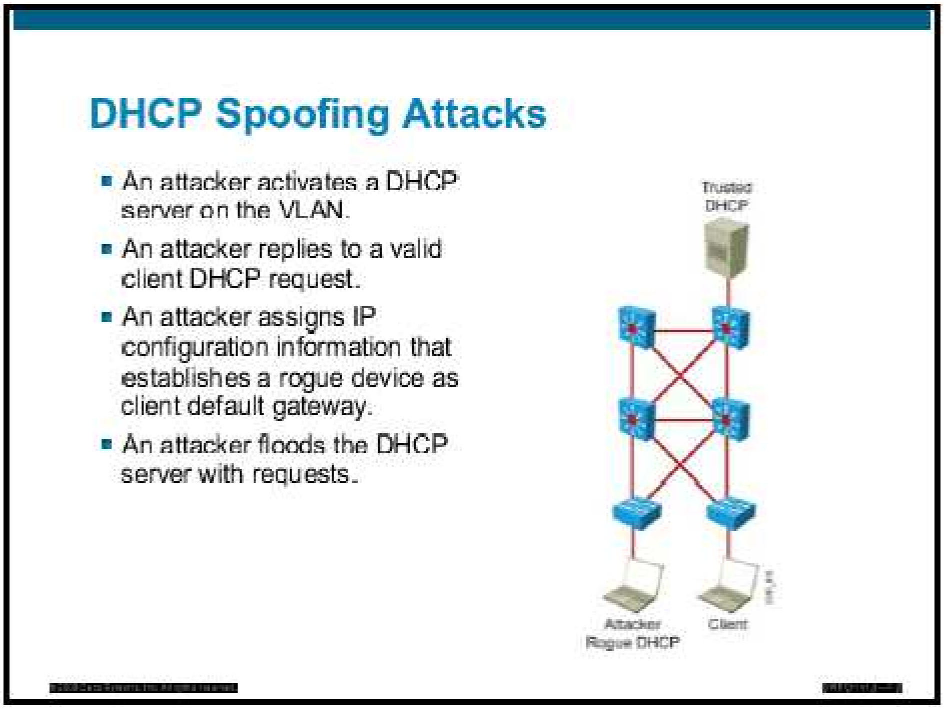

DHCP Spoofing Attacks

- An attacker activates a DHCP server on the VLAN.
- An attacker replies to a valid client DHCP request.
- An attacker assigns IP configuration information that establishes a rogue device as client default gateway.
- An attacker floods the DHCP server with requests.

One of the ways that an attacker can gain access to network traffic is to spoof responses that would be send by a valid DHCP server. The DHCP spoofing device replies to client DHCP requests. The legitimate server may reply also, but if the spoofing device is on the same segment as the client, its reply to the client may arrive first.

The DHCP reply from the intruder offers an IP address and supporting information that designates the intruder as the default gateway or Domain Name System (DNS) server. In the case of a gateway, the clients will then forward packets to the attacking device, which will in turn send them to the desired destination. This is referred to as a man-in-the-middle attack, and it may go entirely undetected as the intruder intercepts the data flow through the network.

The table describes the DHCP spoofing-attack sequence, as shown in the figure.

DHCP Spoofing Attack Sequence

Sequence of Events	Description
1.	Attacker hosts a rogue DHCP server off a switch port.
2.	Client broadcasts a request for DHCP configuration information.
3.	The rogue DHCP server responds before the legitimate DHCP server, assigning attacker-defined IP configuration information.
4.	Host packets are redirected to the attacker’s address as it emulates a default gateway for the erroneous DHCP address that is provided to the client.

DHCP

This subtopic describes DHCP.

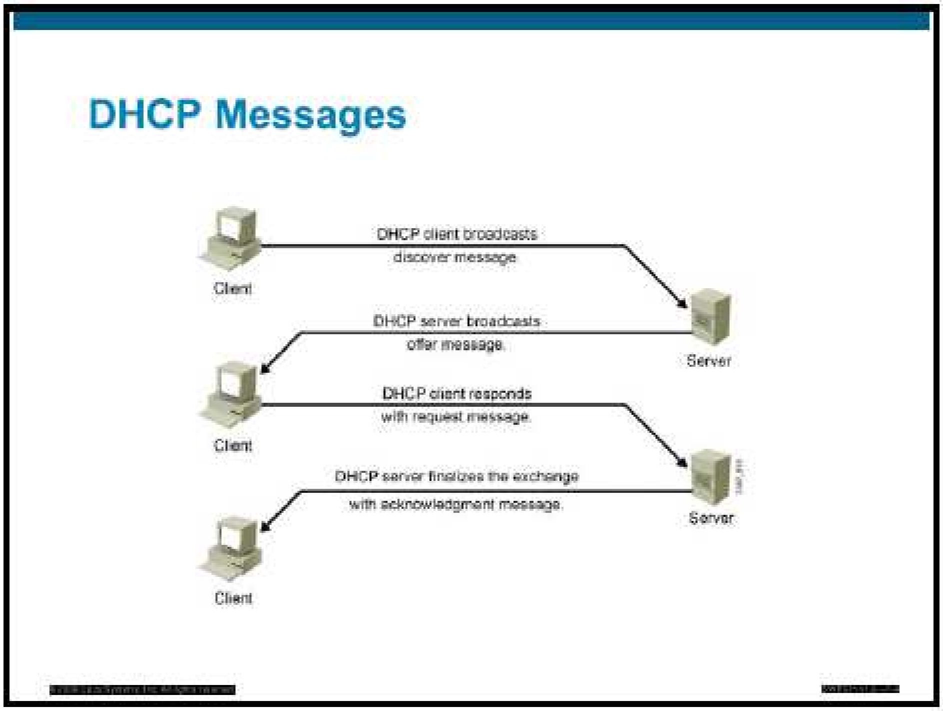

DHCP Messages

Client
DHCP client broadcasts discover message.
Server
DHCP server broadcasts offer message.
Client
DHCP client responds with request message.
Server
DHCP server finalizes the exchange with acknowledgment message.
Client

DHCP uses four messages to provide an IP address to a client:

- DHCP discover broadcast from client
- DHCP offer broadcast to client
- DHCP unicast request from client
- DHCP unicast acknowledgment to client

DHCP Snooping

This topic describes DHCP snooping.

