CCNA: Switching, Routing, and Wireless Essentials

Switch Boot Sequence

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

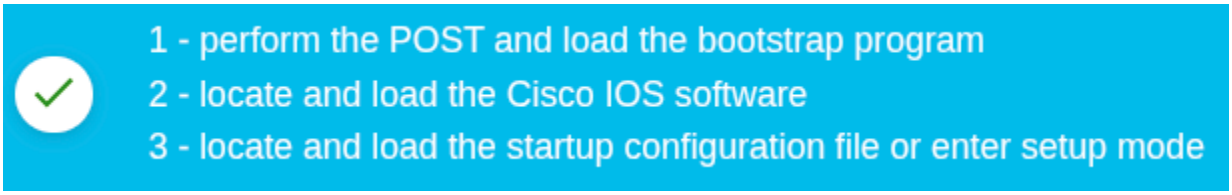

Cisco IOS (Internetwork Operating System) is the proprietary operating system that runs most Cisco routers, switches, and other network devices, acting as the "brain" to control operations, manage data flow, and enable connectivity through a powerful command-line interface (CLI) for configuration and troubleshooting.

The boot system command

- Switch attempts to automatically boot using information in BOOT env variable
    - If not set, switch attempts to load and execute the first executable file it can find

- Catalyst 2960
    - Image file normally contained in a directory with same name as the image file (excluding the .bin file extension)

- IOS operating system then initializes using the Cisco IOS command found in startup-config file
    - Called **config.text**
    - Located in flash

Example:

- BOOT environment variable is set using the **boot system** global configuration mode command

show boot command

- Display what the current IOS boot file is set to

```
S1(config)# boot system
flash:/c2960-lanbasek9-mz.150-2.SE/c2960-lanbasek9-mz.150-2.SE.bin
```

Command	Definition
boot system	The main command
flash:	The storage device
c2960-lanbasek9-mz.150-2.SE/	The path to the file system
c2960-lanbasek9-mz.150-2.SE.bin	The IOS file name

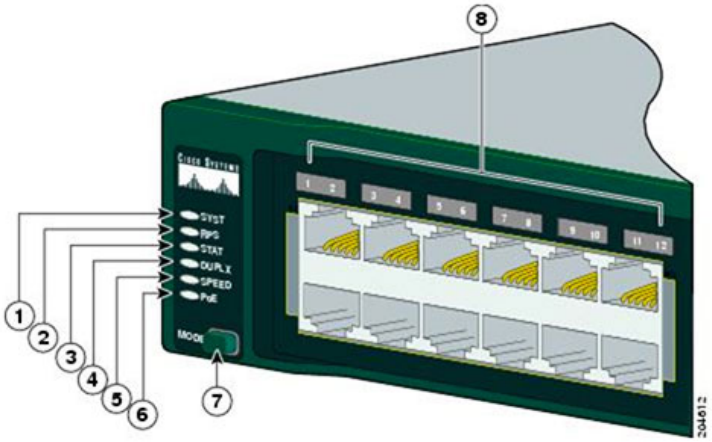

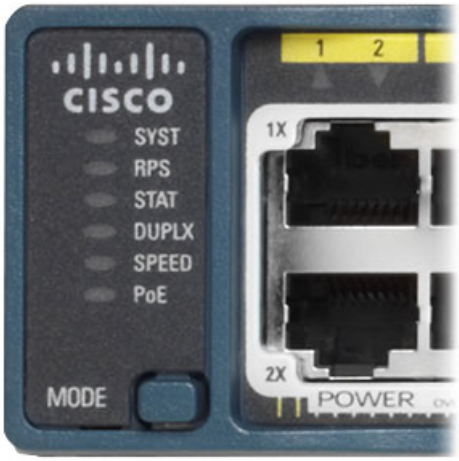

1 - SYST | System LED

- Shows whether the system is receiving power and is functioning properly.
  - LED is off = System is not powered on
  - LED is green = System is operating normally
  - LED is amber = System is receiving power, but is not functioning properly

2 - RPS | Redundant Power System

- Shows the RPS status
  - LED is off = RPS is off
  - LED is green = RPS is connected and ready to provide back up power
  - LED is **blinking** green = RPS is connected, but is unavailable because it is providing power to another device

- LED is amber = RPS is in standby mode
- LED is **blinking** amber = The internal power supply in switch has failed and RPS is providing power

3 - STAT | Port Status LED

- Indicates that the port status mode is selected when LED is green
  - LED is off = no link, or port was administratively shut down
  - LED is green = a link is present
  - LED is **blinking** green = there is activity and port is rx/tx
  - LED is **alternating** green-amber = there is link fault
  - LED is amber = port is blocked to ensure that a loop does not exist in the forwarding domain and is not forwarding data (ports will remain in this state for the first 30 seconds after being activated)
  - LED is **blinking** amber = port is blocked to prevent a possible loop in the forwarding domain

4 - DUPLX | Port Duplex LED

- Indicates that port duplex mode is selected when LED is green
  - Port LEDs that are off are in half-duplex mode
  - LED is green = port is in full-duplex mode

5 - SPEED | Port Speed LED

- LED is off = port is operating at 10mbps
- LED is green = port is operating at 100mbps
- LED is **blinking** green = port is operating at 1000mbps

6 - PoE | Power over Ethernet

- LED is off = indicates the PoE mode is not selected and that none of the ports have been denied power or placed in a fault condition
- LED is **blinking** amber = the PoE mode is not selected, but at least one of the ports has been denied power or has a PoE fault
- LED is green = indicate the PoE mode is selected and the port LEDs will display colors with different meanings
  - Port LED is off = PoE is off
  - Port LED is green = PoE is on
  - Port LED is **alternating** green-amber = PoE is denied b/c providing power to the powered device will exceed the switch power capacity
  - Port LED is **blinking** amber = PoE is off because of a fault, PoE for the port has been disabled

Recovering from a System Crash

- The boot loader has a CLI that provides access to the files stored in flash memory

1. Connect a PC by console cable. Configure terminal emulation software to connect to switch
2. Unplug the switch power cable
3. Reconnect the power cable to the switch and, within 15 seconds, press and hold down the **mode** button while the system LED is still flashing green
4. Continue holding down the **mode** button until the system LED turns briefly amber and then solid green; then release the **mode** button
5. The boot loader **switch**: prompt appears in the terminal emulation software on the PC

Note - Type **help** or **?** at the boot loader prompt to view list of available commands

To view path of switch BOOT env variable

- *set* command
    - Initialize the flash file system using *flash_init* command to view current files in flash

```
switch: set
BOOT=flash:/c2960-lanbasek9-mz.122-55.SE7/c2960-lanbasek9-mz.122-55.SE7.bin
(output omitted)
switch: flash_init
Initializing Flash...
flashfs[0]: 2 files, 1 directories
flashfs[0]: 0 orphaned files, 0 orphaned directories
flashfs[0]: Total bytes: 32514048
flashfs[0]: Bytes used: 11838464
flashfs[0]: Bytes available: 20675584
flashfs[0]: flashfs fsck took 10 seconds.
...done Initializing Flash.
```

    - Enter *dir flash:* command to view the directories and files in flash

```
switch: dir flash:
Directory of flash:/
    2  -rwx  11834846  <date>              c2960-lanbasek9-mz.150-2.SE8.bin
    3  -rwx  2072      <date>              multiple-fs
```

    - Enter *BOOT=flash* command to change the BOOT environment variable path the switch uses to load the new IOS in flash
        - To verify the new BOOT environment variable path, issue *set* command again

- Finally, to load the new IOS, type *boot* command without any arguments

```
switch: BOOT=flash:c2960-lanbasek9-mz.150-2.SE8.bin
switch: set
BOOT=flash:c2960-lanbasek9-mz.150-2.SE8.bin
(output omitted)
switch: boot
```

The boot loader commands support

- Initializing flash
- Formatting flash
- Installing a new IOS
- Changing the BOOT environment variable
- Recovery of lost or forgotten passwords

Switch Management Access

- Switch must have a **switch virtual interface (SVI)** configured with an IPv4 address and subnet mask or an IPv6 address and a prefix length for IPv6

Note - SVI is a virtual interface, not a physical port

- Switch must be configured with a default gateway

Switch Virtual Interface (SVI)

Configuration example

- By default, switch is configured to have its management controlled through VLAN 1
  - All ports are assigned to VLAN 1 by default

Configuration steps

1. Configure the Management Interface
- From VLAN interface configuration mode, an IPv4 address and subnet mask is applied to the management SVI of the switch

Note - SVI for VLAN99 will not appear as up/up until VLAN99 is created and there is a device connected to a switch port associated with VLAN99

Note - The switch may need to be configured for IPv6

example: before configuring IPv6 addressing on Catalyst 2960 IOS v15.0

- Need to enter the global configuration command: *sdm prefer dual-ipv4-and-ipv6 default* then **reload** the switch

Task	IOS Commands
Enter the global configuration mode	S1# configure terminal
Enter interface configuration mode for the SVI	S1(config)# interface vlan 99
Configure the management interface IPv4 address	S1(config-if)# ip address 172.17.99.11 255.255.255.0
Configure the management interface IPv6 address	S1(config-if)# ipv6 address 2001:db8:acad:99::11/64
Enable the management interface	S1(config-if)# no shutdown
Return to the privileged EXEC mode	S1(config-if)# end
Save the running config to the startup config	S1# copy running-config startup-config

2. Configure the Default Gateway
    - Switch should be configured with a default gateway if it will be managed remotely from networks that are not directly connected

Note - It will receive its default gateway information from a router advertisement (RA) message, the switch does not require an IPv6 default gateway

Task	IOS Commands
Enter the global configuration mode	S1# configure terminal
Configure the default gateway for the switch	S1(config)# ip default-gateway 172.17.99.1
Return to the privileged EXEC mode	S1(config)# end
Save the running config to the startup config	S1# copy running-config startup-config

3. Verify Configuration
- *show ip interface brief* and *show ipv6 interface brief* commands are useful for determining the status of both physical and virtual interfaces

Note - An IP address applied to the SVI is only for remote management access to the switch; this does not allow the switch to route layer 3 packets

```
S1# show ip interface brief
Interface     IP-Address     OK? Method   Status    Protocol
Vlan99        172.17.99.11   YES manual   down      down
(output omitted)
S1# show ipv6 interface brief
Vlan99                 [down/down]
    FE80::C27B:BCFF:FEC4:A9C1
    2001:DB8:ACAD:99::11
(output omitted)
```

Duplex Communication

Full-duplex communication

- Increases bandwidth efficiency by allowing both ends of a connection to transmit and receive data simultaneously
- Also known as bidirectional communication
    - Requires microsegmentation

Microsegmented LAN

- Created when a switch port has only one device connected and is operating in full-duplex mode
- There is no collision domain associated with a switch port operating in full-duplex mode

Half-duplex communication

- Unidirectional
- Creates performance issues because data can flow in only one direction at a time– resulting in collisions
- Typically seen in older hardware such as hubs
- Have been replaced by switches that use full-duplex communication by default

Note - Gigabit and 10 Gb NICs require full-duplex to operate

Configure Switch Ports at the Physical Layer

- Manually configured with specific duplex and speed settings

duplex interface configuration mode command

- Manually specify the duplex mode for a switch port

speed interface configuration mode command

- Manually specify the speed

Task	IOS Commands
Enter global configuration mode	S1# configure terminal
Enter interface configuration mode	S1(config)# interface FastEthernet 0/1
Configure the interface duplex	S1(config-if)# duplex full
Configure the interface speed	S1(config-if)# speed 100
Return to the privileged EXEC mode	S1(config-if)# end
Save the running config to the startup config	S1# copy running-config startup-config

The default setting for both duplex and speed for switch ports on Cisco Catalyst 2960 and 3560 switches is **auto**.

10/100/1000 ports operate in either half- or full-duplex mode

- When set to 10 or 100 mbps

10/100/1000 ports operate operate only in full-duplex mode

- When set to 1000 mbps

Autonegotiation

- Useful when the speed and duplex settings of the device connecting to the port are unknown or may change

Note - When connecting to known devices, best practice is to manually set the speed and duplex settings

Note - Mismatched settings for the duplex mode and speed of a switch port can cause connectivity issues. Autonegotiation failure creates mismatched settings

Auto-MDIX

- Automatic medium-dependent interface
- When enabled, the interface automatically detects the required cable connection type (straight-through or crossover) and configures the connection appropriately
- Previously, certain cable types (straight-through or crossover) were required when connecting devices
    - Switch-to-switch or switch-to-router connections required using different Ethernet cables

Without Auto-MDIX

- Straight-through cables must be used to connect to devices such as:
    - **Servers**
    - **Workstations**
    - **Routers**
- Crossover cables must be used to connect to other **switches** or **repeaters**

With Auto-MDIX

- Either type of cable can be used to connect to other devices
- The interface automatically adjusts to communicate successfully

mdix auto interface configuration mode command enables this feature

S1(config-if)# mdix auto

Note - Only enabled by default on Catalyst 2960 and Catalyst 3560 and newer
    - Not available on older Catalyst 2950 and Catalyst 3550

show controllers ethernet-controller command with *phy* keyword

- To examine the auto-MDIX setting for specific interface
- To limit output to lines referencing auto-MDIX
    - *include MDIX* filter

```
S1# show  controllers  ethernet-controller  fa0/1  phy  |  include MDIX
Auto-MDIX                :  On    [AdminState=1    Flags=0x00052248]
```

Switch Verification Commands

Task	IOS Commands
Display interface status and configuration	S1# show interfaces [interface-id]
Display current startup configuration	S1# show startup-config
Display current running configuration	S1# show running-config
Display information about flash file system	S1# show flash
Display system hardware and software status	S1# show version
Display history of command entered	S1# show history
Display IP information about an interface	S1# show IP interface [interface-id] OR S1# show ipv6 interface [interface-id]
Display the MAC address table	S1# show mac-address-table OR S1# show mac address-table

Verify Switch Port Configuration

show running-config command

- Can be used to verify that the switch has been correctly configured

Example:

- Fast Ethernet 0/18 interface is configured with the management VLAN 99
- VLAN 99 is configured with an IPv4 address of 172.17.99.11 255.255.255.0
- The default gateway is set to 172.17.99.1

```
S1# show running-config
Building configuration...
Current configuration : 1466 bytes
!
interface FastEthernet0/18
switchport access vlan 99
switchport mode access
!
(output omitted)
!
interface Vlan99
ip address 172.17.99.11 255.255.255.0
ipv6 address 2001:DB8:ACAD:99::1/64
!
ip default-gateway 172.17.99.1
```

show interfaces command

- Displays status and statistics information on the network interfaces of the switch
- Frequently used when configuring and monitoring network devices

```
S1# show interfaces fastEthernet 0/18
FastEthernet0/18 is up, line protocol is up (connected)
  Hardware is Fast Ethernet, address is 0025.83e6.9092 (bia
0025.83e6.9092)
  MTU 1500 bytes, BW 100000 Kbit/sec, DLY 100 usec,
     reliability 255/255, txload 1/255, rxload 1/255
  Encapsulation ARPA, loopback not set
  Keepalive set (10 sec)
  Full-duplex, 100Mb/s, media type is 10/100BaseTX
```

First line indicates that FastEthernet 0/19 interface is up/up

- Operational

Lower line displays duplex is full and speed is 100mbps

Network Access Layer Issues

show interfaces command

- Useful for detecting common media issues

```
S1# show interfaces fastEthernet 0/18
FastEthernet0/18 is up, line protocol is up (connected)
Hardware is Fast Ethernet, address is 0025.83e6.9092 (bia
0025.83e6.9092)MTU 1500 bytes, BW 100000 Kbit/sec, DLY 100 usec,
```

First parameter (FastEthernet0/18 is up):

- Refers to the hardware layer and indicates whether the interface is receiving a carrier detect signal

Second parameter (line protocol is up):

- Refers to the data link layer and indicates whether the data link layer protocol keepalives are being received

Possible resolutions to output

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

```
S1# show interfaces fastEthernet 0/18
FastEthernet0/18 is up, line protocol is up (connected)
  Hardware is Fast Ethernet, address is 0025.83e6.9092 (bia
0025.83e6.9092)
  MTU 1500 bytes, BW 100000 Kbit/sec, DLY 100 usec,
     reliability 255/255, txload 1/255, rxload 1/255
  Encapsulation ARPA, loopback not set
  Keepalive set (10 sec)
  Full-duplex, 100Mb/s, media type is 10/100BaseTX
  input flow-control is off, output flow-control is unsupported
  ARP type: ARPA, ARP Timeout 04:00:00
  Last input never, output 00:00:01, output hang never
  Last clearing of "show interface" counters never
  Input queue: 0/75/0/0 (size/max/drops/flushes); Total output
drops: 0
  Queueing strategy: fifo
  Output queue: 0/40 (size/max)
  5 minute input rate 0 bits/sec, 0 packets/sec
  5 minute output rate 0 bits/sec, 0 packets/sec
     2295197 packets input, 305539992 bytes, 0 no buffer
     Received 1925500 broadcasts (74 multicasts)
     0 runts, 0 giants, 0 throttles
     3 input errors, 3 CRC, 0 frame, 0 overrun, 0 ignored
     0 watchdog, 74 multicast, 0 pause input
     0 input packets with dribble condition detected
     3594664 packets output, 436549843 bytes, 0 underruns
     8 output errors, 0 collisions, 10 interface resets
     0 unknown protocol drops
     0 babbles, 235 late collision, 0 deferred
```

* Some media errors are not severe enough to cause the circuit to fail but do cause network performance issues

Error Type	Description
Input Errors	Total number of errors. It includes runts, giants, no buffer, CRC, frame, overrun, and ignored counts
Runts	Frames that are discarded because they are smaller than the minimum frame size for the medium. For instance, any Ethernet frame that is l**ess than 64 bytes** is considered a runt.
Giants	Frames that are discarded because they exceed the maximum frame size for the medium. For example, any Ethernet frame that is **greater than 1518 bytes** is considered a giant
CRC	CRC errors are generated when the calculated checksum is not the same as the checksum received
Output Errors	Sum of all errors that prevented the final transmission of datagrams out of the interface that is being examined
Collisions	Number of message retransmitted because of an Ethernet collision
Late Collisions	A collision that occurs **after 512 bits** of the frame have been transmitted

Interface Input and Output Errors

Input Errors

- The sum of all errors in datagrams that were received on the interface being examined
    - Includes runs, giants, CRC, no buffer, frame, overrun, and ignored counts

Runt Frames

- Ethernet frames that are **shorter than** the **64-byte** minimum allowed length are called runts.
- Malfunctioning NICs are the usual cause of excessive runt frames, but can also be caused by collisions

Giants

- Ethernet frames that are **larger than the maximum** allowed size are called giants

CRC Errors

- On the Ethernet and serial interfaces, CRC errors usually indicate a media or cable error.
- Common causes include
    - Electrical interference, loose or damaged connections, or incorrect cabling
- If you see many CRC errors, there is too much noise on the link and you should inspect the cable. You should also search for and eliminate noise sources

Output Errors

- The sum of all errors that prevented the final transmission of datagrams out the interface that is being examined.

Collisions

- Collisions in half-duplex operations are normal. However, you should never see collisions on an interface configured for full-duplex communication

Late Collisions

- A late collision refers to a collision that occurs **after 512 bits** of the frame have been transmitted.
- Causes:
    - Excessive cable lengths are the most common cause
    - Duplex misconfiguration
        - Example: You could have one end of a connection configured for full-duplex and the other for half-duplex
        - Late collision would be seen on interface configured for half-duplex
    - Must configure the same duplex settings on both ends. Properly designed and configured network should never have late collisions

Troubleshooting Network Access Layer Issues

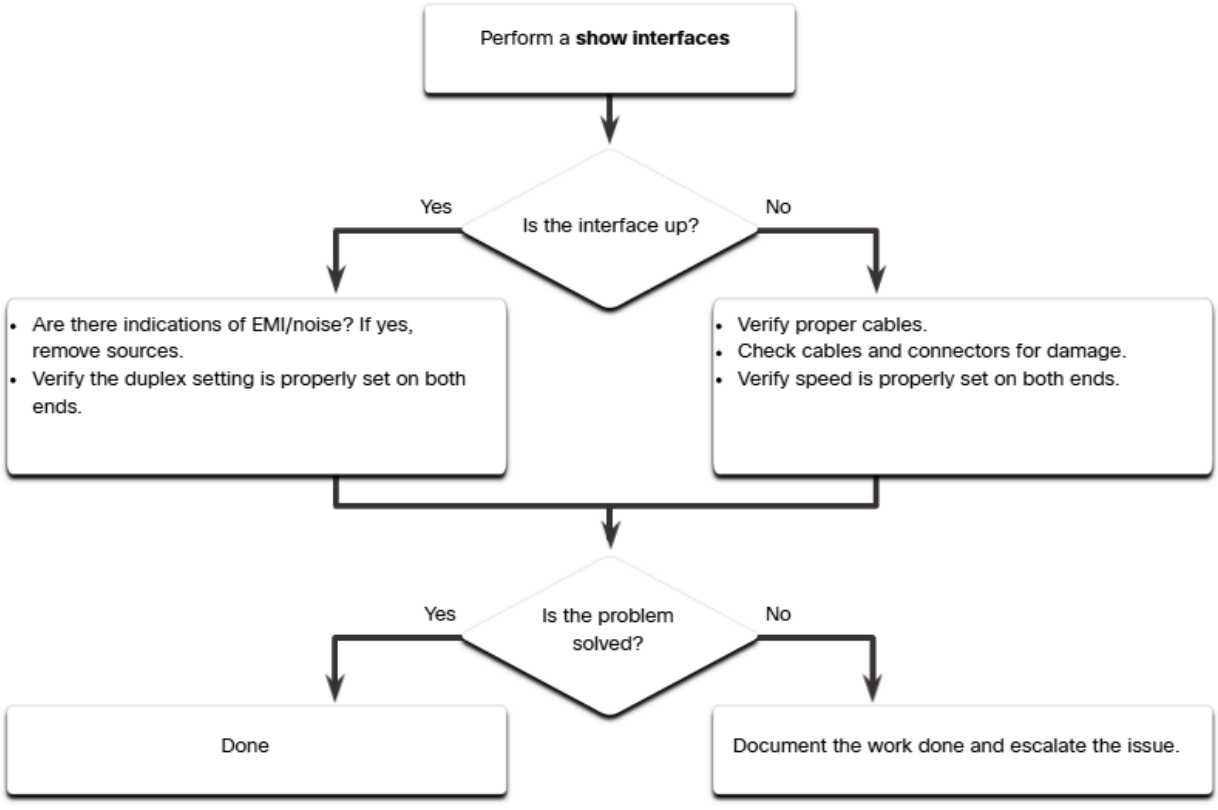

If interface is down

- Check physical cables
- Possible mismatch in speed setting
    - Typically autonegotiated

If interface is up, but issues with connectivity

- *show interfaces* command
    - Check for excessive noise
    - Indications may include an increase in counters for:
        - Runs
        - Giants
        - CRC Errors
    - If there is excessive noise, find and remove source of noise
    - Verify cables do not exceed maximum cable length and check type of cable used
    - Check for collisions
    - If there are collisions or late collisions, verify duplex settings on both ends of connection
        - Usually autonegotiated

Enter configuration mode and set FastEthernet0/1 duplex, speed, and MDIX to auto and save the configuration to NVRAM.

```
S1# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
S1(config)# interface FastEthernet0/1
S1(config-if)# duplex auto
S1(config-if)# speed auto
S1(config-if)# mdix auto
```

End out of interface configuration mode and save the configuration to NVRAM.

```
S1(config-if)# end
%SYS-5-CONFIG_I: Configured from console by console
S1# copy running-config startup-config
```

You have successfully configured duplex, speed, and Auto-MDIX settings on a switch interface and saved the configuration to NVRAM.

Secure Remote Access

Telnet Operation

- Uses TCP port 23
- Older protocol that uses unsecure plaintext transmission of both the login authentication and the data transmitted between the communicating devices.
- Threat actor can monitor packets using Wireshark

SSH Operation

- Secure protocol that uses TCP port 22
- Provides a secure (encrypted) management connection to a remote device.
- Provides security for remote connections by providing strong encryption when a device is authenticated and also for transmitted data between the communicating devices

Verifying the Switch Supports SSH

- To enable SSH on Catalyst 2960
    - Switch must be using a version of IOS software including cryographic (encrypted) features and capabilities
    - Use *show version* command
        - Will display an IOS filename that includes the combination “k9” supports cryptographic (encrypted) features and capabilities

```
S1# show version
Cisco IOS Software, C2960 Software (C2960-LANBASEK9-M), Version
15.0(2)SE7, RELEASE SOFTWARE (fc1)
```

Configure SSH

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

Step 1

Verify SSH support

- Use *show ip ssh* command to verify that the switch supports SH
    - If switch IOS doesn’t support cryptographic features, command is unrecognized

Step 2

Configure the IP domain

- use *ip domain-name domain-name* global configuration mode command

```
S1(config)# ip domain-name cisco.com
```

Step 3

Generate RSA key pairs

Note - not all IOS versions default to SSH version 2, SSH version 1 has known security flaws

- To configure SSH version 2
    - Issue the *ip ssh version 2* global configuration mode command
    - * Generating an RSA key pair automatically enables SSH
    - Use *crypto key generate rsa* global configuration mode command to enable the SSH server on switch and generate an RSA key pair
        - When generating RSA keys, the admin is prompted to enter a modulus length

Note - A longer modulus length is more secure, but takes longer to generate and to use

Note - To delete RSA key pair

- Use *crypto key zeroize rsa* global configuration mode command
- SSH server is automatically disabled

```
S1(config)# crypto key generate rsa
How many bits in the modulus [512]: 1024
```

Step 4

Configure user authentication

- The SSH server can authenticate users locally or using an authentication server
- To use local authentication method, create a username and password pair
    - *username* username *secret* password

```
S1(config)# username admin secret ccna
```

Step 5

Configure the vty lines

- Enable SSH protocol on vty lines using
- *transport input ssh*

Catalyst 2960 has vty lines randing from 0 to 15

- The configuration prevents non-SSH (telnet) connections and limits switch to accept only SSH connections

Use *line vty* then *login local* to require local authentication for SSH connections from local username database

```
S1(config)# line vty 0 15
S1(config-line)# transport input ssh
S1(config-line)# login local
S1(config-line)# exit
```

Step 6

Enable SSH version 2

- By default, SSH supports both versions 1 and 2
- When supporting both versions
    - *show ip ssh* output as supporting version 2

ip ssh version 2

Verify SSH is Operational

Use SSH client such as PUTTY or SSH CLI

```
Login as: admin
Using keyboard-interactive
Authentication.
Password:
S1> enable
Password:
S1#
```

To display the version and configuration data for SSH

- Use *show ip ssh*

```
S1# show ip ssh
SSH Enabled - version 2.0
Authentication timeout: 120 secs; Authentication retries: 3
To check the SSH connections to the device, use the show ssh
command as shown.
S1# show ssh
%No SSHv1 server connections running.
Connection Version Mode Encryption  Hmac                State
Username
0          2.0      IN   aes256-cbc  hmac-sha1     Session
started         admin
0          2.0      OUT  aes256-cbc  hmac-sha1     Session
started         admin
S1#
```

Configure Basic Router Settings

Cisco routers and Cisco switches have many similarities

- Similar model operating system
- Similar command structures
- Many of same commands
- Both have similar initial configuration steps

```
Router# configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.
Router(config)# hostname R1
R1(config)# enable secret class
R1(config)# line console 0
R1(config-line)# password cisco
R1(config-line)# login
R1(config-line)# exit
R1(config)# line vty 0 4
R1(config-line)# password cisco
R1(config-line)# login
R1(config-line)# exit
R1(config)# service password-encryption
R1(config)#
```

Configure a banner (motd)

```
R1(config)# banner motd #Authorized Access Only!#
R1(config)#
```

Save the changes on a router

```
R1# copy running-config startup-config
Destination filename [startup-config]?
Building configuration...
[OK]
```

Enter global configuration mode and name the router R2.

```
Router# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Router(config)# hostname R2
```

Configure class as the secret password.

```
R1(config)# enable secret class
```

Configure cisco as the console line password and require users to login. Then exit line configuration mode.

```
R1(config)# line console 0
R1(config-line)# password cisco
R1(config-line)# login
R1(config-line)# exit
```

Configure cisco as the vty password for lines 0 through 4 and require users to login.

```
R1(config)# line vty 0 4
R1(config-line)# password cisco
R1(config-line)# login
```

Exit line configuration mode and encrypt all plaintext passwords.

```
R1(config-line)# exit
R1(config)# service password-encryption
```

Enter the banner Authorized Access Only! and use # as the delimiting character.

```
R1(config)# banner motd #Authorized Access Only!#
```

Exit global configuration mode and save the configuration.

```
R1(config)# exit
R1# copy running-config startup-config
Destination filename [startup-config]?
Building configuration...

[OK]
```

You successfully configured R2 with initial settings.

Dual Stack Topology

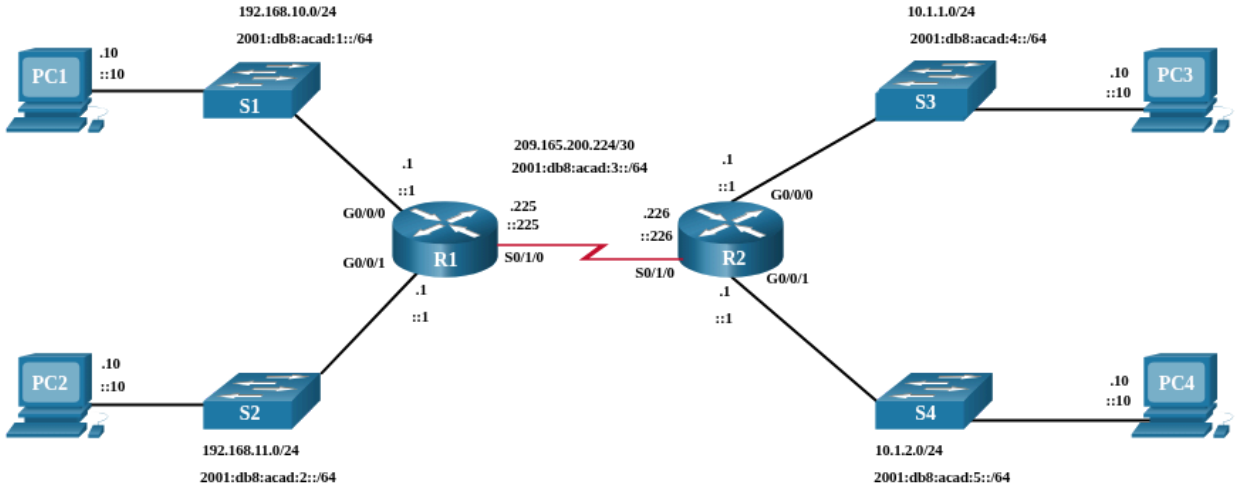

Configure Router Interfaces

Example:

- G2 ISRs have one or two integrated gigabit ethernet interfaces
    - High-Speed WAN interface card (HWIC) slots to accommodate other types of network interfaces
        - Serial, DSL, cable interfaces

To be available, an interface must be:

- Configured with at least one IP address - Use the *ip address* ip-address subnet-mask
    - *ipv6 address* ipv6-address/prefix
- Activated - By default, LAN and WAN interfaces are not activated (shutdown).
    - To enable an interface, it must be activated using the *no shutdown* command
        - Interface must also be connected to another device (hub, switch, router) for physical layer to be active
- Description - Optionally, the interface could also be configured with a short description up to 240 characters

```
R1(config)# interface gigabitethernet 0/0/0
R1(config-if)# ip address 192.168.10.1 255.255.255.0
R1(config-if)# ipv6 address 2001:db8:acad:1::1/64
R1(config-if)# description Link to LAN 1
R1(config-if)# no shutdown
R1(config-if)# exit
R1(config)# interface gigabitethernet 0/0/1
R1(config-if)# ip address 192.168.11.1 255.255.255.0
R1(config-if)# ipv6 address 2001:db8:acad:2::1/64
```

```
R1(config-if)# description Link to LAN 2
R1(config-if)# no shutdown
R1(config-if)# exit
R1(config)# interface serial 0/0/0
R1(config-if)# ip address 209.165.200.225 255.255.255.252
R1(config-if)# ipv6 address 2001:db8:acad:3::225/64
R1(config-if)# description Link to R2
R1(config-if)# no shutdown
R1(config-if)# exit
R1(config)#
```

Configure GigabitEthernet 0/0/0.

- Use g0/0/0 to enter interface configuration mode.
- Configure the IPv4 address 10.1.1.1 and subnet mask 255.255.255.0.
- Configure the IPv6 address 2001:db8:acad:4::1/64.
- Describe the link as Link to LAN 3.
- Activate the interface.

```
Router(config)#interface g0/0/0
Router(config-if)#ip address 10.1.1.1 255.255.255.0
Router(config-if)#ipv6 address 2001:db8:acad:4::1/64
Router(config-if)#description Link to LAN 3
Router(config-if)#no shutdown
%LINK-3-UPDOWN: Interface GigabitEthernet0/0/0, changed state to up
%LINEPROTO-5-UPDOWN: Line protocol on Interface GigabitEthernet0/0/0,
changed state to up
```

Configure GigabitEthernet 0/0/1.

- Use g0/0/1 to enter interface configuration mode.
- Configure the IPv4 address 10.1.2.1 and subnet mask 255.255.255.0.
- Configure the IPv6 address 2001:db8:acad:5::1/64.
- Describe the link as Link to LAN 4.
- Activate the interface.

```
Router(config-if)#interface g0/0/1
Router(config-if)#ip address 10.1.2.1 255.255.255.0
Router(config-if)#ipv6 address 2001:db8:acad:5::1/64
```

```
Router(config-if)#description Link to LAN 4
Router(config-if)#no shutdown
%LINK-3-UPDOWN: Interface GigabitEthernet0/0/1, changed state to up
%LINEPROTO-5-UPDOWN: Line protocol on Interface GigabitEthernet0/0/1,
changed state to up
```

Configure Serial 0/0/0.

- Use s0/0/0 to enter interface configuration mode.
- Configure the IPv4 address 209.165.200.226 and subnet mask 255.255.255.252.
- Configure the IPv6 address 2001:db8:acad:3::226/64.
- Describe the link as Link to R1.
- Activate the interface.

```
Router(config-if)#interface s0/0/0
Router(config-if)#ip address 209.165.200.226 255.255.255.252
Router(config-if)#ipv6 address 2001:db8:acad:3::226/64
Router(config-if)#description Link to R1
Router(config-if)#no shutdown
%LINK-3-UPDOWN: Interface Serial0/0/0, changed state to up
%LINEPROTO-5-UPDOWN: Line protocol on Interface Serial0/0/0, changed
state to up
```

You successfully configured the R2 router interfaces.

IPv4 Loopback Interfaces

- A logical interface that is internal to the router
- Not assigned to a physical port
- Can never be connected to any other device
- Is considered a software interface that is automatically placed in an up state, as long as router is functioning

Useful in testing and managing a Cisco IOS device because it ensures that at least one interface will always be available

- Commonly used in lab environments to create additional interfaces
    - Can create multiple loopback interfaces on a router to simulate more networks for configuration practice and testing purposes

Enable and assign a loopback address:

```
Router(config)# interface loopback number
```

```
Router(config-if)# ip address ip-address subnet-mask
```

Multiple loopback interfaces can be enabled on a router.

- IPv4 address for each loopback interface must be unique and unused

```
R1(config)# interface loopback 0
R1(config-if)# ip address 10.0.0.1 255.255.255.0
R1(config-if)# exit
R1(config)#
%LINEPROTO-5-UPDOWN: Line protocol on Interface Loopback0,
changed state to up
```

Interface Verification Commands

Quickly identify the status of an interface

- *show ip interface brief* and *show ipv6 interface brief* - These display a summary for all interfaces including the IPv4 or IPv6 address of the interface and current operational status
- *show running-config interface* interface-id - This displays the commands applied to the specified interface
- *show ip route* and *show ipv6 route* - These display the contents of the IPv4 or IPv6 routing table stored in RAM.
    - In Cisco IOS 15, active interfaces should appear in the routing table with two related entries
        - code ‘C’ (Connected)
        - code ‘L’ (Local)
    - Previous IOS versions, only a single entry with code ‘C’ will appear

Verify Interface Status

Quickly reveal the status of all interfaces on the router

- *show ip interface brief*
- *show ipv6 interface brief*

Verify that the interfaces are active and operational as indicated by the status “up” and protocol “up”

```
R1# show ip interface brief
Interface              IP-Address      OK? Method Status                Protocol
GigabitEthernet0/0/0   192.168.10.1    YES manual up                    up
GigabitEthernet0/0/1   192.168.11.1    YES manual up                    up
Serial0/1/0            209.165.200.225 YES manual up                    up
Serial0/1/1            unassigned      YES unset  administratively down down
R1# show ipv6 interface brief
GigabitEthernet0/0/0   [up/up]
    FE80::7279:B3FF:FE92:3130
    2001:DB8:ACAD:1::1
GigabitEthernet0/0/1   [up/up]
    FE80::7279:B3FF:FE92:3131
    2001:DB8:ACAD:2::1
Serial0/1/0            [up/up]
    FE80::7279:B3FF:FE92:3130
    2001:DB8:ACAD:3::1
Serial0/1/1            [down/down]     Unassigned
```

Verify IPv6 Link Local and Multicase Addresses

show ipv6 interface brief

- Displays two configured IPv6 addresses per interface
    - One address is the IPv6 global unicast address that was manually entered
    - Other address which begins with FE80, is link-local unicast address for the interface
- A link-local address is automatically added to an interface whenever a global unicast address is assigned
    - An IPv6 network interface is required to have a link-local address, but not necessarily a global unicast address

show ipv6 interface gigabitethernet 0/0/0 command displays the interface status and all of the IPv6 addresses belonging to the interface

```
R1# show ipv6 interface gigabitethernet 0/0/0
GigabitEthernet0/0/0 is up, line protocol is up
  IPv6 is enabled, link-local address is FE80::7279:B3FF:FE92:3130
  No Virtual link-local address(es):
  Global unicast address(es):
    2001:DB8:ACAD:1::1, subnet is 2001:DB8:ACAD:1::/64
  Joined group address(es):
    FF02::1
    FF02::1:FF00:1
    FF02::1:FF92:3130
  MTU is 1500 bytes
  ICMP error messages limited to one every 100 milliseconds
  ICMP redirects are enabled
  ICMP unreachables are sent
  ND DAD is enabled, number of DAD attempts: 1
  ND reachable time is 30000 milliseconds (using 30000)
  ND advertised reachable time is 0 (unspecified)
  ND advertised retransmit interval is 0 (unspecified)
  ND router advertisements are sent every 200 seconds
  ND router advertisements live for 1800 seconds
  ND advertised default router preference is Medium
```

Verify Interface Configuration

```
R1 show running-config interface gigabitethernet 0/0/0
Building configuration...
Current configuration : 158 bytes
!
interface GigabitEthernet0/0/0
    description Link to LAN 1
    ip address 192.168.10.1 255.255.255.0
    negotiation auto
    ipv6 address 2001:DB8:ACAD:1::1/64
end
R1#
```

Gather more detailed interface information

- *show interfaces* - Displays interface information and packet flow count for all interfaces on the device
- *show ip interface* and *show ipv6 interface* - Displays the IPv4 and IPv6 related information for all interfaces on a router

Verify Routes

show ip route and *show ipv6 route* commands

- Reveals the three directly connected network entries and the three local host route interface entries

Local host route

- Has an administrative distance of 0
- Has a /32 mask for IPv4
- Has a /128 mask for IPv6
- For routes on the router that owns the IP address
- Used to allow the router to process packets destined to that IP

```
R1# show ip route
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP

Gateway of last resort is not set
      192.168.10.0/24 is variably subnetted, 2 subnets, 2 masks
C        192.168.10.0/24 is directly connected, GigabitEthernet0/0/0
L        192.168.10.1/32 is directly connected, GigabitEthernet0/0/0
      192.168.11.0/24 is variably subnetted, 2 subnets, 2 masks
C        192.168.11.0/24 is directly connected, GigabitEthernet0/0/1
L        192.168.11.1/32 is directly connected, GigabitEthernet0/0/1
      209.165.200.0/24 is variably subnetted, 2 subnets, 2 masks
C        209.165.200.224/30 is directly connected, Serial0/1/0
L        209.165.200.225/32 is directly connected, Serial0/1/0
R1# show ipv6 route
IPv6 Routing Table - default - 7 entries
Codes: C - Connected, L - Local, S - Static, U - Per-user Static route

C   2001:DB8:ACAD:1::/64 [0/0]
     via GigabitEthernet0/0/0, directly connected
L   2001:DB8:ACAD:1::1/128 [0/0]
     via GigabitEthernet0/0/0, receive
C   2001:DB8:ACAD:2::/64 [0/0]
     via GigabitEthernet0/0/1, directly connected
L   2001:DB8:ACAD:2::1/128 [0/0]
     via GigabitEthernet0/0/1, receive
C   2001:DB8:ACAD:3::/64 [0/0]
     via Serial0/1/0, directly connected
L   2001:DB8:ACAD:3::1/128 [0/0]
     via Serial0/1/0, receive
L   FF00::/8 [0/0]
     via Null0, receive
R1#
```

A ‘C’ next to a route within the routing table indicates that this is a directly connected network

- When the router interface is configured with a global unicast address and is up/up, the IPv6 prefix and prefix length are added to the IPv6 routing table as a connected route
- IPv6 global unicast address applied is also installed in routing table as local route
    - Has /128 prefix
    - Local routes are used by routing table to efficiently process packets with interface address of router as the destination

Filter Show Command Output

- Commands that generate multiple screens of output are, by default, paused after 24 lines.
- At the end of the paused output
    - The --More-- text displays
- Pressing **Enter** displays next line
- Pressing **spacebar** displays the next **set** of lines

terminal length command

- Specify the number of lines to be displayed
    - 0 (zero) prevents the router from pausing between screens of output

Filtering the *show* output

- Can be used to display specific sections of output
- Enable the filtering command:
    - enter a pipe (|) after *show* command then enter a filtering parameter and a filtering expression

Filtering Commands

Section	Shows the entire section that starts with the filtering expression
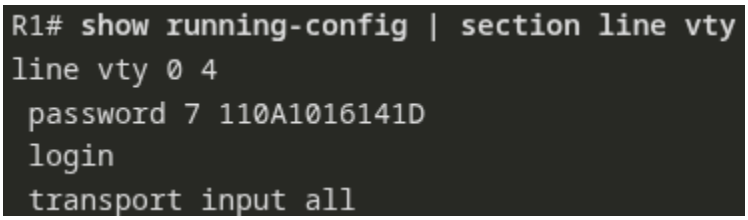	
include	Includes all output lines that match the filtering expression

```
R1# show ip interface brief
Interface              IP-Address      OK? Method Status                Protocol
GigabitEthernet0/0/0   192.168.10.1    YES manual up                    up
GigabitEthernet0/0/1   192.168.11.1    YES manual up                    up
Serial0/1/0            209.165.200.225 YES manual up                    up
Serial0/1/1            unassigned      NO  unset  down                  down
R1#
R1# show ip interface brief | include up
GigabitEthernet0/0/0   192.168.10.1    YES manual up                    up
GigabitEthernet0/0/1   192.168.11.1    YES manual up                    up
Serial0/1/0            209.165.200.225 YES manual up                    up
```

exclude	Excludes all output lines that match the filtering expresion

```
R1# show ip interface brief
Interface              IP-Address      OK? Method Status                Protocol
GigabitEthernet0/0/0   192.168.10.1    YES manual up                    up
GigabitEthernet0/0/1   192.168.11.1    YES manual up                    up
Serial0/1/0            209.165.200.225 YES manual up                    up
Serial0/1/1            unassigned      NO  unset  down                  down
R1#
R1# show ip interface brief | exclude unassigned
Interface              IP-Address      OK? Method Status                Protocol
GigabitEthernet0/0/0   192.168.10.1    YES manual up                    up
GigabitEthernet0/0/1   192.168.11.1    YES manual up                    up
Serial0/1/0            209.165.200.225 YES manual up                    up
```

begin	Shows all of the output lines from a certain point, starting with the line that matches the filtering expression

```
R1# show ip route | begin Gateway
Gateway of last resort is not set
      192.168.10.0/24 is variably subnetted, 2 subnets, 2 masks
C        192.168.10.0/24 is directly connected, GigabitEthernet0/0/0
L        192.168.10.1/32 is directly connected, GigabitEthernet0/0/0
      192.168.11.0/24 is variably subnetted, 2 subnets, 2 masks
C        192.168.11.0/24 is directly connected, GigabitEthernet0/0/1
L        192.168.11.1/32 is directly connected, GigabitEthernet0/0/1
      209.165.200.0/24 is variably subnetted, 2 subnets, 2 masks
C        209.165.200.224/30 is directly connected, Serial0/1/0
L        209.165.200.225/32 is directly connected, Serial0/1/0
```

Enter the command to filter the show running-config output for the 'line con' section.

```
R1# show running-config | section line con

line con 0
 password 7 05080F1C2243
 transport input none
```

Enter the command to filter for 'down' interfaces in the brief listing.

```
R1# show ip interface brief | include down

Serial0/1/1   unassigned   NO  unset  down  down
```

Enter the command to exclude 'up' interfaces in the brief listing.

```
R1# show ip interface brief | exclude up

Interface              IP-Address      OK? Method Status                Protocol
Serial0/1/1              unassigned      NO  unset  down                  down
```

Enter the command to filter the show running-config output to begin at the word 'line'.

```
R1# show running-config | begin line

line con 0
 password 7 05080F1C2243
 transport input none
 stopbits 1
line vty 0 4
 password 7 110A1016141D
 login
 transport input all
```

Command History Feature

Recall commands in history buffer

- CTRL+P
- Up Arrow key

Begins with most recent command

Return to more recent commands in history buffer

- CTRL+N
- Down Arrow key

By default, command history is enabled and system captures last 10 command lines

show history privileged EXE command

- To display the contents of the buffer

terminal history size user EXEC command

- To increase or decrease the size of the buffer

```
R1# terminal history size 200
R1# show history
  show ip int brief
  show interface g0/0/0
  show ip route
  show running-config
  show history
  terminal history size 200
```

Learning Brief

Configure a Switch with Initial Settings

After a Cisco switch is powered on, it goes through a five-step boot sequence. The BOOT environment variable is set using the **boot system** global configuration mode command. The IOS is located in a distinct folder and the folder path is specified. Use the switch LEDs to monitor switch activity and performance: SYST, RPS, STAT, DUPLX, SPEED, and PoE. The boot loader provides access into the switch if the operating system cannot be used because of missing or damaged system files. The boot loader has a command line that provides access to the files stored in flash memory. To prepare a switch for remote management access, the switch must be configured with an IP address and a subnet mask. To manage the switch from a remote network, the switch must be configured with a default gateway. To configure the switch SVI, you must first configure the management interface, then configure the default gateway, and finally, verify your configuration.

Configure Switch Ports

Full-duplex communication increases effective bandwidth by allowing both ends of a connection to transmit and receive data simultaneously. Half-duplex communication is unidirectional. Switch ports can be manually configured with specific duplex and speed settings. Use autonegotiation when the speed and duplex settings of the device connecting to the port are unknown or may change. When auto-MDIX is enabled, the interface automatically detects the required cable connection type (straight-through or crossover) and configures the connection appropriately. There are several **show** commands to use when verifying switch configurations. Use the **show running-config** command and the **show interfaces** command to verify a switch port configuration. The output from the **show interfaces** command is also useful for detecting common network access layer issues because it displays the line and data link protocol status. The reported input errors from the **show interfaces** command include: runt frames, giants, CRC errors, along with collisions and late collisions. Use **show interfaces** to determine if your network has no connection or a bad connection between a switch and another device.

Secure Remote Access

Telnet (using TCP port 23) is an older protocol that uses unsecure plaintext transmission of both the login authentication (username and password) and the data transmitted between the communicating devices. SSH (using TCP port 22) is a secure protocol that provides an encrypted management connection to a remote device. SSH provides security for remote connections by providing strong encryption when a device is authenticated (username and password) and also for the transmitted data between the communicating devices. Use the **show version** command on the switch to see which IOS the switch is currently running. An IOS filename that includes the combination “k9” supports cryptographic features and capabilities. To configure SSH you must verify that the switch supports it, configure the IP domain, generate RSA key pairs, configure use authentication, configure the VTY lines, and enable SSH version 2. To verify that SSH is operational, use the **show ip ssh** command to display the version and configuration data for SSH on the device.

Basic Router Configuration

The following initial configuration tasks should always be performed: name the device to distinguish it from other routers and configure passwords, configure a banner to provide legal notification of unauthorized access, and save the changes on a router. One distinguishing feature between switches and routers is the type of interfaces supported by each. For example, Layer 2 switches support LANs and, therefore, have multiple FastEthernet or Gigabit Ethernet ports. The dual stack topology is used to demonstrate the configuration of router IPv4 and IPv6 interfaces. Routers support LANs and WANs and can interconnect different types of networks; therefore, they support many types of interfaces. For example, G2 ISRs have one or two integrated Gigabit Ethernet interfaces and High-Speed WAN Interface Card (HWIC) slots to accommodate other types of network interfaces, including serial, DSL, and cable interfaces. The IPv4 loopback interface is a logical interface that is internal to the router. It is not assigned to a physical port and can never be connected to any other device.

Verify Directly Connected Networks

Use the following commands to quickly identify the status of an interface: **show ip interface brief** and **show ipv6 interface brief** to see summary all interfaces (IPv4 and IPv6 addresses and operational status), **show running-config interface interface-id** to see the commands applied to a specified interface, and **show ip route** and **show ipv6 route** to see the contents of the IPv4 or IPv6 routing table stored in RAM. The output of the **show ip interface brief** and **show ipv6 interface brief** commands can be used to quickly reveal the status of all interfaces on the router. The **show ipv6 interface gigabitethernet 0/0/0** command displays the interface status and all of the IPv6 addresses belonging to the interface. Along with the link local address and global unicast address, the output includes the multicast addresses assigned to the interface. The output of the **show running-config interface** command displays the current commands applied to a specified interface. The **show interfaces** command displays interface information and packet flow count for all interfaces on the device. Verify interface configuration using the **show ip interface** and **show ipv6 interface** commands, which display the IPv4 and IPv6 related information for all interfaces on a router. Verify routes using the **show ip route** and **show ipv6 route** commands. Filter show command output using the pipe (|) character. Use filter expressions: section, include, exclude, and begin. By default, command history is enabled, and the system captures the last 10 command lines in its history buffer. Use the **show history** privileged EXEC command to display the contents of the buffer.

LAB 1.1.7 - Basic Switch Configuration | **Appendix A: Initialize and Reload a Switch**

a. Console into the switch and enter privileged EXEC mode.

```
Switch> enable
Switch#
```

b. Use the **show flash** command to determine if any VLANs have been created on the switch.

```
Switch# show flash:
Directory of flash:/

    2  -rwx        1919   Mar 1 1993 00:06:33 +00:00  private-config.text
    3  -rwx        1632   Mar 1 1993 00:06:33 +00:00  config.text
    4  -rwx       13336   Mar 1 1993 00:06:33 +00:00  multiple-fs
    5  -rwx    11607161   Mar 1 1993 02:37:06 +00:00
c2960-lanbasek9-mz.150-2.SE.bin
    6  -rwx         616   Mar 1 1993 00:07:13 +00:00  vlan.dat

32514048 bytes total (20886528 bytes free)
```

c. If the **vlan.dat** file was found in flash, then delete this file.

```
Switch# delete vlan.dat
Delete filename [vlan.dat]?
```

d. You are prompted to verify the filename. If you have entered the name correctly, press Enter; otherwise, you can change the filename.

You are prompted to confirm deletion of this file. Press Enter to confirm.

```
Delete flash:/vlan.dat? [confirm]
Switch#
```

e. Use the **erase startup-config** command to erase the startup configuration file from NVRAM. You are prompted to remove the configuration file. Press Enter to confirm.

```
Switch# write erase
Erasing the nvram filesystem will remove all configuration files! Continue?
[confirm]
[OK]
Erase of nvram: complete
Switch#
```

f. Reload the switch to remove any old configuration information from memory. You will then receive a prompt to confirm reloading of the switch. Press Enter to proceed.

```
Switch# reload
Proceed with reload? [confirm]
```

Note: You may receive a prompt to save the running configuration prior to reloading the switch. Respond by typing **no** and press Enter.

```
System configuration has been modified. Save? [yes/no]: no
```

g. After the switch reloads, you should see a prompt to enter the initial configuration dialog. Respond by entering **no** at the prompt and press Enter.

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

Switching Concepts

Switching in Network

Two terms associated with **frames** entering and leaving interface

- Ingress
    - This is used to describe the port where a frame enters the device
- Egress
    - This is used to describe the port that frames will use when leaving the device

LAN Switch

- Maintains a table that is referenced when forwarding traffic through switch
    - Only intelligence is ability to use its table to forward traffic
- Forwards traffic based on **ingress port** and **destination MAC address** of an ethernet frame
- There is only one master switching table that describes a strict association between MAC address and ports
    - Ethernet frame with given destination address always exists same egress port

Note - An ethernet frame will never be forwarded out the same port upon received

Switch MAC Address Table

- Switch uses destination MAC addresses to direct network communications

To know which port to use to transmit frames, it must first learn which devices exist on each port

- Builds table called **MAC address table**
    - Stored in **content addressable memory (CAM)**
        - Special type of memory used in high-speed searching applications
    - Sometimes called the CAM table

The Switch Learn and Forward Method

Learn - Examining the Source MAC Address

- If the **source MAC address** does not exist in MAC address table, the **MAC address** and **incoming port number** are added to the table
- If the **source MAC address** does not exist, the switch **updates the refresh timer** for that entry
    - By default, most Ethernet switches keep an entry in table for **five minutes**
- If the source MAC address does not exist in the table but on a different port, the switch treats this as a new entry
    - The entry is replaced using the MAC address, but with the more current port number

Forward - Examining the Destination MAC Address

If the **destination MAC** address is a **unicast address**, the switch will look for a match between the destination MAC address of the frame and an entry in its MAC address table

- If the destination MAC address is in the table, it will forward the frame out of the specified port
- If the **destination MAC** address is **not** in the table, the switch will forward the frame out **all ports except the incoming port**
    - This is called an **unknown unicast**
    - If the destination MAC address is a **broadcast** or a **multicast**, the frame is also flooded out all ports except the incoming port

Switching Forwarding Methods

- Makes Layer 2 forwarding decisions very quickly

ASICs

- Application-specific-integrated-circuits
- Reduces the frame-handling time within the device and allow the device to manage an increased number of frames without degrading performance

One of two methods to switch frames

Store-and-forward switching

- This method makes a forwarding decision on a frame after it has received the **entire frame** and **checked** the frame for errors using a mathematical error-checking mechanism
    - Cyclic redundancy check (CRC)
- Cisco's primary LAN switching method

Cut-through switching

- This method begins the forwarding process **after** the destination MAC address of an incoming frame and the **egress port** have been **determined**

Store-and-Forward Switching

- Two primary characteristics

Error checking

- After receiving the entire frame on the ingress port, the switch compares the **frame check sequence (FCS)** value in the last field of the datagram against its own FCS calculations.

- The FCS is an error checking process that helps to ensure that the frame is free of physical and data-link errors
    - If the frame is error-free, the switch forwards the frame.
    - Otherwise, the frame is dropped

Automatic buffering

- The ingress port buffering process used by store-and-forward switches provides the flexibility to support any mix of Ethernet speeds
- Drops frames that do not pass the FCS check
    - Does not forward invalid frames

Example:

Handing an incoming frame traveling into a 100 Mbps Ethernet port that must be sent out a 1 Gbps interface would require using the store-and-forward method

- With any mismatch in speeds between the ingress and egress ports, the switch:
    1. Stores the entire frame in a buffer
    2. Computes the FCS check
    3. Forwards it to the egress port buffer
    4. Sends it

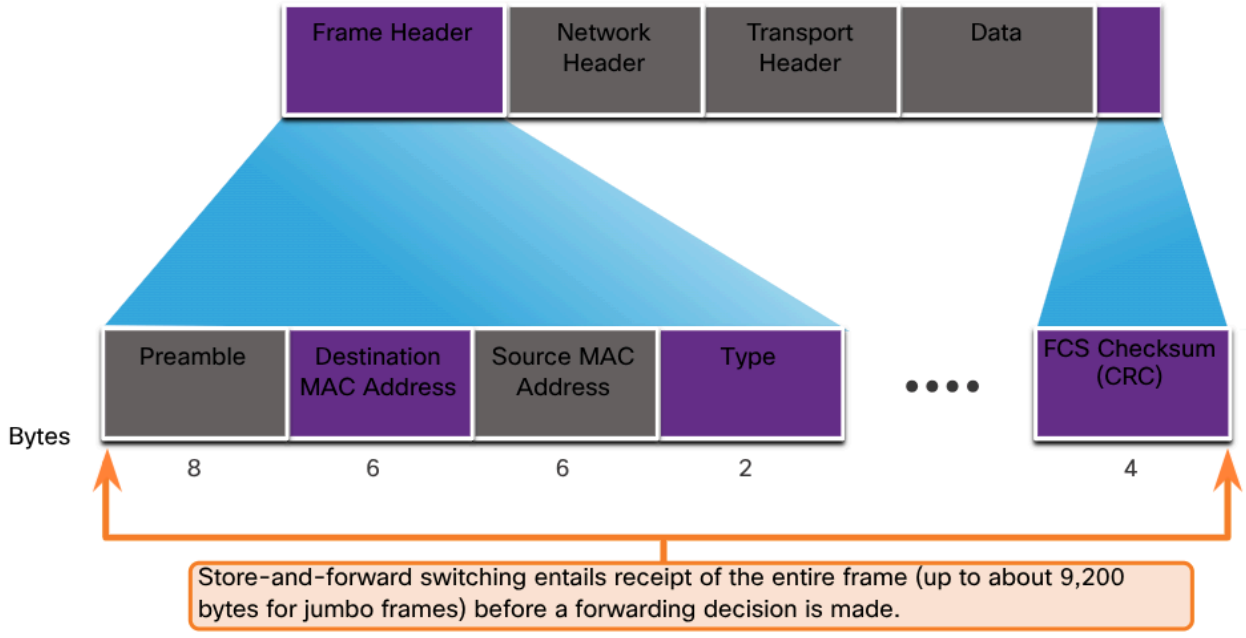

Cut-Through Switching

- May forward invalid frames because no FCS check is performed
  - Forwarding frames with errors
  - If there is a high error rate (invalid frames) in the network, this method can have negative impact on bandwidth
- Has ability to perform **rapid frame switching**
  - The switch can make a forwarding decision as soon as it has looked up the destination MAC address of the frame in its MAC address table
- Switch does not have to wait for rest of frame to enter the ingress port before making forwarding decision

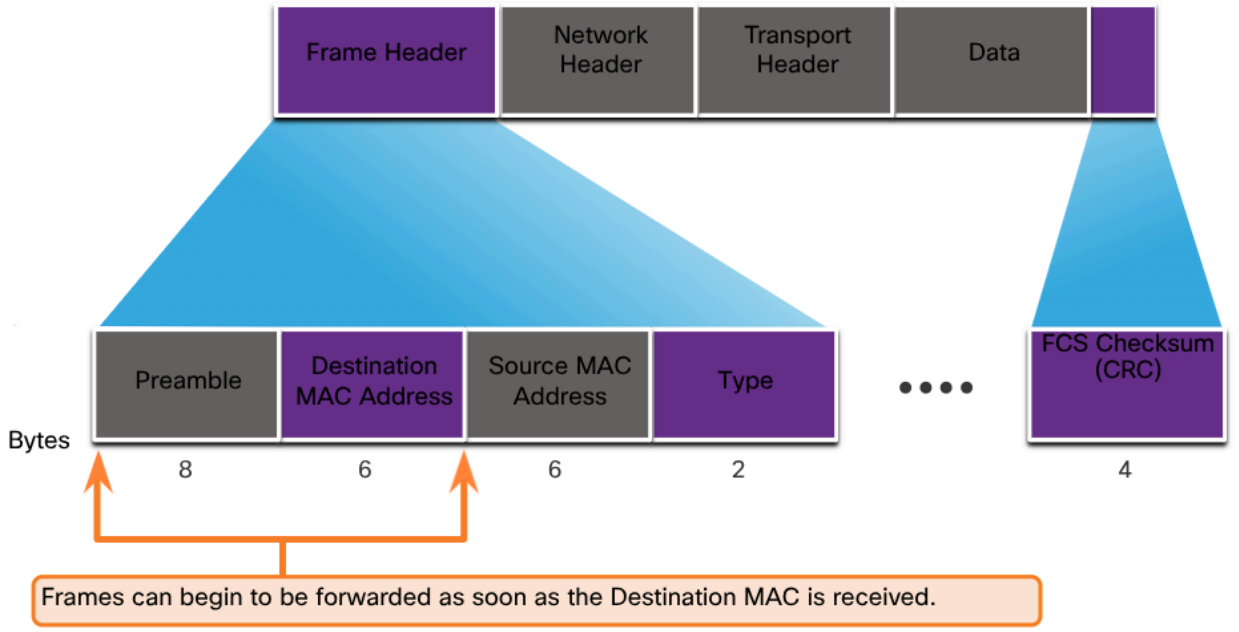

Fragment free switching

- Modified form of cut-through switching in which the switch only starts forwarding the frame
- Provides better error checking than cut-through, with practically no increase in latency
  - Makes it more appropriate for extremely demanding, high-performance computing (HPC) applications that require process-to-process latencies of 10 microseconds or less

Collision Domains

Legacy hub-based ethernet segments

- Network devices competed for the shared medium
- Segments that share the same bandwidth between devices
    - Known as **collision domains**

Half-duplex - Each segment is in its own collision domain

Full-duplex - No collisions

By default, Ethernet switch ports will autonegotiate full-duplex when adjacent device can also operate in full-duplex

- If switch port is connected to device operating in half-duplex, switch port will operate in half-duplex
    - Switch port will be a part of a collision domain

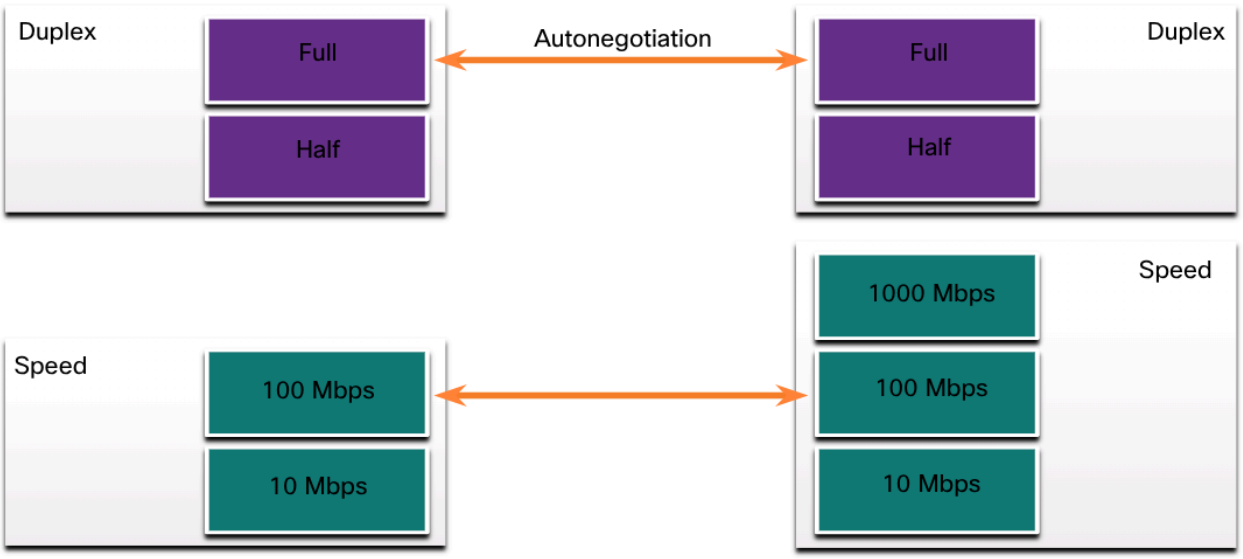

Broadcast Domains

- Collection of interconnected switches forms a single broadcast domain
- **Only** a network layer device can **divide** a Layer 2 broadcast domain
- **Routers** are used to segment broadcast domains, but will also segment a collision domain

Layer 2 broadcast domain is referred to as the MAC broadcast domain

- Consists of all devices on the LAN that receive broadcast frames from host

When a switch receives a broadcast frame, it forwards the frame out each of its ports, **except** the ingress port

Broadcasts

- Sometimes necessary for initially locating other devices and network services, but also **reduce** network efficiency

Network bandwidth

- Used to propagate the broadcast traffic
    - Too many broadcasts and a heavy traffic load on network can result in **congestion**, which slows down network performance

When two switches are connected together, broadcast domain is increased

Alleviate Network Congestion

Lan Switches

- Have special characteristics that help alleviate network congestion
- By default, interconnected switch ports attempt to establish a link in full-duplex
    - Each full-duplex port provides the full bandwidth to the device or devices connected to the port
    - Required for 1 Gbps speeds and higher

Characteristics

- Fast port speeds - Ethernet switch port speeds vary by model and purpose
    - Switches with faster port speeds cost more but can reduce congestion
- Fast internal switching - Switches use a fast internal bus or shared memory to provide high performance
- Large frame buffers - Switches use large memory buffers to temporarily store more received frames before having to start dropping them.
    - This enables **ingress** traffic from a **faster** port to be forwarded to a **slower egress** port without losing frames

- High port density - Lowers overall costs because it reduces the number of switches required.
    - Example: If 96 access ports required, it's less expensive to buy two 48-port instead of four 24-port switches
    - Also keeps traffic local, alleviating congestion

Frame Forwarding

The decision on how a switch forwards traffic is based on the flow of that traffic. The term ingress describes the port where a frame enters a device. The term egress describes the port that frames will use when leaving the device. An Ethernet frame will never be forwarded out the port where it entered. For a switch to know which port to use to transmit a frame, it must first learn which devices exist on each port. As the switch learns the relationship of ports to devices, it builds a table called a MAC address table. Every frame that enters a switch is checked for new information to learn by examining the source MAC address of the frame and port number where the frame entered the switch. If the destination MAC address is a unicast address, the switch will look for a match between the destination MAC address of the frame and an entry in its MAC address table. Switch forwarding methods include store-and-forward and cut-through. Store-and-forward uses error-checking and automatic buffering. Cut-through does not error check. Instead it performs rapid frame switching. This means the switch can make a forwarding decision as soon as it has looked up the destination MAC address of the frame in its MAC address table.

Switching Domains

If an Ethernet switch port is operating in half-duplex, each segment is in its own collision domain. There are no collision domains when switch ports are operating in full-duplex. By default, Ethernet switch ports will autonegotiate full-duplex when the adjacent device can also operate in full-duplex. A collection of interconnected switches forms a single broadcast domain. Only a network layer device, such as a router, can divide a Layer 2 broadcast domain. The Layer 2 broadcast domain is referred to as the MAC broadcast domain. The MAC broadcast domain consists of all devices on the LAN that receive broadcast frames from a host. When a switch receives a broadcast frame, it forwards the frame out each of its ports, except the ingress port where the broadcast frame was received. Each device connected to the switch receives a copy of the broadcast frame and processes it. Switches can: interconnect LAN segments, use a MAC address table to determine egress ports, and can lessen or eliminate collisions entirely. Characteristics of switches that alleviate network congestion are fast port speeds, fast internal switching, large frame buffers, and high port density.

VLANs

Definition

- Provides segmentation and organizational flexibility in a switched network.
- Based on logical connections, instead of physical connections

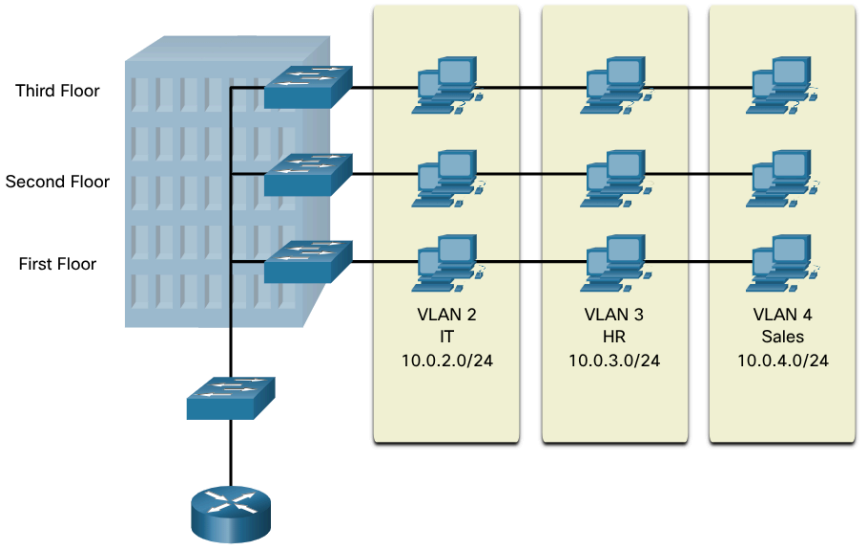

- Allows an administrator to segment networks based on factors such as function, team, or application

- Each VLAN is considered a separate logical network

- * Any switch port can belong to a VLAN

Unicast, broadcast, and multicast packets are forwarded and flooded only to end devices within the VLAN where the packets are sourced.

- Packets destined for devices that do not belong to the VLAN must be forwarded through a device that supports routing

Note - Multiple IP subnets can exist on a switched network, without the use of multiple VLANs

- However, the devices will be in the **same** Layer 2 broadcast domain
    - For example: ARP request, will be received by all devices on the switched network, even by those not intended to receive the broadcast

A VLAN creates a logical broadcast domain that can span multiple physical LAN segments

- Improves network performance by separating large broadcast domains into smaller ones
- If a device in one VLAN sends a broadcast Ethernet frame, all devices in the VLAN receive the frame, but devices in other VLANs do not
- Network administrators can implement access and security policies according to specific groupings of users

Note - Each switch port can be assigned to only **one** VLAN

- Except for a port connected to an IP phone or to another switch

Benefits

- Each VLAN corresponds to an IP network
    - VLAN design must take into consideration the implementation of a hierarchical network-addressing scheme

Hierarchical Network Addressing

- IP network numbers are applied to network segments or VLANs in a way that take sthe network as a whole into consideration
- Blocks of contiguous network addresses are reserved for and configured on devices in a specific area of the network

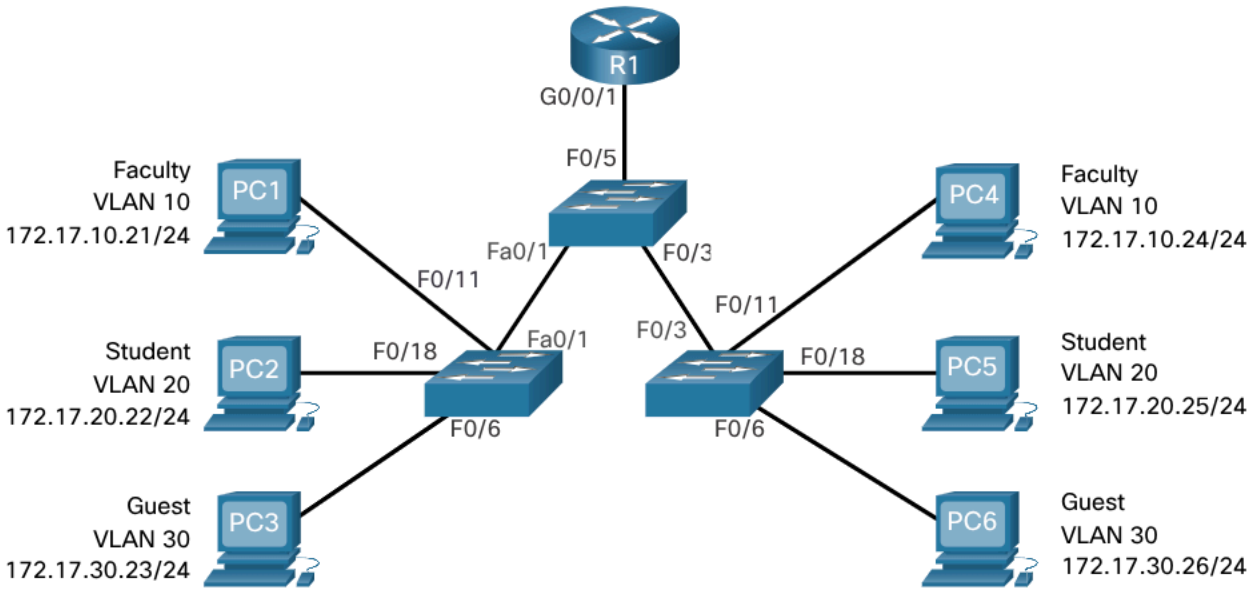

Benefit	Description
Smaller broadcast domains	• Dividing a network into VLANs reduces the number of devices in the broadcast domain
Improved Security	• Only users in the same VLAN can communicate together • Only users in the same VLAN can communicate without the services of a router. ○ The router may have a security feature such as an access control list to restrict communications between VLANs
Improved IT efficiency	• VLANs simplify network management because users with similar network requirements can be configured on the same VLAN • VLANs can be named to make them

	easier to identify
Reduced cost	VLANs reduce the need for expensive network upgrades and use the existing bandwidth and uplinks more efficiently, resulting in cost savings
Better performance	Smaller broadcast domains reduce unnecessary traffic on the network and improve performance
Simpler project and application management	● VLANs aggregate users and network devices to support business or geographic requirements ● Having separate functions makes managing a project or working with a specialized application easier ○ Example: An application is an e-learning development platform for faculty

Types of VLANs

Default VLAN

- On Cisco switch: VLAN 1
    - Therefore, all switch ports are on VLAN 1 unless configured otherwise
    - By default, all Layer 2 control traffic is associated with VLAN 1

Important

- All ports are assigned to VLAN 1 by default
- The native VLAN is VLAN 1 by default
- The management VLAN is VLAN 1 by default
- VLAN 1 cannot be renamed nor deleted

```
Switch# show vlan brief
VLAN Name               Status    Ports
---- ----------------- -------  --------------------
1    default            active   Fa0/1, Fa0/2, Fa0/3, Fa0/4
                                 Fa0/5, Fa0/6, Fa0/7, Fa0/8
                                 Fa0/9, Fa0/10, Fa0/11, Fa0/12
                                 Fa0/13, Fa0/14, Fa0/15, Fa0/16
                                 Fa0/17, Fa0/18, Fa0/19, Fa0/20
                                 Fa0/21, Fa0/22, Fa0/23, Fa0/24
                                 Gi0/1, Gi0/2
1002 fddi-default                      act/unsup
1003 token-ring-default                act/unsup
1004 fddinet-default                   act/unsup
1005 trnet-default                     act/unsup
```

Data VLAN

- VLANs configured to separate user-generated traffic
- Referred to as user VLANs because they separate the network into groups of users or devices

Note - Voice and network management traffic should **not** be permitted on **data** VLANs

Native VLAN

- User traffic from a VLAN must be tagged with its VLAN ID when sent to another switch

Trunk ports

- Used between switches to support the transmission of tagged traffic
  - 802.1Q trunk port inserts a **4-byte tag** in ethernet frame header to identify the VLAN the frame belongs to
  - 802.1Q trunk ports places untagged traffic on the native VLAN (Cisco: VLAN 1)

Switch may also have to send untagged traffic across a trunk link

- Untagged traffic is generated by a switch and may also come from legacy devices

Note - Best practice is to configure native VLAN as an unused VLAN, distinct from VLAN 1 and others

Management VLAN

- Data VLAN configured specifically for network management traffic including SSH, Telnet, HTTPS, HTTP, and SNMP
- By default, VLAN 1 is configured as the management VLAN on a layer 2 switch

Voice VLAN

Separate VLAN is needed to support Voice Over IP (VoIP)

- Assured bandwidth to ensure voice quality
- Transmission priority over other types of network traffic
- Ability to be routed around congested areas on the network
- Delay of less than 150ms across the network

Entire network has to be designed to support VoIP

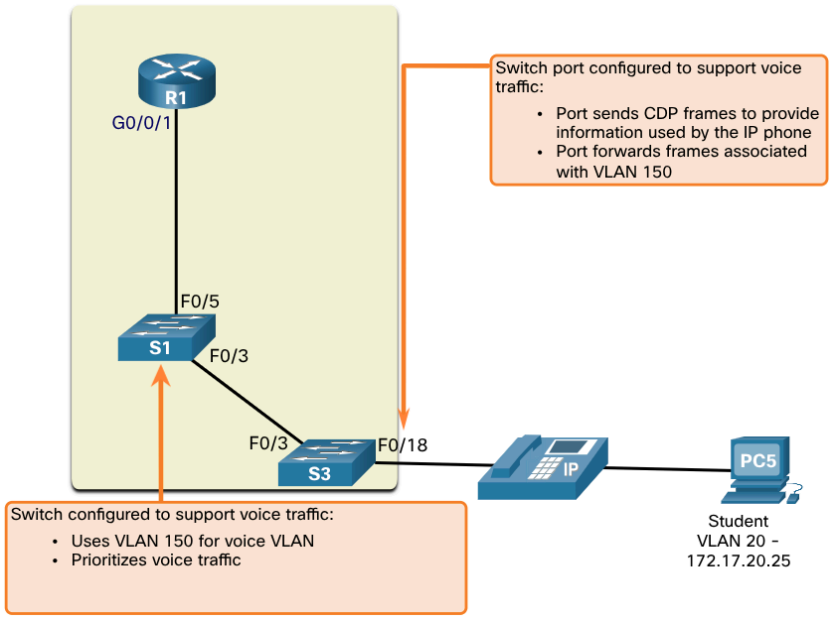

VLANs in a Multi-Switched Environment

Defining VLAN Trunks

- Allows all VLAN traffic to propagate between switches
- Enables devices connected to different switches but in the same VLAN to communicate without going through a router
- Point-to-point link between two network devices that carries more than one VLAN
    - Extends VLANs across an entire network

Cisco supports IEEE 802.1Q for coordinating trunks on:

- Fast Ethernet
- Gigabit Ethernet
- 10-Gigabit Ethernet interfaces

VLAN Trunk (cont)

- Does not belong to a specific VLAN
    - Conduit for multiple VLANs between switches and routers
- Can also be used between a network device and server or another device that is equipped with an appropriate 802.1Q-capable NIC
    - By default, on Cisco Catalyst switch, all VLANs are supported on a trunk port

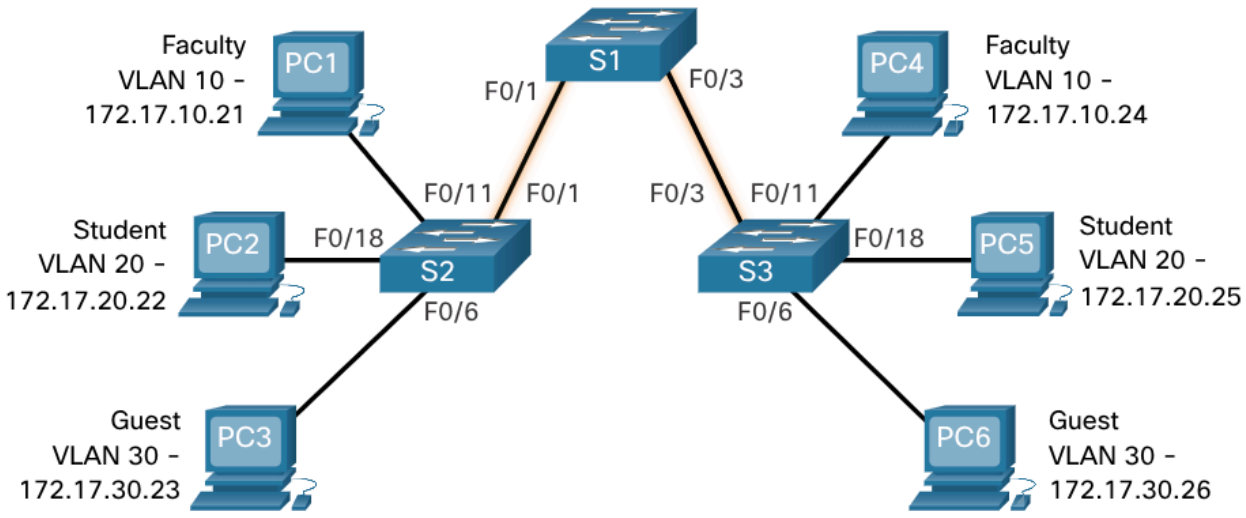

Network without VLANs

- When switch receives a broadcast frame, it forwards the frame out all other ports except for ingress

Network with VLANs

- VLANs associated with and configured on individual switch ports
- Equivalent to an IP network (or subnet)
- VLANs are configured on switch
    - Whereas IP addressing is configured on device

Note - When VLANs are implemented on a switch, the transmission of unicast, multicast, and broadcast traffic from a host in a particular VLAN are restricted to the devices that are in that VLAN

VLAN Identification with a Tag

- Standard ethernet frame header does **not** contain information about the VLAN
    - When ethernet frames are placed on a trunk, information about the VLANs to which they belong to must be added

Tagging

- Accomplished by using the IEEE 802.1Q header, specified in the IEEE 802.1Q standard
- Includes a **4-byte** tag inserted within the original ethernet frame header

When switch receives a frame on a port configured in **access** mode and assigned a VLAN, the switch inserts a VLAN tag in the frame header, recalculates the Frame Check Sequence (FCS), and sends the tagged frame out of a trunk port

VLAN Tag Field Details

- Type - A 2-byte value called the **tag protocol ID (TPID)** value.
    - Ethernet: it is set to hexadecimal 0x8100
- User Priority - A 3-bit value that supports level or service implementation
- Canonical Format Identifier (CFI) - A 1-bit identifier that enables **Token Ring** frames to be carried across Ethernet links
- VLAN ID (VID) - A 12-bit VLAN identification number that supports up to 4096 VLAN IDs

After switch inserts the tag control information fields, it recalculates the FCS values and inserts the new FCS into the frame

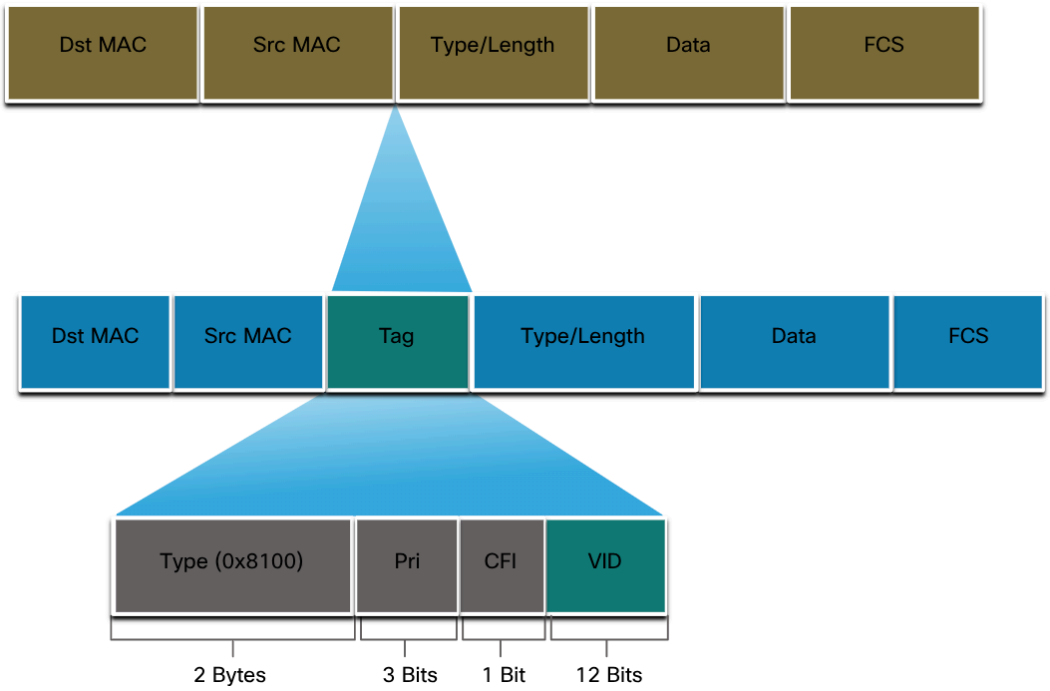

Native VLANs and 802.1Q Tagging

IEEE 802.1Q standard specifies a native VLAN for trunk links

- Defaults to VLAN 1

When an untagged frame arrives on a trunk port it is assigned to the native VLAN.