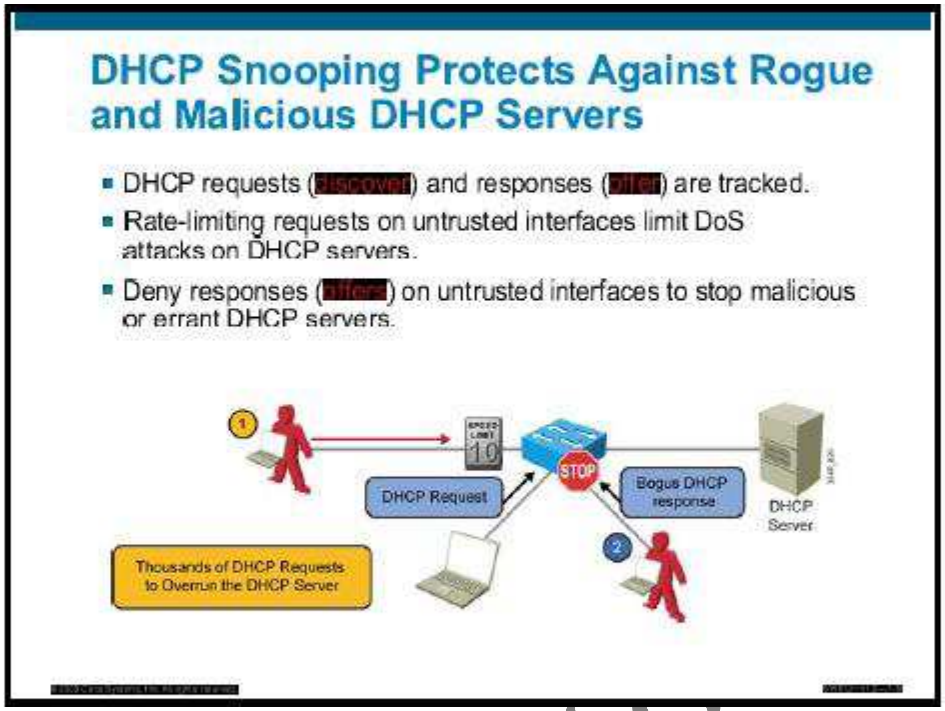

DHCP Snooping Protects Against Rogue and Malicious DHCP Servers

- DHCP requests (discover) and responses (offer) are tracked.
- Rate-limiting requests on untrusted interfaces limit DoS attacks on DHCP servers.
- Deny responses (offers) on untrusted interfaces to stop malicious or errant DHCP servers.

In some cases, an intruder can attach a server to the network and have it assume the role of the DHCP server for that segment. This allows the intruder to give out false DHCP information for the default gateway and domain name servers, which points clients to the hacker. This misdirection allows the hacker to become a man-in-the-middle and to gain access to confidential information, such as username and password pairs, while the end user is unaware of the attack. DHCP snooping can prevent this situation. DHCP snooping is a per-port security mechanism that is used to differentiate an untrusted switch port that is connected to an end user from a trusted switch port that is connected to a DHCP server or to another switch. It can be enabled on a pre-VLAN basis. DHCP snooping allows only authorized DHCP servers to respond to DHCP requests on client ports, thereby mitigating the effect of DHCP denial-of-service (DOS) attacks from an individual client or access port.

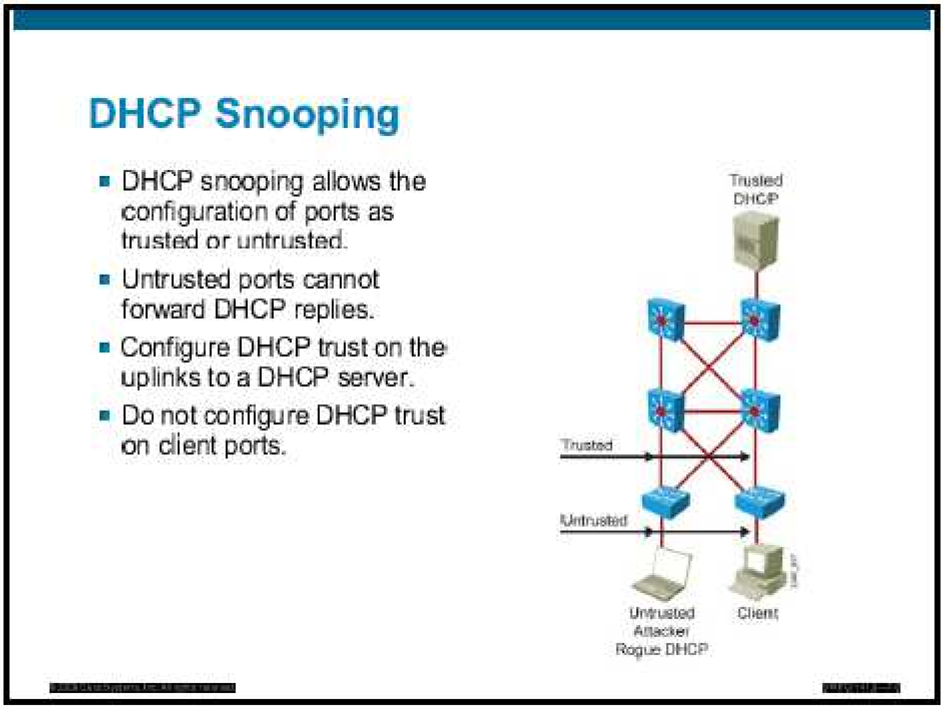

DHCP Snooping

- DHCP snooping allows the configuration of ports as trusted or untrusted.
- Untrusted ports cannot forward DHCP replies.
- Configure DHCP trust on the uplinks to a DHCP server.
- Do not configure DHCP trust on client ports.

DHCP snooping is a Cisco Catalyst feature that determines which switch ports can respond to DHCP requests. Ports are identified as trusted and untrusted. Trusted ports can source all DHCP messages, whereas untrusted ports can source requests only. Trusted ports host a DHCP server or can be an uplink toward the DHCP server. If a rogue device on an untrusted port attempts to send a DHCP response packet into the network, the port is shut down. This feature can be coupled with DHCP option 82, in which switch information, such as the port ID of the DHCP request, can be inserted into DHCP request packet.

Untrusted ports are those that are not explicitly configured as trusted. A DHCP binding table is built for untrusted ports. Each entry contains the client MAC address, IP address, lease time, binding type, VLAN number, and port ID that are recorded as clients make DHCP requests. The table is then used to filter subsequent DHCP traffic. From a DHCP snooping perspective, untrusted access ports should not send any DHCP server responses, such as DHCPOFFER, DHCPACK, or DHCPNAK.

Sequence of Configuration	**Description**
1.	Configure global DHCP snooping.
2.	Configure trusted ports.
3.	Configure Option 82 insertion off (default enabled by Step 2).
4.	Configure rate limiting on untrusted ports.
5.	Configure DHCP snooping for the selected VLANs.

Configure DHCP Snooping

This topic describes DHCP snooping configuration.

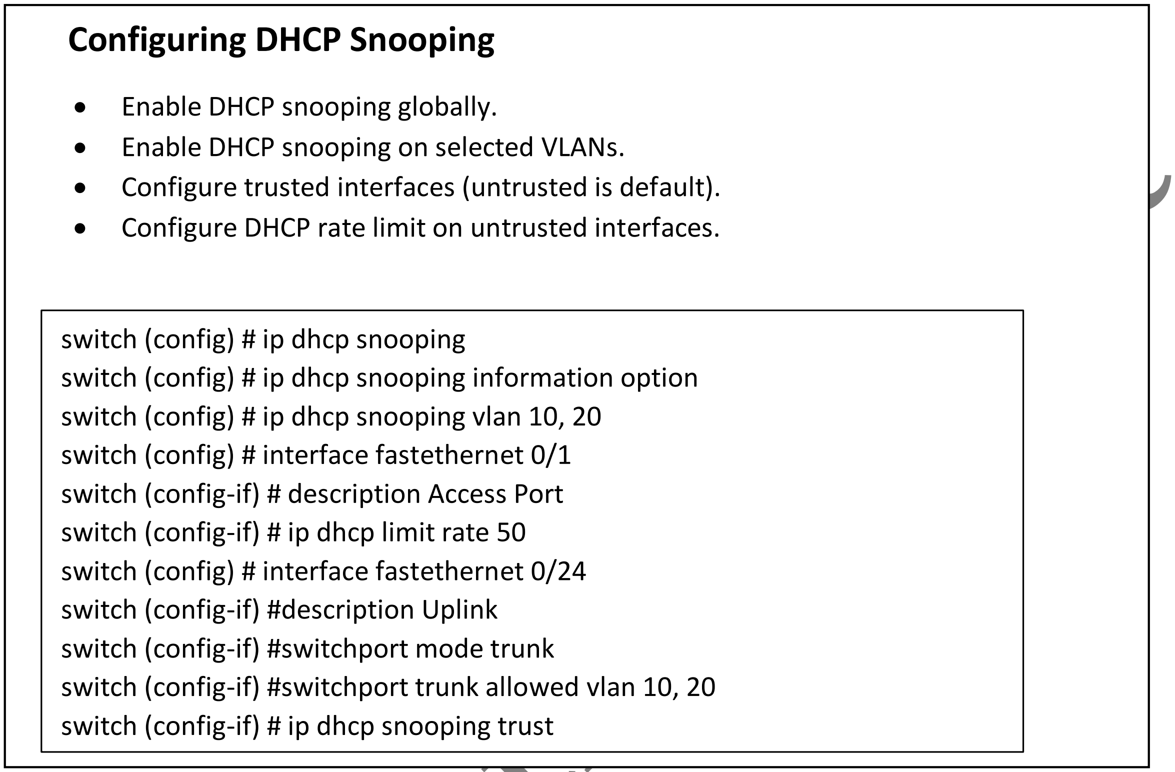

Configuring DHCP Snooping

- Enable DHCP snooping globally.
- Enable DHCP snooping on selected VLANs.
- Configure trusted interfaces (untrusted is default).
- Configure DHCP rate limit on untrusted interfaces.

```
switch (config) # ip dhcp snooping
switch (config) # ip dhcp snooping information option
switch (config) # ip dhcp snooping vlan 10, 20
switch (config) # interface fastethernet 0/1
switch (config-if) # description Access Port
switch (config-if) # ip dhcp limit rate 50
switch (config) # interface fastethernet 0/24
switch (config-if) #description Uplink
switch (config-if) #switchport mode trunk
switch (config-if) #switchport trunk allowed vlan 10, 20
switch (config-if) # ip dhcp snooping trust
```

To enable DHCP snooping, use the commands listed in the table.