- Example: Management frames sent between switches

If link between two switches is a trunk, the switch sends the untagged traffic on the native VLAN

Tagged Frames on the Native VLAN

- Devices that support trunking add a VLAN tag to native VLAN traffic
- Control traffic sent on native VLAN should not be tagged
- If 802.1Q trunk port receives a tagged frame with the VLAN ID that is same as native VLAN, it drops the frame

Untagged Frames on the Native VLAN

- When a Cisco switch trunk port receives untagged frames (unusual), it forward those frames to the native VLAN
- If there are no devices associated with native VLAN (not unusual) and there are no other trunk ports (not unusual), the frame is dropped
- * Default native VLAN is VLAN 1
    - When configuring 802.1Q trunk port, a default Port VLAN ID (PVID) is assigned the value of the native VLAN ID

- All untagged traffic coming in or out of 802.1Q port is forwarded based on the PVID value

Example: If VLAN 99 is native, the PVID is 99 and all untagged traffic is forwarded to VLAN 99

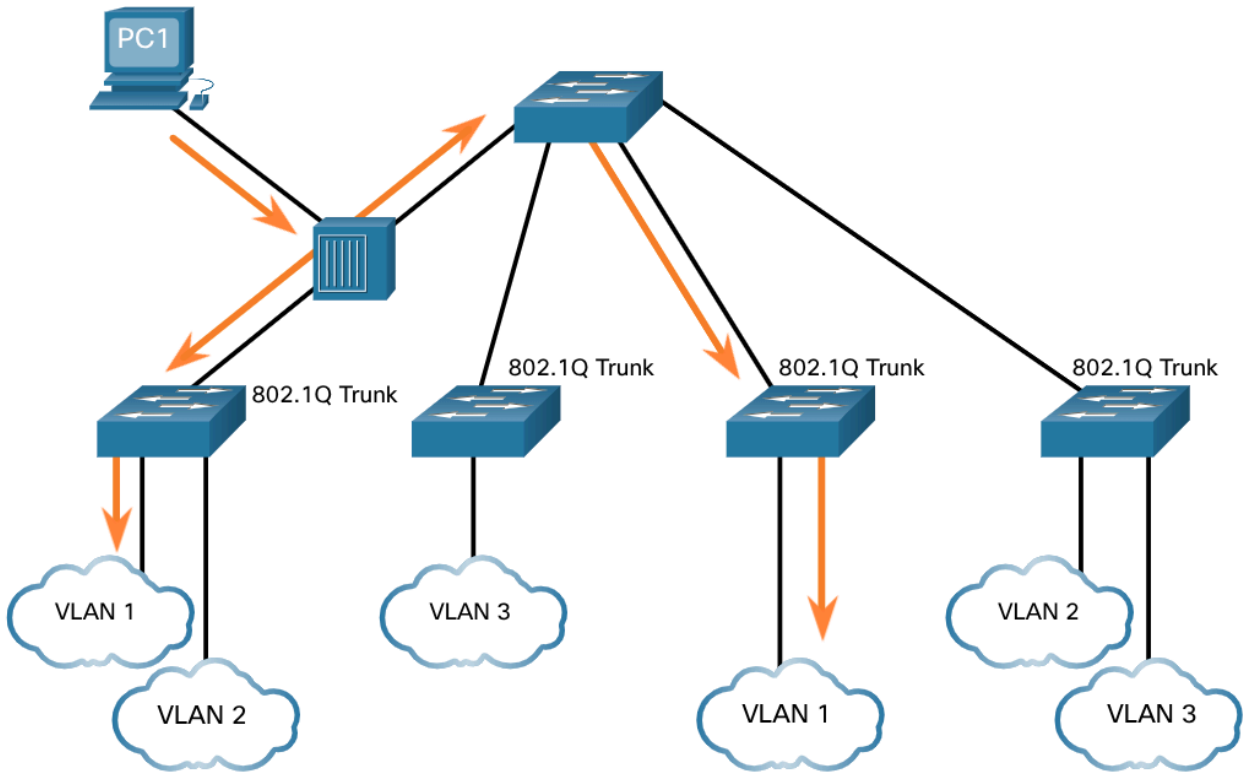

Voice VLAN Tagging

Separate voice VLAN is required to support VoIP

- Enables quality of server (QoS) and security policies to be applied to voice traffic

Note - An IP host can connect to the IP Phone to gain network connectivity as well (dual mac)

- The access port connected to the Cisco IP phone can be configured to use two separate VLANs
  - One VLAN is for traffic, the other is a data VLAN to support the host traffic
  - The link between the switch and the IP phone simulates a trunk link to carry both voice VLAN traffic and data VLAN traffic

The Cisco IP phone contains an integrated three-port 10/100 switch

- Port 1 connects to the switch or other VoIP device
- Port 2 is an internal 10/100 interface that carries the IP phone traffic
- Port 3 (access port) connects to a PC or other device

The switch access port sends CDP packets instructing the attached IP phone to send voice traffic in one of three ways

- Voice VLAN traffic must be tagged with an appropriate Layer 2 **class of service** (CoS) priority value
- Access VLAN traffic can also be tagged with a Layer 2 CoS priority value
- Access VLAN is not tagged (no Layer 2 CoS priority value)

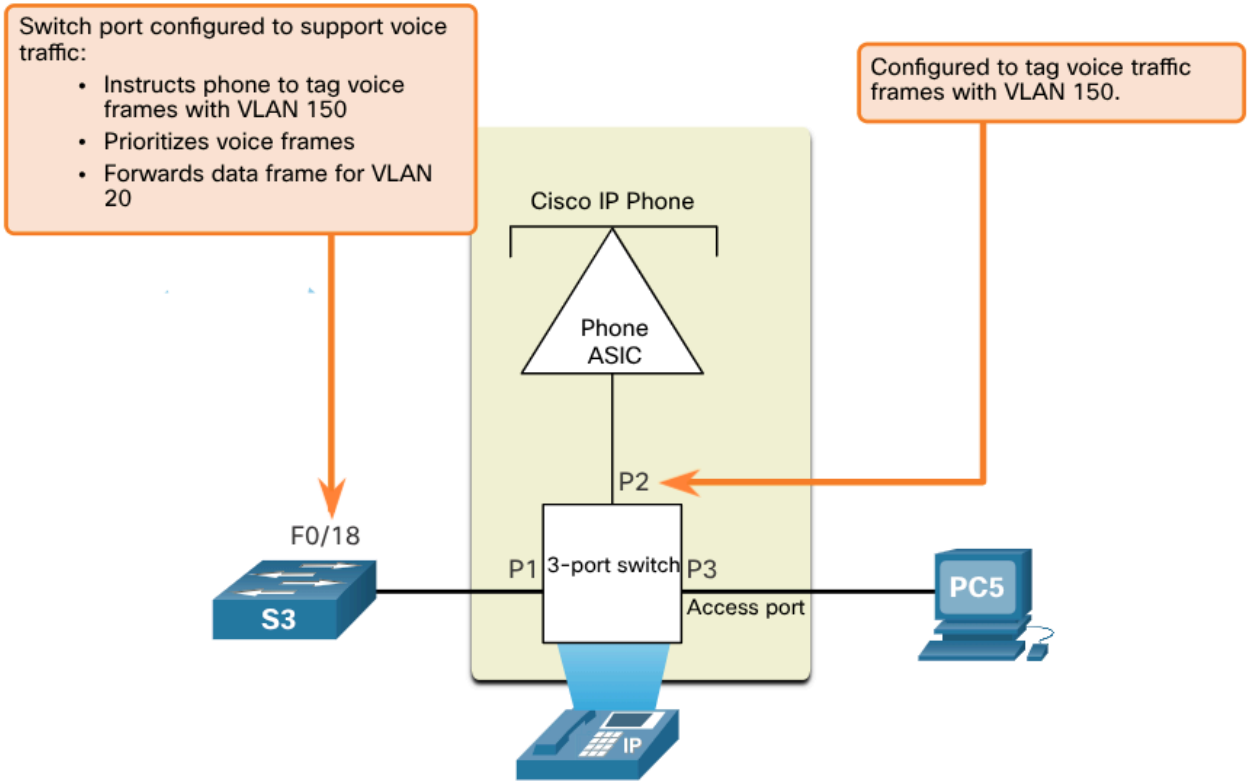

Voice VLAN Verification Example

Example for *show interface fa0/18 switchport* command

```
S1# show interfaces fa0/18 switchport
Name: Fa0/18
Switchport: Enabled
Administrative Mode: static access
Operational Mode: static access
Administrative Trunking Encapsulation: negotiate
Operational Trunking Encapsulation: native
Negotiation of Trunking: Off
Access Mode VLAN: 20 (student)
```

```
Trunking Native Mode VLAN: 1 (default)
Administrative Native VLAN tagging: enabled
Voice VLAN: 150 (voice)
```

VLAN Configuration

Vlan Ranges on Catalyst Switches

Catalyst 2960 | 3650

- Over 4,000 VLANs
  - Normal range: 1 - 1,005
  - Extended range: 1,006 - 4,094

```
Switch# show vlan brief
VLAN Name                Status    Ports
---- ----------------- -------  --------------------
1    default         active   Fa0/1, Fa0/2, Fa0/3, Fa0/4
                              Fa0/5, Fa0/6, Fa0/7, Fa0/8
                              Fa0/9, Fa0/10, Fa0/11, Fa0/12
                              Fa0/13, Fa0/14, Fa0/15, Fa0/16
                              Fa0/17, Fa0/18, Fa0/19, Fa0/20
                              Fa0/21, Fa0/22, Fa0/23, Fa0/24
                              Gi0/1, Gi0/2
1002 fddi-default                 act/unsup
1003 token-ring-default           act/unsup
1004 fddinet-default                 act/unsup
1005 trnet-default                   act/unsup
```

Normal Range VLANs

- They are used in all small- and medium-sized business and enterprise networks
- They are identified by a VLAN ID between 1 and 1005
- IDs 1002 - 1005 are reserved for legacy network technologies
  - IE: Token Ring, Fiber Distributed Data Interface
- IDs 1 and 1002 - 1005 are automatically created and **cannot be removed**
- Configurations are stored in the switch flash memory in a VLAN database file called **vlan.dat**
- When configured, **VLAN trunking protocol (VTP)**, helps synchronize the VLAN database between switches

Extended Range VLANs

- They are used by service providers to service multiple customers and by global enterprises large enough to need extended range VLAN IDs
- They are identified by a VLAN ID between 1006 - 4094
- Configurations are saved, by default, in the running configuration
- They support fewer VLAN features than normal range VLANs
- Requires VTP transparent mode configuration to support extended range VLANs

Note - 4096 is the upper boundary for the number of VLANs available on Catalyst switches, because there are **12 bits** in the VLAN ID field of the IEEE 802.1Q header

VLAN Creation Commands

- Configuration details are stored in flash memory on the switch in file called vlan.dat
    - Persistent and does not require the **copy running-config startup-config** command
    - Other details are often configured on a Cisco switch at the same time VLANs are created
        - Good practice to save running configuration changes to startup config

Task	IOS Command
Enter global configuration mode	Switch# configure terminal
Create a VLAN with a valid ID number	Switch(config)# vlan *vlan-id*
Specify a unique name to identify the VLAN	Switch(config-vlan)# name *vlan-name*
Return to the privileged EXE mode	Switch(config-vlan)# end

VLAN Creation Example

** PC2 not associated with VLAN yet

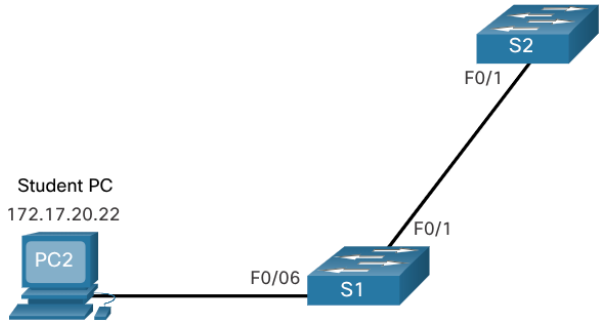


```
S1# configure terminal
S1(config)# vlan 20
S1(config-vlan)# name student
S1(config-vlan)# end
```

Note - A series of VLAN IDs can be entered separated by **commas**, or range of VLAN IDs separated by **hyphens** using the **vlan** *vlan-id* command

- Example: vlan 100,102,105-107

VLAN Port Assignment Commands

switchport mode access command

- Optional, but strongly recommended as security best practice
- The interface changes to strictly access mode
    - Indicates that the port belongs to a single VLAN and will not negotiate to become a trunk link

Task	IOS Command
Enter global configuration mode	Switch# configure terminal
Enter interface configuration mode	Switch(config)# interface *interface-id*
Set the port to access mode	Switch(config-if)# switchport mode access
Assign the port to a VLAN	Switch(config-if)# switchport access vlan *vlan-id*
Return to the privileged EXEC mode	Switch(config-if)# end

Note - *interface range* command to simultaneously configure multiple interfaces

VLAN Port Assignment Example

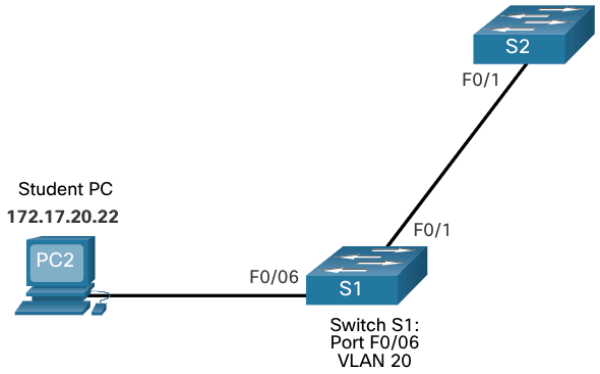


```
S1# configure terminal
S1(config)# interface fa0/6
S1(config-if)# switchport mode
access
S1(config-if)# switchport
access vlan 20
S1(config-if)# end
```

- VLANs are configured on the switch port, not on end device
- When VLAN 20 is configured on the other switches
    - Must configure the other student computers to be in same subnet (172.17.20.0/24)

Data and Voice VLANs

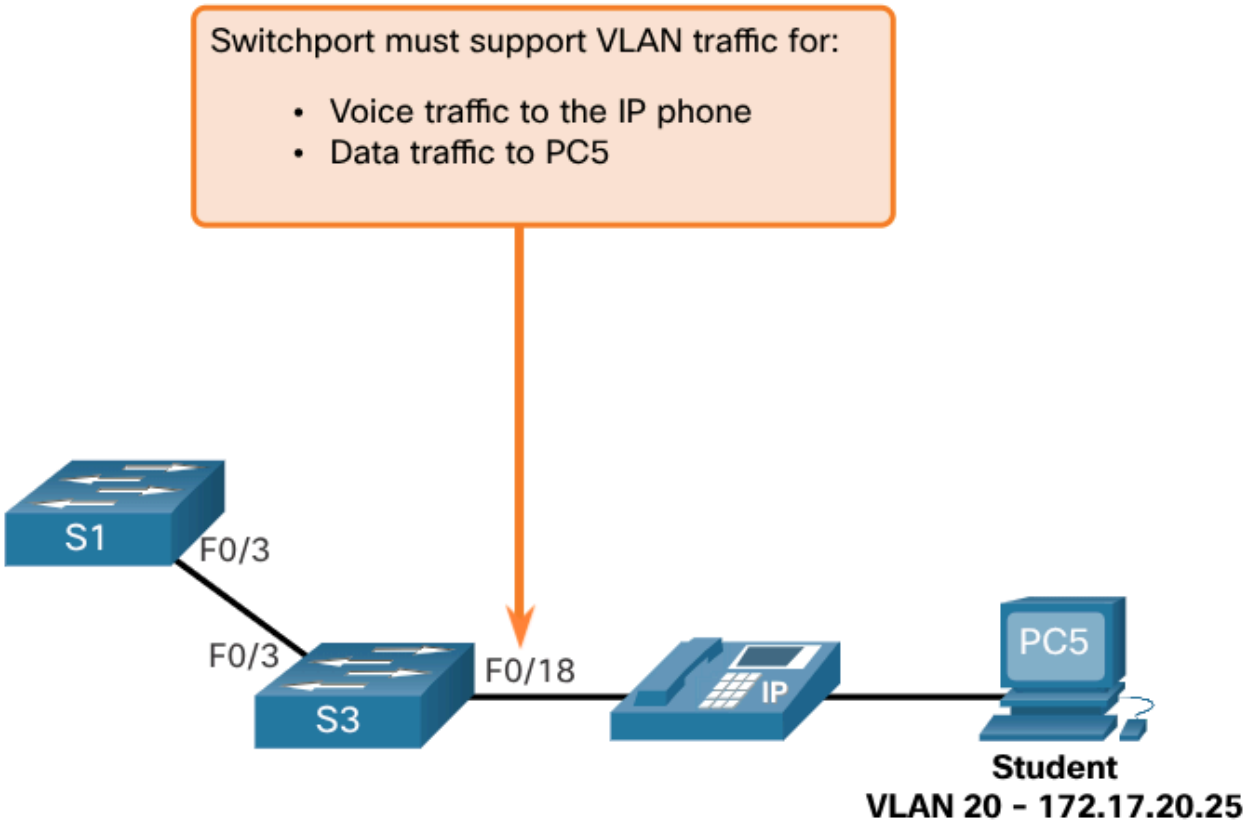

Data and Voice VLAN Example

Assign a voice VLAN to a port

- ***switchport voice vlan*** *vlan-id*

Voice Traffic must be labeled as trusted as soon as it enters the network

Set the trusted state of an interface and indicate which fields of the packet are use to classify traffic

- **mls qos trust [cos | device cisco-phone | dscp | ip-precedence]**

```
S3(config)# vlan 20
S3(config-vlan)# name student
S3(config-vlan)# vlan 150
S3(config-vlan)# name VOICE
S3(config-vlan)# exit
S3(config)# interface fa0/18
S3(config-if)# switchport mode access
S3(config-if)# switchport access vlan 20
S3(config-if)# mls qos trust cos
S3(config-if)# switchport voice vlan 150
S3(config-if)# end
S3#
```

- Creates two VLANs (VLAN 20, and VLAN 150)
- Then assigns F0/18 interface of S3 as a switchport in VLAN 20
- Also assigns voice traffic to VLAN 150 and enables QoS classification based on the class of service (CoS) assigned by the IP phone

The **switchport access vlan** command forces the creation of a VLAN if it does not already exist on the switch

Example: VLAN 30 is not present in **show vlan brief** output

- If **switchport access vlan 30** command is entered on any interface with no previous configuration, switch displays:

```
% Access VLAN does not exist. Creating vlan 30
```

Verify VLAN Information

- Can be validated using Cisco IOS **show** commands

show vlan command displays list of all configured VLANs

- Can also be used with options
- Complete syntax: **show vlan [brief | id *vlan-id* | name *vlan-name* | summary]**

Task	Command Option
Display VLAN name, status, and its ports one VLAN per line	brief
Display information about the identified VLAN ID number	id *vlan-id*
Display information about the identified VLAN name - *vlan-name* is an ASCII string from 1 to 32 characters	name *vlan-name*
Display VLAN summary information	summary

show vlan summary command

- Displays the count of all configured VLANs

```
S1# show vlan summary
Number of existing VLANs              : 7
Number of existing VTP VLANs          : 7
Number of existing extended VLANS     : 0
```

Other useful commands

- **show interfaces** *interface-id* **switchport**
- **show interfaces vlan** *vland-id*

Example:

show interfaces fa0/18 switchport command

- Can be used to confirm that the FastEthernet 0/18 port has been correctly assigned to data and voice VLANs

```
S1# show interfaces fa0/18 switchport
Name: Fa0/18
Switchport: Enabled
Administrative Mode: static access
Operational Mode: static access
Administrative Trunking Encapsulation: dot1q
Operational Trunking Encapsulation: native
Negotiation of Trunking: Off
Access Mode VLAN: 20 (student)
Trunking Native Mode VLAN: 1 (default)
Voice VLAN: 150
Administrative private-vlan host-association: none
(Output omitted)
```

Change VLAN Port Membership

If switch access port has been incorrectly assigned to a VLAN

- Simply re-enter **switchport access vlan** *vlan-id* command with correct VLAN ID

Example:

Assume– Fa0/18 incorrectly configured to be on default VLAN 1 instead of VLAN 20

- **switchport access vlan 20**

To change back to default VLAN 1

- **no switchport access vlan** command

```
S1(config)# interface fa0/18
S1(config-if)# no switchport access vlan
S1(config-if)# end
S1#
S1# show vlan brief
VLAN Name                 Status    Ports
---- ------------------ --------- -------------------------------
1    default            active    Fa0/1, Fa0/2, Fa0/3, Fa0/4
                                  Fa0/5, Fa0/6, Fa0/7, Fa0/8
                                  Fa0/9, Fa0/10, Fa0/11, Fa0/12
                                  Fa0/13, Fa0/14, Fa0/15, Fa0/16
                                  Fa0/17, Fa0/18, Fa0/19, Fa0/20
                                  Fa0/21, Fa0/22, Fa0/23, Fa0/24
                                  Gi0/1, Gi0/2
```

```
20   student              active
1002 fddi-default           act/unsup
1003 token-ring-default  act/unsup
1004 fddinet-default     act/unsup
1005 trnet-default       act/unsup
```

Note - VLAN 20 is still active, even though no ports are assigned to it

show interfaces f0/18 switchport output

- Can also be used to verify that the access VLAN F0/18 has been reset to VLAN 1

```
S1# show interfaces fa0/18 switchport
Name: Fa0/18
Switchport: Enabled
Administrative Mode: static access
Operational Mode: static access
Administrative Trunking Encapsulation: negotiate
Operational Trunking Encapsulation: native
Negotiation of Trunking: Off
Access Mode VLAN: 1 (default)
Trunking Native Mode VLAN: 1 (default)
```

Delete VLANs

no vlan *vlan-id*

- Used to remove a VLAN from switch vlan.dat file

** **Caution**: Before deleting a VLAN, reassign all member ports to a different VLAN first

- Any ports not moved to an active VLAN are unable to communicate with other hosts after VLAN is deleted and until they are assigned to an active VLAN

delete flash:vlan.dat

- Delete entire vlan.dat file

delete vlan.dat

- Abbreviated version
- Can be used if vlan.dat file has not been moved from default location
- Effectively places switch into factory default condition regarding VLAN configurations

Note - To restore a Catalyst switch to factory default, unplug all cables **except console and power cable** from switch. Then enter **erse startup-config** command followed by **delete vlan.dat** command

- Enter global configuration mode.
- Create VLAN 20.
- Name the VLAN **student**.
- Return to privileged EXEC mode.

```
S1# configure terminal

Enter configuration commands, one per line. End with CNTL/Z.

S1(config)# vlan 20
S1(config-vlan)# name student
S1(config-vlan)# end

*Mar 31, 08:55:14.5555: %SYS-5-CONFIG_I: Configured from console by console
```

```
S1# show vlan brief

VLAN Name                             Status    Ports
---- -------------------------------- --------- -------------------------------
1    default                          active    Fa0/1, Fa0/2, Fa0/3, Fa0/4
                                                Fa0/5, Fa0/6, Fa0/7, Fa0/8
                                                Fa0/9, Fa0/10, Fa0/11, Fa0/12
                                                Fa0/13, Fa0/14, Fa0/15, Fa0/16
                                                Fa0/17, Fa0/18, Fa0/19, Fa0/20
                                                Fa0/21, Fa0/22, Fa0/23, Fa0/24
                                                Gi0/1, Gi0/2
20   student                          active
1002 fddi-default                     act/unsup
1003 token-ring-default               act/unsup
1004 fddinet-default                  act/unsup
1005 trnet-default                    act/unsup
```

Complete the following steps to create a voice VLAN:

- Enter global configuration mode.
- Create VLAN 150.
- Name the VLAN **VOICE**.
- Return to global configuration mode.

```
S1# configure terminal
S1(config)# vlan 150
S1(config-vlan)# name VOICE
S1(config-vlan)# exit
```

Complete the following steps to assign the data and voice VLANs to a port:

- Enter interface configuration mode. Use **fa0/18** as the interface designation.
- Configure the port as an access port.
- Assign the data VLAN 20 to the port.
- Enable QoS settings with the **mls qos trust cos** command.
- Assign the voice VLAN 150 to the port.
- Return to privileged EXEC mode.

```
S1(config)# interface fa0/18
S1(config-if)# switchport mode access
S1(config-if)# switchport access vlan 20
S1(config-if)# mls qos trust cos
S1(config-if)# switchport voice vlan 150
S1(config-if)# end
```

```
S1# show vlan brief

LAN Name                             Status    Ports
---- --------------------------------- --------- -------------------------------
1    default                          active    Fa0/1, Fa0/2, Fa0/3, Fa0/4
                                                Fa0/5, Fa0/6, Fa0/7, Fa0/8
                                                Fa0/9, Fa0/10, Fa0/11, Fa0/12
                                                Fa0/13, Fa0/14, Fa0/15, Fa0/16
                                                Fa0/17, Fa0/19, Fa0/20, Fa0/21
                                                Fa0/22, Fa0/23, Fa0/24, Gi0/1
                                                Gi0/2
20   student                          active    Fa0/18
150  VOICE                            active    Fa0/18
1002 fddi-default                     act/unsup
1003 token-ring-default               act/unsup
1004 fddinet-default                  act/unsup
```

```
S1# configure terminal
S1(config)# interface fa0/18
S1(config-if)# no switchport access vlan
S1(config-if)# do show vlan brief

VLAN Name                             Status    Ports
---- --------------------------------- --------- -------------------------------
1    default                          active    Fa0/1, Fa0/2, Fa0/3, Fa0/4
                                                Fa0/5, Fa0/6, Fa0/7, Fa0/8
                                                Fa0/9, Fa0/10, Fa0/11, Fa0/12
                                                Fa0/13, Fa0/14, Fa0/15, Fa0/16
                                                Fa0/17, Fa0/18, Fa0/19, Fa0/20
                                                Fa0/21, Fa0/22, Fa0/23, Fa0/24
                                                Gi0/1, Gi0/2
20   student                          active
150  VOICE                            active    Fa0/18
```

```
1002 fddi-default                     act/unsup
1003 token-ring-default               act/unsup
1004 fddinet-default                  act/unsup
1005 trnet-default                    act/unsup
```

- Enter interface configuration mode. Use **fa0/11** as the interface designation.
- Assign VLAN 20 to the port.
- Return to privileged EXEC mode.

```
S1(config-if)# interface fa0/11
S1(config-if)# switchport mode access
S1(config-if)# switchport access vlan 20
S1(config-if)# end
```

```
S1# show vlan brief

LAN Name                             Status    Ports
---- -------------------------------- --------- -------------------------------
1    default                          active    Fa0/1, Fa0/2, Fa0/3, Fa0/4
                                                Fa0/5, Fa0/6, Fa0/7, Fa0/8
                                                Fa0/9, Fa0/10, Fa0/12, Fa0/13
                                                Fa0/14, Fa0/15, Fa0/16, Fa0/17
                                                Fa0/19, Fa0/20, Fa0/21, Fa0/22
                                                Fa0/23, Fa0/24, Gig0/1, Gig0/2
20   student                          active    Fa0/11
150  VOICE                            active    Fa0/18
1002 fddi-default                     act/unsup
1003 token-ring-default               act/unsup
1004 fddinet-default                  act/unsup
1005 trnet-default                    act/unsup
```

```
S1# show vlan name student

VLAN Name                             Status    Ports
---- -------------------------------- --------- -------------------------------
20   student                          active    Fa0/11, Fa0/18
VLAN Type  SAID       MTU   Parent RingNo BridgeNo Stp  BrdgMode Trans1 Trans2
---- ----- ---------- ----- ------ ------ -------- ---- -------- ------ ------
20   enet  100020     1500  -      -      -        -    -        0      0
```

```
S1# show vlan summary

Number of existing VLANs : 7
Number of existing VTP VLANs : 7
Number of existing extended VLANS : 0
```

```
S1# show interface fa0/11 switchport

Name: Fa0/11
Switchport: Enabled
Administrative Mode: static access
Operational Mode: down
Administrative Trunking Encapsulation: dot1q
Operational Trunking Encapsulation: native
Negotiation of Trunking: Off
Access Mode VLAN: 20 (Students)
Trunking Native Mode VLAN: 1 (default)
Voice VLAN: none
(output omitted)
```

Trunk Configuration Commands

VLAN trunk is a Layer 2 link between two switches that carries traffic for all VLANs

- Unless the allowed VLAN list is restricted manually or dynamically

Task	IOS Command
Enter global configuration mode	Switch# configure terminal
Enter interface configuration mode	Switch(config)# interface *interface-id*
Set the port to permanent trunking mode	Switch(config-if)# switchport mode trunk
Sets the native VLAN to something other than VLAN 1	Switch(config-if)# switchport trunk native vlan *vlan-id*
Specify the list of VLANs to be allowed on the trunk link	Switch(config-if)# switchport trunk allowed vlan *vlan-list*
Return to the privileged EXEC mode	Switch(config-if)# end

Trunk Configuration Example

F0/1 port on switch S1 is configured as trunk port and forwards traffic for VLANs 10, 20, 30
VLAN 99 is configured as native VLAN

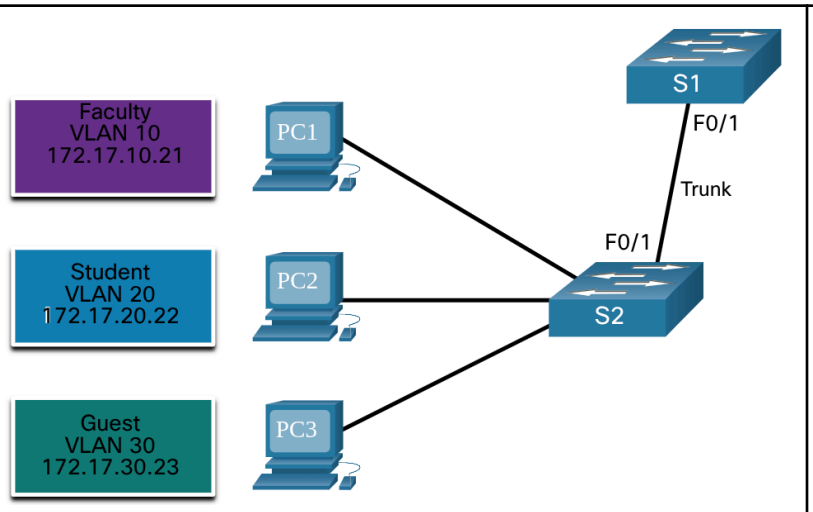

Subnets associated with each VLAN:

VLAN 10 - Faculty/staff
172.17.10.0/24

VLAN 20 - Students
172.17.20.0/24

VLAN 30 - Guests
172.17.30.0/24

VLAN 99 - Native
172.17.99.0/24

```
S1(config)# interface fastEthernet 0/1
S1(config-if)# switchport mode trunk
S1(config-if)# switchport trunk native vlan 99
S1(config-if)# switchport trunk allowed vlan 10,20,30,99
S1(config-if)# end
```

Note - This configuration assumes the use of Catalyst 2960 switches which automatically use 802.1Q encapsulation on trunk links

- Other switches may require manual configuration of encapsulation
- * Always configure both ends of a trunk link with same native VLAN
    - If 802.1Q trunk configuration is not same on both ends, Cisco IOS software reports errors

Verify Trunk Configuration

- Verified with **show interfaces** *interface-id* **switchport** command

```
S1# show interfaces fa0/1 switchport
Name: Fa0/1
Switchport: Enabled
Administrative Mode: trunk
Operational Mode: trunk
Administrative Trunking Encapsulation: dot1q
Operational Trunking Encapsulation: dot1q
Negotiation of Trunking: On
Access Mode VLAN: 1 (default)
Trunking Native Mode VLAN: 99 (VLAN0099)
Voice VLAN: none
Administrative private-vlan host-association: none
Administrative private-vlan mapping: none
Administrative private-vlan trunk native VLAN: none
Administrative private-vlan trunk encapsulation: dot1q
Administrative private-vlan trunk normal VLANs: none
Administrative private-vlan trunk associations: none
Administrative private-vlan trunk private VLANs: none
Operational private-vlan: none
Trunking VLANs Enabled: 10,20,30,99
Pruning VLANs Enabled: 2-1001
(output omitted)
```

Top highlighted area show port f0/1 has its administrative mode set to **trunk**

- Port is in trunking mode

Next highlighted area verifies that native VLAN is VLAN 99

Bottom highlighted area shows VLANs 10, 20, 30, and 99 are enabled on trunk

show interfaces trunk command

- Command for verifying trunk interfaces

Reset the Trunk to the Default State

Remove allowed VLANs and reset native VLAN of trunk

- **no switchport trunk allowed vlan**
- **no switchport trunk native vlan**

When reset to default state, trunk allows all VLANs and uses VLAN 1 as native

```
S1(config)# interface fa0/1
S1(config-if)# no switchport trunk allowed vlan
S1(config-if)# no switchport trunk native vlan
S1(config-if)# end
```

show interfaces fa0/1 switchport command reveals trunk has been reconfigured to default state

```
S1# show interfaces fa0/1 switchport
Name: Fa0/1
Switchport: Enabled
Administrative Mode: trunk
Operational Mode: trunk
Administrative Trunking Encapsulation: dot1q
Operational Trunking Encapsulation: dot1q
Negotiation of Trunking: On
Access Mode VLAN: 1 (default)
Trunking Native Mode VLAN: 1 (default)
Administrative Native VLAN tagging: enabled
Voice VLAN: none
Administrative private-vlan host-association: none
Administrative private-vlan mapping: none
Administrative private-vlan trunk native VLAN: none
Administrative private-vlan trunk Native VLAN tagging: enabled
Administrative private-vlan trunk encapsulation: dot1q
Administrative private-vlan trunk normal VLANs: none
Administrative private-vlan trunk associations: none
Administrative private-vlan trunk mappings: none
Operational private-vlan: none
Trunking VLANs Enabled: ALL
Pruning VLANs Enabled: 2-1001
(output omitted)
```

show interfaces f0/1 switchport command

- reveals that the f0/1 interface is now in static access mode

```
S1(config)# interface fa0/1
S1(config-if)# switchport mode access
S1(config-if)# end
S1# show interfaces fa0/1 switchport
Name: Fa0/1
```

```
Switchport: Enabled
Administrative Mode: static access
Operational Mode: static access
Administrative Trunking Encapsulation: dot1q
Operational Trunking Encapsulation: native
Negotiation of Trunking: Off
Access Mode VLAN: 1 (default)
Trunking Native Mode VLAN: 1 (default)
Administrative Native VLAN tagging: enabled
(output omitted)
```

Introduction to DTP

Dynamic Trunking Protocol (DTP)

- Can speed up the configuration process
- Is a Cisco proprietary protocol, automatically enabled on Catalyst 2960 and Catalyst 3650
- Manages trunk negotiation
    - Only if port on neighbor switch is configured in a trunk mode that supports DTP
- * Switches from other vendors do not support DTP *

Trunking Negotiation

- Managed by DTP, operating on a point-to-point basis only
- Ethernet trunk interfaces support different trunking modes
    - An interface can be set to trunking or nontrunking
        - Or to negotiate trunking with neighbor interface

Caution: Some networking devices might forward DTP frames improperly, causing misconfigurations.

- To avoid, turn off DTP on Cisco switch interfaces that are connected to devices that don't support DTP

Note - Default DTP configuration for Catalyst 3960 and 3650 is **dynamic auto**

To enable trunking from Cisco switch to device that doesn't support DTP

- **switchport mode trunk** and **switchport negotiate**
    - Causes interface to become trunk, but will not generate DTP frames

```
S1(config-if)# switchport mode trunk
S1(config-if)# switchport nonegotiate
```

To re-enable dynamic trunking

- **switchport mode dynamic auto**

```
S1(config-if)# switchport mode dynamic auto
```

If the ports connecting two switches are configured to ignore all DTP advert with **switchport mode trunk** and **switchport nonegotiate** commands

- Ports will stay in trunk port mode

If connecting ports are set to dynamic auto

- They will not negotiate a trunk and will stay in access mode state, creating an inactive trunk link

When configuring a port to be in trunk mode, use **switchport mode trunk** command. Then there is no ambiguity about which state trunk is in; it is always on

Negotiated Interface Modes

switchport mode - full command syntax

```
Switch(config-if)# switchport mode { access | dynamic { auto | desirable } | trunk }
```

Option	Description
Access	● Puts the interface (access port) into permanent nontrunking mode and negotiates to convert the link into a nontrunk link ● The interface becomes a nontrunk interface, regardless of whether the neighboring interface is a trunk interface
dynamic auto	● Makes the interface able to convert the link to a trunk link ● The interface becomes a trunk interface if the neighboring interface is set to trunk or desirable mode ● The default switchport mode for all Ethernet interfaces is **dynamic auto**
dynamic desirable	● Makes the interface actively attempt to convert the link to a trunk link ● The interface becomes a trunk interface if the neighboring interface is set to trunk, desirable, or dynamic auto mode
trunk	● Puts the interface into permanent trunking mode and negotiates to convert the neighboring link into a trunk link ● The interface becomes a trunk interface even if the neighboring interface is not a trunk interface

switchport nonegotiate

- To stop DTP negotiation
- Switch will stop engaging DTP negotiation on this interface
- Can use this command **only** when interface switchport mode is **access** or **trunk**
- Must manually configure neighboring interface as trunk interface to establish a trunk link

Results of a DTP configuration

	Dynamic Auto	**Dynamic Desirable**	**Trunk**	**Access**
Dynamic Auto	Access	Trunk	Trunk	Access
Dynamic Desirable	Trunk	Trunk	Trunk	Access
Trunk	Trunk	Trunk	Trunk	Limited Connectivity
Access	Access	Access	Limited Connectivity	Access

Verify DTP Mode

- Default DTP mode depends on the Cisco IOS Software version and the platform

show dtp interface

- Determines the current DTP mode

```
S1# show dtp interface fa0/1
DTP information for FastEthernet0/1:
TOS/TAS/TNS: ACCESS/AUTO/ACCESS
TOT/TAT/TNT: NATIVE/NEGOTIATE/NATIVE
Neighbor address 1: C80084AEF101
Neighbor address 2: 000000000000
Hello timer expiration (sec/state): 11/RUNNING
Access timer expiration (sec/state): never/STOPPED
Negotiation timer expiration (sec/state): never/STOPPED
Multidrop timer expiration (sec/state): never/STOPPED
FSM state: S2:ACCESS
# times multi & trunk 0
```

```
Enabled: yes
In STP: no
```

Note - Best practice: Set the interface to **trunk** and **nonegotiate** when a trunk link is required

- On links where trunking is not intended, DTP should be turned off

Overview of VLANs

Virtual LANs (VLANs) are a group of devices that can communicate as if each device was attached to the same cable. VLANs are based on logical instead of physical connections. Administrators use VLANs to segment networks based on factors such as function, team, or application. Each VLAN is considered a separate logical network. Any switch port can belong to a VLAN. A VLAN creates a logical broadcast domain that can span multiple physical LAN segments. VLANs improve network performance by separating large broadcast domains into smaller ones. Each VLAN in a switched network corresponds to an IP network; therefore, VLAN design must use a hierarchical network-addressing scheme. Types of VLANs include the default VLAN, data VLANs, the native VLAN, management VLANs. and voice VLANs.

VLANs in a Multi-Switched Environment

A VLAN trunk does not belong to a specific VLAN. It is a conduit for multiple VLANs between switches and routers. A VLAN trunk is a point-to-point link between two network devices that carries more than one VLAN. A VLAN trunk extends VLANs across an entire network. When VLANs are implemented on a switch, the transmission of unicast, multicast, and broadcast traffic from a host in a particular VLAN are restricted to the devices that are in that VLAN. VLAN tag fields include the type, user priority, CFI and VID. Some devices add a VLAN tag to native VLAN traffic. If an 802.1Q trunk port receives a tagged frame with the VID that is the same as the native VLAN, it drops the frame. A separate voice VLAN is required to support VoIP. QoS and security policies can be applied to voice traffic. Voice VLAN traffic must be tagged with an appropriate Layer 2 CoS priority value.

VLAN Configuration

Different Cisco Catalyst switches support various numbers of VLANs including normal range VLANs and extended range VLANs. When configuring normal range VLANs, the configuration details are stored in flash memory on the switch in a file called vlan.dat. Although it is not required, it is good practice to save running configuration changes to the startup configuration. After creating a VLAN, the next step is to assign ports to the VLAN. There are several commands for defining a port to be an access port and assigning it to a VLAN. VLANs are configured on the switch port and not on the end device. An access port can belong to only one data VLAN at a time. However, a port can also be associated to a voice VLAN. For example, a port connected to an IP phone and an end device would be associated with two VLANs: one for voice and one for data. After a VLAN is configured, VLAN configurations can be validated using Cisco IOS show commands. If the switch access port has been incorrectly assigned to a VLAN, then simply re-enter the switchport access vlan *vlan-id* interface configuration command with the correct VLAN ID. The no vlan *vlan-id* global configuration mode command is used to remove a VLAN from the switch vlan.dat file.

VLAN Trunks

A VLAN trunk is an OSI Layer 2 link between two switches that carries traffic for all VLANs. There are several commands to configure the interconnecting ports. To verify VLAN trunk configuration use the show interfaces *interface-ID* switchport command. Use the no switchport trunk allowed vlan and the no switchport trunk native vlan commands to remove the allowed VLANs and reset the native VLAN of the trunk.

Dynamic Trunking Protocol

An interface can be set to trunking or nontrunking, or to negotiate trunking with the neighbor interface. Trunk negotiation is managed by the Dynamic Trunking Protocol (DTP), which operates on a point-to-point basis only, between network devices. DTP is a Cisco proprietary protocol that manages trunk negotiation only if the port on the neighbor switch is configured in a trunk mode that supports DTP. To enable trunking from a Cisco switch to a device that does not support DTP, use the switchport mode trunk and switchport nonegotiate interface configuration mode commands. The switchport mode command has additional options for negotiating the interface mode including access, dynamic auto, dynamic desirable, and trunk. To verify the current DTP mode, issue the show dtp interface command.

Inter-VLAN Routing

- Hosts in one VLAN cannot communicate with hosts in another VLAN unless there is a router or a layer 3 switch to provide routing services

Inter-VLAN routing is the process of forwarding network traffic from one VLAN to another VLAN

Three inter-VLAN routing options:

- **Legacy Inter-VLAN routing** - This is a legacy solution. It does not scale well
- **Router-on-a-stick** - This is an acceptable solution for a small to medium-sized network
- **Layer 3 switch using switched virtual interfaces (SVIs)** - This is the most scalable solution for medium to large organizations

Legacy Inter-VLAN Routing

- The first inter-VLAN routing solution relied on using a router with multiple interfaces
    - Each interface was connected to a switch port in different VLANs
    - Interfaces served as default gateways to the local hosts on the VLAN subnet

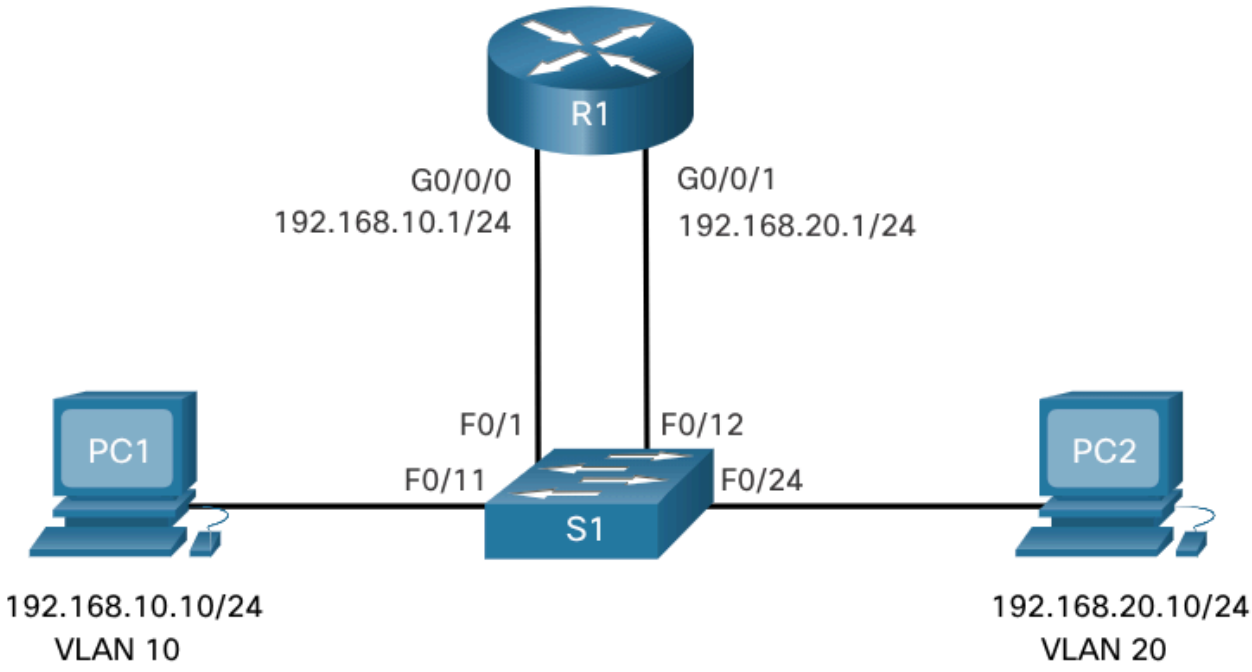

- Fa0/1 port is assigned to VLAN 10 and is connected to R1 G0/0/0 interface
- Fa0/11 port is assigned to VLAN 10 and is connected to PC1
- Fa0/12 port is assigned to VLAN 20 and is connected to the R1 G0/0/1 interface
- Fa0/24 port is assigned to VLAN 20 and is connected to PC2

MAC Address table for S1

Port	MAC Address	VLAN
F0/1	R1 G0/0/0 MAC	10
F0/11	PC1 MAC	10
F0/12	R1 G0/0/1 MAC	20
F0/24	PC2 MAC	20

When PC1 sends a packet to PC2 on another network →

1. Forwards it to its default gateway
    a. 192.168.10.1
2. R1 receives the packet on its G0/0/0 interface
    a. Examines the destination address of the packet
3. R1 routes the packet out its G0/0/1 interface to F0/12 port in VLAN 20 on S1
4. S1 forwards the frame to PC2

Legacy inter-VLAN routing using physical interfaces works, but has significant limitation

- Not reasonably scalable bc routers have a limited number of physical interfaces

Note - No longer implemented

Router-on-a-Stick Inter-VLAN Routing

- Only requires one physical interface to route traffic between multiple VLANs

Example setup:

Cisco IOS router interface is configured as 802.1Q trunk and connected to a trunk port on a Layer 2.

- The router interface is configured using subinterfaces to identify routable VLANs

The Configured subinterfaces are software-based virtual interfaces.

- Each is associated with a single physical interface

The subinterfaces are configured in software on a router

- Each subinterface is independently configured with an IP address and VLAN assignment

Subinterfaces are configured for different subnets that correspond to their VLAN assignment

This facilitates logical routing.

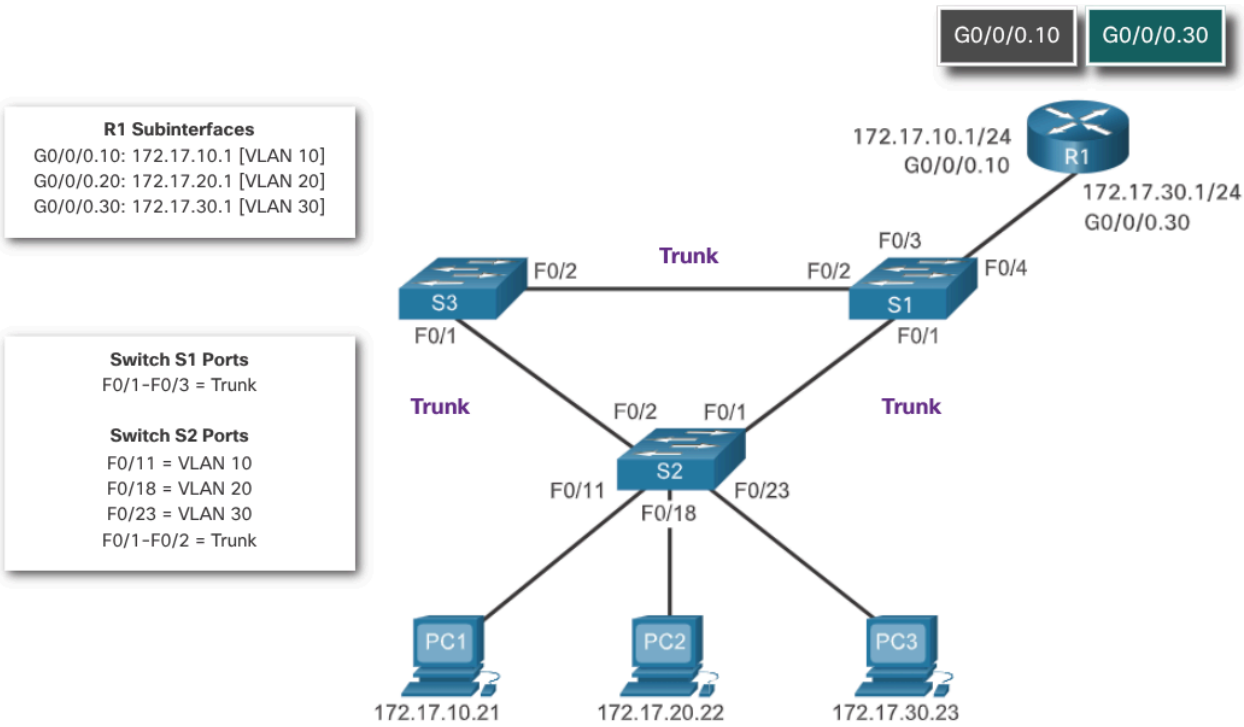

Note - This method does **not** scale beyond 50 VLANs

Inter-VLAN Routing on a Layer 3 Switch

Layer 3 switches and Switched Virtual Interfaces (SVI)

Switched Virtual Interfaces (SVI)

- Virtual interface that is configured on a Layer 3 switch

Note - Layer 3 switch is also called a multilayer switch as it operates at Layer 2 and Layer 3

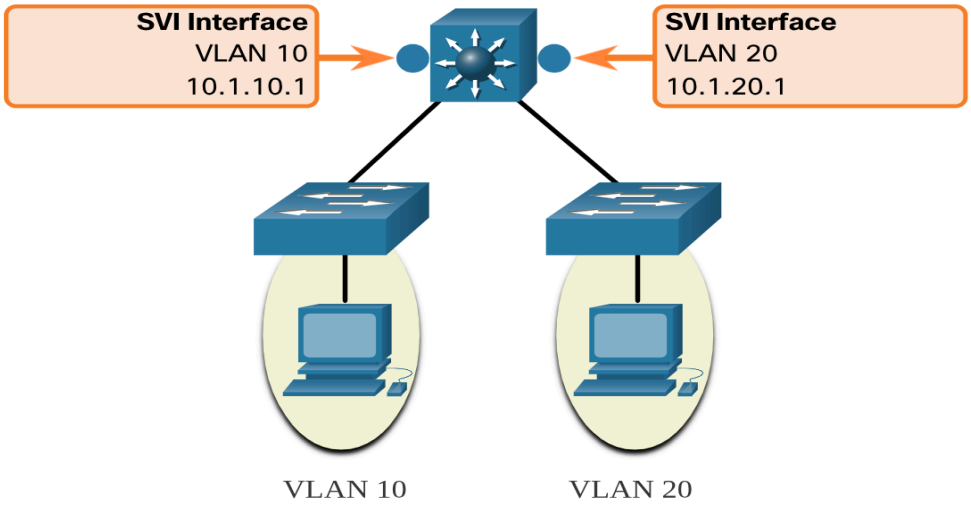

Inter-VLAN SVIs

- Created the same way that a management VLAN interface is configured
- SVIs is created for a VLAN that exists on the switch
- SVI performs the same functions for the VLAN as a router interface would
    - Specifically, provides Layer 3 processing for packets that are sent to or from all ports associated with that VLAN

Advantages of Layer 3 switches

- Much faster than router-on-a-stick because everything is hardware switched and routed
- No need for external links from the switch to the router for routing
- Not limited to one link because Layer 2 EtherChannels can be used as trunk links between the switches to increase bandwidth
- Latency is much lower because data does not need to leave the switch in order to be routed to a different network
- More commonly deployed in a campus LAN than routers

Disadvantages

- Expensive

Router on a stick scenario

Assume R1, S1, S2 have initial basic configuration.

PC1 and PC2 can't ping bc separate networks

Only S1 and S2 can ping, but unreachable by PC1 and PC2 bc separate networks

To enable devices to ping each other, the switches must be configured with VLANs and trunking, and router must be configured for inter-VLAN routing

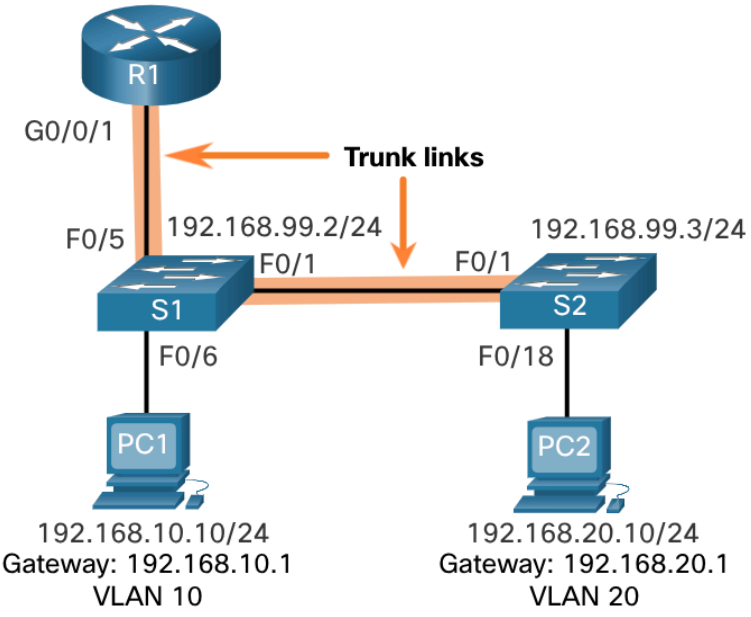

Subinterface	VLAN	IP Address
G0/0/1.10	10	192.168.10.1/24

G0/0/1.20	20	192.168.20.1/24
G0/0/1.99	99	192.168.99.1/24

S1 VLAN and Trunking Configuration

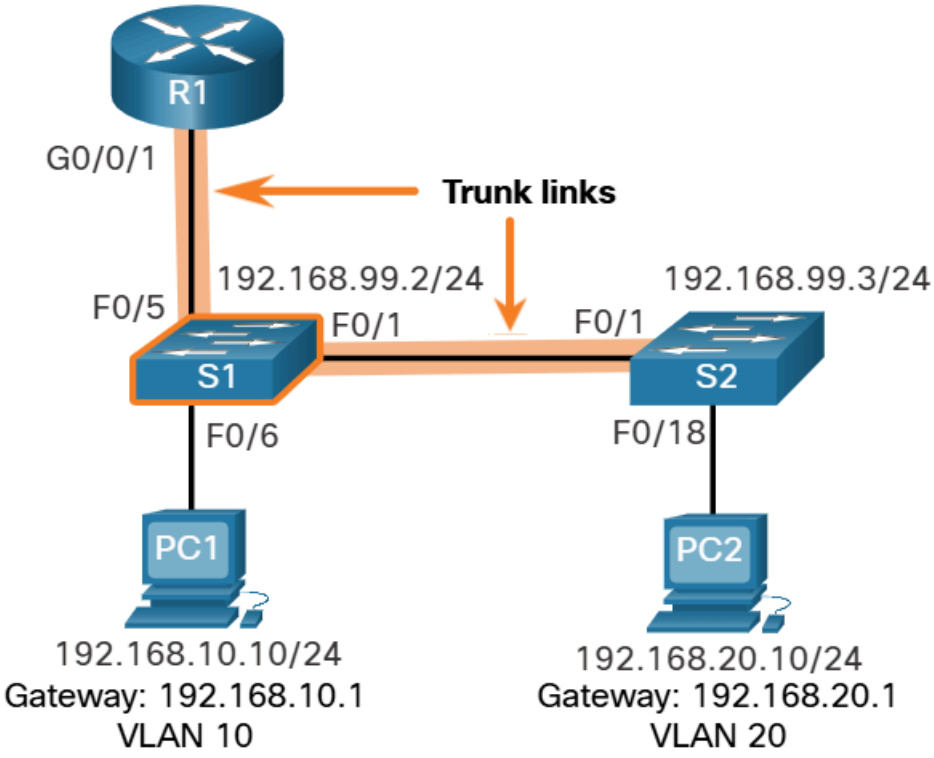

Step 1 -

Create and name the VLANs

- Only created after you exit out of VLAN subconfiguration mode

```
S1(config)# vlan 10
S1(config-vlan)# name LAN10
S1(config-vlan)# exit
S1(config)# vlan 20
S1(config-vlan)# name LAN20
S1(config-vlan)# exit
S1(config)# vlan 99
S1(config-vlan)# name Management
S1(config-vlan)# exit
S1(config)#
```

Step 2

Create the management interface

- Created on VLAN 99 along with the default gateway of R1

```
S1(config)# interface vlan 99
S1(config-if)# ip add 192.168.99.2 255.255.255.0
S1(config-if)# no shut
S1(config-if)# exit
S1(config)# ip default-gateway 192.168.99.1
S1(config)#
```

Step 3

Configure Access Ports

- Port F0/6 connecting to PC1 is configured as an access port in VLAN 10

```
S1(config)# interface fa0/6
S1(config-if)# switchport mode access
S1(config-if)# switchport access vlan 10
S1(config-if)# no shut
S1(config-if)# exit
S1(config)#
```

Step 4

Configure Trunking Ports

- F0/1 connecting to S2 and F0/5 connecting to R1 are configured as trunk ports

```
S1(config)# interface fa0/1
S1(config-if)# switchport mode trunk
S1(config-if)# no shut
S1(config-if)# exit
S1(config)# interface fa0/5
S1(config-if)# switchport mode trunk
S1(config-if)# no shut
S1(config-if)# end
*Mar  1 00:23:43.093: %LINEPROTO-5-UPDOWN: Line protocol on
Interface FastEthernet0/1, changed state to up
```

```
*Mar  1 00:23:44.511: %LINEPROTO-5-UPDOWN: Line protocol on
Interface FastEthernet0/5, changed state to up
```

S2 VLAN and Trunking Configuration

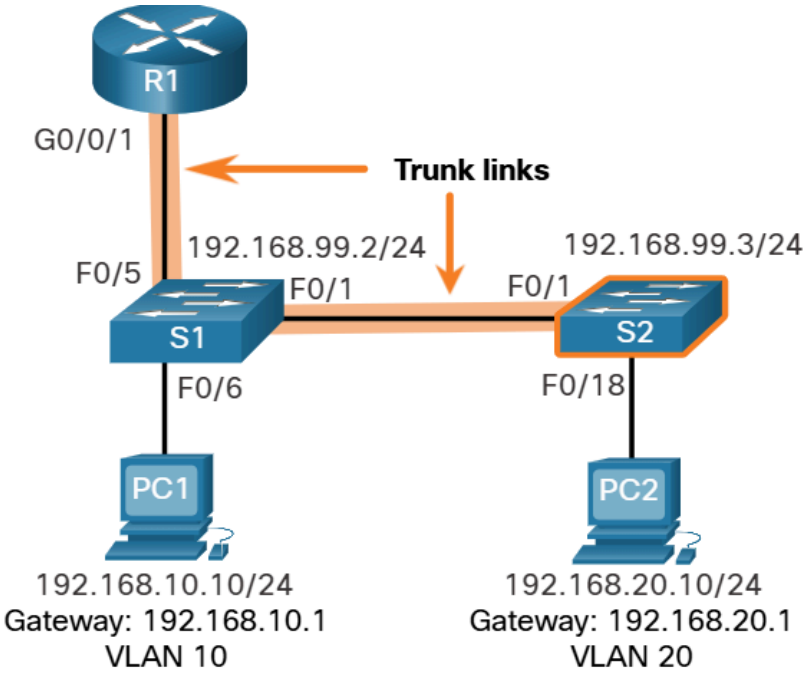


```
S2(config)# vlan 10
S2(config-vlan)# name LAN10
S2(config-vlan)# exit
S2(config)# vlan 20
S2(config-vlan)# name LAN20
S2(config-vlan)# exit
S2(config)# vlan 99
S2(config-vlan)# name Management
S2(config-vlan)# exit
S2(config)#
S2(config)# interface vlan 99
S2(config-if)# ip add 192.168.99.3 255.255.255.0
S2(config-if)# no shut
S2(config-if)# exit
S2(config)# ip default-gateway 192.168.99.1
S2(config)# interface fa0/18
S2(config-if)# switchport mode access
S2(config-if)# switchport access vlan 20
S2(config-if)# no shut
```

```
S2(config-if)# exit
S2(config)# interface fa0/1
S2(config-if)# switchport mode trunk
S2(config-if)# no shut
S2(config-if)# exit
S2(config-if)# end
*Mar  1 00:23:52.137: %LINEPROTO-5-UPDOWN: Line protocol on Interface FastEthernet0/1,
changed state to up
```

R1 Subinterface Configuration

- Router on a stick requires you to create a subinterface for each VLAN to be routed

Subinterface

- Created using **interface** *interface_id.subinterface_id* command
    - Syntax: Physical interface followed by a period and subinterface number
        - Not required, but customary to match subinterface number with VLAN number

Subinterface commands

- **encapsulation dot1q** *vlan_id* **[native]**
    - Configures the subinterface to respond to 802.1Q encapsulated traffic from the specified *vlan_id*
    - **native** keyword option is only appended to set the native VLAN to something other than VLAN 1
- **ip address** *ip-address subnet-mask*
    - Configures the IPv4 address of subinterface
    - Typically serves as default gateway for the identified VLAN

Repeat process for each VLAN to be routed

- Each router subinterface must be assigned an IP address on a unique subnet for routing to occur

When all subinterfaces have been created, enable the physical interface using the **no shutdown**.

- If physical interface is disabled, all subinterfaces are disabled

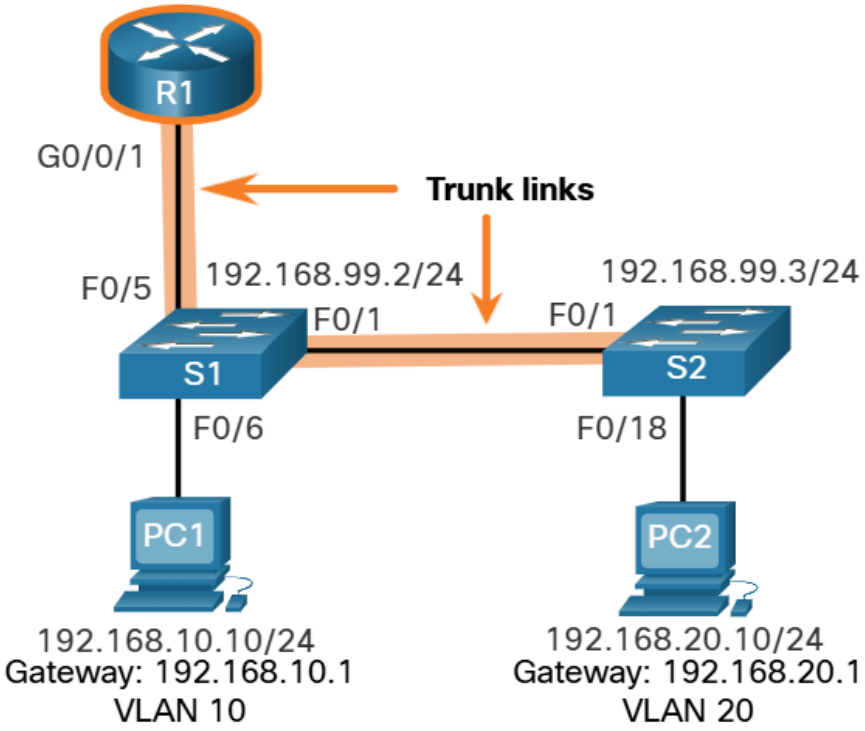


```
R1(config)# interface G0/0/1.10
R1(config-subif)# description Default Gateway for VLAN 10
R1(config-subif)# encapsulation dot1Q 10
R1(config-subif)# ip add 192.168.10.1 255.255.255.0
R1(config-subif)# exit
R1(config)#
R1(config)# interface G0/0/1.20
R1(config-subif)# description Default Gateway for VLAN 20
R1(config-subif)# encapsulation dot1Q 20
R1(config-subif)# ip add 192.168.20.1 255.255.255.0
R1(config-subif)# exit
R1(config)#
R1(config)# interface G0/0/1.99
R1(config-subif)# description Default Gateway for VLAN 99
R1(config-subif)# encapsulation dot1Q 99
R1(config-subif)# ip add 192.168.99.1 255.255.255.0
R1(config-subif)# exit
R1(config)#
R1(config)# interface G0/0/1
R1(config-if)# description Trunk link to S1
R1(config-if)# no shut
```

```
R1(config-if)# end
R1#
*Sep 15 19:08:47.015: %LINK-3-UPDOWN: Interface GigabitEthernet0/0/1,
changed state to down
*Sep 15 19:08:50.071: %LINK-3-UPDOWN: Interface GigabitEthernet0/0/1,
changed state to up
*Sep 15 19:08:51.071: %LINEPROTO-5-UPDOWN: Line protocol on Interface
GigabitEthernet0/0/1, changed state to up
R1#
```

Verify Connectivity Between PC1 and PC2

Router on a stick is complete after the switch trunk and router subinterfaces have been configured

ping command

- Verify connectivity

ipconfig command

- Windows host ip configuration info

Router on a Stick Inter-VLAN Routing Verification

show command

- Used to verify and troubleshoot the router on a stick configuration
    - **show ip route**
    - **show ip interface brief**
    - **show interfaces**
    - **show interfaces trunk**

show ip route

```
R1# show ip route | begin Gateway
Gateway of last resort is not set
      192.168.10.0/24 is variably subnetted, 2 subnets, 2 masks
C        192.168.10.0/24 is directly connected,
GigabitEthernet0/0/1.10
```

```
L        192.168.10.1/32 is directly connected,
GigabitEthernet0/0/1.10
      192.168.20.0/24 is variably subnetted, 2 subnets, 2 masks
C        192.168.20.0/24 is directly connected,
GigabitEthernet0/0/1.20
L        192.168.20.1/32 is directly connected,
GigabitEthernet0/0/1.20
      192.168.99.0/24 is variably subnetted, 2 subnets, 2 masks
C        192.168.99.0/24 is directly connected,
GigabitEthernet0/0/1.99
L        192.168.99.1/32 is directly connected,
GigabitEthernet0/0/1.99
R1#
```

show ip interface brief

```
R1# show ip interface brief | include up
GigabitEthernet0/0/1   unassigned      YES unset  up                    up
Gi0/0/1.10             192.168.10.1    YES manual up                    up
Gi0/0/1.20             192.168.20.1    YES manual up                    up
Gi0/0/1.99             192.168.99.1    YES manual up                    up
R1#
```

show interfaces

```
R1# show interfaces g0/0/1.10
GigabitEthernet0/0/1.10 is up, line protocol is up
  Hardware is ISR4221-2x1GE, address is 10b3.d605.0301 (bia 10b3.d605.0301)
  Description: Default Gateway for VLAN 10
  Internet address is 192.168.10.1/24
  MTU 1500 bytes, BW 100000 Kbit/sec, DLY 100 usec,
     reliability 255/255, txload 1/255, rxload 1/255
  Encapsulation 802.1Q Virtual LAN, Vlan ID  10.
  ARP type: ARPA, ARP Timeout 04:00:00
  Keepalive not supported
  Last clearing of "show interface" counters never
R1#
```

show interfaces trunk

```
S1# show interfaces trunk
Port        Mode             Encapsulation  Status
Native vlan
Fa0/1       on               802.1q         trunking      1
Fa0/5       on               802.1q         trunking      1
Port        Vlans allowed on trunk
Fa0/1       1-4094
Fa0/5       1-4094
Port        Vlans allowed and active in management domain
Fa0/1       1,10,20,99
Fa0/5       1,10,20,99
Port        Vlans in spanning tree forwarding state and not
pruned
Fa0/1       1,10,20,99
Fa0/5       1,10,20,99
S1#
```

Layer 3 Switch Inter-VLAN Routing

Router on a Stick Inter-VLAN routing is simple to implement for a **small** to **medium-sized** organization

- Does not scale easily

Enterprise campus LANs use Layer 3 switches

- Layer 3 switches use hardware-based switching to achieve higher-packet processing rates
- Commonly implemented in enterprise distribution layer writing closets

Capabilities:

- Route from one VLAN to another using multiple switched virtual interfaces (SVIs)
- Convert a Layer 2 switchport to a Layer 3 interface (routed port)
    - A routed port is similar to physical interface on a Cisco IOS router

Layer 3 switches use SVIs

- SVIs are configured using the same **interface vlan** *vlan-id* command to create the management SVI on a Layer 2 switch
- Layer 3 SVI must be created for each of the routable VLANs

Layer 3 Switch Scenario

D1 is connected to two hosts on different VLANs

- PC1 is in VLAN 10
- PC2 is in VLAN 20

D1 VLAN IP Addresses

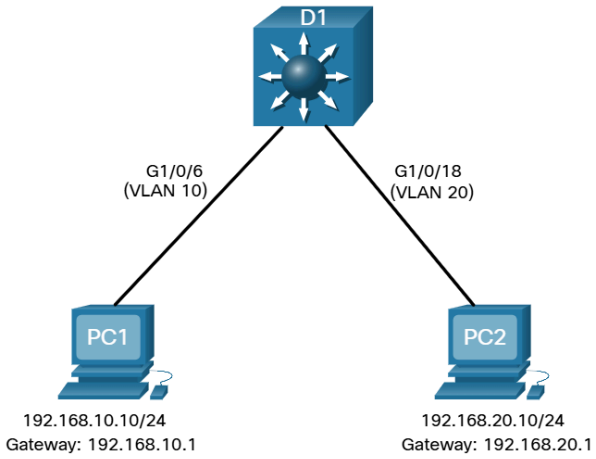

VLAN Interface	IP Address
10	192.168.10.1/24
20	192.168.20.1/24

Layer 3 Switch Configuration

Step 1 -

Create the VLANs

```
D1(config)# vlan 10
D1(config-vlan)# name LAN10
D1(config-vlan)# vlan 20
D1(config-vlan)# name LAN20
D1(config-vlan)# exit
D1(config)#
```

Step 2 -

Create the SVI VLAN interfaces

- Configure the SVI for VLANs 10 and 20
- The IP addresses that are configured will serve as the default gateways to the hosts in the respective VLANs

```
D1(config)# interface vlan 10
D1(config-if)# description Default Gateway SVI for 192.168.10.0/24
D1(config-if)# ip add 192.168.10.1 255.255.255.0
D1(config-if)# no shut
D1(config-if)# exit
D1(config)#
D1(config)# int vlan 20
D1(config-if)# description Default Gateway SVI for 192.168.20.0/24
D1(config-if)# ip add 192.168.20.1 255.255.255.0
D1(config-if)# no shut
D1(config-if)# exit
D1(config)#
*Sep 17 13:52:16.053: %LINEPROTO-5-UPDOWN: Line protocol on
Interface Vlan10, changed state to up
*Sep 17 13:52:16.160: %LINEPROTO-5-UPDOWN: Line protocol on
Interface Vlan20, changed state to up
```

Step 3 -

Configure Access Ports

- Configure the access ports connecting to the hosts and assign them to their respective VLANs

```
D1(config)# interface GigabitEthernet1/0/6
D1(config-if)# description Access port to PC1
D1(config-if)# switchport mode access
D1(config-if)# switchport access vlan 10
D1(config-if)# exit
D1(config)#
D1(config)# interface GigabitEthernet1/0/18
D1(config-if)# description Access port to PC2
D1(config-if)# switchport mode access
D1(config-if)# switchport access vlan 20
D1(config-if)# exit
```

Step 4 -

Enable IP Routing

- Enable IPv4 routing with **ip routing** command to allow traffic to be exchanged between VLANs 10 and 20
- The command must be configured to enable inter-VLAN routing on a Layer 3 switch for IPv4

```
D1(config)# ip routing
D1(config)#
```

Layer 3 Switch Inter-VLAN Routing Verification

- Simpler to configure than Router on a Stick

Verify Connectivity Between PC1 and PC2

ping command

- Verify connectivity

ipconfig command

- Windows host ip configuration info

Routing on a Layer 3 Switch

If VLANs are to be reachable by other Layer 3 devices, then they must be **advertised** using **static** or **dynamic routing**

- To enable routing on Layer 3 switch, a routed port must be configured

A routed port is created on a Layer 3 switch by **disabling** the switchport feature on a Layer 2 port that is connected to another Layer 3 device

- * Configuring the **no switchport** on a Layer 2 port converts it into a Layer 3 interface
- Then the interface can be configured with an IPv4 configuration to connect a router or another Layer 3 switch

Routing Scenario on a Layer 3 Switch

- D1 Layer 3 switch is now connected to R1
  - R1 and D1 are both in an **Open Shortest Path First (OSPF)** routing protocol domain
  - The G0/0/1 interface of R1 has also been configured and enabled
  - R1 is using OSPF to advertise its two networks
    - 10.10.10.0/24
    - 10.20.20.0/24

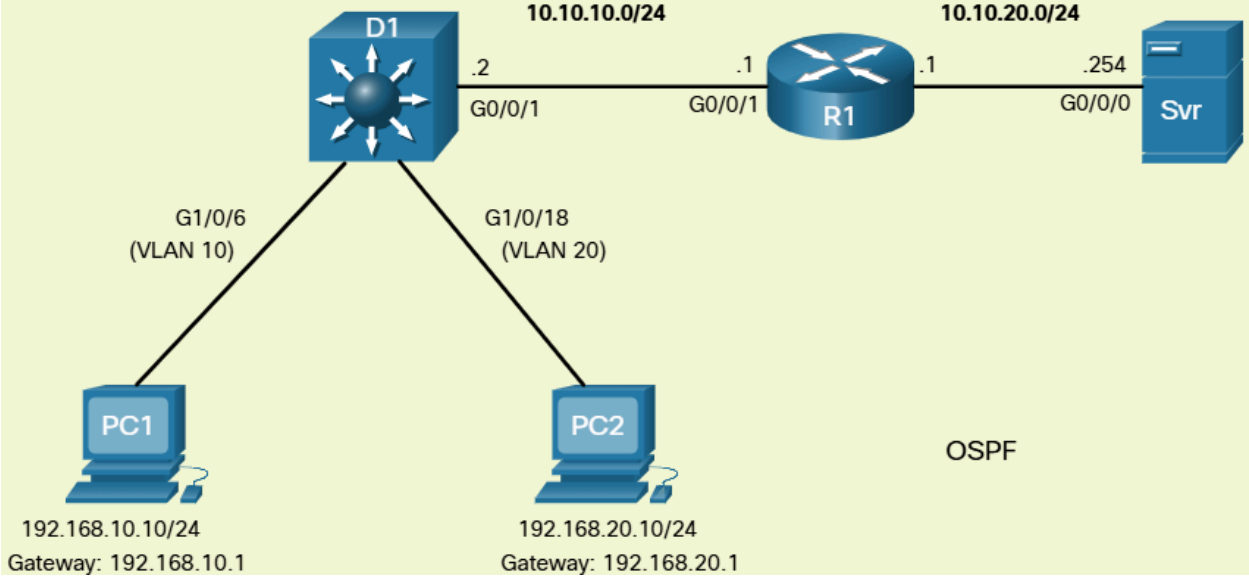

Routing Configuration on a Layer 3 Switch

Step 1 -

Configure the routed port

- Configure the G0/0/1to be a routed port
- Assign it an IPv4
- Enable it

```
D1(config)# interface GigabitEthernet0/0/1
D1(config-if)# description routed Port Link to R1
D1(config-if)# no switchport
D1(config-if)# ip address 10.10.10.2 255.255.255.0
D1(config-if)# no shut
D1(config-if)# exit
D1(config)#
```

Step 2 -

Enable routing

- Ensure IPv4 routing is enabled with **ip routing** command

```
D1(config)# ip routing
D1(config)#
```

Step 3 -

Configure routing

- Configure the OSPF routing protocol to advertise the VLAN 10 and VLAN 20 networks, along with the network that is connected to R1
    - Notice the message informing that an adjacency has been established with R1

```
D1(config)# router ospf 10
D1(config-router)# network 192.168.10.0 0.0.0.255 area 0
D1(config-router)# network 192.168.20.0 0.0.0.255 area 0
D1(config-router)# network 10.10.10.0 0.0.0.3 area 0
D1(config-router)# ^Z
D1#
*Sep 17 13:52:51.163: %OSPF-5-ADJCHG: Process 10, Nbr
```

```
10.20.20.1 on GigabitEthernet0/0/1 from LOADING to FULL,
Loading Done
D1#
```

Step 4 -

Verify routing

- Verify the routing table on D1
    - Notice D1 now has a route to 10.20.20.0/24 network

```
D1# show ip route | begin Gateway
Gateway of last resort is not set
      10.0.0.0/8 is variably subnetted, 3 subnets, 3 masks
C        10.10.10.0/30 is directly connected, GigabitEthernet0/0/1
L        10.10.10.2/32 is directly connected, GigabitEthernet0/0/1
O        10.10.20.0/24 [110/2] via 10.10.10.1, 00:00:06,
GigabitEthernet0/0/1
      192.168.10.0/24 is variably subnetted, 2 subnets, 2 masks
C        192.168.10.0/24 is directly connected, Vlan10
L        192.168.10.1/32 is directly connected, Vlan10
      192.168.20.0/24 is variably subnetted, 2 subnets, 2 masks
C        192.168.20.0/24 is directly connected, Vlan20
L        192.168.20.1/32 is directly connected, Vlan20
D1#
```

Step 5 -

Verify connectivity

- At this time, PC1 and PC2 are able to ping the server connected to R1

```
C:\Users\PC1> ping 10.20.20.254
Pinging 10.20.20.254 with 32 bytes of data:
Request timed out.
Reply from 10.20.20.254: bytes=32 time<1ms TTL=127
Reply from 10.20.20.254: bytes=32 time<1ms TTL=127
Reply from 10.20.20.254: bytes=32 time<1ms TTL=127
Ping statistics for 10.20.20.254:
    Packets: Sent = 4, Received = 3, Lost = 1 (25% loss).
Approximate round trip times in milli-seconds:
    Minimum = 1ms, Maximum = 2ms, Average = 1ms
```

```
C:\Users\PC1>
!================================================
C:\Users\PC2> ping 10.20.20.254
Pinging 10.20.20.254 with 32 bytes of data:
Reply from 10.20.20.254: bytes=32 time<1ms TTL=127
Reply from 10.20.20.254: bytes=32 time<1ms TTL=127
Reply from 10.20.20.254: bytes=32 time<1ms TTL=127
Reply from 10.20.20.254: bytes=32 time<1ms TTL=127
Ping statistics for 10.20.20.254:
    Packets: Sent = 4, Received = 4, Lost = 0 (0% loss).
Approximate round trip times in milli-seconds:
    Minimum = 1ms, Maximum = 2ms, Average = 1ms
C:\Users\PC2>
```

Common Inter-VLAN Issues

Number of reasons why an inter-VLAN configuration may not work

- All related to connectivity issues
- First, check physical layer

Issue Type	How to Fix	How to Verify

Missing VLANs	● Create (or re-create) the VLAN if it does not exist
	● Ensure host port is assigned to the correct VLAN
		`show vlan [brief]`
		`show interfaces switchport`
		`ping`

Switch Trunk Port Issues	● Ensure trunks are configured correctly ● Ensure port is a trunk port and enabled	`show interfaces trunk` `show running-config`

Switch Access Port Issues	● Assign correct VLAN to access port
	● Ensure port is an access port and enabled
	● Host is incorrectly configured in the wrong subnet
		`show interfaces switchport`
		`show running-config interface`
		`ipconfig`

Router Configuration Issues	● Router subinterference IPv4 address is incorrectly configured ● Router subinterference is assigned to the VLAN ID	`show ip interface brief` `show interfaces`

Troubleshoot Inter-VLAN Routing Scenario

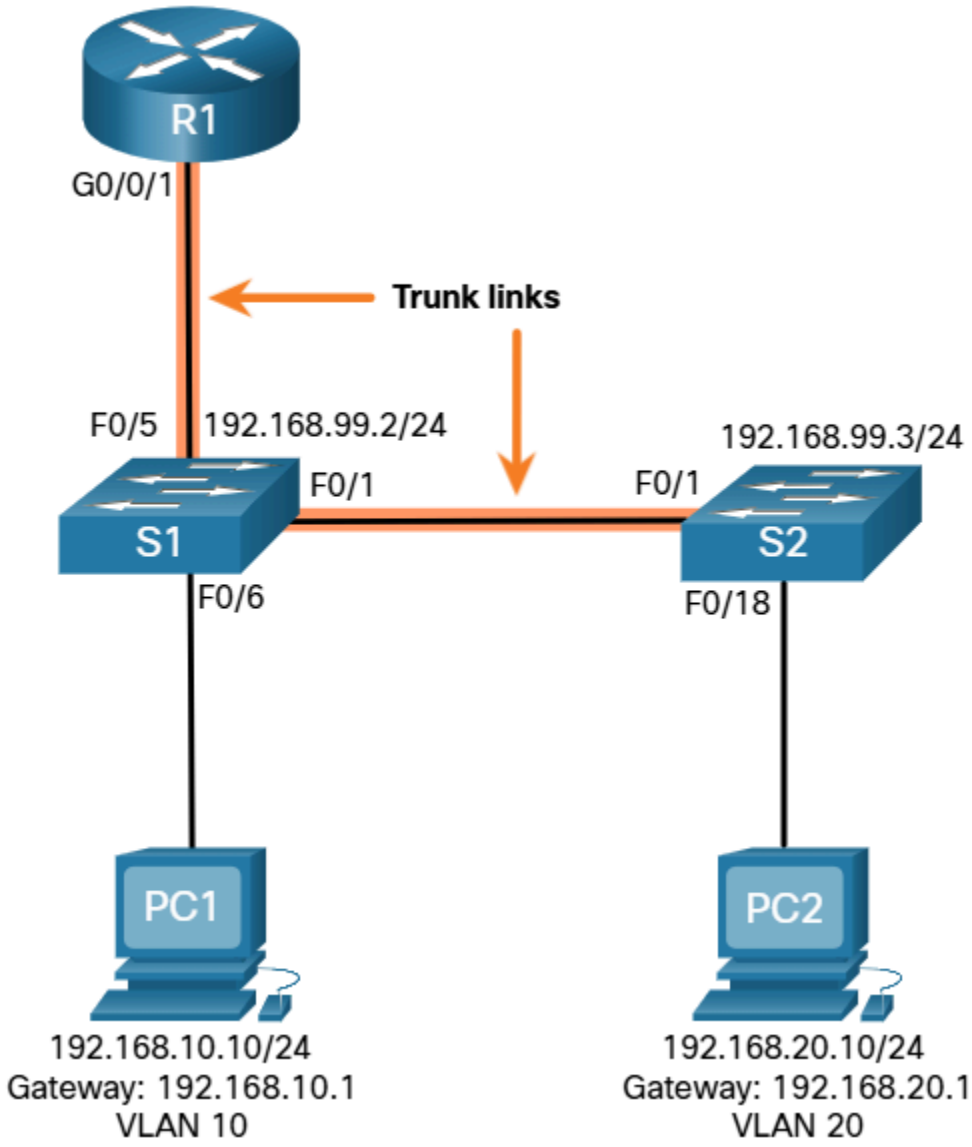

Router R1 Subinterference

Subinterference	VLAN	IP Address
G0/0/0.10	10	192.168.10.1/24
G0/0/0.20	20	192.168.20.1/24
G0/0/0.30	99	192.168.99.1/24

Missing VLANs

- Causes inter-VLAN connectivity issues
  - Could be missing
  - Not created
  - Accidentally deleted
  - Not allowed on trunk link

```
S1# show vlan brief
VLAN Name                              Status    Ports
---- --------------------------------- --------- -------------------------------
1    default                           active    Fa0/2, Fa0/3, Fa0/4, Fa0/7
                                                 Fa0/8, Fa0/9, Fa0/10, Fa0/11
                                                 Fa0/12, Fa0/13, Fa0/14, Fa0/15
                                                 Fa0/16, Fa0/17, Fa0/18, Fa0/19
                                                 Fa0/20, Fa0/21, Fa0/22, Fa0/23
                                                 Fa0/24, Gi0/1, Gi0/2
10   LAN10                             active    Fa0/6
20   LAN20                             active
99   Management                          active
1002 fddi-default                        act/unsup
1003 token-ring-default                  act/unsup
1004 fddinet-default                     act/unsup
1005 trnet-default                       act/unsup
S1#
```

Assume VLAN 10 is accidentally deleted

```
S1(config)# no vlan 10
S1(config)# do show vlan brief
VLAN Name                              Status    Ports
---- --------------------------------- --------- -------------------------------
1    default                           active    Fa0/2, Fa0/3, Fa0/4, Fa0/7
                                                 Fa0/8, Fa0/9, Fa0/10, Fa0/11
                                                 Fa0/12, Fa0/13, Fa0/14, Fa0/15
                                                 Fa0/16, Fa0/17, Fa0/18, Fa0/19
                                                 Fa0/20, Fa0/21, Fa0/22, Fa0/23
                                                 Fa0/24, Gi0/1, Gi0/2
20   LAN20                             active
99   Management                          active
1002 fddi-default                        act/unsup
1003 token-ring-default                  act/unsup
1004 fddinet-default                     act/unsup
1005 trnet-default                       act/unsup
S1(config)#
```

Now that VLAN 10 is missing, Port Fa0/6 was not reassigned to the default VLAN

- When you delete a VLAN, any ports assigned to that VLAN becomes inactive
- Port remains associated with VLAN until you assign them to a new VLAN or recreate the missing VLAN

show interface *interface-id* **switchport** command to verify VLAN membership

```
S1(config)# do show interface fa0/6 switchport
Name: Fa0/6
Switchport: Enabled
Administrative Mode: static access
Operational Mode: static access
Administrative Trunking Encapsulation: dot1q
Operational Trunking Encapsulation: native
Negotiation of Trunking: Off
Access Mode VLAN: 10 (Inactive)
Trunking Native Mode VLAN: 1 (default)
Administrative Native VLAN tagging: enabled
Voice VLAN: none
(Output omitted)
```

Recreating the VLAN would automatically reassign the hosts to it

```
S1(config)# vlan 10
S1(config-vlan)# do show vlan brief
VLAN Name                              Status    Ports
---- --------------------------------- --------- -------------------------------
1    default                           active    Fa0/2, Fa0/3, Fa0/4, Fa0/7
                                                 Fa0/8, Fa0/9, Fa0/10, Fa0/11
                                                 Fa0/12, Fa0/13, Fa0/14, Fa0/15
                                                 Fa0/16, Fa0/17, Fa0/18, Fa0/19
                                                 Fa0/20, Fa0/21, Fa0/22, Fa0/23
                                                 Fa0/24, Gi0/1, Gi0/2
20   LAN20                             active
99   Management                          active
1002 fddi-default                        act/unsup
1003 token-ring-default                  act/unsup
1004 fddinet-default                     act/unsup
1005 trnet-default                       act/unsup
S1(config-vlan)#
```

VLAN was not created as expected

- * You must exit from VLAN sub-configuration mode to create the VLAN

```
S1(config-vlan)# exit
S1(config)# vlan 10
S1(config)# do show vlan brief
VLAN Name                              Status    Ports
---- --------------------------------- --------- -------------------------------
1    default                           active    Fa0/2, Fa0/3, Fa0/4, Fa0/7
                                                 Fa0/8, Fa0/9, Fa0/10, Fa0/11
                                                 Fa0/12, Fa0/13, Fa0/14, Fa0/15
                                                 Fa0/16, Fa0/17, Fa0/18, Fa0/19
                                                 Fa0/20, Fa0/21, Fa0/22, Fa0/23
                                                 Fa0/24, Gi0/1, Gi0/2
10   VLAN0010                            active    Fa0/6
20   LAN20                             active
99   Management                          active
1002 fddi-default                        act/unsup
1003 token-ring-default                  act/unsup
1004 fddinet-default                     act/unsup
1005 trnet-default                       act/unsup
S1(config)#
```

VLAN 10 is now included in list and host connected to Fa0/6 is on VLAN 10

Switch Trunk Port Issues

Another issue for inter-VLAN routing includes misconfigured switch ports

- In legacy inter-VLAN, this could be caused when the connecting router port is not assigned to correct VLAN

Most common cause with stick-on-a-router

- Misconfigured trunk port

Scenario:

PC1 was able to connect to hosts in other VLANs until recently.

Maintenance logs revealed that S1 Layer 2 switch was recently accessed for routine maintenance.

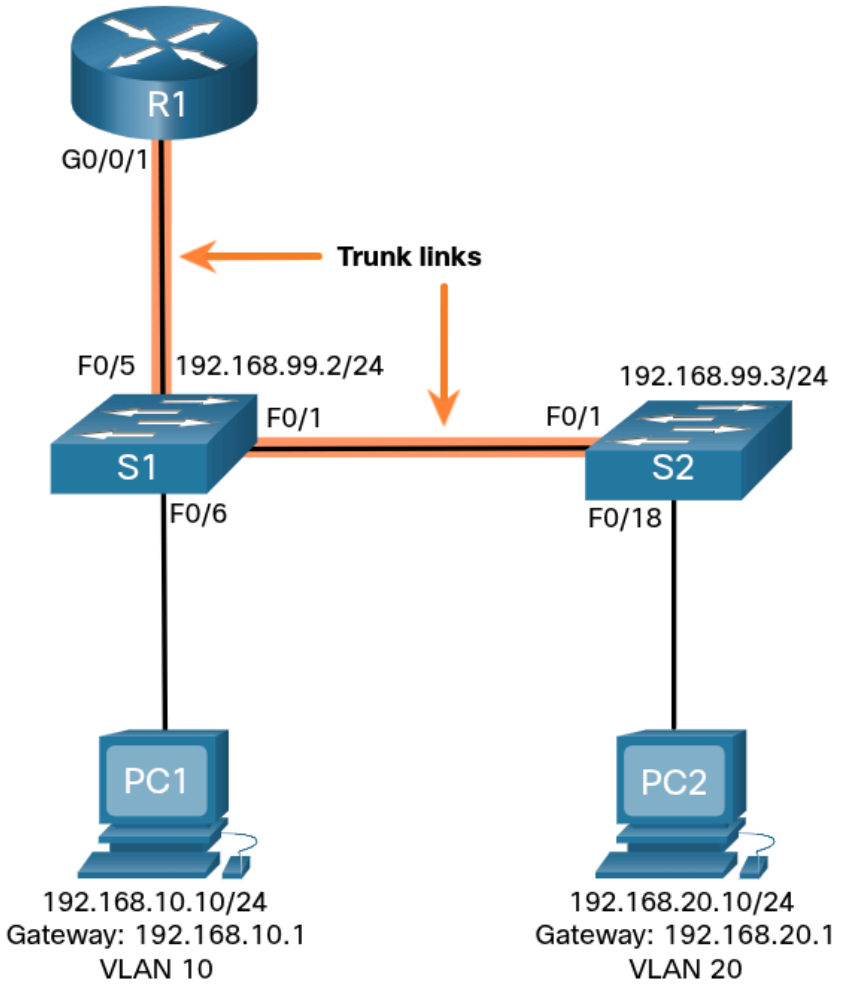

Verify the port connecting to R1 (F0/5) is correctly configured as trunk link

- **show interfaces trunk** command

```
S1# show interfaces trunk
Port        Mode           Encapsulation  Status        Native vlan
Fa0/1       on             802.1q         trunking      1
Port        Vlans allowed on trunk
Fa0/1       1-4094
Port        Vlans allowed and active in management domain
Fa0/1       1,10,20,99
Port        Vlans in spanning tree forwarding state and not pruned
Fa0/1       1,10,20,99
S1#
```

The Fa0/5 port connecting to R1 is missing. Verify interface configuration

- **show running-config interface fa0/5**

```
S1# show running-config | include interface fa0/5
Building configuration...
Current configuration : 96 bytes
!
interface FastEthernet0/5
description Trunk link to R1
switchport mode trunk
shutdown
end
S1#
```

Port was accidentally shut down. Re-enable port and verify trunking status

```
S1(config)# interface fa0/5
S1(config-if)# no shut
S1(config-if)#
*Mar  1 04:46:44.153: %LINK-3-UPDOWN: Interface FastEthernet0/5, changed state to up
S1(config-if)#
*Mar  1 04:46:47.962: %LINEPROTO-5-UPDOWN: Line protocol on Interface FastEthernet0/5,
changed state to up
S1(config-if)# do show interface trunk
Port        Mode           Encapsulation  Status        Native vlan
Fa0/1         on           802.1q       trunking      1
Fa0/5         on           802.1q       trunking      1
Port        Vlans allowed on trunk
Fa0/1         1-4094
Fa0/5         1-4094
Port        Vlans allowed and active in management domain
Fa0/1         1,10,20,99
Fa0/5         1,10,20,99
Port        Vlans in spanning tree forwarding state and not pruned
Fa0/1         1,10,20,99
Fa0/1         1,10,20,99
S1(config-if)#
```

Note - To reduce risk of failed inter-switch link disrupting inter-VLAN routing, redundant links and alternate paths should be part of network design.

Switch Access Port Issues

When problem is suspected with switch access port configuration, use verification commands to examine configuration and identify problem

Assume PC1 has correct IPv4 address and default gateway, but is not able to **ping** its own default gateway

- PC1 is supposed to be connected to a VLAN 10 port

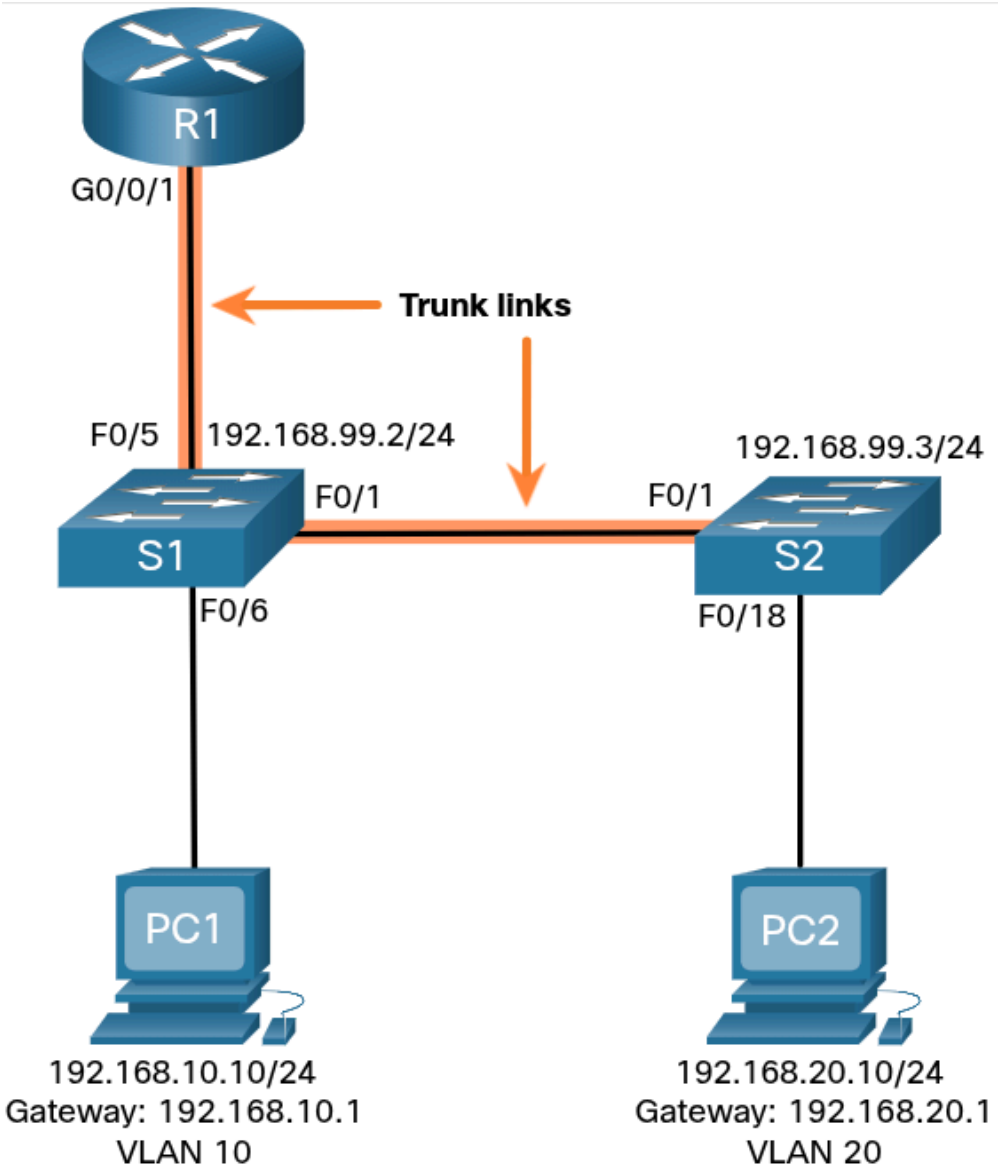

Verify port configuration on S1

- **show interfaces** *interface-id* **switchport** command

```
S1# show interface fa0/6 switchport
Name: Fa0/6
Switchport: Enabled
Administrative Mode: static access
Operational Mode: static access
Administrative Trunking Encapsulation: dot1q
Operational Trunking Encapsulation: native
Negotiation of Trunking: Off
Access Mode VLAN: 1 (default)
Trunking Native Mode VLAN: 1 (default)
Administrative Native VLAN tagging: enabled
Voice VLAN: none
```

Fa0/6 port has been configured as an access port as indicated by “static access”

- Not been configured to be in VLAN 10

```
S1# show running-config interface fa0/6
Building configuration...
Current configuration : 87 bytes
!
interface FastEthernet0/6
description PC-A access port
switchport mode access
end
S1#
```

Assign port Fa0/6 to VLAN 10 and verify port assignment

```
S1# configure terminal
S1(config)# interface fa0/6
S1(config-if)# switchport access vlan 10
S1(config-if)#
S1(config-if)# do show interface fa0/6 switchport
Name: Fa0/6
Switchport: Enabled
Administrative Mode: static access
Operational Mode: static access
Administrative Trunking Encapsulation: dot1q
Operational Trunking Encapsulation: native
```

```
Negotiation of Trunking: Off
Access Mode VLAN: 10 (VLAN0010)
Trunking Native Mode VLAN: 1 (default)
Administrative Native VLAN tagging: enabled
Voice VLAN: none
(Output omitted)
```

Router Configuration Issues

Router on a stick configuration problems are usually related to subinterface misconfigurations

- Example: an incorrect IP address was configured or wrong VLAN ID was assigned to the subinterface

Switch trunk link has been verified

- Then verify the subinterface status using **show ip interface brief** command

```
R1# show ip interface brief
Interface              IP-Address      OK? Method Status                Protocol
GigabitEthernet0/0/0   unassigned      YES unset  administratively down down
GigabitEthernet0/0/1   unassigned      YES unset  up                    up
Gi0/0/1.10          192.168.10.1    YES manual up                    up
Gi0/0/1.20          192.168.20.1    YES manual up                    up
Gi0/0/1.99          192.168.99.1    YES manual up                    up
Serial0/1/0            unassigned      YES unset  administratively down down
Serial0/1/1            unassigned      YES unset  administratively down down
R1#
```

Subinterfaces have been assigned the correct IPv4 addresses and are operational

Verify which VLANs each of the subinterfaces are on

- **show interfaces** command
  - Generates great deal of additional unrequired outputs
  - Can be reduced using IOS command filters

```
R1# show interfaces | include Gig|802.1Q
GigabitEthernet0/0/0 is administratively down, line protocol is down
GigabitEthernet0/0/1 is up, line protocol is up
  Encapsulation 802.1Q Virtual LAN, Vlan ID  1., loopback not set
GigabitEthernet0/0/1.10 is up, line protocol is up
  Encapsulation 802.1Q Virtual LAN, Vlan ID  100.
GigabitEthernet0/0/1.20 is up, line protocol is up
  Encapsulation 802.1Q Virtual LAN, Vlan ID  20.
GigabitEthernet0/0/1.99 is up, line protocol is up
```

```
  Encapsulation 802.1Q Virtual LAN, Vlan ID  99.
R1#
```

Pipe symbol (|) + select keywords help filter command output

Example:

- **include** was used to identify only lines containing letters “Gig” or 802.1Q” will be displayed

Notice G0/0/1.10 interface has been incorrectly assigned to VLAN 100 instead of VLAN 10

```
R1# show running-config interface g0/0/1.10
Building configuration...
Current configuration : 146 bytes
!
interface GigabitEthernet0/0/1.10
description Default Gateway for VLAN 10
encapsulation dot1Q 100
ip address 192.168.10.1 255.255.255.0
end
R1#
```

Configure subinterface G0/0/1.10 to be on correct VLAN

- **encapsulation dot1q 10** command

```
R1# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
R1(config)# interface gigabitEthernet 0/0/1.10
R1(config-subif)# encapsulation dot1Q 10
R1(config-subif)# end
R1#
R1# show interfaces | include Gig|802.1Q
GigabitEthernet0/0/0 is administratively down, line protocol is down
GigabitEthernet0/0/1 is up, line protocol is up
  Encapsulation 802.1Q Virtual LAN, Vlan ID  1., loopback not set
GigabitEthernet0/0/1.10 is up, line protocol is up
  Encapsulation 802.1Q Virtual LAN, Vlan ID  10.
GigabitEthernet0/0/1.20 is up, line protocol is up
  Encapsulation 802.1Q Virtual LAN, Vlan ID  20.
GigabitEthernet0/0/1.99 is up, line protocol is up
R1#
```

4.5.3 What did I learn in this module?

Inter-VLAN Routing Operation

Hosts in one VLAN cannot communicate with hosts in another VLAN unless there is a router or a Layer 3 switch to provide routing services. Inter-VLAN routing is the process of forwarding network traffic from one VLAN to another VLAN. Three options include legacy, router-on-a-stick, and Layer 3 switch using SVIs. Legacy used a router with multiple Ethernet interfaces. Each router interface was connected to a switch port in different VLANs. Requiring one physical router interface per VLAN quickly exhausts the physical interface capacity of a router. The 'router-on-a-stick' inter-VLAN routing method only requires one physical Ethernet interface to route traffic between multiple VLANs on a network. A Cisco IOS router Ethernet interface is configured as an 802.1Q trunk and connected to a trunk port on a Layer 2 switch. The router interface is configured using subinterfaces to identify routable VLANs. The configured subinterfaces are software-based virtual interfaces, associated with a single physical Ethernet interface. The modern method is Inter-VLAN routing on a Layer 3 switch using SVIs. The SVI is created for a VLAN that exists on the switch. The SVI performs the same functions for the VLAN as a router interface. It provides Layer 3 processing for packets being sent to or from all switch ports associated with that VLAN.

Router-on-a-Stick Inter-VLAN Routing

To configure a switch with VLANs and trunking, complete the following steps: create and name the VLANs, create the management interface, configure access ports, and configure trunking ports. The router-on-a-stick method requires a subinterface to be created for each VLAN to be routed. A subinterface is created using the interface *interface_id subinterface_id* global configuration mode command. Each router subinterface must be assigned an IP address on a unique subnet for routing to occur. When all subinterfaces have been created, the physical interface must be enabled using the no shutdown interface configuration command. From a host, verify connectivity to a host in another VLAN using the ping command. Use ping to verify connectivity with the host and the switch. To verify and troubleshoot use the show ip route, show ip interface brief, show interfaces, and show interfaces trunk commands.

Inter-VLAN Routing using Layer 3 Switches

Enterprise campus LANs use Layer 3 switches to provide inter-VLAN routing. Layer 3 switches use hardware-based switching to achieve higher-packet processing rates than routers. Capabilities of a Layer 3 switch include routing from one VLAN to another using multiple switched virtual interfaces (SVIs) and converting a Layer 2 switchport to a Layer 3 interface (i.e., a routed port). To provide inter-VLAN routing, Layer 3 switches use SVIs. SVIs are configured using the same interface vlan *vlan-id* command used to create the management SVI on a Layer

2 switch. A Layer 3 SVI must be created for each of the routable VLANs. To configure a switch with VLANS and trunking, complete the following steps: create the VLANS, create the SVI VLAN interfaces, configure access ports, and enable IP routing. From a host, verify connectivity to a host in another VLAN using the ping command. Next, verify connectivity with the host using the ping Windows host command. VLANs must be advertised using static or dynamic routing. To enable routing on a Layer 3 switch, a routed port must be configured. A routed port is created on a Layer 3 switch by disabling the switchport feature on a Layer 2 port that is connected to another Layer 3 device. The interface can be configured with an IPv4 configuration to connect to a router or another Layer 3 switch. To configure a Layer 3 switch to route with a router, follow these steps: configure the routed port, enable routing, configure routing, verify routing, and verify connectivity.

Troubleshoot Inter-VLAN Routing

There are a number of reasons why an inter-VAN configuration may not work. All are related to connectivity issues such as missing VLANs, switch trunk port issues, switch access port issues, and router configuration issues. A VLAN could be missing if it was not created, it was accidently deleted, or it is not allowed on the trunk link. Another issue for inter-VLAN routing includes misconfigured switch ports. In a legacy inter-VLAN solution, a misconfigured switch port could be caused when the connecting router port is not assigned to the correct VLAN. With a router-on-a-stick solution, the most common cause is a misconfigured trunk port. When a problem is suspected with a switch access port configuration, use ping and show interfaces *interface-id* switchport commands to identify the problem. Router configuration problems with router-on-a-stick configurations are usually related to subinterface misconfigurations. Verify the subinterface status using the show ip interface brief command.

STP

Spanning Tree Protocol

- Designed specifically to eliminate Layer 2 loops in network

Redundancy

- Hierarchical network design, fixing the problem of a single point of failure, yet creates **Layer 2 loops**
- Can include both physical and logical redundancy

Spanning Tree Protocol (cont)

- Loop-prevention network protocol that allows for redundancy while creating a loop-free Layer 2 topology
- IEEE 802.1D is original IEEE MAC Bridging standard for STP

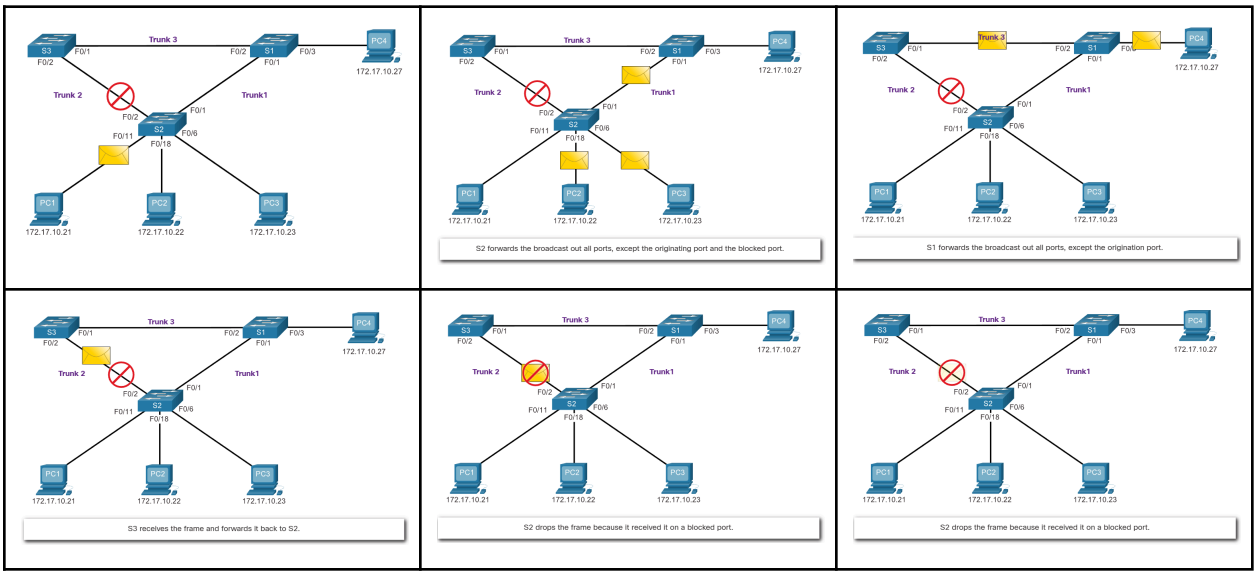

STP Recalculation

- STP compensates for Network Failure

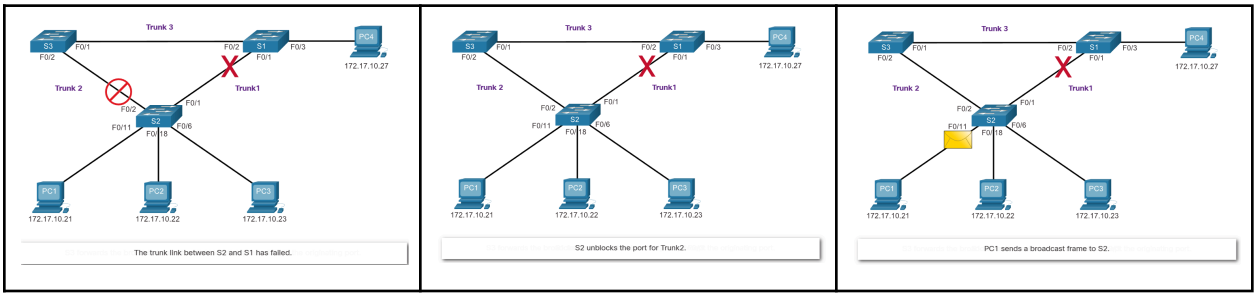

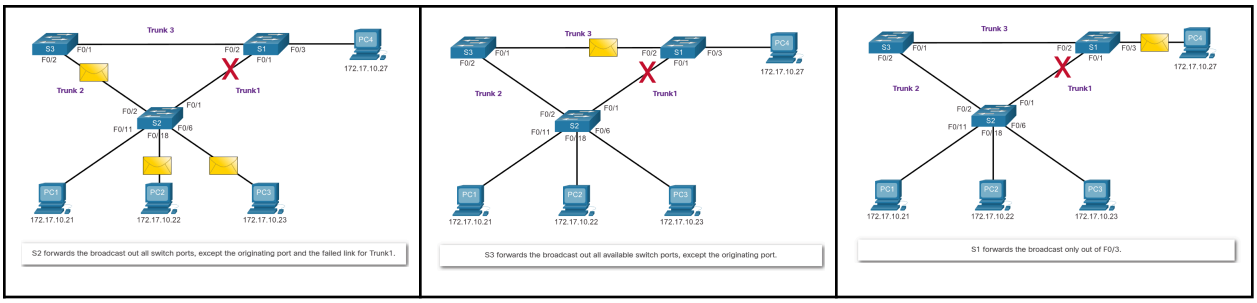

Redundant Switch Links Issues

- Path redundancy provides multiple network services by eliminating single points of failure.
- When multiple paths exist between two devices on network, and there is **no** spanning tree
    - Layer 2 loop occurs

Layer 2 loop

- Can result in MAC address table instability, link saturation, and high CPU utilization on switches and end devices

Layer 2 ethernet

Unlike Layer 3 protocols (IPv4 and IPv6)

- Does not include a mechanism to recognize and eliminate endlessly loop frames

IPv4 and IPv6

- Includes mechanism that limits the number of times a Layer 3 network device can retransmit a packet

A router will decrement TTL (Time to Live) in every IPv4 packet

- Hop Limit field in every IPv6 packet
- When decremented to 0, router will **drop** the packet

Ethernet and Ethernet switches have no comparable mechanism for limiting the number of times a switch retransmits a Layer 2 frame

Layer 2 Loops

Without STP enabled, Layer 2 loops can form

- Causes broadcast, multicast and unknown unicast frames to loop endlessly

Example: ARP Request are forwarded out of all of the switch ports, except ingress port.

- This ensures all devices in the broadcast domain are able to receive the frame.

- If there are more than one path for the frame to be forwarded out of, an endless loop can result.
- When loop occurs, MAC address table on switch will constantly change with the updates from the broadcast frames, resulting in MAC database instability
    - Causing high CPU utilization, which makes switch unable to forward frames

Unknown Unicast frames

- Sent onto a looped network can result in duplicate frames arriving at the destination device
- * This frame is when the switch does not have the destination MAC address in its MAC address table and must forward the frame out all ports, except ingress.