Steps for Enabling DHCP Snooping

Step	**comments**
1. Enable DHCP snooping globally. Switch (config) # **ip dhcp snooping**	By default, the feature is not enabled.
2. Enable DHCP option 82 Switch (config-if) # **Ip dhcp snooping information option**	This is optional for the forwarded DHCP request packet to contain information on the switch port where it originated.
3. Configure DHCP server interfaces or uplink ports as trusted. Switch (config-if) # **ip dhcp snooping trust**	At least one trusted port must be configured. Use the no keyword to revert to untrusted. By default, all ports are untrusted.
4. Configure the number of DHCP packets per second (p/s) that are acceptable on the port. Switch (config-if) # **Ip dhcp snooping limit rate** rate	Configure the number of DHCP p/s that an interface can receive. Normally, the rate limit applies to untrusted interfaces. This step is used to prevent DHCP starvation attacks by limiting the rate of the DHCP requests on untrusted ports.
5. Enable DHCP snooping on specific VLAN(s). Switch (config) # **ip dhcp snooping vlan** number [number]	This step is required for identifying VLANs that will be subject to DHCP snooping.
6. Verify the configuration. Switch # **show ip dhcp snooping**	Verify the configuration.

DHCP Snooping Verification

This subtopic describes DHCP snooping verification.

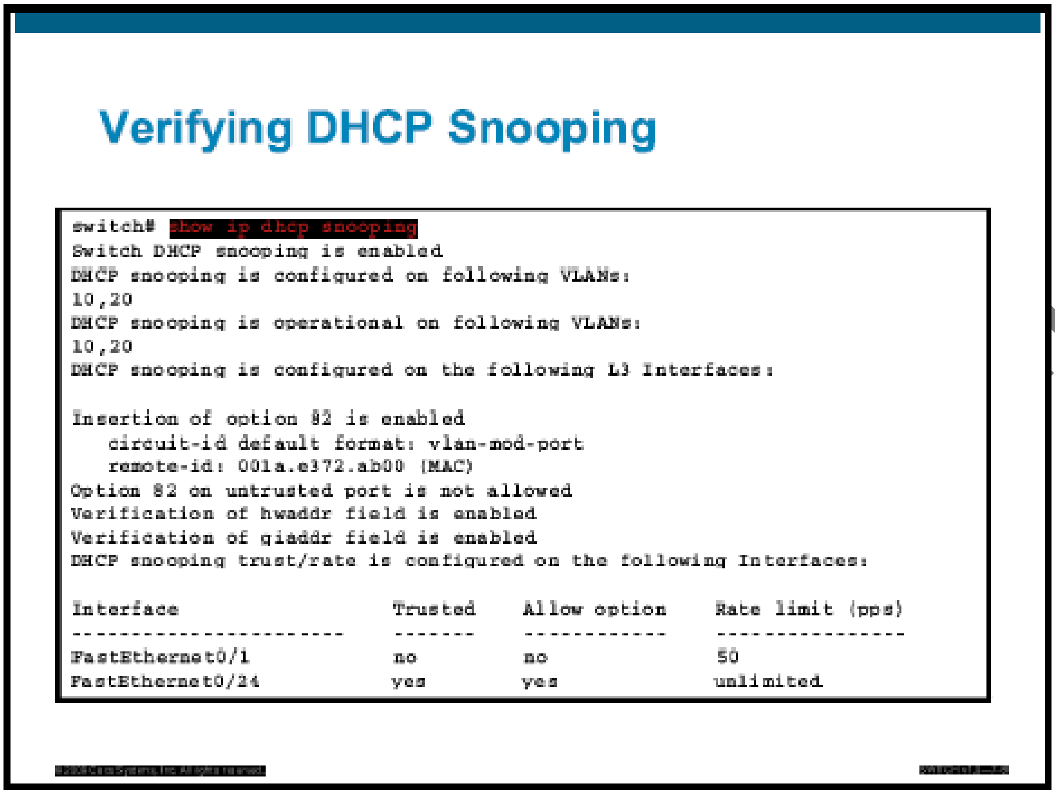

Verifying DHCP Snooping

```
switch# show ip dhcp snooping
Switch DHCP snooping is enabled
DHCP snooping is configured on following VLANs:
10,20
DHCP snooping is operational on following VLANs:
10,20
DHCP snooping is configured on the following L3 Interfaces:

Insertion of option 82 is enabled
   circuit-id default format: vlan-mod-port
   remote-id: 001a.e372.ab00 (MAC)
Option 82 on untrusted port is not allowed
Verification of hwaddr field is enabled
Verification of giaddr field is enabled
DHCP snooping trust/rate is configured on the following Interfaces:

Interface                    Trusted    Allow option    Rate limit (pps)
-----------------------      -------    ------------    ----------------
FastEthernet0/1              no         no              50
FastEthernet0/24             yes        yes             unlimited
```

The **show ip dhcp snooping** family of commands is used to display information about the DHCP snooping configuration.

Only ports that are trusted or that have a rate limit applied will be shown in the output. All other ports are untrusted and are not displayed.

In the output, DHCP snooping is configured for VLANs 10 and 20, and is operational on both of them. Interface FastEthernet0/1 has its rate limited and is not trusted, while interface FastEthernet0/24 does not have any rate limitation and is trusted. All the other ports are untrusted and do not have rate limit. They are not displayed.

ARP Poisoning

This topic describes ARP poisoning.

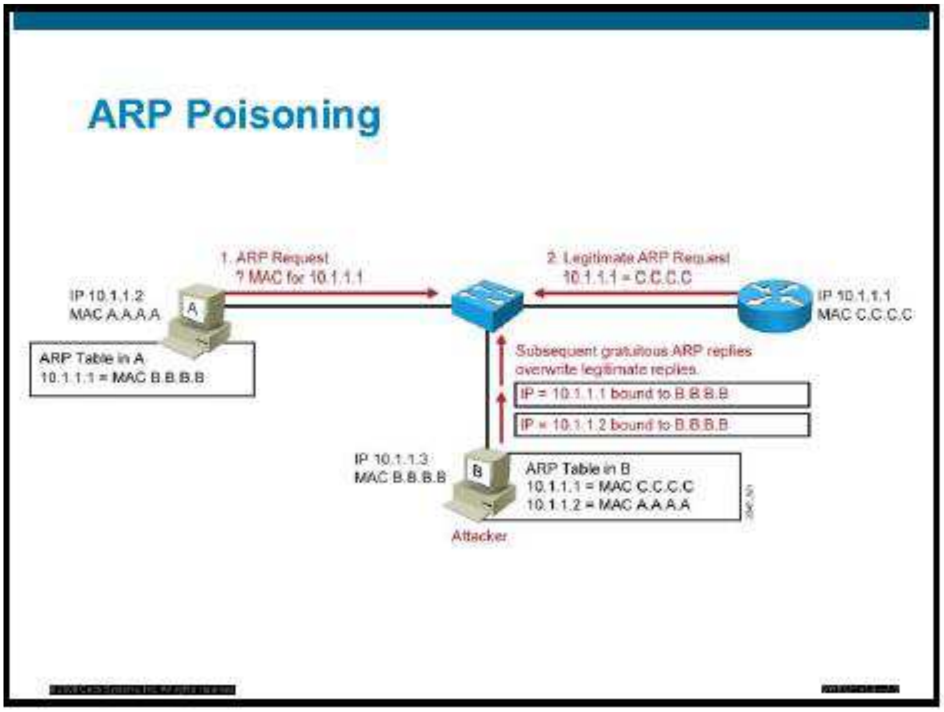

ARP Poisoning

In normal ARP operation, a host sends a broadcast to determine the MAC address of a host with a particular IP address. The device at the IP address replies with its MAC address. The originating host caches the ARP response, using it to populate the destination Layer 2 header of packets that are sent to that IP address.

By spoofing an ARP reply from a legitimate device with a gratuitous ARP, an attacking device appears to be the destination host that is sought by the senders. The ARP reply from the attacker causes the sender to store the MAC address of the attacking system in its ARP cache. All packets that are destined for those IP addresses will be forwarded through the attacker system.

The figure illustrates the sequence of events in an ARP spoofing attack.

An ARP spoofing attack follows the sequence show in the table.

ARP Spoofing Attack

Step or Sequence Number	Description
1.	Host A sends an ARP request for MAC address of C.
2.	Router C replies with its MAC and IP addresses. C also updates its ARP cache.
3.	Host A binds MAC address of C to its IP address in its ARP cache.
4.	Host B (attacker) sends ARP binding MAC address of B to IP address of C.
5.	Host A updates ARP cache with MAC address of B bound to IP address of C.
6.	Host B sends ARP binding MAC address of B to IP address of A.
7.	Router C updates ARP cache with MAC address of B bound to IP address of A.
8.	Packets are diverted through attacker (B).

Dynamic ARP Inspection

This topic describes Dynamic ARP Inspection (DAI).

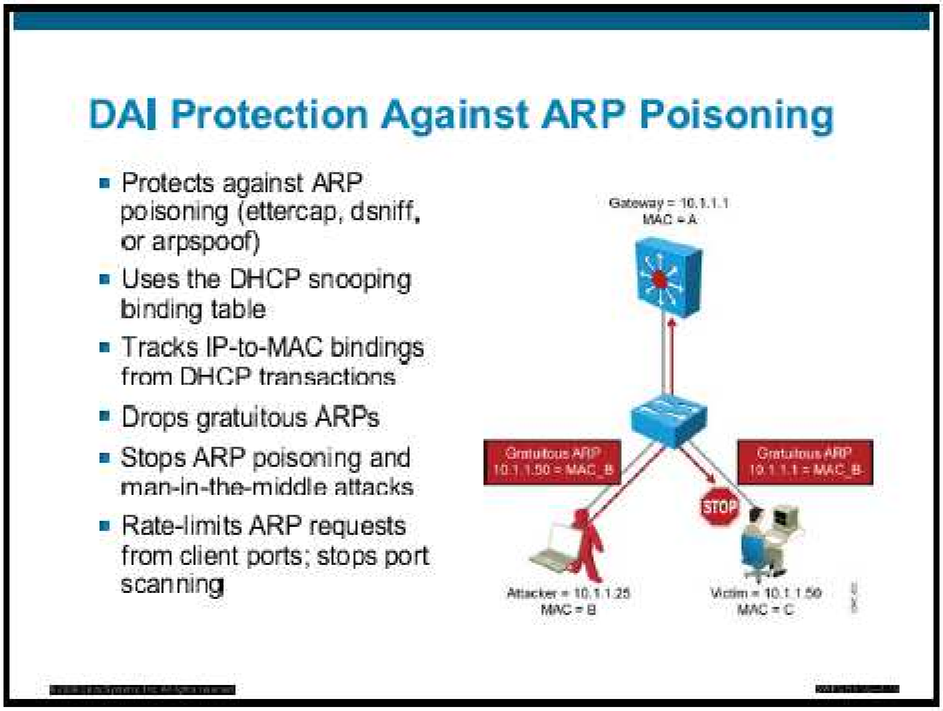

DAI Protection Against ARP Poisoning

- Protects against ARP poisoning (ettercap, dsniff, or arpspoof)
- Uses the DHCP snooping binding table
- Tracks IP-to-MAC bindings from DHCP transactions
- Drops gratuitous ARPs
- Stops ARP poisoning and man-in-the-middle attacks
- Rate-limits ARP requests from client ports; stops port scanning