* Verify a switch is a root bridge - *show spanning-tree*

Broadcast Storm

- An abnormally high number of broadcasts overwhelming the network during a specific amount of time
- Can disable a network within seconds by overwhelming switches and end devices
- Can be caused by a hardware problem such as a faulty NIC or from a Layer 2 loop in the network

Layer 2 loops (cont)

- Likely to have immediate and disabling consequences on the network

Layer 2 multicasts are typically forwarded the same way as a broadcast by the switch

- IPv6 packets are **never** forwarded as a Layer 2 broadcast
- ICMPv6 Neighbor Discovery uses **Layer 2 multicasts**

Note - A host caught in a Layer 2 loop is not accessible to other hosts on the network

- Due to the constant change in its MAC address table, the switch doesn't know out of which port to forward unicast frames

Note - **Spanning tree is enabled**, by **default**, on Cisco switches to **prevent Layer 2 loops** from occurring

The Spanning Tree Algorithm

STP is based on an algorithm invented by Radia Perlman

- Published in a 1985 paper "An Algorithm for Distributed Computation of a Spanning Tree in an Extended LAN

Spanning Tree Algorithm (STA)

- Creates a loop-free topology by selecting a single root bridge where all other switches determine a single least-cost path

STA Scenario Topology	
STA Scenario Topology This STA scenario uses an Ethernet LAN with redundant connections between multiple switches	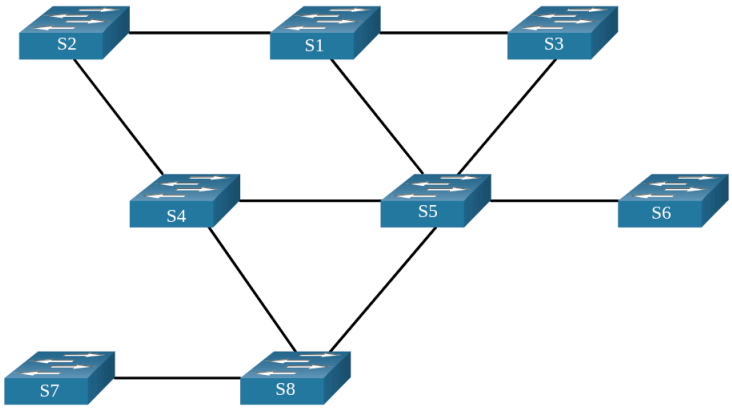
Select the Root Bridge	
Select the Root Bridge The spanning tree algorithm begins by selecting a single root bridge. The figure shows that switch S1 has been selected as the root bridge. In this topology, all links are equal cost (same bandwidth). Each switch will determine a single, least cost path from itself to the root bridge. Note - The STA and STP refers to switches as bridges. This is because in the early days of Ethernet, switches were referred to as bridges.	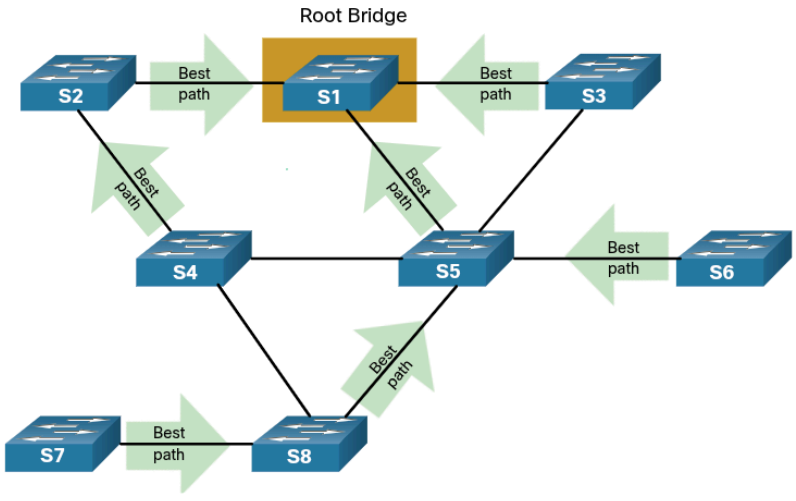

Block Redundant Paths	
Block Redundant Paths STP ensures that there is only one logical path between all destinations on the network by intentionally blocking redundant paths that could cause a loop. When a port is blocked, user data is prevented from entering or leaving that port. Blocking the redundant paths is critical to preventing loops on the network	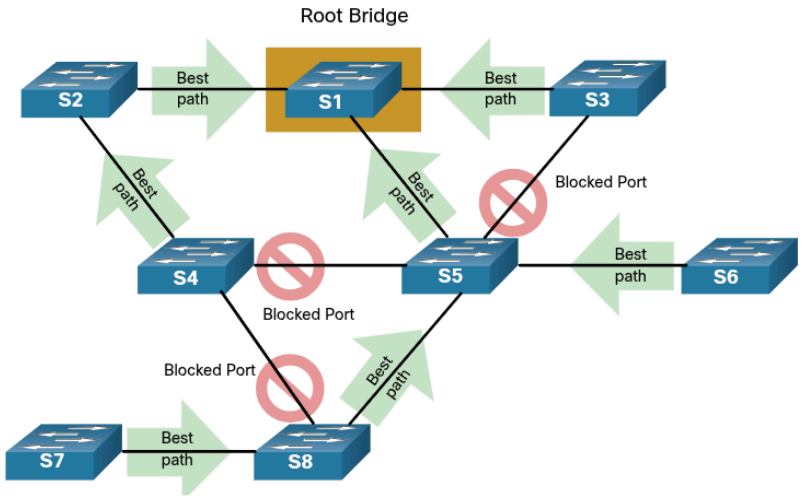 *Switches S4, S5, and S8 have blocked redundant paths to the root bridge*
Loop-Free Topology	
Loop-Free Topology A blocked port has the effect of making that link a non-forwarding link between the two switches. Notice that this creates a topology where each switch has only a single path to the root bridge.	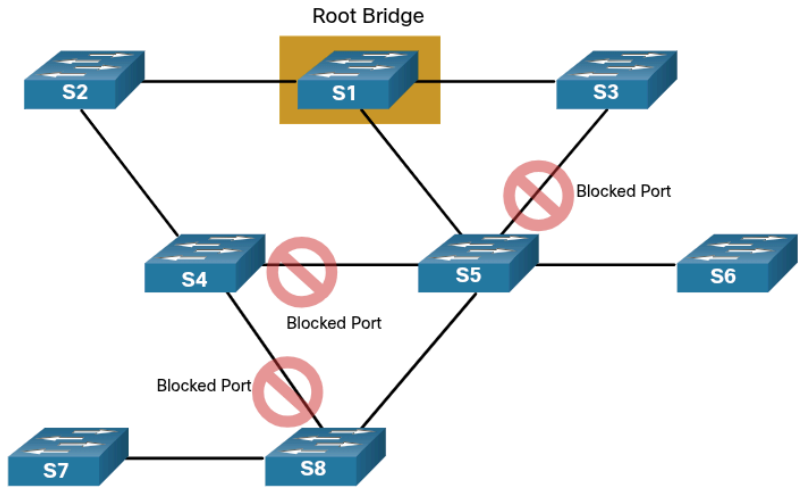 *Each switch not has just one forwarding path to the root bridge*
Link Failure Causes Recalculation	
The physical paths still exist to provide redundancy, but these paths are disabled to prevent the loops from occurring. If the	

path is ever needed to compensate for a network cable or switch failure, STP recalculates the paths and unblocks the necessary ports to allow the redundant path to become Active.

STP recalculations can also occur any time a new switch or new inter-switch link is added to the network.

The figure shows a link failure between switches S2 and S4 causing STP to recalculate.

Notice that the previously redundant link between S4 and S5 is now forwarding to compensate for this failure.

There is only one path between every switch and the root bridge

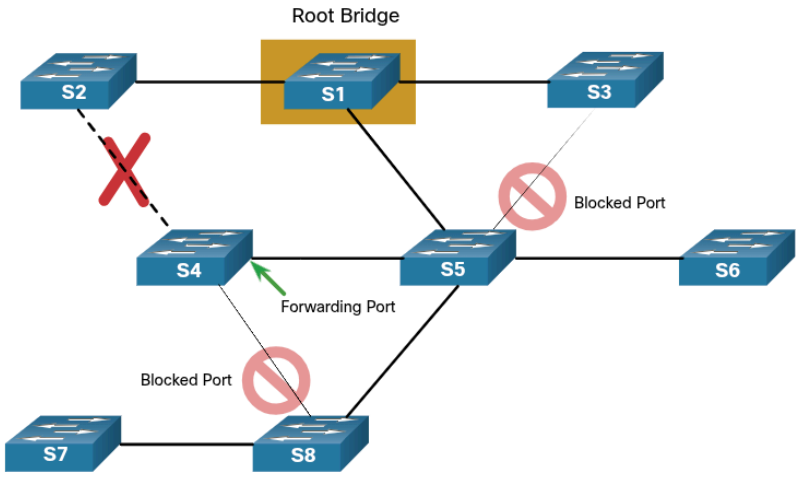

Note - STP prevents loops from occurring by configuring a loop-free path through the network using strategically placed “blocking-state” ports.

The switches running STP are able to compensate for failures by dynamically unblocking the previously blocked ports and permitting traffic to traverse the alternate paths.

STP Operations

Steps to a Loop-Free Topology

1. Elect the root bridge
2. Elect the root ports
3. Elect designated ports
4. Elect alternative (blocked) ports

Bridge Protocol Data Units (BPDUs)

- Switches use this to share information about themselves and their connections during STA and STP functions

- Used to select the root bridge, root ports, designated ports, and alternate ports
- Each BPDU contains a **bridge ID (BID)**

Bridge ID (BID)

- Identifies which switch sent the BPDU
- Involved in making many of the STA decisions including **root bridge** and **port roles**

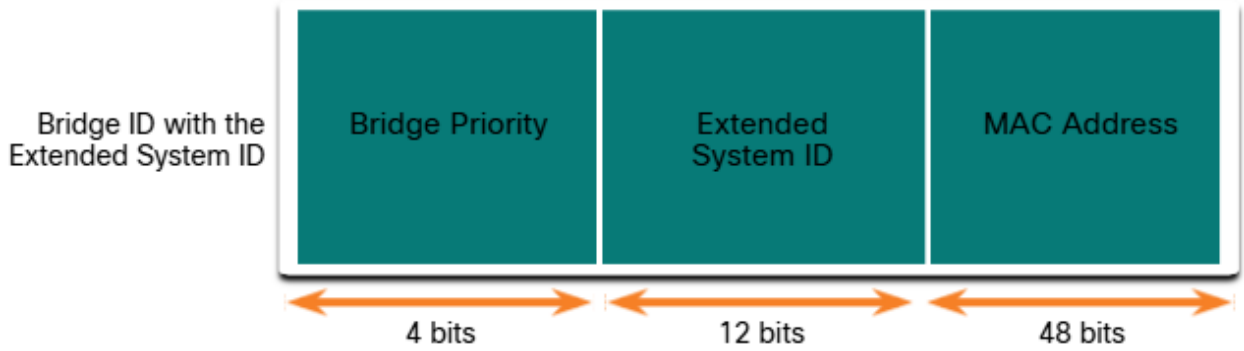

*The BID includes the **Bridge Priority**, the **Extended System ID**, and the **MAC Address** of the switch*

Bridge Priority

- The **default priority value** for all Cisco switches is the decimal value **32768**.
- The range is 0 to 61440 in increments of 4096
- A lower bridge priority is preferable
- A bridge priority of **0** takes precedence over all other bridge priorities

Extended System ID

- Value is a decimal value added to the bridge priority value in the BID to identify the VLAN for the BPDU

Early IEEE 802.1D were designed for networks that didn't use VLANs

- Single common spanning tree across all switches

Example:

Older switches, the extended system ID was **not** included in the BPDUs

802.1D was enhanced to include support for VLANs

- Required that a 12-bit VLAN ID be included in BPDU frame
- VLAN information is included in BPDU frame through use of extended system ID

Note - Extended system ID allows later implementations of STP to have different root bridges for different sets of VLANs. This can allow for redundant, non-forwarding links in a STP topology for one set of VLANs to be used by a different set of VLANs using a different root bridge

MAC Address

- When two switches are configured with the **same** priority and have the **same** extended system ID, the switch having the MAC address with the lowest value, expressed in hexadecimal, will have the lower BID

Elect the Root Bridge

The STA designates a single switch as the root bridge and uses it as the reference point for all path calculations.

Switches exchange BPDUs to build the loop-free topology beginning with selecting the root bridge

Election process

- Determines which switch becomes the root bridge.
- All switches in the broadcast domain participate in the election process.
- After a switch boots, it begins to send out BPDU frames every two seconds

Root ID

- BPDU frames during boot contain the BID of the sending switch and the BID of the root bridge

Note - The switch with the lowest BID will become the root bridge

- At first, all switches declare themselves as the root bridge with their own BID set as the Root ID
- Eventually, the switches learn through the exchange of BPDUs which switch has the lowest BID and will agree on one root bridge

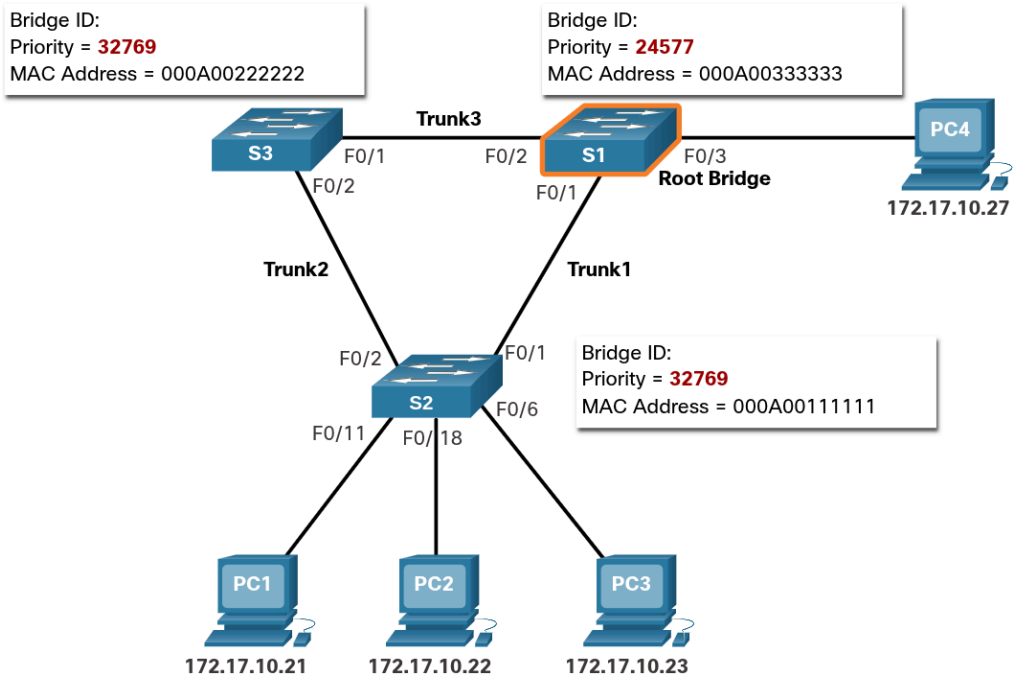

Impact of Default BIDs

- Because the default priority is **32768**, it is possible for two or more switches to have the **same** priority
- The switch with the lowest MAC address will become the root bridge

Note - Recommended that the administrator configure the desired root bridge switch with a lower priority

In the lower example, the priority of all the switches is **32769**. The value is based on the **32768** default bridge priority and the extended system ID (VLAN 1 assignment) associated with each switch (32768+1)

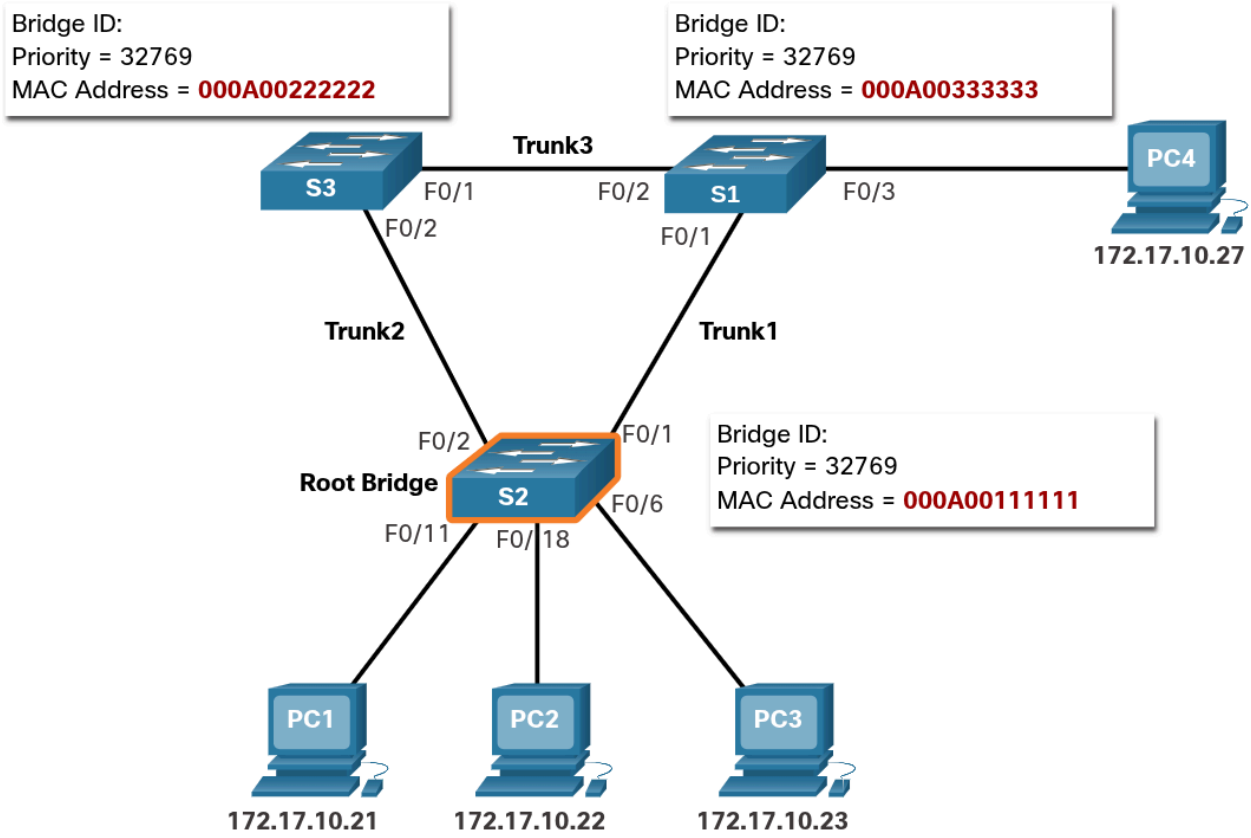

Determine the Root Path Cost

When the root bridge has been elected for a given spanning tree instance, the STA starts the process of determining the best paths to the root bridge from all destinations in the broadcast domain

Internal root path cost

- The path information is determined by the sum of all the individual port costs along the path from the switch to the root bridge

Note - BPDU includes the root path cost

- This is the cost of the path from the sending switch to the root bridge
- When a switch receives the BPDU, it adds the ingress port cost of the segment to determine its internal root path cost

The default port costs are defined at which the port operates

- Cisco switches by default use the values as defined by the IEEE 802.1D standard
    - Also known as short path cost, for STP and RSTP
- The IEEE suggests using values defined in IEEE-802.1w
    - Also known as long path cost, when using 10 Gbps links and faster

Link Speed	STP Cost: IEEE 802.1D-1998	RSTP Cost: IEEE 802.1w-2004
10 Gbps	2	2,000
1 Gbps	4	20,000
100 Mbps	19	200,000
10 Mbps	100	2,000,000

Note - Although switch ports have default port cost associated, the port cost is configurable

Elect the Root Ports

After the root bridge has been determined, the STA algorithm is used to select the root port.

- Every non-root switch will select one root port

Root Port

- Port closest to the root bridge in terms of overall cost (best path) to the root bridge
- Overall cost is known as the **internal root path cost**

Internal root path

- Cost is equal to sum of all port costs along the path to the root bridge
- Paths with lowest cost become preferred, and all other redundant paths are blocked

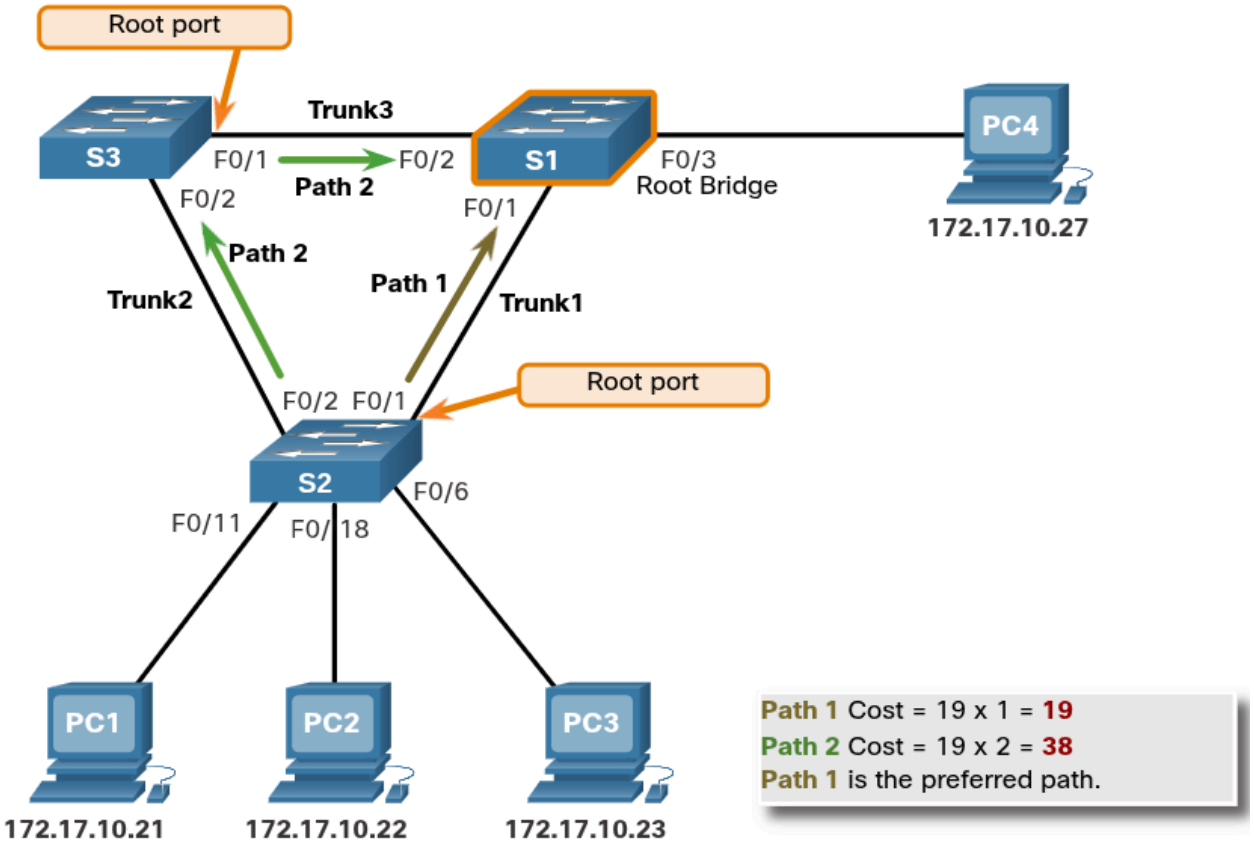

Elect Designated Ports

Loop prevention of spanning tree becomes evident in the next two steps

- After each switch selects a root port, the switches will then select designated ports

Note - Every segment between two switches have one designated port

- The designated port on the segment (with two switches) that has the LOWEST internal root path cost to the root bridge
    - IE: The designated port has the best path to receive traffic leading to the root bridge

Note - What is not a root port or a designated port becomes an alternative or blocked port. The end result is a single path from every switch to the root bridge

How STA selects the designated ports

Designated Ports on Root Bridge

All ports on the root bridge are designated ports. This is because the root bridge has the lowest cost to itself.

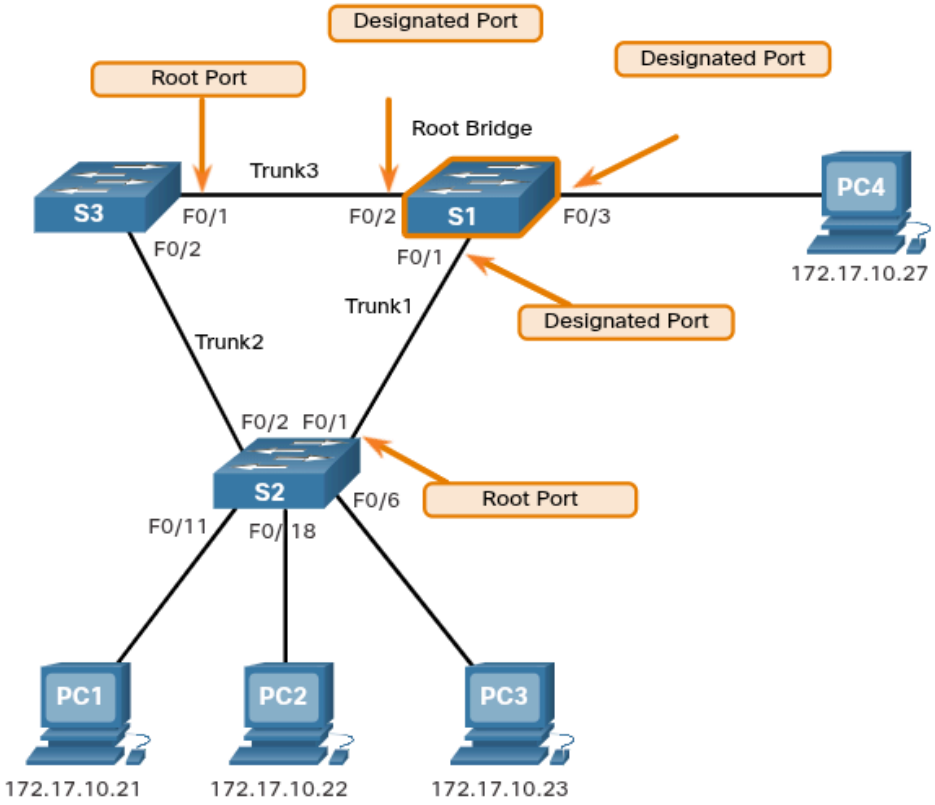

All the ports on the root bridge are designated ports

How STA selects the designated ports

Designated Port When There is a Root Port

If one end of a segment is a root port, then the other end is a designated port.The figure shows switch S4 is connected to S3.

The Fa0/1 interface on S4 is its root port because it has the best and only path to the root bridge.

The Fa0/3 interface on S3 at the other end of the segment would therefore, be the designated port.

Note - All switch ports with end devices (hosts) attached are designated ports

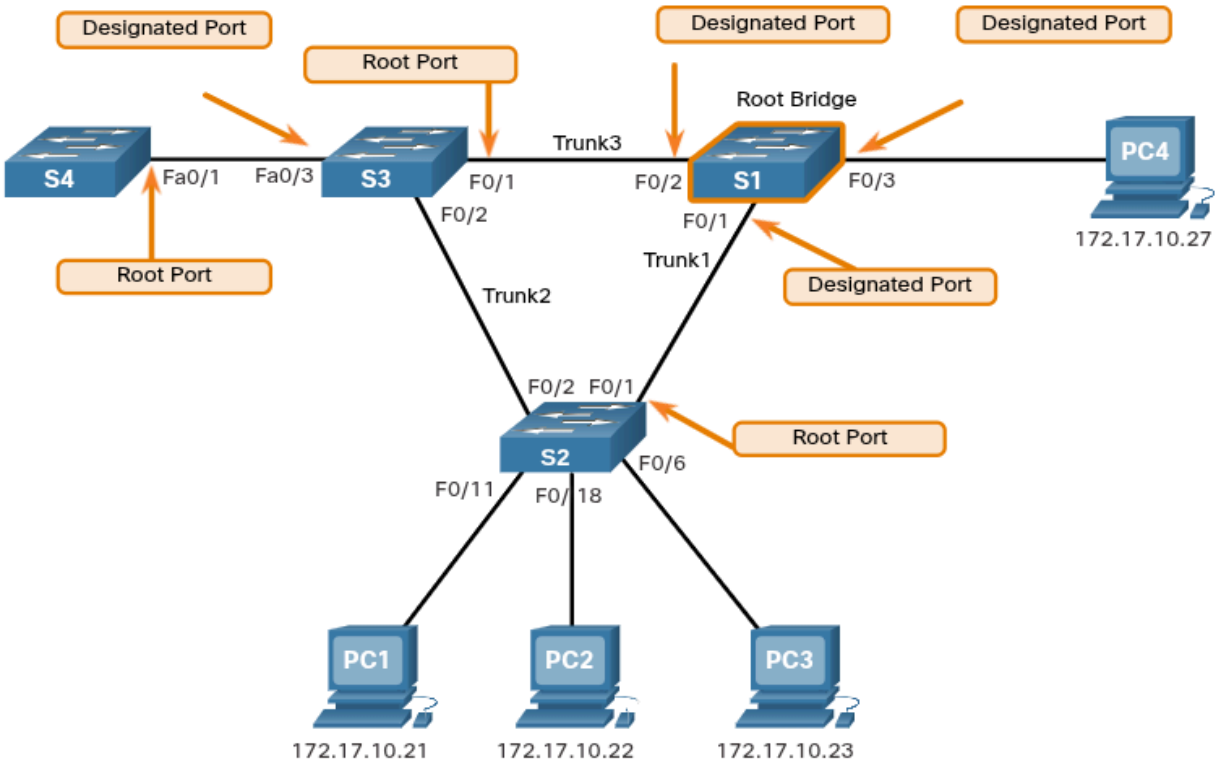

Fa0/1 interface on S4 is a root port because the Fa0/3 interface of S3 is a designated port.

How STA selects the designated ports

Designated Port When There is No Root Port

This leaves only segments between two switches where neither of the switches is the root bridge.

For example, the last segment is the one between S2 and S3. Both S2 and S3 have the same path cost to the root bridge.

The spanning tree algorithm will use the bridge ID as a tie breaker.

Although not shown, S2 has a lower BID. Therefore, the F0/2 port of S2 will be chosen as the designated port. Designated ports are in the **forwarding** state.

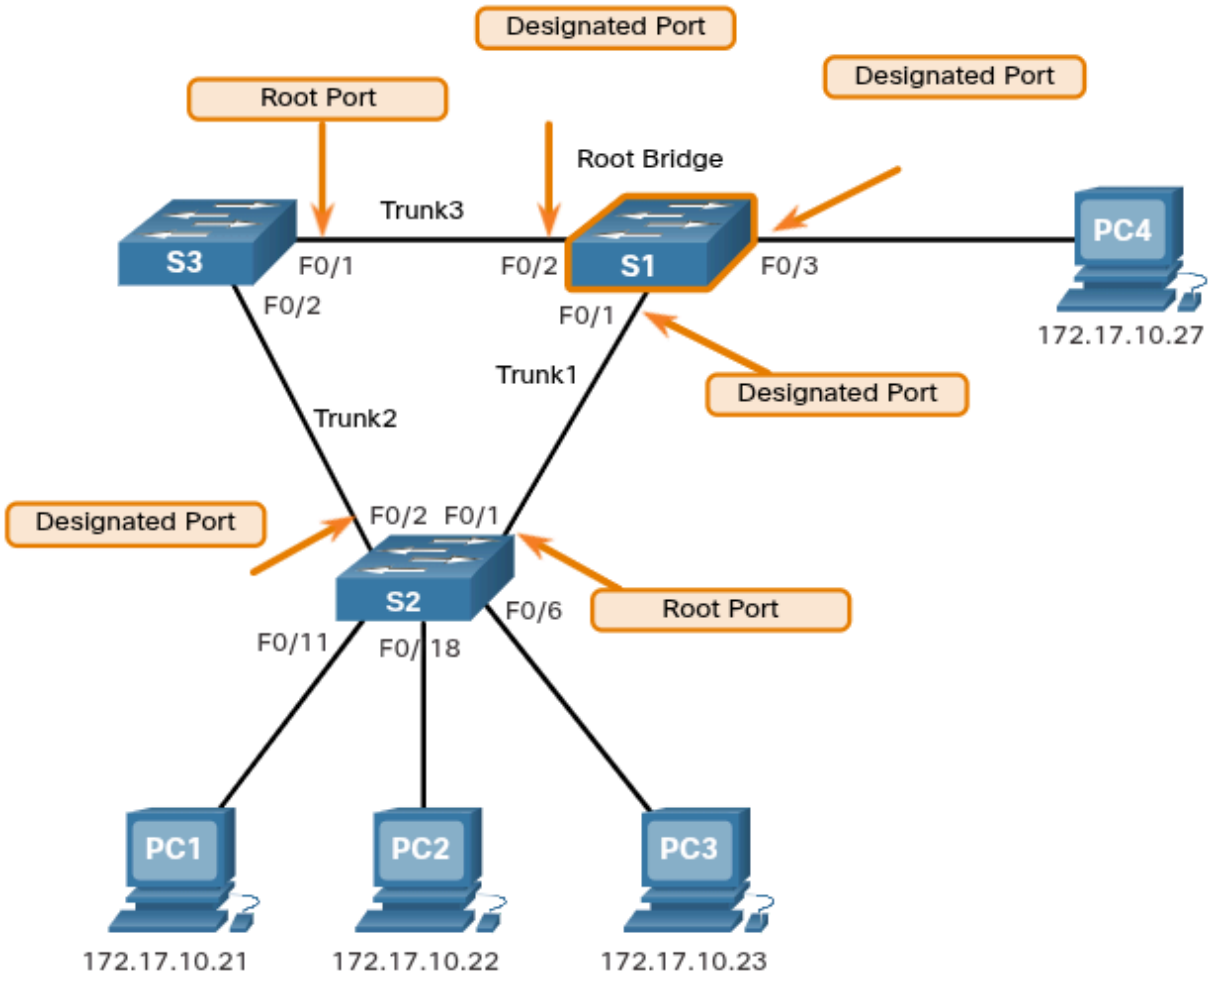

The Fa0/2 interface of S2 is the designated port on the segment with S3.

Elect Alternate (Blocked) Ports

If a port is **not** a root port or a designated port, then it becomes an alternate (or backup) port.
Alternate ports and backup ports are in **discarding** or **blocking** state to prevent loops.

Figure: The STA has configured port F0/2 on S3 in the alternate role.
Port F0/2 on S3 is in **blocking** state and will not forward Ethernet frames.
All other inter-switch ports are in **forwarding** state. This is the loop-prevention part of STP.

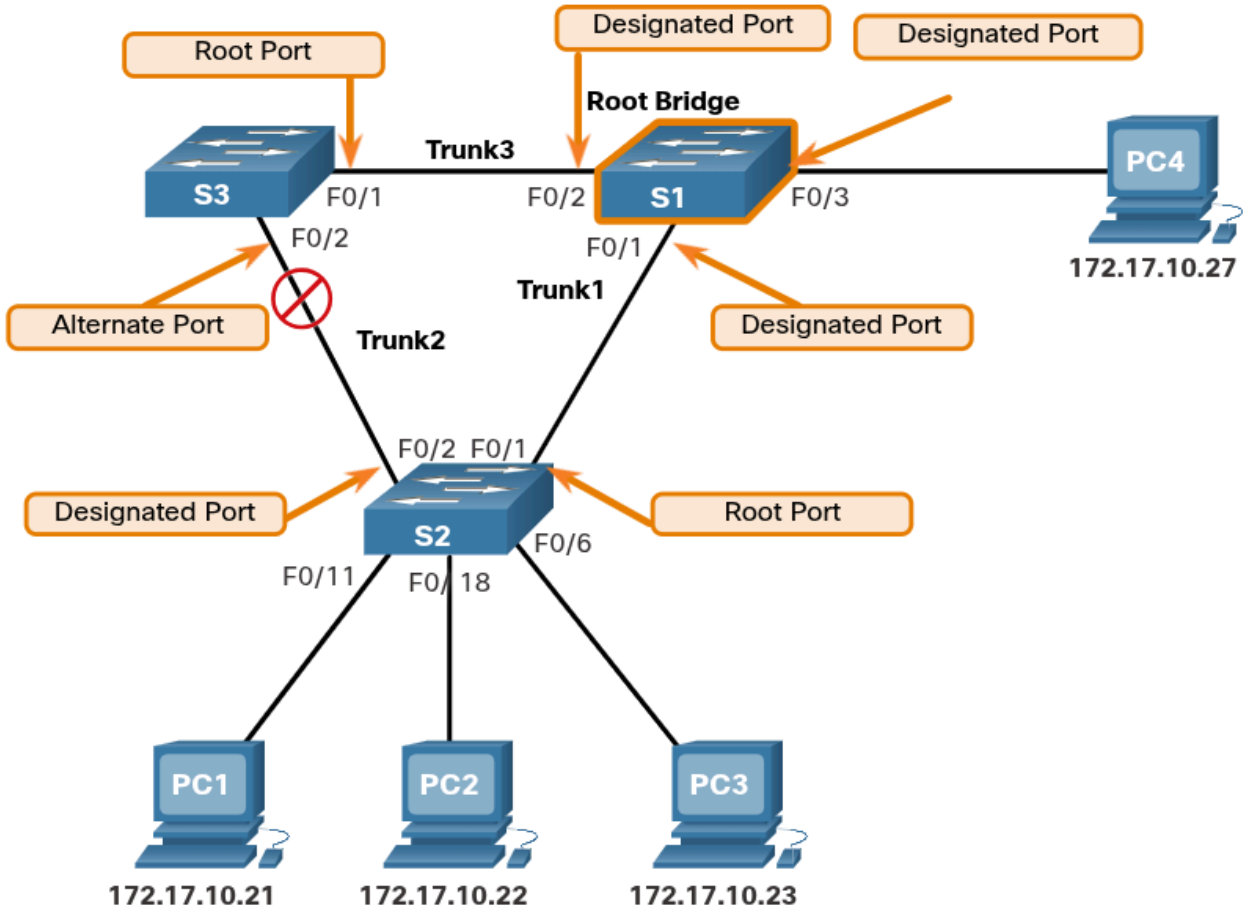

*The Fa0/2 interface of S3 is not a root port or a designated port, so it becomes an **alternate** or **blocked** port.*

Elect a Root Port from Multiple Equal-Cost Paths

Root port and designated ports are based on the lowest path cost to the root bridge.

Scenarios:

- What happens if the switch has multiple equal-cost paths to the root bridge?
    - How does a switch designate a root port?

1. Lowest sender BID
2. Lowest sender port priority
3. Lowest sender port ID

Example and Explanation

Lowest Sender BID

The figure shows a topology with four switches, including switch S1 as the root bridge.

- Port F0/1 on switch S3 and port F0/3 on switch 4 have been selected as root ports because they have the lowest cost path (root path cost) to the root bridge.
- S2 has two ports, F0/1 and F0/2 with equal cost paths to the root bridge

The bridge IDs of the neighboring switches, S3 and S4, will be used to break the tie.

- This is known as sender's BID
- S3 has a BID of 32769.5555.5555.5555 and S4 has a BID of 32769.1111.1111.1111.
    - S4 has a lower BID, the F0/1 port of S2, which is connected to S4, will be root port

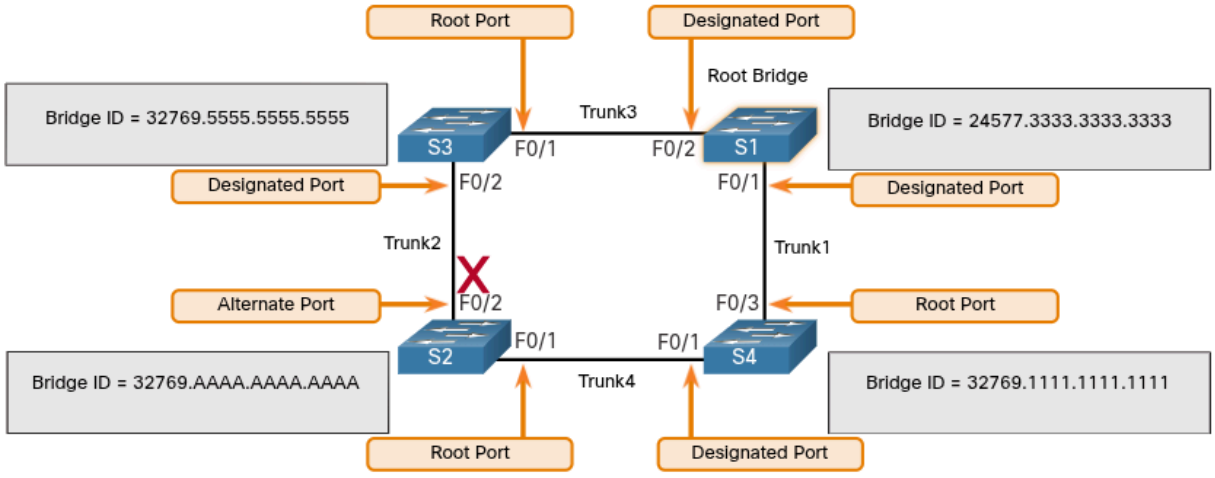

Example and Explanation

Lowest Sender Port Priority

To demonstrate the next two criteria, the topology is changed to one where two switches are connected with two equal-cost paths between them.

- S1 is the root bridge, so both ports are designated ports

S4 has two ports with equal-cost paths to the root bridge

- Because both are connected to the same switch, the sender's BID (S1) is equal.
- The first step is a tie

Next on the list is the sender's (S1) port priority.

- The default port priority is 128, so both ports on S1 have the same port priority
  - This is also a tie.
- If either port on S1 was configured with a lower port priority, S4 would put its adjacent port in **forwarding** state. The other port on S4 would be a **blocking** state.

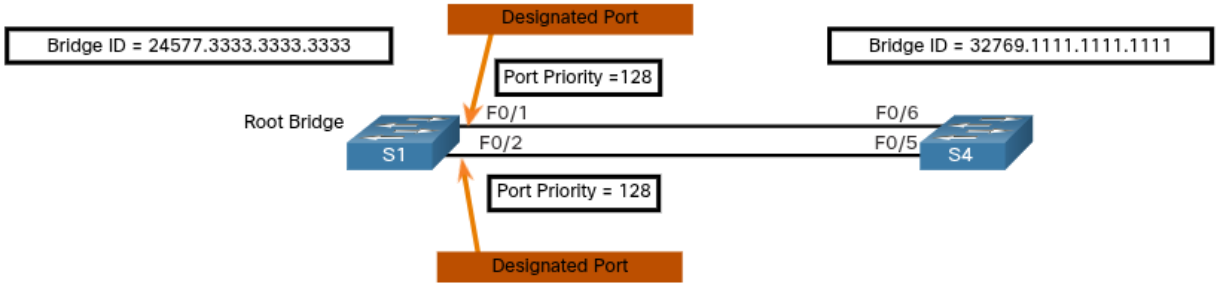

Lowest Sender Port ID

The last tie-breaker is the lowest sender's port ID.

Switch S4 has received BPDUs from port F0/1 and port F0/2 on S1.

Note - the decision is based on the sender's port ID.

Because the port ID of F0/1 on S1 is lower than port F0/2, the port F0/6 on switch S4 will be the root port. This is the port on S4 that is connected to the F0/1 port on S1.

Port F0/5 on S4 will become an alternate port and placed in the **blocking** state, which is the loop-prevention part of STP.

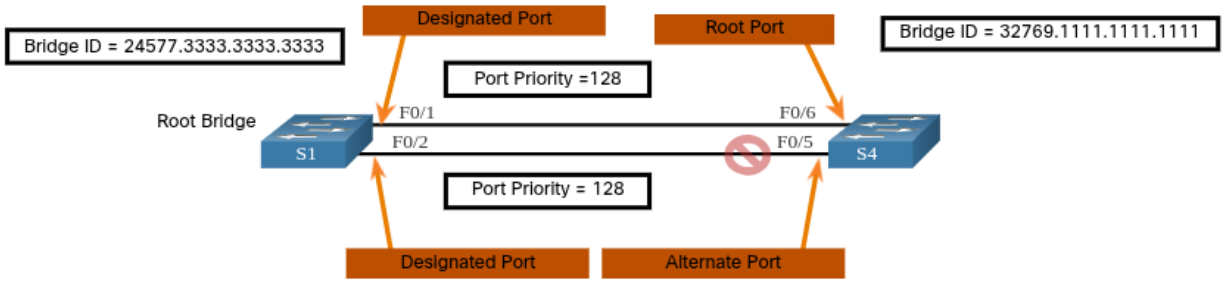

STP Timers and Port States

STP convergence requires three timers:

- **Hello Timer** - The interval between BPDUs
    - The default is 2 seconds but can be modified between **1** and **10 seconds**
- **Forward Delay Timer** - The time that is spent in the **listening** and **learning** state
    - The default is 15 seconds but can be modified between **4** and **30 seconds**
- **Max Age Timer** - The maximum length of time that a switch waits before attempting to stop the STP topology
    - The default is 20 seconds but can be modified between **6** and **40** seconds

Note - The default times can be changed on the root bridge, which dictates the value of these timers for the STP domain.

- STP facilitates the logical loop-free path throughout the broadcast domain.
- The spanning tree is determined through the information learned by the exchange of the BPDU frames between the interconnected switches.
- If a switch port transitions directly from the blocking state to the forwarding state without information about the full topology during the transition, the port can temporarily create a **data loop**

STP has **five ports states**, **four** of which are **operational** port states.

- The disabled state is considered non-operational

Note - To avoid problems with STP, IEEE recommends a maximum diameter of **seven** switches when using the default STP timers

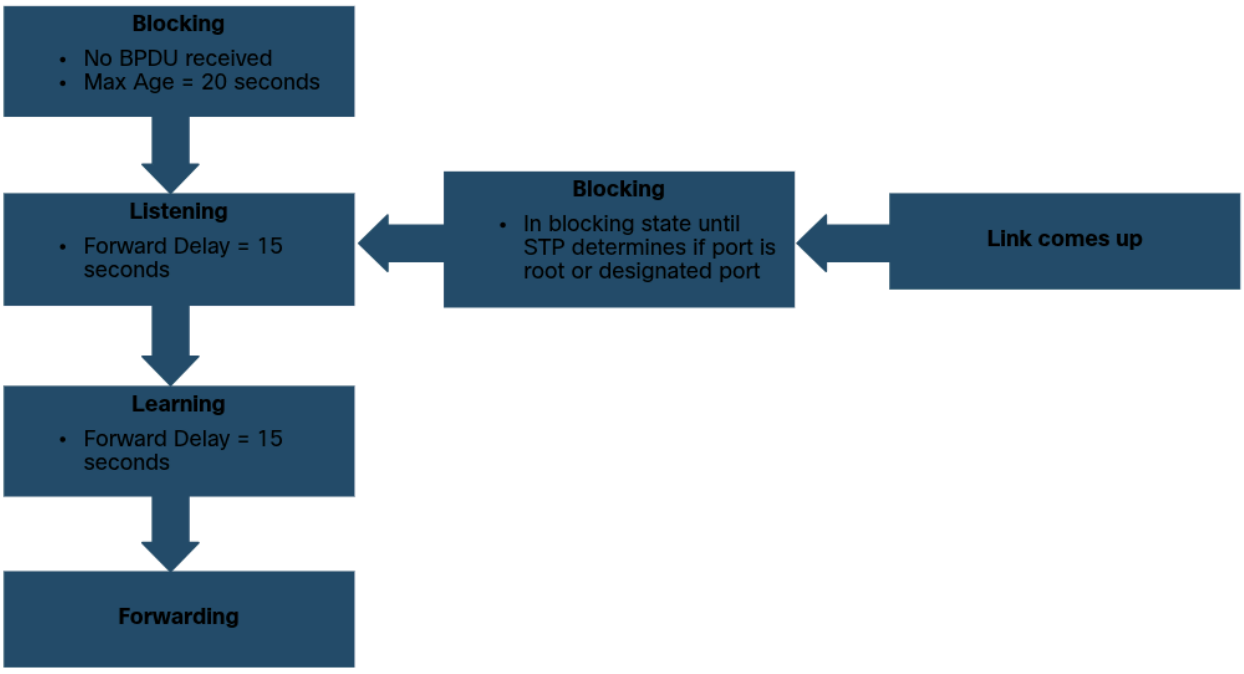

Port State	Description
Blocking	The port is an alternate port and does not participate in frame forwarding. The port receives BPDU frames to determine the location and root ID of the root bridge. The BPDU frames also determine which port roles each switch port should assume in the final active STP topology. - With a **Max Age timer** of **20 seconds**, a switch port that has not received an expected BPDU from a neighbor switch will go into the blocking state.
Listening	After the blocking state, a port will move to the listening state. The port receives BPDUs to determine the path to the root. The switch port also transmits its own BPDU frames and informs adjacent switches that the switch port is preparing to participate in the active topology.
Learning	A switch port transitions to the learning state after the listening state. During the listening state, the switch port receives and processes BPDUs and prepares to participate in frame forwarding. It also begins to populate the MAC address table. - However, in the learning state, user frames are **not forwarded** to the destination
Forwarding	In the forwarding state, a switch port is considered part of the **active topology**. The switch port forwards user traffic and sends and receives BPDU frames.
Disabled	A switch port in the disabled state does not participate in spanning tree and does not forward frames. The disabled state is set when the switch port is **administratively disabled**.

Operational Details of Each Port State			
Port State	**BPDU**	MAC Address Table	Forwarding Data Frames
Blocking	Receive only	No update	No
Listening	Receive and send	No update	No
Learning	Receive and send	Updating table	No
Forwarding	Receive and send	Updating table	Yes
Disabled	None sent or received	No update	No

Per-VLAN Spanning Tree

STP can be configured to operate in an environment with multiple VLANs

Per-VLAN Spanning Tree (PVST) versions of STP

- There is a root bridge elected for each spanning tee instance.
    - This makes it possible to have different root bridges for different sets of VLANs
- STP operates a separate instance of STP for each individual VLAN. If all ports on all switches are members of VLAN 1, then there is only **one** spanning tree instance

Different Versions of STP

Spanning Tree Protocol and acronym STP can be misleading

- Generally use term **Rapid Spanning Tree Protocol (RSTP)** and **Multiple Spanning Tree Protocol (MSTP)**

Latest standard for spanning tree

- Contained in IEEE-802-1D-2004
- IEEE standard for local and metropolitan area networks: Media Access Control (MAC) Bridges
    - This version states that switches and bridges that comply with the standard will use Rapid Spanning Tree Protocol (RSTP) instead of older STP protocol (specified in original 802.1d standard)

STP Variety	Description
STP	This is the original IEEE 802.1D version (802.1D-1998 and earlier) that provides a loop-free topology in a network with redundant links. Also called **Common Spanning tree (CST)**, it assumes one spanning tree instance for the entire bridged network, regardless of the number of VLANs
PVST+	Per-VLAN Spanning Tree (PVST+) is a Cisco enhancement of STP that provides a separate 802.1D spanning tree instance for each VLAN configured in the network. Supports PortFast, UplinkFast, BackboneFast, BPDU guard, BPDU filter, root guard, and loop guard.
RSTP	Rapid Spanning Tree Protocol (RSTP) or IEEE 802.1w is an evolution of STP that provides faster convergence than STP
802.1D-2004	This is an updated version of the STP standard, incorporating IEEE 802.1w

Rapid PVST+	This is a Cisco enhancement of RSTP that uses PVST+ and provides a separate instance of 802.1w per VLAN. Each separate instance supports PortFast, BPDU filter, root guard, and loop guard
MSTP	Multiple Spanning Tree Protocol (MSTP) is an IEEE standard inspired by the earlier Cisco proprietary Multiple Instance STP (MISTP) implementation. MSTP maps multiple VLANs into the **same** spanning tree instance.
MST	Multiple Spanning Tree (MST) is the Cisco implementation of MSTP, which provides up to **16** instances of RSTP and combines many VLANs with the same physical and logical topology into a common RSTP instance. Each instance supports PortFast, BPDU guard, BPDU filter, root guard, and loop guard.

Cisco switches running IOS 15.0 or later, running PVST+ by default

- 15.0+ incorporates many specifications of IEEE 802.1D-2004. Such as an alternate ports in place of the former non-designated ports
- Switches must be explicitly configured for **rapid spanning tree mode** in order to run the rapid spanning tree protocol

RSTP Concepts

- (IEEE 802.1w) supersedes original 802.1D while retaining backward compatibility.
- 802.1w STP terminology remains primarily the same as the original IEEE 802.1D STP terminology.
    - Most parameters have been left unchanged
- Increases the speed of the recalculation of the spanning tree when the Layer 2 network topology changes.
- Can achieve much faster convergence in a properly configured network, sometimes in as little as a few hundred milliseconds.
    - If a port is configured to be an alternate port it can immediately change to a forwarding state without waiting for the network to converge

Note - Rapid PVST+ is the Cisco implementation of RSTP on a per-VLAN basis. With Rapid PVST+ an independent instance of RSTP runs for each VLAN

RSTP Port States and Port Roles

- Port states and Port Roles between STP and RSTP are similar

STP and RSTP Port States

- There are only three port states in RSTP that correspond to the three possible operational states in STP.
- The 802.1D disabled, blocking, and listening states are merged into a unique 802.1w discarding state

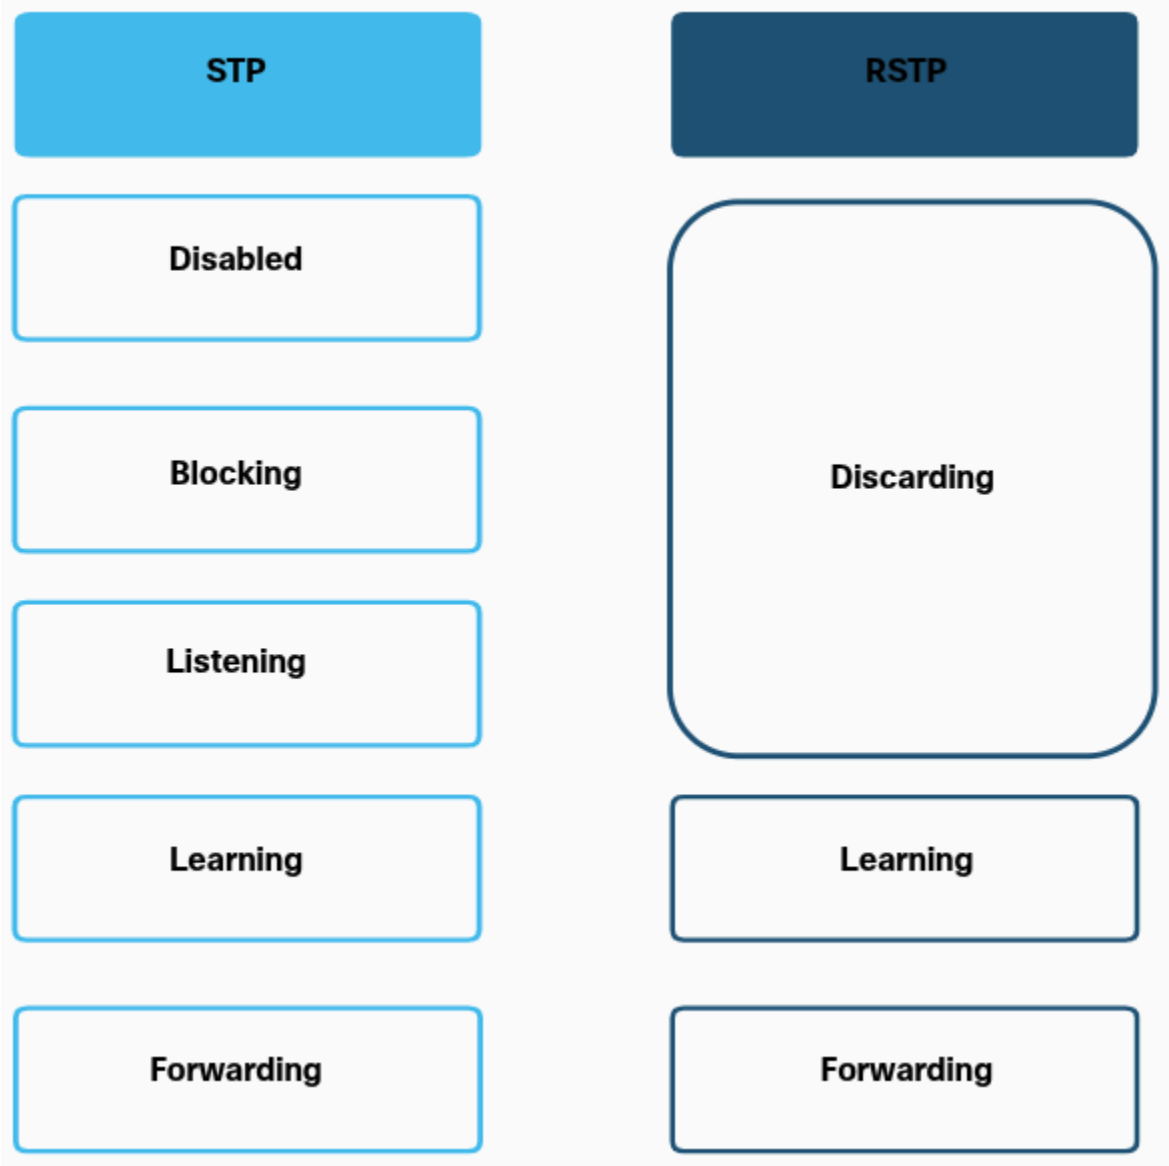

STP and RSTP Port Roles

- Root ports and designated ports are the same for both STP and RSTP.
    - However, there are two RSTP port roles that correspond to the blocking state of STP
    - In STP, a blocked port is defined as not being the designated or root port. RSTP has two port roles for this purpose

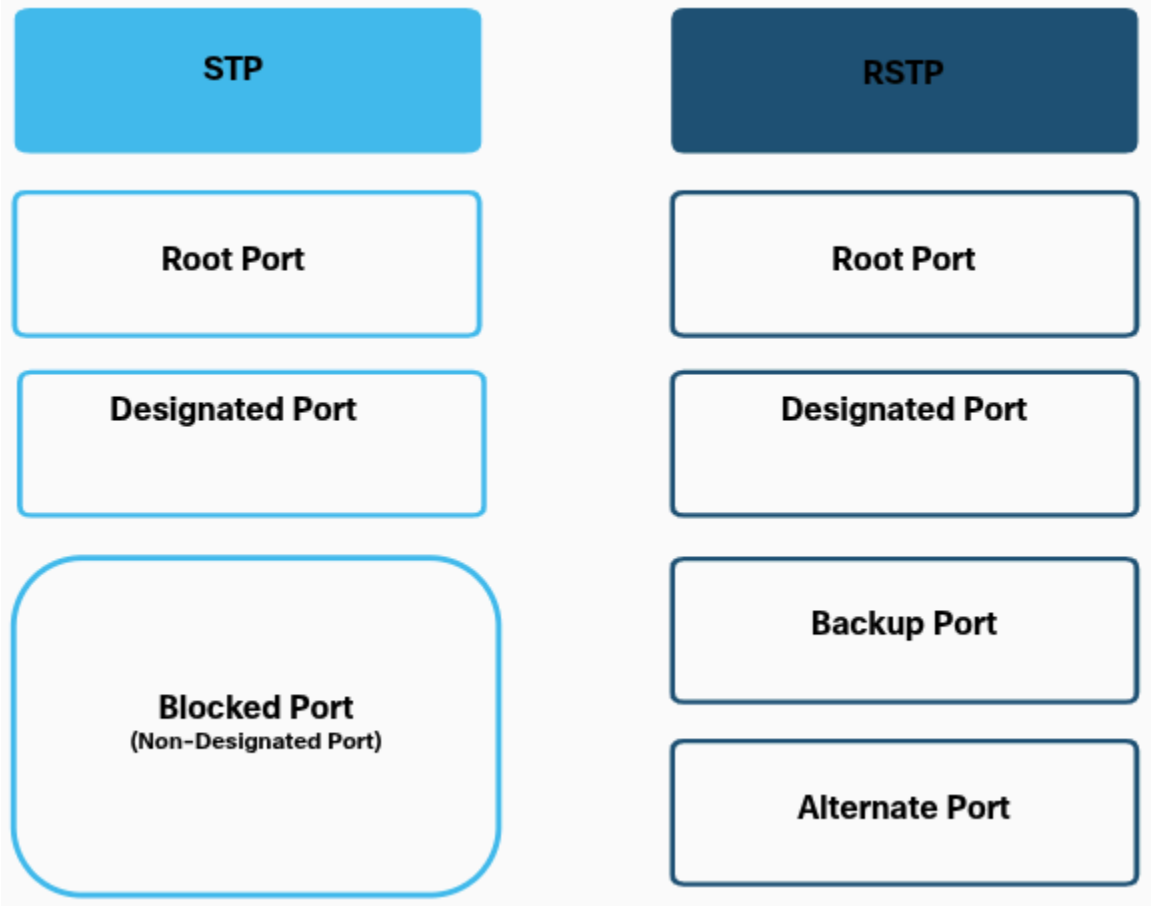

RSTP Alternate and Backup Ports

- The alternate port has an alternate path to the root bridge.
- The backup port is a backup to a shared medium, such as a hub
- A backup port is less common because hubs are now considered legacy devices

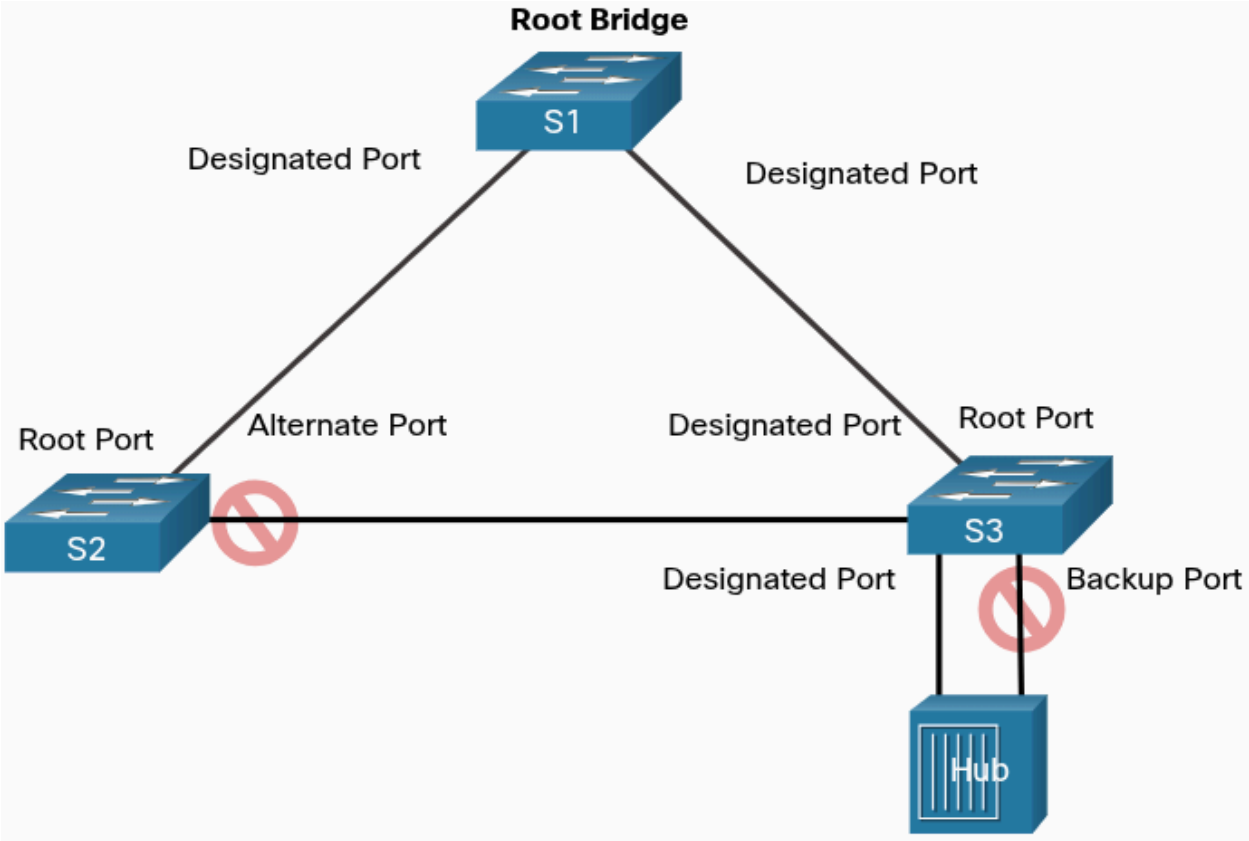

PortFast and BPDU Guard

When a device is connected to a switch port, or when a switch powers up, the switch port goes through both the listening and learning states, each time waiting for the **Forward Delay** timer to expire

- This delay is **15 seconds** for each state, listening and learning, for a total of **30 seconds**
- This delay can present a problem for DHCP clients trying to discover a DHCP server
    - DHCP messages from the connected host will **not** be forwarded for the 30 seconds of Forward Delay timers and the DHCP process may timeout.
    - The result is that an IPv4 client will not receive a valid IPv4 address

Note - Although this may occur with clients sending ICMPv6 Router Solicitation messages, the router will continue to send ICMPv6 Router Advertisement messages so the device will know how to obtain its address information

When a switch port is configured with **PortFast**, that port transitions from **blocking** to **forwarding state** immediately, bypassing the usual 802.1D STP transition states (the listening and learning states) and avoiding a 30 second delay

- You can use PortFast on access ports to allow devices connected to these ports, such as DHCP clients, to access the network immediately, rather than waiting for IEEE 802.1D STP to converge on each VLAN
    - Because the purpose of PortFast is to minimize the time that access ports must wait for spanning tree to converge, it should only be used on access ports.
- If you enable PortFast on a port connecting to another switch, you risk creating a spanning tree loop
- PortFast is **only for use** on switch ports that **connect** to **end devices**

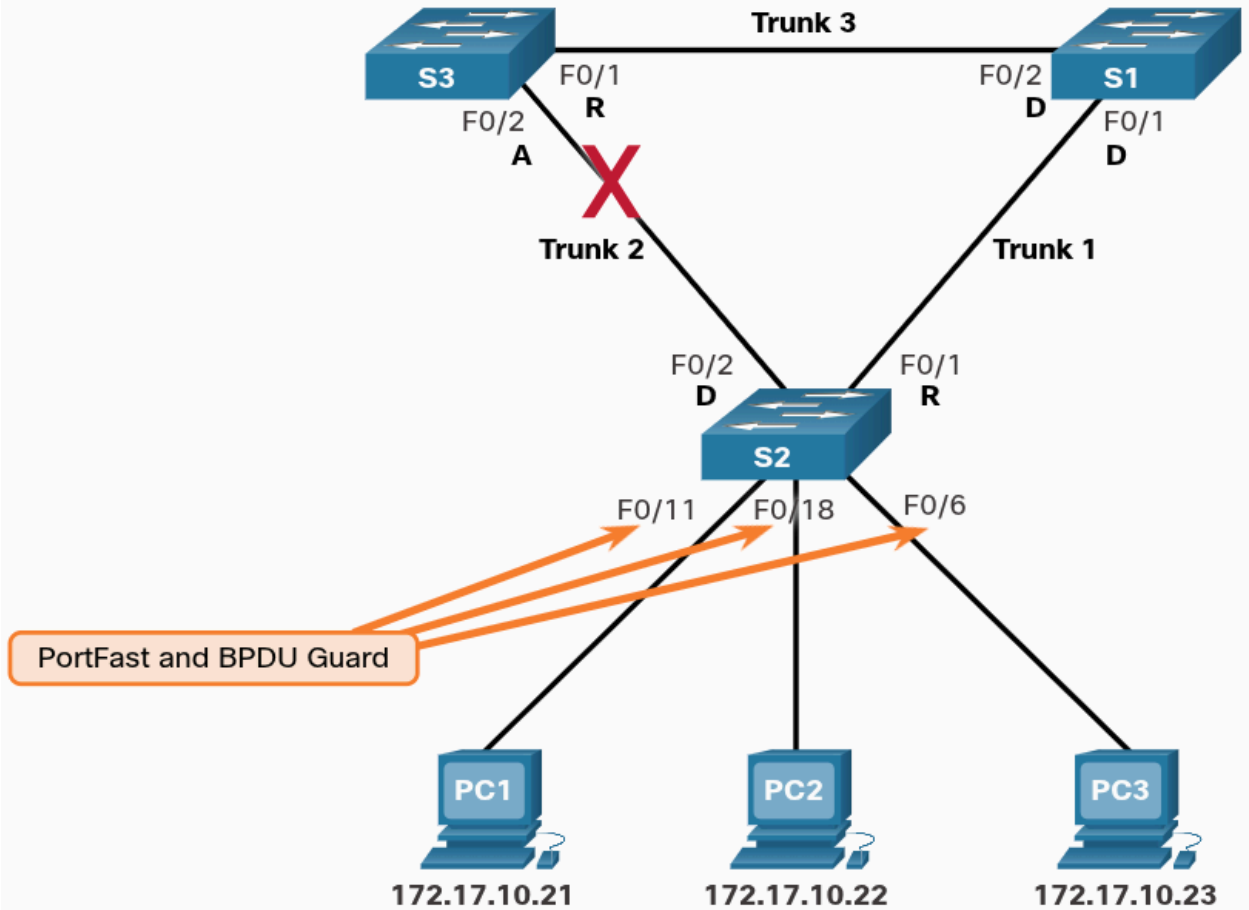

In a valid PortFast configuration, BPDUs should never be received on PortFast-enabled switch ports because that would indicate that another bridge or switch is connected to the port.

- This potentially causes a **spanning tree loop**

To Prevent this type of scenario from occurring, Cisco switches support a feature called **BPDU Guard**

- When enabled, BPDU guard immediately puts the switch port in an **errdisabled (error-disabled)** state on receipt of any BPDU
- This protects against potential loops by effectively shutting down the port

The BPDU guard feature provides a secure response to invalid configurations because an administrator must manually put the interface back into service.

Alternatives to STP

STP was and still is an Ethernet loop-prevention protocol. Ethernet LANs went from a few interconnected switches connected to a single router, to a sophisticated hierarchical network design including access, distribution and core layer switches.

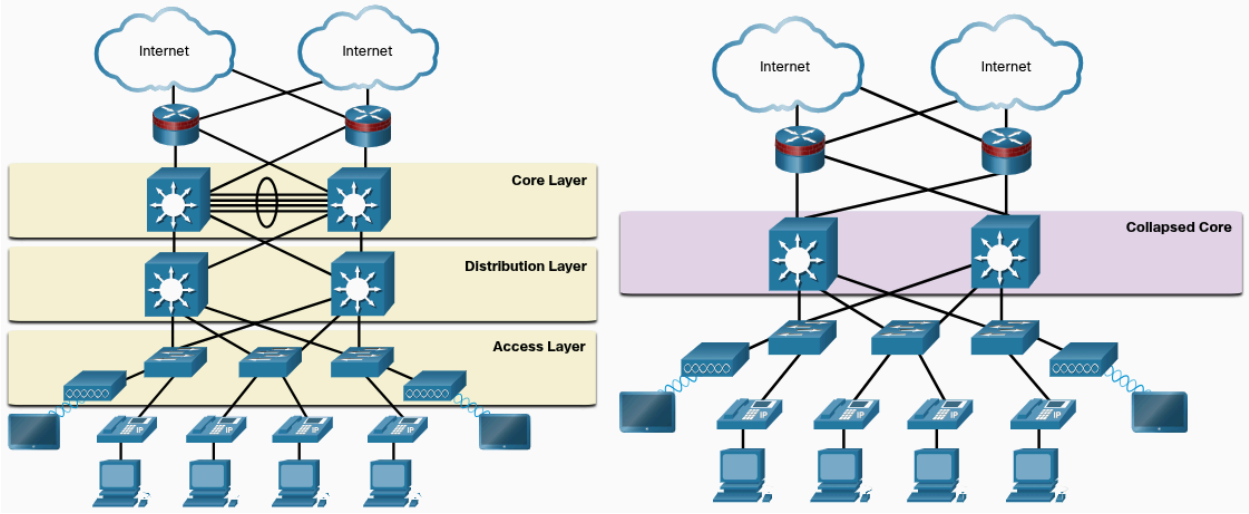

Depending on implementation, Layer 2 may include not only the access layer, but also the distribution and even the core layers.

- May include hundreds of switches, with hundreds or even thousands of VLANs
- STP has adapted, as part of RSTP and MSTP

lac

Note - An important aspect to network design is fast and predictable convergence when there is a failure or change in topology.

- Spanning tree does not offer the same efficiencies and predictabilities provided by routing protocols at Layer 3.

Figure below shows a traditional hierarchical network design with the distribution and core multilayer switches performing routing.

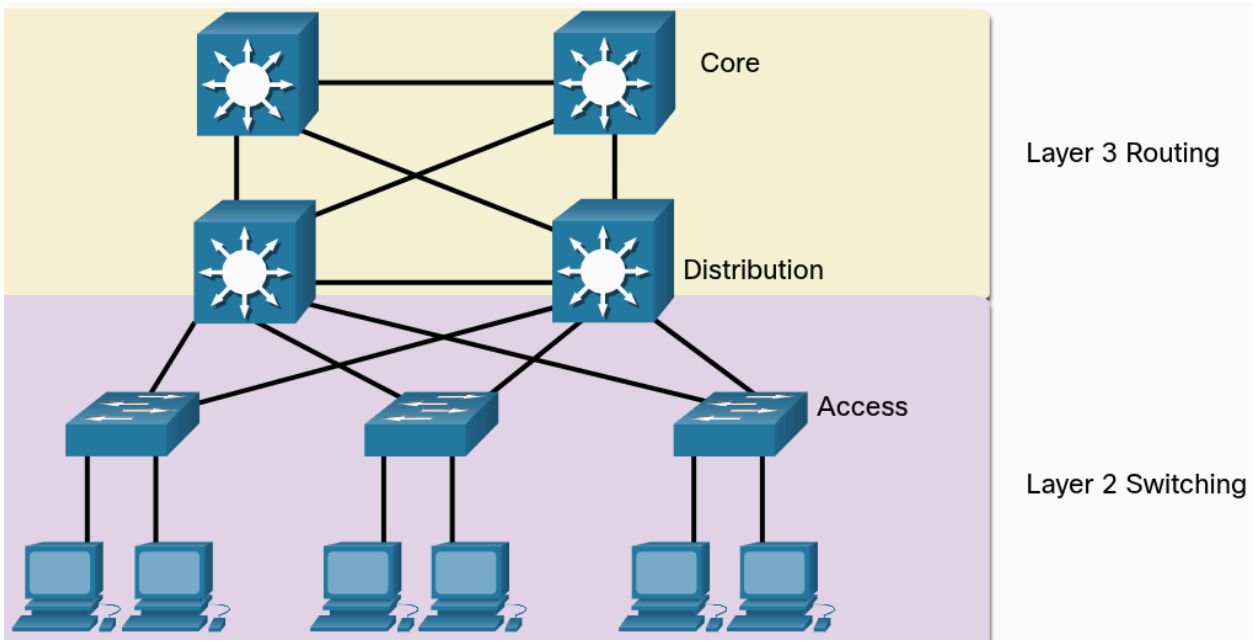

Layer 3 routing allows for redundant paths and loops in the topology, without blocking ports.

- For this reason, some environments are transitioning to Layer 3 everywhere **except** where devices **connect** to the **access layer switch**

The connections between access layer switches and distribution switches would be Layer 3 instead of Layer 2, shown in figure below.

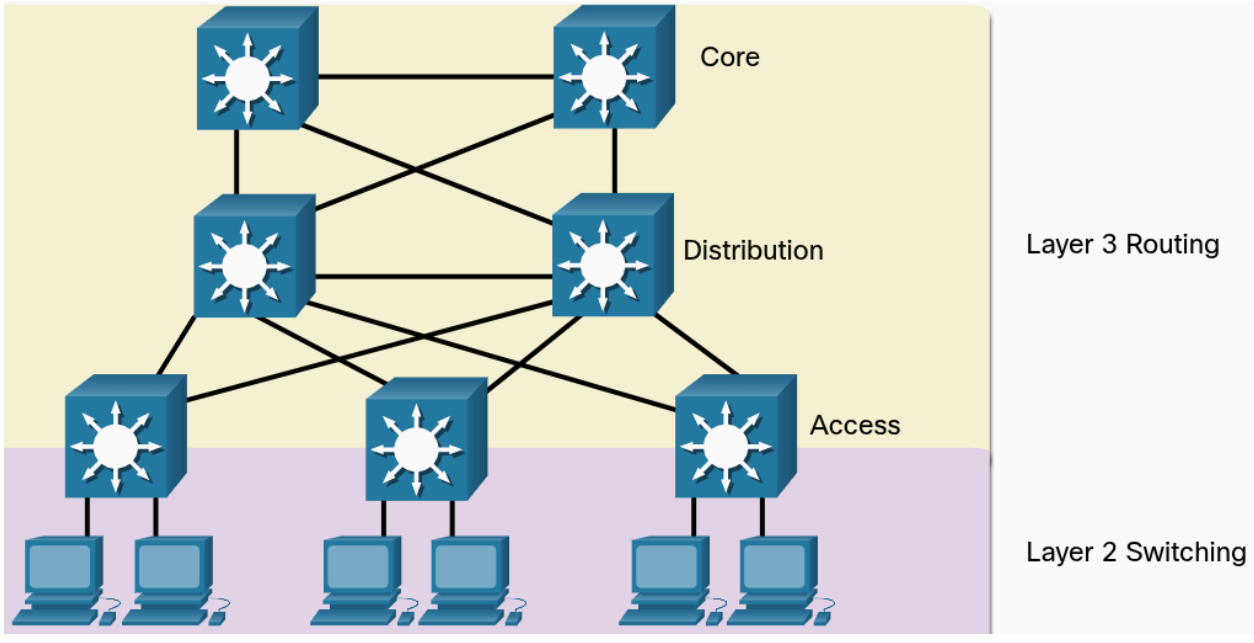

STP will most likely continue to be used as a loop prevention mechanism in enterprise, other technologies are also being used:

- Multi System Link Aggregation (MLAG)
- Shortest Path Bridging (SPB)
- Transparent Interconnect of Lots of Links (TRILL)

5.4.1 What did I learn in this module?

Purpose of STP

Redundant paths in a switched Ethernet network may cause both physical and logical Layer 2 loops. A Layer 2 loop can result in MAC address table instability, link saturation, and high CPU utilization on switches and end-devices. This results in the network becoming unusable. Unlike the Layer 3 protocols, IPv4 and IPv6, Layer 2 Ethernet does not include a mechanism to recognize and eliminate endlessly looping frames. Ethernet LANs require a loop-free topology with a single path between any two devices. STP is a loop-prevention network protocol that allows for redundancy while creating a loop-free Layer 2 topology. Without STP, Layer 2 loops can form, causing broadcast, multicast and unknown unicast frames to loop endlessly, bringing down a network. A broadcast storm is an abnormally high number of broadcasts overwhelming the network during a specific amount of time. Broadcast storms can disable a network within seconds by overwhelming switches and end devices. STP is based on an algorithm invented by Radia Perlman. Her spanning tree algorithm (STA) creates a loop-free topology by selecting a single root bridge where all other switches determine a single least-cost path.

STP Operations

Using the STA, STP builds a loop-free topology in a four-step process: elect the root bridge, elect the root ports, elect designated ports, and elect alternate (blocked) ports. During STA and STP functions, switches use BPDUs to share information about themselves and their connections. BPDUs are used to elect the root bridge, root ports, designated ports, and alternate ports. Each BPDU contains a BID that identifies the switch that sent the BPDU. The BID is involved in making many of the STA decisions including root bridge and port roles. The BID contains a priority value, the MAC address of the switch, and an extended system ID. The lowest BID value is determined by the combination of these three fields. The switch with the lowest BID will become the root bridge. Because the default BID is 32,768 it is possible for two or more switches to have the same priority. In this scenario, where the priorities are the same, the switch with the lowest MAC address will become the root bridge. When the root bridge has been elected for a

given spanning tree instance, the STA determines the best paths to the root bridge from all destinations in the broadcast domain. The path information, known as the internal root path cost, is determined by the sum of all the individual port costs along the path from the switch to the root bridge. After the root bridge has been determined the STA algorithm selects the root port. The root port is the port closest to the root bridge in terms of overall cost, which is called the internal root path cost. After each switch selects a root port, switches will select designated ports. The designated port is a port on the segment (with two switches) that has the internal root path cost to the root bridge. If a port is not a root port or a designated port, then it becomes an alternate (or backup) port. Alternate ports and backup ports are in discarding or blocking state to prevent loops. When a switch has multiple equal-cost paths to the root bridge, the switch will determine a port using the following criteria: lowest sender BID, then the lowest sender port priority, and finally the lowest sender port ID. STP convergence requires three timers: the hello timer, the forward delay timer, and the max age timer. Port states are blocking, listening, learning, forwarding, and disabled. In PVST versions of STP, there is a root bridge elected for each spanning tree instance. This makes it possible to have different root bridges for different sets of VLANs.

Evolution of STP.

The term Spanning Tree Protocol and the acronym STP can be misleading. STP is often used to refer to the various implementations of spanning tree, such as RSTP and MSTP. RSTP is an evolution of STP that provides faster convergence than STP. RSTP port states are learning, forwarding and discarding. PVST+ is a Cisco enhancement of STP that provides a separate spanning tree instance for each VLAN configured in the network. PVST+ supports PortFast, UplinkFast, BackboneFast, BPDU guard, BPDU filter, root guard, and loop guard. Cisco switches running IOS 15.0 or later, run PVST+ by default. Rapid PVST+ is a Cisco enhancement of RSTP that uses PVST+ and provides a separate instance of 802.1w per VLAN. When a switch port is configured with PortFast, that port transitions from blocking to forwarding state immediately, bypassing the STP listening and learning states and avoiding a 30 second delay. Use PortFast on access ports to allow devices connected to these ports, such as DHCP clients, to access the network immediately, rather than waiting for STP to converge on each VLAN. Cisco switches support a feature called BPDU guard which immediately puts the switch port in an error-disabled state upon receipt of any BPDU to protect against potential loops. Over the years, Ethernet LANs went from a few interconnected switches that were connected to a single router, to a sophisticated hierarchical network design. Depending on the implementation, Layer 2 may include not only the access layer, but also the distribution or even the core layers. These designs may include hundreds of switches, with hundreds or even thousands of VLANs. STP has adapted to the added redundancy and complexity with enhancements as part of RSTP and MSTP. Layer 3 routing allows for redundant paths and loops in the topology, without blocking ports. For this reason, some environments are transitioning to Layer 3 everywhere except where devices connect to the access layer switch.

EtherChannel

- Aggregates links between devices into bundles
- Bundles include redundant links.
    - STP may block one of those links, but not all of them

EtherChannel Operation

Link Aggregation

- Allows redundant links between devices that will not be blocked by STP
    - Known as EtherChannel

EtherChannel

- A link aggregation technology that groups multiple physical Ethernet links together into one single logical link
    - Used to provide fault tolerance, load sharing, increased bandwidth, and redundancy between switches, routers, and servers
- Makes it possible to combine the number of physical links between the switches to increase the overall speed of switch-to-switch communication

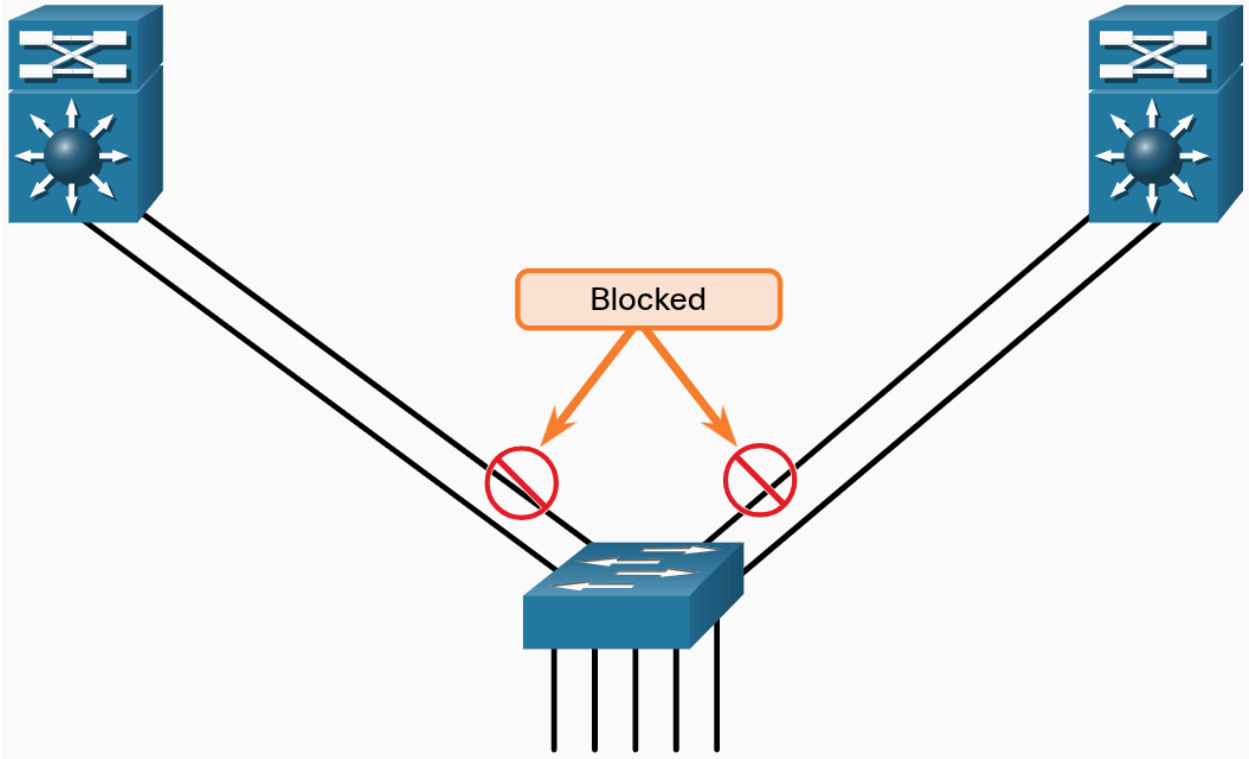

By default, STP will block redundant links

EtherChannel (cont)

- Originally developed by Cisco as LAN switch-to-switch technique of grouping several Fast Ethernet or Gigabit Ethernet ports into one logical channel
- When configured, the resulting virtual interface is called a **port channel**
- The physical interfaces are bundled together into a port channel interface

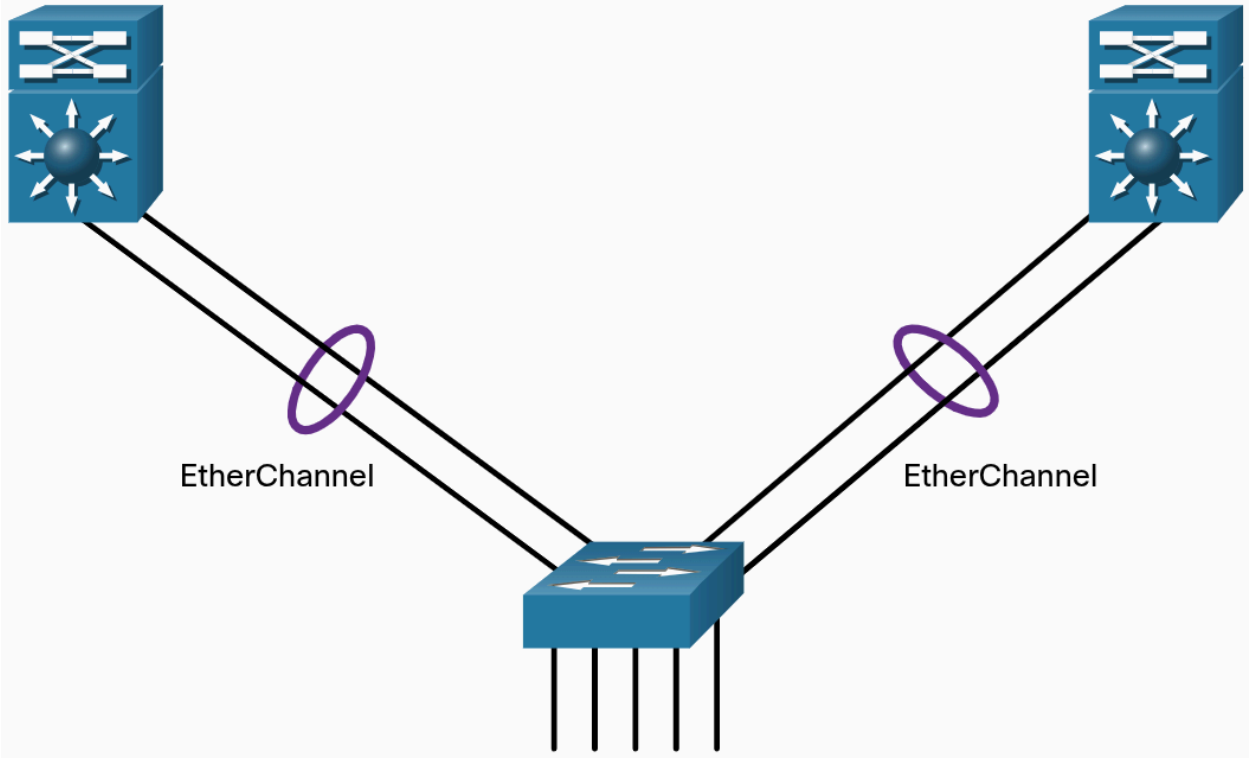

Advantages of EtherChannel

- Most **configuration tasks** can be done on the **EtherChannel interface instead** of on each **individual port**, ensuring configuration consistency throughout the links
- EtherChannel relies on **existing** switch ports. No need to upgrade the link to a faster and more expensive connection to have **more bandwidth**
- **Load balancing** takes place between links that are part of the same EtherChannel
    - Depending on hardware platform, one or more load-balancing methods can be implemented
    - These methods include **source MAC** and **destination MAC** load balancing, or **source IP** and **destination IP** load balancing, across the physical links
- EtherChannel creates an aggregation that is seen as **one** logical link.
    - When **several** EtherChannel **bundles** exist between two switches, **STP** may **block** one of the bundles to prevent switching loops.
    - When STP blocks one of the redundant links, it blocks the entire EtherChannel
        - This blocks all the ports belonging to that EtherChannel link. Where there is only one EtherChannel link, all physical links in the EtherChannel are active because STP sees only one (logical) link.

- EtherChannel provides **redundancy** because the overall link is seen as one logical connection.
    - The loss of one physical link within the channel does **not** create a change in topology
        - Therefore, a spanning tree recalculation is not required
    - Assuming at least one physical link is present, the EtherChannel remains functional, even if its overall throughput decreases because of a lost link within the EtherChannel

Implementation Restrictions

Implementation on Catalyst 2960

- Interface types **cannot** be mixed
    - Example: Fast Ethernet and Gigabit Ethernet **cannot** be mixed within a single EtherChannel
- Currently each EtherChannel can consist of up to **eight** compatibly-configured Ethernet ports. EtherChannel provides **full-duplex** bandwidth up to **800 Mbps** (Fast EtherChannel) or **8 Gbps** (Gigabit EtherChannel) between one switch and another switch or host.
- The Cisco catalyst 2960 Layer 2 switch currently supports up to **six** EtherChannels.
    - However, as new IOSs are developed and platforms change, some cards and platforms may support increased numbers of ports within an EtherChannel link, as well as support an increased number of Gigabit EtherChannels.
- The Individual EtherChannel Group member port confirmation **must be consistent** on **both** devices
    - If the physical ports of one side are configured as trunks, the physical ports of the other side **must also** be configured as trunks within the same native VLAN
    - All ports in each EtherChannel link **must be** configured as Layer 2 ports
- Each EtherChannel has a logical port channel interface. A configuration applied to the port channel interface affects **all** physical interfaces that are assigned to that interface.

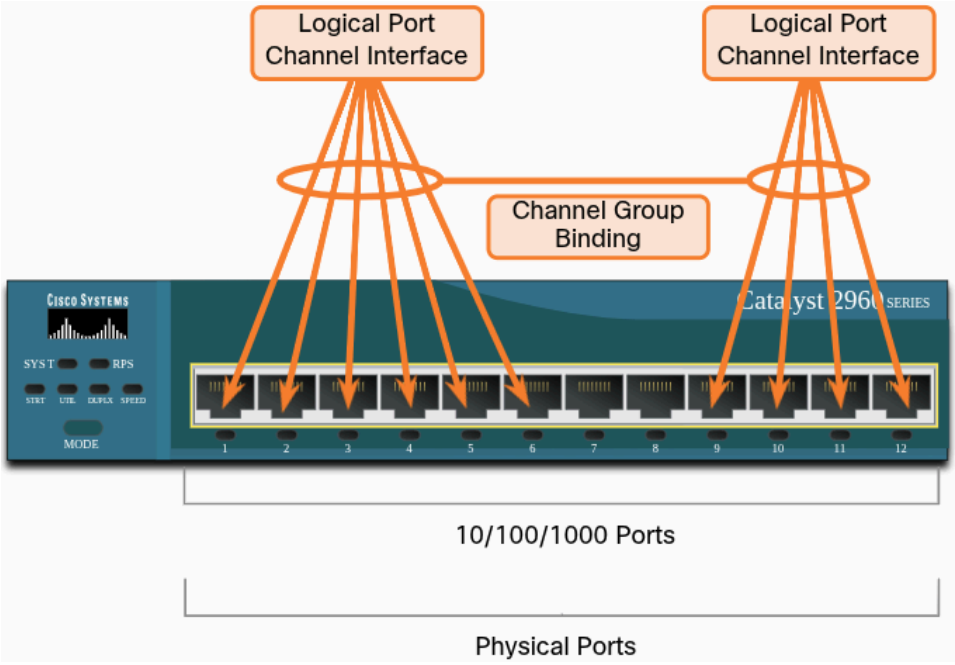

AutoNegotiation Protocols

EtherChannels can be formed through negotiation using one of two protocols

- Port Aggregation Protocol (PAgP)
- Link Aggregation Control Protocol (LACP)

These protocols allow ports with similar characteristics to form a channel through **dynamic negotiation** with adjoining switches.

Note - It is also possible to configure a static or unconditional EtherChannel without PAgP or LACP

PAgP Operation

(Pronounced “Pag - P”) - Cisco proprietary protocol that aids in the **automatic** creation of EtherChannel links.

- When used to configure EtherChannel, PAgP packets are sent between EtherChannel-capable ports to negotiate the forming of a channel.
    - When PAgP identifies matched Ethernet links, it groups the links into an EtherChannel
    - EtherChannel is then added to the spanning tree as a single port.
- When enabled, PAgP also managed the EtherChannel

PAgP packets are sent every **30 seconds**

PAgP checks for configuration consistency and manages link additions and failures between two switches. It ensures that when an EtherChannel is created, all ports have the same type of configuration

Note - In EtherChannel, it is mandatory that **all ports** have the **same speed**, **duplex setting**, and **VLAN information**. Any port-channel modification after the creation of the channel also changes the aggregated channel ports

PAgP helps create the EtherChannel link by detecting the configuration of each side and ensuring that the links are compatible so that the EtherChannel link can be enabled when needed. Modes for PAgP:

- **On** - Forces the interface to channel without PAgP
    - Interfaces configured in the on mode do **not** exchange PAgP packets
- **PAgP desirable** - Places an interface in an **active negotiating state** in which the interface initiates negotiations with other interfaces by sending PAgP packets
- **PAgP auto** - Places an interface in a **passive negotiating state** in which the interface responds to the PAgP packets that it **receives** but does **not** initiate PAgP negotiation

The modes **must be compatible** on each side

- If one side is configured to be in auto mode, it is placed in **passive state**, waiting for the other side to initiate the EtherChannel negotiation
- If the other side is also set to auto, the negotiation **never starts** and the EtherChannel does not form
- If all modes are **disabled** by using the *no* command, or if **no mode** is configured, then the EtherChannel is **disabled**

The **on mode** manually places the interface in an EtherChannel, without any negotiation.

- It works **only** if the other side is also set to on.
- If the other side is set to negotiate parameters through PAgP, **no EtherChannel forms**, because the side that is set to on mode **does not** negotiate

No negotiation between the two switches means there is no checking to make sure that all the links in the EtherChannel are terminating on the other side, or that there is PAgP compatibility on the other switch.

PAgP Mode Settings Example

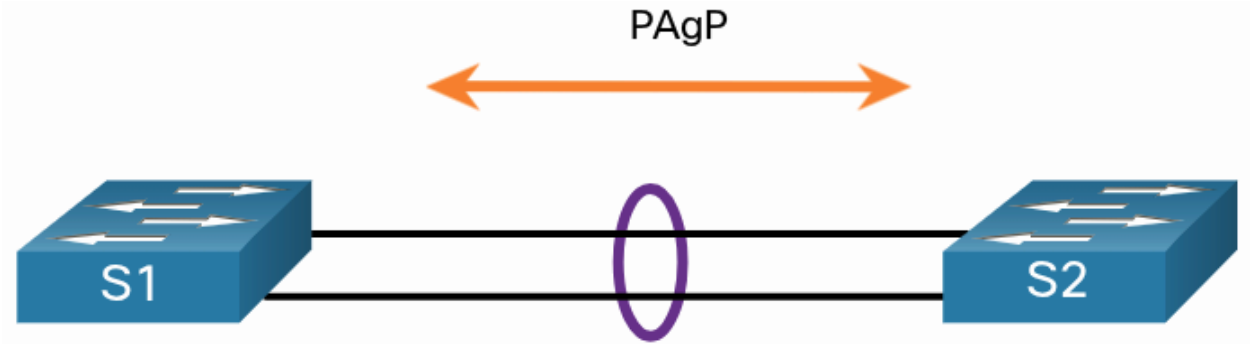

PAgP Modes		
S1	**S2**	Channel Establishment
On	On	Yes
On	Desirable/ Auto	No
Desirable	Desirable	Yes
Desirable	Auto	Yes
Auto	Desirable	Yes
Auto	Auto	No

LACP Operation

Part of an IEEE specification (802.3ad) that allows several physical ports to be bundled to form a single logical channel

- Allows a switch to negotiate an automatic bundle by sending LACP packets to other switch. Performs a function similar to PAgP with Cisco EtherChannel

Because LACP is an IEEE standard, it can be used to facilitate EtherChannels in **multivendor environments**

- On Cisco devices, both protocols are supported

Note - LACP was originally defined as **IEEE 802.3ad**. However, LACP is now defined in the newer **IEEE 802.1AX** standard for local and metropolitan area networks

LACP provides the **same negotiation** as PAgP.