ARP does not have any authentication. It is quite simple for a malicious user to spoof addresses by using tools such as Ettercap, dsniff, and arpspoof to poison the ARP tables of other hosts on the same VLAN. In a typical attack, a malicious user can send unsolicited ARP replies (gratuitous ARP packets) to other hosts on the subnet with the MAC address of the attacker and the IP address of the default gateway. Frames that are intended for default gateways and are sent from hosts with poisoned ARP tables are sent to the hacker (allowing the packets to be sniffed) or to an unreachable host as a Dos attack. ARP poisoning leads to various man-in-the-middle attacks, posing a security threat in the network.

Dynamic ARP inspection helps prevent the man-in-the-middle attacks by not relaying invalid or gratuitous ARP replies out to other ports in the same VLAN. Dynamic ARP inspection intercepts all ARP requests and all replies on the untrusted ports. Each intercepted packet is verified for valid IP-to-MAC bindings, which are gathered via DHCP snooping. Denied ARP packets are either dropped or logged by the switch for auditing so that ARP poisoning attacks are stopped. Incoming ARP packets on the trusted ports are not inspected. Dynamic ARP inspection can also rate-limit ARP requests from client ports to minimize port-scanning mechanisms.

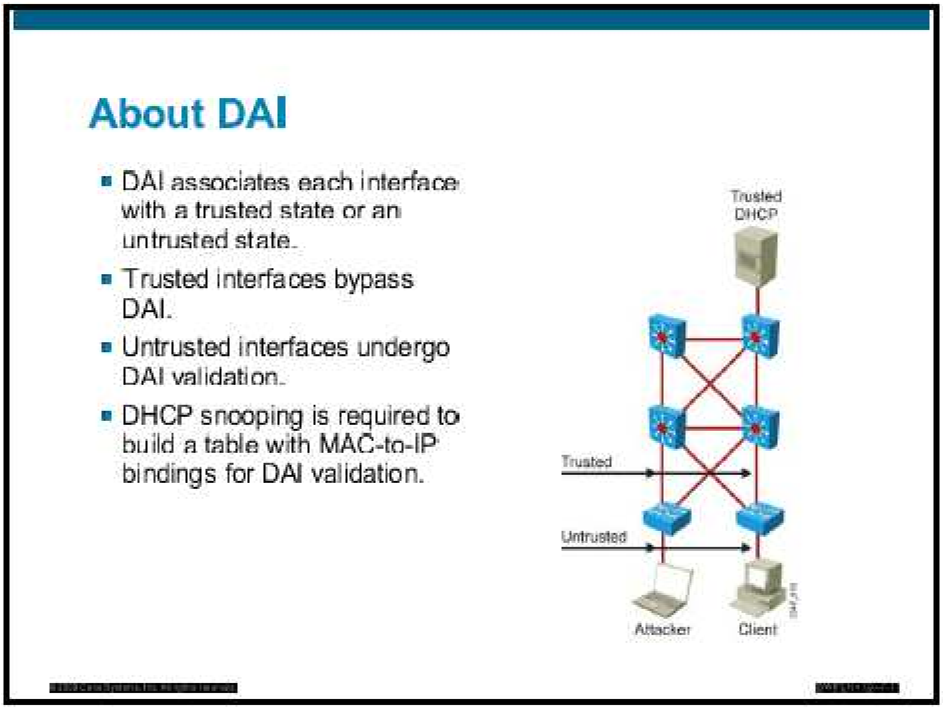

About DAI

- DAI associates each interface with a trusted state or an untrusted state.
- Trusted interfaces bypass DAI.
- Untrusted interfaces undergo DAI validation.
- DHCP snooping is required to build a table with MAC-to-IP bindings for DAI validation.

To prevent ARP spoofing or poisoning, a switch must ensure that only valid ARP requests and responses are relayed. DAI prevents these attacks by intercepting and validating all ARP requests and responses. Each intercepted ARP reply is verified for valid MAC-address-to-IP-address bindings before it is forwarded to a PC to update the ARP cache. ARP replies coming from invalid devices are dropped.

DAI determines the validity of an ARP packet based on a valid MAC-address-to-IP-address bindings database that is built by DHCP snooping. In addition, to handle hosts that use statically configured IP addresses, DAI can validate ARP packets against user-configured ARP access control lists (ACLs).

To ensure that only valid ARP requests and responses are relayed, DAI performs these tasks:

- Forwards ARP packets that are received on a trusted interface without any checks
- Intercepts all ARP packets on untrusted ports
- Verifies that each intercepted packet has a valid IP-to-MAC address binding before forwarding packets that can update the local ARP cache
- Drops, logs, or drops and logs ARP packets with invalid IP-to-MAC address bindings

Configure all access switch ports as untrusted and all switch ports that are connected to other switches as trusted. In this case, all ARP packets that are entering the network would be from an upstream distribution or core switch, bypassing the security check and requiring no further validation.

You can also use DAI to rate-limit the ARP packets and then error-disable the interface if the rate is exceeded.