- LACP helps create the EtherChannel link by detecting the configuration of each side and making sure that they are compatible. Modes for LACP:
    - **On** - Forces the interface to channel **without** LACP.
        - Interfaces in the on mode do **not** exchange LACP packets
    - **LACP active** - Places a port in an active negotiating state.
        - The port initiates negotiation with other ports by sending LACP packets
    - **LACP passive** - Places a port in a passive negotiating state.
        - The port responds to the LACP packets that it receives but does **not** initiate LACP packet negotiation

Just as with PAgP, modes **must** be compatible on both sides for the EtherChannel link to form. The on mode is repeated, because it creates the EtherChannel configuration unconditionally, without PAgP or LACP dynamic negotiation

LACP allows for **eight active links**, and also **eight standby links**

- A standby link will become active should one of the current active link **fails**

LACP Mode Settings Example

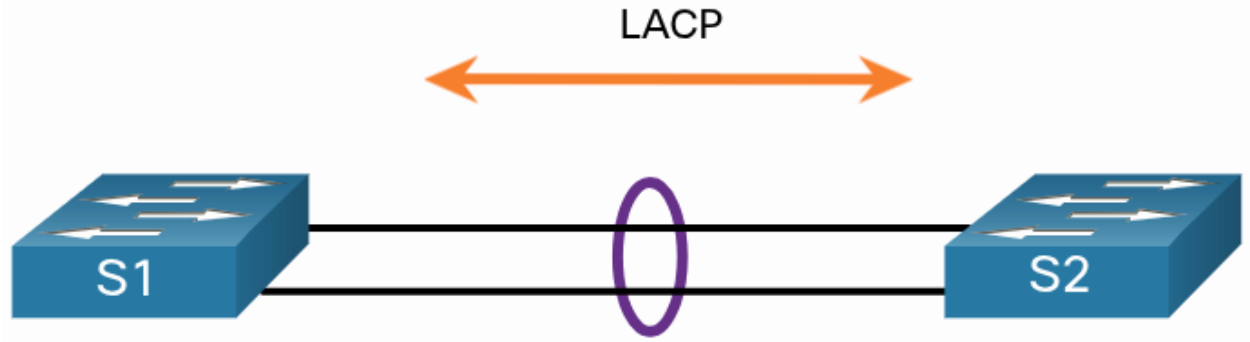

LACP Modes		
S1	**S2**	Channel Establishment
On	On	Yes
On	Active/ Passive	No
Active	Active	Yes
Active	Passive	Yes
Passive	Active	Yes
Passive	Passive	No

Configure EtherChannel

Guidelines

- **EtherChannel Support** - All Ethernet interfaces **must** support EtherChannel with **no** requirement that interfaces be **physically contiguous**
- **Speed and Duplex** - Configure all interfaces in an EtherChannel to operate at the **same speed** and in the **same duplex mode**
- **VLAN match** - All interfaces in the EtherChannel bundle **must** be assigned to the **same VLAN** or to be configured as a **trunk**
- **Range of VLANs** - An EtherChannel supports the **same allowed range of VLANs** on **all** interfaces in a **trunking** EtherChannel
    - IF the allowed range of VLANs is **not** the same, the interfaces do **not** form an EtherChannel, even when they are set to **auto** or **desirable** mode
- **Three interface parameters** much match
    - **Trunking mode**, **Native** VLANs, **Allowed** VLANs

Figure shows configuration that would allow an EtherChannel to form between S1 and S2

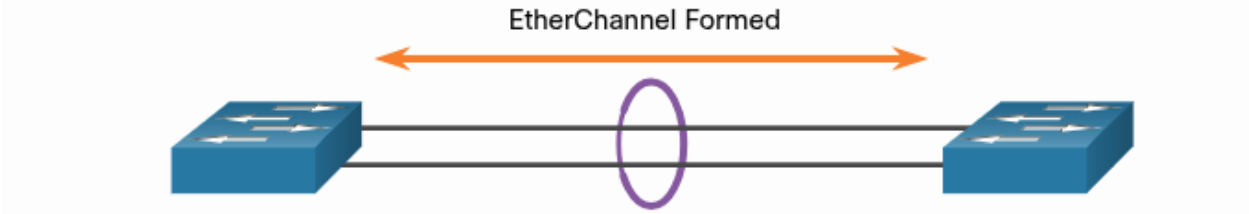

S1 Port Configurations	
Speed	1 Gbps
Duplex	Full
VLAN	10

S2 Port Configurations	
Speed	1 Gbps
Duplex	Full
VLAN	10

An EtherChannel is formed when configuration settings match on both switches

*Next figure, S1 ports are configured as **half duplex**. Therefore, an EtherChannel will **not** form between S1 and S2*

S1 Port Configurations	
Speed	1 Gbps
Duplex	Half
VLAN	20

S2 Port Configurations	
Speed	1 Gbps
Duplex	Full
VLAN	20

*An EtherChannel is **not** formed when configuration settings are **different** on each switch*

If these settings must be changed:

- Configure them in port channel interface configuration mode
  - Any configuration that is applied to the port channel interface also **affects individual** interfaces
  - However, configurations that are applied to the individual interfaces do **not** affect the port channel interface
  - Therefore, making configuration changes to an interface that is part of an EtherChannel link may **cause** interface compatibility issues

The port channel can be configured in access mode, trunk mode (most common), or on a routed port

LACP Configuration Example

EtherChannel is disabled by default and must be configured

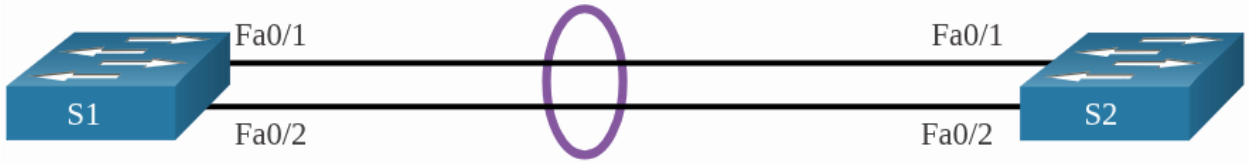

Configuring EtherChannel with LACP requires three steps:

1. Specify the interfaces that compose the EtherChannel group using:
   a. **interface range** *interface* command (global configuration)
   b. The **range** keyword allows you to **select several** interfaces and configure them all together
2. Create the port channel interface with:
   a. **channel-group** *identifier* **mode active** command (interface range configuration)
   b. The identifier specifies a channel group number. The **mode active** keyword identify this as an LACP EtherChannel configuration
3. To change Layer 2 settings on the port channel interface:
   a. Enter port channel interface config mode: **interface port-channel** command– followed by the *interface identifier*

```
S1(config)# interface range FastEthernet 0/1 - 2
S1(config-if-range)# channel-group 1 mode active
Creating a port-channel interface Port-channel 1
S1(config-if-range)# exit
S1(config)# interface port-channel 1
S1(config-if)# switchport mode trunk
S1(config-if)# switchport trunk allowed vlan 1,2,20
```

Enter interface range mode for FastEthernet0/1 and FastEthernet0/2. Use fa 0/1 - 2 as the interface designation.

```
S1(config)#interface range fa 0/1 - 2
```

Use context sensitive help (?) to display the options for the channel-group command.

```
S1(config-if-range)#channel-group ?
<1-6> Channel group number
```

Select channel-group 1 and display the next option.

```
S1(config-if-range)#channel-group 1 ?
mode Etherchannel Mode of the interface
```

Enter the mode keyword and display the next set of options.

```
S1(config-if-range)#channel-group 1 mode ?
active     Enable LACP unconditionally
auto       Enable PAgP only if a PAgP device is   detected
desirable  Enable PAgP unconditionally
on         Enable Etherchannel only
passive    Enable LACP only if a LACP device is detected
```

Configure the channel-group to use LACP unconditionally.

```
S1(config-if-range)#channel-group 1 mode active
*Mar 21 00:02:28.184: %LINEPROTO-5-UPDOWN: Line protocol on Interface
FastEthernet0/1, changed state
to down
*Mar 21 00:02:28.193: %LINEPROTO-5-UPDOWN: Line protocol on Interface
FastEthernet0/2, changed state
to down
*Mar 21 00:02:36.179: %LINEPROTO-5-UPDOWN: Line protocol on Interface
FastEthernet0/1, changed state
to up
*Mar 21 00:02:36.674: %LINEPROTO-5-UPDOWN: Line protocol on Interface
FastEthernet0/2, changed state
to up
*Mar 21 00:04:31.170: %LINEPROTO-5-UPDOWN: Line protocol on Interface
FastEthernet0/1, changed state
to down
```

```
*Mar 21 00:04:31.186: %LINEPROTO-5-UPDOWN: Line protocol on Interface
FastEthernet0/2, changed state
to down
*Mar 21 00:04:33.116: %LINEPROTO-5-UPDOWN: Line protocol on Interface
FastEthernet0/1, changed state
to up
*Mar 21 00:04:34.114: %LINK-3-UPDOWN: Interface Port-channel1, changed
state to up
*Mar 21 00:04:35.037: %LINEPROTO-5-UPDOWN: Line protocol on Interface
FastEthernet0/2, changed state
to up
*Mar 21 00:04:35.121: %LINEPROTO-5-UPDOWN: Line protocol on Interface
Port-channel1, changed state to
up
```

Configure the switchport settings for the port-channel that was created:

- Enter interface configuration mode for port-channel 1
- Configure port-channel 1 as a trunk
- Allow VLANS 1,2,20 to cross the trunk link. Enter the VLANs as shown with no spaces.

```
S1(config-if-range)#interface port-channel 1
S1(config-if)#switchport mode trunk
S1(config-if)#switchport trunk allowed vlan 1,2,20
```

You have successfully configured EtherChannel.

Verify and Troubleshoot EtherChannel

Verify EtherChannel

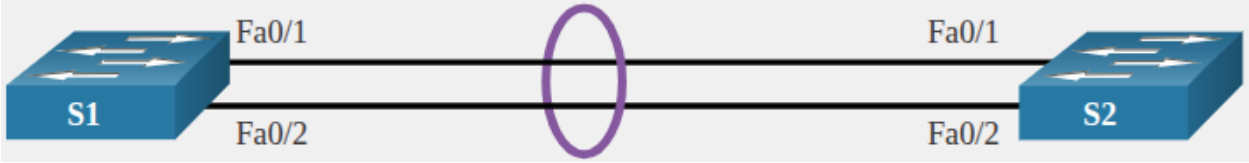

show interfaces port-channel

- Displays the general status of the port channel interface

```
S1# show interfaces port-channel 1
Port-channel1 is up, line protocol is up (connected)
  Hardware is EtherChannel, address is c07b.bcc4.a981 (bia c07b.bcc4.a981)
  MTU 1500 bytes, BW 200000 Kbit/sec, DLY 100 usec,
     reliability 255/255, txload 1/255, rxload 1/255
(output omitted)
```

show etherchannel summary

- Use to display one line of information per port channel
    - When several port channel interfaces are configured on the same device

The interface bundle consists of the Fa0/1 and Fa0/2 interfaces.
This group is a Layer 2 EtherChannel and is in use

```
S1# show etherchannel summary
Flags:  D - down        P - bundled in port-channel
        I - stand-alone s - suspended
        H - Hot-standby (LACP only)
        R - Layer3      S - Layer2
        U - in use      N - not in use, no aggregation
        f - failed to allocate aggregator
        M - not in use, minimum links not met
        m - not in use, port not aggregated due to minimum links not met
        u - unsuitable for bundling
        w - waiting to be aggregated
        d - default port
        A - formed by Auto LAG
Number of channel-groups in use: 1
Number of aggregators:           1
Group  Port-channel  Protocol    Ports
------+-------------+-----------+-----------------------------------------------
1      Po1(SU)         LACP      Fa0/1(P)    Fa0/2(P)
```

show etherchanel port-channel

- Display information about a specific port channel interface

The Port Channel 1 interface consists of two physical interfaces, Fa0/1 and Fa0/2.
It uses LACP in **active mode**.

- It is properly connected to another switch with a compatible configuration, which is why the port channel is said to be in use

```
S1# show etherchannel port-channel
                Channel-group listing:
                ----------------------
Group: 1
----------
                Port-channels in the group:
                ---------------------------
Port-channel: Po1    (Primary Aggregator)
------------
Age of the Port-channel   = 0d:01h:02m:10s
Logical slot/port   = 2/1          Number of ports = 2
HotStandBy port = null
Port state          = Port-channel Ag-Inuse
Protocol            =   LACP
Port security       = Disabled
Load share deferral = Disabled
```

```
Ports in the Port-channel:
Index   Load   Port     EC state        No of bits
------+------+------+------------------+-----------
  0     00     Fa0/1    Active             0
  0     00     Fa0/2    Active             0
Time since last port bundled:    0d:00h:09m:30s    Fa0/2
```

show interfaces etherchannel

* On any physical interface member of an EtherChannel bundle
    - Provides information about the role of the interface in the EtherChannel

Fa0/1 is part of the EtherChannel bundle 1.
The protocol is LACP.

```
S1# show interfaces f0/1 etherchannel
Port state    = Up Mstr Assoc In-Bndl
Channel group = 1           Mode = Active          Gcchange = -
Port-channel  = Po1         GC   =   -             Pseudo port-channel = Po1
Port index    = 0           Load = 0x00            Protocol =   LACP
Flags:  S - Device is sending Slow LACPDUs   F - Device is sending fast LACPDUs.
       A - Device is in active mode.        P - Device is in passive mode.
Local information:
                         LACP port     Admin     Oper    Port
Port      Flags   State     Priority      Key       Number  State
Fa0/1     SA      bndl      32768         0x1       0x1     0x102       0x3D
Partner's information:
                 LACP port                        Admin  Oper   Port    Port
Port      Flags   Priority  Dev ID          Age    key    Key    Number  State
Fa0/1     SA      32768     c025.5cd7.ef00  12s    0x0    0x1    0x102   0x3Dof the
port in the current state: 0d:00h:11m:51sllowed vlan 1,2,20
```

Common Issues with EtherChannel Configurations

All interfaces within an EtherChannel **must** have the same configuration of **speed** and **duplex mode**, **native** and **allowed VLANs on trunks**, and **access VLAN** on **access ports**.

Common EtherChannel issues:

- Assigned Ports in the EtherChannel are **not** part of the same VLAN, or **not** configured as **trunks**. Ports with **different native** VLANs **cannot** form an EtherChannel
- Trunking was configured on some of the ports that make up the EtherChannel, but **not all of them**
    - Not recommended that you configure trunking mode on **individual** ports that make up the EtherChannel
    - When configuring a trunk on EtherChannel, **verify** the **trunking mode** on EtherChannel
- If the allowed range of VLANs is **not** the same, the ports do **not** form an EtherChannel even when PAgP is set to **auto** or **desirable**
- The dynamic negotiation options for PAgP and LACP are **not compatibility configured** on both ends of the EtherChannel\

Note - It is easy to configure PAgP or LACP with DTP

- They are all protocols used to automate behavior on trunk links
- PAgP and LACP are used for **link aggregation (EtherChannel)**
- DTP is used for **automating the creation of trunk links**

** When EtherChannel trunk is configured, typically EtherChannel (PAgP or LACP) is configured **first** and **then** DTP

Troubleshoot EtherChannel Example

Fa0/1 and Fa0/2 on switches S1 and S2 are connected with an EtherChannel. However, the EtherChannel is not operational.

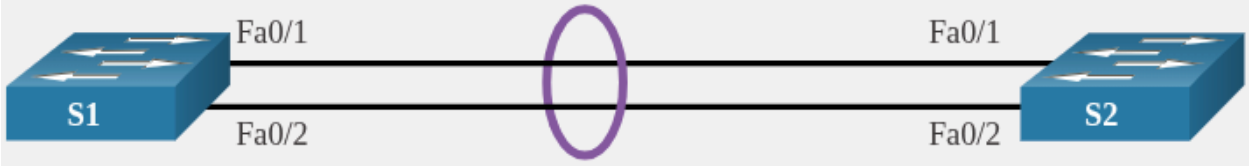

Step 1. View the EtherChannel Summary Information

show etherchannel summary indicates that EtherChannel is down

```
S1# show etherchannel summary
Flags:  D - down        P - bundled in port-channel
        I - stand-alone s - suspended
        H - Hot-standby (LACP only)
        R - Layer3      S - Layer2
        U - in use      N - not in use, no aggregation
        f - failed to allocate aggregator
        M - not in use, minimum links not met
        m - not in use, port not aggregated due to minimum links not met
        u - unsuitable for bundling
        w - waiting to be aggregated
        d - default port
        A - formed by Auto LAG
Number of channel-groups in use: 1
Number of aggregators:           1
Group  Port-channel  Protocol    Ports
------+-------------+-----------+-----------------------------------------------
1      Po1(SD)         -      Fa0/1(D)    Fa0/2(D)
```

Step 2. View Port Channel Configuration

show run | begin interface port-channel

- More detailed output indicates that there are incompatible PAgP modes configured onS1 and S2

```
S1# show run | begin interface port-channel
interface Port-channel1
 switchport trunk allowed vlan 1,2,20
 switchport mode trunk
!
interface FastEthernet0/1
 switchport trunk allowed vlan 1,2,20
 switchport mode trunk
 channel-group 1 mode on
!
interface FastEthernet0/2
 switchport trunk allowed vlan 1,2,20
 switchport mode trunk
 channel-group 1 mode on
!====================================
S2# show run | begin interface port-channel
interface Port-channel1
 switchport trunk allowed vlan 1,2,20
 switchport mode trunk
!
interface FastEthernet0/1
 switchport trunk allowed vlan 1,2,20
 switchport mode trunk
 channel-group 1 mode desirable
!
interface FastEthernet0/2
 switchport trunk allowed vlan 1,2,20
 switchport mode trunk
 channel-group 1 mode desirable
```

Step 3. Correct the Misconfiguration

To correct the issue, the PAgP mode on the EtherChannel is changed to **desirable**

Note - EtherChannel and STP must **interoperate**

- The order in which EtherChannel-related commands are entered is important, which is why you see interface Port-Channel 1 **removed and then re-added** with **channel-group** command, as **opposed to directly changed**
- If one tries to change the configuration directly, STP errors cause the associated ports to go into **blocking** or **errdisabled state**

```
S1(config)# no interface port-channel 1
S1(config)# interface range fa0/1 - 2
S1(config-if-range)# channel-group 1 mode desirable
Creating a port-channel interface Port-channel 1
S1(config-if-range)# no shutdown
S1(config-if-range)# exit
S1(config)# interface port-channel 1
S1(config-if)# switchport mode trunk
S1(config-if)# end
S1#
```

6.4.3 What did I learn in this module?

EtherChannel Operation

To increase bandwidth or redundancy, multiple links could be connected between devices. However, STP will block redundant links to prevent switching loops. EtherChannel is a link aggregation technology that allows redundant links between devices that will not be blocked by STP. EtherChannel groups multiple physical Ethernet links together into one single logical link. It provides fault-tolerance, load sharing, increased bandwidth, and redundancy between switches, routers, and servers. When an EtherChannel is configured, the resulting virtual interface is called a port channel. EtherChannel has several advantages, as well as some restrictions to implementation. EtherChannels can be formed through negotiation using one of two protocols, PAgP or LACP. These protocols allow ports with similar characteristics to form a channel through dynamic negotiation with adjoining switches. Whthren an EtherChannel link is configured using Cisco-proprietary PAgP, PAgP packets are sent between EtherChannel-capable ports to negotiate the forming of a channel. Modes for PAgP are On, PAgP desirable, and PAgP auto. LACP performs a function similar to PAgP with Cisco EtherChannel. Because LACP is an IEEE standard, it can be used to facilitate EtherChannels in multivendor environments. Modes for LACP are On, LACP active, and LACP passive.

Configure EtherChannel

The following guidelines and restrictions are useful for configuring EtherChannel:

- EtherChannel support - All Ethernet interfaces on all modules must support EtherChannel with no requirement that interfaces be physically contiguous, or on the same module.
- Speed and duplex - Configure all interfaces in an EtherChannel to operate at the same speed and in the same duplex mode.

- VLAN match - All interfaces in the EtherChannel bundle must be assigned to the same VLAN or be configured as a trunk.
- Range of VLANs - An EtherChannel supports the same allowed range of VLANs on all the interfaces in a trunking EtherChannel.

Configuring EtherChannel with LACP requires three steps:

Step 1. Specify the interfaces that compose the EtherChannel group using the interface range interface global configuration mode command.

Step 2. Create the port channel interface with the channel-group identifier mode active command in interface range configuration mode.

Step 3. To change Layer 2 settings on the port channel interface, enter port channel interface configuration mode using the interface port-channel command, followed by the interface identifier.

Verify and Troubleshoot EtherChannel.

There are a number of commands to verify an EtherChannel configuration including show interfaces port-channel, show etherchannel summary, show etherchannel port-channel, and show interfaces etherchannel. Common EtherChannel issues include the following:

- Assigned ports in the EtherChannel are not part of the same VLAN, or not configured as trunks. Ports with different native VLANs cannot form an EtherChannel.
- Trunking was configured on some of the ports that make up the EtherChannel, but not all of them.
- If the allowed range of VLANs is not the same, the ports do not form an EtherChannel even when PAgP is set to the auto or desirable mode.
- The dynamic negotiation options for PAgP and LACP are not compatibly configured on both ends of the EtherChannel.

DHCPv4

- Dynamic Host Configuration Protocol (DHCP) dynamically assigns IP address
- DHCPv4 is for IPv4 network

Concepts

DHCPv4 Server and Client

DHCPv4 Server

- Scalable
- In small branch or SOHO, Cisco router can be configured to provide DHCPv4 services without dedicated server
- Dynamically assigns, or leases, IPv4 address from pool of addresses for limited period of time chosen by server or until client no longer needs the address

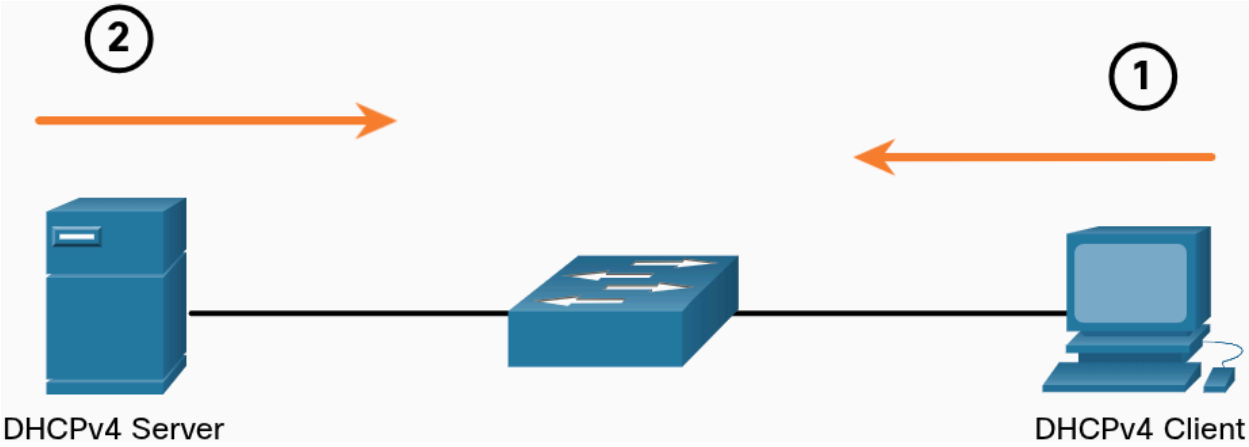

1. *The DHCPv4 lease process begins with the client sending a message requesting the services of a DHCP server*
2. *If there is a DHCPv4 server that receives the message, it will respond with an IPv4 address and possible other network configuration information*

DHCPv4 Operation

- Works in client/ server mode
- Client communicate with DHCPv4 server
    - Server assigns or leases an IPv4 address
- Client connects to the network with leased IPv4 address until lease expires
- Client must contact DHCP server periodically to extend lease
- Lease mechanism ensures clients that move or power off do not keep addresses that they no longer need
- When lease expires, DHCP server returns address to pool to be reallocated

Steps to Obtain a lease

1. DHCP Discover (DHCPDISCOVER)
2. DHCP Offer (DHCPOFFER)
3. DHCP Request (DHCPREQUEST)
4. DHCP Acknowledgement (DHCPACK)

DHCP Discover (DHCPDISCOVER)

- Purpose of DHCPDISCOVER message– to find DHCPv4 servers on network
- Client starts process using a **broadcast DHCPDISCOVER** message with its own MAC address to discover available DHCPv4 servers
- Client uses Layer 2 and Layer 3 broadcast addresses to communicate

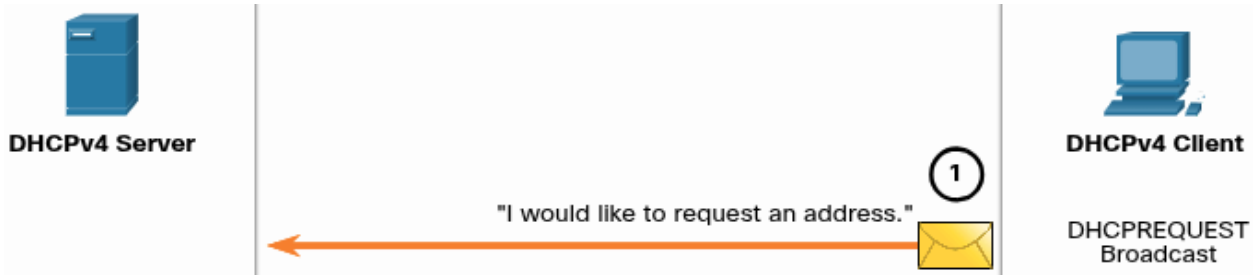

DHCP Offer (DHCPOFFER)

- DHCPv4 server receives DHCPDISCOVER message
    - Reserves available IPv4 address to lease to client
- Creates an **ARP entry** consisting of the MAC address of client and leased IPv4 address
- DHCPv4 server sends binding DHCPOFFER message to client

DHCP Request (DHCPREQUEST)

- Client **receives DHCPOFFER** from server
    - Client **send back DHCPREQUEST** message
    - Message is used for both **lease origination** and **lease renewal**
- When used for lease origination
    - DHCPREQUEST serves as binding acceptance notice to selected server for the parameters it has offered and and **implicit decline** to any **other** servers that may have provided client a binding offer
- DHCPREQUEST message is sent in form of a **broadcast** to inform DHCPv4 server and any other DHCPv4 servers about the accepted offer

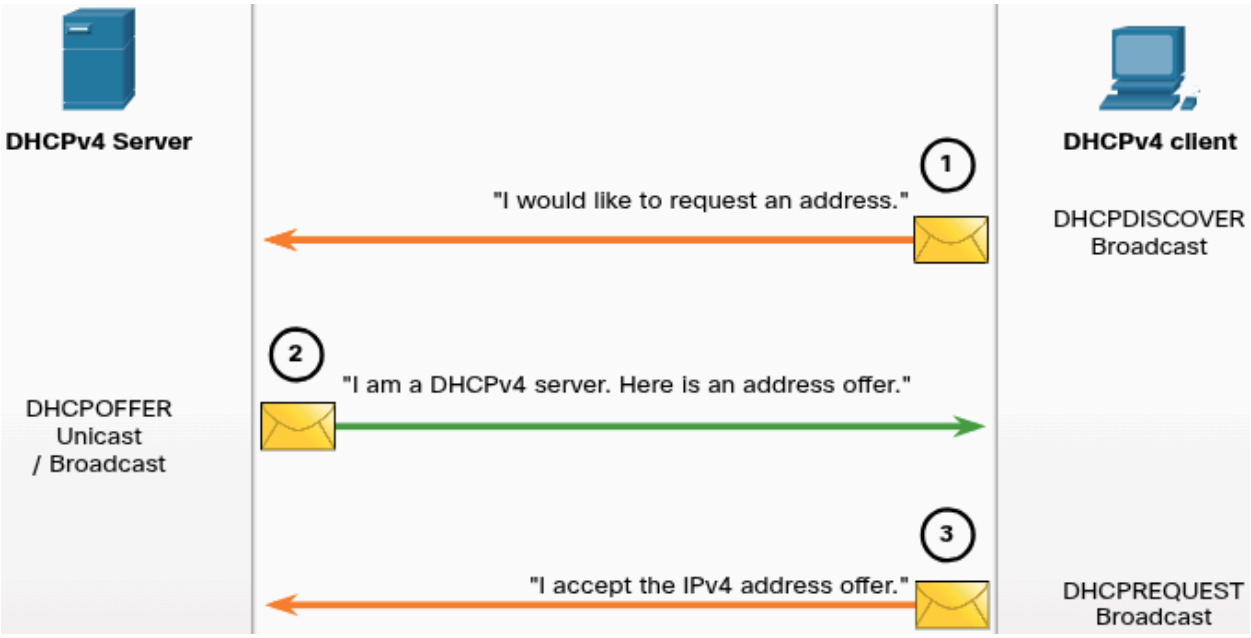

DHCP Acknowledgement (DHCPACK)

- Upon receiving DHCPREQUEST message
  - Server may verify lease information with an **ICMP ping** to address to ensure it is not being used already
  - Server will create a new **ARP entry** for client lease
  - Server will reply with DHCPACK message
- DHCPACK message is a duplicate of DHCPOFFER, except for a change in the message type field
- When client receives DHCPACK message
  - Client logs configuration information and may perform an **ARP lookup** for the assigned address
  - If no reply to ARP, client knows IPv4 address is valid and uses it as own

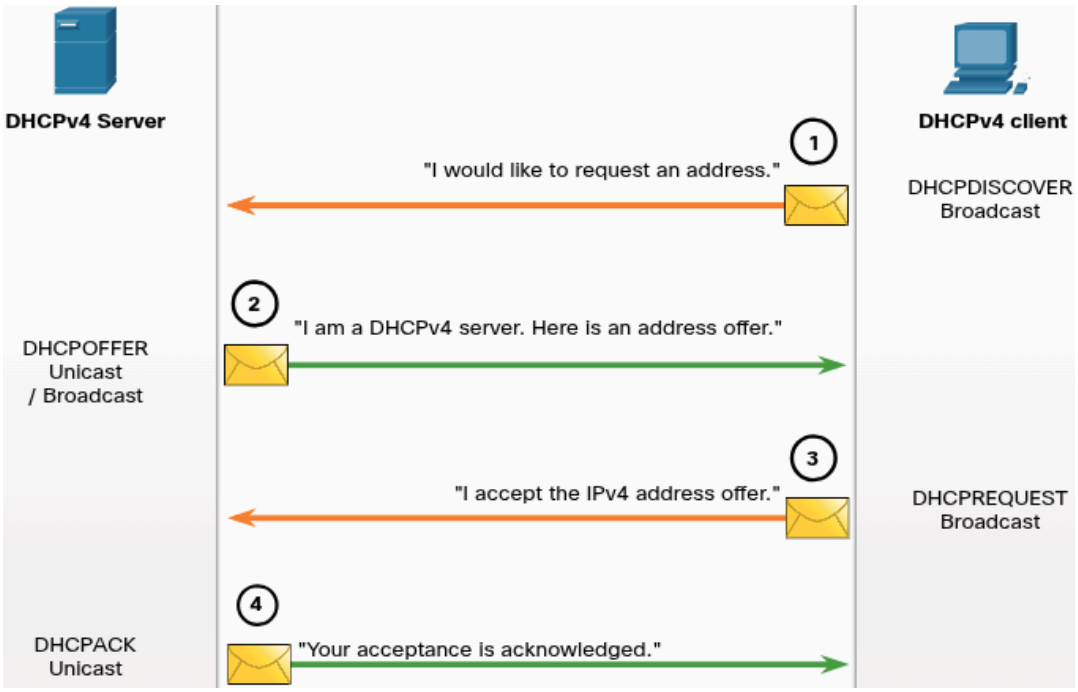

Steps to Renew a Lease

1. DHCP Request (DHCPREQUEST)
   Before lease expires
   - Client sends DHCPREQUEST message to DHCPv4 server that original offered IPv4 address
   - If DHCPACK is not received within specified amount of time
     - Client broadcasts another DHCPREQUEST so that one of the other DHCPv4 servers can extend the lease

2. DHCP Acknowledgement (DHCPACK)
   On receiving DHCPREQUEST message
   - Server verifies lease information by returning a DHCPACK

Note - These messages (primarily DHCPOFFER and DHCPACK) can be sent as **unicast** or **broadcast** accordingly to IETF RFT 2131

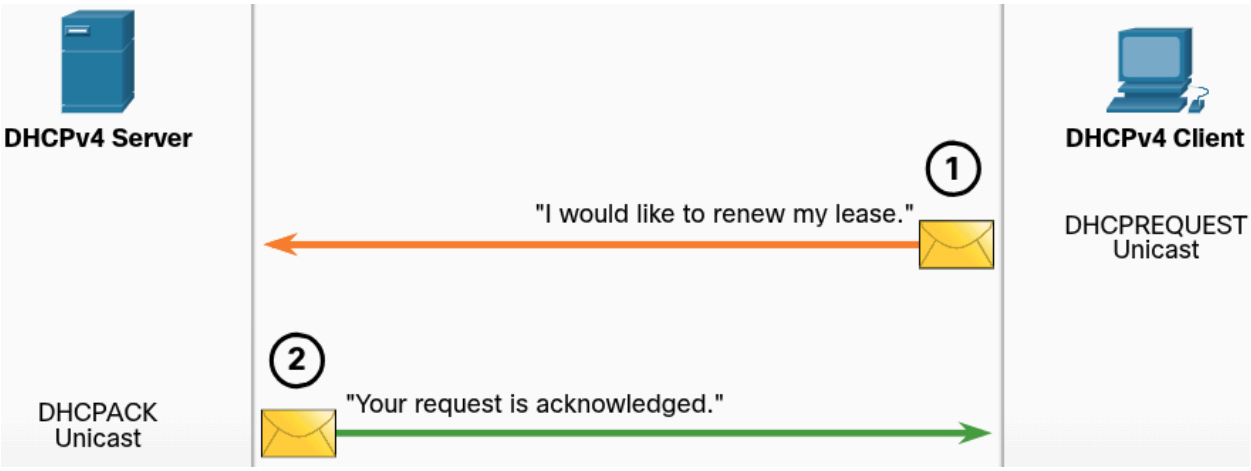

Configure a Cisco IOS DHCPv4 Server

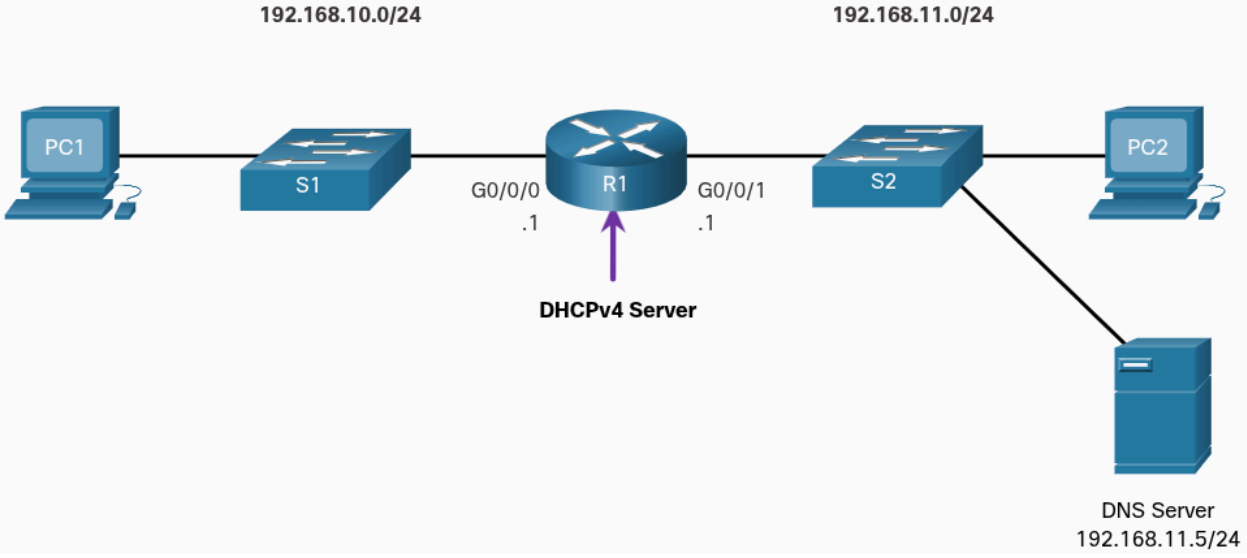

Steps to Configure a Cisco IOS DHCPv4 Server

1. Exclude IPv4 Addresses
2. Define a DHCPv4 pool name
3. Configure the DHCPv4 pool

1. Exclude IPv4 Addresses

- The router functioning as the DHCPv4 server assigns all IPv4 addresses in DHCPv4 address pool unless it is configured to exclude specific addresses.
- Typically, some IPv4 addresses in pool are assigned to network devices that require static address
    - These IPv4 addresses should not be assigned to other devices

```
Router(config)# ip dhcp excluded-address low-address [high-address]
```

A single address or a range of addresses can be excluded by specifying the *low-address* and *high-address* of the range

2. Define a DHCPv4 Pool Name

- The **ip dhcp pool** *pool-name* command creates a pool with the specified name and puts the router in **DHCPv4 configuration mode**, which is identified by prompt: Router(dhcp-config)**#**

```
Router(config)# ip dhcp pool pool-name
Router(dhcp-config)#
```

3. Configure the DHCPv4 Pool

- The address pool and default gateway router must be configured
- Use **network** statement to define the range of available addresses
- Use **default-router** command to define the default gateway router
- The gateway is the LAN interface of the router closest to the client
    - One gateway is required, but you can list up to **eight** addresses
        - If multiple gateways

Other DHCPv4 pool commands are optional

- IPv4 address of DNS server that is available to DHCPv4 client is configured using **dns-server** command
- The **domain-name** command is used to define domain name
- The **lease** command changes duration of DHCPv4 lease
- The **netbios-name-server** command is used to define the NetBIOS WINS server

Tasks to complete DHCPv4 pool configuration	
Task	IOS Command
Define address pool	network *network-number* [*mask* | */ prefix-length*]
Define default router or gateway	default-router *address* [*address2… address8*]
Define DNS server	dns-server *address* [*address2…address8*]
Define domain name	domain-name *domain*
Define duration of DHCP lease	lease {*days* [*hours* [*minutes*]] | infinite}
Define NetBIOS WINS server	netbios-name-server *address* [*address2…. address8*]

Note - Microsoft recommends **not** deploying WINS, instead configure DNS for Windows name resolution and decommission WINS

Configuration Example

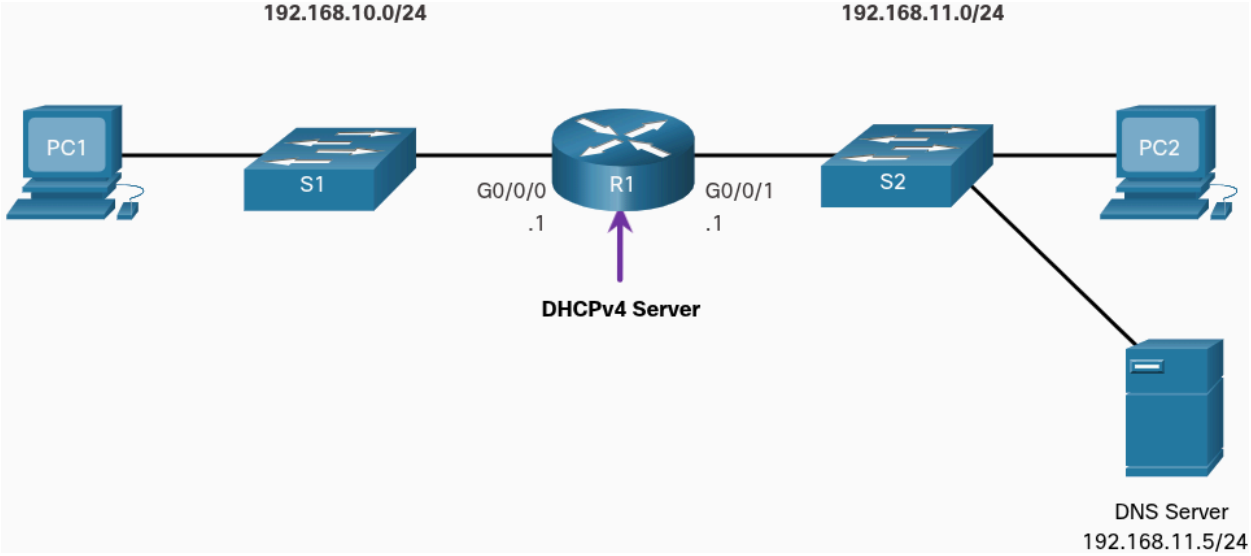

```
R1(config)# ip dhcp excluded-address 192.168.10.1 192.168.10.9
R1(config)# ip dhcp excluded-address 192.168.10.254
R1(config)# ip dhcp pool LAN-POOL-1
R1(dhcp-config)# network 192.168.10.0 255.255.255.0
R1(dhcp-config)# default-router 192.168.10.1
R1(dhcp-config)# dns-server 192.168.11.5
R1(dhcp-config)# domain-name example.com
R1(dhcp-config)# end
R1#
```

DHCPv4 Verification Commands

Command	Description
showing running-config | section dhcp	Displays the DHCPv4 commands configured on the router
show ip dhcp binding	Displays a list of all IPv4 address to MAC address bindings provided by the DHCPv4 service
show ip dhcp server statistics	Displays count information regarding the number of DHCPv4 messages that have been sent and received

Verify DHCPv4 is Operational

- R1 has been configured to provide DHCPv4 services
- PC1 has not been powered up, therefore, does not have IP address

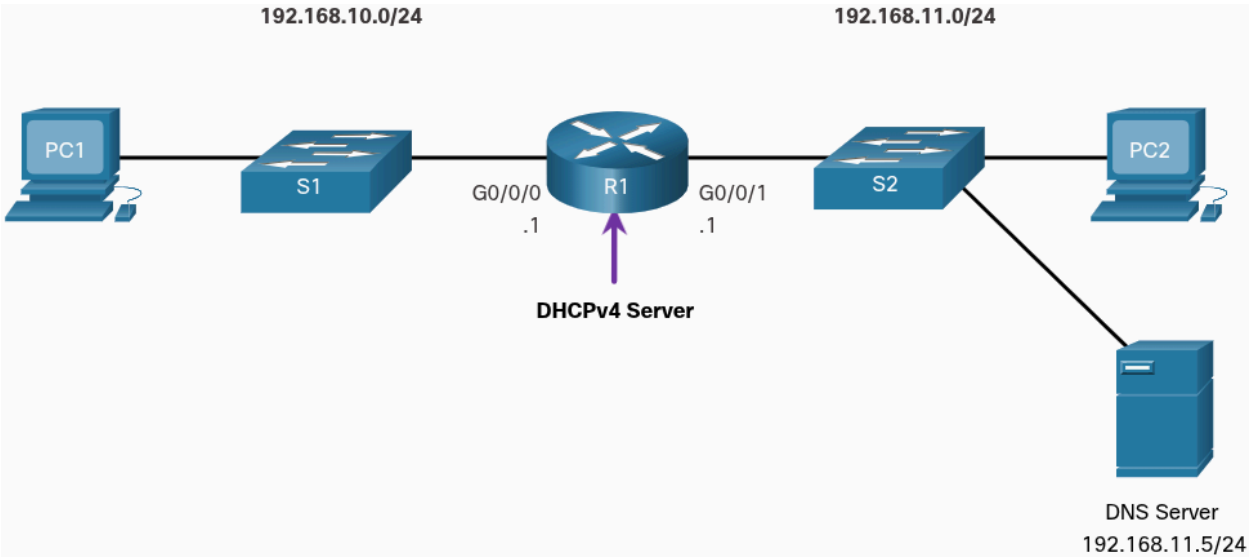

The output for the following commands assumes PC1 has received its IPv4 addressing information from the DHCPv4 server. You may need to enter *ipconfig /renew* on a Windows PC to force it to send out a DHCPDISCOVER message

Verify the DHCPv4 Configuration

- The **show running-config | section dhcp** command displays the DHCPv4 commands configured on R1
- The **| section** parameter displays only the command associated with DHCPv4 configuration

```
R1# show running-config | section dhcp
ip dhcp excluded-address 192.168.10.1 192.168.10.9
ip dhcp excluded-address 192.168.10.254
ip dhcp pool LAN-POOL-1
 network 192.168.10.0 255.255.255.0
 default-router 192.168.10.1
 dns-server 192.168.11.5
 domain-name example.com
```

Verify DHCPv4 Bindings

- The operations of DHCPv4 can be verified using the **show ip dhcp binding** command
- This command displays a list of all IPv4 address to MAC address bindings that hve been provided by DHCPv4 service

```
R1# show ip dhcp binding
Bindings from all pools not associated with VRF:
IP address      Client-ID/              Lease expiration        Type       State
Interface
                Hardware address/
                User name
192.168.10.10   0100.5056.b3ed.d8       Sep 15 2019 8:42 AM    Automatic  Active
GigabitEthernet0/0/0
```

Verify DHCPv4 Statistics

- The output of **show ip dhcp server statistics** is used to verify that messages are being received or sent by router
- This command displays count information regarding the number of DHCPv4 messages that have been sent and received

```
R1# show ip dhcp server statistics
Memory usage         19465
```

```
Address pools         1
Database agents       0
Automatic bindings    2
Manual bindings       0
Expired bindings      0
Malformed messages    0
Secure arp entries    0
Renew messages        0
Workspace timeouts    0
Static routes         0
Relay bindings        0
Relay bindings active         0
Relay bindings terminated     0
Relay bindings selecting      0
Message               Received
BOOTREQUEST           0
DHCPDISCOVER          4
DHCPREQUEST           2
DHCPDECLINE           0
DHCPRELEASE           0
DHCPINFORM            0
```

Verify DHCPv4 Client Received IPv4 Addressing

- The **ipconfig /all** command displays the TCP/IP parameters
- PC1 connected to 192.168.10.0/24, automatically received a DNS suffix, IPv4 address, subnet mask, default gateway, and DNS server address from that pool
    - No DHCP-specific router interface configuration is required

```
C:\Users\Student> ipconfig /all
Windows IP Configuration
   Host Name . . . . . . . . . . . . : ciscolab
   Primary Dns Suffix  . . . . . . . :
   Node Type . . . . . . . . . . . . : Hybrid
   IP Routing Enabled. . . . . . . . : No
   WINS Proxy Enabled. . . . . . . . : No
Ethernet adapter Ethernet0:
   Connection-specific DNS Suffix  . : example.com
   Description . . . . . . . . . . . : Realtek PCIe GBE Family Controller
   Physical Address. . . . . . . . . : 00-05-9A-3C-7A-00
   DHCP Enabled. . . . . . . . . . . : Yes
   Autoconfiguration Enabled . . . . : Yes
   IPv4 Address. . . . . . . . . . . : 192.168.10.10
   Subnet Mask . . . . . . . . . . . : 255.255.255.0
   Lease Obtained  . . . . . . . . . : Saturday, September 14, 2019 8:42:22AM
   Lease Expires   . . . . . . . . . : Sunday, September 15, 2019 8:42:22AM
   Default Gateway . . . . . . . . . : 192.168.10.1
   DHCP Server . . . . . . . . . . . : 192.168.10.1
   DNS Servers . . . . . . . . . . . : 192.168.11.5
```

Disable the Cisco IOS DHCPv4 Server

The DHCPv4 service is enabled by default

- Use **no service dhcp** global configuration mode command
- Use **service dhcp** global configuration mode command to re-enable
    - Enabling the service has **no effect** if the parameters are **not** configured

Note - Clearing the DHCP bindings or stopping and restarting the DHCP service may result in **duplicate IP addresses** being **temporarily** assigned on the network

```
R1(config)# no service dhcp
R1(config)# service dhcp
R1(config)#
```

DHCPv4 Relay

In a complex hierarchical network, enterprise servers are usually located centrally.

- Servers may provide DHCP, DNS, TFTP, and FTP
- Network clients typically not on same subnet
    - In order to locate the servers and receive service, clients often use broadcast messages

PC1 is attempting to acquire IPv4 from DHCPv4 using broadcast message. R1 is not configured as DHCPv4 server and does not forward the broadcast.

- DHCPv4 server is located on different network, PC1 can't receive an IP using DHCP.
- R1 must be configured to relay DHCPv4 messages to DHCPv4 server

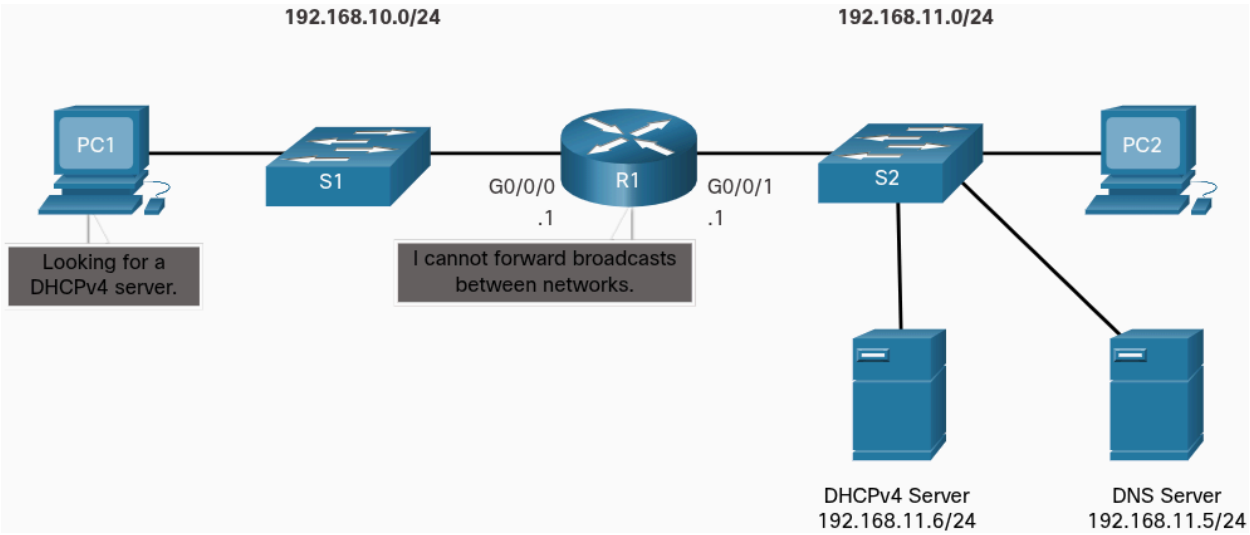

ipconfig /release

- Release all current IPv4 addressing information with **ipconfig /release** command

```
C:\Users\Student> ipconfig /release
Windows IP Configuration
Ethernet adapter Ethernet0:
   Connection-specific DNS Suffix  . :
   Default Gateway . . . . . . . . . :
```

ipconfig /renew

- This command causes PC1 to broadcast a **DHCPDISCOVER** message

PC1 is unable to locate DHCPv4 server.

```
C:\Users\Student> ipconfig /renew
Windows IP Configuration
An error occurred while renewing interface Ethernet0 : unable to
connect to your DHCP server. Request has timed out.
```

ip helper-address

- Configure R1 with **ip helper-address** *address* interface configuration command
- Will cause R1 to **relay** DHCPv4 broadcasts to the DHCPv4 server

```
R1(config)# interface g0/0/0
R1(config-if)# ip helper-address 192.168.11.6
R1(config-if)# end
R1#
```

show ip interface

- When R1 has been configured with DHCPv4 relay agent
    - It accepts broadcast requests for DHCPv4 service
    - Then forwards those requests as **unicast** to IPv4 address 192.168.11.6
- Can use **show ip interface** command to verify configuration

```
R1# show ip interface g0/0/0
GigabitEthernet0/0/0 is up, line protocol is up
  Internet address is 192.168.10.1/24
  Broadcast address is 255.255.255.255
  Address determined by setup command
  MTU is 1500 bytes
  Helper address is 192.168.11.6
(output omitted)
```

ipconfig /all

- PC1 is now able to acquire IPv4 address

```
C:\Users\Student> ipconfig /all
Windows IP Configuration

Ethernet adapter Ethernet0:
   Connection-specific DNS Suffix  . : example.com
   IPv4 Address. . . . . . . . . . . : 192.168.10.10
   Subnet Mask . . . . . . . . . . . : 255.255.255.0
   Default Gateway . . . . . . . . . : 192.168.10.1
```

Other Service Broadcasts Relayed

- The **ip helper-address** command forwards following eight UDP services:
    - Port 37: Time
    - Port 49: TACACS
    - Port 53: DNS
    - Port 67: DHCP/BOOTP server
    - Port 68: DHCP/BOOTP client
    - Port 69: TFTP
    - Port 137: NetBIOS name service
    - Port 138: NetBIOS datagram service

Configure a DHCPv4 Client

Cisco Router as a DHCPv4 Client

- Use **ip address dhcp** interface configuration mode command

Assume ISP has been configured to provide select customers with IP address from 209.165.201.0/27 range after G0/0/1 interface is configured with **ip address dhcp** command

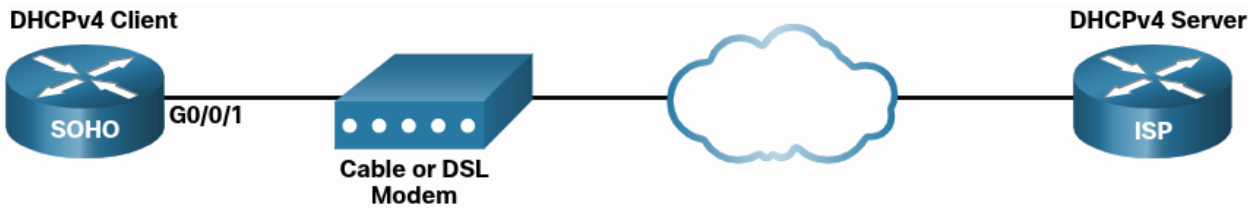

Configuration Example

To configure Ethernet interface as DHCP client

- Use **ip address dhcp** interface configuration mode command

```
SOHO(config)# interface G0/0/1
SOHO(config-if)# ip address dhcp
SOHO(config-if)# no shutdown
Sep 12 10:01:25.773: %DHCP-6-ADDRESS_ASSIGN: Interface
GigabitEthernet0/0/1 assigned DHCP address 209.165.201.12, mask
255.255.255.224, hostname SOHO
```

The **show ip interface g0/0/1** command confirms that the interface is up and that the address was allocated by DHCPv4 server

```
SOHO# show ip interface g0/0/1
GigabitEthernet0/0/1 is up, line protocol is up
  Internet address is 209.165.201.12/27
  Broadcast address is 255.255.255.255
  Address determined by DHCP
(output omitted)
```

Home Router as DHCPv4 Client

- Typically already set to receive IPv4 addressing information automatically from ISP
    - Allows customers to easily set up router and connect to internet

Example:

- The internet connection type is set to **Automatic Configuration - DHCP**
    - Selection used when router is connected to DSL or cable modem and acts as DHCPv4 client, requesting an IPv4 address from ISP

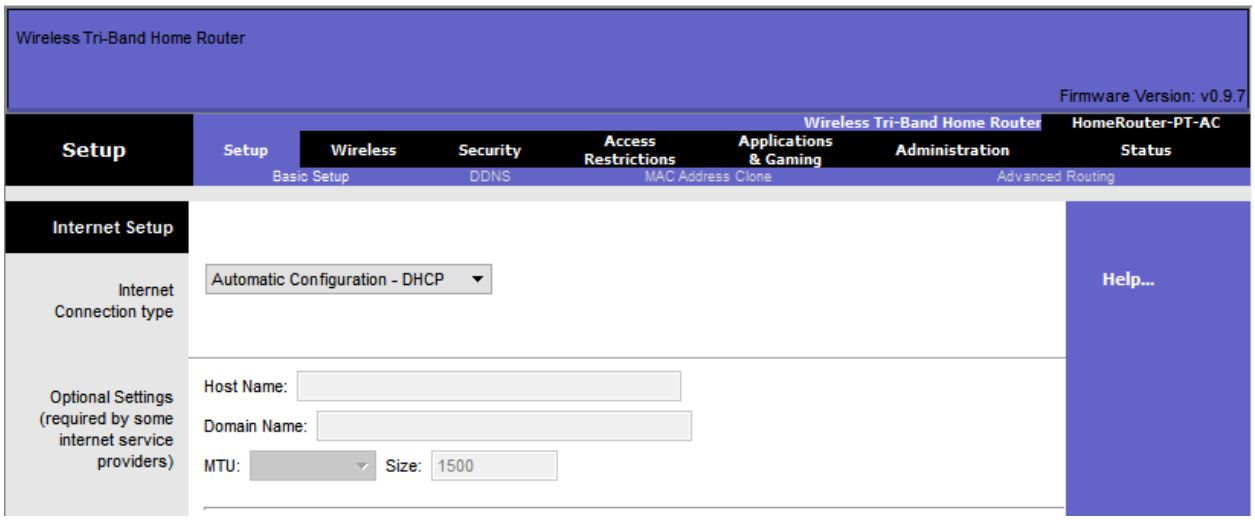

7.4.3 What did I learn in this module?

DHCPv4 Concepts

The DHCPv4 server dynamically assigns, or leases, an IPv4 address to a client from a pool of addresses for a limited period of time chosen by the server, or until the client no longer needs the address. The DHCPv4 lease process begins with the client sending message requesting the services of a DHCP server. If there is a DHCPv4 server that receives the message it will respond with an IPv4 address and possible other network configuration information. The client must contact the DHCP server periodically to extend the lease. This lease mechanism ensures that clients that move or power off do not keep addresses that they no longer need. When the client boots (or otherwise wants to join a network), it begins a four-step process to obtain a lease: DHCPDISCOVER, then DHCPOFFER, then DHCPREQUEST, and finally DHCPACK. Prior to lease expiration, the client begins a two-step process to renew the lease with the DHCPv4 server: DHCPREQUEST then DHCPACK.

Configure a Cisco IOS DHCPv4 Server

A Cisco router running Cisco IOS software can be configured to act as a DHCPv4 server. Use the following steps to configure a Cisco IOS DHCPv4 server: exclude IPv4

addresses, define a DHCPv4 pool name, and configure the DHCPv4 pool. Verify your configuration using the show running-config | section dhcp, show ip dhcp binding, and show ip dhcp server statistics commands. The DHCPv4 service is enabled, by default. To disable the service, use the no service dhcp global configuration mode command. In a complex hierarchical network, enterprise servers are usually located centrally. These servers may provide DHCP, DNS, TFTP, and FTP services for the network. Network clients are not typically on the same subnet as those servers. In order to locate the servers and receive services, clients often use broadcast messages. A PC is attempting to acquire an IPv4 address from a DHCPv4 server using a broadcast message. If the router is not configured as a DHCPv4 server, it will not forward the broadcast. If the DHCPv4 server is located on a different network, the PC cannot receive an IP address using DHCP. The router must be configured to relay DHCPv4 messages to the DHCPv4 server. The network administrator releases all current IPv4 addressing information using the ipconfig /release command. Next, the network administrator attempts to renew the IPv4 addressing information with the ipconfig /renew command. A better solution is to configure R1 with the ip helper-address *address interface* configuration command. The network administrator can use the show ip interface command to verify the configuration. The PC is now able to acquire an IPv4 address from the DHCPv4 server as verified with the ipconfig /all command. By default, the ip helper-address command forwards the following eight UDP services:

- Port 37: Time
- Port 49: TACACS
- Port 53: DNS
- Port 67: DHCP/BOOTP server
- Port 68: DHCP/BOOTP client
- Port 69: TFTP
- Port 137: NetBIOS name service
- Port 138: NetBIOS datagram service

Configure a DHCPv4 Client

The Ethernet interface is used to connect to a cable or DSL modem. To configure an Ethernet interface as a DHCP client, use the ip address dhcp interface configuration mode command. Home routers are typically already set to receive IPv4 addressing information automatically from the ISP. The internet connection type is set to Automatic Configuration - DHCP. This selection is used when the router is connected to a DSL or cable modem and acts as a DHCPv4 client, requesting an IPv4 address from the ISP.

IPv6 GUA Assignment

To use either stateless address autoconfiguration (SLAAC) or DHCPv6, review **global unicast addresses (GUAs)** and **link-local addresses (LLAs)**

- On router, IPv6 global unicast address (GUA) is **manually** configured using the **ipv6 address** *ipv6-address/prefix-length* interface configuration command

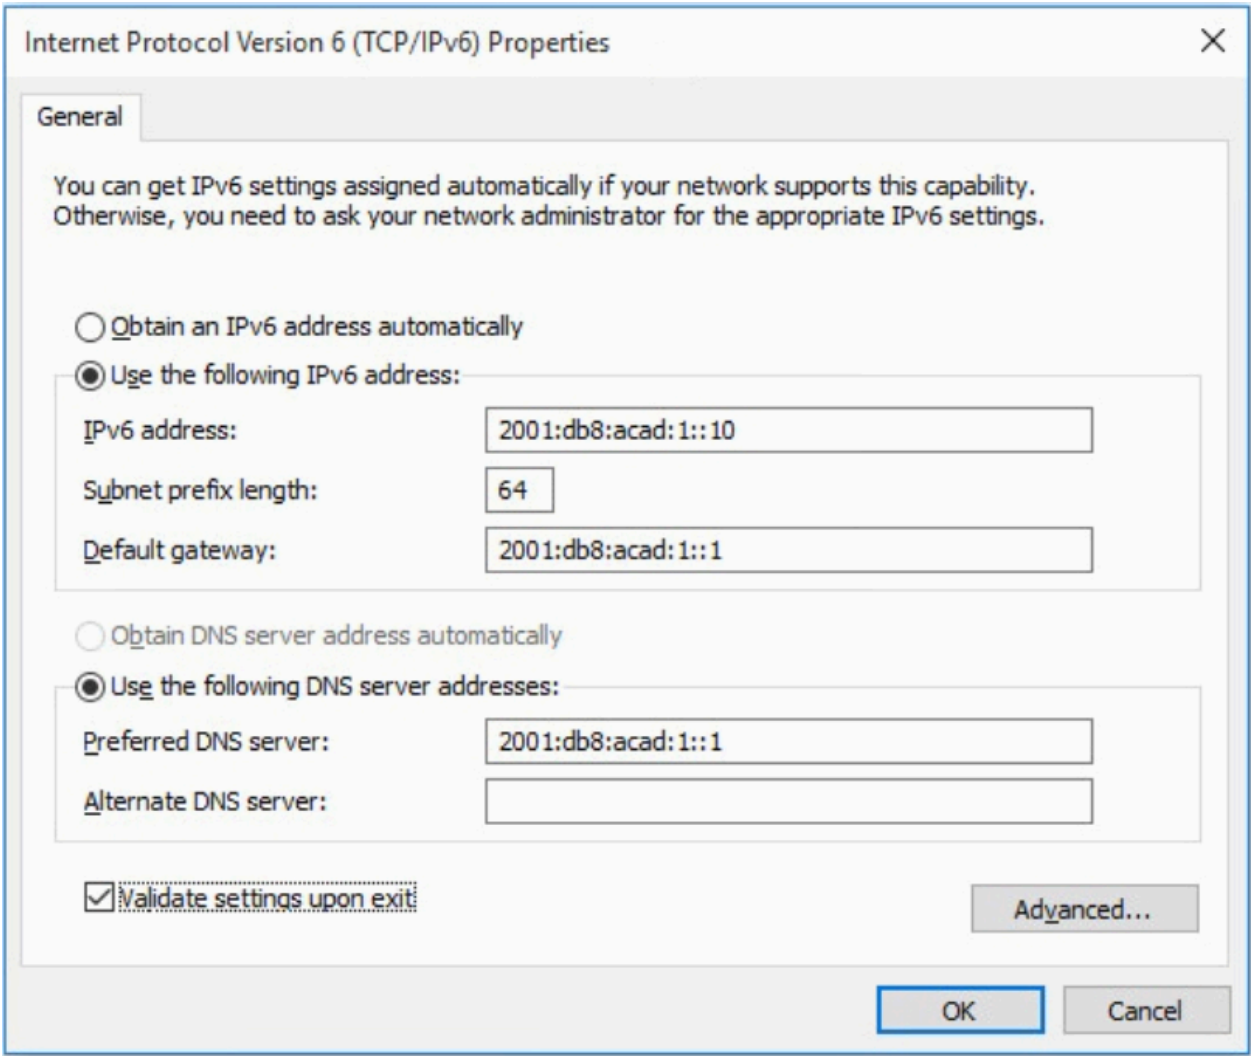

Manually entering IPv6 GUA is time consuming and somewhat error prone

- Most windows hosts are enabled to dynamically acquire an IPv6 GUA configuration

IPv6 Host Link-Local Address

- When automatic IPv6 addressing is selected, host will attempt to automatically obtain and configure IPv6 address information on interface
- Host will use one of three methods defined by the **Internet Control Message Protocol version 6 (ICMPv6) Router Advertisement (RA) message** received on interface.

An IPv6 router that is on the **same link** as the **host** sends out RA messages that suggests to the hosts how to obtain IPv6

- IPv6 **link-local** is **automatically created** by host when it boots and Ethernet interface is active

ipconfig shows an automatically generated link-local address (LLA)

Note - Host operating systems will at times show a link-local address appended with a "%" and a number

- This is known as a **Zone ID** or **Scope ID**. Used by the OS to associate the LLA with a specific interface

```
C:\PC1> ipconfig
Windows IP Configuration
Ethernet adapter Ethernet0:
   Connection-specific DNS Suffix  . :
   IPv6 Address. . . . . . . . . . . :
   Link-local IPv6 Address . . . . . :
fe80::fb:1d54:839f:f595%21
   IPv4 Address. . . . . . . . . . . : 169.254.202.140
   Subnet Mask . . . . . . . . . . . : 255.255.0.0
   Default Gateway . . . . . . . . . :
C:\PC1>
```

IPv6 GUA Assignment

- By default, IPv6-enabled router advertises its IPv6 information
  - Allows a host to dynamically create or acquire IPv6 configuration
- Can be assigned dynamically using **stateless** and **stateful** services

All stateless and stateful methods use ICMPv6 RA messages to suggest to host how to create or acquire its IPv6 config

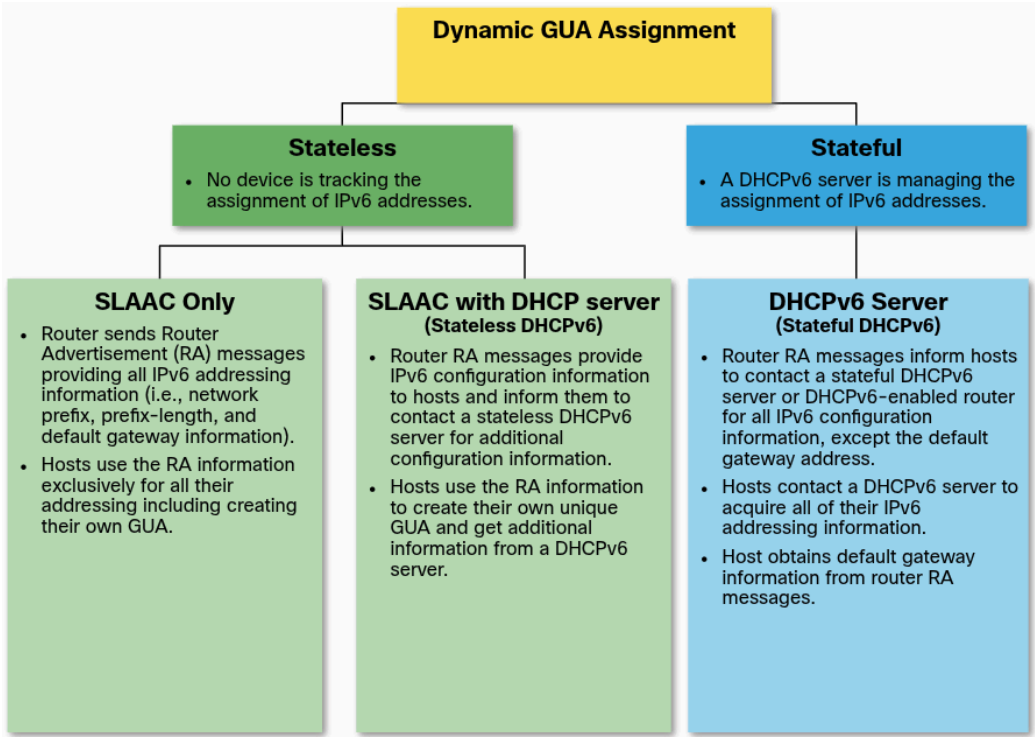

Three RA Message Flags

Decision of how host will obtain IPv6 GUA depends on settings within RA message

- **A flag** - This is the **Address Autoconfiguration** flag.
    - Use **Stateless** Address Autoconfiguration **(SLAAC)** to create an IPv6 GUA
- **O flag** - This is the **Other Configuration** flag.
    - Other information is available from a **stateless** DHCPv6 server
- **M flag** - This is the **Managed Address Configuration** flag.
    - Use a **stateful** DHCPv6 server to obtain an IPv6 GUA

Using different combinations of the A, O, and M flags, RA messages inform the host about the dynamic options available

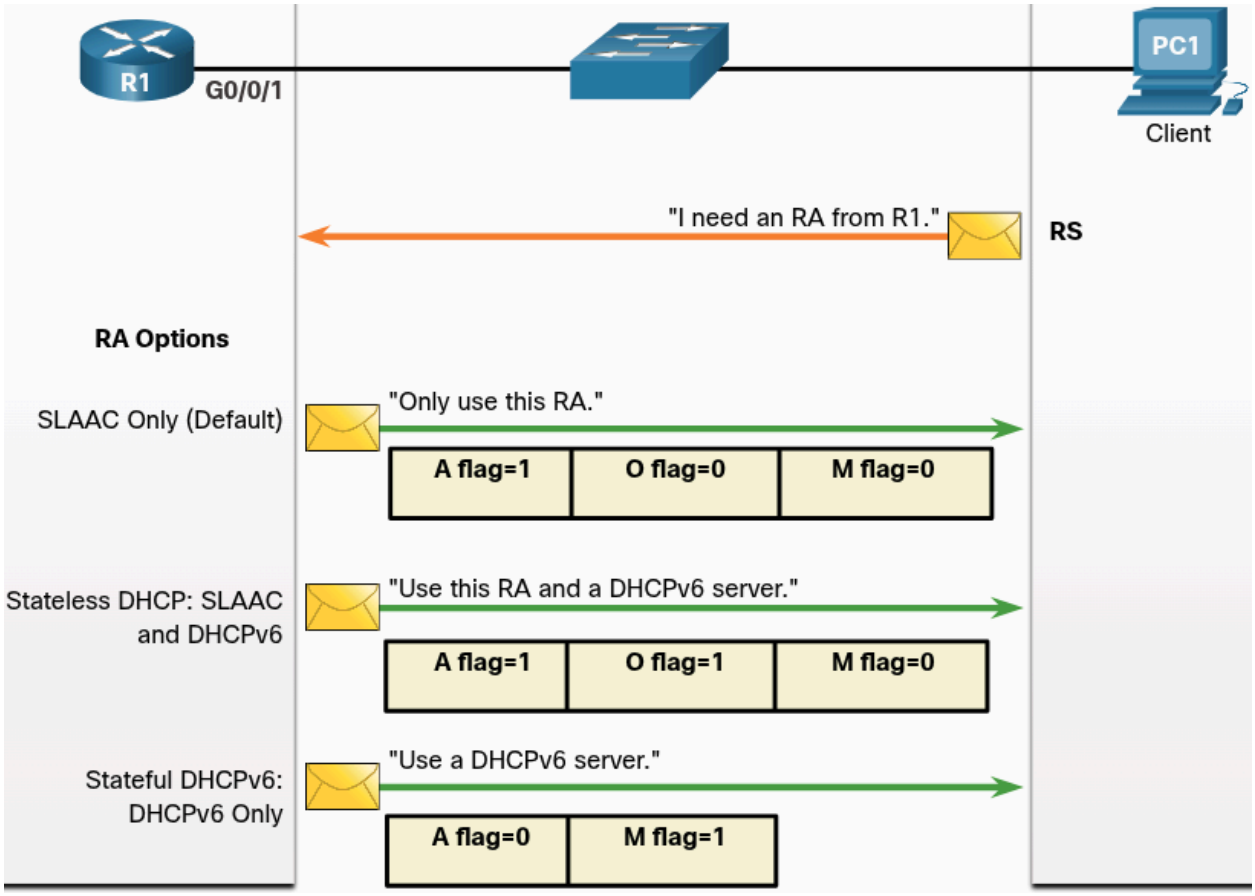

SLAAC and DHCPv6

Dynamically addressing protocols for IPv6 network.

SLAAC

- Every device in an IPv6 network needs a **GUA**
- SLAAC enables hosts to create their own unique IPv6 global unicast address **without** services of DHCPv6 server
- Uses ICMPv6 RA messages to provide addressing and other configuration information normally provided by DHCP server
    - A host configures its IPv6 address based on information sent in RA
    - RA messages are sent by an IPv6 router every **200 seconds**
- A host can also send a **Router Solicitation (RS)** message requesting that an IPv6-enabled router send the host an RA
- Can be deployed as SLAAC only, or SLAAC with DHCPv6

Enabling SLAAC

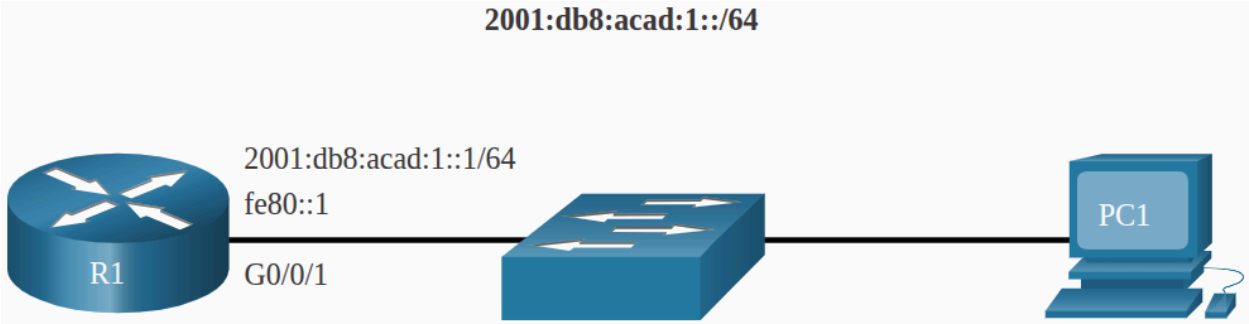

Assume R1 GigabitEthernet 0/0/1 has been configured with the indicated IPv6 GUA and link-local addresses

Verify IPv6 Addresses

- **show ipv6 interface** command displays current settings on G0/0/1 interface

R1 has been assigned the following IPv6 addresses:

- **Link-local IPv6 address** - fe80::1
- **GUA and subnet** - 2001:db8:acad:1::1 and 2001:db8:acad:1::/64
- **IPv6 all-nodes group** - ff02::1

```
R1# show ipv6 interface G0/0/1
GigabitEthernet0/0/1 is up, line protocol is up
  IPv6 is enabled, link-local address is FE80::1
  No Virtual link-local address(es):
  Description: Link to LAN
  Global unicast address(es):
    2001:DB8:ACAD:1::1, subnet is 2001:DB8:ACAD:1::/64
  Joined group address(es):
    FF02::1
    FF02::1:FF00:1
(output omitted)
```

Enabling IPv6 Routing

- Although router interface has IPv6 config, it is still not yet enabled to send RAs containing address configuration information to hosts using SLAAC

To enable, a router must join the IPv6 all-routers group using **ipv6 unicast-routing** global config command

```
R1(config)# ipv6 unicast-routing
R1(config)# exit
R1#
```

Verify SLAAC is Enabled

- The IPv6 all-routers group responds to the IPv6 multicast address ff02::2
- Use **show ipv6 interface** command to verify if a router is enabled

An IPv6-enabled Cisco router sends RA messages to the IPv6 all-nodes multicast address **ff02::1** every **200 seconds**

```
R1# show ipv6 interface G0/0/1 | section Joined
  Joined group address(es):
    FF02::1
    FF02::2
    FF02::1:FF00:1
R1#
```

SLAAC Only Method

- Enabled by default when the **ipv6 unicast-routing** command is configured
- All enabled Ethernet interfaces with an IPv6 GUA configured will start sending RA messages with the **A flag** set to **1**, and the **O** and **M flags** set to **0**
- **O=0** and **M=0** flags instruct the client to use the information in the RA message **exclusively**
    - The RA includes the prefix, prefix-length, DNS server, MTU, and default gateway information
- No further information available from DHCPv6 server

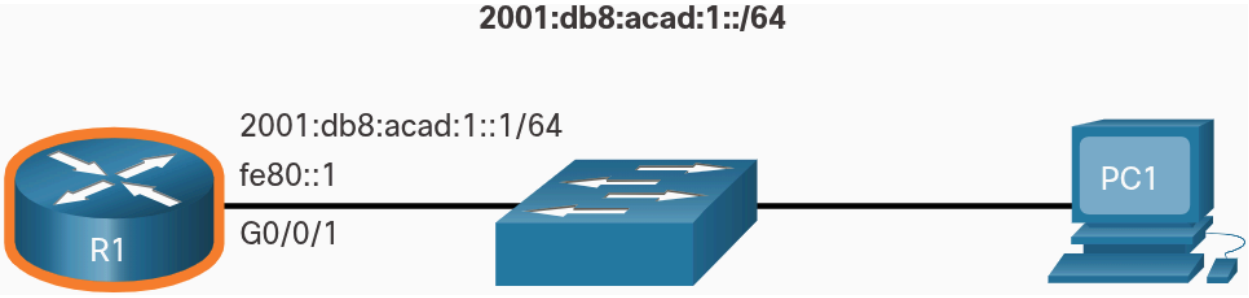

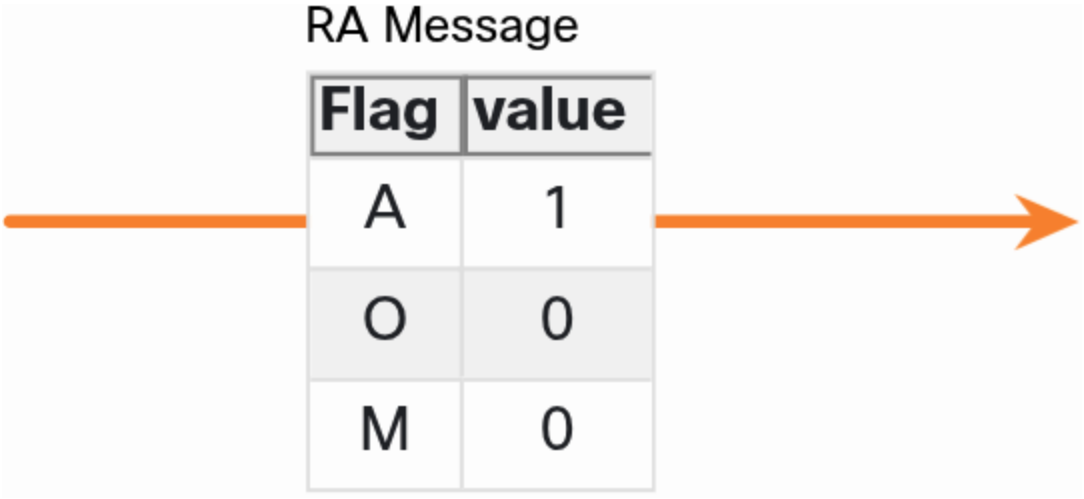

Example: PC1 is enabled to obtain IPv6 address information automatically

- Because of the settings of the A, O, and M flags, PC1 performs SLAAC only, using information contained in RA message sent by R1

The default gateway address is the source IPv6 address of RA message, which is the LLA for R1

- Can only be obtained automatically from RA message
- DHCPv6 server does **not** provide this information

```
C:\PC1> ipconfig
Windows IP Configuration
Ethernet adapter Ethernet0:
   Connection-specific DNS Suffix  . :
   IPv6 Address. . . . . . . . . . . :
2001:db8:acad:1:1de9:c69:73ee:ca8c
   Link-local IPv6 Address . . . . . :
fe80::fb:1d54:839f:f595%21
   IPv4 Address. . . . . . . . . . . : 169.254.202.140
   Subnet Mask . . . . . . . . . . . : 255.255.0.0
   Default Gateway . . . . . . . . . : fe80::1%21
C:\PC1>
```

ICMPv6 RS Messages

- A router sends RA messages every 200 seconds.
    - However it will also send an RA message if it receives an RS message from host

When a client is configured to obtain its addressing information automatically, it sends an RS message to the IPv6 all-routers multicast address of ff02::2

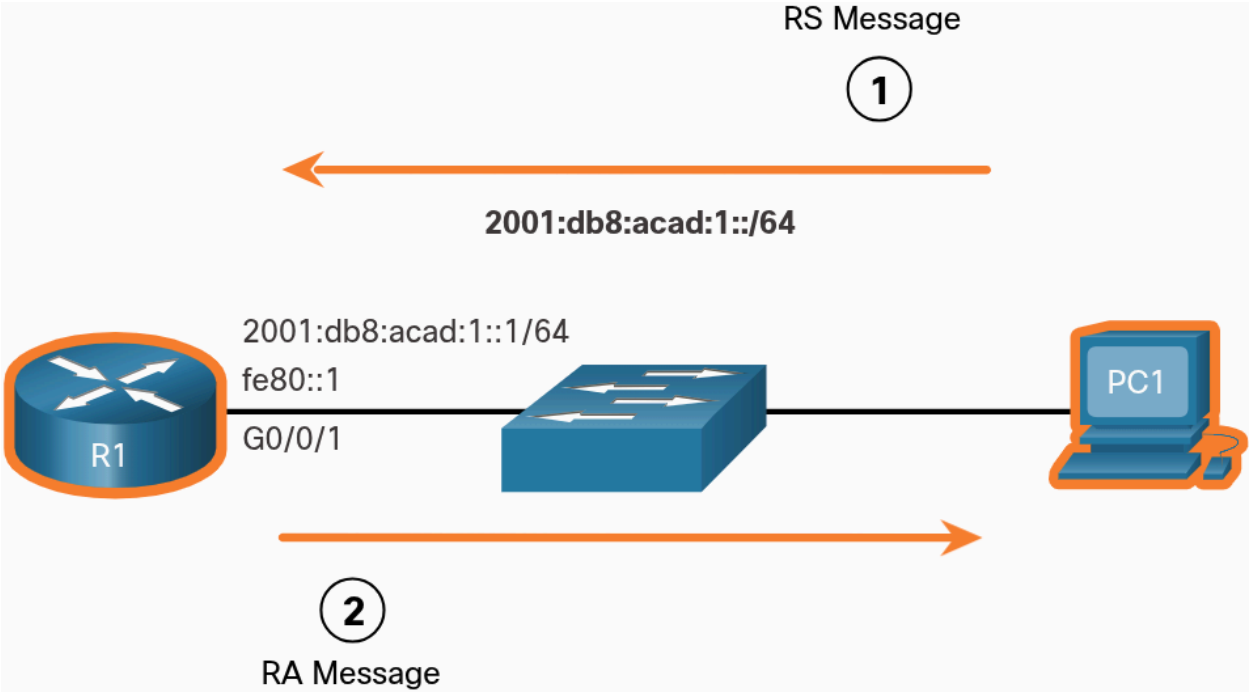

*1. PC1 has just booted and has not yet received an RA message. Therefore, it sends an RS message to the IPv6 all-routers multicast address of **ff02::2** requesting an **RA***
*2. R1 is part of the IPv6 all-routers group and received the RS message. it generates an RA containing the local network prefix and prefix length (IE 2001:db8:acad:1::/64). It then sends the RA message to the IPv6 all-nodes multicast address of **ff02::1**. PC1 uses this information to create a **unique IPv6 GUA***

Host Process to Generate Interface ID

- Using SLAAC, a host typically acquires its **64-bit IPv6 subnet** information from the **router RA.**
    - However, it must **generate** the **remainder** 64-bit interface identifier (ID) using one of two methods:

- **Randomly Generated** - The 64-bit interface ID is randomly generated by the client operating system. This is the method now used by Windows 10 hosts
- **EUI-64** - The host creates an interface ID using its **48-bit MAC address**
    - The host inserts the hex value of the fffe in the middle of the address, and flips the seventh bit of the interface ID
        - This changes the value of the **second hexadecimal digit** of the interface ID
    - Some operating systems default to the randomly generated interface ID instead of the EUI-64 method, due to privacy concerns
        - This is because the Ethernet MAC address of the host is used by EUI-64 to create the interface ID

Note - Windows, Linux, and Mac OS allow for the user to modify the generation of the interface ID to be easily randomly generated or to use EUI-64

Example: **ipconfig** output, the Windows 10 PC1 host used the IPv6 subnet information contained in the R1 RA and randomly generated a 64-bit interface ID

```
C:\PC1> ipconfig
Windows IP Configuration
Ethernet adapter Ethernet0:
   Connection-specific DNS Suffix  . :
   IPv6 Address. . . . . . . . . . . :
2001:db8:acad:1:1de9:c69:73ee:ca8c
   Link-local IPv6 Address . . . . . :
fe80::fb:1d54:839f:f595%21
   IPv4 Address. . . . . . . . . . . : 169.254.202.140
   Subnet Mask . . . . . . . . . . . : 255.255.0.0
   Default Gateway . . . . . . . . . : fe80::1%6
C:\PC1>
```

Duplicate Address Detection

- Enables the host to create IPv6 address
    - No guarantee the address is unique on network

SLAAC is stateless, therefore a host has the option to verify that a newly created IPv6 address is unique before use

Duplicate Address Detection (DAD)

- Used by host to ensure IPv6 GUA is unique
- Implemented using ICMPv6
- To perform, host sends an ICMPv6 Neighbor Solicitation (NS) message with a specially constructed multicast address, called a **solicited-note multicast address**
    - Duplicates the **last** 24 bits of IPv6 address of host
- If no other devices respond with a NA message, the address is virtually guaranteed to be unique and can be used by host
- If an NA is received by host, then the address is **not** unique and operating system has to determine a new interface ID to use

The Internet Engineering Task Force (IETF) recommends that DAD is used on **all** IPv6 unicast addresses regardless of whether it is created using SLAAC only, obtained using stateful DHCPv6, or manually configured.

- DAD is not mandatory because a 64-bit interface ID provides 18 quintillion possibilities and the chances of duplication is remote
    - However, most operating systems perform DAD on all IPv6 unicast addresses, regardlessh

DHCPv6
Operation Steps

Stateless
- Uses parts of SLAAC to ensure that all the necessary information is supplied to host

Stateful
- Does not require SLAAC

Note - DHCPv6 is defined in RFC 3315

The host begins the DHCPv6 client/server communications after stateless DHCPv6 or stateful DHCPv6 is indicated in the RA

- Server to client DHCPv6 messages use UDP destination port 546
- Client to server DHCPv6 messages use UDP destination port 547

The steps for DHCPv6 operations

1. The host sends an RS message
2. The router responds with an RA message
3. The host sends a DHCPv6 SOLICIT message
4. The DHCPv6 server responds with an ADVERTISE message
5. The host responds to the DHCPv6 server
6. The DHCPv6 server sends a REPLY message

Step 1. Host sends an RS message

PC1 sends an RS message to all IPv6-enabled routers

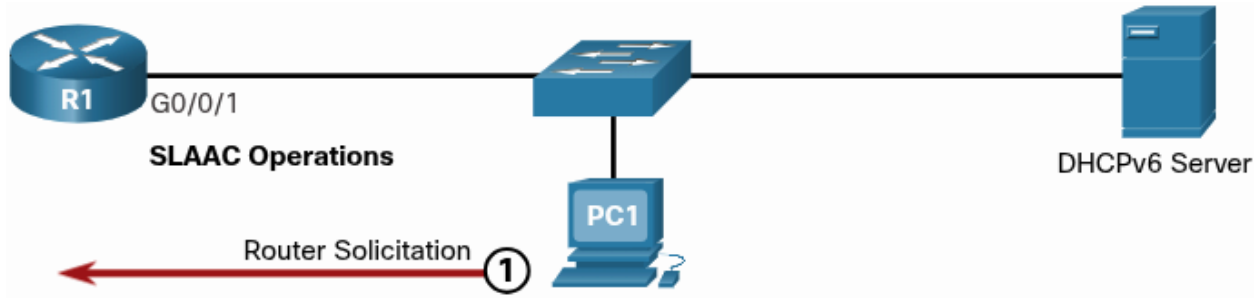

Step 2. Router responds with an RA message

R1 receives the RS and responds with an RA indicating that the client is to initiate communication with a DHCPv6 server

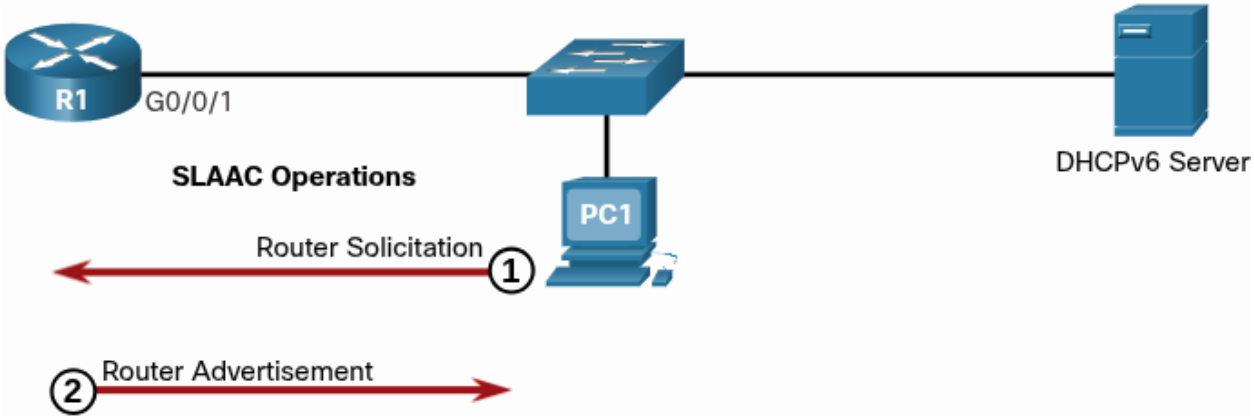

Step 3. Host sends a DHCPv6 SOLICIT message

The client, now a DHCPv6 client, needs to locate a DHCPv6 server and sends a DHCPv6 SOLICIT message to the reserved IPv6 multicast all-DHCPv6-servers address of **ff02::1:2**

- The multicast address has link-local scope, which means routers do not forward the messages to other networks

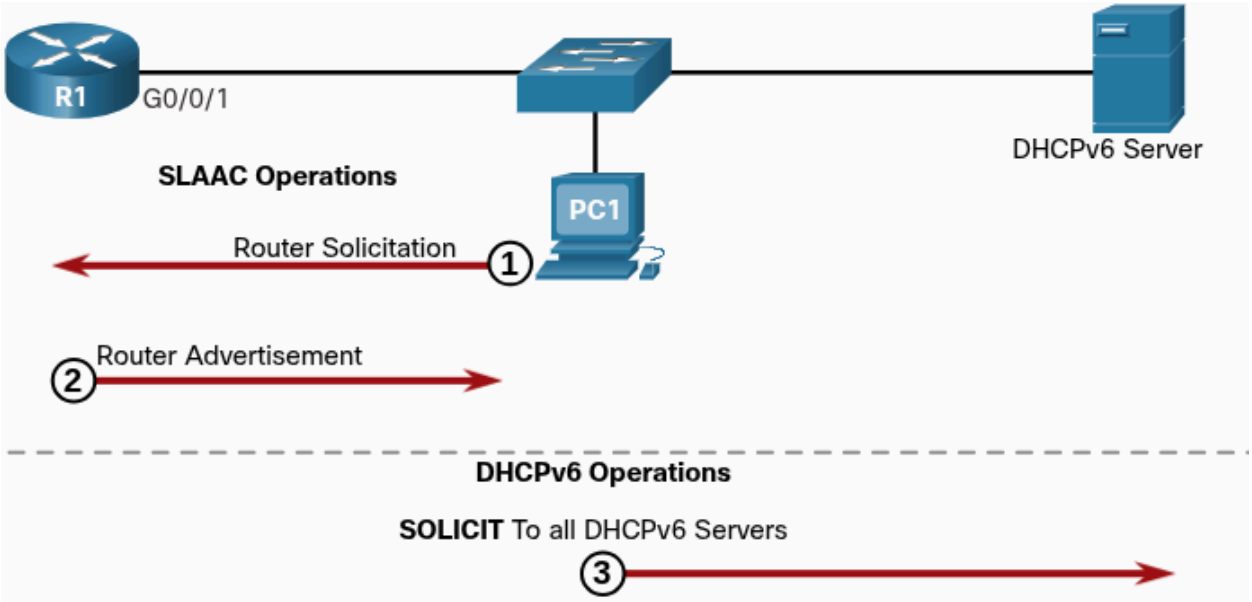

Step 4. DHCPv6 server responds with an ADVERTISE message

One or more DHCPv6 servers respond with a DHCPv6 ADVERTISE unicast message.

- The ADVERTISE message informs the DHCPv6 client that the server is available for DHCPv6 service

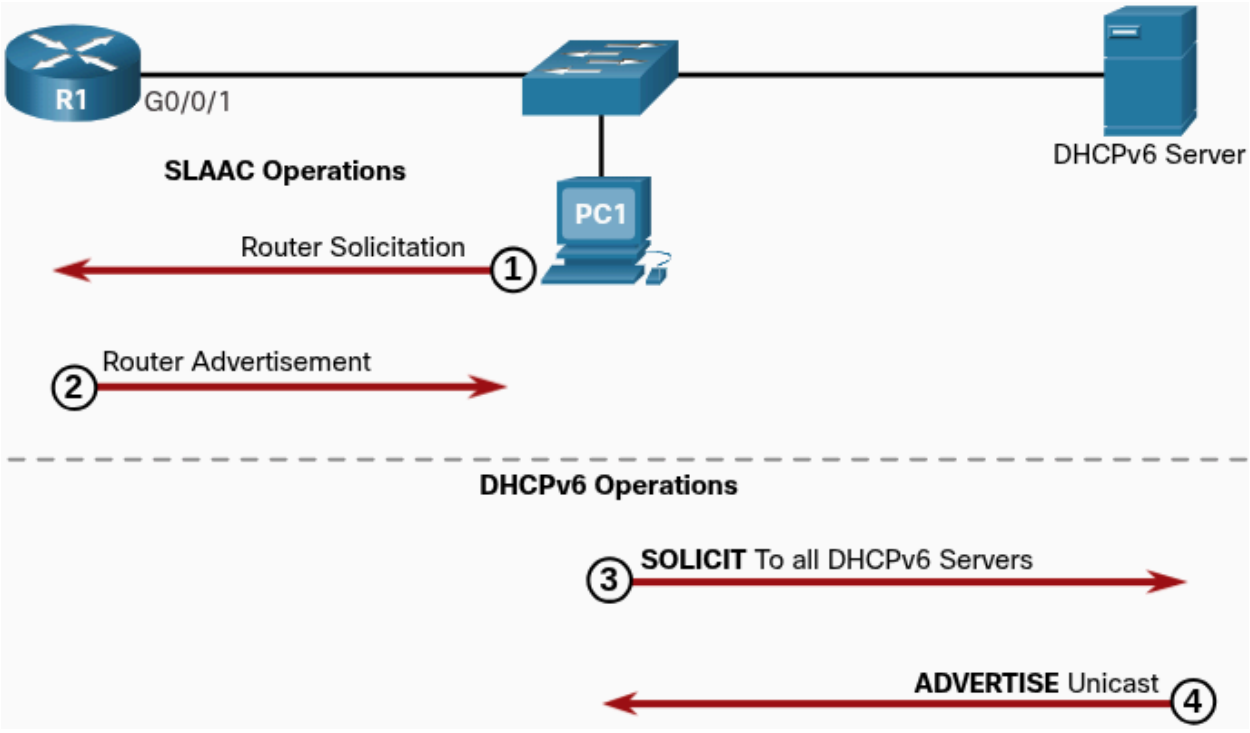

Step 5. Host responds to the DHCPv6 server

The PC1 response depends on whether it is using stateful or stateless DHCPv6:

- **Stateless DHCPv6 client** - The client creates an IPv6 address using the prefix in the RA message and a self-generated Interface ID.
  The client then sends a DHCPv6 INFORMATION-REQUEST message to the DHCPv6 server requesting additional configuration parameters (IE: DNS server address)
- **Stateful DHCPv6 client** - The client sends a DHCPv6 REQUEST message to the DHCPv6 server to obtain all necessary IPv6 configuration parameters

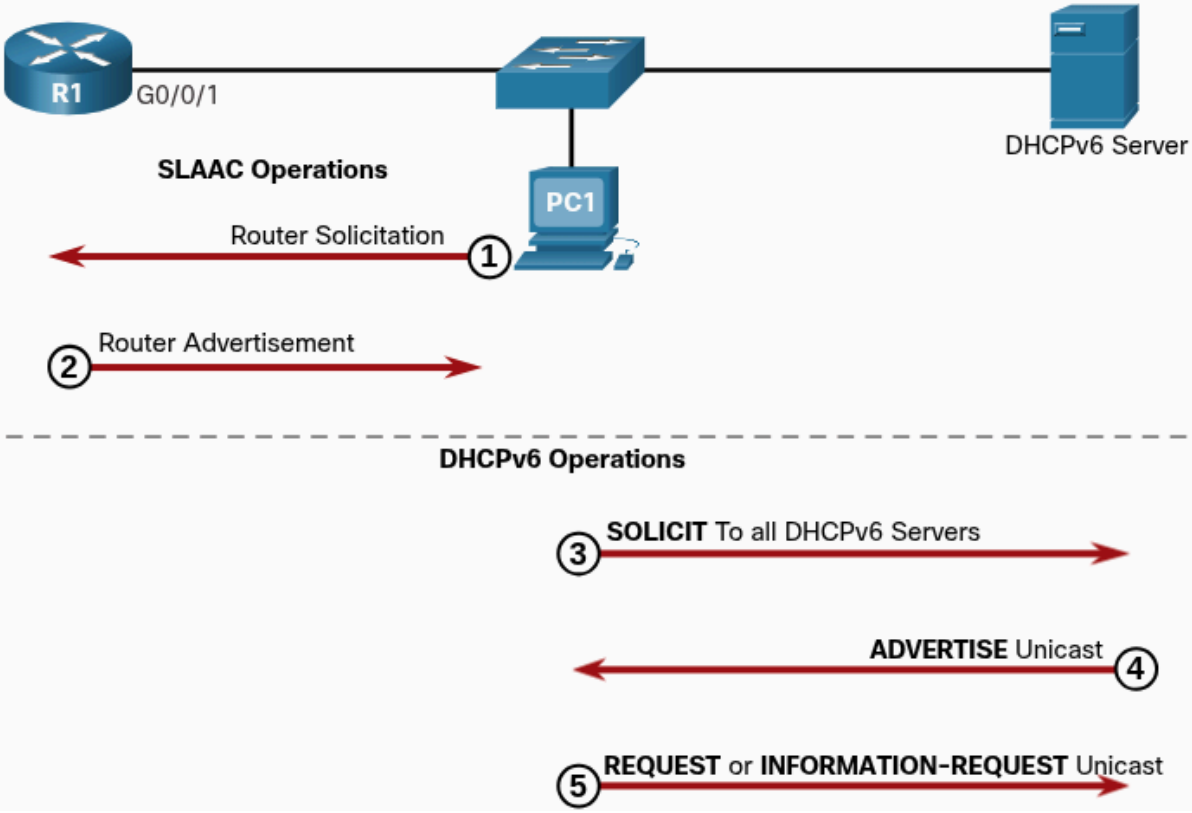

Step 6. DHCPv6 sends a REPLY message

The server sends a DHCPv6 REPLY unicast message to the client. The content of the message varies depending on if it is replying to a REQUEST or INFORMATION-REQUEST message

Note - The client will use the source IPv6 link-local address of the RA as its default gateway

- A DHCPv6 server does not provide this information

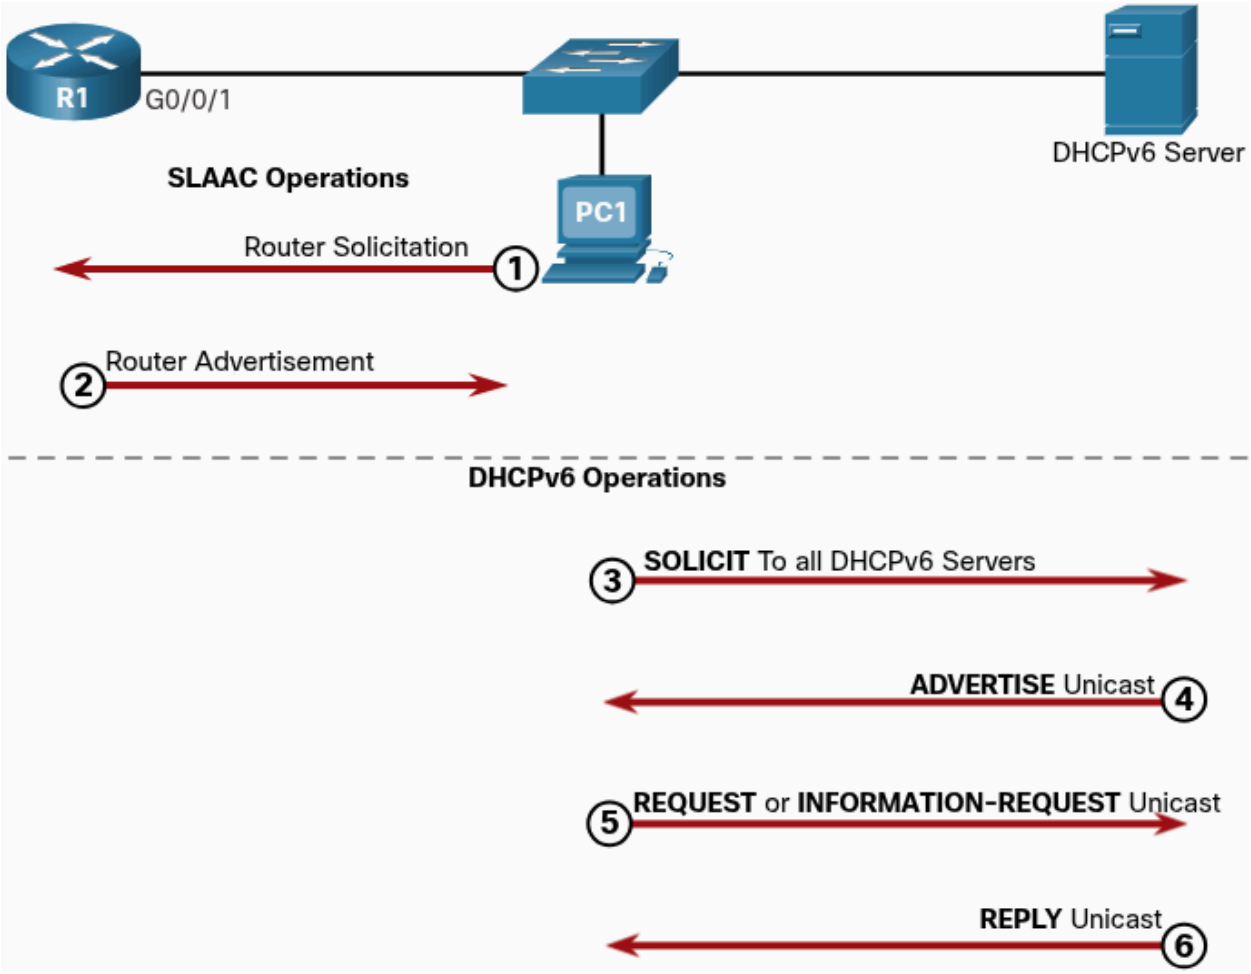

Stateless DHCPv6 Operation

- Only providing information that is identical for all devices on the network such as the IPv6 address of a DNS server
- Known as stateless because the server is **not** maintaining any client state information (IE: list of available and allocated IPv6 addresses)
    - Only provides configuration parameters for clients, not IPv6 addresses

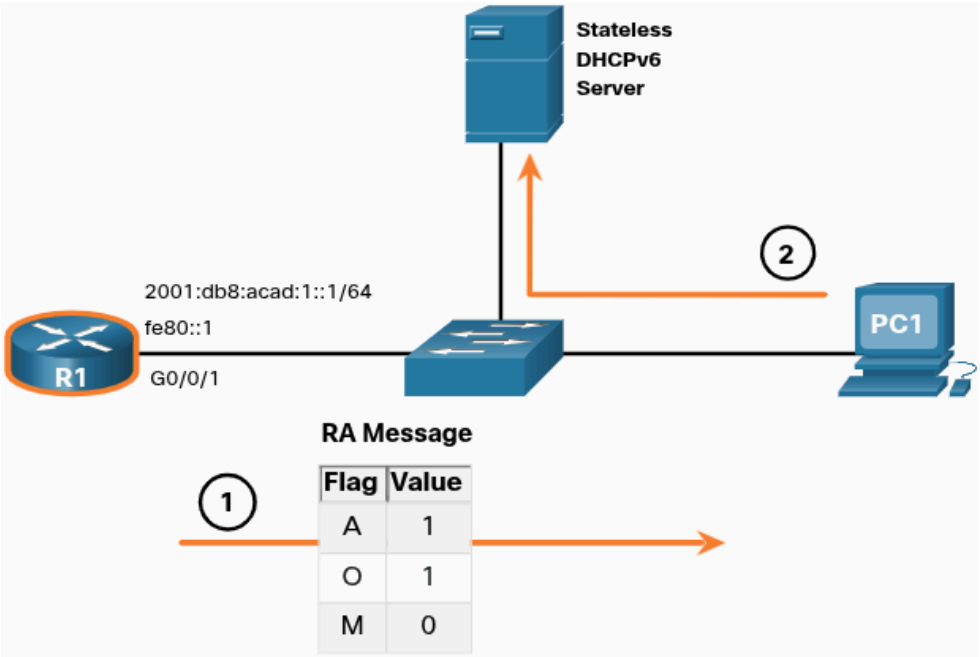

1. *PC1 receives a stateless DHCP RA message. The RA message contains the network prefix and prefix length. The M flag for stateful DHCP is set to the default value 0. The A=1 flag tells the client to use SLAAC. The O=1 flag informs the client that additional configuration information is available from a stateless DHCPv6 server*
2. *The client sends a DHCPv6 SOLICIT message looking for a stateless DHCPv6 server to obtain additional information (IE: DNS server addresses)*

Enable Stateless DHCPv6 on an Interface

- **ipv6 nd other-config-flag** interface configuration command
    - This sets the O flag to 1

The highlighted output confirms that the RA will tell receiving hosts to use stateless autoconfigure (A flag = 1) and contact a DHCPv6 server to obtain another configuration information (O flag = 1)

Note - Use **no ipv6 nd other-config-flag** to reset interface to default SLAAC only option (O flag = 0)

```
R1(config-if)# ipv6 nd other-config-flag
R1(config-if)# end
R1#
R1# show ipv6 interface g0/0/1 | begin ND
  ND DAD is enabled, number of DAD attempts: 1
  ND reachable time is 30000 milliseconds (using 30000)
  ND advertised reachable time is 0 (unspecified)
  ND advertised retransmit interval is 0 (unspecified)
  ND router advertisements are sent every 200 seconds
  ND router advertisements live for 1800 seconds
  ND advertised default router preference is Medium
  Hosts use stateless autoconfig for addresses.
  Hosts use DHCP to obtain other configuration.
R1#
```

Stateful DHCPv6 Operation

- RA message tells client to obtain all addressing information from a stateful DHCPv6 server, except the default gateway address which is the source IPv6 link-local address of the RA
    - Sends a **REQUEST**
- Known as stateful because the DHCPv6 server **maintains** IPv6 state information
    - Similar to DHCPv4 server allocating addresses for IPv4

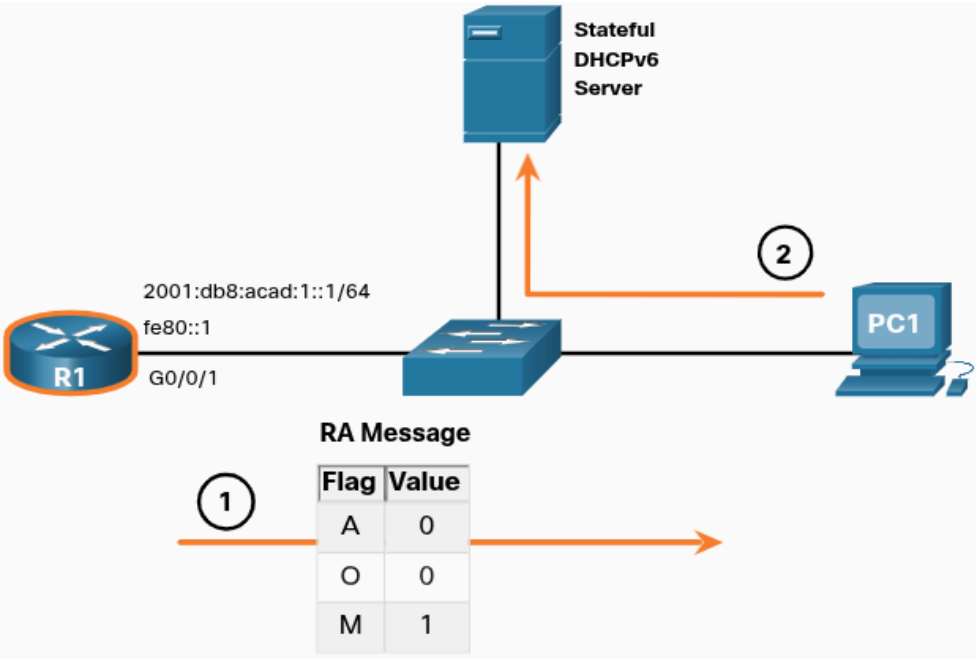

1. *PC1 receives a DHCPv6 RA message with the O flag set to 0 and the M flag set to 1, indicating to PC1 that it will receive all its IPv6 addressing information from a stateful DHCPv6 server*
2. *PC1 sends a DHCPv6 SOLICIT message looking for a stateful DHCPv6 server*

Note - If A=1 and M=1, some operating systems such as Windows will create an IPv6 address using SLAAC and obtain a different address from the stateful DHCPv6 server

- In most cases, recommended to manually set the A flag to 0

Enable Stateful DHCPv6 on an Interface

- **ipv6 nd managed-config-flag** interface configuration command
    - This sets the M flag to 1
- **ipv6 nd prefix default no-autoconfig** interface command **disables** SLAAC by setting the A flag to 0

The highlighted output confirms that the RA will tell the host to obtain all IPv6 configuration information from DHCPv6 server (M flag = 1)

```
R1(config)# int g0/0/1
R1(config-if)# ipv6 nd managed-config-flag
R1(config-if)# ipv6 nd prefix default no-autoconfig
R1(config-if)# end
R1#
R1# show ipv6 interface g0/0/1 | begin ND
  ND DAD is enabled, number of DAD attempts: 1
  ND reachable time is 30000 milliseconds (using 30000)
  ND advertised reachable time is 0 (unspecified)
  ND advertised retransmit interval is 0 (unspecified)
  ND router advertisements are sent every 200 seconds
  ND router advertisements live for 1800 seconds
  ND advertised default router preference is Medium
  Hosts use DHCP to obtain routable addresses.
R1#
```

Configure DHCPv6 Server

DHCPv6 Router Roles

- **DHCPv6 Server** - Router provides stateless or stateful DHCPv6 services
- **DHCPv6 Client** - Router interface acquires an IPv6 IP configuration from a DHCPv6 server
- **DHCPv6 Relay Agent** - Router provides DHCPv6 forwarding services when the client and the server are located on different networks

Configure a **Stateless** DHCPv6 **Server**

- Requires that the router advertise the IPv6 network addressing information in RA messages
    - However, client **must** contact a DHCPv6 server for **more** information

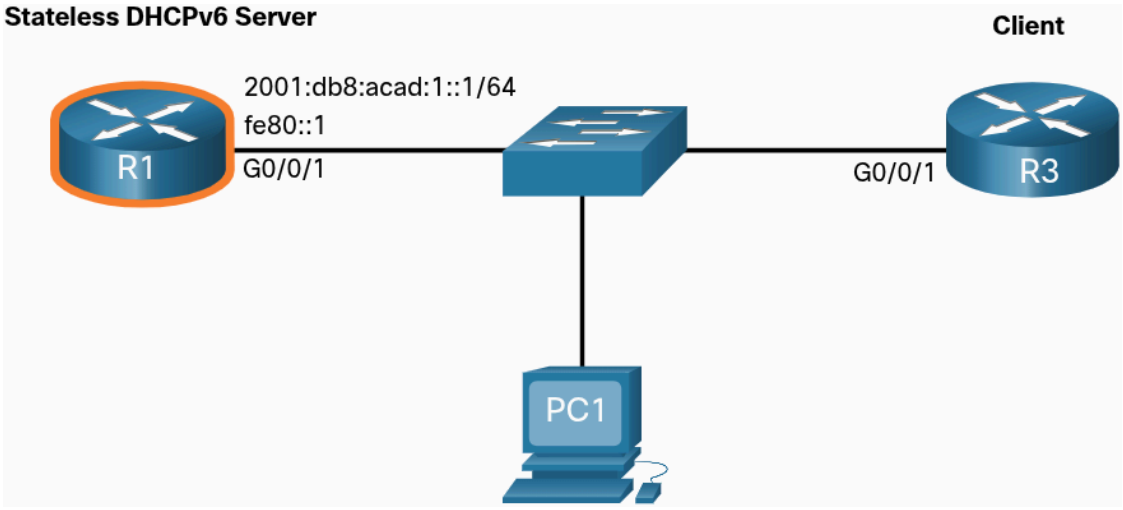

Example explained: R1 will provide SLAAC services for the host IPv6 configuration and DHCPv6 services

Five steps to configure and verify a router as a stateless DHCPv6 server

Step 1. Enable IPv6 routing
Step 2. Define a DHCPv6 pool name
Step 3. Configure the DHCPv6 pool
Step 4. Bind the DHCPv6 pool to an interface
Step 5. Verify that the hosts have received IPv6 addressing information

Step 1. Enable IPv6 routing

- **ipv6 unicast-routing** command is required to enable IPv6 routing.
    - Not necessary for router to be stateless DHCPv6 server, but required for the router to source ICMPv6 RA messages

```
R1(config)# ipv6 unicast-routing
R1(config)#
```

Step 2. Define a DHCPv6 pool name

- Create the DHCPv6 pool using the **ipv6 dhcp pool** *POOL-NAME* global config command
    - This enters DHDCPv6 pool sub-configuration mode as identified by the **Router(config-dhcpv6)#** prompt

Note - The pool name does not have to be uppercase, but using uppercase name makes it easier to see in configuration

```
R1(config)# ipv6 dhcp pool IPV6-STATELESS
R1(config-dhcpv6)#
```

Step 3. Configure the DHCPv6 pool

- R1 will be configured to provide additional DHCP information including DNS server address and domain name

```
R1(config-dhcpv6)# dns-server 2001:db8:acad:1::254
R1(config-dhcpv6)# domain-name example.com
R1(config-dhcpv6)# exit
R1(config)#
```

Step 4. Bind the DHCPv6 pool to an interface

- DHCPv6 pool has to be bound to the interface using **ipv6 dhcp server** *POOL-NAME* interface config command

The router responds to stateless DHCPv6 requests on this interface with the information contained in the pool

- The O flag needs to be manually changed from 0 to 1 using **ipv6 nd other-config-flag**
- RA messages sent on this interface indicate that additional information is available from a stateless DHCPv6 server
- The A flag is 1 by default, telling clients to use SLAAC to create their own GUA

```
R1(config)# interface GigabitEthernet0/0/1
R1(config-if)# description Link to LAN
R1(config-if)# ipv6 address fe80::1 link-local
R1(config-if)# ipv6 address 2001:db8:acad:1::1/64
R1(config-if)# ipv6 nd other-config-flag
R1(config-if)# ipv6 dhcp server IPV6-STATELESS
R1(config-if)# no shut
R1(config-if)# end
R1#
```

Step 5. Verify hosts received IPv6 addressing information

- Verify stateless DHCP on Windows host, use **ipconfig /all**

Notice output displays PC1 created its IPv6 GUA using 2001:db8:acad:1::/64 prefix
Also notice the default gateway is the IPv6 link-local address of R1

- This confirms PC1 derived its IPv6 configuration from the RA of R1

The highlighted output confirms that PC1 has learned the domain name and DNS server address information from the stateless DHCPv6 server

```
C:\PC1> ipconfig /all
Windows IP Configuration
Ethernet adapter Ethernet0:
   Connection-specific DNS Suffix  . : example.com
   Description . . . . . . . . . . . : Intel(R) 82574L Gigabit
Network Connection
   Physical Address. . . . . . . . . : 00-05-9A-3C-7A-00
   DHCP Enabled. . . . . . . . . . . : Yes
   Autoconfiguration Enabled . . . . : Yes
   IPv6 Address. . . . . . . . . . . :
2001:db8:acad:1:1de9:c69:73ee:ca8c (Preferred)
   Link-local IPv6 Address . . . . . :
```

```
fe80::fb:1d54:839f:f595%21(Preferred)
   IPv4 Address. . . . . . . . . . . : 169.254.102.23
(Preferred)
   Subnet Mask . . . . . . . . . . . : 255.255.0.0
   Default Gateway . . . . . . . . . : fe80::1%6
   DHCPv6 IAID . . . . . . . . . . . : 318768538
   DHCPv6 Client DUID. . . . . . . . :
00-01-00-01-21-F3-76-75-54-E1-AD-DE-DA-9A
   DNS Servers . . . . . . . . . . . : 2001:db8:acad:1::254
   NetBIOS over Tcpip. . . . . . . . : Enabled
C:\PC1>
```

Configure a **Stateless** DHCPv6 **Client**

- A router can also be a DHCPv6 client and get an IPv6 configuration from a DHCPv6 server, such as a router functioning as a DHCPv6 server

R1 is a stateless DHCPv6 server

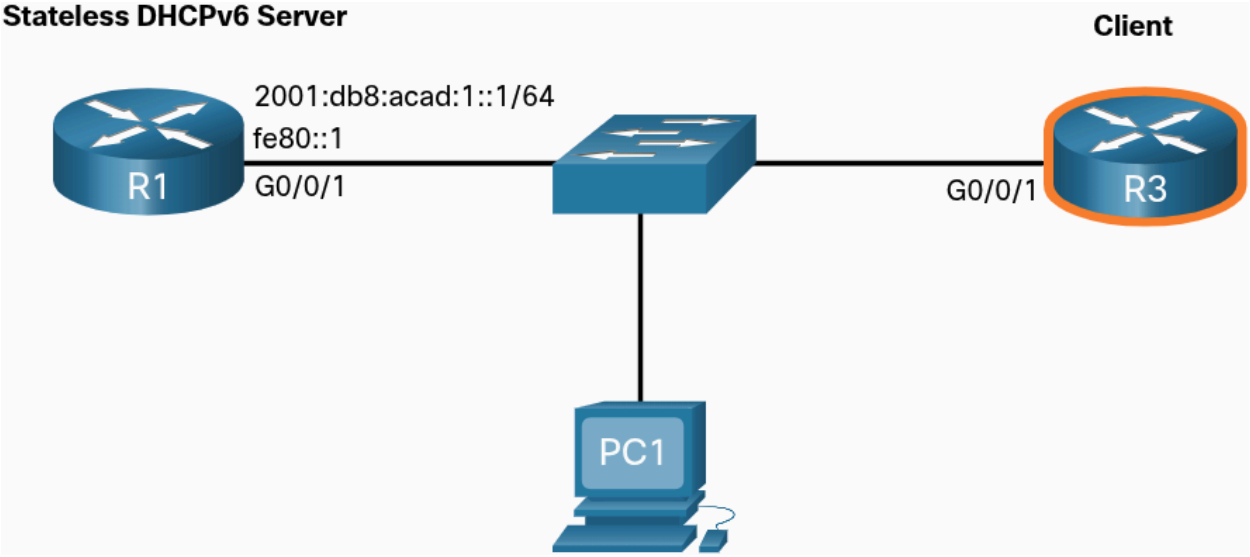

Five Steps to configure and verify a router as stateless DHCPv6 client

Step 1. Enable IPv6 routing
Step 2. Configure the client router to create an LLA
Step 3. Configure the client router to use SLAAC
Step 4. Verify that the client router is assigned a GUA
Step 5. Verify that the client router received other necessary DHCPv6 information

Step 1. Enable IPv6 routing

- DHCPv6 client router needs to have **ipv6 unicast-routing** enabled

```
R3(config)# ipv6 unicast-routing
R3(config)#
```

Step 2. Configure client router to create an LLA

- The client router needs to have a link-local address. An IPv6 link-local address is created on a router interface when a global unicast address is configured
    - Can also be created **without** a GUIA using **ipv6 enable** interface configuration command.

Cisco IOS uses EUI-64 to create a randomized Interface ID

```
R3(config)# interface g0/0/1
R3(config-if)# ipv6 enable
R3(config-if)#
```

Step 3. Configure client router to use SLAAC

- The client router needs to be configured to use SLAAC to create an IPv6 configuration
- **ipv6 address autoconfig** command enables the automatic configuration of IPv6 addressing using SLAAC

```
R3(config-if)# ipv6 address autoconfig
R3(config-if)# end
R3#
```

Step 4. Verify client router is assigned a GUA

- Use **show ipv6 interface brief** command to verify the host configuration

Note - It may take the interface a few seconds to complete process

```
R3# show ipv6 interface brief
GigabitEthernet0/0/0   [up/up]
    unassigned
GigabitEthernet0/0/1   [up/up]
    FE80::2FC:BAFF:FE94:29B1
    2001:DB8:ACAD:1:2FC:BAFF:FE94:29B1
Serial0/1/0            [up/up]
    unassigned
Serial0/1/1            [up/up]
    unassigned
```

```
R3#
```

Step 5. Verify client router received other DHCPv6 information

- The **show ipv6 dhcp interface g0/0/1** command confirms that the DNS and domain names were also learned by R3

```
R3# show ipv6 dhcp interface g0/0/1
GigabitEthernet0/0/1 is in client mode
  Prefix State is IDLE (0)
  Information refresh timer expires in 23:56:06
  Address State is IDLE
  List of known servers:
    Reachable via address: FE80::1
    DUID: 000300017079B3923640
    Preference: 0
    Configuration parameters:
      DNS server: 2001:DB8:ACAD:1::254
      Domain name: example.com
      Information refresh time: 0
  Prefix Rapid-Commit: disabled
  Address Rapid-Commit: disabled
R3#
```

Configure a **Stateful** DHCPv6 **Server**

- Requires that the IPv6 enabled router tells the host to contact a DHCPv6 server to obtain all necessary IPv6 network addressing information

Figure displays, R1 will provide stateful DHCPv6 services to all hosts on local network.
Configuring a stateful DHCPv6 server is similar to configure a stateless server

- Significant difference is that a stateful DHCPv6 server also includes IPv6 addressing information similar to DHCPv4 server

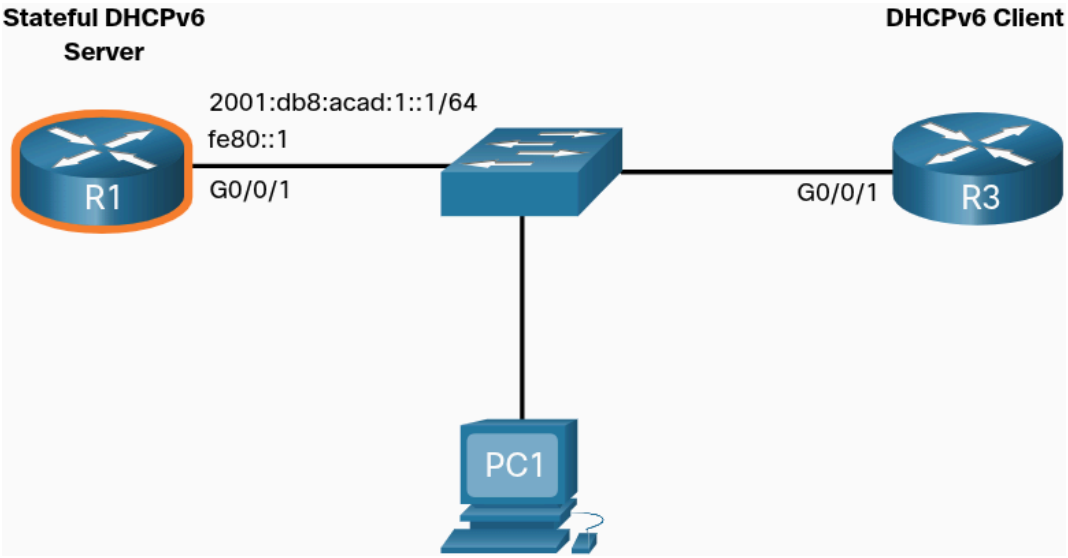

Five steps to configure and verify a router as a stateful DHCPv6 server

Step 1. Enable IPv6 routing
Step 2. Define a DHCPv6 pool name
Step 3. Configure the DHCPv6 pool
Step 4. Bind the DHCPv6 pool to an interface
Step 5. Verify that that hosts have received IPv6 addressing information

Step 1. Enable IPv6 routing

- **ipv6 unicast-routing** command is required to enable IPv6 routing

```
R1(config)# ipv6 unicast-routing
R1(config)#
```

Step 2. Define a DHCPv6 pool name

- Create the DHCPv6 pool using **ipv6 dhcp pool** *POOL-NAME* global config command

```
R1(config)# ipv6 dhcp pool IPV6-STATEFUL
R1(config-dhcpv6)#
```

Step 3. Configure the DHCPv6 pool

R1 will be configured to provide IPv6 addressing, DNS server address, and domain name.

- With stateful DHCPv6, all addressing and other configuration parameters must be assigned by the DHCPv6 server

The **address prefix** command is used to indicate the pool of addresses to be allocated by the server

- Other information provided by the stateful DHCPv6 server typically includes DNS server address and domain name

Setting DNS server to Google's public DNS server

```
R1(config-dhcpv6)# address prefix 2001:db8:acad:1::/64
R1(config-dhcpv6)# dns-server 2001:4860:4860::8888
R1(config-dhcpv6)# domain-name example.com
R1(config-dhcpv6)#
```

Step 4. Bind the DHCPv6 pool to an interface

The DHCPv6 pool has to be bound to the interface using **ipv6 dhcp server** *POOL-NAME* interface config command

- The M flag is manually changed from 0 to 1 using the interface command **ipv6 nd managed-config-flag**
- The A flag is manually changed from 1 to 0 using the interface command **ipv6 nd prefix default no-autoconfig**
    - Can be left at 1, but some client operating systems such as Windows will create a GUA using SLAAC and get sa GUA from a stateful DHCPv6 server
    - Setting the A lag to 0 tells the client not to use SLAAC to create a GUA
- The **ipv6 dhcp server** command binds the DHCPv6 pool to the interface. R1 will not respond with the information contained in the pool when it receives stateful DHCPv6 requests on this interface

Note - You can use **no ipv6 nd managed-config-flag** command to set M flag back to its default of 0

- The **no ipv6 nd prefix default no-autoconfig** command sets the A flag back to its default of 1

```
R1(config)# interface GigabitEthernet0/0/1
R1(config-if)# description Link to LAN
R1(config-if)# ipv6 address fe80::1 link-local
R1(config-if)# ipv6 address 2001:db8:acad:1::1/64
R1(config-if)# ipv6 nd managed-config-flag
R1(config-if)# ipv6 nd prefix default no-autoconfig
R1(config-if)# ipv6 dhcp server IPV6-STATEFUL
R1(config-if)# no shut
R1(config-if)# end
```

```
R1#
```

Step 5. Verify hosts received IPv6 addressing information

- Verify on Windows with **ipconfig /all**

```
C:\PC1> ipconfig /all
Windows IP Configuration
Ethernet adapter Ethernet0:
   Connection-specific DNS Suffix  . : example.com
   Description . . . . . . . . . . . : Intel 82574L Gigabit Network Connection
   Physical Address. . . . . . . . . : 00-05-9A-3C-7A-00
   DHCP Enabled. . . . . . . . . . . : Yes
   Autoconfiguration Enabled . . . . : Yes
   IPv6 Address. . . . . . . . . . . : 2001:db8:acad:1:a43c:fd28:9d79:9e42 (Preferred)
   Lease Obtained. . . . . . . . . . : Saturday, September 27, 2019, 10:45:30 AM
   Lease Expires . . . . . . . . . . : Monday, September 29, 2019 10:05:04 AM
   Link-local IPv6 Address . . . . . : fe80::192f:6fbc:9db:b749%6(Preferred)
   Autoconfiguration IPv4 Address. . : 169.254.102.73 (Preferred)
   Subnet Mask . . . . . . . . . . . : 255.255.0.0
   Default Gateway . . . . . . . . . : fe80::1%6
   DHCPv6 IAID . . . . . . . . . . . : 318768538
   DHCPv6 Client DUID. . . . . . . . : 00-01-00-01-21-F3-76-75-54-E1-AD-DE-DA-9A
   DNS Servers . . . . . . . . . . . : 2001:4860:4860::8888
   NetBIOS over Tcpip. . . . . . . . : Enabled
C:\PC1>
```

Configure a **Stateful** DHCPv6 **Client**

- Client router needs to have **ipv6 unicast-routing** enabled and an IPv6 link-local address to send and receive IPv6 messages

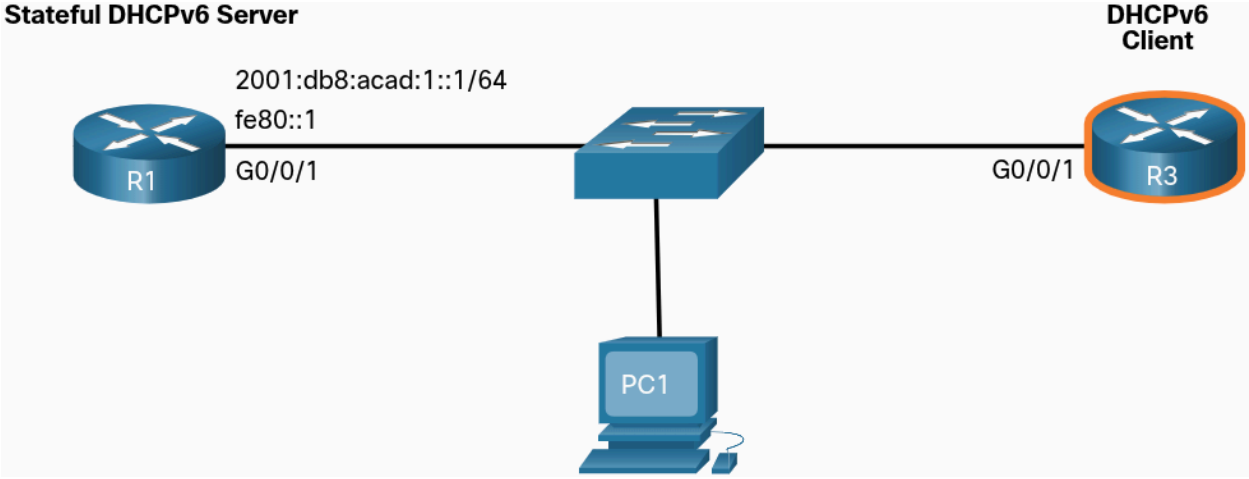

Five steps to configure and verify router as stateful DHCPv6 client

Step 1. Enable IPv6 routing
Step 2. Configure the client router to create an LLA
Step 3. Configure the client router to use DHCPv6
Step 4. Verify that the client router is assigned a GUA
Step 5. Verify that the client router received other necessary DHCPv6 information

Step 1. Enable IPv6 routing

- Client router needs to have **ipv6 unicast-routing** enabled

```
R3(config)# ipv6 unicast-routing
R3(config)#
```

Step 2. Configure client router to create an LLA

- The **ipv6 enable** command is configured on the R3 G0/0/1 interface
- This enables the router to create an IPv6 LLA **without** needing a GUA

```
R3(config)# interface g0/0/1
R3(config-if)# ipv6 enable
R3(config-if)#
```

Step 3. Configure client router to use DHCPv6

- The **ipv6 address dhcp** command configures R3 to solicit its IPv6 addressing information from a DHCPv6 server

```
R3(config-if)# ipv6 address dhcp
R3(config-if)# end
R3#
```

Step 4. Verify client router is assigned a GUA

- Use **show ipv6 interface brief** command to verify host configuration

```
R3# show ipv6 interface brief
GigabitEthernet0/0/0    [up/up]
    unassigned
GigabitEthernet0/0/1    [up/up]
    FE80::2FC:BAFF:FE94:29B1
    2001:DB8:ACAD:1:B4CB:25FA:3C9:747C
Serial0/1/0             [up/up]
    unassigned
Serial0/1/1             [up/up]
    unassigned
R3#
```

Step 5. Verify client router received other DHCPv6 information

- The **show ipv6 dhcp interface g0/0/1** command confirms DNS and domain names were learned by R3

```
R3# show ipv6 dhcp interface g0/0/1
GigabitEthernet0/0/1 is in client mode
  Prefix State is IDLE
  Address State is OPEN
  Renew for address will be sent in 11:56:33
  List of known servers:
    Reachable via address: FE80::1
    DUID: 000300017079B3923640
    Preference: 0
    Configuration parameters:
      IA NA: IA ID 0x00060001, T1 43200, T2 69120
        Address: 2001:DB8:ACAD:1:B4CB:25FA:3C9:747C/128
                 preferred lifetime 86400, valid lifetime 172800
                 expires at Sep 29 2019 11:52 AM (172593 seconds)
      DNS server: 2001:4860:4860::8888
      Domain name: example.com
      Information refresh time: 0
  Prefix Rapid-Commit: disabled
  Address Rapid-Commit: disabled
R3#
```

DHCPv6 Server Verification Commands

Verify DHCPv6 operation on router

- **show ipv6 dhcp pool**
- **show ipv6 dhcp binding**

show ipv6 dhcp pool

- Verifies the name of the DHCPv6 pool and its parameters
- Also identifies the number of active clients

Example:

- The IPV6-STATEFUL pool currently has 2 clients, which reflects PC1 and R3 receiving their IPv6 global unicast address

* When a router is providing stateful DHCPv6 services, it also maintains the database of assigned IPv6 addresses

```
R1# show ipv6 dhcp pool
DHCPv6 pool: IPV6-STATEFUL
  Address allocation prefix: 2001:DB8:ACAD:1::/64 valid 172800 preferred 86400
(2 in use, 0 conflicts)
  DNS server: 2001:4860:4860::8888
  Domain name: example.com
  Active clients: 2
R1#
```

show ipv6 dhcp binding

- Display the IPv6 link-local address of the client and the global unicast address assigned by the server

Output displays current stateful binding on R1. The first client is PC1 and second client is R3

- Information is maintained by a stateful DHCPv6 server
    - Stateless DHCPv6 server would **not** maintain this information

```
R1# show ipv6 dhcp binding
Client: FE80::192F:6FBC:9DB:B749
  DUID: 0001000125148183005056B327D6
  Username : unassigned
  VRF : default
  IA NA: IA ID 0x03000C29, T1 43200, T2 69120
    Address: 2001:DB8:ACAD:1:A43C:FD28:9D79:9E42
             preferred lifetime 86400, valid lifetime 172800
             expires at Sep 27 2019 09:10 AM (171192 seconds)
Client: FE80::2FC:BAFF:FE94:29B1
  DUID: 0003000100FCBA9429B0
  Username : unassigned
  VRF : default
  IA NA: IA ID 0x00060001, T1 43200, T2 69120
    Address: 2001:DB8:ACAD:1:B4CB:25FA:3C9:747C
             preferred lifetime 86400, valid lifetime 172800
             expires at Sep 27 2019 09:29 AM (172339 seconds)
R1#
```

Configure a DHCPv6 Relay Agent

- Created for when DHCPv6 server is located on different network than client
- Configuration similar to configuration of IPv4 router as DHCPv4 relay

Figure: R3 is configured as stateful DHCPv6 server.

- PC1 is on 2001:db8:acad:2::/64 and requires services of stateful DHCPv6 server to acquire IPv6 configuration
- R1 needs to be configured as DHCPv6 Relay Agent

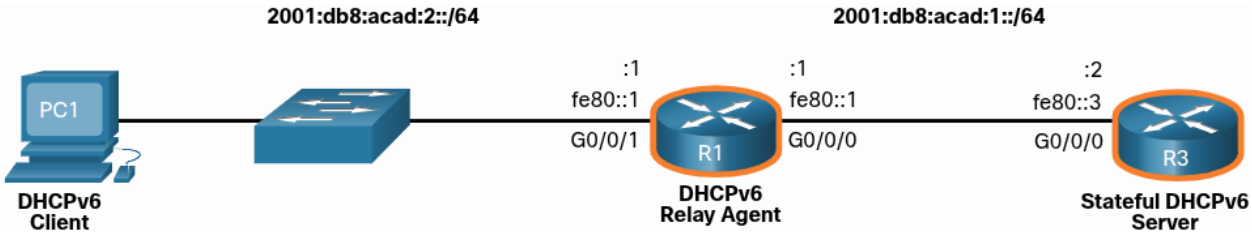

Command syntax to configure router as DHCPv6 relay agent:
Router(config-if)# ipv6 dhcp relay destination *ipv6-address* [*interface-type interface-number*]

Command is configured on the interface facing the DHCPv6 clients and specifies the DHCPv6 server address and egress interface to reach server.

- ** The egress interface is only required when the next-hop address is LLA

```
R1(config)# interface gigabitethernet 0/0/1
R1(config-if)# ipv6 dhcp relay destination 2001:db8:acad:1::2
G0/0/0
R1(config-if)# exit
R1(config)#
```

Verify the DHCPv6 Relay Agent

- **show ipv6 dhcp interface**
- **show ipv6 dhcp binding**

Windows:

- **ip config /all**

show ipv6 dhcp interface

```
R1# show ipv6 dhcp interface
GigabitEthernet0/0/1 is in relay mode
  Relay destinations:
    2001:DB8:ACAD:1::2
    2001:DB8:ACAD:1::2 via GigabitEthernet0/0/0
R1#
```

show ipv6 dhcp binding

- On R3, use this command to verify if any hosts have been assigned an IPv6 config

Notice that a client link-local address has been assigned an IPv6 GUA. Assume this is PC1

```
R3# show ipv6 dhcp binding
Client: FE80::5C43:EE7C:2959:DA68
  DUID: 0001000124F5CEA2005056B3636D
  Username : unassigned
  VRF : default
  IA NA: IA ID 0x03000C29, T1 43200, T2 69120
    Address: 2001:DB8:ACAD:2:9C3C:64DE:AADA:7857
```

```
            preferred lifetime 86400, valid lifetime 172800
            expires at Sep 29 2019 08:26 PM (172710 seconds)
R3#
```

ipconfig /all to confirm on windows

```
C:\PC1> ipconfig /all
Windows IP Configuration
Ethernet adapter Ethernet0:
   Connection-specific DNS Suffix  . : example.com
   Description . . . . . . . . . . . : Intel(R) 82574L Gigabit
Network Connection
   Physical Address. . . . . . . . . : 00-05-9A-3C-7A-00
   DHCP Enabled. . . . . . . . . . . : Yes
   Autoconfiguration Enabled . . . . : Yes
   IPv6 Address. . . . . . . . . . . :
2001:db8:acad:2:9c3c:64de:aada:7857 (Preferred)
   Link-local IPv6 Address . . . . . :
fe80::5c43:ee7c:2959:da68%6(Preferred)
   Lease Obtained  . . . . . . . . . : Saturday, September 27,
2019, 11:45:30 AM
   Lease Expires . . . . . . . . . . : Monday, September 29, 2019
11:05:04 AM
   IPv4 Address. . . . . . . . . . . : 169.254.102.73 (Preferred)
   Subnet Mask . . . . . . . . . . . : 255.255.0.0
   Default Gateway . . . . . . . . . : fe80::1%6
   DHCPv6 IAID . . . . . . . . . . . : 318768538
   DHCPv6 Client DUID. . . . . . . . :
00-01-00-01-21-F3-76-75-54-E1-AD-DE-DA-9A
   DNS Servers . . . . . . . . . . . : 2001:4860:4860::8888
   NetBIOS over Tcpip. . . . . . . . : Enabled
C:\PC1>
```

8.5.2 What did I learn in this module?

IPv6 GUA Assignment

On a router, an IPv6 global unicast addresses (GUA) is manually configured using the ipv6 address *ipv6-address/prefix-length* interface configuration command. When automatic IPv6 addressing is selected, the host will attempt to automatically obtain and configure IPv6 address information on the interface. The IPv6 link-local address is automatically created by the host when it boots and the Ethernet interface is active. By default, an IPv6-enabled router advertises its IPv6 information enabling a host to dynamically create or acquire its IPv6 configuration. The IPv6 GUA can be assigned dynamically using stateless and stateful services. The decision of how

a client will obtain an IPv6 GUA depends on the settings within the RA message. An ICMPv6 RA message includes three flags to identify the dynamic options available to a host:

- A flag - This is the Address Autoconfiguration flag. Use SLAAC to create an IPv6 GUA.
- O flag - This is the Other Configuration flag. Get Other information from a stateless DHCPv6 server.
- M flag - This is the Managed Address Configuration flag. Use a stateful DHCPv6 server to obtain an IPv6 GUA.

SLAAC

The SLAAC method enables hosts to create their own unique IPv6 global unicast address without the services of a DHCPv6 server. SLAAC, which is stateless, uses ICMPv6 RA messages to provide addressing and other configuration information that would normally be provided by a DHCP server. SLAAC can be deployed as SLAAC only, or SLAAC with DHCPv6. To enable the sending of RA messages, a router must join the IPv6 all-routers group using the ipv6 unicast-routing global config command. Use the show ipv6 interface command to verify if a router is enabled. The SLAAC only method is enabled by default when the ipv6 unicast-routing command is configured. All enabled Ethernet interfaces with an IPv6 GUA configured will start sending RA messages with the A flag set to 1, and the O and M flags set to 0. The A = 1 flag suggests to the client to create its own IPv6 GUA using the prefix advertised in the RA. The O =0 and M=0 flags instructs the client to use the information in the RA message exclusively. A router sends RA messages every 200 seconds. However, it will also send an RA message if it receives an RS message from a host. Using SLAAC, a host typically acquires its 64-bit IPv6 subnet information from the router RA. However, it must generate the remainder 64-bit interface identifier (ID) using one of two methods: randomly generated, or EUI-64. The DAD process is used by a host to ensure that the IPv6 GUA is unique. DAD is implemented using ICMPv6. To perform DAD, the host sends an ICMPv6 NS message with a specially constructed multicast address, called a solicited-node multicast address. This address duplicates the last 24 bits of IPv6 address of the host.

DHCPv6

The host begins the DHCPv6 client/server communications after stateless DHCPv6 or stateful DHCPv6 is indicated in the RA. Server to client DHCPv6 messages use UDP destination port 546, while client to server DHCPv6 messages use UDP destination port 547. The stateless DHCPv6 option informs the client to use the information in the RA message for addressing, but additional configuration parameters are available from a DHCPv6 server. This is called stateless DHCPv6 because the server is not maintaining any client state information. Stateless DHCPv6 is enabled on a router interface using the ipv6 nd other-config-flag interface configuration command. This sets the O flag to 1. In stateful DHCPv6, the RA message tells the client to obtain all addressing information from a stateful DHCPv6 server, except the default gateway address which is the source IPv6 link-local address of the RA. It is called stateful because the DHCPv6 server maintains IPv6 state information. Stateful DHCPv6 is enabled on a router interface using the ipv6 nd managed-config-flag interface configuration command. This sets the M flag to 1.

Configure DHCPv6 Server

A Cisco IOS router can be configured to provide DHCPv6 server services as one of the following three types: DHCPv6 server, DHCPv6 client, or DHCPv6 relay agent. The stateless DHCPv6 server option requires that the router advertise the IPv6 network addressing information in RA messages. A router can also be a DHCPv6 client and get an IPv6 configuration from a DHCPv6 server. The stateful DHCP server option requires that the IPv6-enabled router tells the host to contact a DHCPv6 server to acquire all required IPv6 network addressing information. For a client router to be a DHCPv6 router, it needs to have ipv6 unicast-routing enabled and an IPv6 link-local address to send and receive IPv6 messages. Use the show ipv6 dhcp pool and show ipv6 dhcp binding commands to verify DHCPv6 operation on a router. If the DHCPv6 server is located on a different network than the client, then the IPv6 router can be configured as a DHCPv6 relay agent using the ipv6 dhcp relay destination *ipv6-address [interface-type interface-number]* command. This command is configured on the interface facing the DHCPv6 clients and specifies the DHCPv6 server address and egress interface to reach the server. The egress interface is only required when the next-hop address is an LLA. Verify the DHCPv6 relay agent is operational with the show ipv6 dhcp interface and show ipv6 dhcp binding commands.

First Hop Redundancy Protocols

Default Gateway Limitations

If a router or router interface (that serves as a default gateway) fails, the hosts configured with that default gateway are isolated from outside networks.

First Hop Redundancy Protocols (FHRPs)

- A mechanism to provide alternate default gateways in switched networks where two or more routers are connected to the same VLANs

In a switched network, each client receives only **one** default gateway. There is no way to use a secondary gateway, even if a second path exists to carry packets off the local segment.

Figure below:

- R1 is responsible for routing packets from PC1. If R1 becomes unavailable, the routing protocols can dynamically converge.
- R2 now routes packets from outside networks that would have gone through R1
    - However, traffic from inside network associated with R1 configured with R1 as the default gateway, are still sent to R1 and **dropped**

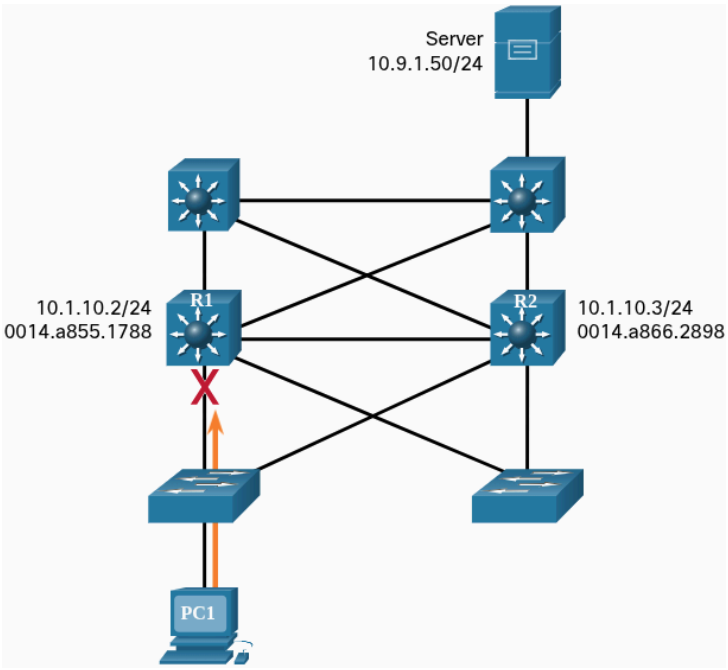

PC1 is unable to reach the default gateway

End devices are typically configured with a single IPv4 address for default gateway.

- This address does **not** change when the network topology changes

If the default gateway IPv4 address cannot be reached, the local device is unable to send packets off the local network segment, effectively disconnecting it from other networks.

- Even if a redundant router exist that could serve as a default gateway for that segment, there is no dynamic method by which these devices can determine the address of a new default gateway

Note - IPv6 devices receive their default gateway address dynamically from the ICMPv6 Router Advertisement. However, IPv6 devices benefit with a faster failover to the new default gateway when using FHRP.

Router Redundancy

One way to prevent a single point of failure at the default gateway is to implement a virtual router.

- Multiple routers are configured to work together to present the illusion of a single router to the hosts on the LAN

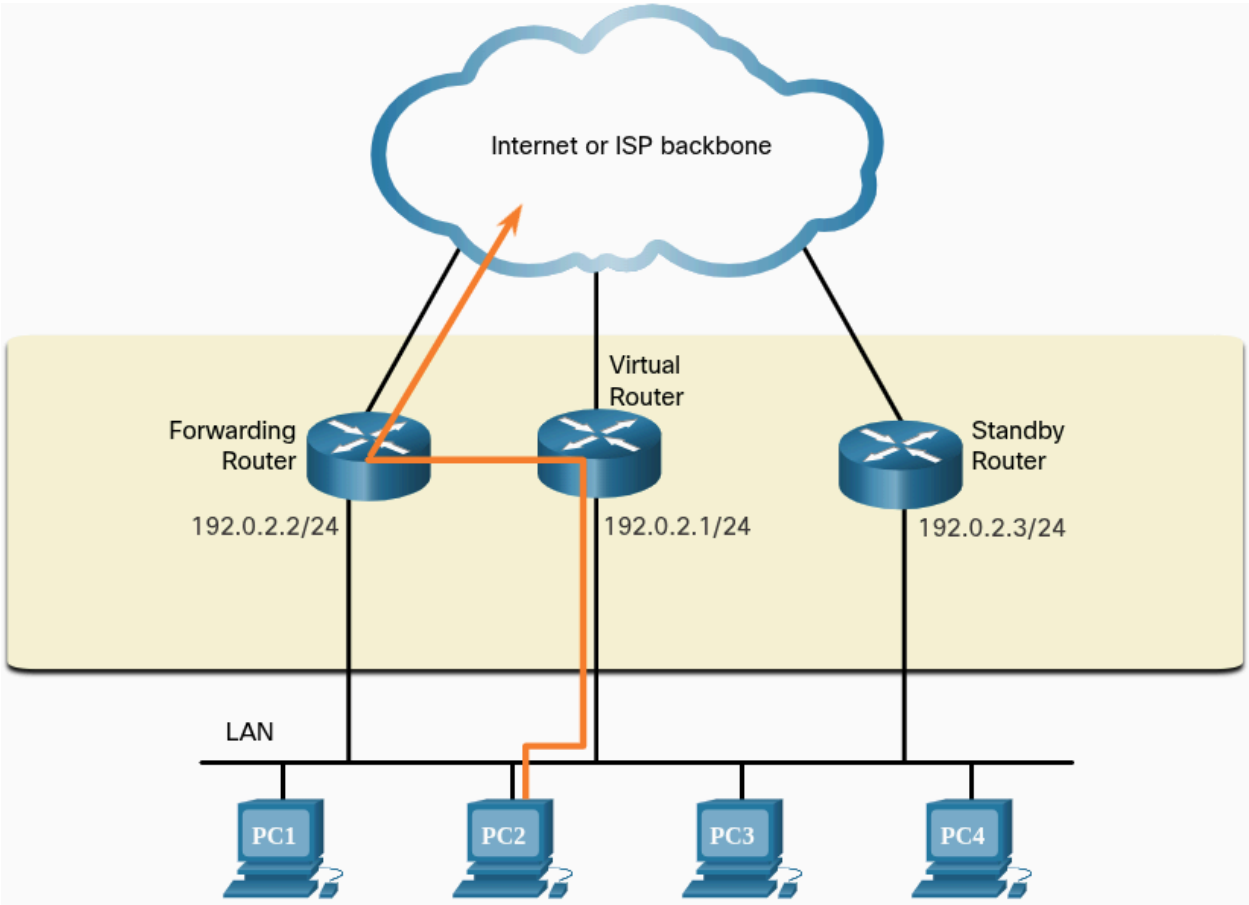

The IPv4 address of the **virtual router** is configured as the **default gateway**.

When frames are sent from host devices to the default gateway, the hosts use **ARP** to resolve the MAC address that is associated with the IPv4 address of the default gateway.

The ARP resolution returns the MAC address of the virtual router.

Frames that are sent to the MAC address of the virtual router can then be physically processed by the currently active router within the virtual router group
A protocol is used to identify two or more routers as the devices that are responsible for processing frames that are sent to the MAC or IP address of a single virtual router. Host devices send traffic to the address of the virtual router. The physical router that forwards this traffic is transparent to the host devices.

A redundancy protocol provides the mechanism for determining which router should take the active role in forwarding traffic. It also determines when the forwarding role must be taken over by a standby router. The transition from one forwarding router to another is transparent to the end devices.

The ability of a network to dynamically recover from the failure of a device acting as a default gateway is known as **first-hop redundancy**.

Steps for Router Failover

When the active router fails, the redundancy protocol transitions the standby router to the new active router role

1. The standby router stops seeing Hello messages from the forwarding router
2. The standby router assumes the role of the forwarding router
3. Because the new forwarding router assumes both the IPv4 and MAC addresses of the virtual router, the host devices see no disruption in service

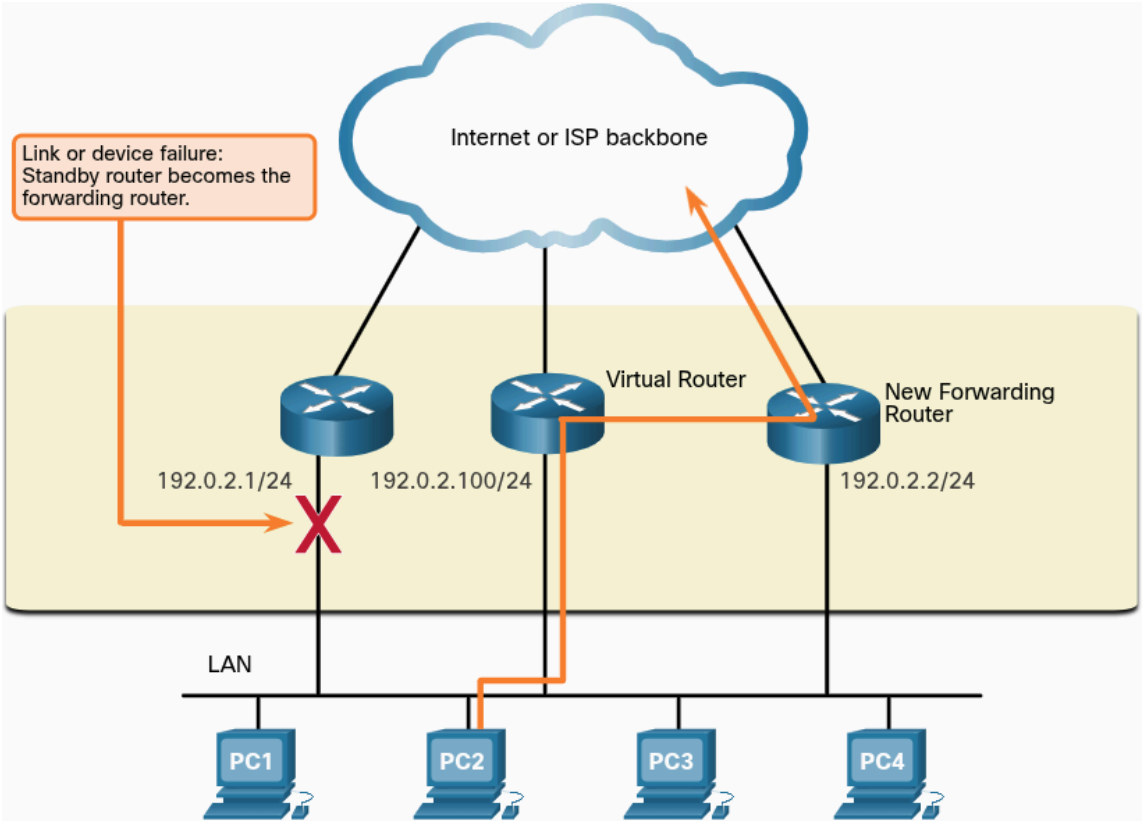

FHRP Options

FHRP Options	Description
Hot Standby Router Protocol (HSRP)	HSRP is a Cisco-proprietary FHRP that is designed to allow for transparent failover of a first-hop IPv4 device. HSRP is used in a group of routers for selecting an active device and a standby device.

	In a group of device interfaces, the active device is the device that is used for routing packets; the function of the HSRP standby router is to monitor the operational status of the HSRP group and to quickly assume packet-forwarding responsibility if the active router fails
HSRP for IPv6	This is a Cisco-proprietary FHRP that provides the same functionality of HSRP, but in an IPv6 environment. An HSRP IPv6 group has a virtual MAC address derived from the HSRP group number and a virtual IPv6 link-local address derived from the HSRP virtual MAC Address. Periodic router advertisements (RAs) are sent for the HSRP virtual IPv6 link-local address when the HSRP group is active. When the group becomes active, these RAs stop after a final RA is sent
Virtual Router Redundancy Protocol version 2 (VRRPv2)	This is a non-proprietary election protocol that dynamically assigns responsibility for one or more virtual routers to the VRRP routers on an IPv4 LAN. This allows several routers on a multiaccess link to use the same virtual IPv4 address. A VRRP router is configured to run the VRRP protocol in conjunction with one or more other routers attached to a LAN. In a VRRP configuration, one router is elected as the virtual router master, with the other routers acting as backups, in case the virtual router master fails
VRRPv3	This provides the capability to support IPv4 and IPv6 addresses. The VRRPv3 works in multi-vendor environments and is more scalable than VRRPv2
Gateway Load Balancing Protocol (GLBP)	This is a Cisco-proprietary FHRP that protects data traffic from a failed router or circuit, like HSRP and VRRP, while also allowing load balancing (also called load sharing) between a group of redundant routers

GLBP for IPv6	This is a Cisco-proprietary FHRP that provides the same functionality of GLBP, but in an IPv6 environment. GLBP for IPv6 provides automatic router backup for IPv6 hosts configured with a single default gateway on a LAN. Multiple first-hop routers on the LAN combine to offer a single virtual first-hop IPv6 router while sharing the IPv6 packet forwarding load
ICMP Router Discovery Protocol (IRDP)	Specified in RFC 1256, IRDP is a legacy FHRP solution. IRDP allows IPv4 hosts to locate routers that provide IPv5 connectivity to other (nonlocal) IP networks

HSRP Priority and Preemption

The role of active and standby routers is determined during the HSRP election process

- By default, the router with the **numerically highest IPv4 address** is elected as the active router
- It is always better to control how your network will operate under normal conditions

HSRP Priority

- Can be used to determine the active router
- The router with the **highest HSRP priority** will become the active router
    - By **default**, the HSRP is **100**
    - If the priorities are **equal**, the router with the **numerically highest IPv4 address** is elected as the active router

To configure router to be active router, use **standby priority** interface command

- The range of HSRP priority is **0** to **255**

HSRP Preemption

- By default, after a router becomes the active router, it will **remain** the active router **even if** another router comes online with a **higher HSRP priority**

To force a new HSRP election process to take place when a higher priority router comes online

- Preemption must be enabled using the **standby preempt** interface command

Preemption

- The ability of an HSRP router to trigger the re-election process
- With **preemption enabled**, a router that comes online with a **higher HSRP priority** will **assume** the **role** of the active router

- Only allows a router to become the active router **if** it has a higher priority
    - A router enabled for preemption, with equal priority but a higher IPv5 address will not preempt an active router

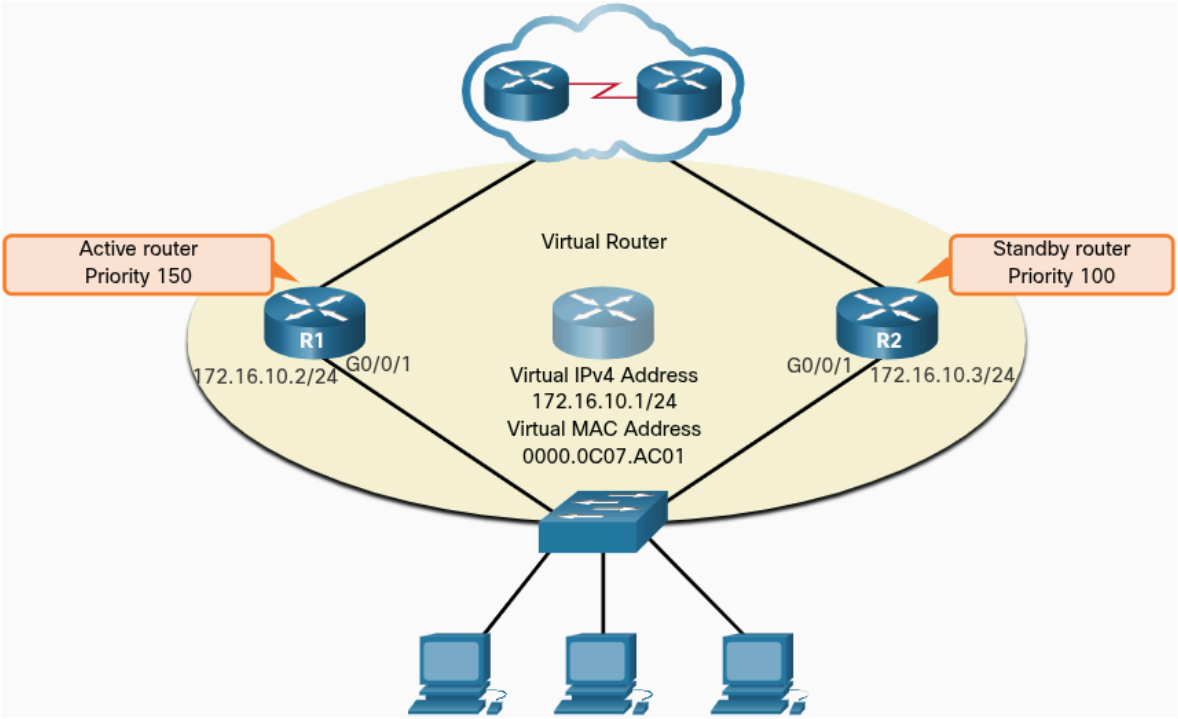

R1 has been configured with the HSRP priority of 150 while R2 has the default HSRP priority of 100.
Preemption has been enabled on R1. With a higher priority, R1 is the **active** router and R2 is the **standby** router.
Due power failure affecting **only** R1, the active router is no longer available and the standby router, R2, assumes the role of the active router. After power is restored, R1 comes back online.

- R1 has higher priority and preemption is enabled, it will force a new election process. R1 will re-assume the role of active router and R2 falls back to role of standby router

Note - with **preemption disabled**, the router that boots up **first** will become active router if there are no other routers online during election process

HSRP States and Timers

A router can either be the **active** HSRP router responsible for forwarding traffic for the segment, or it can be a **passive** HSRP router on standby, ready to assume the active role if the active router fails.

When an **interface** is **configured with** HSRP or is **first activated with** an existing HSRP configuration, the router sends and receives HSRP **hello packets** to begin the process of determining which state it will assume in the HSRP group

HSRP State	Description
Initial	This state is entered through a configuration change or when an interface first becomes available
Learn	The router has not determined the virtual IP address and has not yet seen a hello message from the active route. In this state, the router waits to hear from the active router
Listen	The router knows the virtual IP address, but the router is neither the active router nor the standby router. It listens for hello messages from those routers
Speak	The router sends periodic hello messages and actively participates in the election of the active and/or standby router
Standby	The router is a candidate to become the next active router and sends periodic hello messages
Active	The router won the election

The active and standby HSRP routers send hello packets to the HSRP group multicast address every **3 seconds** by default.
d
The standby router will become active if it does not receive a hell message from the active router after **10 seconds**

- Timer settings can be set lower to speed up the failover or preemption
- To avoid increased CPU usage and unnecessary standby state changes, do **not** set hell timer **below 1 second** or the **hold timer below 4 seconds**

9.3.1 What did I learn in this module?

First Hop Redundancy Protocols

If a router or router interface that serves as a default gateway fails, the hosts configured with that default gateway are isolated from outside networks. FHRP provides alternate default gateways in switched networks where two or more routers are connected to the same VLANs. One way to prevent a single point of failure at the default gateway, is to implement a virtual router. With a virtual router, multiple routers

are configured to work together to present the illusion of a single router to the hosts on the LAN. When the active router fails, the redundancy protocol transitions the standby router to the new active router role. These are the steps that take place when the active router fails:

1. The standby router stops seeing Hello messages from the forwarding router.
2. The standby router assumes the role of the forwarding router.
3. Because the new forwarding router assumes both the IPv4 and MAC addresses of the virtual router, the host devices see no disruption in service.

The FHRP used in a production environment largely depends on the equipment and needs of the network. These are the options available for FHRPs:

- HSRP and HSRP for IPv6
- VRRPv2 and VRRPv3
- GLBP and GLBP for IPv6
- IRDP

HSRP

HSRP is a Cisco-proprietary FHRP designed to allow for transparent failover of a first-hop IP device. HSRP is used in a group of routers for selecting an active device and a standby device. In a group of device interfaces, the active device is the device that is used for routing packets; the standby device is the device that takes over when the active device fails, or when pre-set conditions are met. The function of the HSRP standby router is to monitor the operational status of the HSRP group and to quickly assume packet-forwarding responsibility if the active router fails. The router with the highest HSRP priority will become the active router. Preemption is the ability of an HSRP router to trigger the re-election process. With preemption enabled, a router that comes online with a higher HSRP priority will assume the role of the active router. HSRP states include initial, learn, listen, speak, and standby.

Endpoint Security

Network Attacks Today

- **Distributed Denial of Service (DDoS)** - Coordinated attack from many devices, called zombies, with the intention of degrading or halting public access to an organization's website and resources
- **Data Breach** - An attack in which an organization's data servers or hosts are compromised to steal confidential information
- **Malware** - An attack in which an organization's hosts are infected with malicious software that cause a variety of problems
    - Ransomware such as WannaCry– Encrypts the data on a host and locks access to it until a ransom is paid

Network Security Devices

VPN-Enabled Router

- Provides a secure connection to remote users across a public network and into the enterprise network
    - VPN services can be integrated into the firewall

Next Generation Firewall (NGFW)

- Provides stateful packet inspection, application visibility and control
    - A next-generation intrusion prevention system (NGIPS)
    - Advanced malware protection (AMP)
    - URL filtering

Network Access Control (NAC)

- Includes authentication, authorization, and account (AAA) services
- Might be incorporated into an appliance that can manage access policies across a wide variety of users and device types
    - Example: Cisco Identity Services Engine (ISE)

Endpoint Protection

- Switches
- Wireless LAN Controllers (WLCs)
- Access Points (AP)
    - Most susceptible to LAN-related attacks

If internal host is infiltrated, it can become a starting point for a threat actor to gain access to critical system devices

Endpoints

- Hosts commonly consist of laptops, desktops, services, IP phones, and BYODs
- Susceptible to malware-related attacks that originate through emails or web browsing
- Typically uses traditional host-based security features such as antivirus/ antimalware, host-based wireless, and host-based intrusion prevention systems (HIPSs)

Today, endpoints are best protected by combination of NAC, host-based AMP software, email security appliance (ESA), and web security appliance (WSA)

Advanced Malware Protection (AMP) products include endpoint solutions such as Cisco AMP for endpoints

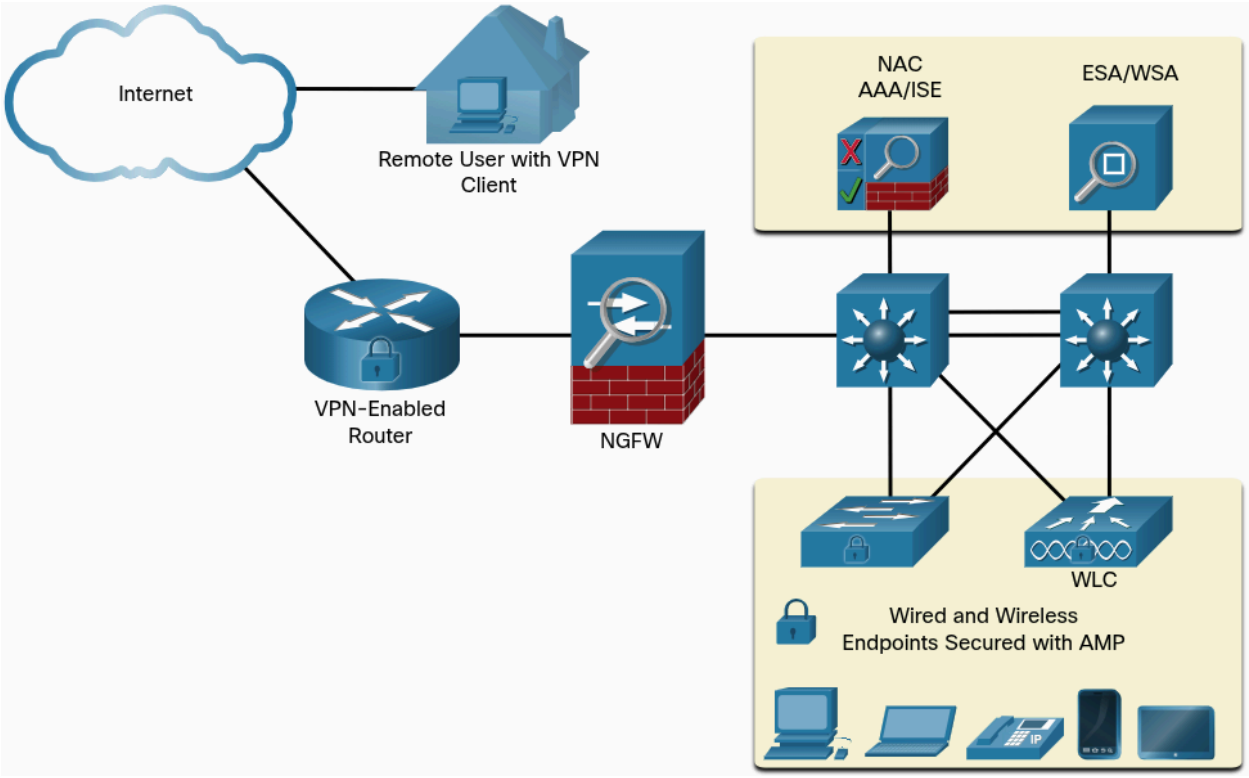

Cisco Email Security Appliance

Phishing

- Particularly virulent form of spam
- Entices user to click a link or open an attachment
- Spear phishing targets high-profile employees or executives that may have elevated login credentials
- 95% of all attacks on enterprise network

Cisco ESA

- Device that is designed to monitor Simple Mail Transfer Protocol (SMTP)
- Constantly updated by real-time feeds from Cisco Talos
    - Detects and correlates threats and solutions by using a worldwide database monitoring system
    - Pulled by Cisco ESA every **three** to **five minutes**
- Block known threats
- Remediate against stealth malware that evaded initial detection
- Discard emails with bad links
- Block access to the newly infected sites
- Encrypt content in outgoing email to prevent data loss

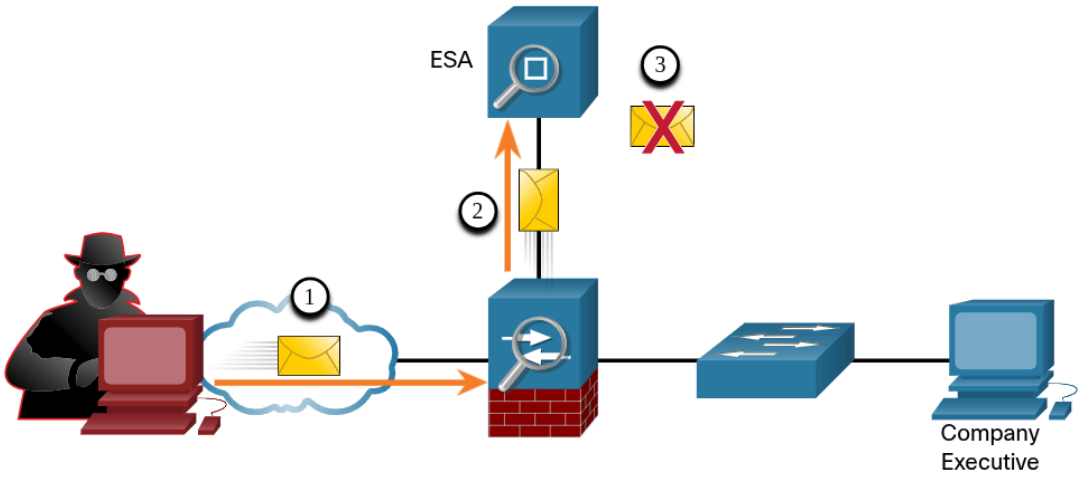

1. Threat actor sends a phishing attack to an important host on the network
2. The firewall forwards all email to the ESA
The ESA analyzes the email, logs it, and if it is malware, discards it

Cisco Web Security Appliance

- Mitigation technology for web-based threats
- Helps organization address the challenges of securing and controlling web traffic
- Combines advanced malware protection, application visibility and control, acceptable use policy controls, and reporting

Cisco WSA provides complete control over how users access the internet

- Certain features and applications, such as chat, messaging, video and audio, can be allowed, restricted with time and bandwidth limits, or blocked
- Can perform blacklisting of URLs, URL-filtering, malware scanning, URL categorization, Web application filtering, and encryption and decryption of web traffic

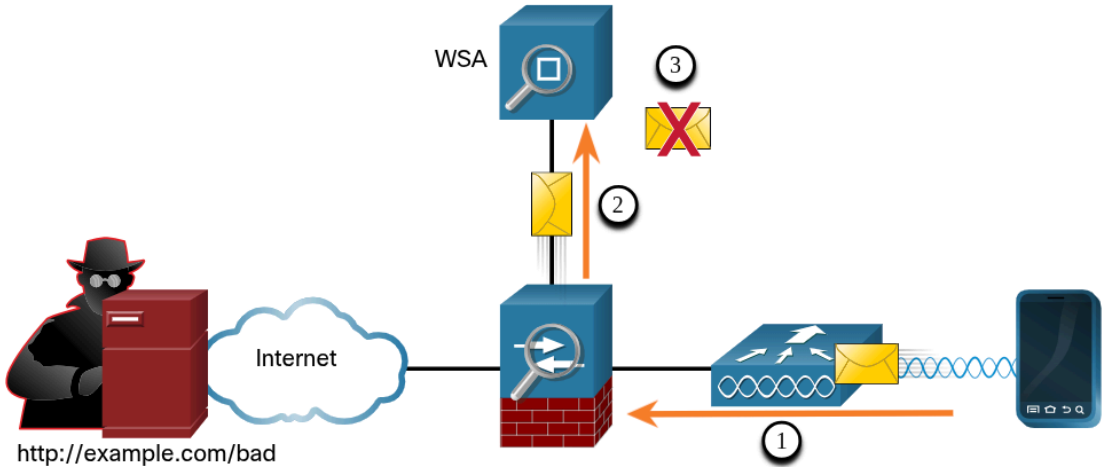

1. A user attempts to connect to a website
2. The firewall forwards the website request to the WSA
3. The WSA evaluates the URL and determines it is a known blacklisted site. The WSA discards the packet and sends an access denied message to the user

Access Control

Authentication with a Local Password

- Simplest method of remote access authentication
- Configure a login and password combination on console, vty lines, and aux ports
- Weakest and least secure
    - Provides no accountability and password is sent in plaintext

```
R1(config)# line vty 0 4
R1(config-line)# password ci5c0
R1(config-line)# login
```

SSH is a more secure form of remote access

- Requires a username and password, both are encrypted during transmission
- Username and password can be authenticated by local database method
- Provides more accountability bc username is recorded when a user logs in

```
R1(config)# ip domain-name example.com
R1(config)# crypto key generate rsa general-keys modulus 2048
R1(config)# username Admin secret Str0ng3rPa55w0rd
R1(config)# ip ssh version 2
R1(config)# line vty 0 4
R1(config-line)# transport input ssh
R1(config-line)# login local
```

Local database method has limitations:

- User accounts must be configured locally on each device. In large enterprise environment with multiple routers and switches to manage, it can take time to implement and change local databases on each device
- The local database configuration provides no fallback authentication method
    - For example: What if administrator forgets the username and password for device? With no backup method available for authentication, password recovery becomes only option

Better solution is to have all devices refer to the same database of usernames and passwords from a central server

AAA Components

- Stands for Authentication, Authorization, and Accounting.
- Similar to using a credit card
    - Credit card identifies who can use it, how much user can spend, and keeps an account of what items and services user purchased
- Provides primary framework to set up access control on a network device
- Way to control who is permitted to access a network (authenticate), what they can do while they are there (authorize), and to audit what actions they performed while accessing the network (accounting)

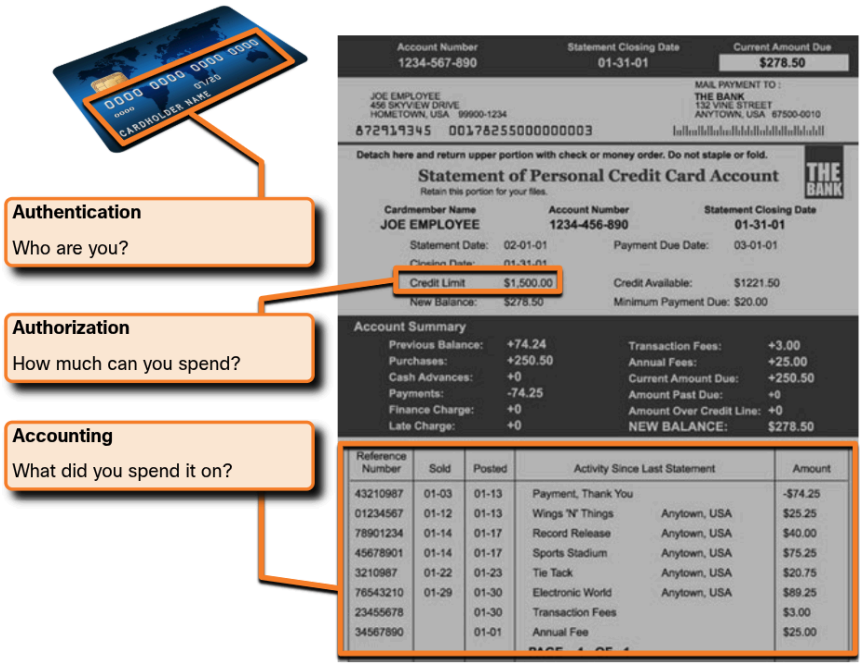

Authentication

Local and server-based are two common methods

Local AAA Authentication

- Stores usernames and passwords locally in network device such as Cisco router. User authenticate against the local database. Ideal for small networks.

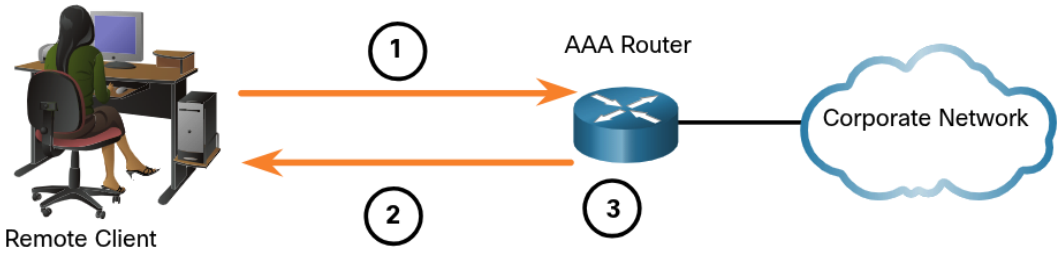

1. Client establishes a connection with the router
2. The AAA router prompts user for a username and password
3. The router authenticates the username and password using local database and user is provided access to the network based on information in local database

Server-Based AAA Authentication

- Router accesses a central AAA server
- AAA server contains the usernames and passwords for all users
- Router uses either the Remote Authentication Dial-In User Service (RADIUS) or Terminal Access Controller Access Control System (TACACS+) protocols to communicate with the AAA server
    - When there are multiple routers and switches, server-based AAA is more appropriate

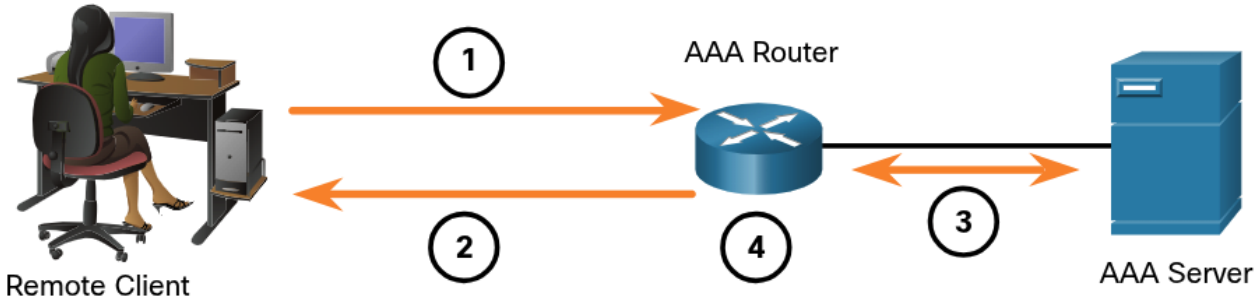

1. The client establishes a connection with the router
2. The AAA router prompts the user for a username and password
3. The router authenticates the username and password using a AAA server
4. The user is provided access to the network based on information in the remote AAA server

Authorization

- Automatic and does not require users to perform additional steps after authentication
- Authorization governs what users can and cannot do on the network after they are authenticated
- Uses a set of attributes that describes the user's access to the network
    - Used by the AAA server to determine privileges and restrictions for that user

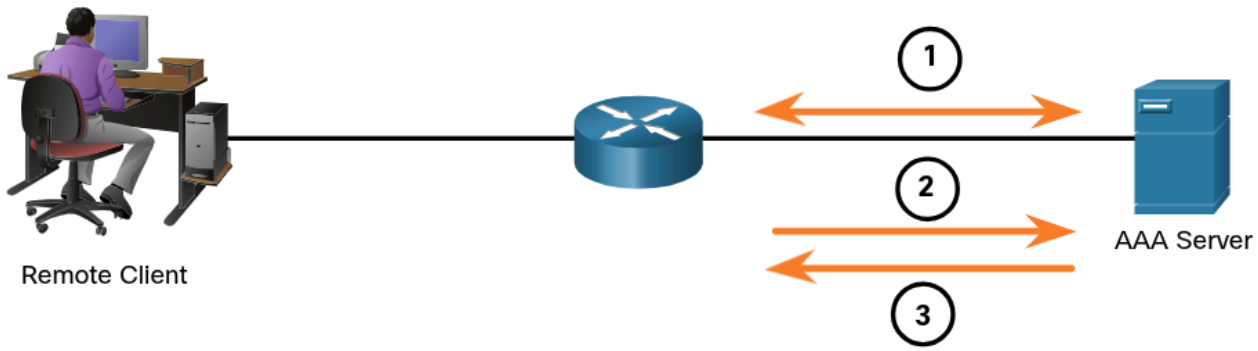

1. When a user has been authenticated, a session is established between the router and the AAA server
2. The router requests authorization from the AAA server for the client's requested service
3. The AAA server returns a PASS/FAIL response for authorization

Accounting

- Collects and reports usage data
  - This data can be used for such purposes as auditing or billing
  - Might include the start and stop connection times, executed commands, number of packets, and number of bytes
- Primary use is to combine it with AAA authentication

The AAA server keeps a detailed log of exactly what the authenticated user does on the device

- This includes all EXEC and configuration commands issued by user
- Useful when troubleshooting devices
- Also provides evidence for when individuals perform malicious acts

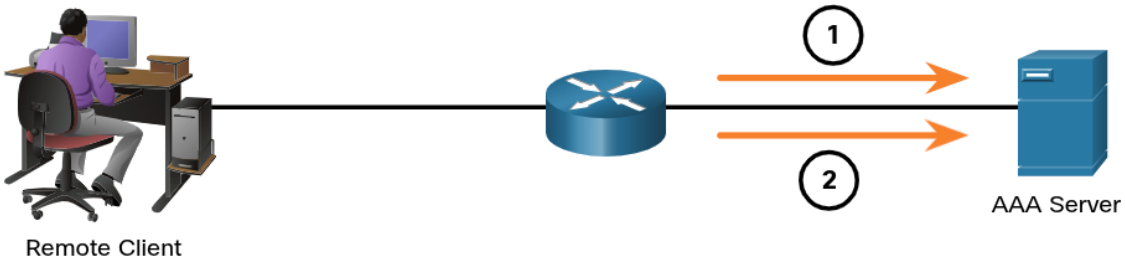

1. When a user has been authenticated, the AAA accounting process generates a start message to begin the accounting process
2. When the user finishes, a stop message is recorded and the accounting process ends

802.1X

- IEEE 802.1X standard is a port-based access control and authentication protocol
  - Restricts unauthorized workstations from connecting to a LAN through publicly accessible switch ports

The authentication server authenticates each workstation that is connected to a switch port before making available any services offered by the switch or the LAN

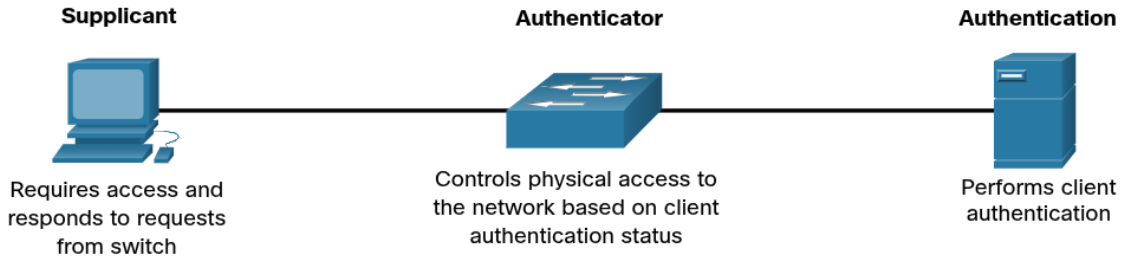

- **Client (Supplicant)** - A device running 802.1X-compliant client software, which is available for wired or wireless devices
- **Switch (Authenticator)** - The switch acts as an intermediary between the client and the authentication server. It requests identifying information from the client, verifies that information with the authentication server, and relays a response to the client. Another device that could act as authenticator is a wireless access point
- **Authentication server** - The server validates the identity of the client and notifies the switch or wireless access point that the client is or is not authorized to access the LAN and switch services

Layer 2 Vulnerabilities

Recall - OSI model is divided into seven layers working independently of each other.

Security is routinely implemented to protect elements in Layer 3 up through Layer 7

- Using VPNs, firewalls, and IPS devices

If Layer 2 is compromised, then all Layers above are also affected

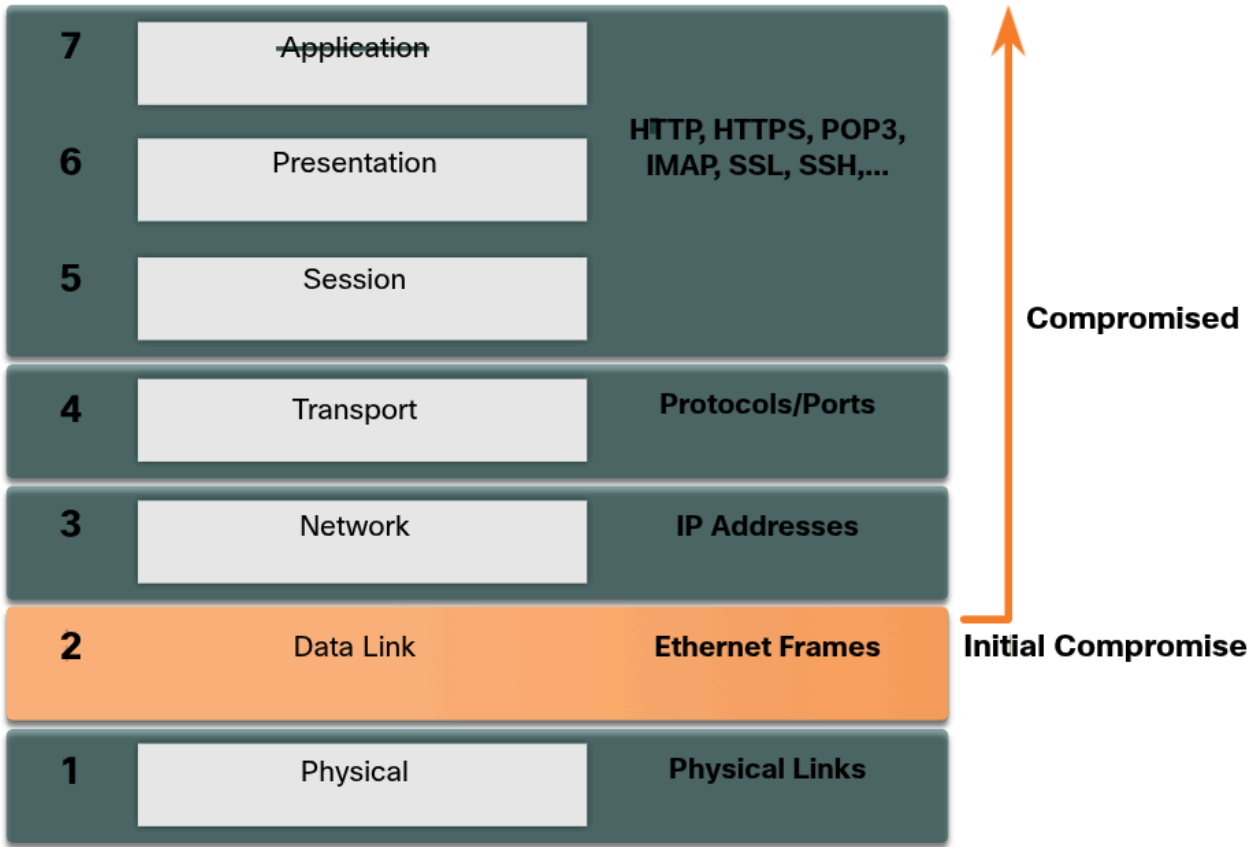

Switch Attack Categories

Layer 2 is weak link.

- This is because LANs were traditionally under the administrative control of a single organization
- Inherently trust all persons and devices connected to our LAN
    - BYOD and more sophisticated attacks make LANs more vulnerable to penetration

First step in mitigating attacks on Layer 2 infrastructure is to understand the underlying operation of Layer 2 and the threats posted by Layer 2 infrastructure.

Layer 2 Attacks

Category	Examples
MAC Table Attacks	Include MAC address flooding attacks
VLAN Attacks	Includes VLAN hopping and VLAN double-tagging attacks. It also includes attacks between devices on a common VLAN
DHCP Attacks	Includes DHCP starvation and DHCP spoofing attacks
ARP Attacks	Includes ARP spoofing and ARP poisoning attacks
Address Spoofing Attacks	Includes MAC address and IP address spoofing attacks
STP Attacks	Includes Spanning Tree Protocol manipulation attacks

Layer 2 Attack Mitigation

Solution	Description
Port Security	Prevents many types of attacks including MAC address flooding attacks and DHCP starvation attacks
DHCP Snooping	Prevents DHCP starvation and DHCP spoofing attacks
Dynamic ARP Inspection (DAI)	Prevents ARP spoofing and ARP poisoning attacks
IP Source Guard (IPSG)	Prevents MAC and IP address spoofing attacks

These Layer 2 solutions will not be effective if the management protocols are not secured

- Example: Management protocols Syslog, Simple Network Management Protocol (SNMP), Trivial File Transfer Protocol (TFTP), telnet, File Transfer Protocol (FTP), etc

- Always use secure variants of these protocols such as: SSH, Secure Copy Protocol (SCP), Secure FTP (SFTP), and Secure Socket Layer/Transport Layer Security (SSL/TLS)
- Consider using out-of-band management network to manage devices
- Use a dedicated management VLAN where nothing but management traffic resides
- Use ACLs to filter unwanted access

MAC Address Table Attack

Switch Operation Review

Recall - To make forwarding decisions, a Layer 2 switch builds on a table based on the source MAC addresses in received frames

- MAC address tables are stored in memory and are used to more efficiently forward frames

```
S1# show mac address-table dynamic
          Mac Address Table
-------------------------------------------
Vlan    Mac Address       Type        Ports
----    -----------       --------    -----
   1    0001.9717.22e0     DYNAMIC    Fa0/4
   1    000a.f38e.74b3     DYNAMIC    Fa0/1
   1    0090.0c23.ceca     DYNAMIC    Fa0/3
   1    00d0.ba07.8499     DYNAMIC    Fa0/2
S1#
```

MAC Address Table Flooding

- All MAC tables have a fixed size
    - A switch can run out of resources in which to store MAC addresses
- Flooding take advantage of this limitation by bombarding the switch with fake source MAC addresses until switch MAC address table is full

When flooding occurs, the switch treats the frame as an unknown unicast and begins to flood all incoming traffic out all ports on the same VLAN without referencing the MAC table

- This now allows threat actor to capture all of the frames sent from one host to another on the local LAN or local VLAN

Note - Traffic is flooded only within the local LAN or VLAN. The threat actor can only capture traffic within the local LAN or VLAN to which the threat actor is connected

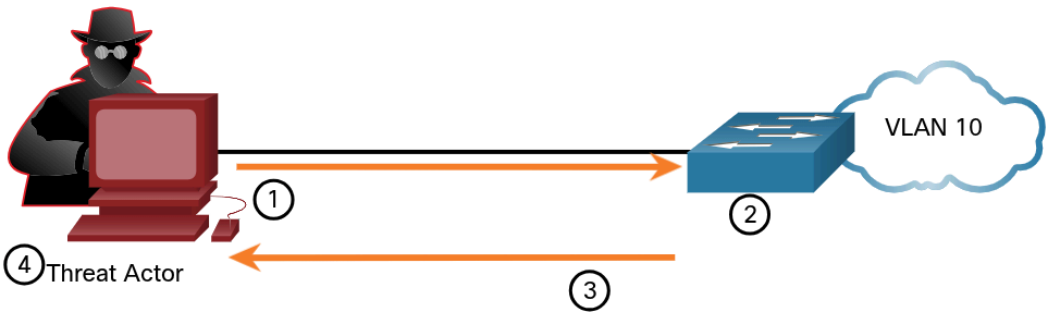

*1. The threat actor is connected to VLAN 10 and uses **macof** to rapidly generate many random source and destination MAC and IP addresses*
2. Over a short period of time, the switch's MAC table fills up
*3. When the MAC table is full, the switch begins to flood all frames that it receives. As long as **macof** continues to run, the MAC table remains full and the switch continues to flood all incoming frames out every port associated with VLAN 10*
4. The threat actor then uses packet sniffing software to capture frames from any and all devices connected to VLAN 10

*If the threat actor stops **macof** from running or is discovered and stopped, the switch eventually ages out the older MAC address entries from the table and begins to act like a switch again.*

MAC Address Table Attack Mitigation

What makes tools such as **macof** so dangerous is that an attacker can create a MAC table overflow attack very quickly

- Example: A catalyst 6500 switch can store 132,000 MAC addresses in its MAC address table
    - **macof** can flood a switch up to 8,000 bogus frames per second; creating a MAC address table overflow attack in a matter of seconds

```
# macof -i eth1
36:a1:48:63:81:70 15:26:8d:4d:28:f8 0.0.0.0.26413 > 0.0.0.0.49492: S 1094191437:1094191437(0) win 512
16:e8:8:0:4d:9c da:4d:bc:7c:ef:be 0.0.0.0.61376 > 0.0.0.0.47523: S 446486755:446486755(0) win 512
18:2a:de:56:38:71 33:af:9b:5:a6:97 0.0.0.0.20086 > 0.0.0.0.6728: S 105051945:105051945(0) win 512
e7:5c:97:42:ec:1 83:73:1a:32:20:93 0.0.0.0.45282 > 0.0.0.0.24898: S 1838062028:1838062028(0) win 512
62:69:d3:1c:79:ef 80:13:35:4:cb:d0 0.0.0.0.11587 > 0.0.0.0.7723: S 1792413296:1792413296(0) win 512
c5:a:b7:3e:3c:7a 3a:ee:c0:23:4a:fe 0.0.0.0.19784 > 0.0.0.0.57433: S 1018924173:1018924173(0) win 512
88:43:ee:51:c7:68 b4:8d:ec:3e:14:bb 0.0.0.0.283 > 0.0.0.0.11466: S 727776406:727776406(0) win 512
b8:7a:7a:2d:2c:ae c2:fa:2d:7d:e7:bf 0.0.0.0.32650 > 0.0.0.0.11324: S 605528173:605528173(0) win 512
e0:d8:1e:74:1:e 57:98:b6:5a:fa:de 0.0.0.0.36346 > 0.0.0.0.55700: S 2128143986:2128143986(0) win 512
```

- They not only affect the local switch, they can also affect other connected Layer 2 switches. When the MAC address table of a switch is full, it starts flooding out all ports including those connected to other Layer 2 switches
- To mitigate, network administrators must implement port security

Port security will only allow a specified number of source MAC addresses to be learned on the port.