Dynamic ARP Inspection Configuration

This subtopic describes DAI configuration

Configuring DAI

- Enable DHCP snooping globally.
- Enable DHCP snooping on selected VLANs.
- Enable ARP inspection on selected VLANs.
- Configure trusted interfaces (untrusted is default).

```
Switch (config) # ip dhcp snooping
Switch (config) # ip dhcp snooping vlan 10, 20
Switch (config) # ip arp inspection vlan 10, 20
Switch (config) # interface fastethernet 0/1
Switch (config-if) # ip dhcp limit rate 50
Switch (config) # interface fastethernet 0/24
Switch (config-if) # description uplink
Switch (config-if) # switchport mode trunk
Switch (config-if) # switchport trunk allowed vlan 10, 20
Switch (config-if) # Ip dhcp snooping trust
Switch (config-if) # Ip arp inspection trust
```

The table describes the commands you use to configure DAI.

DAI Commands

Command	Description
Switch (config) # **ip arp inspection vlan** vlan id [, vlan id]	Enables DAI on a VLAN or range of VLANs
Switch (config-if) # **ip arp inspection trust**	Enables DAI on an interface and sets the interface as a trusted interface
Switch (config) # **Ip arp inspection validate** { [src-mac] [dst-mac] [ip] }	Configures DAI to drop ARP packets when the IP addresses are invalid, or when the MAC addresses in the body of the ARP packets do not match the addresses that are specified in the Ethernet header

It is generally advisable to configure all access switch ports as untrusted and to configure all uplink ports that are connected to other switches as trusted.This example of DAI implementation illustrates the configuration that is required on switch 2 with port FastEthernet 3/3 as the uplink port toward the DHCP server. To mitigate the chances of ARP spoofing, these procedures are recommended:

Step 1. Implement protection against DHCP spoofing.
Step 2. Enable DAI

IP Source Guard

This subtopic describes IP Source Guard.

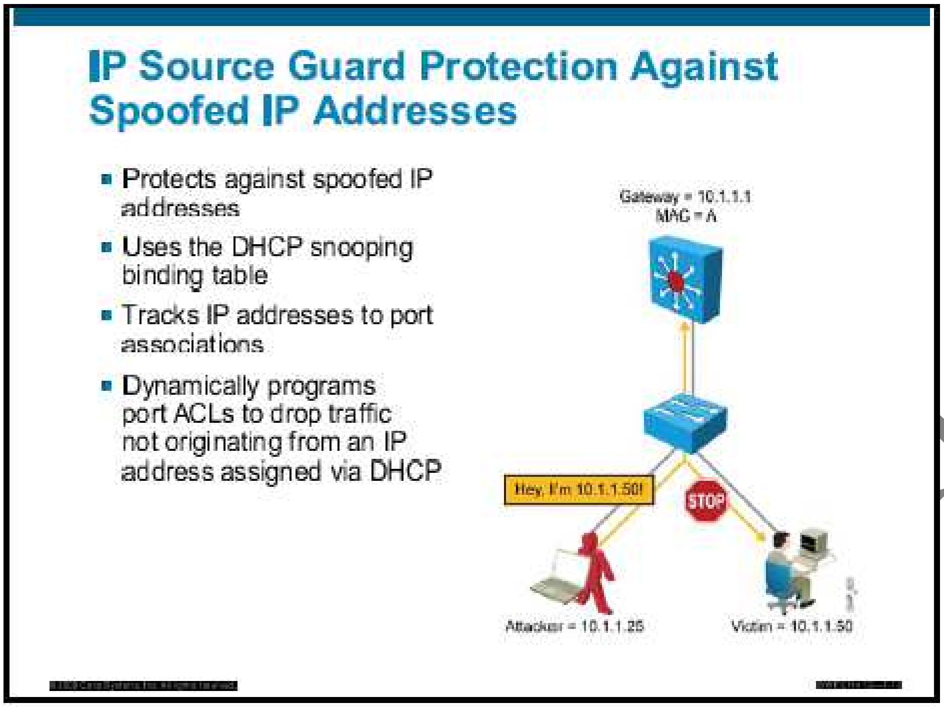

IP Source Guard Protection Against Spoofed IP Addresses

- Protects against spoofed IP addresses
- Uses the DHCP snooping binding table
- Tracks IP addresses to port associations
- Dynamically programs port ACLs to drop traffic not originating from an IP address assigned via DHCP

IP Source Guard prevents a malicious host from attacking the network by hijacking the IP address of its neighbor. IP Source Guard provides per-port IP traffic filtering of the assigned source IP addresses at wire speed. It dynamically maintains per-port VLAN ACLs based on IP-to-MAC-to-switch port bindings. The bindings table is populated either by the DHCP snooping feature or through static configuration of entries. IP Source Guard is typically deployed for untrusted switch ports in the access layer.