LAN Attacks

VLAN and DHCP Attacks

VLAN Hopping Attacks

- Enables traffic from one VLAN to be seen by another VLAN **without** the aid of a router
- A threat actor configures the host to **spoof 802.1Q and DTP signaling** to trunk with the connecting switch
    - Can send and receive traffic on any VLAN, effectively **hopping between VLANs**

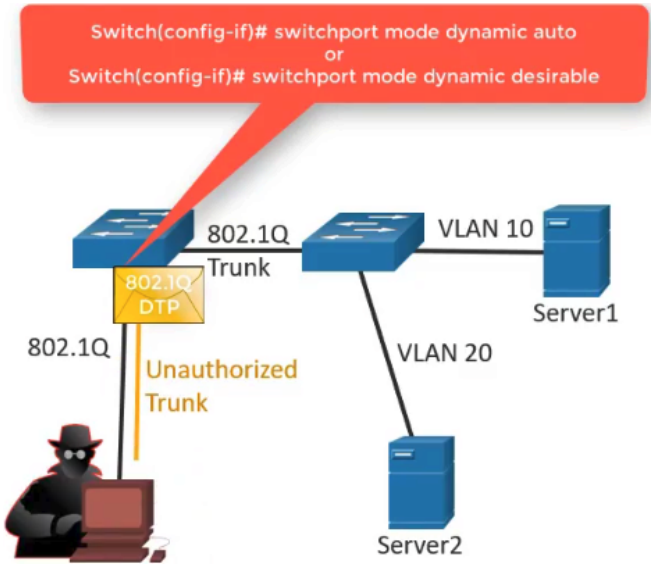

VLAN Double-Tagging Attack

- A threat actor may **embed a hidden 802.1Q tag** inside the frame that already has an 802.1Q tag
    - This tag allows the frame to go to a VLAN that the **original 802.1Q tag did not specify**
    - Only works when the attack is connected to a port residing in the **same VLAN as the native VLAN** of the trunk port

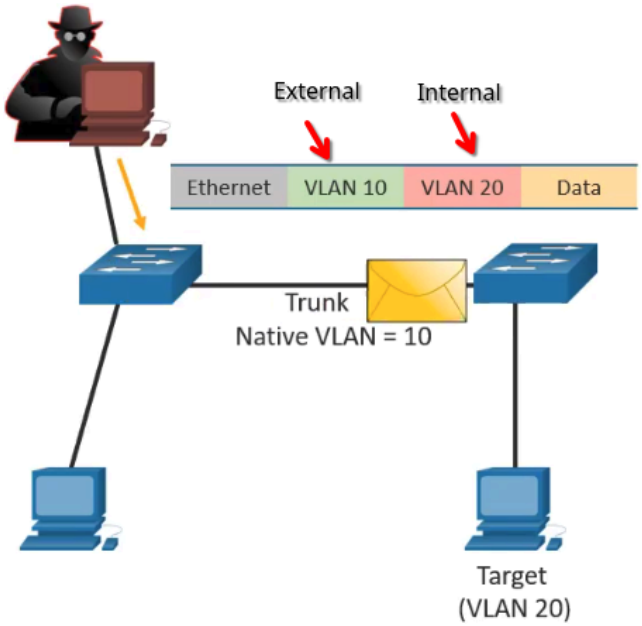

This scenario only works when the threat actor is on the same VLAN as the native VLAN between the switches

Can be prevented with a couple of security guidelines

1. Get rid of trunking on any access ports that go to end devices
2. “Auto trunking”-- Dynamic desirable/ Dynamic auto should be disabled
    a. Use of manual static trunking as needed
3. Native VLANs should only be used on trunk links
    a. Should never span to user device

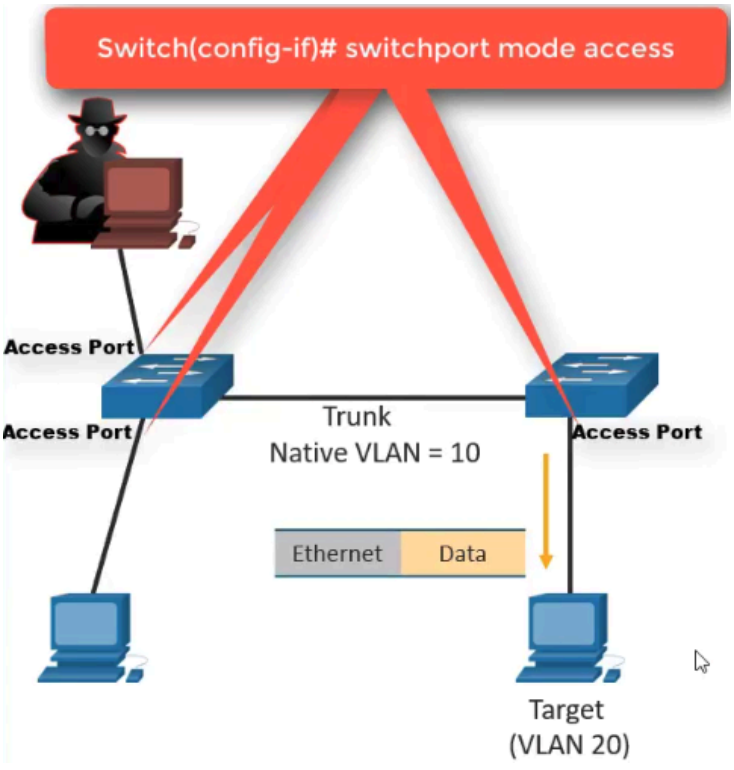

Step 1 -

- The threat actor sends a double-tagged 802.1Q frame to the switch
- The outer header has the VLAN tag of the threat actor, which is the same as the native VLAN of the trunk port
- The inner tag is the victim VLAN

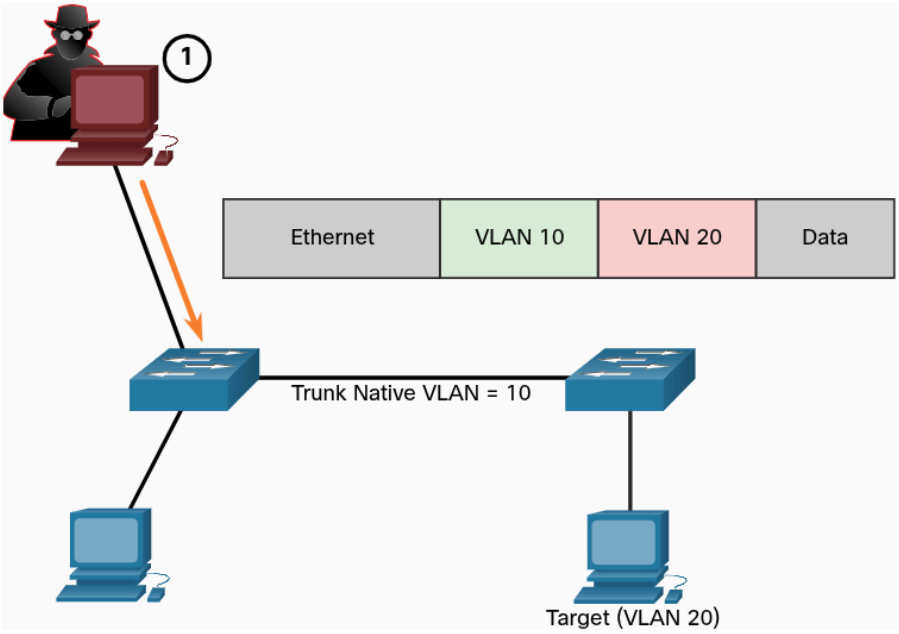

Step 2 -

- The frame arrives on the first switch, which looks at the first 4-byte 802.1Q tag
- The switch sees that the frame is destined for VLAN 10, which is the native VLAN
- The switch forwards the packet out all VLAN 10 ports after stripping the VLAN 10 tag
- The frame is not retagged because it is part of the native VLAN
    - At this point, the VLAN 20 tag is still intact and has not been inspected by first the switch

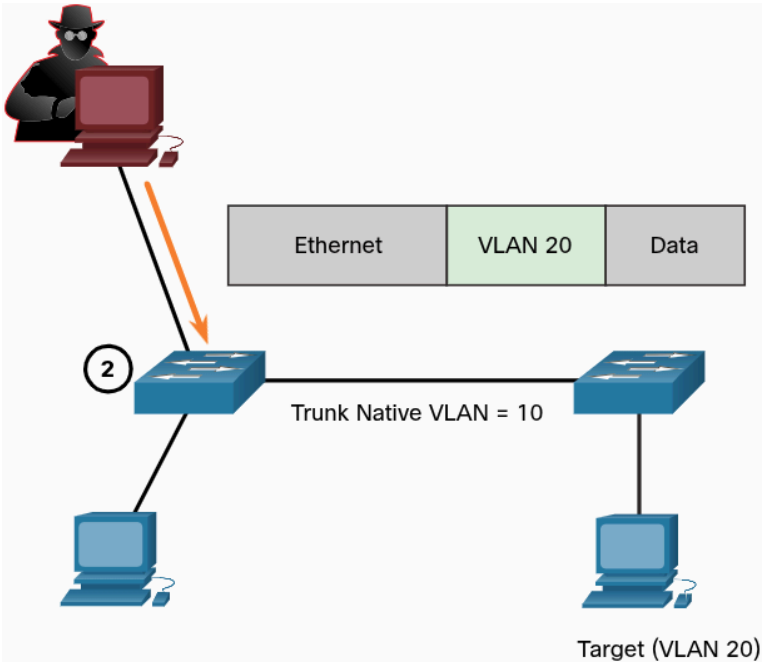

Step 3 -

- The frame arrives at the second switch which has no knowledge that it was supposed to be for VLAN 10.
    - Native VLAN traffic is not tagged by the sending switch as specified in the 802.1Q specification
- The second switch looks only at the inner 802.1Q tag that the threat actor inserted and sees that the frame is destined for VLAN 20, the target VLAN
- The second switch sends the frame on to the target or floods it, depending on whether there is an existing MAC address table entry for the target

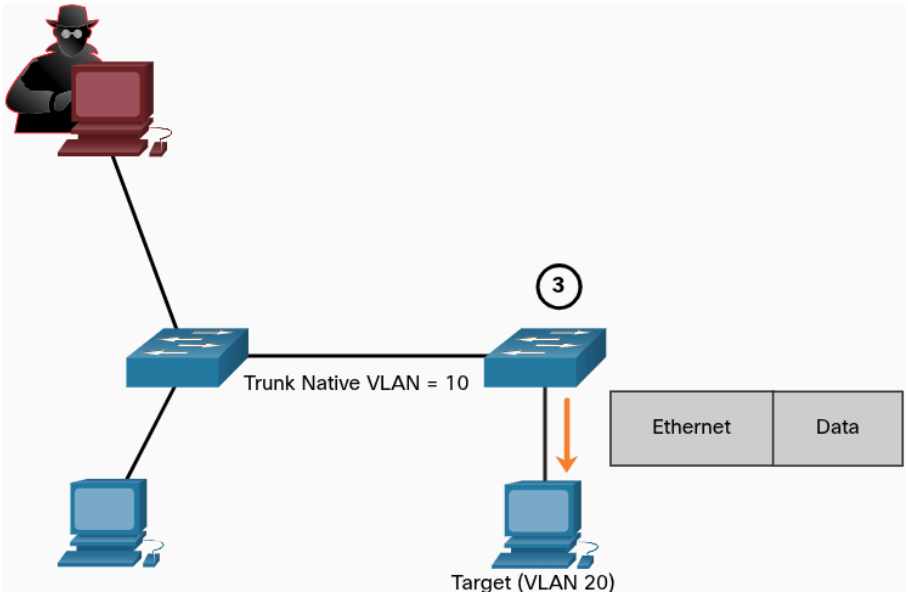

A VLAN double-tagging attack is unidirectional and works only when the attacker is connected to a port residing in the same VLAN as the native VLAN of the trunk port

- The idea is that double tagging allows the attacker to send data to hosts or servers on a VLAN that otherwise would be blocked by some type of access control configuration
- Presumably the return traffic will also be permitted, thus giving the attacker the ability to communicate with devices on the normally blocked VLAN

VLAN Attack Mitigation

VLAN hopping and VLAN double-tagging attacks can be prevented by implementing the following trunk security guidelines:

- Disable trunking on all access ports
- Disable auto trunking on trunk links so that trunks must be manually enabled
- Be sure that the native VLAN is only used for trunk links

DHCP Messages

DHCP servers dynamically provide IP configuration information including:

- IP address
- Subnet mask
- Default gateway
- DNS Servers
- Etc

DHCP uses **DORA** for as the message exchange between client and server

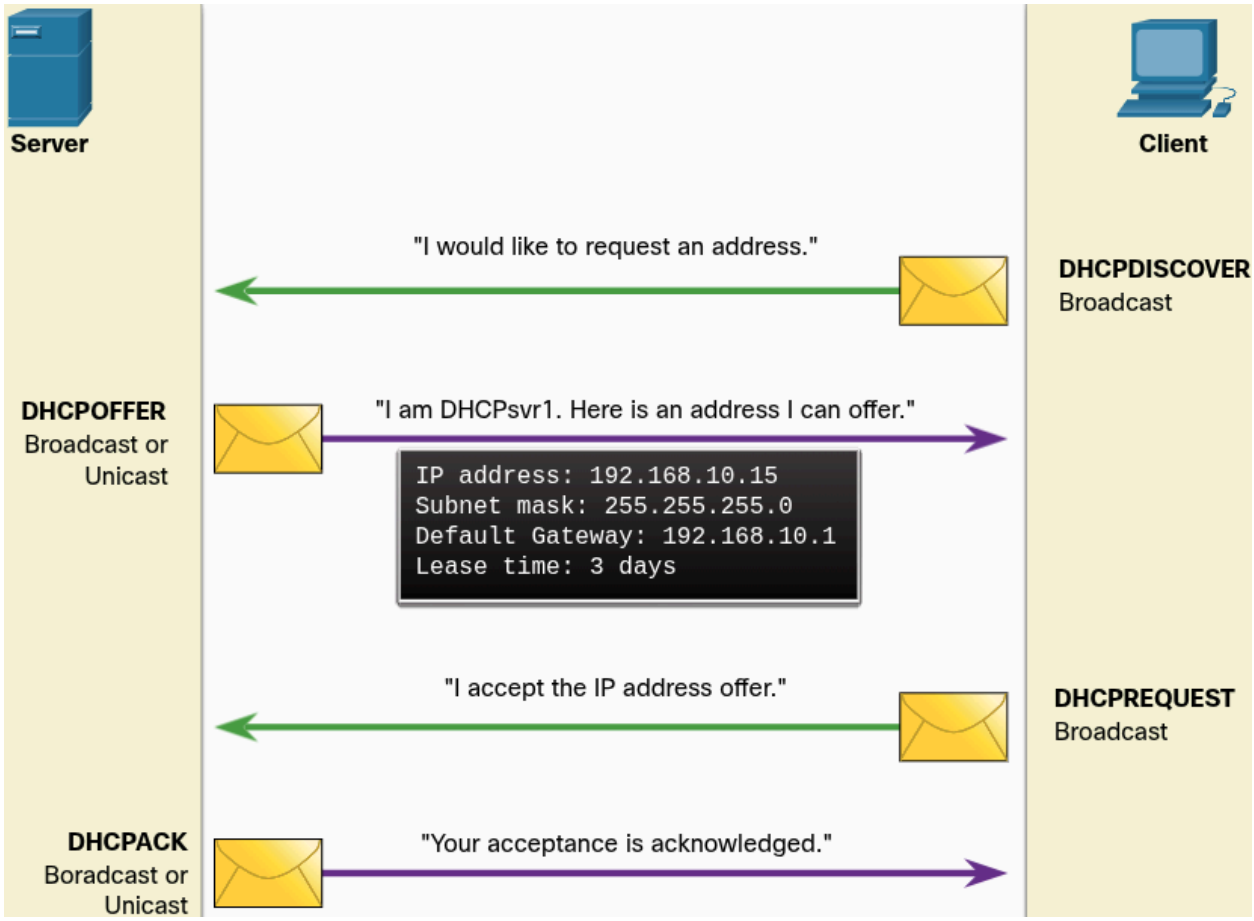

DHCP Attacks

DHCP Starvation Attack

- The goal is to create a DoS for connecting clients
- Tool example: Gobbler
    - Has ability to look at the entire scope of leasable IP addresses and tries to lease them all, specifically, it creates DHCP discover messages with bogus MAC addresses
- A threat actor creates **DHCP discover messages with bogus MAC addresses**

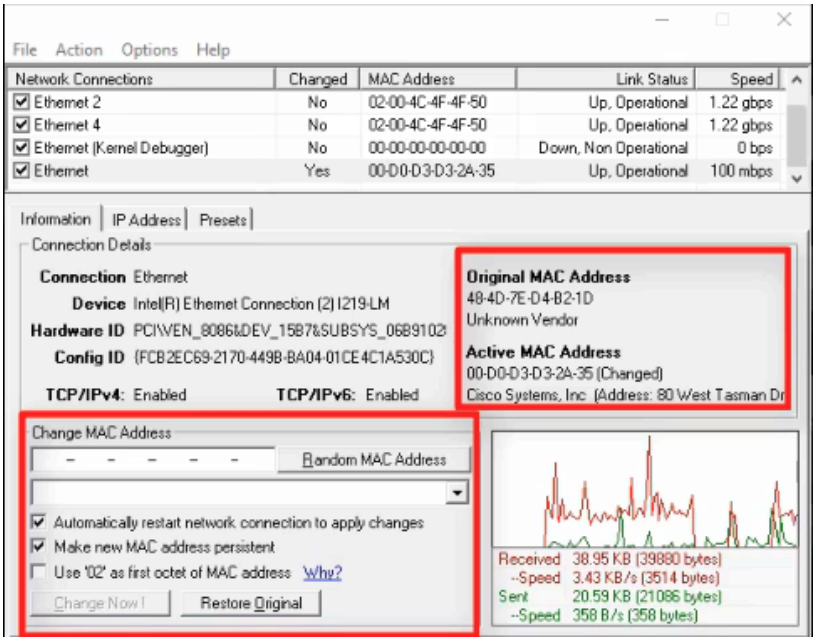

DHCP Spoofing Attack

- A rogue DHCP server is connected to the network and **provides false IP configuration parameters** to legitimate clients:
    - **Wrong default gateway** - The rogue server provides an invalid gateway or the IP address of its host to create a man-in-the-middle attack
        - This may go entirely undetected as the intruder intercepts the data flow through the network
    - **Wrong DNS server** - The rogue server provides an incorrect DNS server address pointing the user to a nefarious website
    - **Wrong IP address** - The rogue server provides an invalid IP address effectively creating a DoS attack on the DHCP client

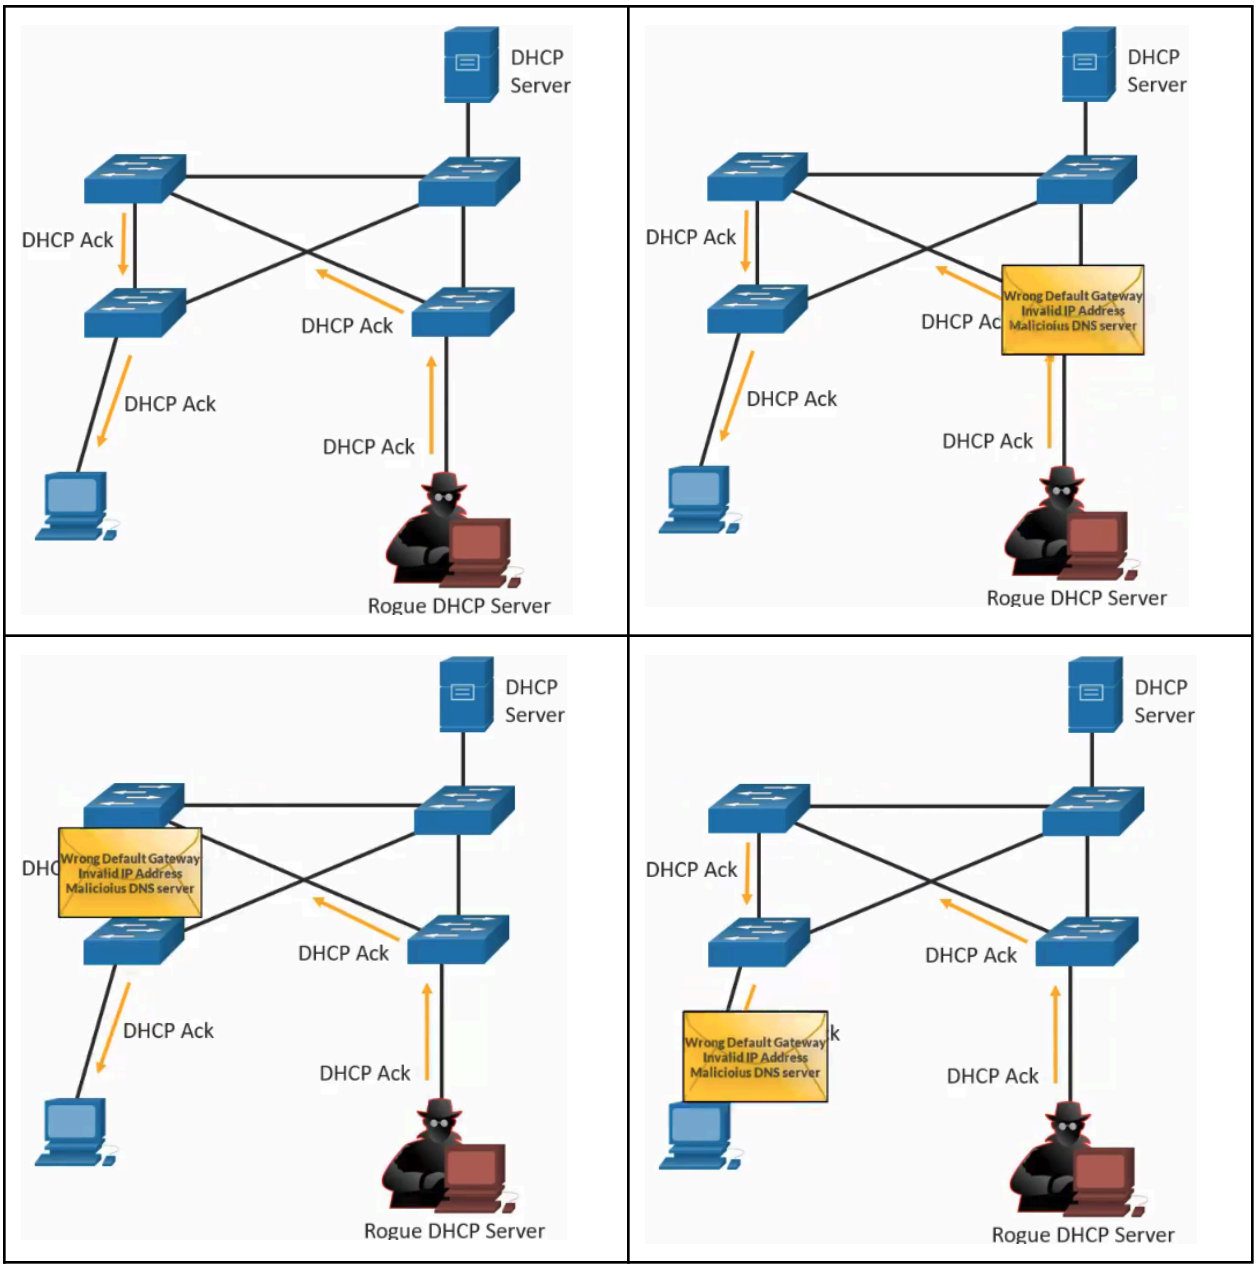
DHCP Server
DHCP Ack
DHCP Ack
DHCP Ack
DHCP Ack
Rogue DHCP Server
DHCP Server
DHCP Ack
Wrong Default Gateway
Invalid IP Address
Malicious DNS server
DHCP Ack
DHCP Ack
Rogue DHCP Server
DHCP Server
Wrong Default Gateway
Invalid IP Address
Malicious DNS server
DHCP Ack
DHCP Ack
DHCP Ack
Rogue DHCP Server
DHCP Server
DHCP Ack
Wrong Default Gateway
Invalid IP Address
Malicious DNS server
DHCP Ack
DHCP Ack
Rogue DHCP Server

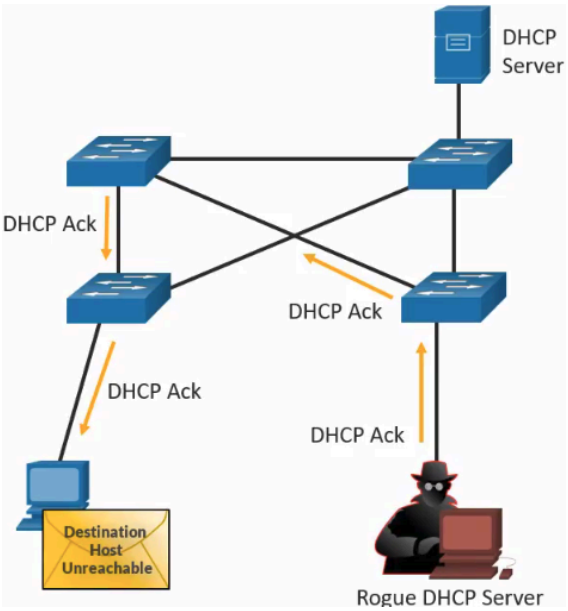
DHCP Server
DHCP Ack
DHCP Ack
DHCP Ack
DHCP Ack
Destination Host Unreachable
Rogue DHCP Server

Step 1 -

Threat Actor Connects Rogue DHCP Server

A threat actor successfully connects a rogue DHCP server to a switch port on the same subnet and VLANs as the target clients. The goal of the rogue server is to provide clients with false IP configuration information

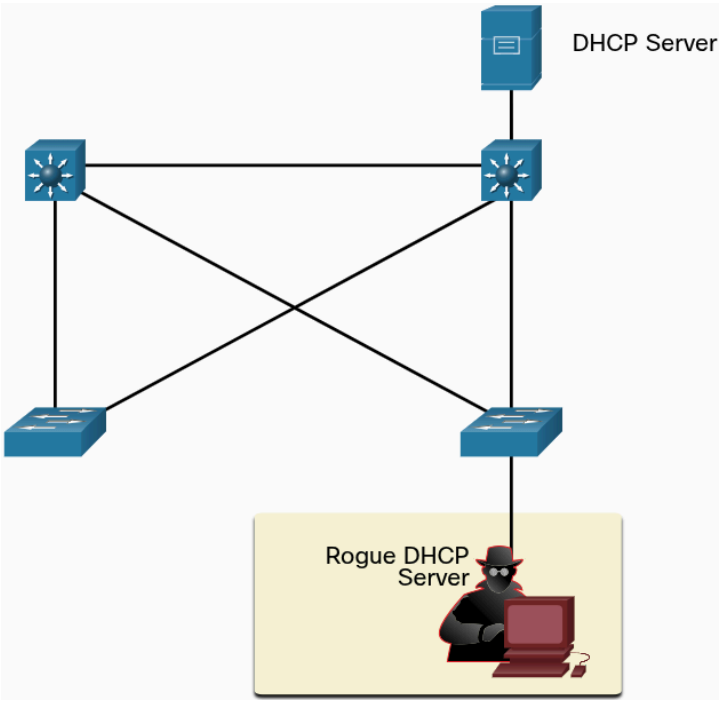

Step 2 -

Client Broadcasts DHCP Discovery Messages

A legitimate client connects to the network and requires IP configuration parameters. Therefore, the client broadcasts a DHCP Discovery request looking for a response from a DHCP server. Both servers will receive the message and respond

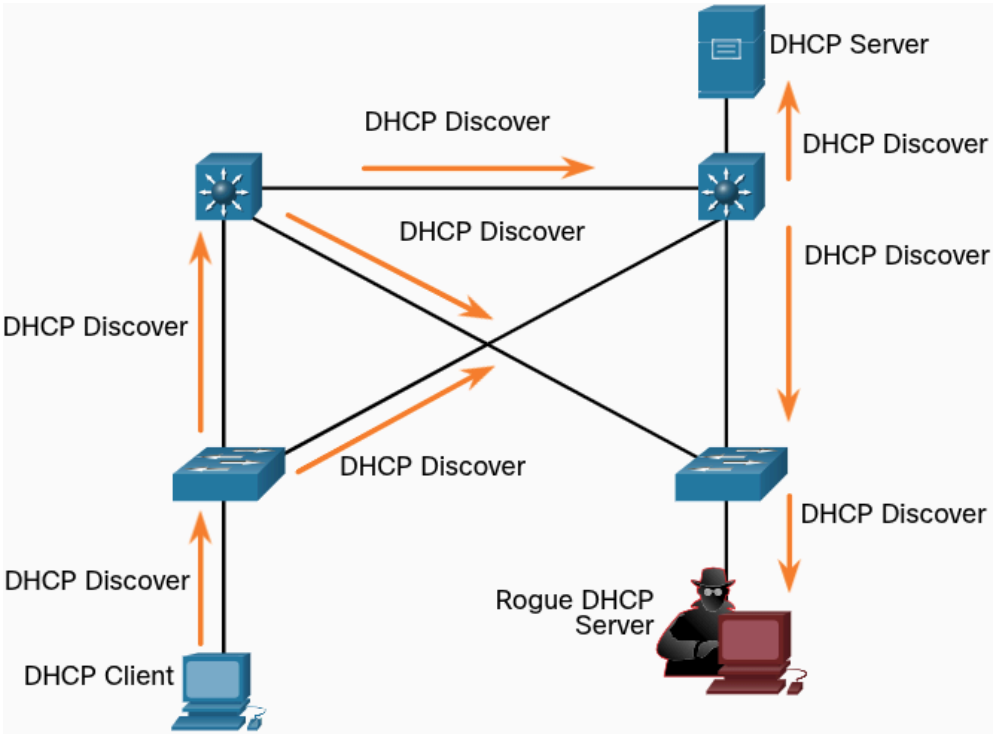

Step 3 -

Legitimate and Rogue DHCP Reply

The legitimate DHCP server responds with valid IP configuration parameters. However the rogue server also responds with a DHCP offer containing IP configuration parameters defined by the threat actor. The client will reply to the first offer received

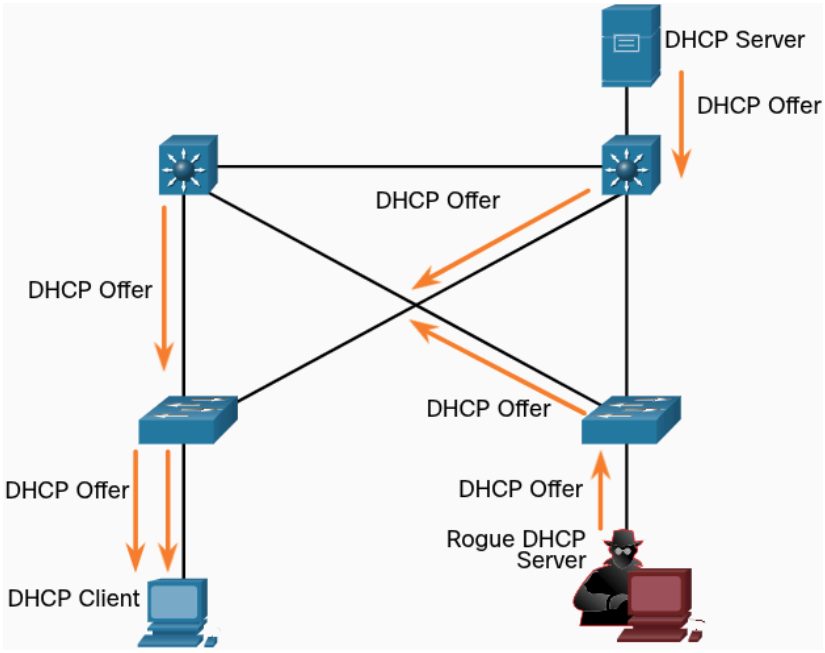

Step 4 -

Client Accepts Rogue DHCP Offer

The rogue offer was received first, and therefore, the client broadcasts a DHCP request accepting the IP parameters defined by the threat actor. The legitimate and rogue server will receive the request

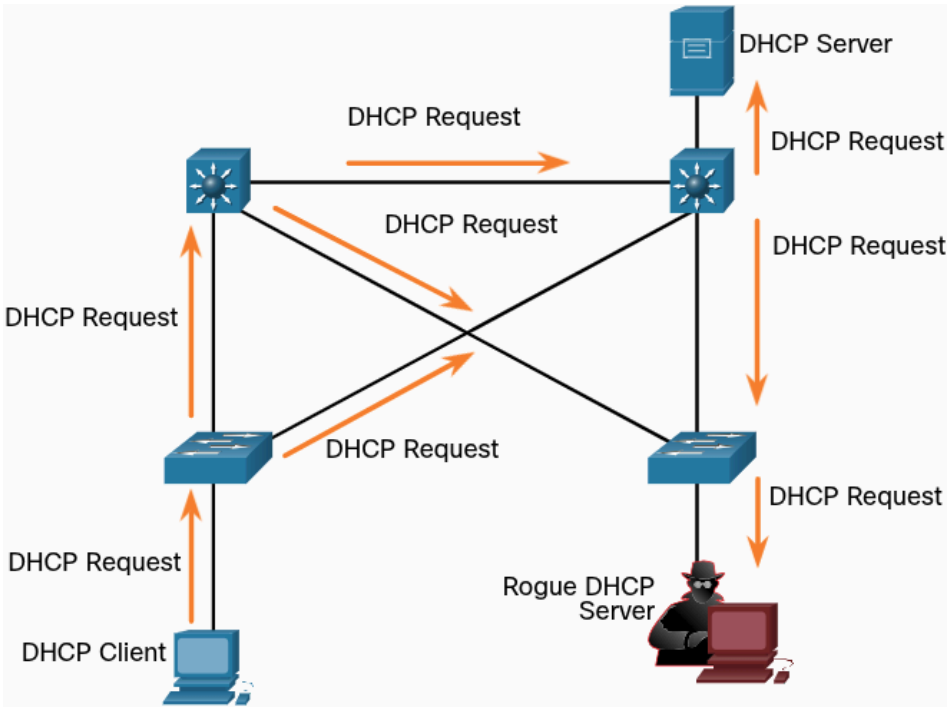

Step 5 -

Rogue Server Acknowledges

The rogue server unicasts a reply to the client to acknowledge its request. The legitimate server will cease communicating with the client

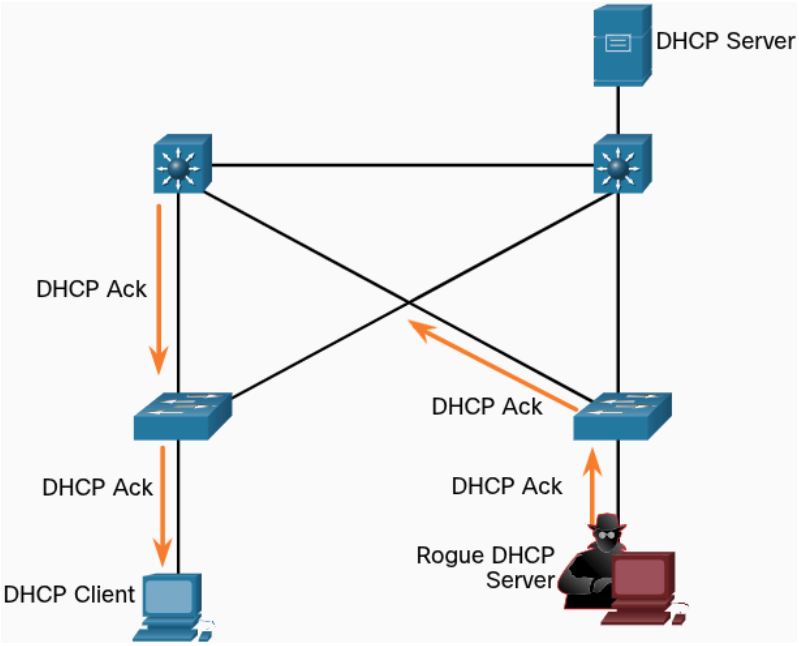

ARP Attacks

Recall - Hosts broadcast ARP Requests to determine the MAC address of a host with a particular IPv4 address

- Typically done to discover the mac address of the default gateway.
- All hosts on the subnet receive and process the ARP Request
- The host with the matching IPv4 address in the ARP Request sends an ARP reply

According to ARP RFC– A client is allowed to send an unsolicited ARP Request called a “gratuitous ARP”. When a host sends a gratuitous ARP, other hosts on the subnet store the MAC address and IPv4 address contained in the gratuitous ARP in their ARP tables

The problem is that an attacker can send a gratuitous ARP message containing a spoofed MAC address to a switch, and the switch would update its MAC table accordingly. Therefore, any host can claim to be the owner of any IP and MAC address combination they choose

- In a typical attack, a threat actor can send unsolicited ARP Replies to other hosts on the subnet with the MAC address of the threat actor and the IPv5 address of the default gateway

Tools to create man-in-the-middle attacks: dsniff, Cain & Abel, ettercap, Yersinia, etc.
IPv6 uses ICMPv6 Neighbor Discovery Protocol for Layer 2 address resolution

- Ipv6 includes strategies to mitigate Neighbor Advertisement spoofing, similar to way IPv6 prevents a spoofed ARP Reply

ARP spoofing and ARP poisoning are mitigated by implementing DAI

Step 1 -

Normal State with Converge MAC Tables

Each device has an accurate MAC table with the correct IPv4 and MAC addresses for the other devices on the LAN

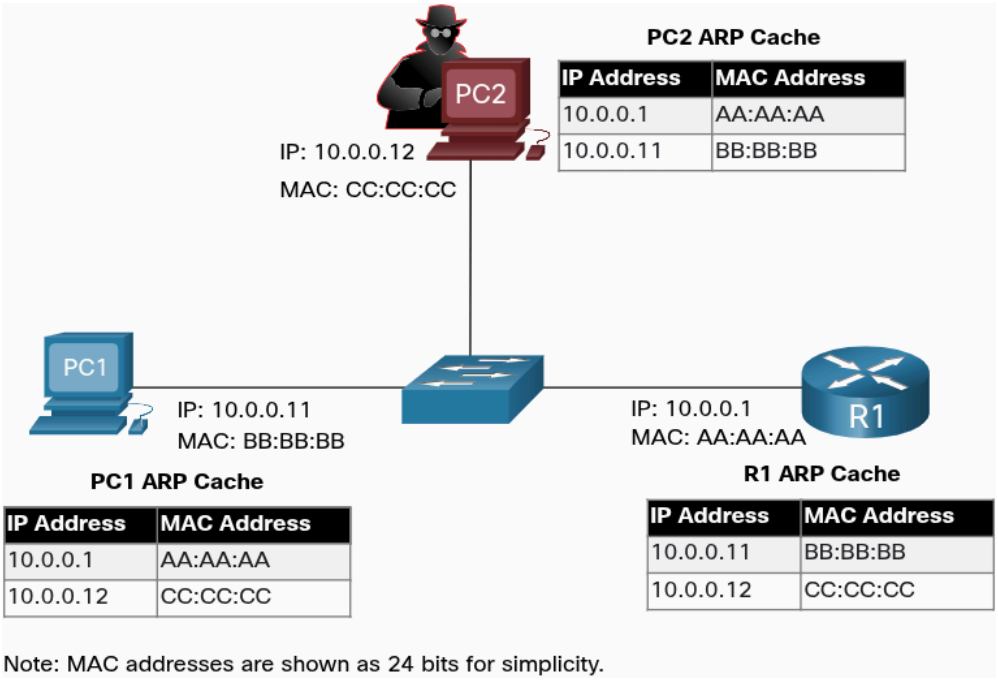

Step 2 -

ARP Spoofing Attack

Threat actor sends two spoofed gratuitous ARP Replies in an attempt to replace R1 as default gateway

1. The first one informs all devices on the LAN that the threat actor's MAC address (CC:CC:CC) maps to the R1's IPv4 address, 10.0.0.1
2. The second informs all devices on the LAN that the threat actor's MAC address (CC:CC:CC) maps to PC1's IPv4 address, 10.0.0.11

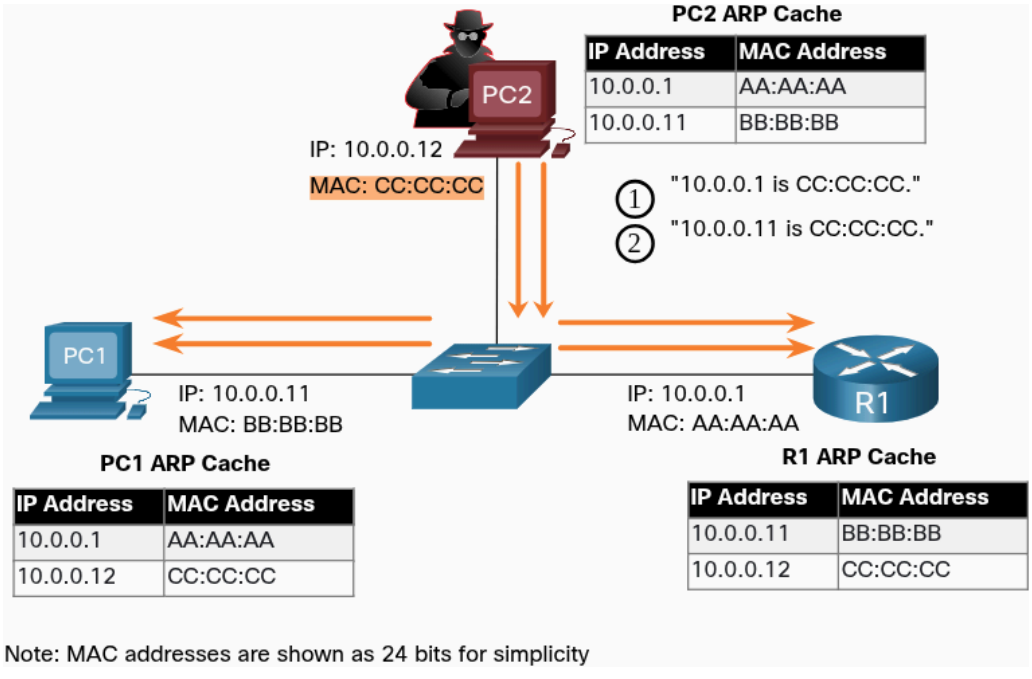

Step 3 -

ARP Poisoning Attack with Man-in-the-Middle Attack

R1 and PC1 remove the correct entry for each other's MAC address and replace it with PC2's MAC address

- The threat actor has now poisoned the ARP caches of all devices on the subnet.
- ARP poisoning leads to various man-in-the-middle attacks, posing a serious security threat to the network

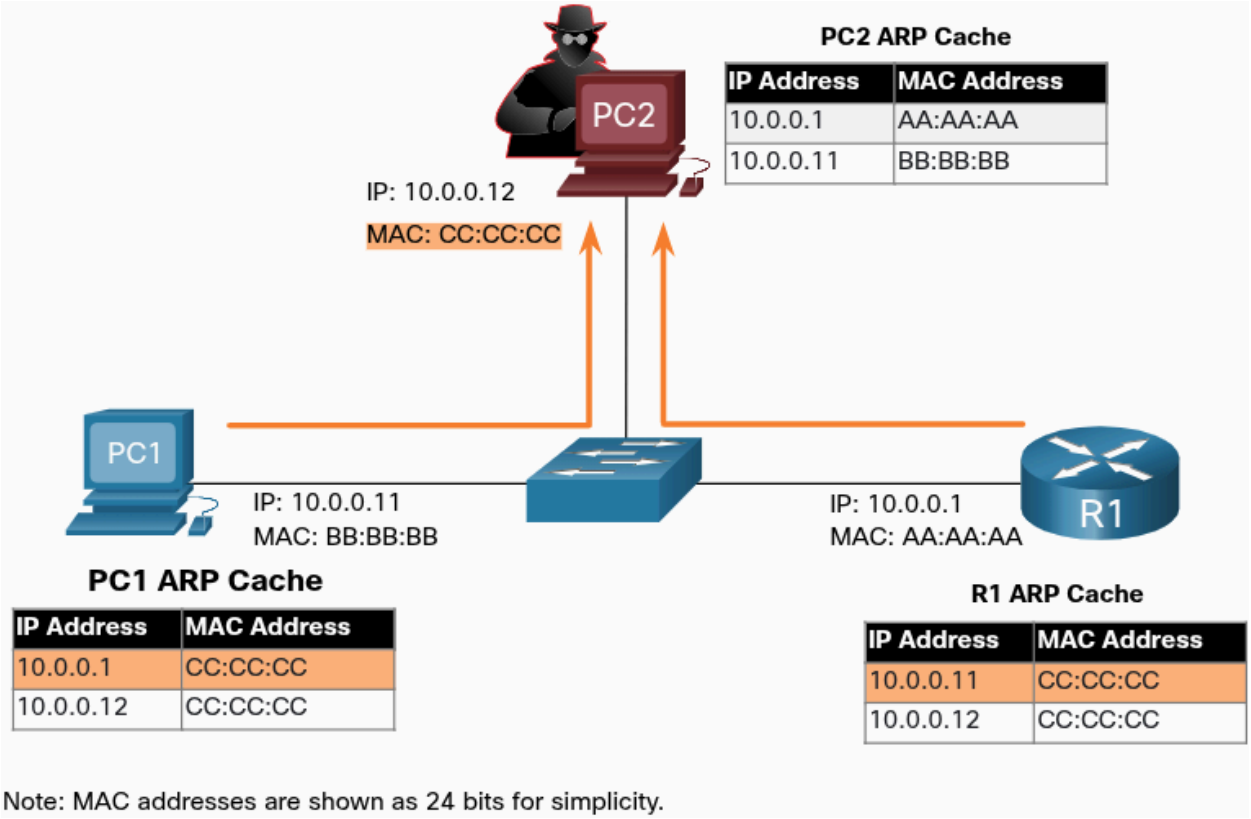

Note: MAC addresses are shown as 24 bits for simplicity.

Address Spoofing Attack

IP Address spoofing

- When the threat actor hijacks a valid IP address of another device on the subnet, or uses a random IP address.
    - Difficult to mitigate, especially when it is used inside a subnet in which the IP belongs

MAC Address spoofing

- Occurs when the threat actors alter the MAC address of their host to match another known MAC address of a target host
- The attacking host then sends a frame throughout the network with the newly-configured MAC address
- When the switch receives the frame, it examines the source MAC address
- The switch overwrites the current MAC table entry and assigns the MAC address to the new port
- It then forwards frames destined for the target host to the attacking host

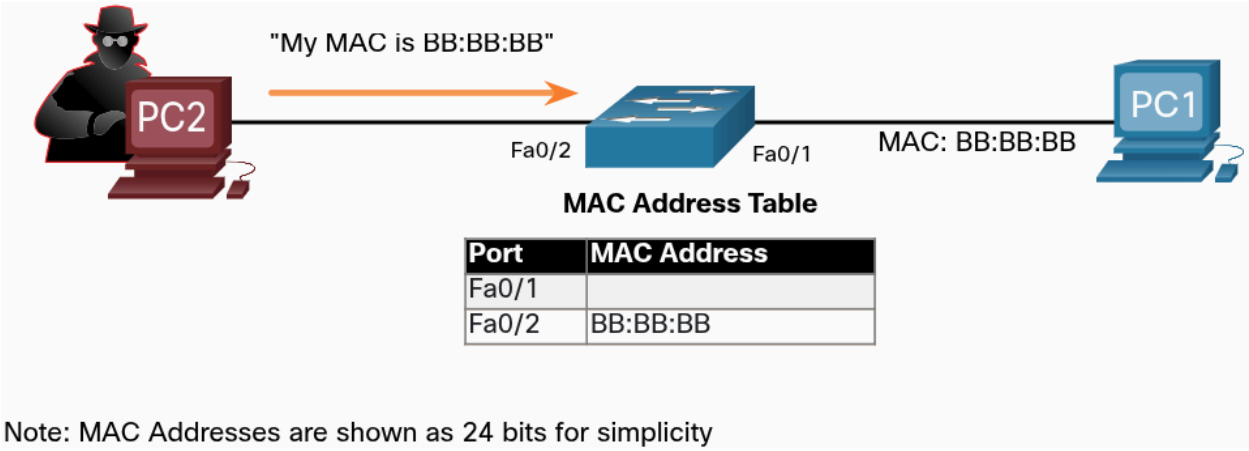

When the target host sends traffic, the switch will correct the error, realigning the MAC address to the original port

- To stop the switch from returning the port assignment to its correct state, the threat actor can create a program or script that will constantly send frames to the switch so that the switch maintains the incorrect or spoofed information.
- There is no security mechanism at Layer 2 that allows a switch to verify the source of MAC addresses, which is what makes it so vulnerable to spoofing

IP and MAC address spoofing can be mitigated by implementing IPSG

STP Attack

- Attackers can manipulate STP to conduct an attack by spoofing the root bridge and changing the topology of a network
- Attackers can make their host appear as root bridges, therefore, capture all traffic for the immediate switched domain

To conduct an STP Manipulation attack

- The attacking host broadcasts STP bridge protocol data units (BDPUs) containing configuration and topology changes that will force spanning-tree recalculations
    - The BPDUs sent by the attacking host announce a lower bridge priority in attempt to be elected as the root bridge

Note - These issues can occur when someone adds an Ethernet switch to the network without any malicious intent

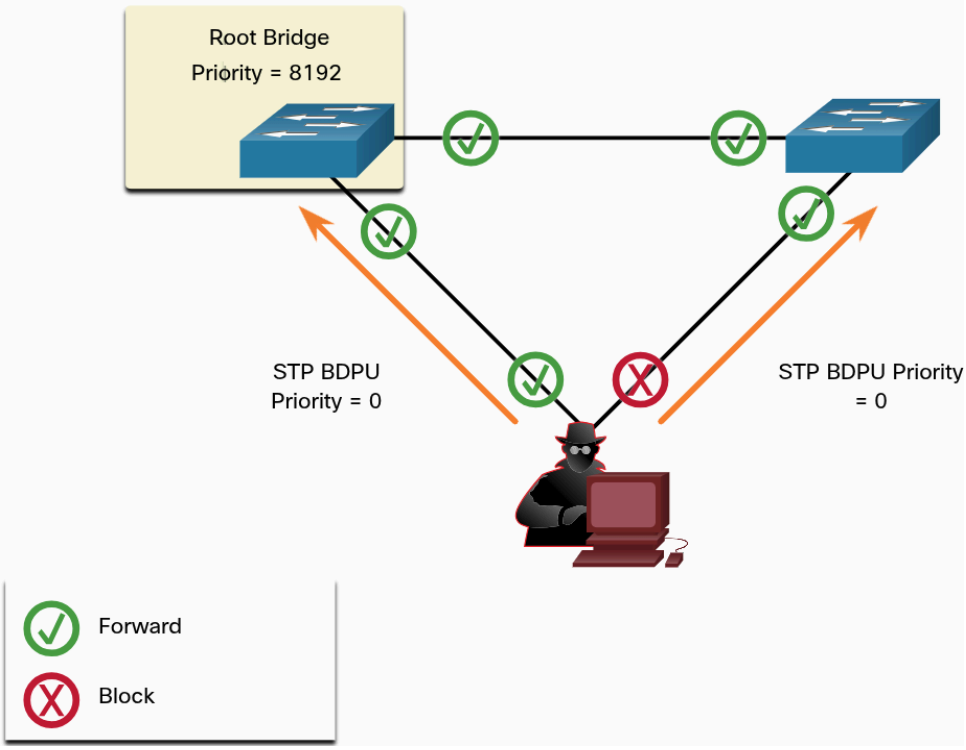

If successfully, the attacking host becomes the root bridge and can now capture a variety of frames that would otherwise not be accessible

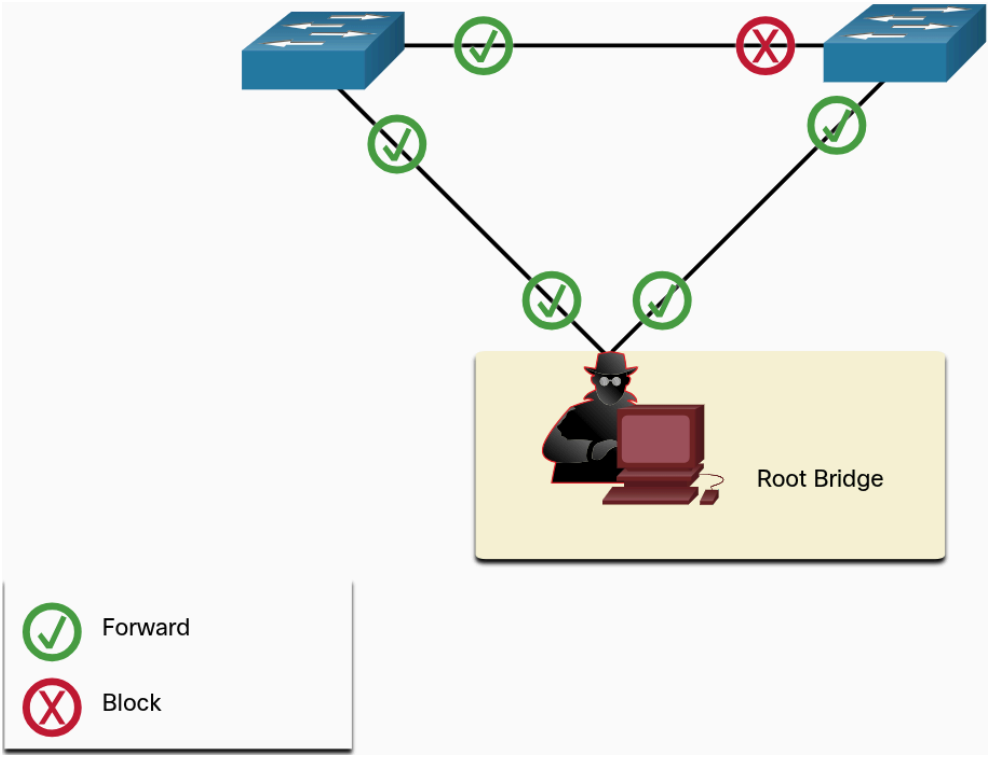

The STP attack is mitigated by implementing BPDU Guard on all access ports

CDP Reconnaissance

- A proprietary Layer 2 link discovery protocol
    - Enabled on all Cisco devices by default
- Can automatically discover other CDP-enabled devices and help auto-configure their connection
- Network administrators use CDP to help configure and troubleshoot network devices

CDP Information is sent out CDP-enabled ports in a periodic, unencrypted multicast.

- Includes the IP address of device, IOS software version, platform, capabilities, and native VLAN
- The device receiving the CDP message updates its CDP database
- Extremely useful in network troubleshooting
    - Example: CDP can be used to verify Layer 1 and 2 connectivity
    - If an administrator cannot ping a directly connected interface, but still receives CDP information, the problem is most likely related to the Layer 3 configuration

However, the information provided can also be used by threat actor to discover network infrastructure vulnerabilities

Figure displays sample Wireshark capture of contents of a CDP packet. Attacker is able to identify the Cisco IOS software version used by device

- This allows the attacker to determine whether there were any security vulnerabilities specific to that particular version of IOS

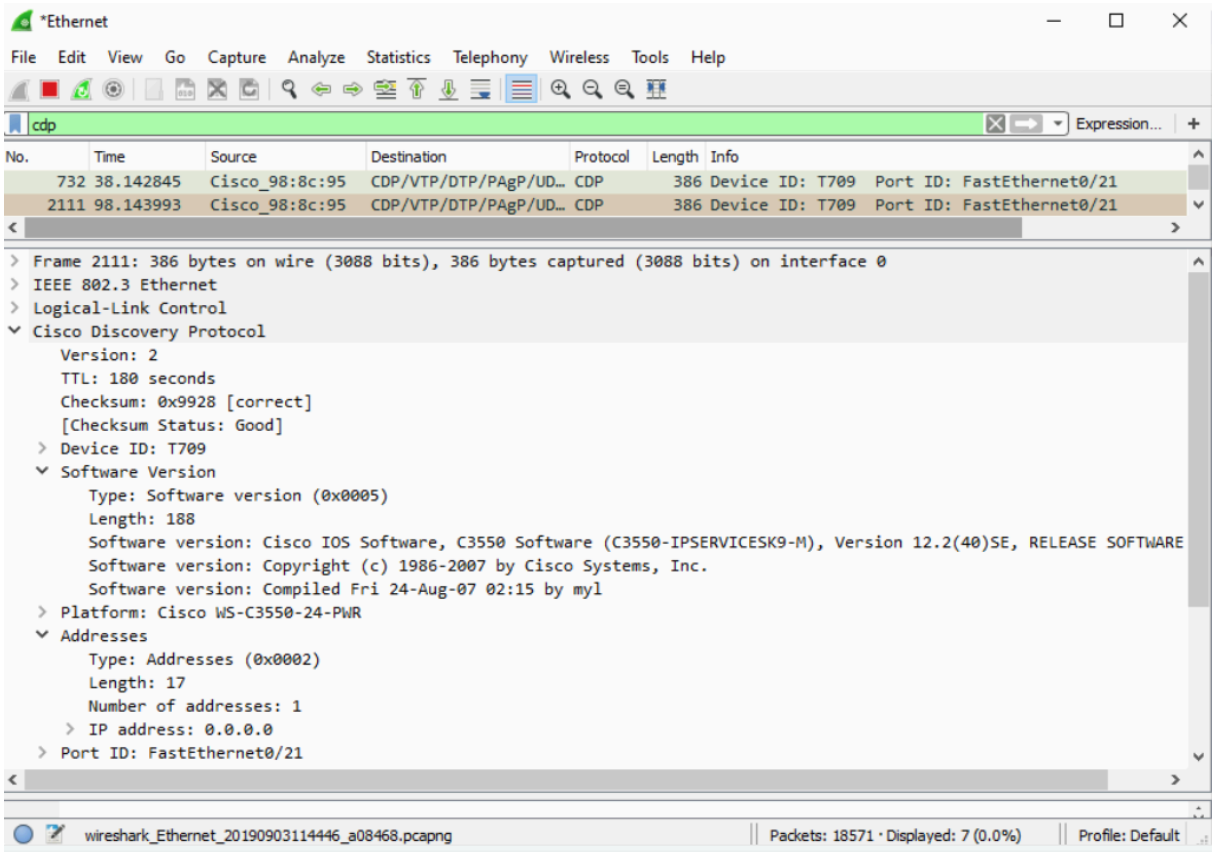

CDP Broadcasts are sent unencrypted and unauthenticated

- An attacker could interfere with the network infrastructure by sending crafted CDP frames containing bogus device information to directly-connected Cisco devices

To mitigate the exploitation of CDP, limit the use of CDP on devices or ports

- Example: Disable CDP on edge ports that connect to untrusted devices

To disable CDP globally on a device, use the **no cdp enable** interface configuration command
To enable CDP on a port, use the **cdp enable** interface configuration command

Note - Link Layer Discovery Protocol (LLDP) is also vulnerable to reconnaissance attacks.
Configure **no lldp run** to disable LLDP globally
To disable LLDP on the interface, configure **no lldp transmit** and **no lldp receive**

10.6.1 What did I learn in this module?

Endpoints are particularly susceptible to malware-related attacks that originate through email or web browsing, such as DDOS, date breaches, and malware. These endpoints have typically used traditional host-based security features, such as antivirus/antimalware, host-based firewalls, and host-based intrusion prevention systems (HIPSs). Endpoints are best protected by a combination of NAC, host-based AMP software, an email security appliance (ESA), and a web security appliance (WSA). Cisco WSA can perform blacklisting of URLs, URL-filtering, malware scanning, URL categorization, Web application filtering, and encryption and decryption of web traffic.

AAA controls who is permitted to access a network (authenticate), what they can do while they are there (authorize), and to audit what actions they performed while accessing the network (accounting). Authorization uses a set of attributes that describes the user's access to the network. Accounting is combined with AAA authentication. The AAA server keeps a detailed log of exactly what the authenticated user does on the device. The IEEE 802.1X standard is a port-based access control and authentication protocol that restricts unauthorized workstations from connecting to a LAN through publicly accessible switch ports.

If Layer 2 is compromised, then all layers above it are also affected. The first step in mitigating attacks on the Layer 2 infrastructure is to understand the underlying operation of Layer 2 and the Layer 2 solutions: Port Security, DHCP Snooping, DAI, and IPSG. These won't work unless management protocols are secured.

MAC address flooding attacks bombard the switch with fake source MAC addresses until the switch MAC address table is full. At this point, the switch treats the frame as an unknown unicast and begins to flood all incoming traffic out all ports on the same VLAN without referencing the MAC table. The threat actor can now capture all of the frames sent from one host to another on the local LAN or local VLAN. The threat actor uses macof to rapidly generate many random source and destination MAC and IP. To mitigate MAC table overflow attacks, network administrators must implement port security.

A VLAN hopping attack enables traffic from one VLAN to be seen by another VLAN without the aid of a router. The threat actor configures a host to act like a switch to take advantage of the automatic trunking port feature enabled by default on most switch ports.

A VLAN double-tagging attack is unidirectional and works only when the threat actor is connected to a port residing in the same VLAN as the native VLAN of the trunk port. Double tagging allows the threat actor to send data to hosts or servers on a VLAN that otherwise would be blocked by some type of access control configuration. Return traffic will also be permitted, letting the threat actor communicate with devices on the normally blocked VLAN.

VLAN hopping and VLAN double-tagging attacks can be prevented by implementing the following trunk security guidelines:

- Disable trunking on all access ports.
- Disable auto trunking on trunk links so that trunks must be manually enabled.
- Be sure that the native VLAN is only used for trunk links.

DHCP Attack: DHCP servers dynamically provide IP configuration information including IP address, subnet mask, default gateway, DNS servers, and more to clients. Two types of DHCP attacks are DHCP starvation and DHCP spoofing. Both attacks are mitigated by implementing DHCP snooping.

ARP Attack: A threat actor sends a gratuitous ARP message containing a spoofed MAC address to a switch, and the switch updates its MAC table accordingly. Now the threat actor sends unsolicited ARP Requests to other hosts on the subnet with the MAC Address of the threat actor and the IP address of the default gateway. ARP spoofing and ARP poisoning are mitigated by implementing DAI.

Address Spoofing Attack: IP address spoofing is when a threat actor hijacks a valid IP address of another device on the subnet or uses a random IP address. MAC address spoofing attacks occur when the threat actors alter the MAC address of their host to match another known MAC address of a target host. IP and MAC address spoofing can be mitigated by implementing IPSG.

STP Attack: Threat actors manipulate STP to conduct an attack by spoofing the root bridge and changing the topology of a network. Threat actors make their hosts appear as root bridges; therefore, capturing all traffic for the immediate switched domain. This STP attack is mitigated by implementing BPDU Guard on all access ports

CDP Reconnaissance: CDP information is sent out CDP-enabled ports in a periodic, unencrypted multicast. CDP information includes the IP address of the device, IOS software version, platform, capabilities, and the native VLAN. The device receiving the CDP message updates its CDP database. the information provided by CDP can also be used by a threat actor to discover network infrastructure vulnerabilities. To mitigate the exploitation of CDP, limit the use of CDP on devices or ports.

Implement Port Security

Secure Unused Ports

Simple method:

- Disable all unused ports on a switch
- Navigate to each unused port and issue the Cisco IOS **shutdown** command
- Reactivate port with **no shutdown** command

To configure a range of ports, use **interface range** command

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

Example: shutdown ports F0/8 through F0/24 on S1

```
S1(config)# interface range fa0/8 - 24
S1(config-if-range)# shutdown
%LINK-5-CHANGED: Interface FastEthernet0/8, changed state to administratively down
(output omitted)
%LINK-5-CHANGED: Interface FastEthernet0/24, changed state to administratively down
S1(config-if-range)#
```

Mitigate MAC Address Table Attacks

Simple method:

- Enable port security

Port Security

- Limits the number of valid MAC addresses allowed on a port
- Allows an administrator to manually configure MAC addresses for a port or to permit the switch to dynamically learn a limited number of MAC addresses
    - When a port configured with port security receives a frame, the source MAC address of the frame is compared to the list of secure source MAC addresses that were manually configured or dynamically learned on the port

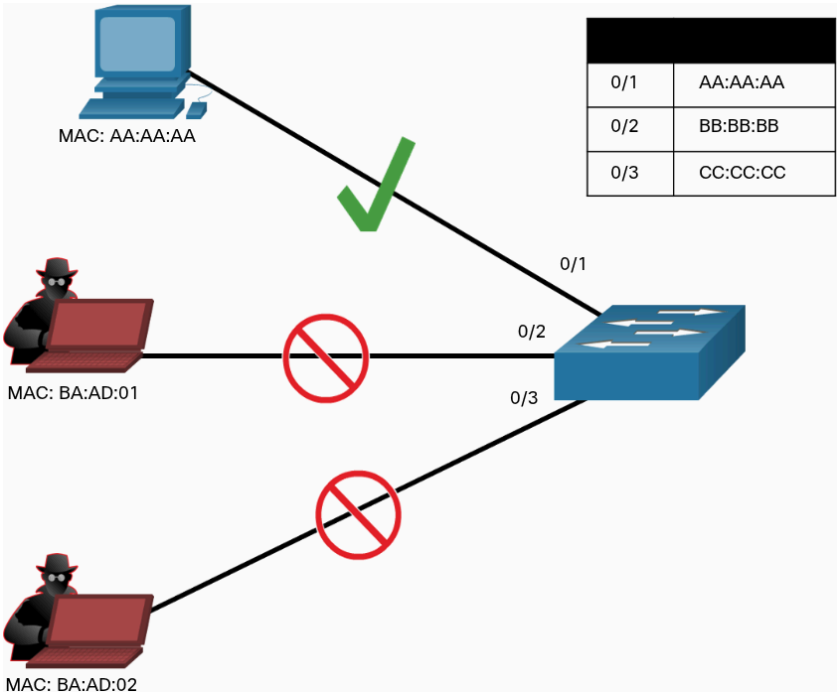

Enable Port Security

- In example below, the **switchport port-security** command was rejected
    - Port security can only be configured on **manually configured** access ports or **manually configured** trunk ports
- By default, Layer 2 switch ports are set to dynamic auto (trunking on). Therefore, in the example, the port is configured with the **switchport mode access** interface configuration command

```
S1(config)# interface f0/1
S1(config-if)# switchport port-security
Command rejected: FastEthernet0/1 is a dynamic port.
S1(config-if)# switchport mode access
S1(config-if)# switchport port-security
S1(config-if)# end
S1#
```

show port-security interface command

- Displays the current port security settings for selected port

In example below:

Notice how port security is enabled, port status is **Secure-down**, which means **there are no devices attached and no violation has occurred**

- Violation mode is Shutdown, and how the maximum number of MAC addresses is 1

If a device is connected to the port, the switch port status would display **Secure-up** and the switch will automatically add the device's MAC address as a secure MAC

```
S1# show port-security interface f0/1
Port Security              : Enabled
Port Status                : Secure-down
Violation Mode             : Shutdown
Aging Time                 : 0 mins
Aging Type                 : Absolute
SecureStatic Address Aging : Disabled
Maximum MAC Addresses      : 1
Total MAC Addresses        : 0
Configured MAC Addresses   : 0
Sticky MAC Addresses       : 0
Last Source Address:Vlan   : 0000.0000.0000:0
Security Violation Count   : 0
S1#
```

Note - If an active port is configured with **switchport port-security** command and more than one device is connected to that port, the port will transition to the **error-disabled** state

After port security is enabled, other port security specifics can be configured

```
S1(config-if)# switchport port-security ?
  aging         Port-security aging commands
  mac-address   Secure mac address
  maximum       Max secure addresses
  violation     Security violation mode
S1(config-if)# switchport port-security
```

Limit and Learn Mac Addresses

To set a maximum number of MAC addresses allowed on a port

```
Switch(config-if)# switchport port-security maximum value
```

The default port security value is 1. The maximum number of secure MAC addresses that can be configured depends the switch and the IOS

- In example below, the maximum is 8192

```
S1(config)# interface f0/1
S1(config-if)# switchport port-security maximum ?
  <1-8192>  Maximum addresses
S1(config-if)# switchport port-security maximum
```

The switch can be configured to learn about MAC addresses on a secure port in one of three ways:

1. **Manually configured**
   The administrator manually configures a static MAC address(es)

   ```
   Switch(config-if)# switchport port-security mac-address mac-address
   ```

2. **Dynamically Learned**
   When the **switchport port-security** command is entered, the current source MAC for the device connected to the port is automatically secured but is not added to the startup configuration.
   - If the switch is rebooted, the port will have to re-learn the device's MAC address

3. **Dynamically Learned - Sticky**
   The administrator can enable the switch to dynamically learn the MAC address and "stick" them to the running configuration

   ```
   Switch(config-if)# switchport port-security mac-address sticky
   ```

   Saving the running configuration will commit the dynamically learned MAC address to NVRAM

Example below:

- Demonstrates a complete port security configuration for F0/1 with a host connected to it.
- The administrator specifies a maximum of 2 MAC addresses
    - Manually configures one secure MAC address
    - Then configures the port to dynamically learn additional secure MAC addresses up to 2 secure MAC address maximum

Use **show port-security interface** and the **show port-security address** command to verify the configuration

```
*Mar  1 00:12:38.179: %LINK-3-UPDOWN: Interface FastEthernet0/1, changed state to up
*Mar  1 00:12:39.194: %LINEPROTO-5-UPDOWN: Line protocol on Interface FastEthernet0/1,
changed state to up
S1#conf t
Enter configuration commands, one per line.  End with CNTL/Z.
S1(config)#
S1(config)# interface fa0/1
S1(config-if)# switchport mode access
S1(config-if)# switchport port-security
S1(config-if)# switchport port-security maximum 2
S1(config-if)# switchport port-security mac-address aaaa.bbbb.1234
S1(config-if)# switchport port-security mac-address sticky
S1(config-if)# end
S1# show port-security interface fa0/1
Port Security              : Enabled
Port Status                : Secure-up
Violation Mode             : Shutdown
Aging Time                 : 0 mins
Aging Type                 : Absolute
SecureStatic Address Aging : Disabled
Maximum MAC Addresses      : 2
Total MAC Addresses        : 2
Configured MAC Addresses   : 1
Sticky MAC Addresses       : 1
Last Source Address:Vlan   : a41f.7272.676a:1
Security Violation Count   : 0
S1# show port-security address
             Secure Mac Address Table
------------------------------------------------------------------------
Vlan    Mac Address       Type                          Ports   Remaining Age
                                                                    (mins)
----    -----------       ----                          -----   -------------
1    a41f.7272.676a    SecureSticky                 Fa0/1        -
1    aaaa.bbbb.1234    SecureConfigured             Fa0/1        -
------------------------------------------------------------------------
Total Addresses in System (excluding one mac per port)      : 1
Max  Addresses limit in System (excluding one mac per port) : 8192
S1#
```

The output of **show port-security interface** command verifies that port security is enabled, there is a host connected to the port (IE: secure-up), a total of 2 MAC addresses will be allowed, and S1 has learned one MAC address statically and one MAC address dynamically (IE: sticky)

The output of **show port-security address** command lists the two learned MAC addresses

Port Security Aging

Can be used to set the aging time for static and dynamic secure addresses on a port

- **Absolute** - The secure addresses on the port are deleted after the specified aging time
- **Inactivity** - The secure addresses on the port are deleted only if they are inactive for the specified aging time

Use aging to remove secure MAC addresses on a secure port without manually deleting the existing secure MAC addresses.

- Aging time limits can also be increased to ensure past secure MAC addresses remain, even while new MAC addresses are added
- Statically configured secured addresses can be enabled or disabled on a per-port basis

Use the **switchport port-security aging** command to enable or disable static aging for the secure port, or to set the aging time or type

```
Switch(config-if)# switchport port-security aging {  static | time time |
type  {absolute | inactivity}
```

Parameter	Description
static	Enable aging for statically configured secure addresses on this port
time *time*	Specify the aging time for this port. The range is 0 to 1440 minutes. If the time is 0, aging is disabled for this port
type absolute	Set the absolute aging time. All the secure addresses on this port age out exactly after the time (in minutes) specified and are removed from the secure address list
type inactivity	Set the inactivity aging type. The secure addresses on this port age out only if there is no data traffic from the secure source address for the specified time period

Example below shows an administrator configuring the aging type to 10 minutes of inactivity and by using the **show port-security interface** command to verify the configuration

```
S1(config)# interface fa0/1
S1(config-if)# switchport port-security aging time 10
S1(config-if)# switchport port-security aging type inactivity
S1(config-if)# end
S1# show port-security interface fa0/1
Port Security              : Enabled
Port Status                : Secure-up
Violation Mode               : Shutdown
Aging Time                   : 10 mins
Aging Type                   : Inactivity
SecureStatic Address Aging : Disabled
Maximum MAC Addresses      : 2
Total MAC Addresses        : 2
Configured MAC Addresses   : 1
Sticky MAC Addresses         : 1
Last Source Address:Vlan   : a41f.7272.676a:1
Security Violation Count   : 0
S1#
```

Port Security Violation Modes

If the MAC address of a device attached to the port differs from the list of secure addresses, then a port violation occurs

- By default, the port enters the **error-disabled** state

To set the port security violation mode

```
Switch(config-if)# switchport port-security violation {  protect |
restrict |  shutdown}
```

Security Violation Mode Descriptions

Mode	Description
shutdown (default)	The port transitions to the error-disabled state immediately, turns off the port LED, and sends a syslog message. It increments the violation counter. When a secure port is in the error-disabled state, an administrator must re-enable it by entering the **shutdown** and **no shutdown** commands
restrict	The port drops packets with unknown source addresses until you remove a sufficient number of secure MAC addresses to drop below the maximum value or increase the maximum value. This mode causes the Security Violation counter to increment and generates a syslog message
protect	This is the least secure of the security violation modes. The port drops packets with unknown MAC source addresses until you remove a sufficient number of secure MAC addresses to drop below the maximum value or increase the maximum value. No syslog message is sent.

Security Violation Mode Comparison

Violation Mode	Discards Offending Traffic	Sends Syslog Message	Increase Violation Counter	Shuts Down Port
Protect	Yes	No	No	No
Restrict	Yes	Yes	Yes	No
Shutdown	Yes	Yes	Yes	Yes

Example below: Shows an administrator changing the security violation to “restrict”

- The output **show port-security interface** command confirms changes

```
S1(config)# interface f0/1
S1(config-if)# switchport port-security violation restrict
S1(config-if)# end
S1#
S1# show port-security interface f0/1
Port Security              : Enabled
Port Status                : Secure-up
Violation Mode           : Restrict
Aging Time                 : 10 mins
Aging Type                 : Inactivity
SecureStatic Address Aging : Disabled
Maximum MAC Addresses      : 2
Total MAC Addresses        : 2
Configured MAC Addresses   : 1
Sticky MAC Addresses     : 1
Last Source Address:Vlan   : a41f.7272.676a:1
Security Violation Count   : 0
S1#
```

Ports in error-disabled state

What happens when the port security violation is shutdown and a port violation occurs?

- The port is physically shutdown and placed in the error-disabled state, and no traffic is sent or received on the port

Figure below: The port security violation is changed back to the default shutdown setting. Then the host with MAC address of a41f.7272.676a is disconnected and a new host is plugged into F0/1

```
S1(config)# int fa0/1
S1(config-if)# switchport port-security violation shutdown
S1(config-if)# end
S1#
*Mar  1 00:24:15.599: %LINEPROTO-5-UPDOWN: Line protocol on Interface FastEthernet0/1, changed
state to down
*Mar  1 00:24:16.606: %LINK-3-UPDOWN: Interface FastEthernet0/1, changed state to down
*Mar  1 00:24:19.114: %LINK-3-UPDOWN: Interface FastEthernet0/1, changed state to up
*Mar  1 00:24:20.121: %LINEPROTO-5-UPDOWN: Line protocol on Interface FastEthernet0/1, changed
state to up
S1#
*Mar  1 00:24:32.829: %PM-4-ERR_DISABLE: psecure-violation error detected on Fa0/1, putting
Fa0/1 in err-disable state
*Mar  1 00:24:32.838: %PORT_SECURITY-2-PSECURE_VIOLATION: Security violation occurred, caused
by MAC address a41f.7273.018c on port FastEthernet0/1.
*Mar  1 00:24:33.836: %LINEPROTO-5-UPDOWN: Line protocol on Interface FastEthernet0/1, changed
state to down
*Mar  1 00:24:34.843: %LINK-3-UPDOWN: Interface FastEthernet0/1, changed state to down
S1#
```

Note - The port protocol and link status are changed to down and the port LED is turned off

Example below: The **show interface** command identifies the port status as **err-disabled**. The output of **show port-security** interface command now shows the port status as Secure-shutdown instead of Secure-up. The Security Violation counter increments by 1.

```
S1# show interface fa0/1 | include down
FastEthernet0/18 is down, line protocol is down (err-disabled)
(output omitted)
S1# show port-security interface fa0/1
Port Security              : Enabled
Port Status                : Secure-shutdown
Violation Mode             : Shutdown
Aging Time                 : 10 mins
Aging Type                 : Inactivity
SecureStatic Address Aging : Disabled
Maximum MAC Addresses      : 2
Total MAC Addresses        : 2
Configured MAC Addresses   : 1
Sticky MAC Addresses       : 1
Last Source Address:Vlan   : a41f.7273.018c:1
Security Violation Count   : 1
S1#
```

The administrator should determine what caused the security violation if an unauthorized device is connected to a secure port, the security threat is eliminated before re-enabling the port.

Example below: The first host is reconnected to F0/1. To re-enable the port, first use the **shutdown** command, then, use the **no shutdown** command to make the port operational

```
S1(config)# interface fa0/1
S1(config-if)# shutdown
S1(config-if)#
*Mar  1 00:39:54.981: %LINK-5-CHANGED: Interface FastEthernet0/1, changed state to
administratively down
S1(config-if)# no shutdown
S1(config-if)#
*Mar  1 00:40:04.275: %LINK-3-UPDOWN: Interface FastEthernet0/1, changed state to up
*Mar  1 00:40:05.282: %LINEPROTO-5-UPDOWN: Line protocol on Interface FastEthernet0/1, changed
state to up
S1(config-if)#
```

Verify Port Security

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

Port Security for All Interfaces

To display port security settings for the switch, use the **show port-security** command

- Example below indicates that only one port is configured with the switchport port-security command

```
S1# show port-security
Secure Port  MaxSecureAddr  CurrentAddr  SecurityViolation  Security Action
                (Count)        (Count)          (Count)
---------------------------------------------------------------------------
      Fa0/1              2            2                  0          Shutdown
---------------------------------------------------------------------------
Total Addresses in System (excluding one mac per port)  : 1
Max Addresses limit in System (excluding one mac per port) : 8192
S1#
```

Port Security for a Specific Interface

Use the **show port-security interface** command to view details for a specific interface

```
S1# show port-security interface fastethernet 0/1
Port Security               : Enabled
Port Status                 : Secure-up
Violation Mode              : Shutdown
Aging Time                  : 10 mins
Aging Type                  : Inactivity
SecureStatic Address Aging  : Disabled
Maximum MAC Addresses    : 2
Total MAC Addresses      : 2
Configured MAC Addresses    : 1
Sticky MAC Addresses        : 1
Last Source Address:Vlan    : a41f.7273.018c:1
Security Violation Count    : 0
S1#
```

Verify Learned MAC Addresses

To verify that MAC addresses are “sticking”, use **show run** command

```
S1# show run interface fa0/1
Building configuration...
Current configuration : 365 bytes
!
interface FastEthernet0/1
switchport mode access
switchport port-security maximum 2
switchport port-security mac-address sticky
switchport port-security mac-address sticky a41f.7272.676a
switchport port-security mac-address aaaa.bbbb.1234
switchport port-security aging time 10
```

```
switchport port-security aging type inactivity
switchport port-security
end
S1#
```

Verify Secure MAC Addresses

To display all secure MAC addresses that are manually configured or dynamically learned on all switch interfaces, use the **show port-security address** command

```
S1# show port-security address
            Secure Mac Address Table
------------------------------------------------------------------------------
Vlan    Mac Address      Type                          Ports   Remaining Age
                                                                    (mins)
----    -----------      ----                          -----   -------------
   1    a41f.7272.676a    SecureSticky                  Fa0/1        -
   1    aaaa.bbbb.1234    SecureConfigured              Fa0/1        -
------------------------------------------------------------------------------
Total Addresses in System (excluding one mac per port)  : 1
Max Addresses limit in System (excluding one mac per port) : 8192
S1#
```

Mitigate VLAN Attacks

VLAN Attacks Review

- Spoofing DTP messages from the attacking host to cause the switch to enter trunking mode.
    - From here, the attacker can send traffic tagged with the target VLAN, and the switch then delivers the packets to the destination
- Introducing a rogue switch and enabling trunking.
    - The attacker can then access all VLANs on the victim switch from the rogue switch.
- Another type of VLAN hopping attack is a double-tagging (or double-encapsulated) attack.
    - This attack takes advantage of the way hardware on most switches operate

Steps to Mitigate VLAN Hopping Attacks

1. Disable DTP (auto trunking) negotiations on non-trunking ports by using the **switchport mode access** interface configuration command
2. Disable unused ports and put them in an unused VLAN
3. Manually enable the trunk link on a trunking port by using the **switchport mode trunk** command
4. Disable DTP (auto trunking) negotiations on trunking ports by using the **switchport nonegotiate** command
5. Set the native VLAN to a VLAN other than VLAN 1 by using **switchport trunk native vlan** *vlan_number* command

Example below, assume:

- F0/1 through F0/16 are active access ports
- F0/17 through F0/20 are not currently in use
- F0/21 through F0/25 are trunk ports

VLAN hopping can be mitigated by implementing:

```
S1(config)# interface range fa0/1 - 16
S1(config-if-range)# switchport mode access
S1(config-if-range)# exit
S1(config)#
S1(config)# interface range fa0/17 - 20
S1(config-if-range)# switchport mode access
S1(config-if-range)# switchport access vlan 1000
S1(config-if-range)# shutdown
S1(config-if-range)# exit
S1(config)#
S1(config)# interface range fa0/21 - 24
S1(config-if-range)# switchport mode trunk
S1(config-if-range)# switchport nonegotiate
S1(config-if-range)# switchport trunk native vlan 999
S1(config-if-range)# end
S1#
```

- F0/1 to 0/16 are access ports and therefore trunking is disabled by explicitly making them access ports
- F0/17 to 0/20 are unused ports and are disabled and assigned to an unused VLAN
- F0/21 to 0/24 are trunk links and are manually enabled as trunks with DTP disabled
    - The native VLAN is also changed from default VLAN 1 to an unused VLAN 999

Mitigate DHCP Attacks

DHCP Attack Review

The goal of DHCP starvation attack is to create a Denial of Service (DoS) for connecting clients.

- Requires an attack tool such as Gobbler

Recall - DHCP starvation attacks can be effectively mitigated by using port security because Gobbler uses a unique source MAC address for each DHCP request sent.

Mitigating DHCP spoofing attacks requires more protection.

- Gobbler could be configured to use the actual interface MAC address as the source Ethernet address, but specify a different Ethernet address in the DHCP payload
    - This would render port security ineffective because the source MAC address would be legitimate

DHCP spoofing attacks can be mitigated by using DHCP snooping on trusted ports

DHCP Snooping

Does not rely on source MAC addresses

- DHCP snooping determines whether DHCP messages are from an administratively configured trusted or untrusted source
- It then filters DHCP messages and rate-limits DHCP traffic from untrusted sources

Trusted devices: devices under administrative control IE switches, routers, servers

- Any device beyond the firewall or outside network is an untrusted source.
- All access ports are generally treated as untrusted sources

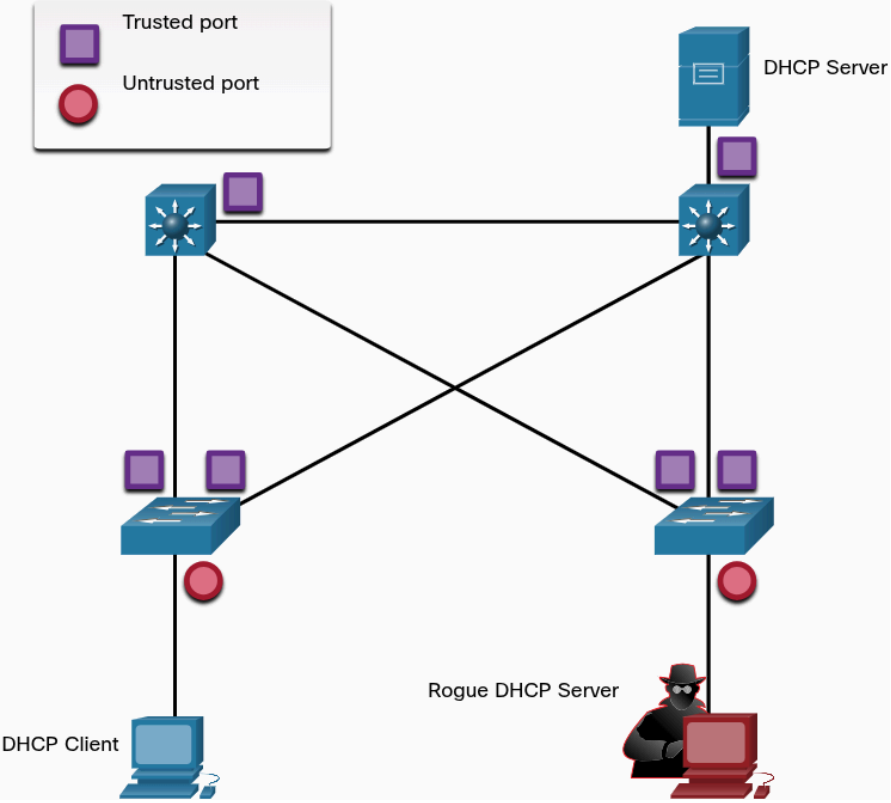

Notice: The rogue DHCP server would be on an untrusted port after enabling DHCP snooping. All interfaces are treated as untrusted by default.

- Trusted interfaces are typically trunk links and ports directly connected to a legit DHCP server. These interfaces must be explicitly configured as trusted

DHCP snooping binding table

- A DHCP table is built that includes the source MAC address of a device on an untrusted port and the IP address assigned by the DHCP server to that device. The MAC address and IP address are bound together

Steps to Implement DHCP Snooping

1. Enable DHCP snooping by using the **ip dhcp snooping** global configuration command
2. On trusted ports, use the **ip dhcp snooping trust** interface configuration command
3. Limit the number of DHCP discover messages that can be received per second on untrusted ports by using the **ip dhcp snooping limit rate** interface configuration command
4. Enable DHCP snooping by VLAN, or by a range of VLANs, by using the **ip dhcp snooping** *vlan* global configuration command

DHCP Snooping Configuration Example

Notice that F0/5 is an untrusted port because it connects to a PC. F0/1 is a trusted port because it connects to the DHCP server

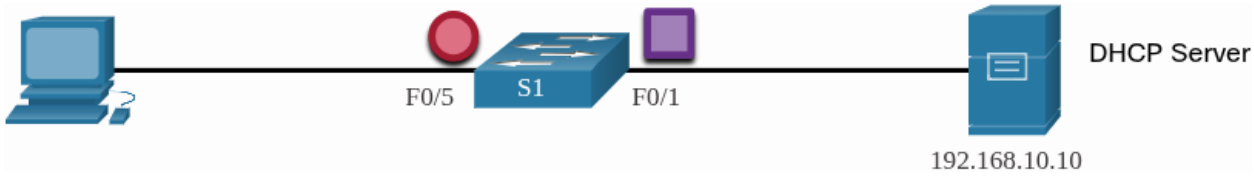

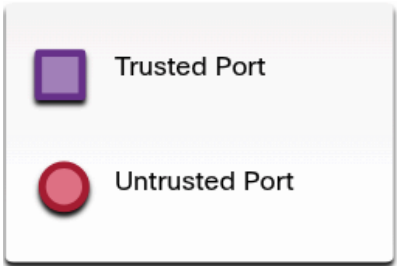

Example below: How to configure DHCP snooping on S1.

- Notice how DHCP snooping is first enabled.
- Then the upstream interface to the DHCP server is explicitly trusted.
- Next, the range of FastEthernet ports from F0/5 to F0/24 are untrusted by default, so a rate limit is set to six packets per second.
- Finally, DHCP snooping is enabled on VLANs 5, 10, 50, 51, and 52

```
S1(config)# ip dhcp snooping
S1(config)# interface f0/1
S1(config-if)# ip dhcp snooping trust
S1(config-if)# exit
S1(config)# interface range f0/5 - 24
S1(config-if-range)# ip dhcp snooping limit rate 6
S1(config-if-range)# exit
S1(config)# ip dhcp snooping vlan 5,10,50-52
S1(config)# end
S1#
```

Use the **show ip dhcp snooping** privileged EXEC command to verify DHCP snooping and **show ip dhcp snooping binding** to view the clients that have received DHCP information

Note - DHCP snooping is also required by Dynamic ARP Inspection (DAI)

```
S1# show ip dhcp snooping
Switch DHCP snooping is enabled
DHCP snooping is configured on following VLANs:
5,10,50-52
DHCP snooping is operational on following VLANs:
none
DHCP snooping is configured on the following L3 Interfaces:
Insertion of option 82 is enabled
   circuit-id default format: vlan-mod-port
   remote-id: 0cd9.96d2.3f80 (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            yes        yes               unlimited
  Custom circuit-ids:
FastEthernet0/5            no         no              6
  Custom circuit-ids:
FastEthernet0/6            no         no              6
  Custom circuit-ids:
S1# show ip dhcp snooping binding
MacAddress         IpAddress    Lease(sec) Type          VLAN Interface
------------------ ---------------- ---------- ------------- ----
--------------------
00:03:47:B5:9F:AD  192.168.10.11   193185  dhcp-snooping 5    FastEthernet0/5
```

Mitigate ARP Attacks

Dynamic ARP Inspection

A threat actor can send unsolicited ARP requests to other hosts on the subnet with the MAC address of the threat actor and the IP address of the default gateway. To prevent ARP spoofing and the resulting ARP poisoning, a switch must ensure that only valid ARP Request and Replies are relayed.

Dynamic ARP Inspection (DAI) requires DHCP snooping and helps prevent ARP attacks by:

- Not relaying invalid or gratuitous ARP Requests out to other ports in the same VLAN
- Intercepting all ARP Request and Replies on untrusted ports
- Verifying each intercept packet for a valid IP-to-MAC binding
- Dropping and logging ARP Requests coming from invalid sources to prevent ARP poisoning
- Error-disabling the interface if the configured DAI number of ARP packets is exceeded

DAI Implementation Guidelines

To mitigate the chances of ARP spoofing and ARP poisoning:

- Enable DHCP snooping globally
- Enable DHCP snooping on select VLANs
- Enable DAI on select VLANs
- Configure trusted interfaces for DHCP snooping and ARP inspection

Generally advisable to configure all access switch ports are untrusted and to configure all uplink ports that are connected to other switches as trusted

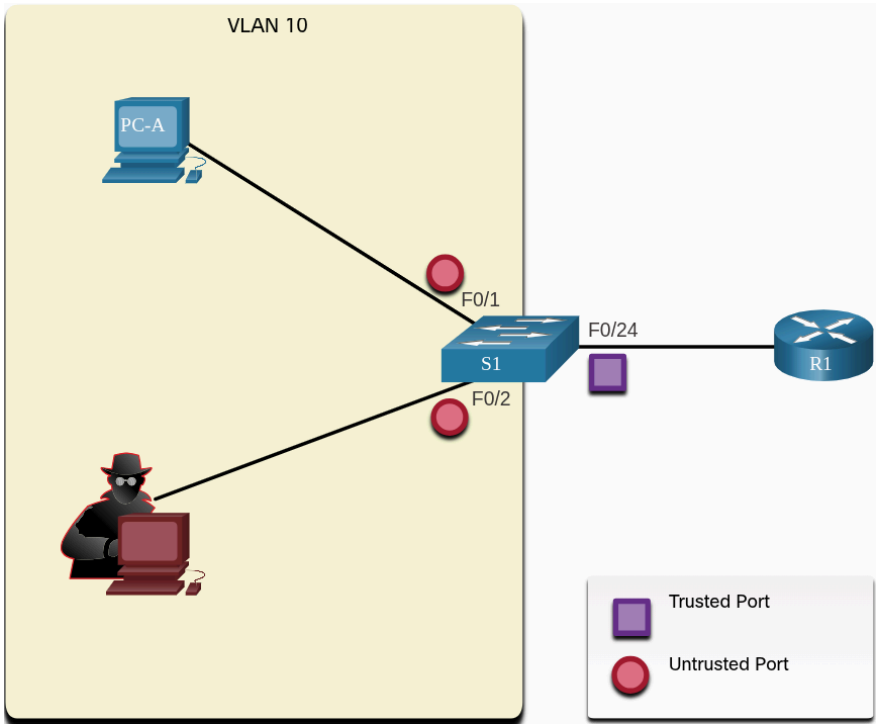

DAI Configuration Example

Example below:

- DHCP snooping is enabled bc DAI requires the DHCP snooping binding table to operate
- Next, DHCP snooping and ARP inspection are enabled for the PCs on VLAN10
    - The uplink port to the router is trusted, and therefore, is configured as trusted fro DHCP snooping and ARP inspection

```
S1(config)# ip dhcp snooping
S1(config)# ip dhcp snooping vlan 10
S1(config)# ip arp inspection vlan 10
S1(config)# interface fa0/24
S1(config-if)# ip dhcp snooping trust
S1(config-if)# ip arp inspection trust
```

DAI can also be configured to check for both destination or source MAC and IP addresses:

- **Destination MAC** - Checks the destination MAC address in the Ethernet header against the target MAC address in ARP body
- **Source MAC** - Checks the source MAC address in the Ethernet header against the sender MAC address in the ARP body
- **IP Address** - Checks the ARP body for invalid and unexpected IP addresses including address 0.0.0.0, 255.255.255.255, and all IP multicast addresses

The **ip arp inspection validate {[src-mac] [dst-mac][ip]}** global configuration command

- Used to configure DAI to drop ARP packets when the IP addresses are invalid
- Can be used when the MAC addresses in the body of the ARP packets do not match the addresses that are specified in the Ethernet header

Example below: notice how only one command can be configured.
Therefore, entering multiple **ip arp inspection validate** commands overwrites the previous commands

- To include more than one validation method, enter them on the same command line

```
S1(config)# ip arp inspection validate ?
dst-mac  Validate destination MAC address
  ip         Validate IP addresses
  src-mac  Validate source MAC address
S1(config)# ip arp inspection validate src-mac
S1(config)# ip arp inspection validate dst-mac
S1(config)# ip arp inspection validate ip
S1(config)# do show run | include validate
ip arp inspection validate ip
S1(config)# ip arp inspection validate src-mac dst-mac ip
S1(config)# do show run | include validate
ip arp inspection validate src-mac dst-mac ip
S1(config)#
```

Mitigate STP Attacks

PortFast and BPDU Guard

Recall - Network attackers can manipulate the Spanning Tree Protocol (STP) to conduct an attack by spoofing the root bridge and changing the topology of a network

To mitigate STP manipulation attacks, use PortFast and Bridge Protocol Data Unit (BPDU) Guard

- **PortFast** - Portfast immediately brings an interface configured as an access port to the forwarding state from a blocking state, bypassing the listening and learning states.
    - Apply to all end-user ports. PortFasts should only be configured on ports attached to end devices

- **BPDU Guard** - BPDU guard immediately error disables a port that receives a BPDU. Like PortFast, BPDU guard should only be configured on interfaces attached to the end devices

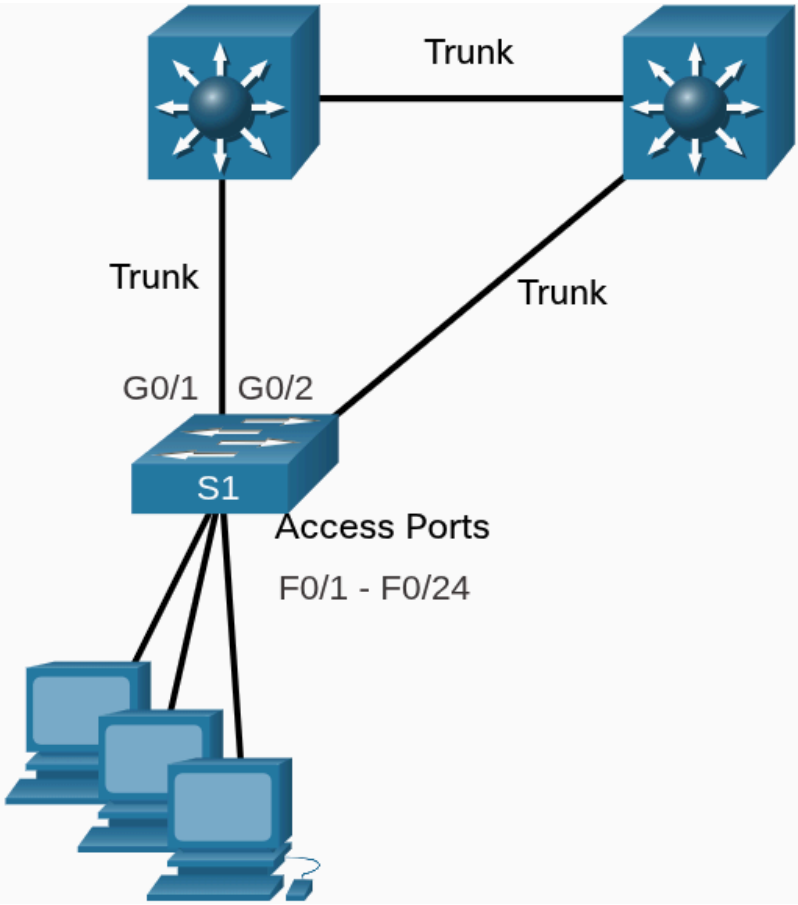

Configure PortFast

PortFast bypasses the STP listening and learning states to minimize the time that access ports must wait for STP to converge

- If PortFast is enabled on a port connecting to another switch, there is a risk of creating a spanning-tree loop

PortFast can be enabled on an interface by using the **spanning-tree portfast** interface configuration command.

- Alternatively, PortFast can be configured globally on all access ports by using the **spanning-tree portfast default** global configuration command

To verify whether PortFast is enabled globally you can use either the **show running-config | begin span** command or the **show spanning-tree summary** command.
To verify if PortFast is enabled on an interface, use the **show running-config interface** *type/number* command

The **show spanning-tree interface** *type/number* **detail** command can also be used

Notice that when PortFast is enabled, warning messages are displayed

```
S1(config)# interface fa0/1
S1(config-if)# switchport mode access
S1(config-if)# spanning-tree portfast
%Warning: portfast should only be enabled on ports connected to a
single
host. Connecting hubs, concentrators, switches, bridges, etc... to
this
interface when portfast is enabled, can cause temporary bridging
loops.
Use with CAUTION
%Portfast has been configured on FastEthernet0/1 but will only
have effect when the interface is in a non-trunking mode.
S1(config-if)# exit
S1(config)# spanning-tree portfast default
%Warning: this command enables portfast by default on all interfaces.
You
should now disable portfast explicitly on switched ports leading to
hubs,
switches and bridges as they may create temporary bridging loops.
S1(config)# exit
S1# show running-config | begin span
spanning-tree mode pvst
spanning-tree portfast default
spanning-tree extend system-id
!
interface FastEthernet0/1
switchport mode access
spanning-tree portfast
!
interface FastEthernet0/2
!
interface FastEthernet0/3
!
interface FastEthernet0/4
!
interface FastEthernet0/5
!
(output omitted)
S1#
```

Configure BPDU Guard

Even though PortFast is enabled, the interface will still listen for BPDUs

- Unexpected BPDUs might be accidental, or part of an unauthorized attempt to add a switch to the network

If any BPDUs are received on a BPDU Guard enabled port, that port is put into error-disabled state

- This means the port is shut down and must be manually re-enabled or automatically recovered through the **errdisable recovery cause bpduguard** global command

BPDU Guard can be enabled on a port by using the **spanning-tree bpduguard enable** interface configuration command

- * Alternatively, use the **spanning-tree portfast bpduguard default** global command to globally enable BPDU guard on all PortFast-enabled ports

To display information about the state of spanning tree, use the **show spanning-tree summary** command

Example below: PortFast default and BPDU Guard are both enabled as the default state for ports configured as access mode

Note - Always enable BPDU Guard on all PortFast-enabled ports

```
S1(config)# interface fa0/1
S1(config-if)# spanning-tree bpduguard enable
S1(config-if)# exit
S1(config)# spanning-tree portfast bpduguard default
S1(config)# end
S1# show spanning-tree summary
Switch is in pvst mode
Root bridge for: none
Extended system ID           is enabled
Portfast Default             is enabled
PortFast BPDU Guard Default  is enabled
Portfast BPDU Filter Default is disabled
Loopguard Default            is disabled
EtherChannel misconfig guard is enabled
UplinkFast                   is disabled
BackboneFast                 is disabled
Configured Pathcost method used is short
(output omitted)
S1#
```

11.6.3 What did I learn in this module?

All switch ports (interfaces) should be secured before the switch is deployed for production use. The simplest and most effective method to prevent MAC address table overflow attacks is to enable port security. By default, Layer 2 switch ports are set to dynamic auto (trunking on). The switch can be configured to learn about MAC addresses on a secure port in one of three ways: manually configured, dynamically learned, and dynamically learned - sticky. Port security aging can be used to set the aging time for static and dynamic secure addresses on a port. Two types of aging are supported per port: absolute and inactivity. If the MAC address of a device attached to the port differs from the list of secure addresses, then a port violation occurs. By default, the port enters the error-disabled state. When a port is shutdown and placed in the error-disabled state, no traffic is sent or received on that port. To display port security settings for the switch, use the show port-security command.

To mitigate VLAN hopping attacks:

Step 1. Disable DTP negotiations on non-trunking ports.

Step 2. Disable unused ports.

Step 3. Manually enable the trunk link on a trunking port.

Step 4. Disable DTP negotiations on trunking ports.

Step 5. Set the native VLAN to a VLAN other than VLAN 1.

The goal of a DHCP starvation attack is to create a Denial of Service (DoS) for connecting clients. DHCP spoofing attacks can be mitigated by using DHCP snooping on trusted ports. DHCP snooping determines whether DHCP messages are from an administratively-configured trusted or untrusted source. It then filters DHCP messages and rate-limits DHCP traffic from untrusted sources. Use the following steps to enable DHCP snooping:

Step 1. Enable DHCP snooping.

Step 2. On trusted ports, use the ip dhcp snooping trust interface configuration command.

Step 3. Limit the number of DHCP discovery messages that can be received per second on untrusted ports.

Step 4. Enable DHCP snooping by VLAN, or by a range of VLANs.

Dynamic ARP inspection (DAI) requires DHCP snooping and helps prevent ARP attacks by:

- Not relaying invalid or gratuitous ARP Replies out to other ports in the same VLAN.
- Intercepting all ARP Requests and Replies on untrusted ports.
- Verifying each intercepted packet for a valid IP-to-MAC binding.
- Dropping and logging ARP Replies coming from invalid to prevent ARP poisoning.
- Error-disabling the interface if the configured DAI number of ARP packets is exceeded.

To mitigate the chances of ARP spoofing and ARP poisoning, follow these DAI implementation guidelines:

- Enable DHCP snooping globally.
- Enable DHCP snooping on selected VLANs.
- Enable DAI on selected VLANs.
- Configure trusted interfaces for DHCP snooping and ARP inspection.

As a general guideline, configure all access switch ports as untrusted and all uplink ports that are connected to other switches as trusted.

DAI can also be configured to check for both destination or source MAC and IP addresses:

- Destination MAC - Checks the destination MAC address in the Ethernet header against the target MAC address in ARP body.
- Source MAC - Checks the source MAC address in the Ethernet header against the sender MAC address in the ARP body.
- IP address - Checks the ARP body for invalid and unexpected IP addresses including addresses 0.0.0.0, 255.255.255.255, and all IP multicast addresses.

To mitigate Spanning Tree Protocol (STP) manipulation attacks, use PortFast and Bridge Protocol Data Unit (BPDU) Guard:

- PortFast - PortFast immediately brings an interface configured as an access or trunk port to the forwarding state from a blocking state, bypassing the listening and learning states. Apply to all end-user ports. PortFast should only be configured on ports attached to end devices. PortFast bypasses the STP listening and learning states to minimize the time that access ports must wait for STP to converge. If PortFast is enabled on a port connecting to another switch, there is a risk of creating a spanning-tree loop.
- BPDU Guard - BPDU guard immediately error disables a port that receives a BPDU. Like PortFast, BPDU guard should only be configured on interfaces attached to end devices. BPDU Guard can be enabled on a port by using the spanning-tree bpduguard enable interface configuration command. Alternatively, Use the spanning-tree portfast bpduguard default global configuration command to globally enable BPDU guard on all PortFast-enabled ports.

WLAN Concepts

Introduction to Wireless

Wireless LAN (WLAN)

- A type of wireless network commonly used in homes, offices, and campus environments.

Types of Wireless

- Based on IEEE standards

Wireless Personal-Area Networks (WPAN)

- Uses low powered transmitters for a short-range network
    - usually 20 to 30 ft (6 to 9 meters)
- Bluetooth and ZigBee based devices are commonly used in WPANs
- Based on 802.15 standard and a 2.4Ghz radio frequency

Wireless LANs (WLAN)

- Uses transmitters to cover a medium-sized network, usually up to 300 feet.
- Suitable for use in a home, office, and even campus environment
- Based on 802.11 standard
- 2.4Ghz or 5Ghz radio frequency

Wireless MANs (WMAN)

- Uses transmitters to provide wireless server over a larger geographic area
- Suitable for providing wireless access to a metropolitan city or specific district
- Use specific licensed frequencies

Wireless Wide-Area Networks (WWMANs)

- Uses transmitters to provide coverage over an extensive geographic area.
- Suitable for national and global communications
- Use specific licensed frequencies

Wireless Technologies

- Uses the unlicensed radio spectrum to send and receive data
- The unlicensed spectrum is accessible to anyone who has a wireless router and wireless technology in the device they are using

Bluetooth

An IEEE 802.15 WPAN standard that uses a device-pairing process to communicate over distances up to 300 ft (100m).

- **Bluetooth Low Energy (BLE)** - Supports multiple network technologies including mesh topology to large scale network devices
- **Bluetooth Basic Rate/Enhanced Rate (BR/EDR)** - Supports point to point topologies and is optimized for audio streaming

WiMAX (Worldwide Interoperability for Microwave Access)
An alternative to broadband wired internet connections, competing with DSL and Cable

- Typically used in areas that are not yet connected to a DSL or cable provider

- This is an IEEE 802.16 WWAN standard that provides high-speed wireless broadband access of up to 30 miles (50 km)
- Operates in similar way to Wi-Fi, but at higher speeds, over greater distances, and for a greater number of users
- Uses a network of WiMAX towers similar to cell phone towers
- Transmitters and cellular transmitters may share space on the same tower

Cellular Broadband
Cellular 4G/5G are wireless mobile networks primarily used by cellular phones but can be used in automobiles, tablets, and laptops

- Multi-access networks carrying both data and voice communications
- A cell site is created by a cellular tower transmitting signals in a given area
    - Interconnecting cell sites form the cellular network

Two types of cellular networks:

- **Global System for Mobile (GSM)**
    - Internationally recognized
- **Code Division Multiple Access (CDMA)**
    - Primarily used in the US

4th Generation mobile network (4F) is current mobile network

- Delivers speeds that are 10 times the previous 3G networks

New 5G holds promise of delivering 100 times faster speeds than 4G and connecting more devices to network than ever before

Satellite Broadband
Provides network access to remote sites through the use of a directional satellite dish that is aligned with a specific geostationary Earth orbit satellite

- Usually more expensive and requires a clear line of sight
- Typically used by rural homeowners and businesses where cable and DSL are not available

802.11 Standards

Defines how radio frequencies are used for wireless links

- Most of the standards specify that wireless devices have one antenna to transmit and receive wireless signals on the specified radio frequency (2.4Ghz or 5Gh)
- Some newer standards that transmit and receive at higher speeds require access points (APs) and wireless clients to have multiple antennas using the multiple-input and multiple-output (MIMO) technology

MIMO
Uses multiple antennas as both the transmitter and receive to improve communication performance

- Up to eight transmit and receive antennas can be used to increase throughput

IEEE WLAN Standard	**Radio Frequency**	Description
802.11	2.4 Ghz	● Speeds of up to 2 Mbps
802.11a	5 Ghz	● Speeds of up to 54 Mbps ● Small coverage area ● Less effective at penetrating building structures ● Not interoperable with 802.11b and 802.11g
802.11b	2.4 Ghz	● Speeds up to 11 Mbps ● Longer range than 802.11a ● Better able to penetrate building structures
802.11g	2.4 Ghz	● Speeds up to 54 Mbps ● Backward compatible with 802.11b with reduced bandwidth capacity
802.11n	2.4 Ghz 5 Ghz	● Data rates range from 150 Mbps to 600 Mbps with a distance range of up to 70 m (230 feet) ● APs and wireless clients require multiple antennas using MIMO technology ● Backward compatible with 802.11a/b/g devices with limiting data rates
802.11ac	5 Ghz	● Provides data rates ranging from 450 Mbps to 1.3 Gbps (1300 Mbps) using MIMO technology

		• Up to eight antennas can be supported • Backwards compatible with 802.11a/n devices with limiting data rates
802.11ax	2.4 Ghz 5 Ghz	• Latest standard released in 2019 • Also known as Wi-Fi 6 or High-Efficiency Wireless (HEW) • Provides improved power efficiency, higher data rates, increased capacity, and handles many connected devices • Currently operates using 2.4 Ghz and 5 Ghz but will use 1 Ghz and 7 Ghz when those frequencies become available

Radio Frequencies

All wireless devices operate in the radio waves range of the electromagnetic spectrum.

- WLAN networks operate in the 2.4 Ghz frequency band and the 5 Ghz band
- Wireless LAN devices have transmitters and receivers tuned to specific frequencies of the radio waves range

Following bands are allocated to 802.11 wireless LANs

- 2.4Ghz (UHF) - 802.11b/g/n/ax
- 5Ghz (SHF) - 802.11a/n/ac/ax

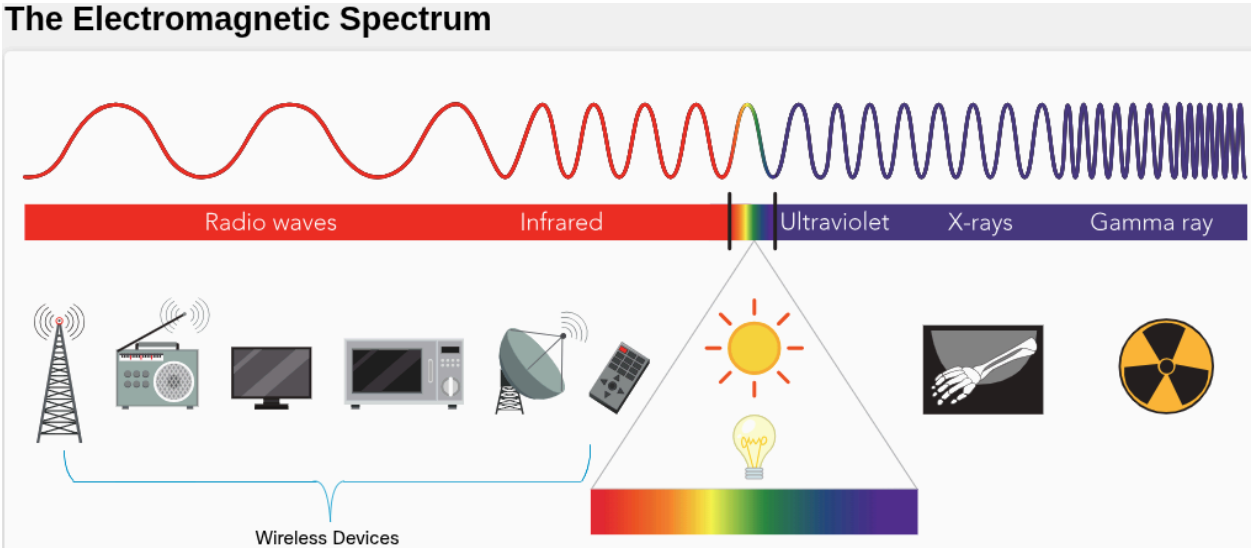

Wireless Standards Organizations

International Telecommunications Union (ITU)

- Regulates the allocation of the radio frequency spectrum and satellite orbits through the ITU-R
    - ITU-R stands for ITU Radiocommunication Sector

IEEE

- Specifies how a radio frequency is modulated to carry information
- Maintains the standards for local and metropolitan area networks (MAN) with IEEE 802 LAN/MAN family of standards
- The dominant standards in the IEEE 802 family are 802.3 Ethernet and 802.11 WLAN

Wi-Fi Alliance

- Global, non-profit, industry trade association devoted to promoting the growth and acceptance of WLANS
- An association of vendors whose objectives is to improve the interoperability of products based on 802.11 standard by certifying vendors for conformance to industry norms and adherence to standards

WLAN Components

Wireless NICs

- End devices with wireless NICs
- A network device, such as a wireless router and Wireless AP

USB Wireless Adapter

Wireless Home Router

- **Access Point** - Provides 802.11a/b/g/n/ac wireless access
- **Switch** - Provides a four-port, full-duplex, 10/100/1000 Ethernet switch to interconnect wired devices
- **Router** - Provides a default gateway for connecting to other network infrastructures

Commonly implemented as a small business or residential wireless access device.

- Advertises its wireless services by sending beacons containing its **shared service set identifier (SSID)**

- Devices discover the SSID and attempt to associate and authenticate with it

Most wireless routers provide advanced features, such as high-speed access, support for video streaming, IPv6 addressing, quality of service (QoS), configuration utilities, and USB ports tro connect printers or portable drives

- Wi-Fi range extenders
    - A device can connect wirelessly to the extenders, which boosts its communications to be repeated to the wireless router

Wireless Access Points

Provides dedicated wireless access to users

AP Categories

Autonomous APs

Standalone devices configured using a CLI or GUI

- Useful in situations where only a couple of APs are required
- A home router is an example
    - The entire AP configuration resides on device
- If demands increase, more APs would be required
- Each AP would operate independent of other APs and each AP would require manual configuration and management

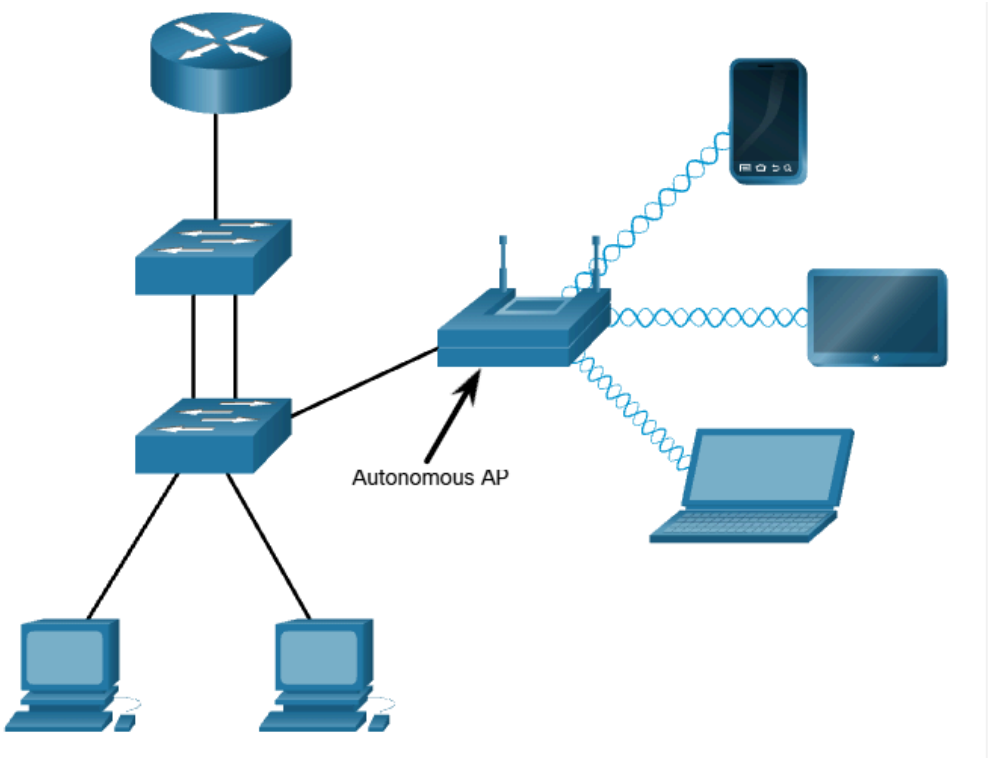

Controller-based APs

Require no initial configuration and often called lightweight APs (LAPs)

- Uses the Lightweight Access Point Protocol (LWAPP) to communicate with a WLAN Controller (WLC)
- Useful in situations where many APs are required in network
- Each AP is automatically configured and managed by the WLC

Example below: WLC has four ports connected to switching infrastructure. The four ports are configured as a link aggregation group (LAG) to bundle them together.

- LAG provides redundancy and load-balancing
- All ports on switch connected to the WLC need to be trunking and configured with EtherChannel on
- LAG does not operate exactly like EtherChannel
    - WLC does not support Port Aggregation Protocol (PaGP) or Link Aggregation Control Protocol (LACP)

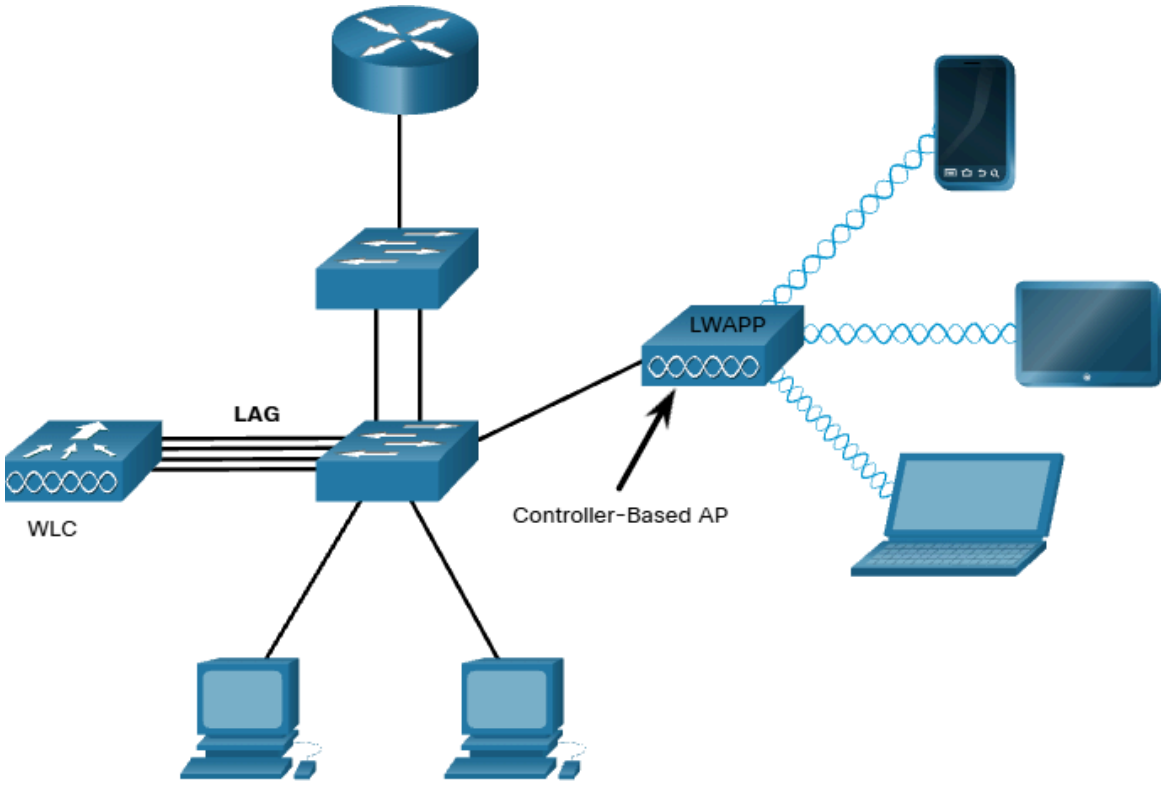

Wireless Antennas

Omnidirectional Antennas
Provides 360-degree coverage

- Ideal for horses, open office, conference rooms, outside

Directional Antennas
Focuses the radio signal in a given direction

- Enhances the signal to and from the AP in the direction the antenna is pointing
- Provides a stronger signal strength in one direction and reduced signal strength in all other directions

Multiple Input Multiple Output MIMO Antennas
Uses multiple antennas to increase available bandwidth for IEEE 802.11n/ac/ax

- Up to eight transmit and receive antennas can be used to increase throughput

WLAN Operation

Wireless Topology Modes

Ad hoc mode

When two devices connect wirelessly in a peer-to-peer (P2P) manner without using APs or wireless routers

- Example: Connecting directly to each other via Bluetooth or Wi-Fi Direct
- Referred as an independent basic service set (IBSS)

Infrastructure mode

When wireless clients interconnect via a wireless router or AP, such as WLANs

- APs connect to the network infrastructure using the wired distribution system

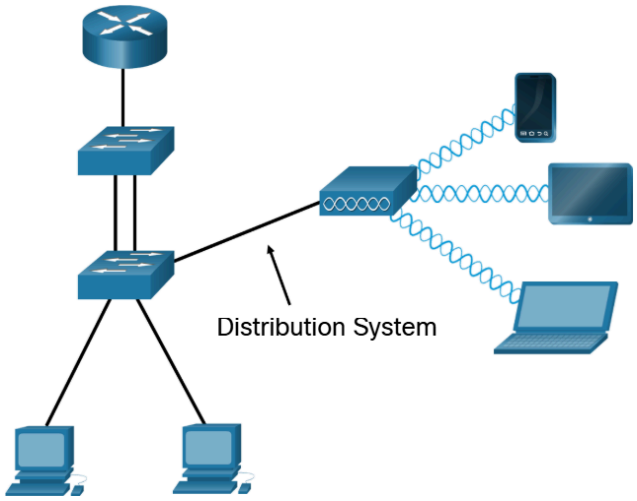

Tethering

A variation of Ad hoc topology– when a smart phone or tablet with cellular data access is enabled to create a personal hotspot

- Usually a temporary quick solution that enables a smart phone to provide the wireless services of a Wi-Fi router

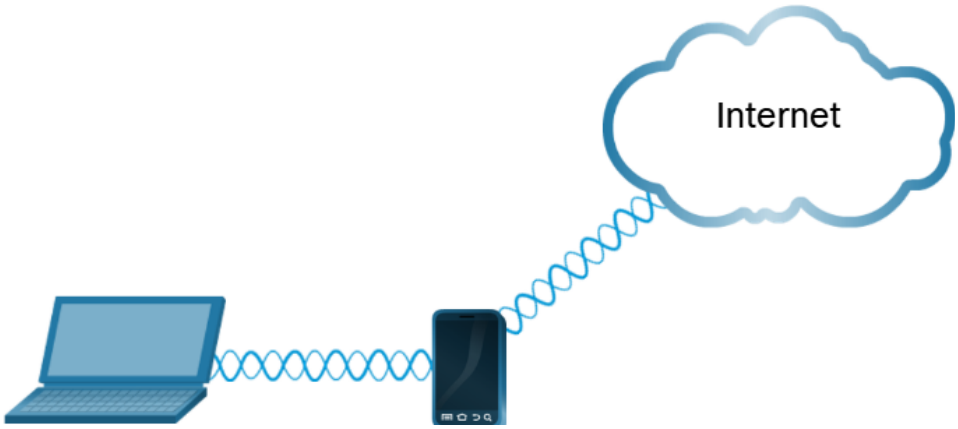

BSS and ESS

Basic Service Set (BSS)

Consists of a single AP interconnecting all associated wireless clients.

- Example displays circles depict the coverage area for the BSS, called Basic Service Set (BSA)
- If client moves out of its BSA, it can no longer directly communicate with other wireless clients within the BSA

The Layer 2 MAC address of the AP is used to uniquely identify each BSS, called the Basic Service Set Identifier (BSSID)

- BSSID is formal name of BSS and is always associated with only one AP

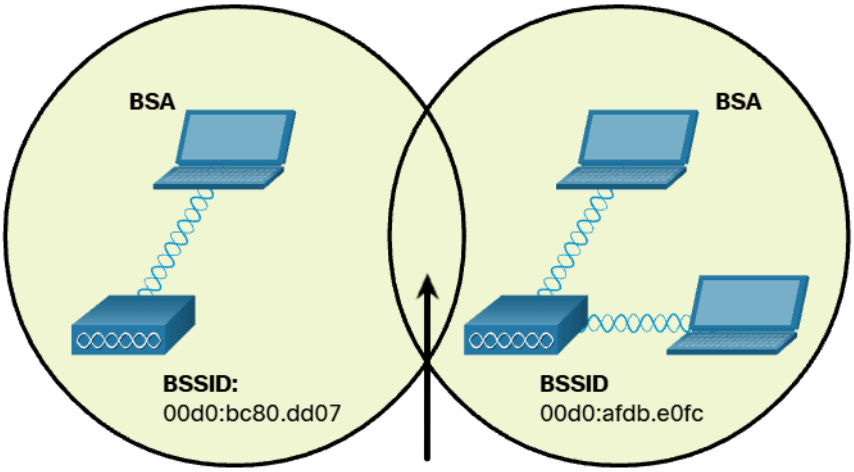

10 to 15% overlapping to provide roaming between the BSAs without loss of connectivity

Extended Service Set (ESS)

Two or more BSSs join through a common distribution system (DS) into an ES

- The union of two or more BSSs interconnected by a wired DS
- Each ESS is identified by an SSID and each BSS is identified by its BSSID

Wireless clients in one BSA can now communicate with wireless clients in another BSA within the same ESA

- Roaming mobile wireless clients may move from one BSA to another (within the same ESS) and seamlessly connect

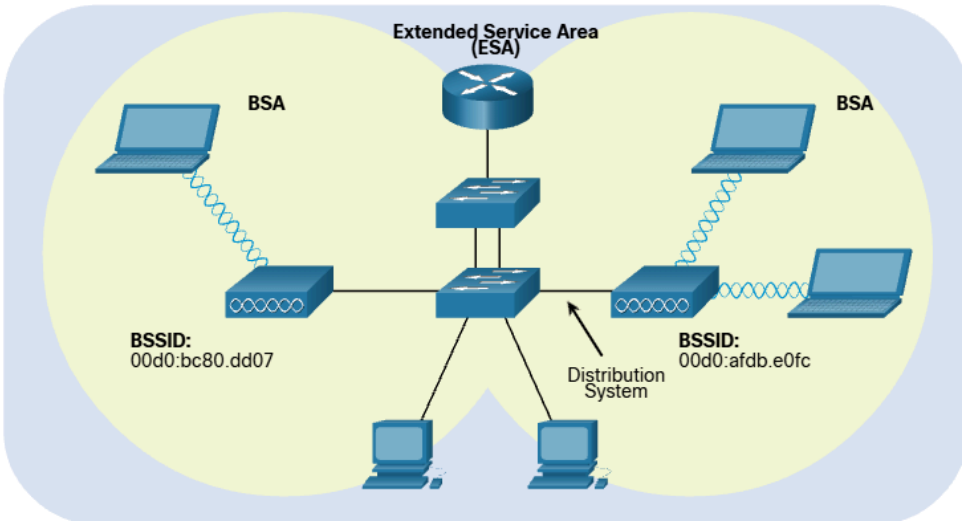

802.11 Frame Structure

- Layer 2 frames consist of a header, payload, and Frame Check Sequence (FCS)
- Similar to Ethernet frame format, except contains more fields

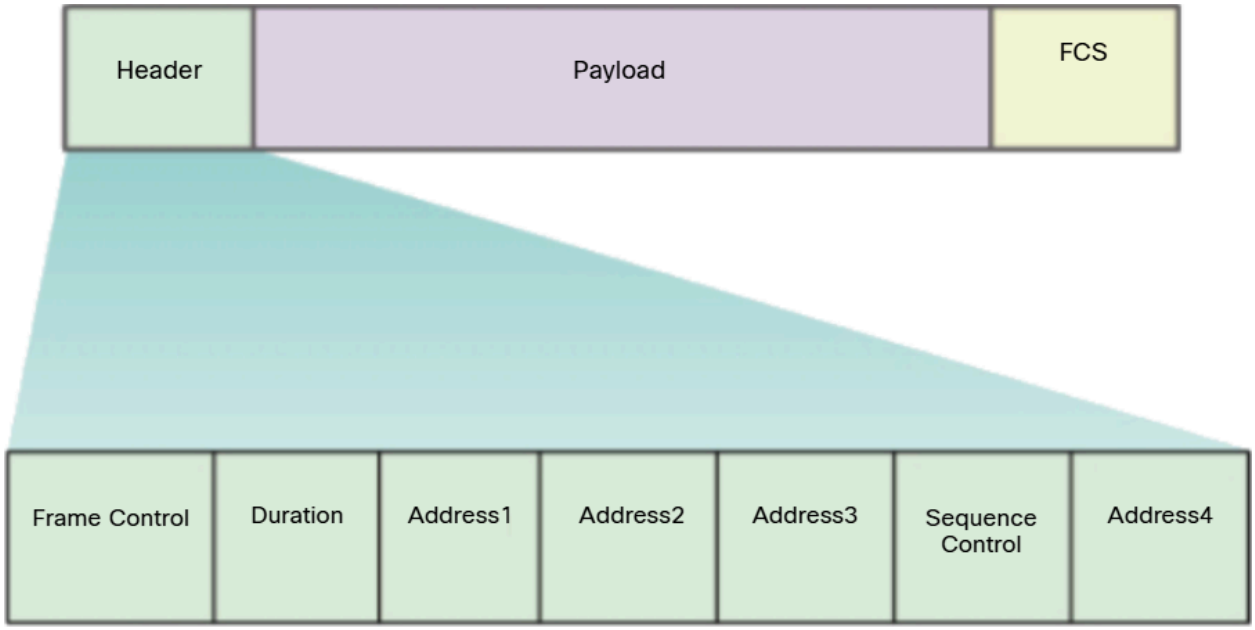

- **Frame Control** - Identifies the type of wireless frame and contains **four** subfields for protocol version, frame type, address type, power management, and security settings
- **Duration** - Typically used to indicate the remaining direction needed to receive the next frame transmission

From a wireless device:

- **Address 1 Receiver Address** - MAC address of the AP
- **Address 2 Transmitter Address** - MAC address of the sender
- **Address 3 SA/DA/BSSID** - MAC address of the destination which could be a wireless device or wired device

From the AP

- **Address 1 Receiver Address** - MAC address of the sender
- **Address 2 Transmitter Address** - MAC address of the AP
- **Address 3 SA/DA/BSSID** - MAC address of the wireless destination
- **Sequence Control** - Contains information to control sequencing and fragmented frames
- **Address4** - Usually missing because its used only in ad hoc mode
- **Payload** - Contains the data for transmission
- **FCS** - Used for Layer 2 error control

CSMA/CA

Note - WLANs are half-duplex, shared media configurations

- **Half-duplex**: only one client can transmit or receive at any given moment
- **Shared media**: wireless clients can all transmit and receive on the same radio channel
    - **Problem**: A wireless client cannot hear while it is sending, makes it impossible to detect a collision

Resolution - WLANs use carrier sense multiple access with collision avoidance (CSMA/CA) as method to determine how and when to send data on network

1. Listens to the channel to see if it is idle, which means that it senses no other traffic is currently on the channel. This channel is also called the carrier.
2. Sends a request to send (RTS) message to the AP to request dedicated access to the network
3. Receives a clear to send (CTS) message from the AP granting access to send
4. If the wireless client does not receive a CTS message, it waits a random amount of time before restarting the process
5. After it receives the CTS, it transmits the data
6. All transmissions are acknowledged. If a wireless client does not receive an acknowledgement, it assumes a collision occurred and restarts the process

Wireless Client and AP Association

To communicate, must first associate with an AP or wireless router

- Discover a wireless AP
- Authenticate with AP
- Associate with AP

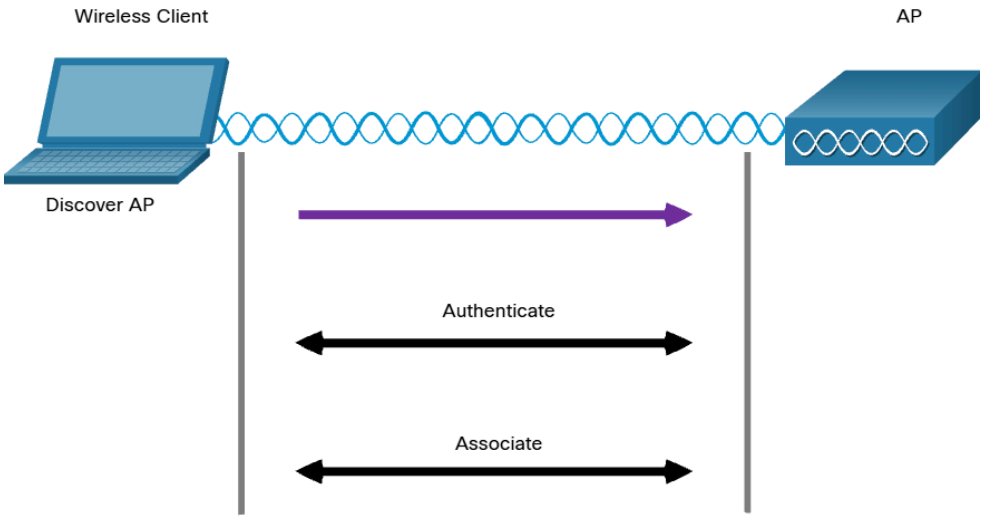

In order to have a successful association, a wireless client and AP must agree on specific parameters. Parameters must then be configured on the AP and subsequently on the client to enable negotiation of successful association

- **SSID** - The SSID name appears in the list of available wireless networks on a client.
    - In larger organizations that use multiple VLANs to segment traffic, each SSID is mapped to one VLAN. Depending on the network configuration, several APs on a network can share a common SSID
- **Password** - Required from the wireless client to authenticate to the AP
- **Network mode** - Refers to the 802.11a/b/g/n/ac/ad WLAN standards.
    - APs and wireless routers can operate in a Mixed mode meaning that they can simultaneously support clients connecting via multiple standards
- **Security mode** - Refers to the security parameter settings, such as WEP, WPA, or WPA2
    - Always enable the highest security level supported
- **Channel settings** - Refers to the frequency bands used to transmit wireless data
    - Wireless routers and APs can scan the radio frequency channels and automatically select an appropriate channel setting
    - The channel can also be set manually if there is interference with another AP or wireless device

Passive and Active Discover Mode

Passive mode

The AP openly advertises its service by periodically sending broadcast beacon frames containing the SSID, support standards, and security settings

- Primary purpose of beacon is to allow wireless clients to learn which networks and APs are available in a given area
- This allows wireless clients to choose which network and AP to use

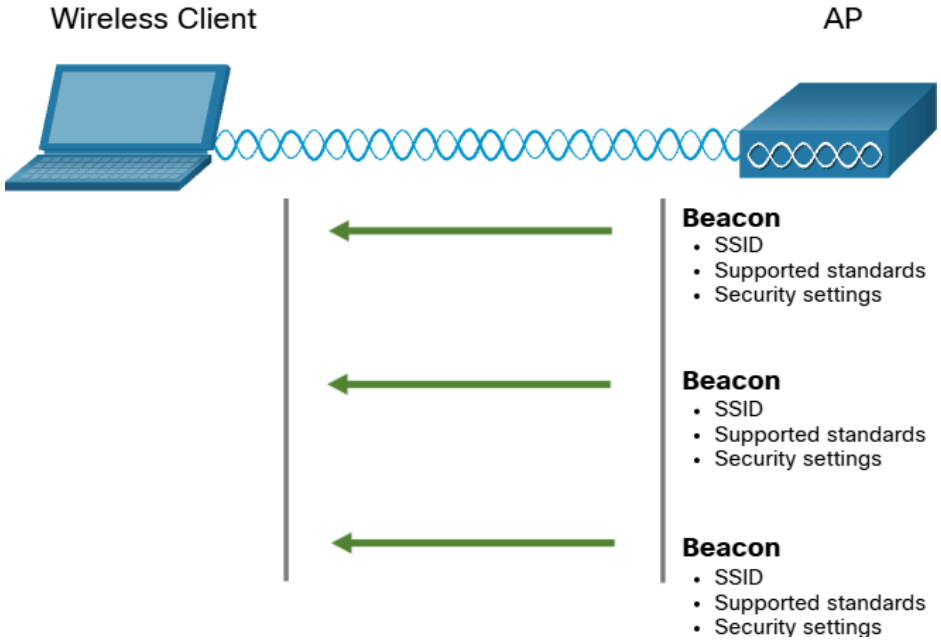

Active mode

Wireless clients must know the name of the SSID.

- Wireless clients initiates the process by broadcasting a probe request frame on multiple channels
    - The probe request includes the SSID name and standards supported
- APs configured with the SSID will send a probe response that includes the SSID, supported standards, and security settings
- Active mode may be required if an AP or wireless router is configured to not broadcast beacon frames
- Wireless client can also send a probe request without an SSID name to discover nearby WLAN networks
- APs configured to broadcast beacon frames would respond to the wireless client with a probe response and provide the SSID name
    - APs with the broadcast SSID feature disabled do not respond

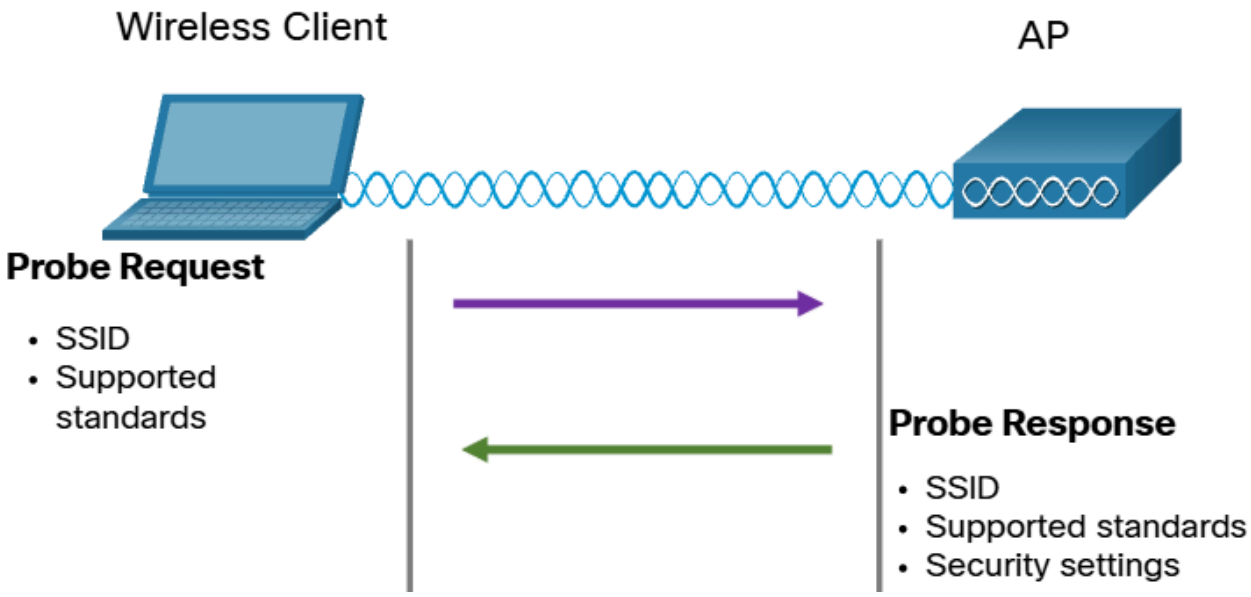

Introduction to CAPWAP

- IEEE standard protocol that enables a WLC to manage multiple APs and WLANs
- Responsible for the encapsulation and forwarding of WLAN client traffic between an AP and a WLC
- Based on LWAPP but adds additional security with Datagram Transport Layer Security (DTLS)
- Establishes tunnels on User Datagram Protocol (UDP) ports
- Can operate either over IPv4 or IPv6
    - Uses IPv4 by default
- IPv4 and IPv6 both use UDP port 5246 and 5247
    - Port 5246 is for CAPWAP control messages used by the WLC to manage the AP

- Port 5247 is used by CAPWAP to encapsulate data packets traveling to and from wireless clients
- CAPWAP tunnels use different IP protocols in the packet header
    - IPv4 uses IP protocol 17
    - IPv6 uses IP protocol 136

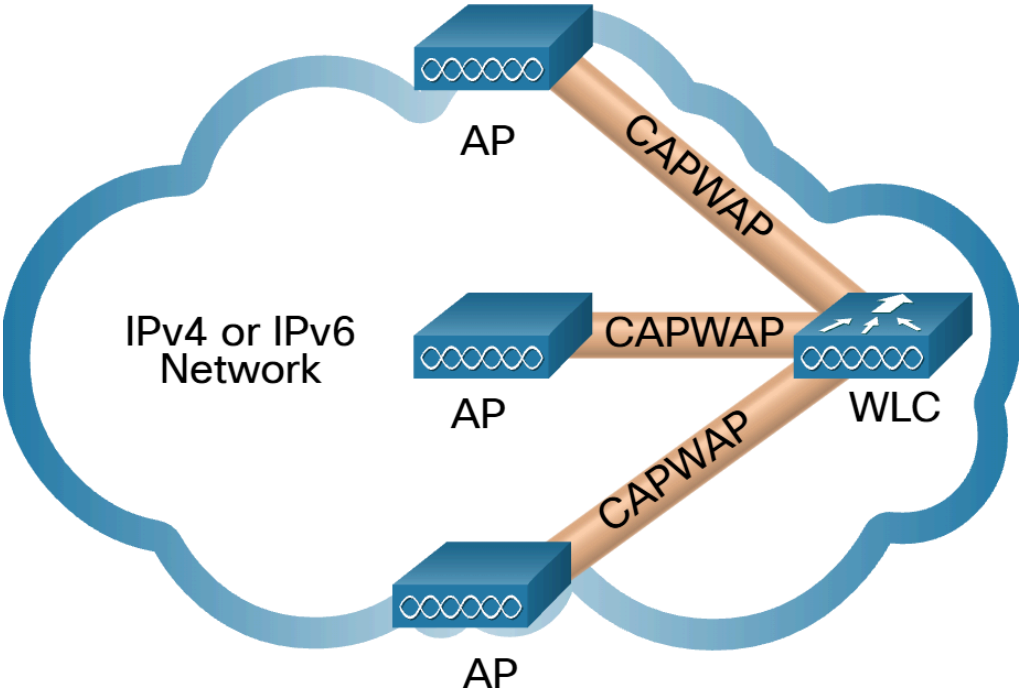

Split MAC Architecture

A key component of CAPWAP is the concept of split media access control (MAC). The CAPWAP split MAC concept does all the functions normally performed by individual APs and distributes them between two functional components

- AP MAC Functions
- WLC MAC Functions

AP MAC Functions	**WLC MAC Functions**
Beacons and probe responses	Authentication
Packet acknowledgements and retransmissions	Association and re-association of roaming clients
Frame queueing and packet prioritization	Frame translation to other protocols
MAC layer data encryption and decryption	Termination of 802.11 traffic on a wired interface

DTLS Encryption

DTLS is a protocol which provides security between the AP and the WLC

- Allows them to communicate using encryption and prevents eavesdropping or tampering
- Is enabled by default to secure the CAPWAP control channel but is disabled by default for the data channel

All CAPWAP management and control traffic exchanged between an AP and WLC is encrypted and secured by default to provide control plane privacy and prevent Man-in-the-Middle (MITM) attacks

CAPWAP data encryption is optional and is enabled per AP.

- Requires a DTLS license to be installed on the WLC prior to being enabled on an AP
- When enabled, all WLAN client traffic is encrypted at the AP before being forwarded to the WLC and vice versa

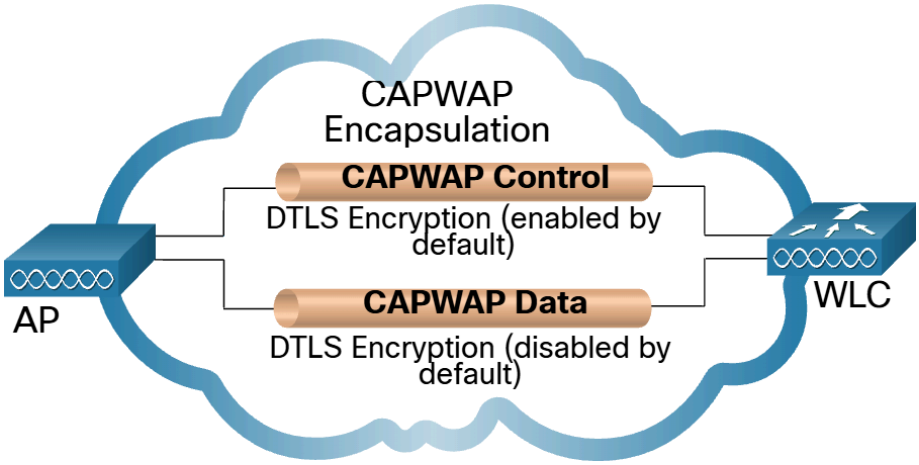

FlexConnect APs

- A wireless solution for branch office and remote office deployments
- Lets you configure and control access points in a branch office from the corporate office through a WAN link, without deploying a controller in each office

- **Connected mode** - The WLC is reachable
  - The flexConnect AP has CAPWAP connectivity with its WLC and can send traffic through the CAPWAP tunnel
  - The WLC performs all its CAPWAP functions
- **Standalone mode** - The WLC is unreachable
  - The flexConnect has lost or failed to establish CAPWAP connectivity with its WLC.
  - A FlexConnect AP can assume some of the WLC functions such as client data traffic locally and performing client authentication locally

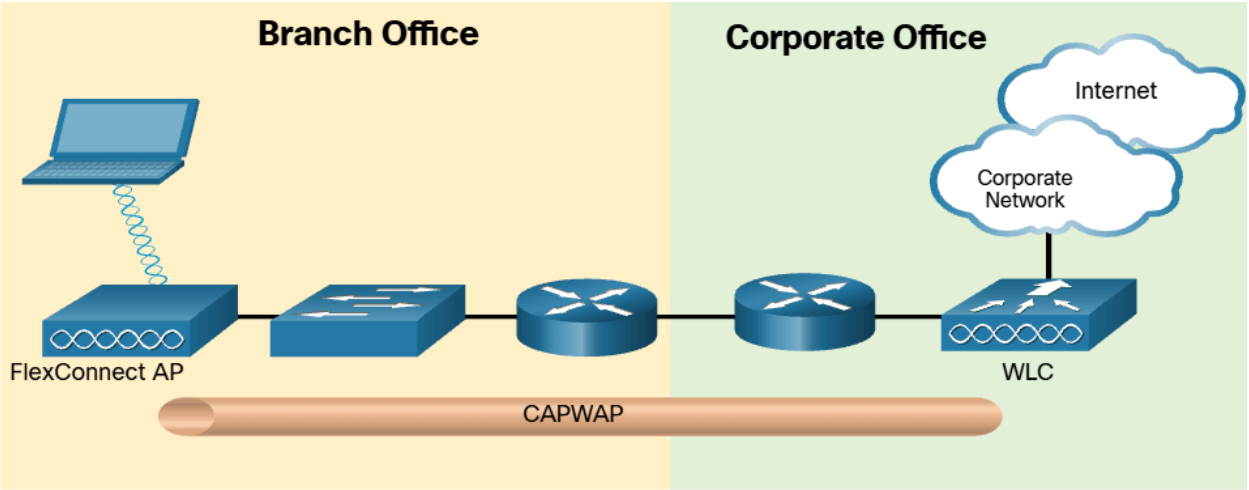

Channel Management

Frequency Channel Saturation

Wireless LAN devices have transmitters and receivers tuned to specific frequencies of radio waves to communicate

- Common practice is for frequencies to be allocated as ranges
- Such ranges are then split into smaller ranges called channels

If the demand for a specific channel is too high, that channel is likely to become oversaturated

- The saturation of the wireless medium degrades the quality of the communication

Techniques to improve wireless communication and alleviate saturation

Direct-Sequence Spread Spectrum (DSSS)

This is a modulation technique designed to spread a signal over a larger frequency band

- Spread spectrum techniques were developed during war time to make it more difficult for enemies to intercept or jam a communication signal
    - It does this by spreading the signal over a wider frequency which effectively hides the discernable peak of the signal
- A properly configured receiver can reverse the DSSS modulation and re-construct the original signal

DSSS is used by 802.11b devices to avoid interference from other devices using the same 2.4Ghz frequency

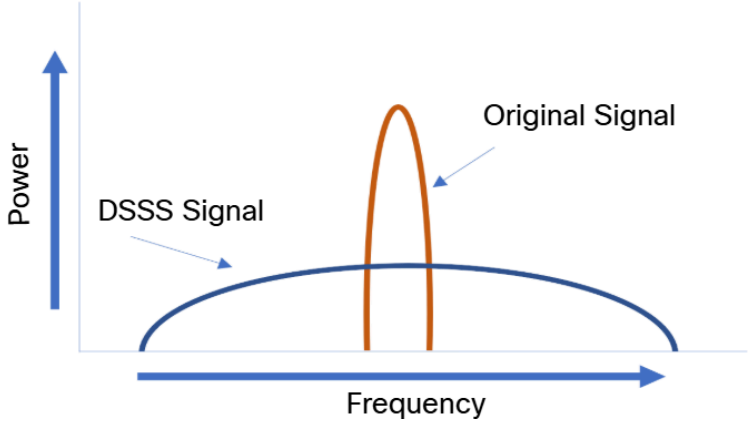

Frequency-Hopping Spread and Spectrum (FHSS)
This relies on spread spectrum methods to communicate.

- It transmits radio signals by rapidly switching a carrier signal among many frequency channels

The sender and receiver must be synchronized to “know” which channel to jump to

- This channel hopping process allows for a more efficient usage of the channels, decreasing channel congestion.

- FHSS was used by the original 802.11 standard.
- Walkie-talkies and 900 MHz cordless phones also use FHSS
    - Bluetooth uses a variation of FHSS

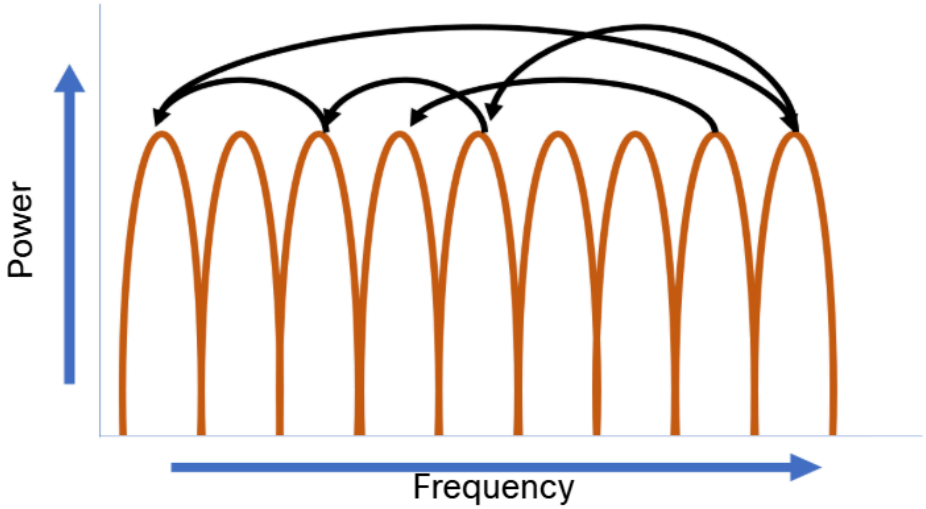

Orthogonal Frequency-Division Multiplexing (OFDM)
This is a subset of frequency division multiplexing in which a single channel uses multiple sub-channels on adjadcent frequencies

- Sub-channels in an OFDM system are precisely orthogonal to one another which allow the sub-channels to overlap without interference
- OFDM is used by a number of communication systems including 802.11a/g/n/ac

The new 802.11ax uses variation of OFM called Orthogonal frequency-division multiaccess (OFDMA)

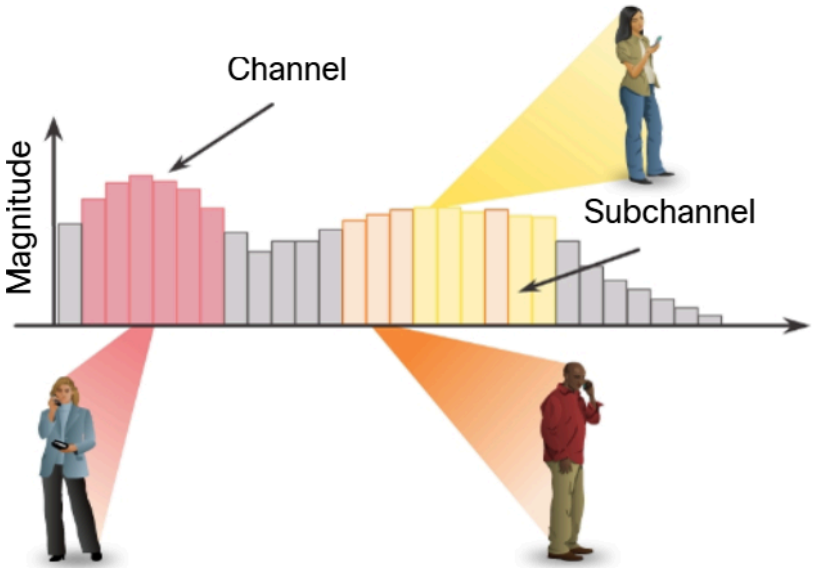

Channel Selection

Best practice for WLANs requiring multiple APs is to use non-overlapping channels

- Example: 802.11b/g/n standards operate in the 2.4 Ghz to 2.5 Ghz
    - The 2.4 Ghz band is subdivided into multiple channels
    - Each channel is allotted 22 Mhz bandwidth and is separated from the next channel by 5 Mhz

The 802.11b standard identifies 11 channels for North America

- 13 in Europe
- 14 in Japan

2.4 Ghz Overlapping Channels in North America

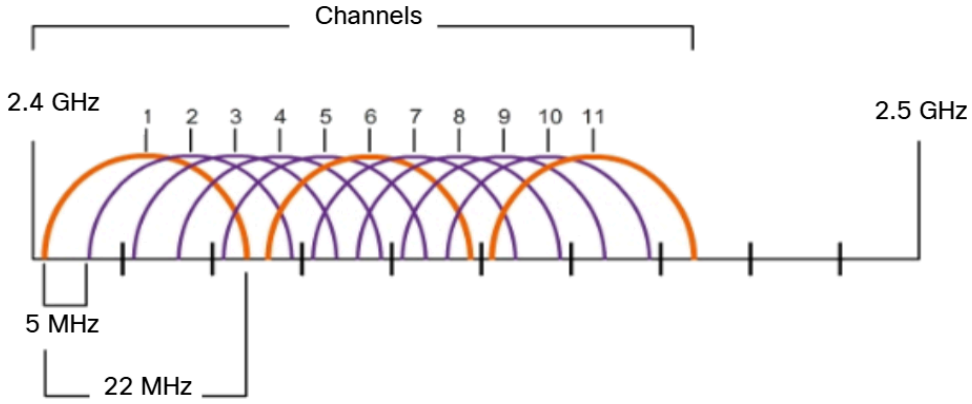

Interference occurs when one signal overlaps a channel reserved for another signal, causing possible distortion.

- Best practice for 24 Ghz WLANs that require multiple APs is to use non-overlapping channels, although modern APs will do this automatically
    - If there are three adjacent APs, use channels 1, 6, and 11

2.4 Ghz Non-Overlapping Channels for 802.11b/g/n

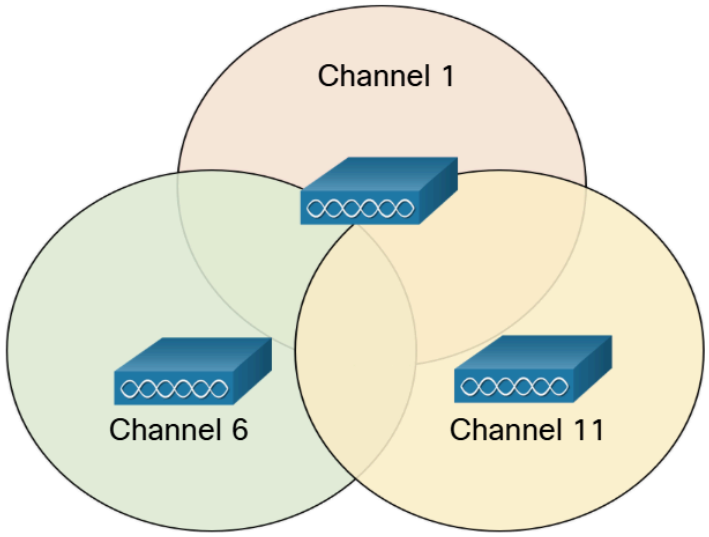

For the 5 Ghz standards 802.11a/n/ac, there are 24 channels

- Divided into three sections
- Each channel is separated from the next channel by 20 Mhz

Figure below: 24 Unlicensed National Information Infrastructure (U-NNI) 24 channels for the 5 Ghz band

- Although slight overlap at tails of each channel's frequency, the channels do not interfere with one another.

5 Ghz wireless can provide faster data transmission for wireless clients in heavily populated wireless networks because of the large amount of non-overlapping wireless channels

5 Ghz First Eight Non-Interfering Channels

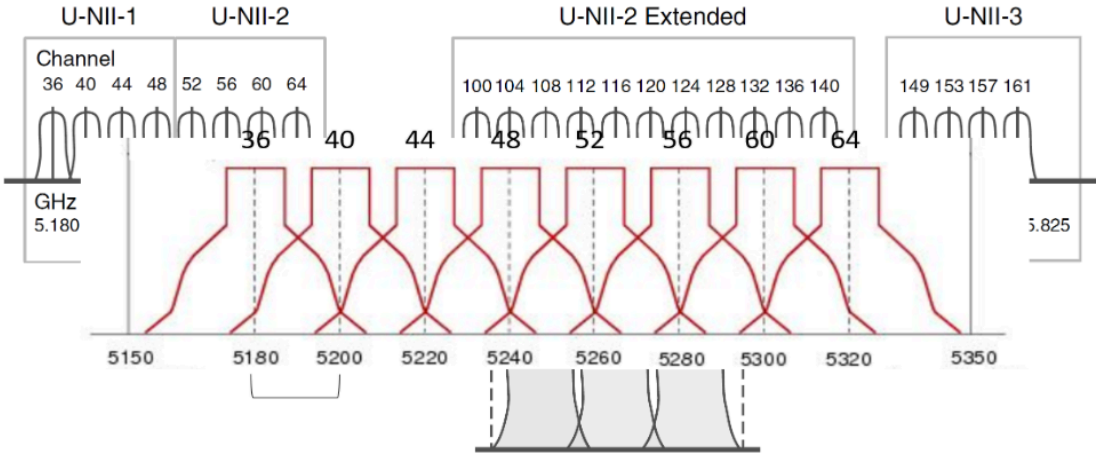

As with 2.4 Ghz WLANs, choose non-interfering channels when configuring multiple 5 Ghz APs that are adjacent to each other

5 Ghz Non-Interfering Channels for 802.11a/n/ac

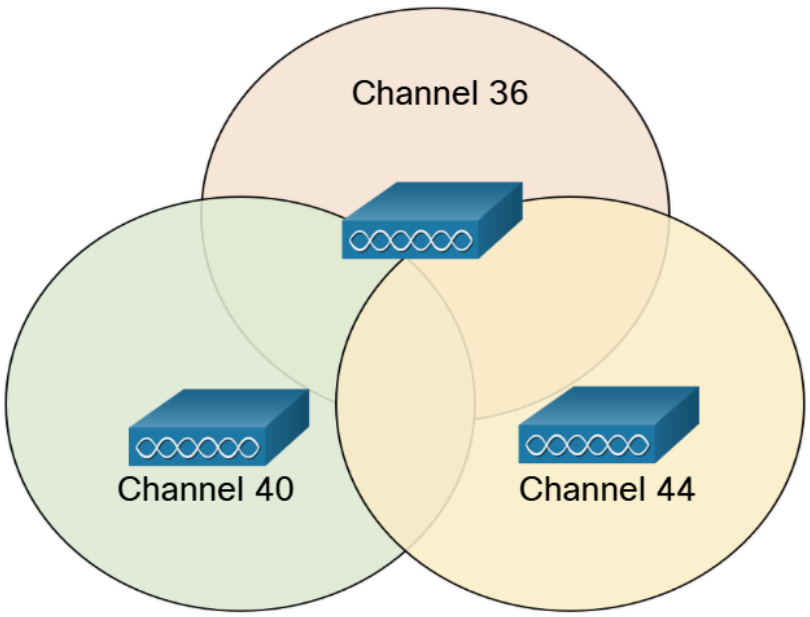

Plan a WLAN Deployment

The number of users supported by a WLAN depends on the geographical layout of the facility, including the number of bodies and devices that can fit in a space, the data rates users expect, the use of non-overlapping channels by multiple APs in an ESS, and transmit power settings.

- When planning the location of APs, the approximate circular coverage area is important
- If APs are to use existing wiring or if there are locations where APs cannot be placed, not these locations on the map
- Note all potential sources of interference which can include microwave ovens, wireless video cameras, fluorescent lights, motion detectors, or any other device that uses the 2.4 Ghz range
- Position APs above obstructions
- Position APs vertically near the ceiling in the center of each coverage area, if possible
- Position APs in locations where users are expected to be
    - Example: conference rooms are typically a better location for APs than a hallway
- If an IEEE 802.1 network has been configured for mixed mode, the wireless clients may experience slower than normal speeds in order to support the older wireless standards

When estimating the expected coverage area of an AP, realize that this value varies depending on the WLAN standard or mix of standards that are deployed, the nature of the facility, and the transmit power that the AP is configured for. Always consult the specifications for the AP when planning for coverage areas

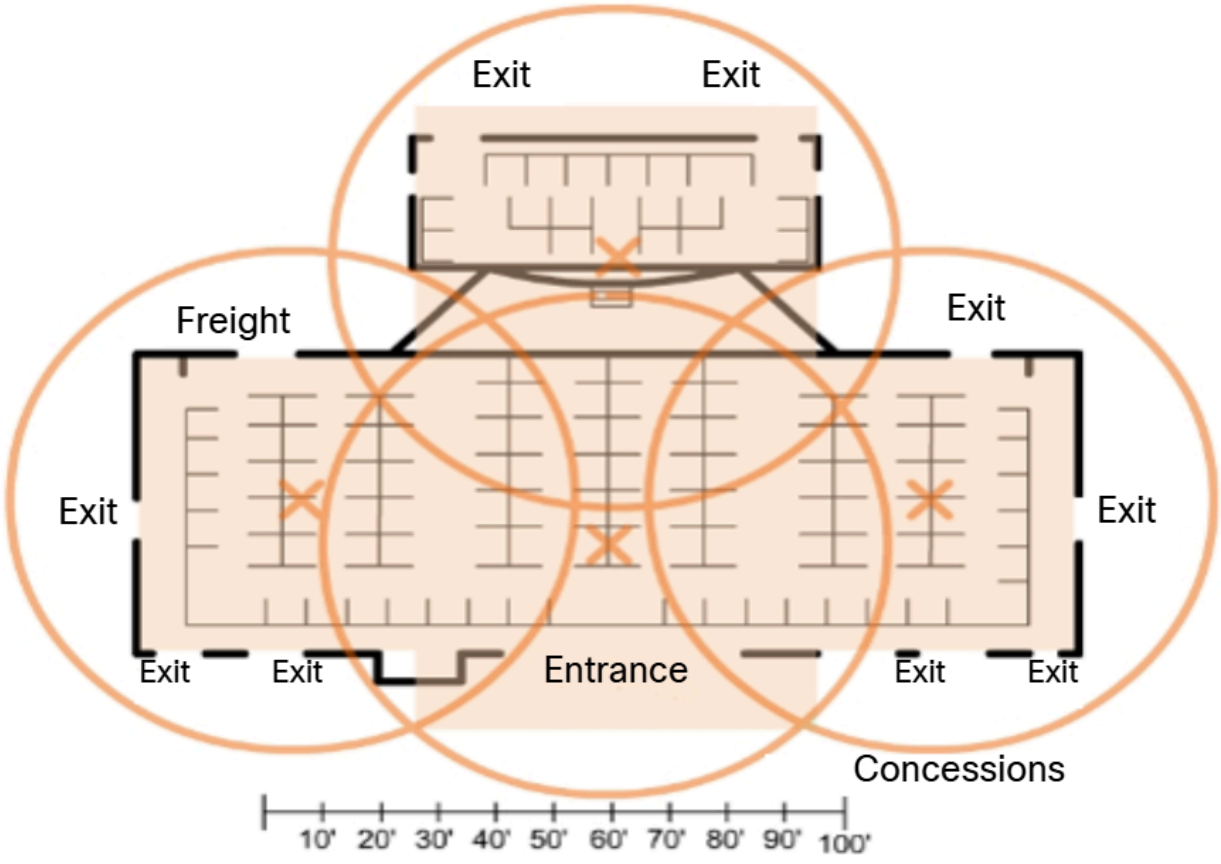

WLAN Threats

Attacks can be generated by outsiders, disgruntled employees, and even unintentionally by employees

- **Interception of data** - Wireless data should be encrypted to prevent it from being read by eavesdroppers
- **Wireless intruders** - Unauthorized users attempting to access network resources can be deterred through effective authentication techniques
- **Denial of Service (DoS) Attacks** - Access to WLAN services can be compromised either accidentally or maliciously. Various solutions exist depending on the source of the DoS attack
- **Rogue APs** - Unauthorized APs installed by a well-intentioned user or for malicious purposes can be detected using management software

DoS Attacks

Wireless DoS attacks can be the result of:

- **Improperly configured devices** - Configuration errors can disable the WLAN
    - Example: An administrator could accidentally alter a configuration and disable the network, or an intruder with administrator privileges could intentionally disable a WLAN
- **A malicious user intentionally interfering with wireless communication** - Their goal is to disable the wireless network completely or to the point where no legitimate device can access the medium
- **Accidental interference** - WLANs are prone to interference from other wireless devices including microwave ovens, cordless phones, baby monitors, etc
    - 2.4 Ghz band is more prone to interference than 5 Ghz band

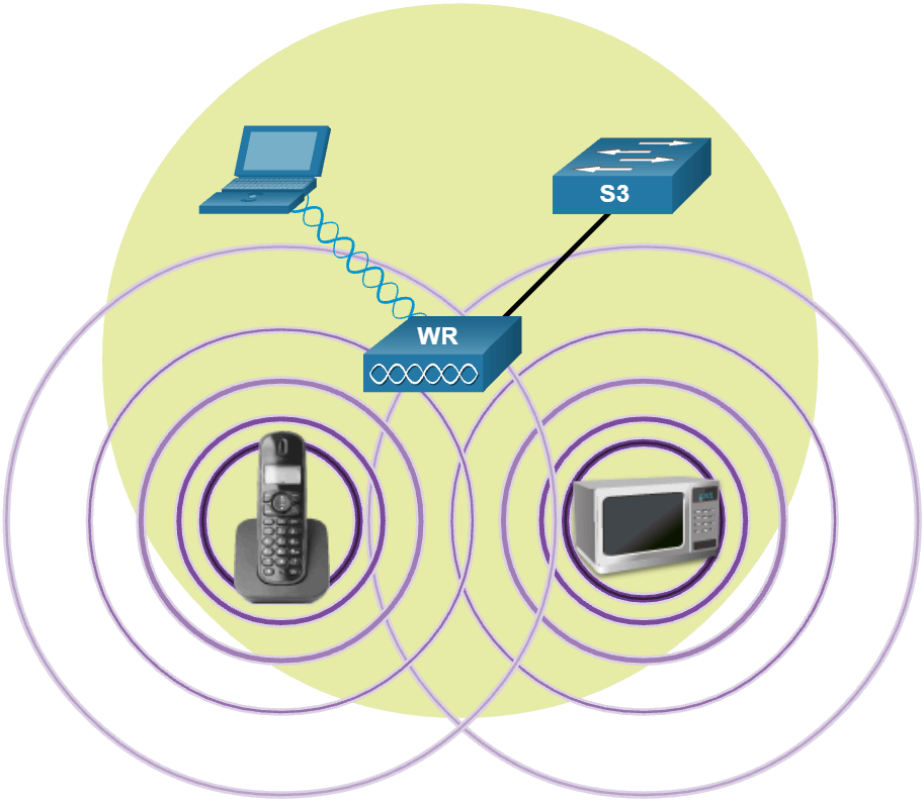

To minimize the risk of a DoS attack due to improperly configured devices and malicious attacks, harden all devices, keep passwords secure, create backups, and ensure that all configuration changes are incorporated off-hours.

Monitor the WLAN for any accidental interference problems and address them as they appear. Because the 2.4 Ghz band is used by other devices types, the 5 Ghz should be used in areas prone to interference

Rogue Access Points

- An AP or wireless router that has been connected to a corporate network without explicit authorization and against corporate policy
    - Anyone with access to the premises can install (maliciously or non-maliciously) an inexpensive wireless router that can potentially allow access to a secure network resource.
- Once connected, the rogue AP can be used by an attacker to capture MAC addresses, capture data packets, gain access to network resources, or launch a man-in-the-middle attack

A personal network hotspot could also be used as a rogue AP

- Example: A user with secure network access enables their authorized Windows host to become a Wi-Fi AP
    - Doing so circumvents the security measures the other unauthorized devices can now access network resources as a shared device

To prevent the installation of rogue APs, organizations must configure WLCs with rogue AP policies, and use monitoring software to actively monitor the radio spectrum for unauthorized APs

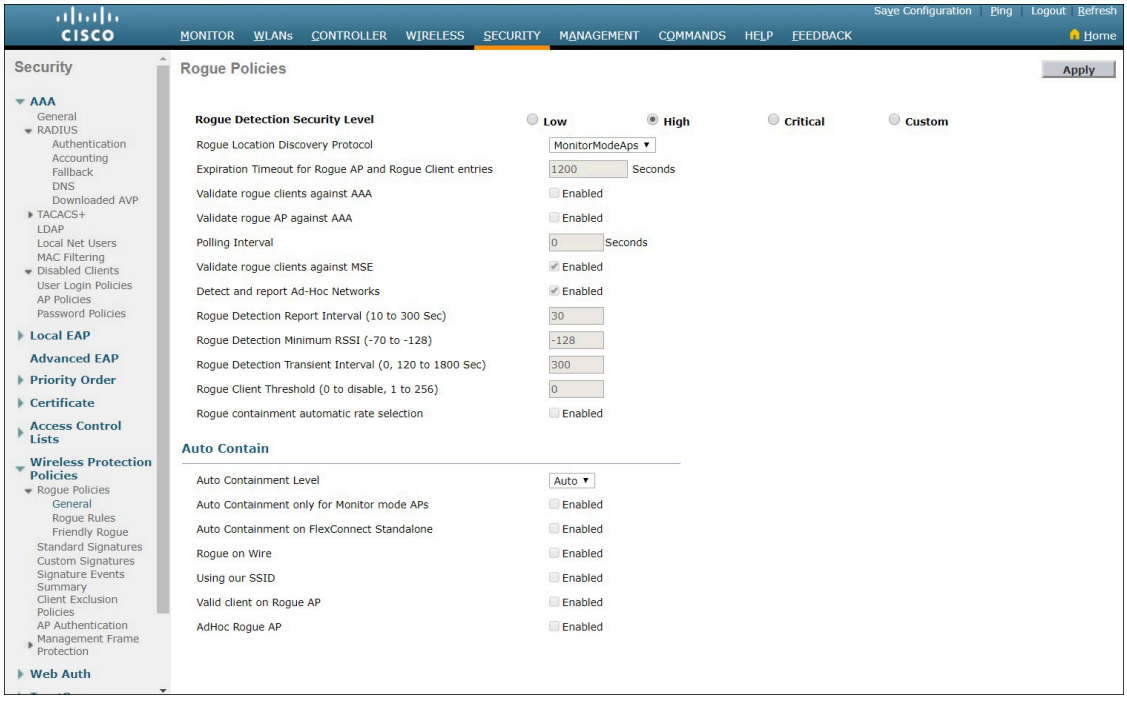

Man-in-the-Middle Attack

- The hacker is positioned in between two legitimate entities in order to read or modify the data that passes between the two parties
- A popular wireless MITM attack is called “evil twin AP”
    - An attacker introduces a rogue AP and configures it with the same SSID as a legitimate AP

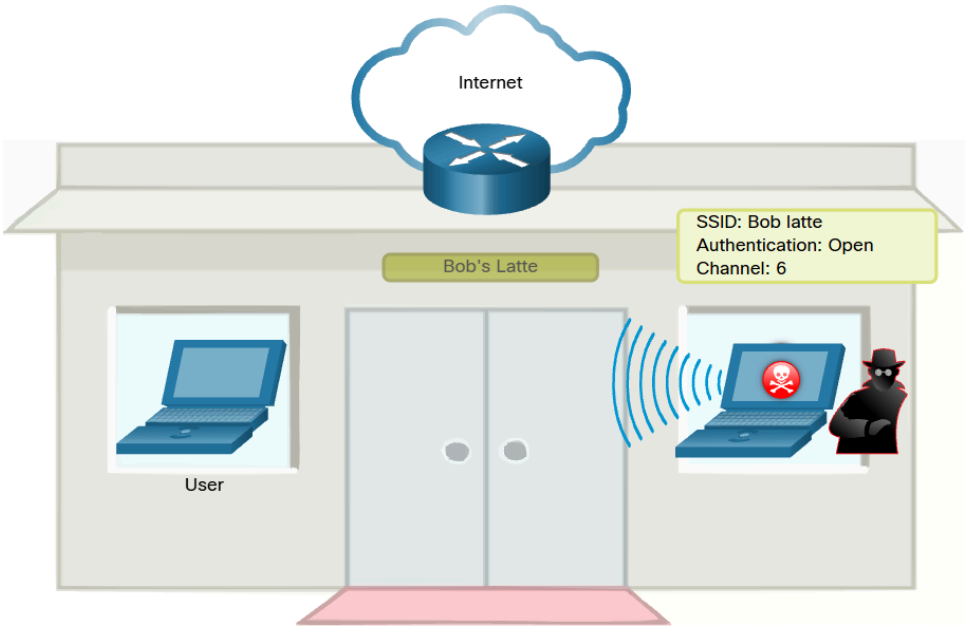

Wireless clients attempting to connect to a WLAN would see two APs with the same SSID offering wireless access

- Those near the rogue AP find the stronger signal and most likely associate with it
- User traffic is now sent to the rogue AP, which in turn captures the data and forwards it to the legitimate AP
- Return traffic from the legitimate AP is sent to the rogue AP, captures, and then forwarded to the unsuspecting user
- The attacker can steal the user’s passwords, personal information, gain access to their device, and compromise the system

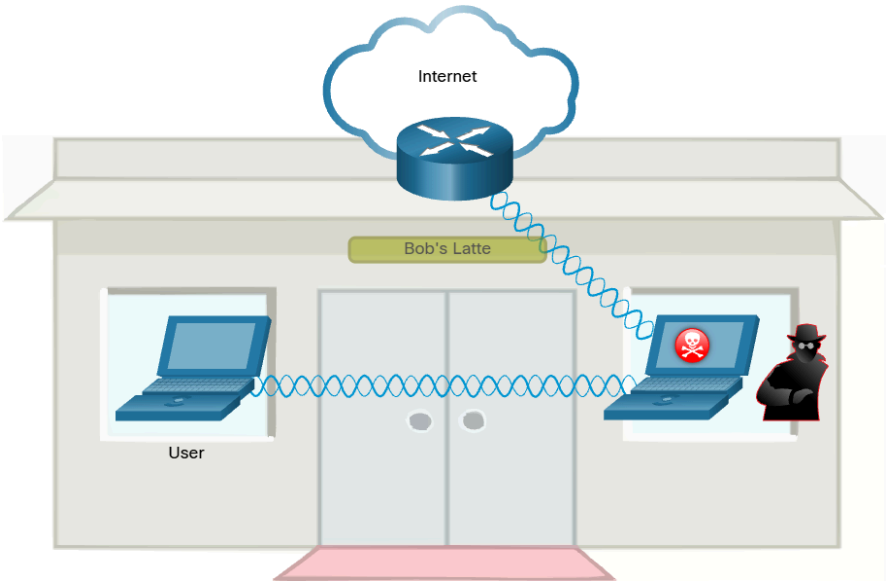

Defeating an attack like an MITM attack depends on the sophistication of the WLAN infrastructure and the vigilance in monitoring activity on the network

- The process begins with identifying legitimate devices on the WLAN
- To do this, users must be authenticated
- After all of the legitimate devices are known, the network can be monitors for abnormal devices or traffic

Secure WLANs

SSID Cloaking and MAC Address Filtering

SSID Cloaking

APs and some wireless routers allow the SSID beacon frame to be disabled. Wireless clients must manually configure the SSID to connect to the network.

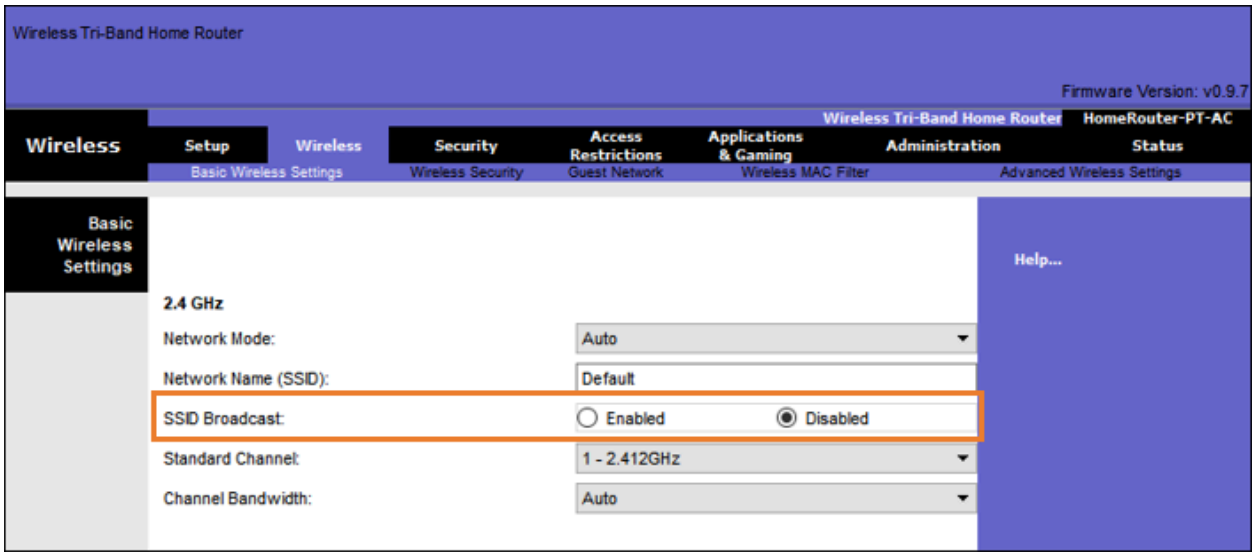

MAC Address Filtering

An administrator can manually permit or deny clients wireless access based on their physical MAC hardware address

- In figure below, the router is configured to permit two MAC addresses
- Devices with different MAC addresses will not be able to join the 2.4Ghz WLAN

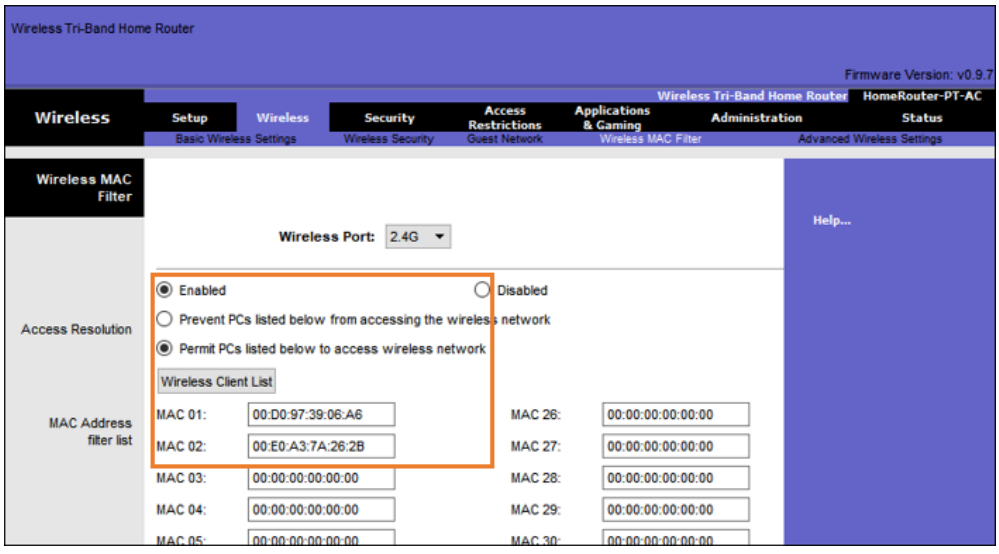

802.11 Original Authentication Methods

Neither SSID clocking nor MAC address filtering is truly secured.

- SSIDs are easily discovered even if APs do not broadcast them
- MAC addresses can be spoofed

Two types of authentication were introduced with original 802.11 standard:

- **Open system authentication** - Any wireless client should easily be able to connect and should only be used in situations where security is of no concern, such as those providing free internet access like cafes, hotels, and remote areas.
    - The wireless client is responsible for providing security such as using a virtual private network (VPN) to connect securely. VPNs provide authentication and encryption services

- **Shared key authentication** - Provides mechanisms, such as WEP, WPA, WPA2, and WPA3 to authenticate and encrypt data between a wireless client and AP. However, the password must be pre-shared between both parties to connect.

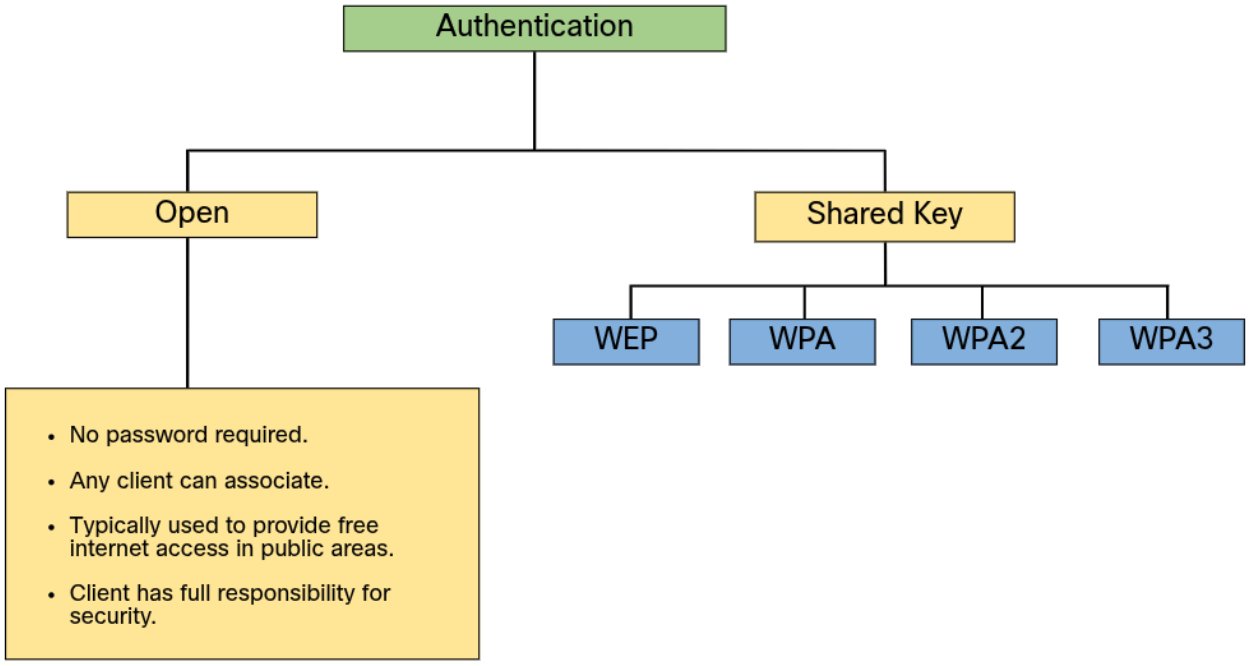

Shared Key Authentication Methods

Authentication Method	Description
Wired Equivalent Privacy (WEP)	The original 802.11 specification designed to secure the data using the Rivest Cipher 4 (RC4) encryption method with a static key. However, the key never changes when exchanging packets. This makes it easy to hack. WEP is no longer recommended and should never be used
Wi-Fi Protected Access (WPA)	A Wi-Fi Alliance standard that uses WEP, but secures the data with the much stronger Temporal Key Integrity Protocol (TKIP) encryption algorithm. TKIP changes the key for each packet, making it much more difficult to hack
WPA2	The current industry standard for securing wireless networks. It uses the Advanced Encryption Standard (AES) for encryption. AES is currently considered the strongest encryption protocol
WPA3	The next generation of Wi_fi security. All WPA3-enabled devices use the latest security methods, disallow outdated legacy protocols, and require the use of Protected Management Frames (PMF)

Authenticating a Home User

Two WPA2 authentication methods:

- **Personal** - Intended for home and small office networks, users authenticate using a pre-shared key (PSK). Wireless clients authenticate with the wireless router using a pre-shared password. No special authentication server is required.
- **Enterprise** - Intended for enterprise networks but requires a Remote Authentication Dial-In User Service (RADIUS) authentication server. Although more complicated to set up, it provides additional security. The device must be authenticated by the RADIUS server and then users must authenticate using 802.1X standard, which uses the Extensible Authentication Protocol (EAP) for authentication

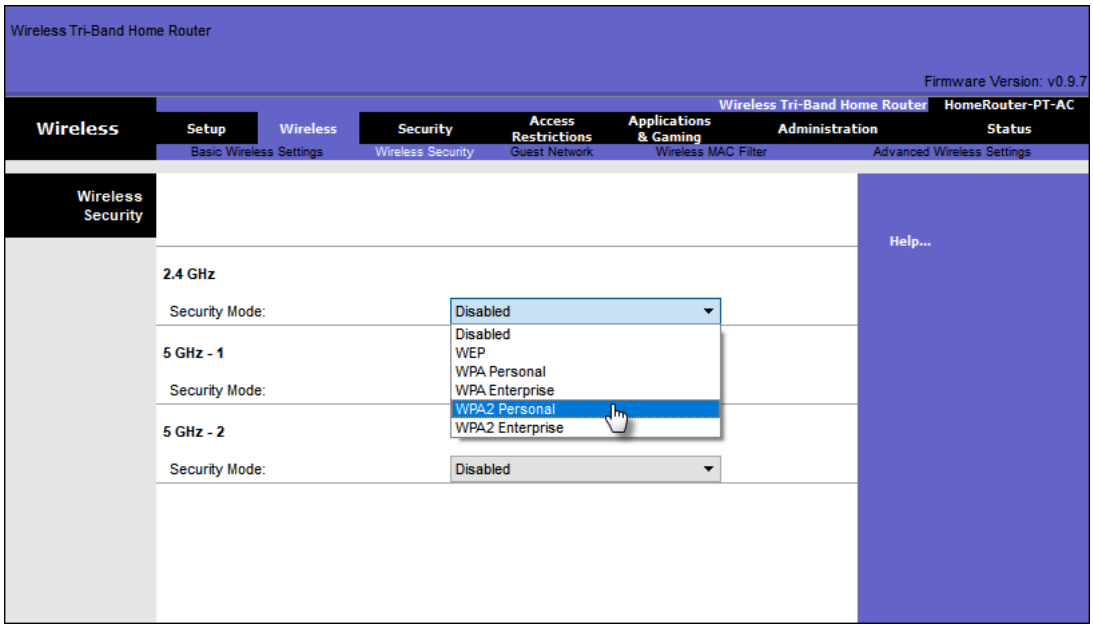

Encryption Methods

WPA and WPA2 standards use the following encryption protocols:

- **Temporal Key Integrity Protocol (TKIP)** - The encryption method used by WPA. It provides support for legacy WLAN equipment by addressing the original flaws associated with the 802.11 WEP encryption method. It makes use of WEP, but encrypts the Layer 2 payload using TKIP, and carries out a Message Integrity Check (MIC) in the encrypted packet to ensure the message has been altered.
- **Advanced Encryption Standard (AES)** - AES is the encryption method used by WPA2. It is the preferred method because it is a far stronger method of encryption. It uses the Counter Cipher Mode with Blocking Chaining Message Authentication Code Protocol (CCMP) that allows destination hosts to recognize if the encrypted and non-encrypted bits have been altered

Figure below: The administrator is configuring the wireless router to use WPA2 with AES encryption on the 2.4 Ghz band

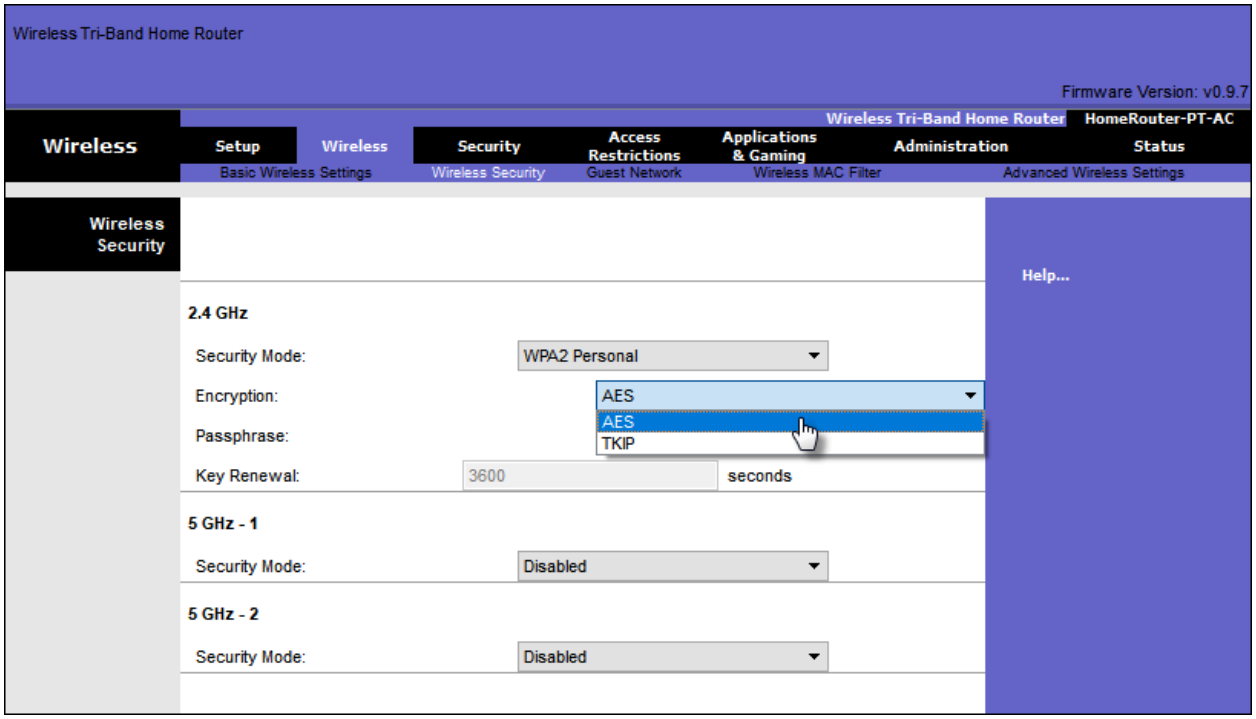

Authentication in the Enterprise

In networks with stricter security requirements, additional authentication or login is required to grant wireless clients access. The Enterprise security mode choice requires an Authentication, Authorization, and Accounting (AAA) RADIUS server

- **RADIUS Server IP Address** - This is the reachable address of the RADIUS server
- **UDP port numbers** - Officially assigned:
    - UDP port 1812 for RADIUS Authentication
        - Also UDP port 1645
    - UDP port 1813 for RADIUS Accounting
        - Also UDP port 1646

Figure below: The administrator is configuring the wireless router with WPA2 Enterprise authentication using AES encryption. The RADIUS server IPv4 address is configured as well with a strong password to be used between the wireless router and the RADIUS server

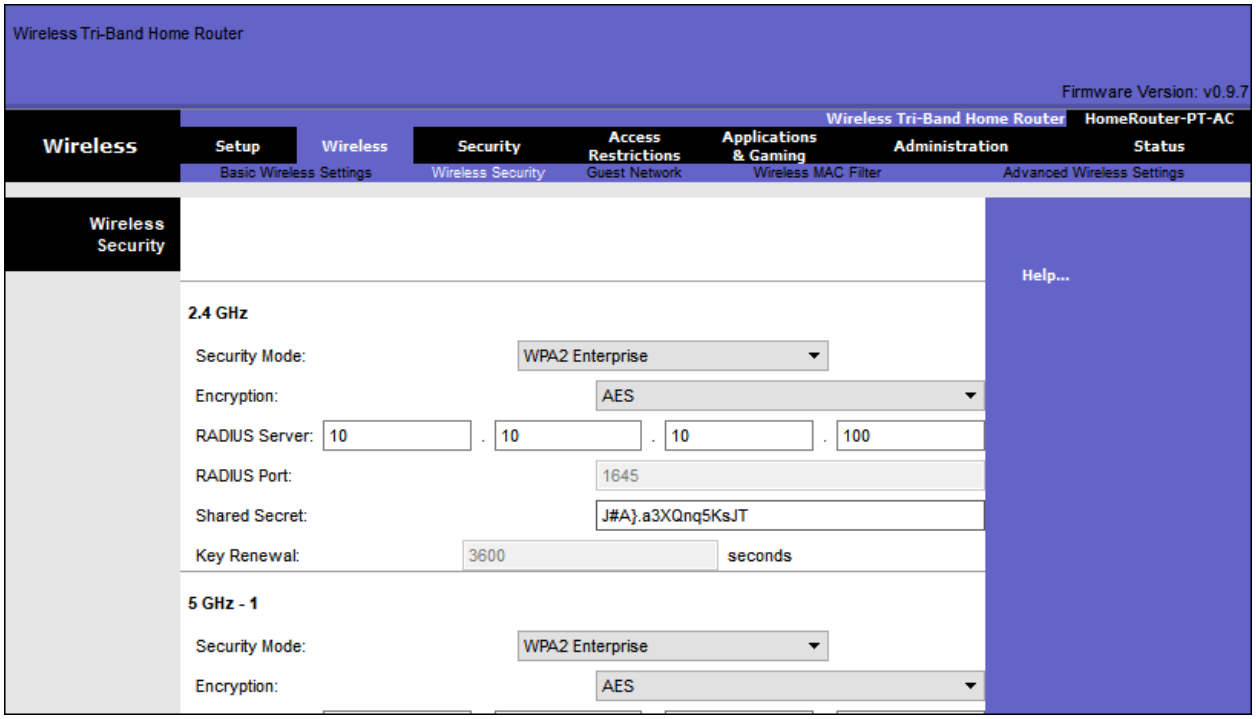

The shared key is not a parameter that must be configured on a wireless client.

- It is only required on the AP to authenticate with the RADIUS server

User authentication and authorization is handled by the 802.1X standard, which provides a centralized, server-based authentication of end users.

The 802.1X login process uses EAP to communicate with the AP and RADIUS server

- EAP is a framework for authenticating network access
    - It provides a secure authentication mechanism and negotiate a secure private key which can then be used for a wireless encryption session using TKIP or AES encryption

WPA3

Features:

- WPA3-Personal
- WPA3-Enterprise
- Open Networks
- Internet of Things (IoT) Onboarding

WPA3-Personal

In WPA2-Personal, threat actors can listen in on the “handshake” between a wireless client and the AP and use a brute force attack to try and guess the PSK. WPA3-Personal thwarts this attack by using Simultaneous Authentication of Equals (SAW), a feature specified in the IEEE 802.11-2016. THe PSK is never exposed, making it impossible for the threat actor to guess.