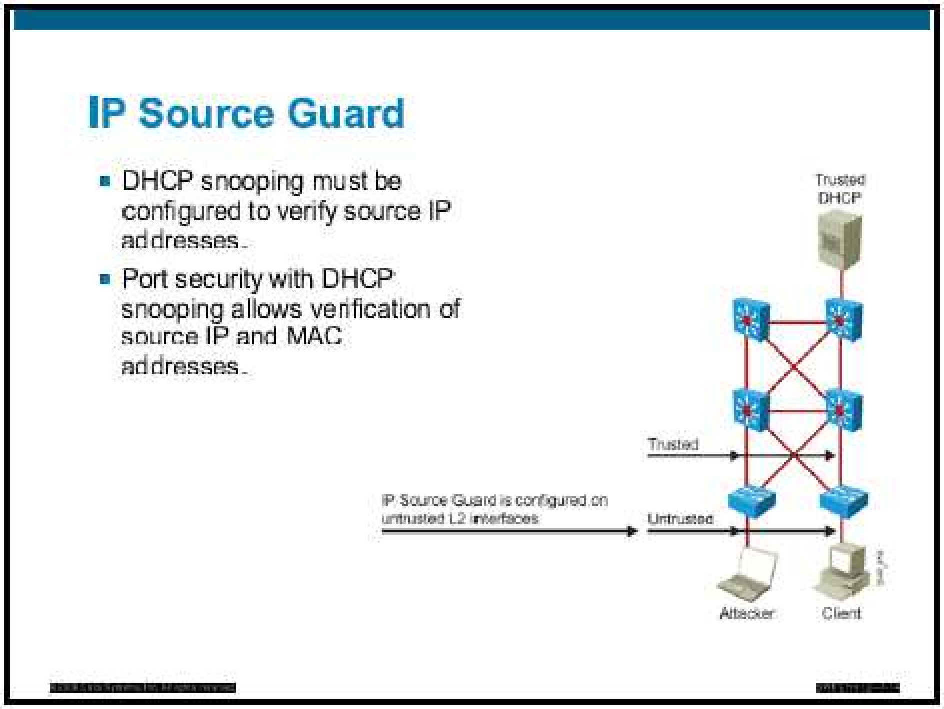

IP Source Guard

- DHCP snooping must be configured to verify source IP addresses.
- Port security with DHCP snooping allows verification of source IP and MAC addresses.

IP Source Guard is similar to DHCP snooping. You can enable IP Source Guard on a DHCP snooping untrusted Layer 2 port to prevent IP address spoofing. To start, all IP traffic on the port is blocked except for DHCP packets that are captured by the DHCP snooping process.

When a client receives a valid IP address from the DHCP server, or when a static IP source binding is configured by the user, a per-port and VLAN access control list (PVACL) is installed on the port.

This process restricts the client IP traffic to those source IP addresses that are configured in the binding: any IP traffic with a source IP address other than that in the IP source binding will be filtered out. This filtering limits the ability of a host to attack the network by claiming the IP address of a neighbor host.

IP Source Guard supports only the Layer 2 port, including both access and trunk. For each untrusted Layer 2 port, there are two levels of IP traffic security filtering, as follows:

- **Source IP address filter:** IP traffic is filtered based on its source IP address. Only IP traffic with a source IP address that matches the IP source binding entry is permitted.

  An IP source address filter is changed when a new IP source entry binding is created or deleted on the port. The port PVACL will be recalculated and reapplied in the hardware to reflect the IP source binding on the port, a default PVACL that denies all IP traffic is installed on the port. Similarly, when the IP filter is disabled, any IP source filter PVACL will be removed from the interface.

- **Source IP and MAC address filter:** IP traffic is filtered based on its source IP address in addition to its MAC address; only IP traffic with source IP and MAC addresses that match the IP source binding entry are permitted.

IP Source Guard Configuration

This subtopic describes IP Source Guard configuration.

Catalyst Integrated Security configuration

```
sw (config) # ip dhcp snooping
sw (config) # ip dhcp snooping vlan 10, 20
sw (config) # ip arp inspection vlan 10, 20
sw (config) # interface fastethernet 0/1
sw (config-if) # description Access Port
sw (config-if) # switchport mode access
sw (config-if) # switchport port-security maximum 2
sw (config-if) # switchport port-security violation restrict
sw (config-if) # switchport port-security
sw (config-if) # ip dhcp limit rate 50
sw (config-if) # ip verify source port-security
sw (config) # interface fastethernet 0/24
sw (config-if) # description uplink
sw (config-if) # switchport mode trunk
sw (config-if) # switchport trunk allowed vlan 10, 20
sw (config-if) # ip dhcp snooping trust
sw (config-if) # ip arp inspection trust
```

The table describes the procedure for enabling IP Source Guard.

DHCP snooping is required for learning valid IP address and MAC address pairs.

IP Source Guard Configuration Commands

	Command	Purpose
Step 1	switch (config) # **ip dhcp snooping**	Enables DHCP snooping, globally. You can use the no keyword to disable DHCP snooping.
Step 2	switch (config) # **ip dhcp snooping vlan** number [number]	Enables DHCP snooping on your VLANs.
Step 3	switch (config-if) # **ip dhcp snooping vlan** number [number]	Configures the interface as trusted or untrusted. You can use the no keyword to configure an interface to receive only messages from within the network.
Step 4	switch (cofig-if) # **ip verify source vlan dhcp-snooping port-security**	Enables IP Source Guard, source IP, and source MAC address filtering on the port.
Step 5	switch (config-if) # **switchport port-security limit rate invalid-source-mac N**	(optional) Sets the rate limit for bad packets. This rate limit also applies to the port where DHCP snooping security mode is enabled as filtering the IP and MAC address.
Step 6	switch (config) # **ip source binding** ip-addr **ip vlan** number **interface** interface	Configures a static IP binding on the port.
Step 7	switch (config) # **end**	Exits configuration mode.

Note: - The static IP source binding can be configured on a switch port only. If you issue the IP **source binging** VLAN interface command on a Layer 3 port, you will receive this error message: Static IP source binding can be configured on the switch port only.

Summary

This topic summarizes the key points that were discussed in this lesson.

Summary

- DHCP spoofing attacks send unauthorized replies to DHCP queries.
- DHCP snooping is used to counter a DHCP spoofing attack.
- DHCP snooping is easily implemented on a Cisco Catalyst switch.
- ARP spoofing can be used to redirect traffic to an unauthorized device on the network.
- DAI in conjunction with DHCP snooping can be used to counter ARP spoofing attacks.
-