WPA3-Enterprise

Still uses 802.1X/EAP authentication. However, it requires the use of a 192-bit cryptographic suite and eliminates the mixing of security protocols for previous 802.11 standards.
WPA3-Enterprise adheres to the Commercial National Security Algorithm (CNSA) Suite which is commonly used in high security Wi-Fi networks

Open Networks

Open networks in WPA2 send user traffic to unauthenticated, clear text. In WPA3, open or public Wi-Fi networks still do not use any authentication. However, they do use Opportunistic Wireless Encryption (OWE) to encrypt all wireless traffic

IoT Onboarding

Although WPA2 included Wi-Fi Protected Setup (WPS) to quickly onboard devices without configuring them first, WPS is vulnerable to a variety of attacks and is not recommended. Furthermore, IoT devices are typically headless, meaning they have no built-in GUI for configuration, and needed any easy way to get connected to the wireless network.

- The Device Provisioning Protocol (DPP) was designed to address this need.

The key is typically stamped on the outside of the device or its packaging as a Quick Response (QR) code. The network administrator can scan the QR code and quickly onboard the device. Although not strictly part of the WPA3 standard, DPP will replace WPS over time

12.8.1 What did I learn in this module?

A Wireless LAN (WLAN) is a type of wireless network that is commonly used in homes, offices, and campus environments. Wireless networks are based on IEEE standards and can be classified into four main types: WPAN, WLAN, WMAN, and WWAN. Wireless LAN technologies uses the unlicensed radio spectrum to send and receive data. Examples of this technology are Bluetooth, WiMAX, Cellular Broadband, and Satellite Broadband. The IEEE 802.11 WLAN standards define how radio frequencies are used for wireless links. WLAN networks operate in the 2.4 GHz frequency band and the 5 GHz band. Standards ensure interoperability between devices that are made by different manufacturers. Internationally, the three organizations influencing WLAN standards are the ITU-R, the IEEE, and the Wi-Fi Alliance.

To communicate wirelessly, most devices include integrated wireless NICs that incorporate a radio transmitter/receiver. The wireless router serves as an access point, a switch, and a router. Wireless clients use their wireless NIC to discover nearby APs advertising their SSID. Clients then attempt to associate and authenticate with an AP. After being authenticated, wireless users have access to network resources. APs can be categorized as either autonomous APs or controller-based APs. There are three types of antennas for business class APs: omnidirectional, directional, and MIMO.

The 802.11 standard identifies two main wireless topology modes: Ad hoc mode and Infrastructure mode. Tethering is used to provide quick wireless access. Infrastructure mode defines two topology building blocks: A Basic Service Set (BSS) and an Extended Service Set (ESS). All 802.11 wireless frames contain the following fields: frame control, duration, address 1, address 2, address 3, sequence control, address 4, payload, and FCS. WLANs use CSMA/CA as the method to determine how and when to send data on the network. Part of the 802.11 process is discovering a WLAN and subsequently connecting to it. Wireless devices discover a wireless AP, authenticate with it, and then associate with it. Wireless clients connect to the AP using a scanning process which may be passive or active.

CAPWAP is an IEEE standard protocol that enables a WLC to manage multiple APs and WLANs. The CAPWAP split MAC concept does all of the functions normally performed by individual APs and distributes them between two functional components: AP MAC functions and WLC MAC functions. DTLS is a protocol which provides security between the AP and the WLC. FlexConnect is a wireless solution for branch office and remote office deployments. You configure and control access points in a branch office from the corporate office through a WAN link, without deploying a controller in each office. There are two modes of operation for the FlexConnect AP: connected and standalone.

Wireless LAN devices have transmitters and receivers tuned to specific frequencies of radio waves to communicate. Frequencies are allocated as ranges. Ranges are then split into smaller ranges called channels: DSSS, FHSS, and OFDM. The 802.11b/g/n standards operate in the 2.4 GHz to 2.5GHz spectrum. The 2.4 GHz band is subdivided into multiple channels. Each channel is allotted 22 MHz bandwidth and is separated from the next channel by 5 MHz. When planning the location of APs, the approximate circular coverage area is important.

Wireless networks are susceptible to threats, including: data interception, wireless intruders, DoS attacks, and rogue APs. Wireless DoS attacks can be the result of: improperly configured devices, a malicious user intentionally interfering with the wireless communication, and accidental interference. A rogue AP is an AP or wireless router that has been connected to a corporate network without explicit authorization. When connected, a threat actor can use the rogue AP to capture MAC addresses, capture data packets, gain access to network resources, or launch a MITM attack. In a MITM attack, the threat actor is positioned in between two legitimate entities to read or modify the data that passes between the two parties. A popular wireless MITM attack is called the “evil twin AP” attack, where a threat actor introduces a rogue AP and configures it with the same SSID as a legitimate AP. To prevent the installation of rogue APs, organizations must configure WLCs with rogue AP policies.

To keep wireless intruders out and protect data, two early security features are still available on most routers and APs: SSID cloaking and MAC address filtering. There are four shared key authentication techniques available: WEP, WPA, WPA2, and WPA3 (Devices with WPA3 are not yet readily available). Home routers typically have two choices for authentication: WPA and WPA2. WPA2 is the stronger of the two. Encryption is used to protect data. The WPA and WPA2 standards use the following encryption protocols: TKIP and AES. In networks that have stricter security requirements, an additional authentication or login is required to grant wireless clients access. The Enterprise security mode choice requires an Authentication, Authorization, and Accounting (AAA) RADIUS server.

WLAN Configuration

The wireless Router

Sometimes called an integrated router

- Include a switch for wired clients
- A port for an internet connection (sometimes labeled “WAN”)
- Wireless components for wireless client access

Cisco Meraki MX64W

Figure displays topology depicting the physical connection of a wired laptop to wireless router, which is then connected to a cable or DSL modem for internet

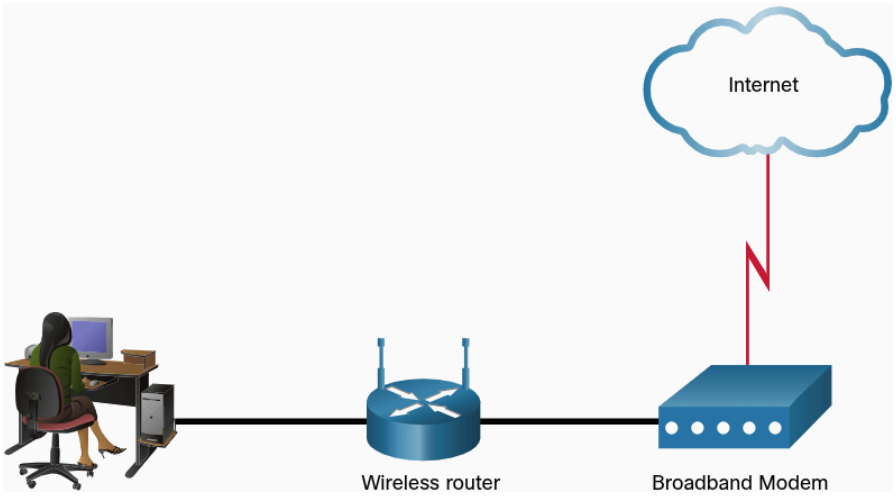

These wireless routers typically provide WLAN security, DHCP services, integrated Name Address Translation (NAT), quality of service (QoS), etc. Feature set will vary based on router model

Note - Cable or DSL modem configuration is usually done by the service provider’s representative either on-site or remotely through a walkthrough with you on the phone.

Log in to the Wireless Router

- Ready OOBE
- Should configure accordingly

To gain access to WebUI

- Type in router's default IP address found in documentation or search internet into URL bar of browser
- Log in using documentation's admin default username and password

Basic Network Setup

1. Log in to the router from web browser
2. Change the default administrative password
3. Log in with the new administrative password
4. Change the default DHCP IPv4 addresses
5. Renew the IP address
6. Log in to the router with the new IP address

Log in to the router from a web browser

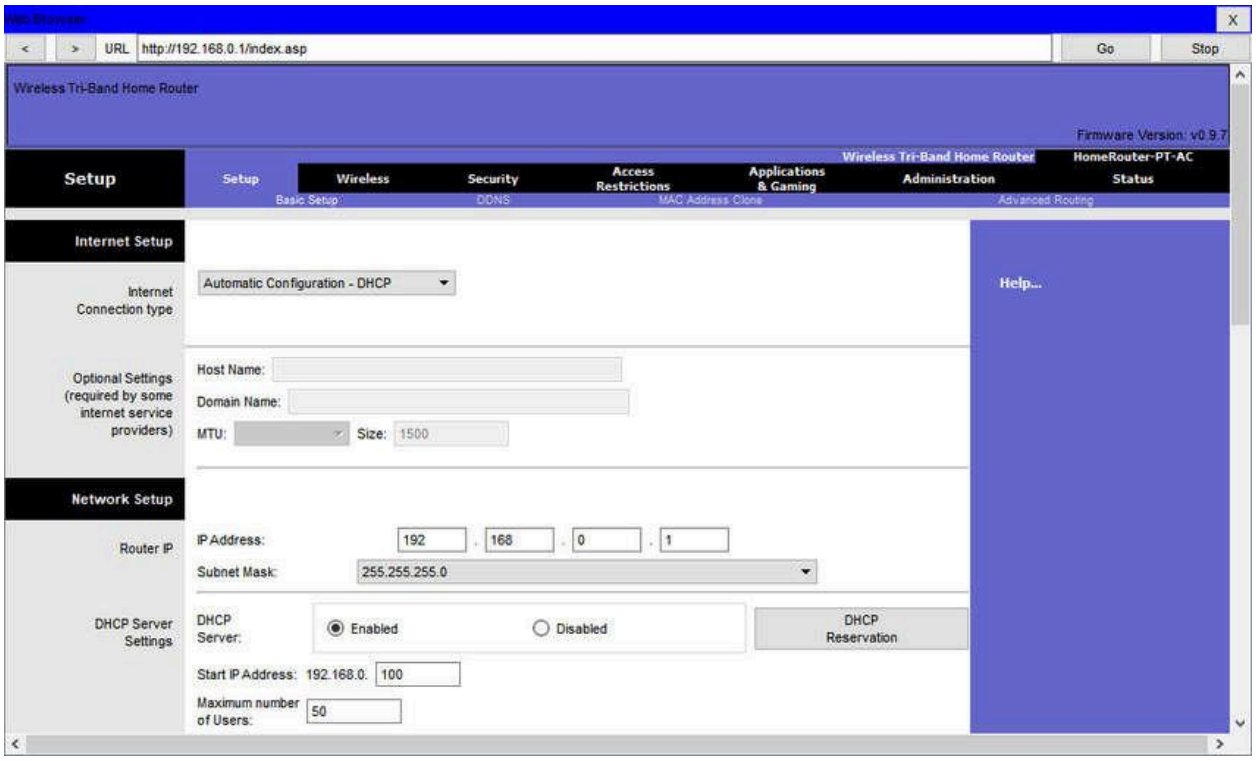

Change the default administrative password

Find administration port of router's GUI

- Username remains the same

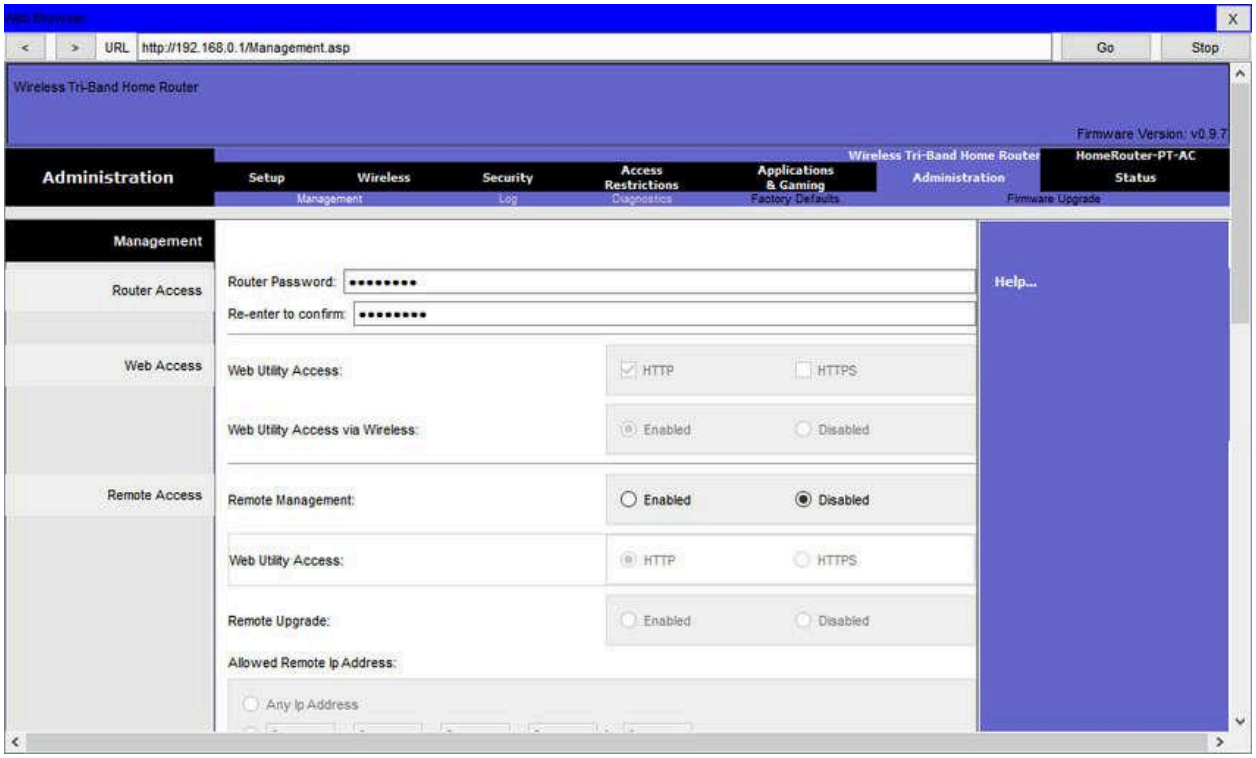

Log in with new administrative password

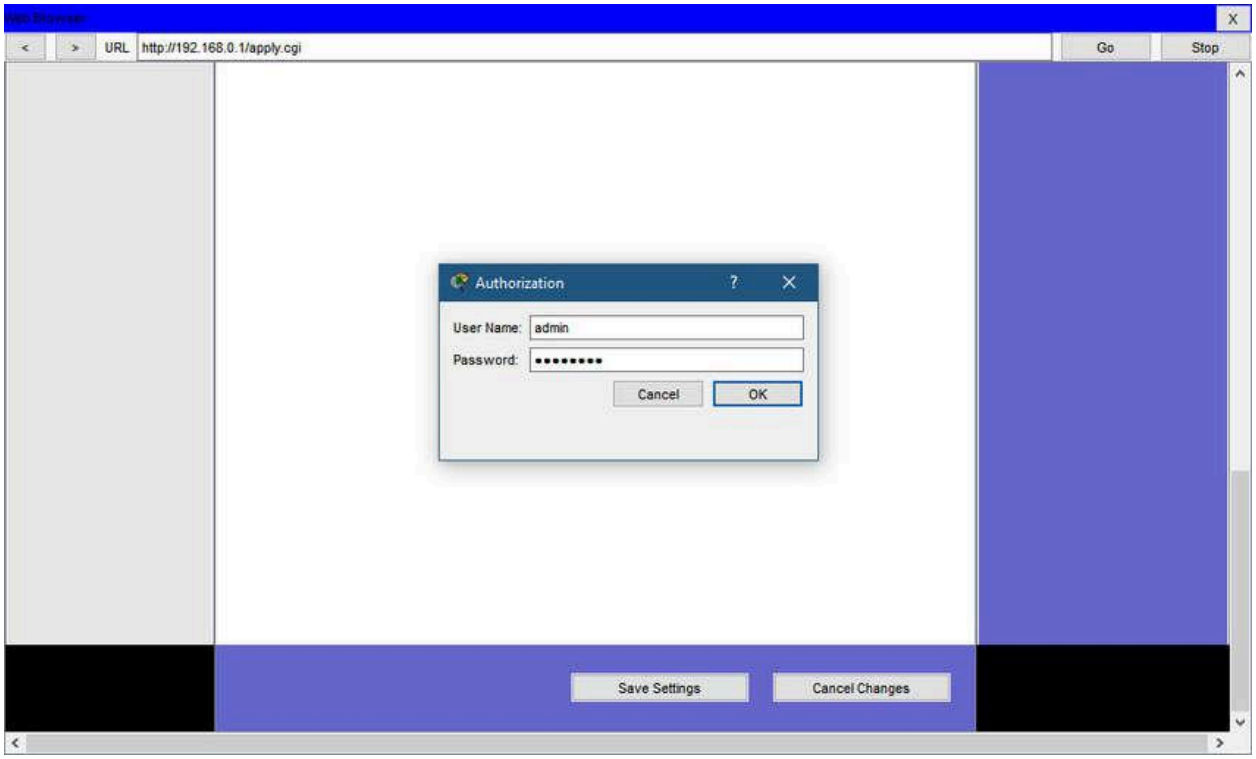

Change the default DHCP IPv4 addresses

Change the default router IPv4 address. Best practice to use private IPv4 addressing inside the network

- Example: 10.10.10.1

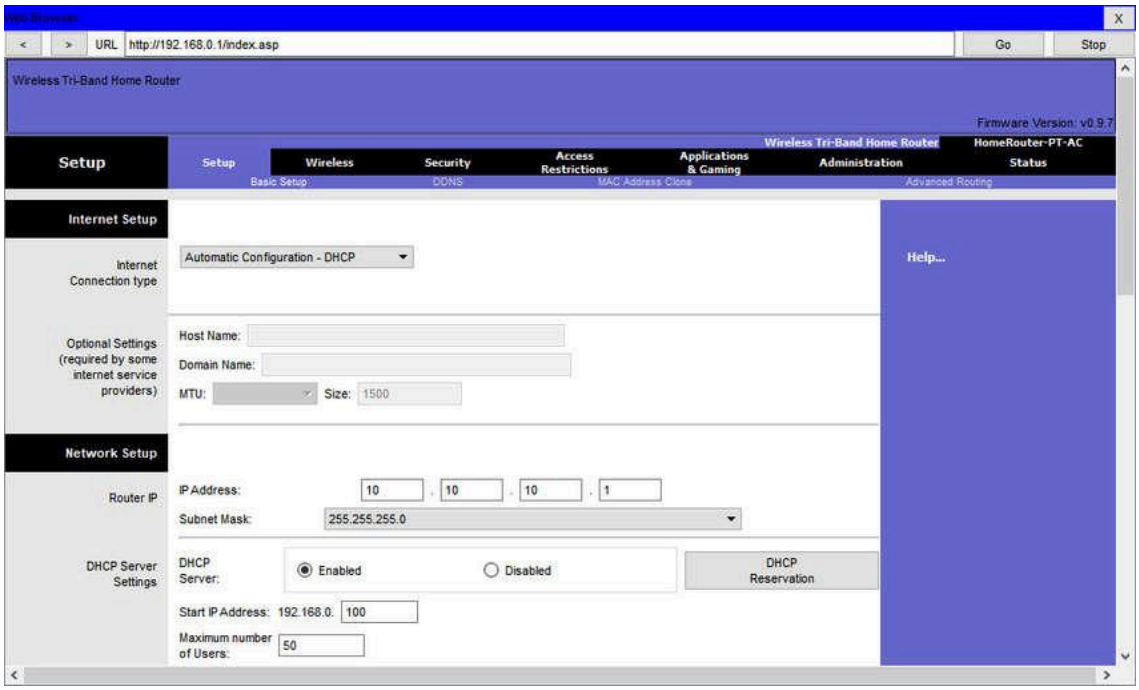

Renew the IP address

When you click save, you will temporarily lose access to router. Open command prompt and renew IP address

- **ipconfig /renew**

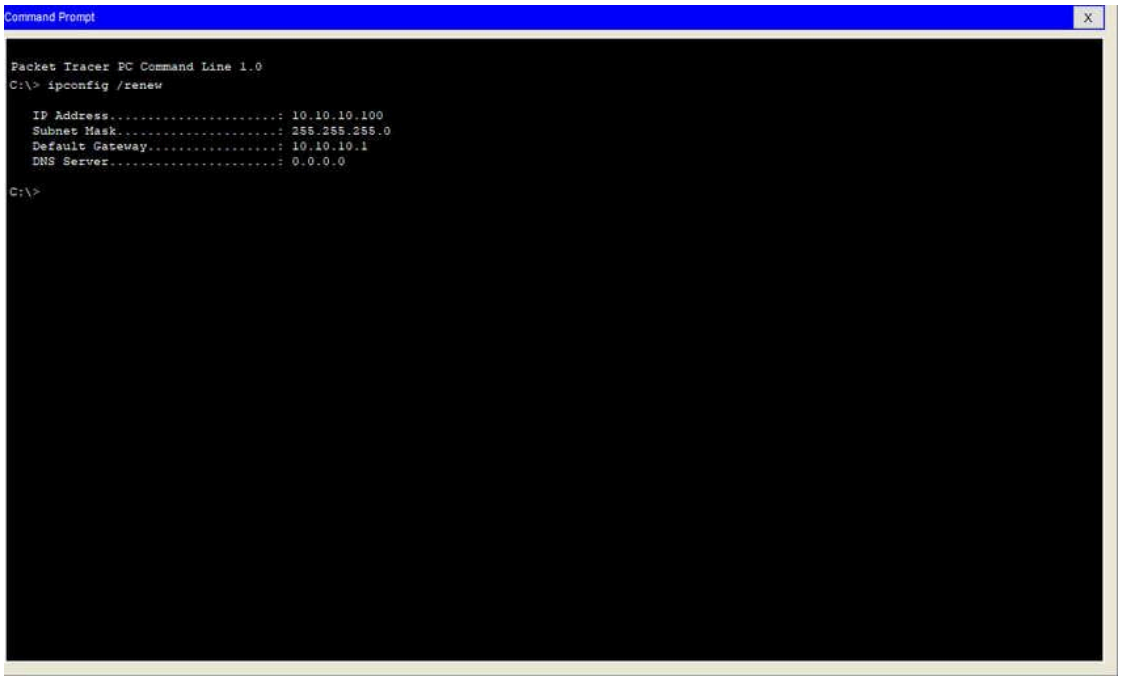

Log in router with new IP address

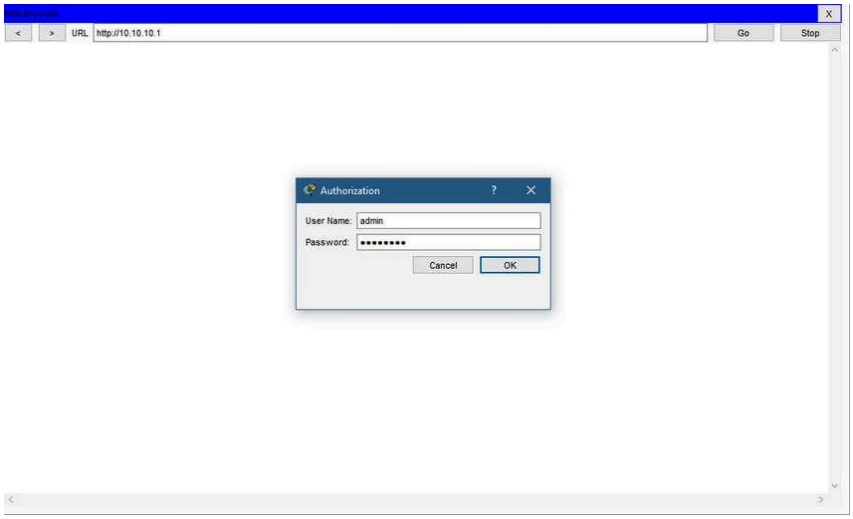

Basic Wireless Setup

1. View the WLAN defaults
2. Change the network mode
3. Configure the SSID
4. Configure the channel
5. Configure the security mode
6. Configure the passphrase

View the WLAN defaults

OOBE, wireless router provides wireless access to devices using a default wireless network name and password

- SSID
- Locate wireless settings to change the defaults

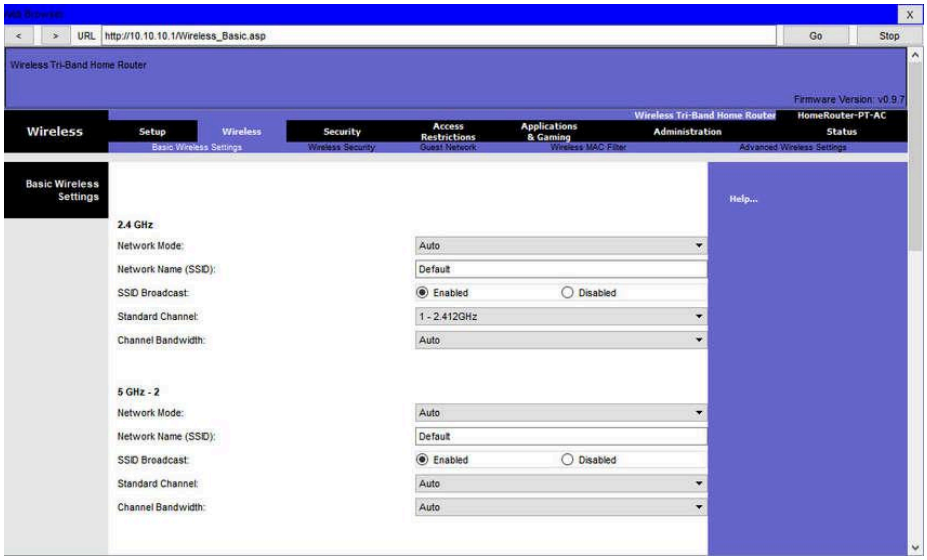

Change the network mode

Some wireless routers allow selection of which 802.11 standard to implement

- Today's configured for legacy or mixed mode support 802.11a, 802.11n, 802.11ac NICs

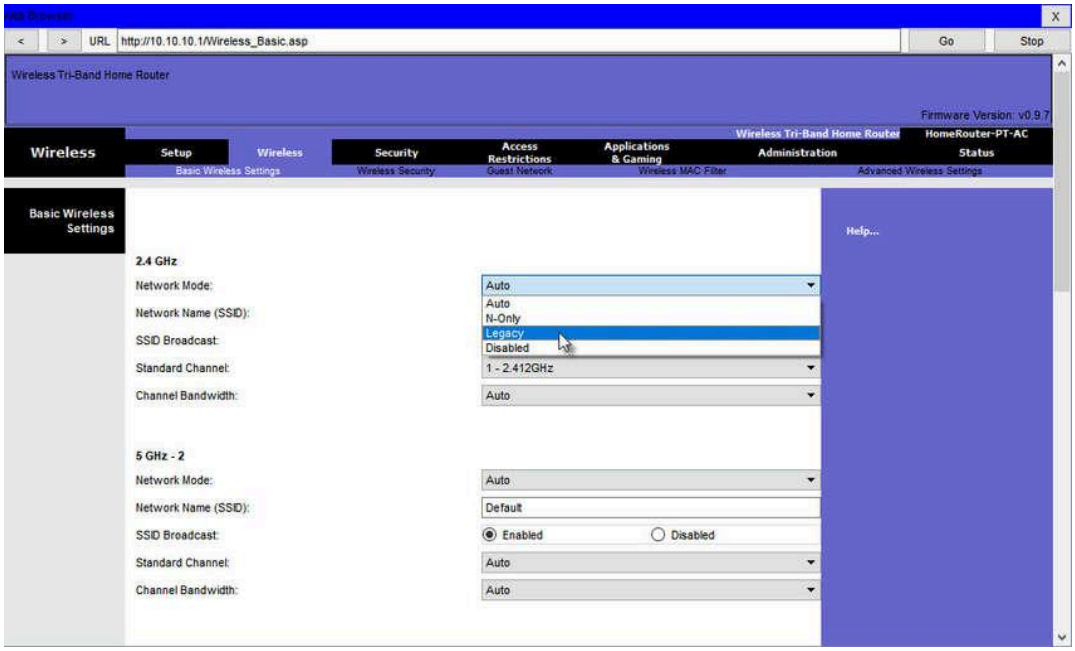

Configure the SSID

Assign the SSID to the WLANs. The router announces its presence by sending broadcasts advertising its SSID

- This allows wireless hosts to automatically discover the name of the wireless network
- If disabled, you must manually enter the SSID on each wireless device

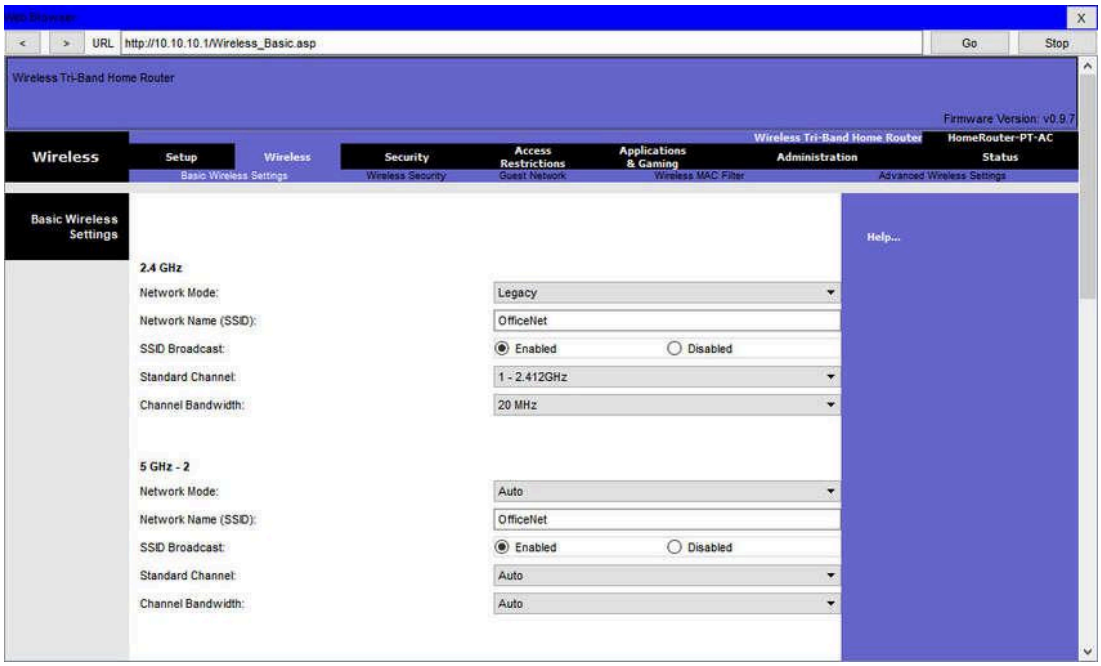

Configure the channel

Devices configured with the same channel within the 2.4Ghz ban may overlap and cause distortion, slowing down the wireless performance and potentially break network connections.

- Solution: Configure non-overlapping channels and access points that are near to each other
- Specifically channels 1, 6, 11 are non overlapping

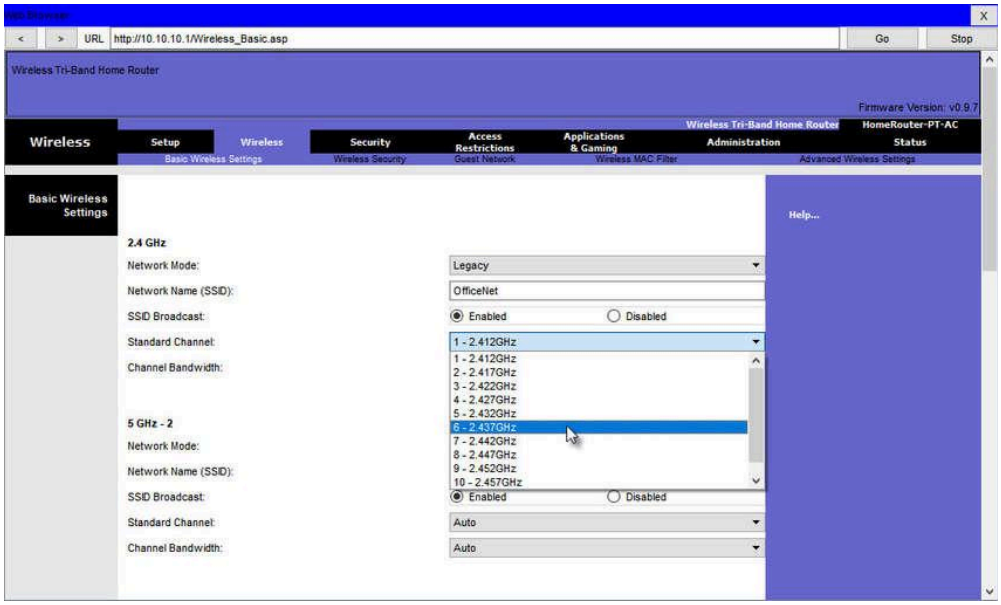

Configure the security mode

OOBE, wireless router may have no WLAN security configured

- Example below: The personal version of Wi-Fi Protected Access version 2 (WPA2-Personal) is selected for all three WLANs
- WPA2 with Advanced Encryption Standard (AES) encryption is ~~currently~~ strongest security mode

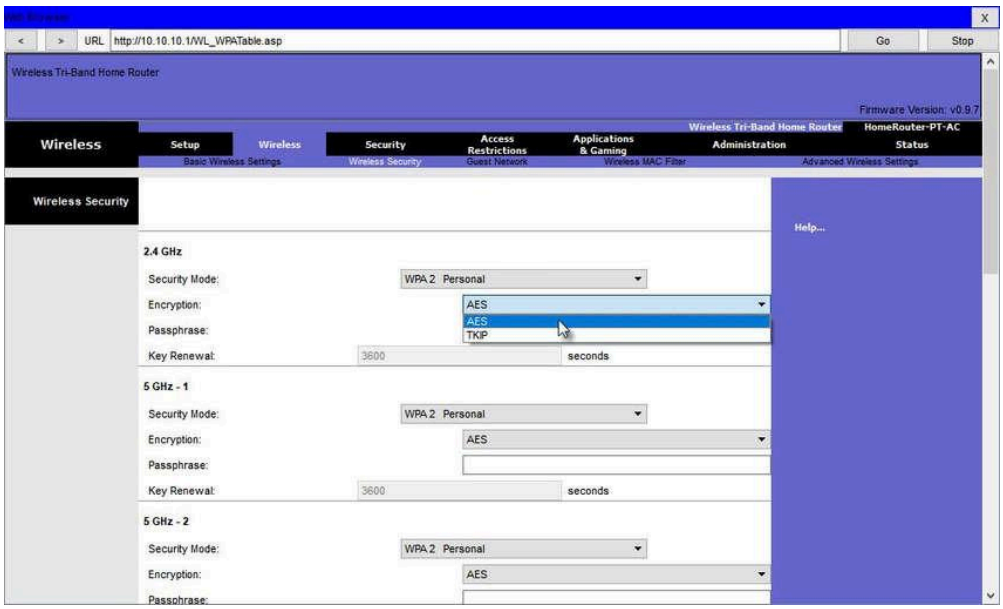

Configure the passphrase

WPA2 personal uses a passphrase to authenticate wireless clients

- Easier to use in a small office or home environment bc does not require authentication server
- Larger organizations implement WPA2 enterprise and require wireless clients to authenticate with a username and password

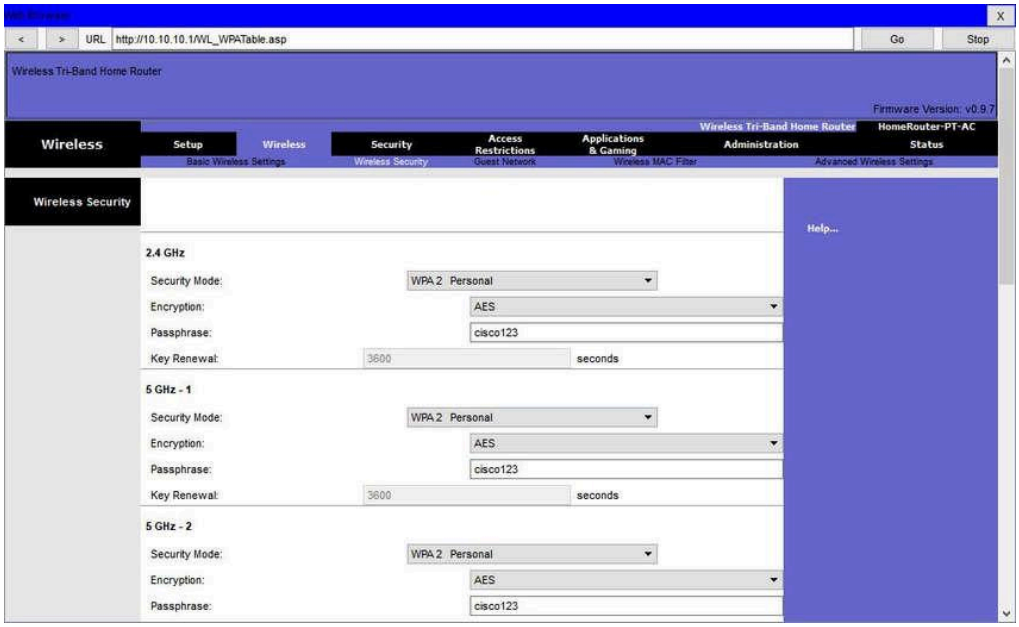

Configure a Wireless Mesh Network

- When want to extend range beyond approx 45 meters indoors and 90 meters outdoors, add wireless access points

Example below: Two access points are configured with same WLAN settings

- Notice channels selected are 1 and 11 so access points do not interfere with channel 6 configured previously on wireless router

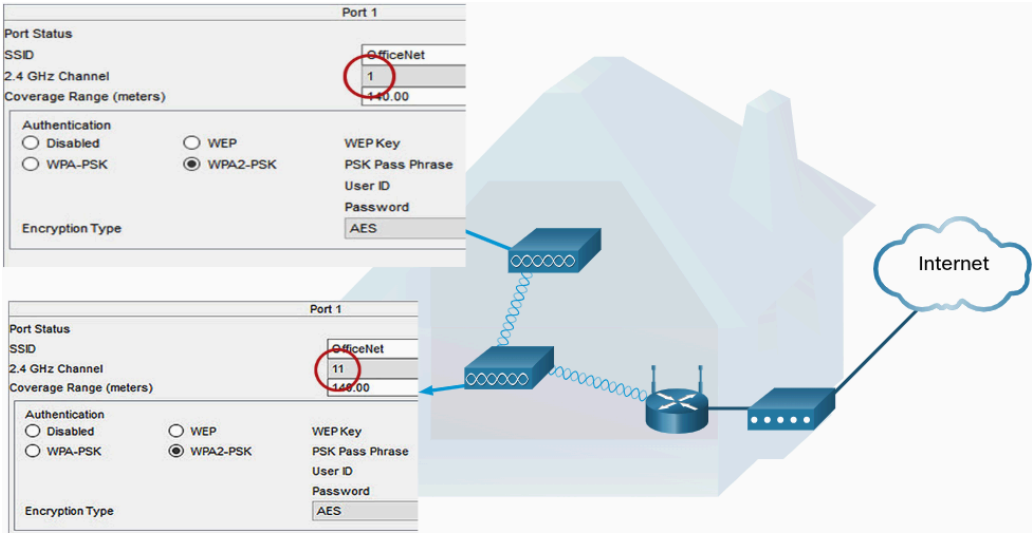

Extending a WLAN in a small office or home has become increasingly easier. Manufacturers have made creating a wireless mesh network (WMN) simple through smartphone apps)

- Buy system, disperse access points, plug them in, download app, and configure WMN

NAT for IPv4

Translates network private IP to public IP

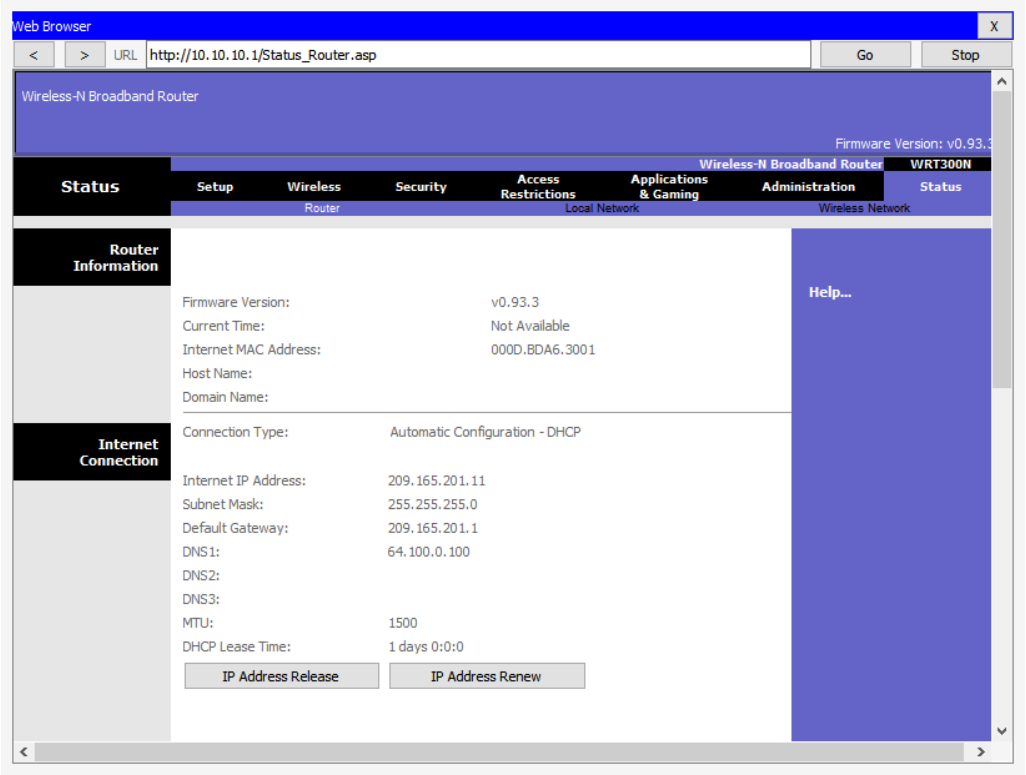

The 209.165.201.11 IPv4 address is publicly routable on the internet.

- Any address with the 10 in the first octet is a private IPv5 address and cannot be routed on the internet
- Router will use a process called Network Address Translation (NAT) to convert private IPv4 addresses to internet-routable IPv4 addresses

With NAT, a private (local) source IPv4 address is translated to a public (global) address

- The process is reversed for incoming packets
- The router is able to translate many internal IPv5 addresses into public addresses, by using NAT

Some ISPs use private addressing to connect to customer devices. However, eventually your traffic will leave the provider's network and be routed on the internet.

NAT makes looking up your public IP by tracking the source port numbers for every session established by a device. If your ISP has IPv6 enabled, you will see a unique IPv6 address for each device.

Quality of Service

By configuring this, can guarantee that certain traffic types, such as voice and video, are prioritized over traffic that is not as time-sensitive.

- On some wireless routers, traffic can also be prioritized on specific ports

Figure below is simplified mockup of a QoS interface based on a Netgear GUI

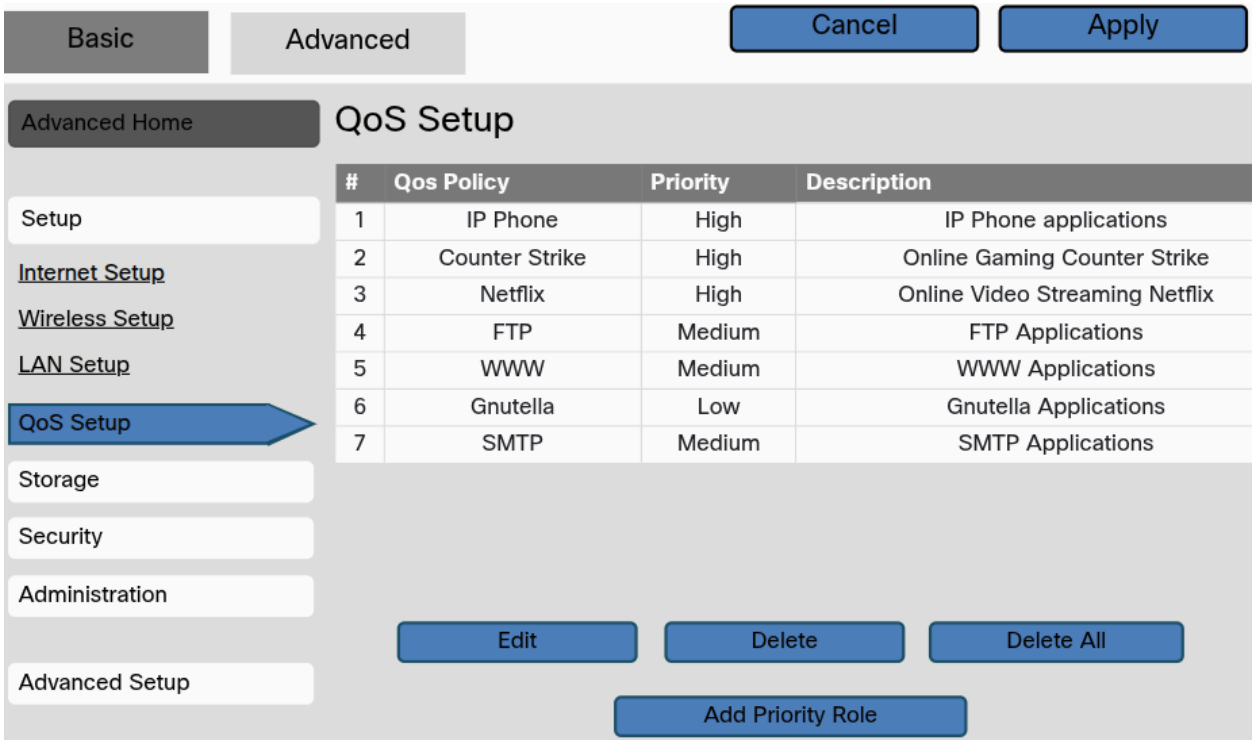

Port Forwarding

Routers will typically block TCP and UDP ports to prevent unauthorized access in and out of a LAN.

- Port forwarding is a rule-based method of directing traffic between devices on separate networks.

When traffic reaches the router, the router determines if the traffic should be forwarded to a certain device based on the port number found with the traffic

Example: Router might be configured to forward port 80, associated with HTTP

- When router receives packet with destination port of 80, router forwards traffic to the server inside the network that serves web pages

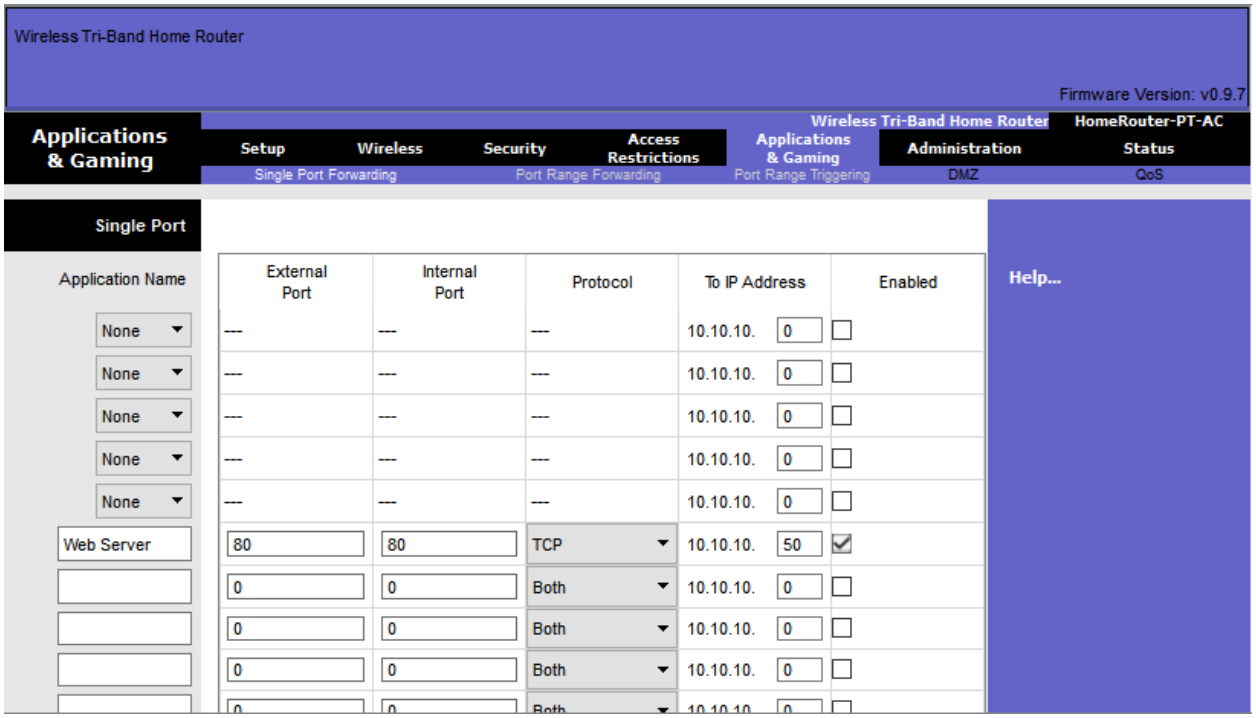

Port triggering allows the router to temporarily forward data through inbound ports to a specific device. You can use port triggering to forward data to a computer only when a designated port range is used to make an outbound request

- Example: A game uses ports 27000 to 27100 for connecting with other players
    - These are trigger ports
- Example: A chat client might use port 56 for connecting the same players so that they can interact with each other.
    - If there is a gaming traffic on an outbound port within the triggered port range, inbound chat traffic on port 56 is forwarded to the computer that is being used to play the video game and chat with friends
    - When the game is over and the triggered ports are no longer in use, port 56 is no longer allowed to send traffic of any type to this computer

Configure a Basic WLAN on the WLC

WLC Topology

The access point (AP) is a controller-based AP as opposed to an autonomous AP

- Recall that controller-based APs require no initial configuration and are often called lightweight APs (LAPs)
    - LAPs use the Lightweight Access Point Protocol (LWAPP) to communicate with a WLAN controller (WLC)
- Controller-based APs are useful in situations where many APs are required in the network
    - As more APs are added, each AP is automatically configured and managed by the WLC

Topology

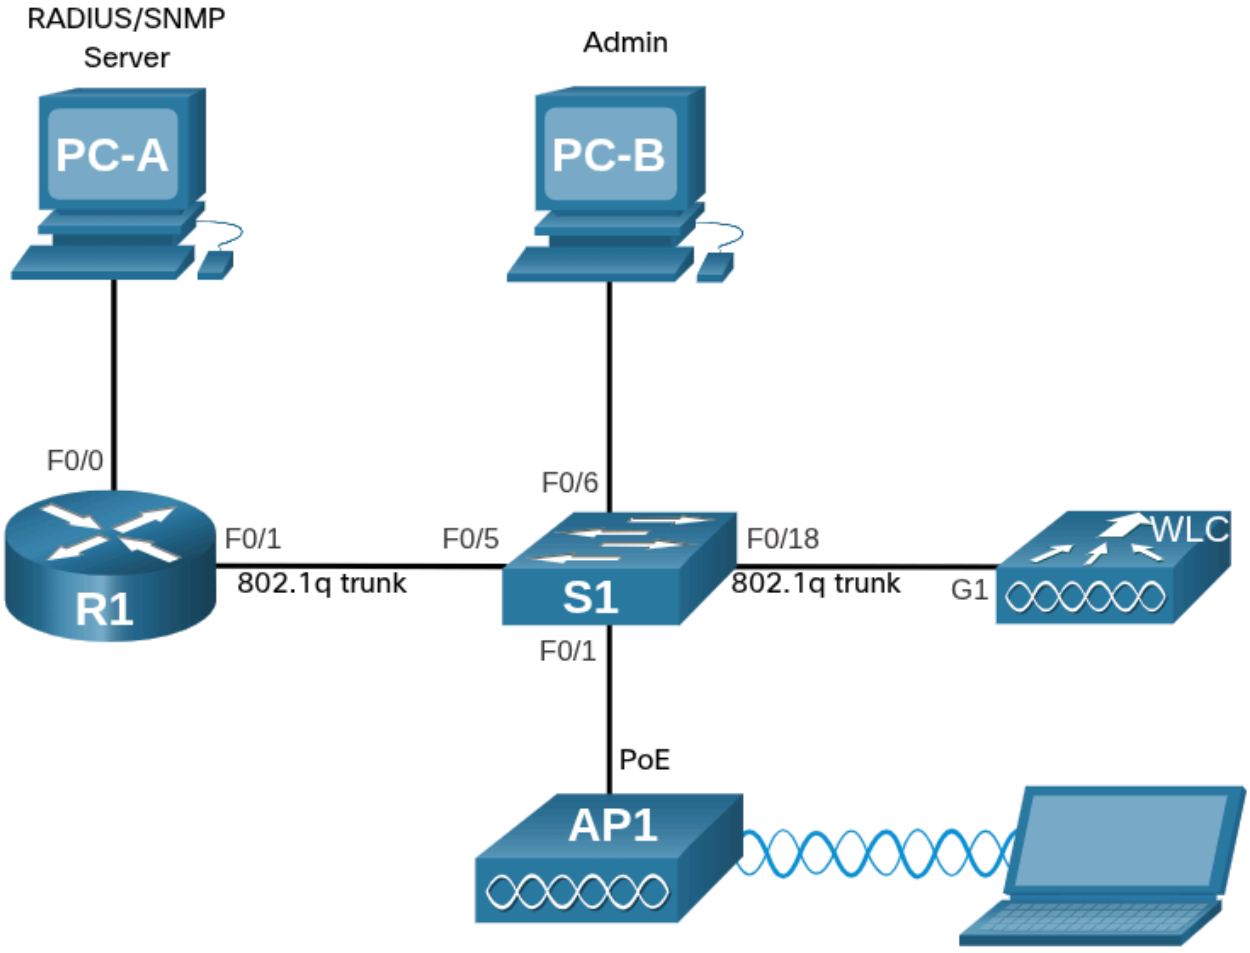

The AP is PoE, which means it is powered over the Ethernet cable that is attached to the switch

Addressing Table

Device	Interface	IP Address	Subnet Mask
R1	F0/0	172.16.1.1	255.255.255.0
R1	F0/1.1	192.168.200.1	255.255.255.0
S1	VLAN 1	DHCP	
WLC	Management	192.168.200.254	255.255.255.0
AP1	Wired 0	192.168.200.3	255.255.255.0
PC-A	NIC	172.16.1.254	255.255.255.0
PC-B	NIC	DHCP	
Wireless Laptop	NIC	DHCP	

Log in to the WLC

Configuring a wireless LAN controller (WLC) is not that much different from configuring a wireless router

- A WLC controls APs and provides more services and management capabilities

Note - Image below is GUI from Cisco 3504 Wireless Controller

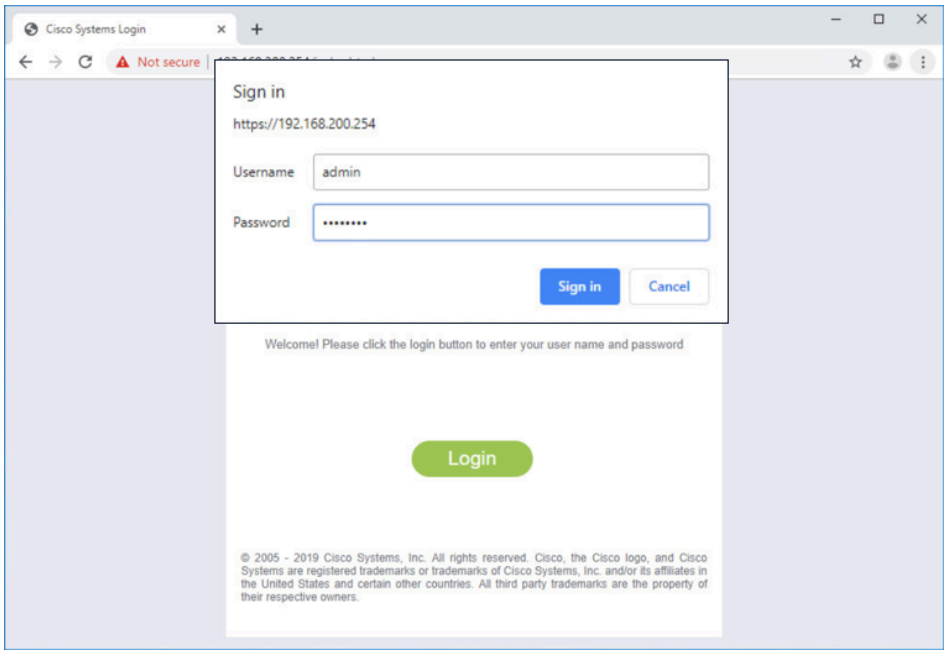

The **Network Summary** page is a dashboard that provides a quick overview of the number of configured wireless networks, associated access points (APs), and active clients

- Can also see number of rogue access points and clients

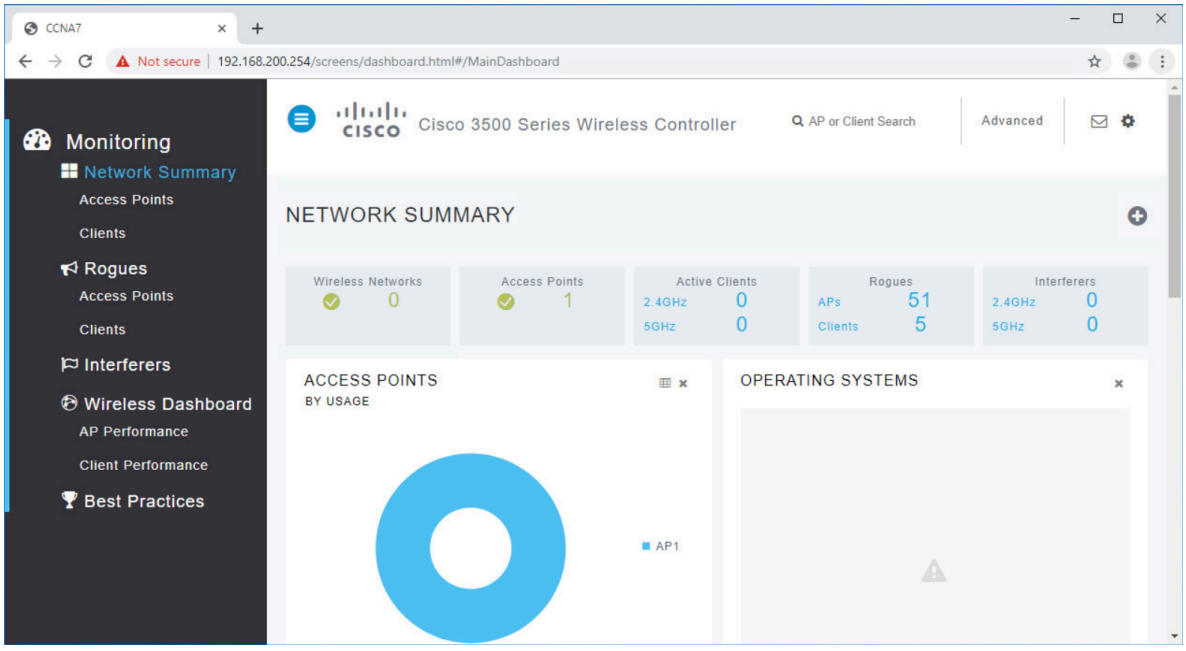

View AP Information

Click **Access Points** from the left menu to view an overall picture of the AP's system information and performances

- The AP is using IP address 192.168.200.3
- Because Disco Discovery Protocol (CDP) is active on this network, the WLC knows that the AP is connected to the F0/1 port on switch

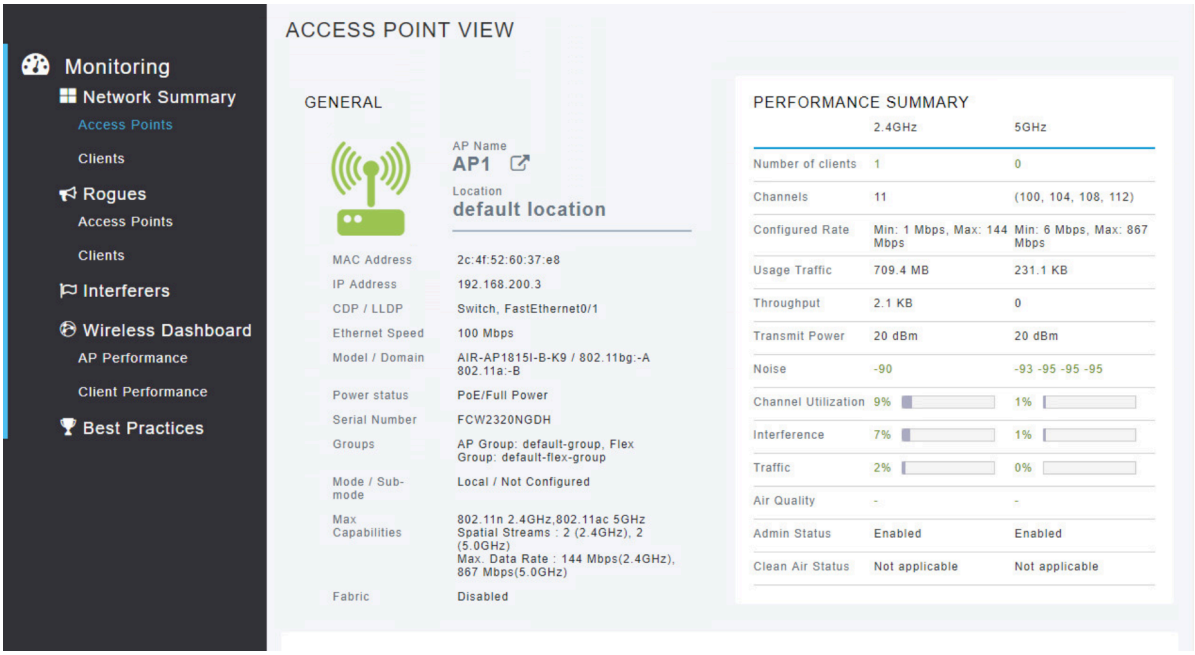

This AP in the topology is a Cisco Aironet 1815i which means you can use the command-line and a limited set of familiar IOS commands

- Example below: the network administrator pinged the default gateway, ping WLC, and verified the wired interface

```
AP1# ping 192.168.200.1
Sending 5, 100-byte ICMP Echos to 192.168.200.1, timeout is 2 seconds
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max =
1069812.242/1071814.785/1073817.215 ms
AP1# ping 192.168.200.254
Sending 5, 100-byte ICMP Echos to 192.168.200.254, timeout is 2 seconds
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max =
1055820.953/1057820.738/1059819.928 ms
AP1# show interface wired 0
wired0    Link encap:Ethernet  HWaddr 2C:4F:52:60:37:E8
          inet addr:192.168.200.3  Bcast:192.168.200.255  Mask:255.255.255.255
          UP BROADCAST RUNNING PROMISC MULTICAST  MTU:1500  Metric:1
          RX packets:2478 errors:0 dropped:3 overruns:0 frame:0
          TX packets:1494 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:80
          RX bytes:207632 (202.7 KiB)  TX bytes:300872 (293.8 KiB)
AP1#
```

Advanced Settings

For Cisco 3504 Wireless Controller

- Click **Advanced** in upper right-hand corner to access the advanced **Summary** page

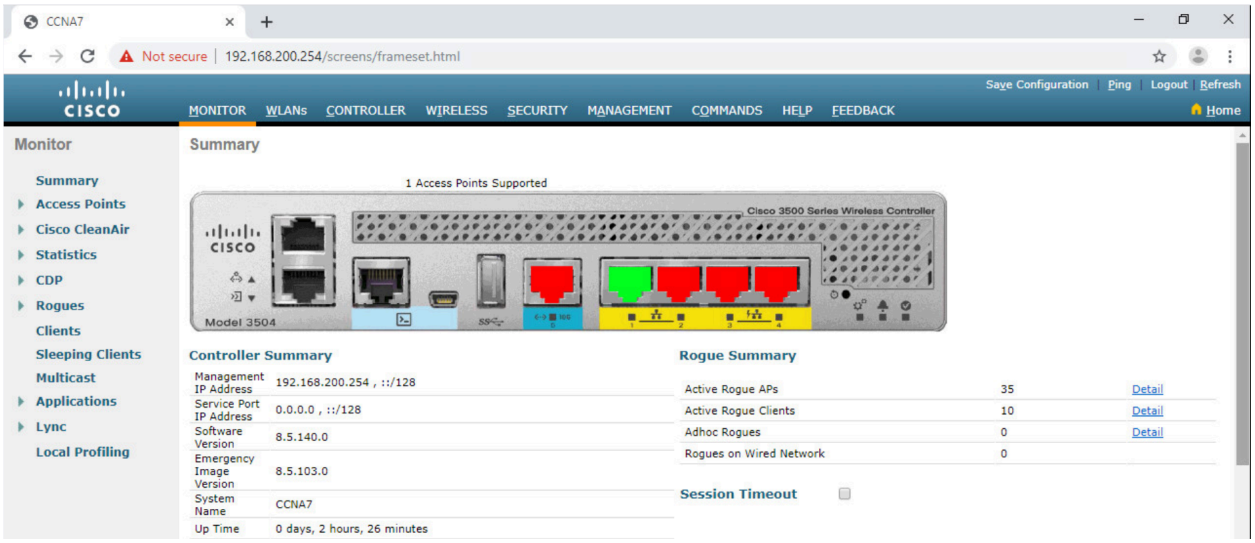

Configure a WLAN

Wireless LAN controllers have ports and interfaces

- The resembling switch ports are virtual interfaces
- Created in software and very similar to VLAN interfaces
- Each interface will carry traffic from a WLAN is configured on the WLC as a different VLAN

Cisco 3504 WLC

- Can support 150 access points and 4096 VLANs, however only has five physical ports
    - Each physical port can support many APs and WLANs
- The ports are essentially trunk ports that can carry traffic from multiple VLANs to a switch for distribution to multiple APs
    - Each AP can support multiple WLANs

Basic WLAN configuration on the WLC:

1. Create the WLAN
2. Apply and Enable the WLAN
3. Select the interface
4. Secure the WLAN
5. Verify the WLAN is Operational
6. Monitor the WLAN
7. View Wireless Client Information

Create the WLAN

Administrator is creating a new WLAN that will use **Wireless_LAN** as name and SSID

- ID is arbitrary value used to identify the WLAN in display output on WLC

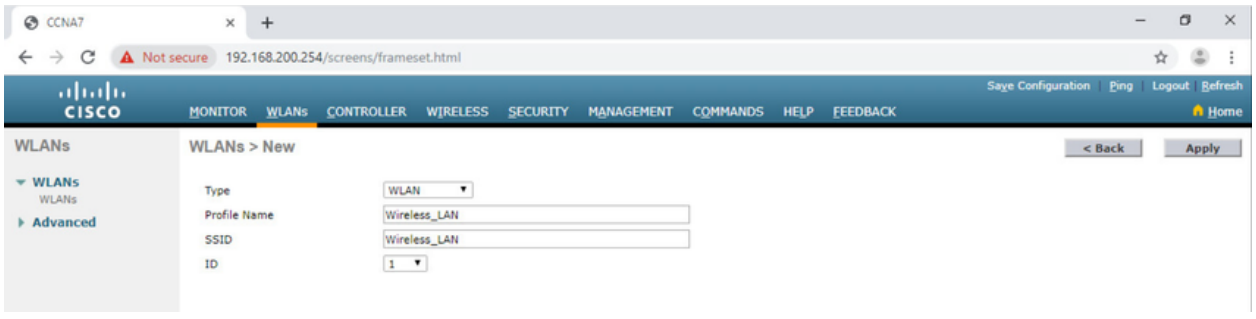

Apply and Enable the WLAN

After clicking **Apply**, the network administrator must enable the WLAN before it can be accessed by users.

- The Enable checkbox allows the network administrator to configure a variety of features for the WLAN, as well as additional WLANs, before enabling them for wireless client access
- From here, the network administrator can configure a variety of settings for the WLAN including security, QoS, policies, and other advanced settings

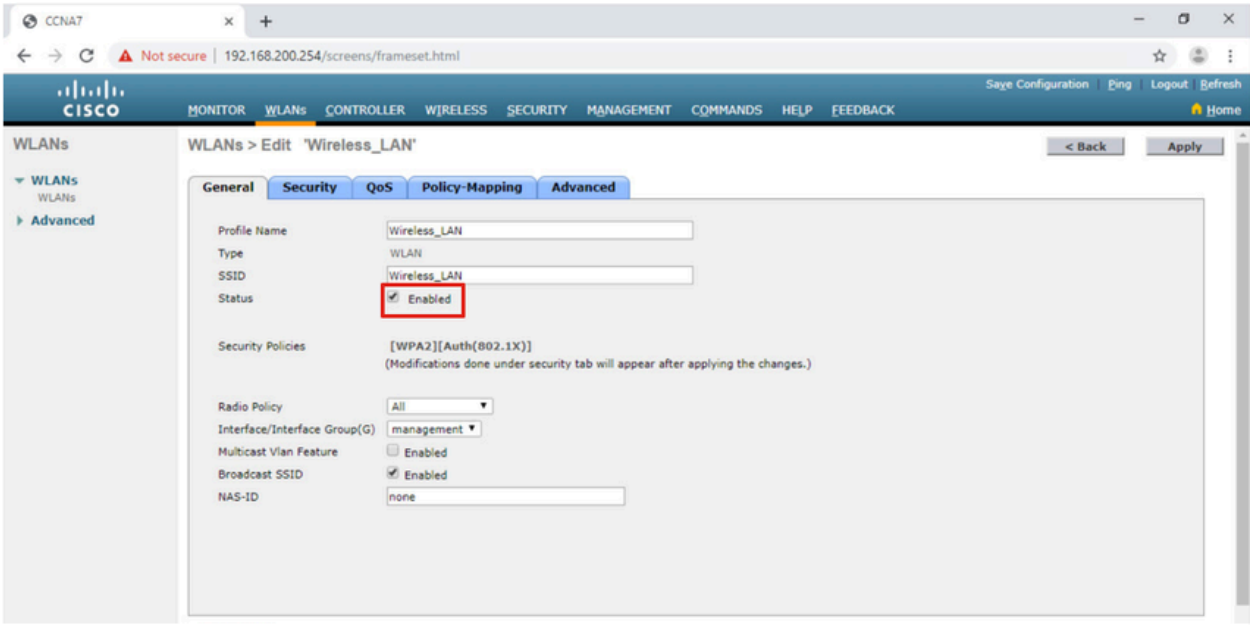

Select the Interface

When you create a WLAN, you must select the interface that will carry the WLAN traffic

- Figure below: shows the selection of an interface that has already been created on the WLC

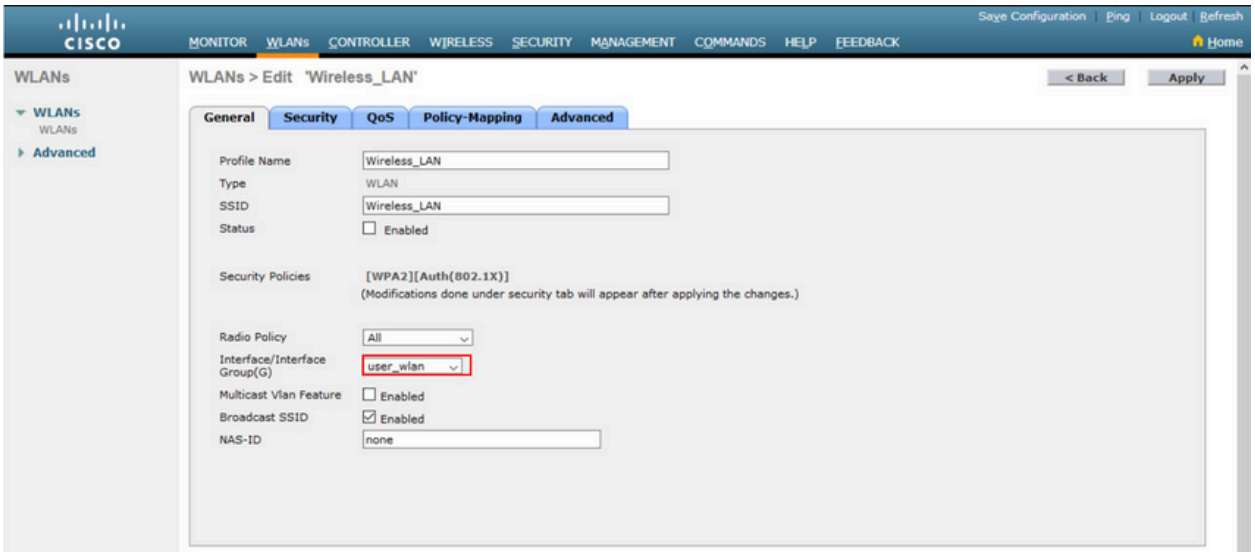

Secure the WLAN

Click the Security tab to access all the available options for securing the LAN. The network admin wants to secure Layer 2 with **WPA2-PSK**

- WPA2 and 802.1X are set by default

In the Layer 2 Security drop down box, verify that **WPA+WPA2** is selected.

Click PSK and enter the pre-shared key. Then click **Apply**.

- This will enable the WLAN with WPA2-PSK authentication.
- Wireless clients that know the pre-shared key can now associate and authenticate with the AP

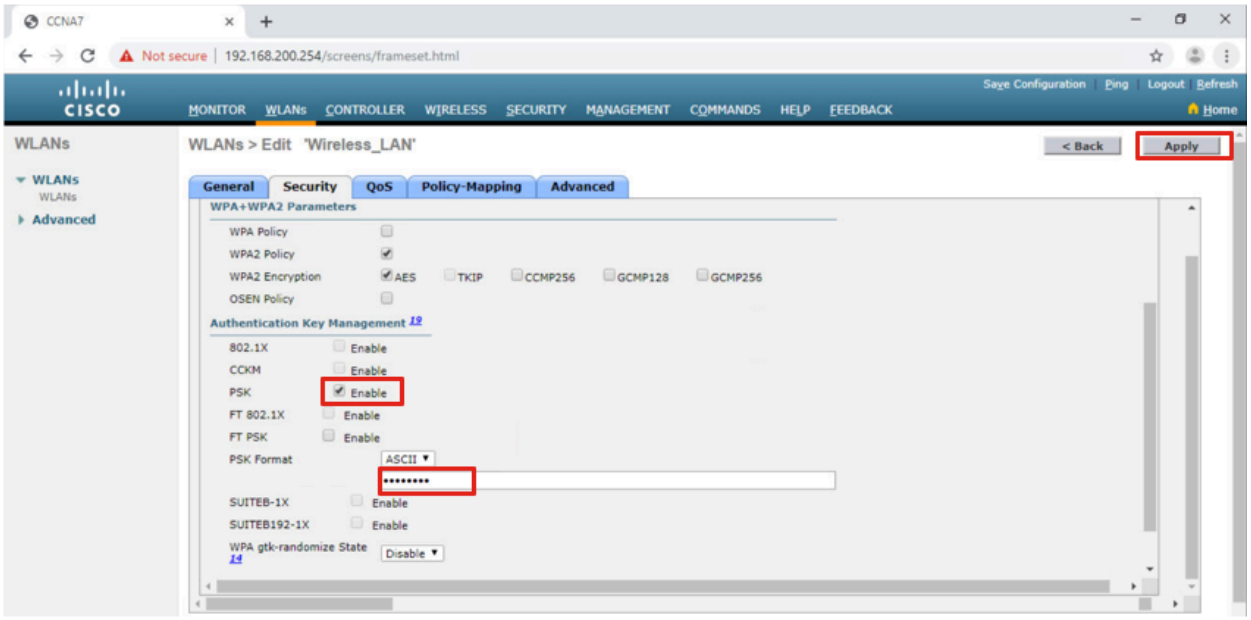

Verify the WLAN is Operational

Click **WLANs** in the menu on the left to view the newly configured WLAN

- In figure, you can verify that WLAN ID 1 is configured with **Wireless_LAN** as name and SSID, it is enabled, and is using WPA2 PSK security

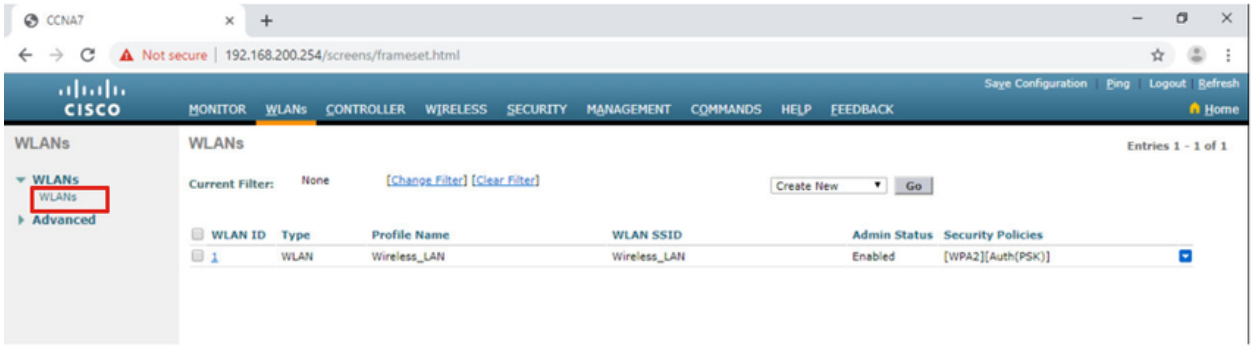

Monitor the WLAN

Click the **Monitor** tab at the top to access the advanced **Summary** page again. Here you can see that the **Wireless_LAN** now has one client using its services

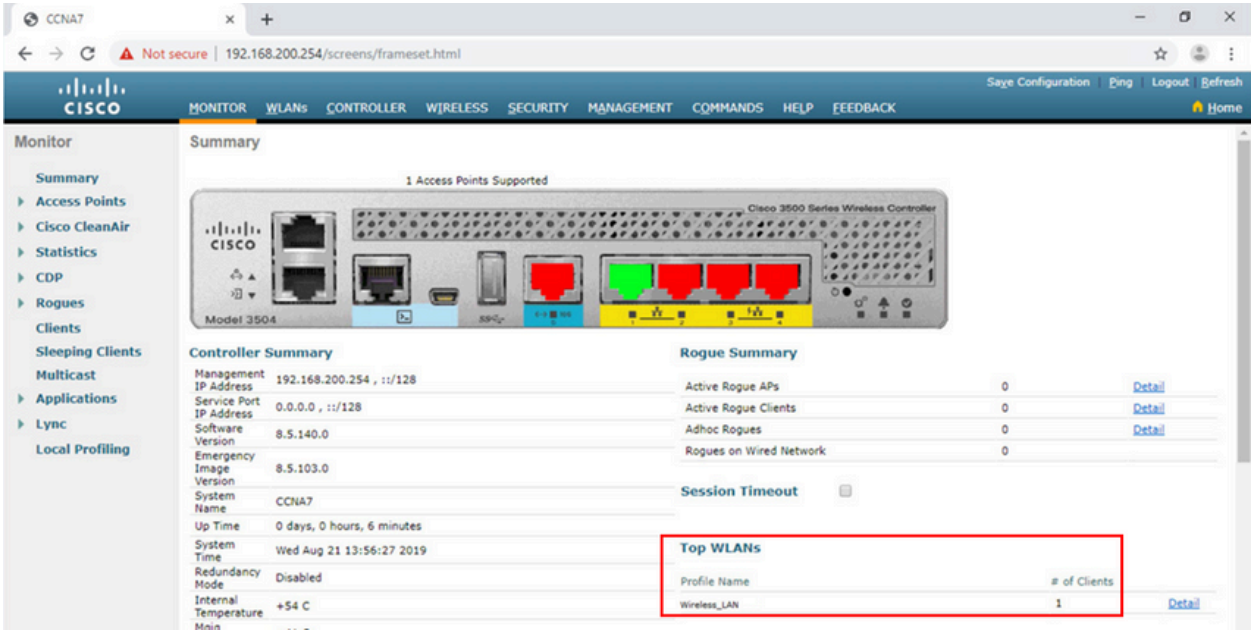

View Wireless Client Details

Click **Clients** in the left menu to view more information about the clients connected to the WLAN.

- One client is attached to **Wireless_LAN** through AP1 and was given the IP address 192.168.5.2
- DHCP services in this topology are provided by the router

Configure a WPA2 Enterprise WLAN on the WLC

SNMP and RADIUS

Simple Network Management Protocol (SNMP)

- Used to monitor the network

The network administrator wants the WLC to forward all SNMP log messages, called traps, to the SNMP server.
The network administrator also wants to use a RADIUS server for authentication, authorization, and accounting (AAA) services.

- Instead of entering a publicly known pre-shared key to authenticate, as they do with WPA2-PSK, users will enter their own username and password credentials
    - Credentials will be verified by RADIUS server
    - This way, individual user access can be tracked and audited if necessary and user accounts can be added or modified from a central location

Remote Authentication Dial-In User Service (RADIUS)

- This server is required for WLANs that are using WPA2 Enterprise authentication

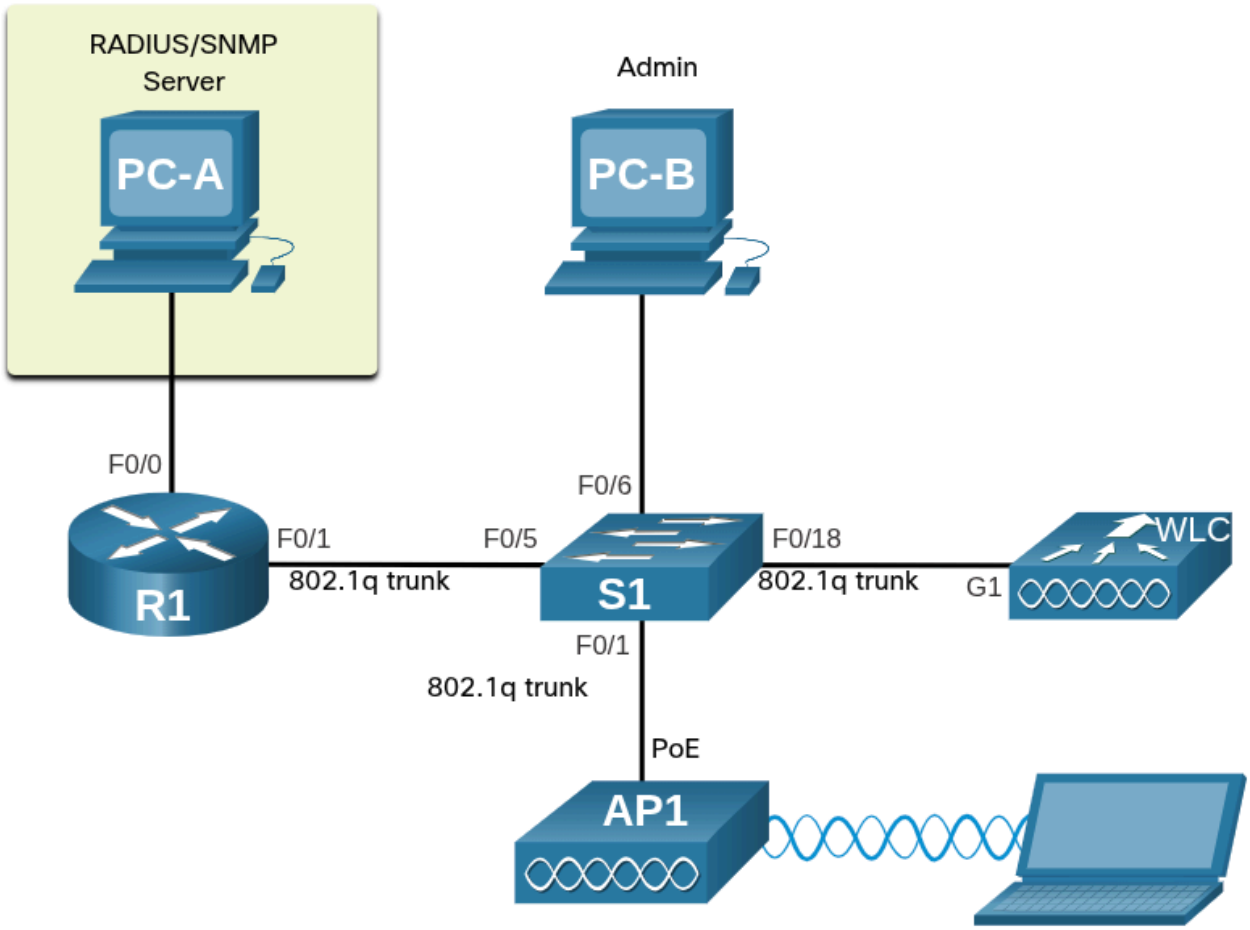

Configure SNAP Server Information

1. Click the **MANAGEMENT** tab to access a variety of management features. SNMP is listed at the top of the menu on the left click.
2. Click **SNMP** to export the sub-menus, and then **Trap Receivers**.
3. Click **New…** to configure the new SNMP trap receiver

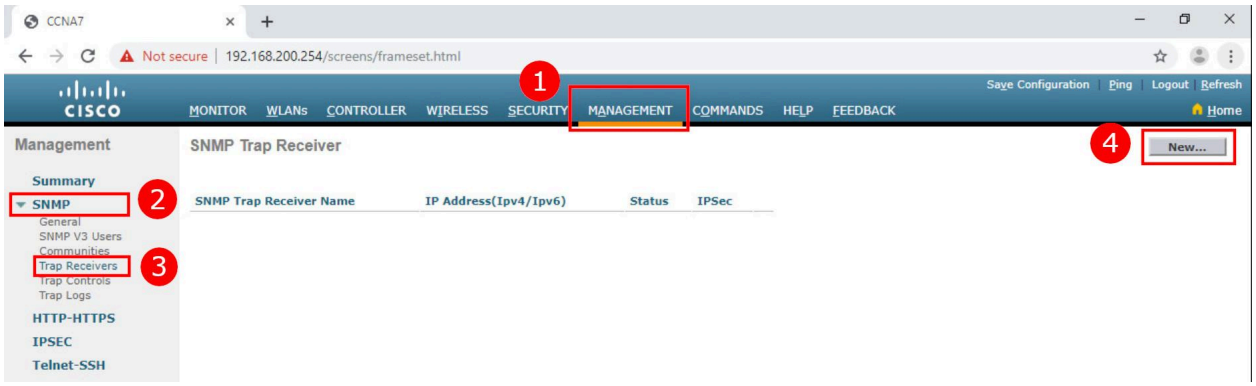

1. Click **MANAGEMENT**
2. Click **SNMP**
3. Click **Trap Receivers**
4. Click **New…**

Enter the SNMP Community name and the IP address (IPv4 or IPv6) for the SNMP server. Click **Apply**. The WLC will not forward SNMP log messages to the SNMP server

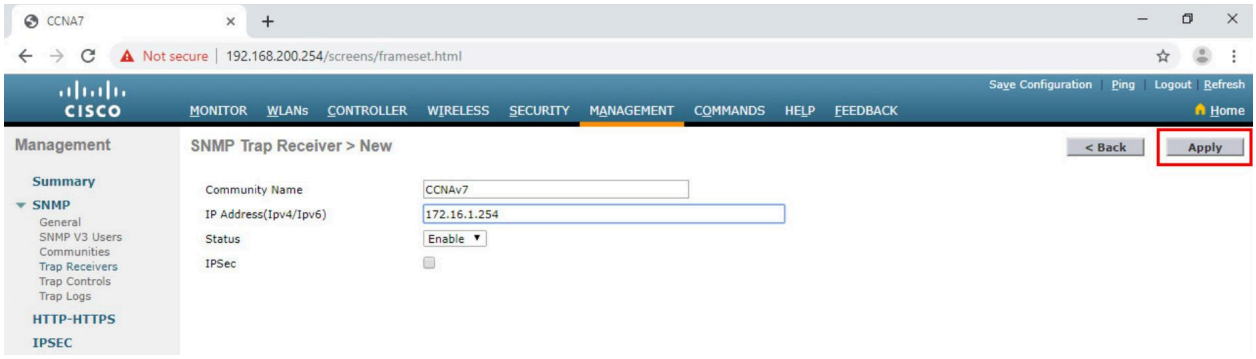

Configure RADIUS Server Information

In example configuration: The network admin wants to configure a WLAN using WPA2 Enterprise, as opposed to WPA2 Personal or WPA2 PSK. Authentication will be handled by the RADIUS server running on PC-A

To configure the WLC with the RADIUS server information

1. Click the **SECURITY** tab > **RADIUS** > **Authentication**.
2. No RADIUS servers are currently configured
    a. Click **New…** to add PC-A as the RADIUS server

1. Click **SECURITY**
2. Click **RADIUS**
3. Click **Authentication**
4. Click **New…**

Enter the IPv4 address for PC-A and the shared secret. This is the password used between the WLC and the RADIUS server. It is not for users. Click **Apply**

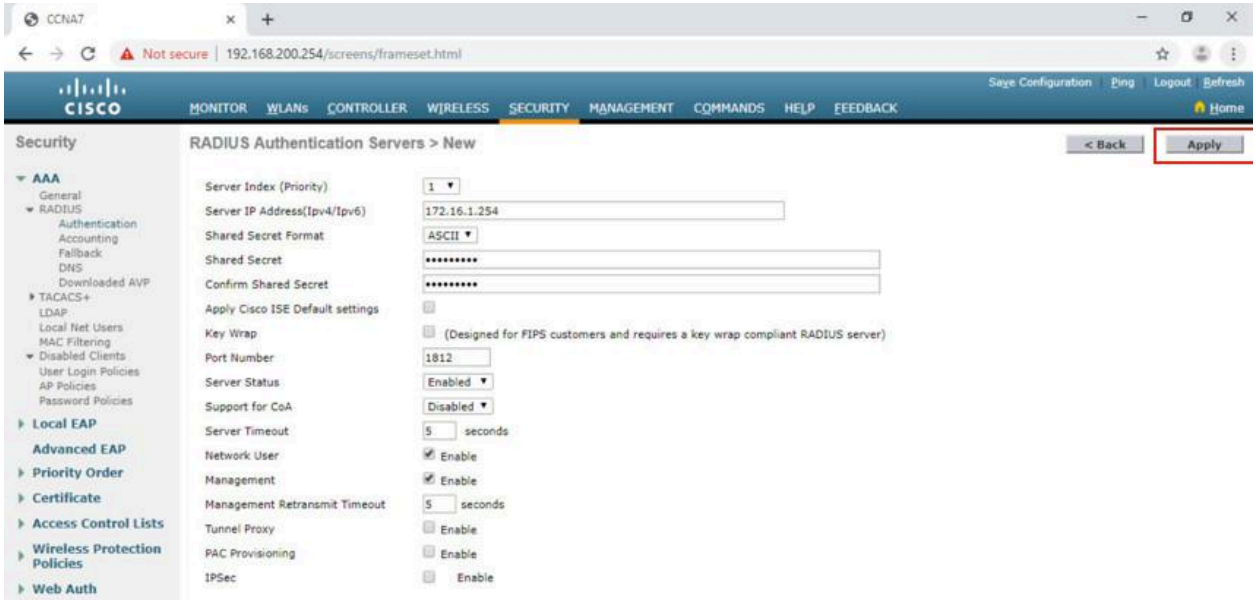

After clicking **Apply**, the list of configured **RADIUS Authentication Servers** refreshes with the new server listed

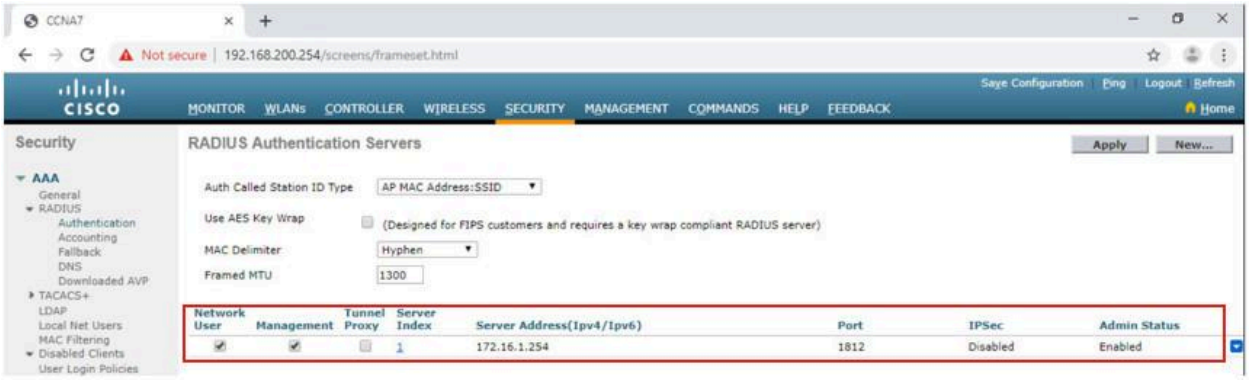

Topology with VLAN 5 Addressing

Each WLAN configured on the WLC needs its own virtual interface. The WLC has five physical ports for data traffic. Each physical port can be configured to support multiple WLANs, each on its own virtual interface.

- Physical ports can also be aggregated to create high-bandwidth links

The network admin has decided that the new WLAN will use interface VLAN 5 and network 192.168.5.0/24

- R1 already has a subinterface configured and active for VLAN 5

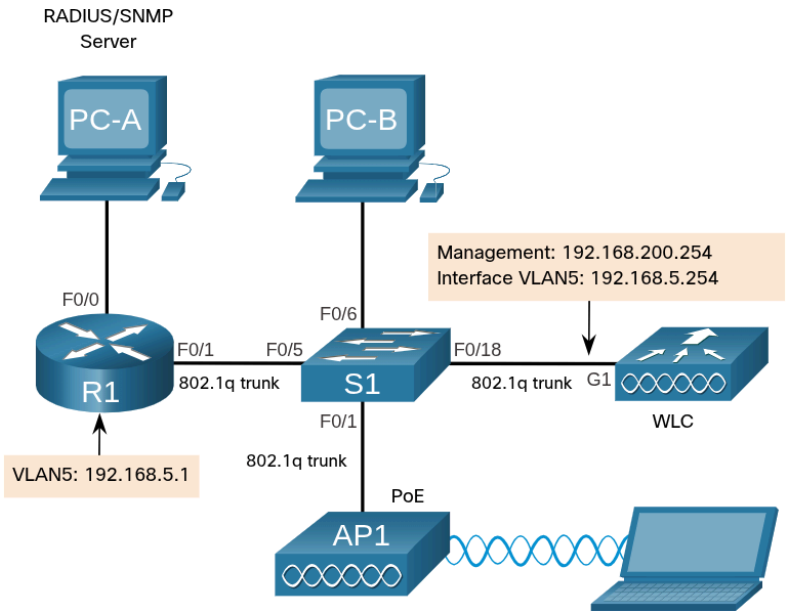


```
R1# show ip interface brief
Interface                  IP-Address      OK? Method Status                Protocol
FastEthernet0/0            172.16.1.1      YES manual up                    up
FastEthernet0/1            unassigned      YES unset  up                    up
FastEthernet0/1.1          192.168.200.1   YES manual up                    up
FastEthernet0/1.5          192.168.5.1  YES manual up                    up
(output omitted)
R1#
```

Configure a New Interface

Vlan interface configuration on WLC:

1. Create a new interface
2. Configure the VLAN name and ID
3. Configure the port and interface address
4. Configure the DHCP server address
5. Apply and Confirm
6. Verify Interfaces

Create a new interface

Click **CONTROLLER > Interfaces > New…**

Configure the VLAN name and ID

Network admin configures the interface name as **vlan5** and the VLAN ID as **5**.

- Clicking **Apply** will create the new interface

Configure the port and interface address

On the **Edit** page for the interface, configure the physical port number. G1 in the topology is Port Number 1 on the WLC.

Then configure the VLAN 5 interface addressing

In figure below: VLAN 5 is assigned IPv4 address 192.168.5.254/24

- R1 is the default gateway at IPv4 address 192.168.5.1

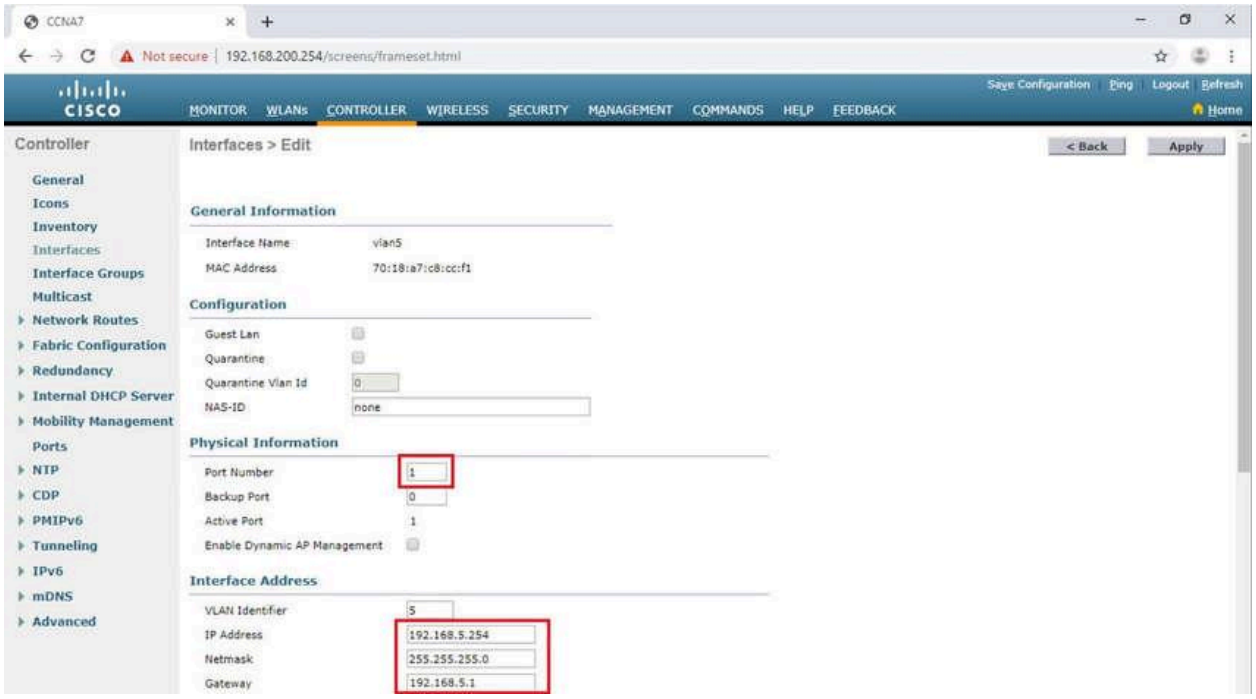

Configure the DHCP Server address

In larger enterprises, WLCs will be configured to forward DHCP messages to a dedicated DHCP server.

Scroll down the page to configure the primary DHCP server as IPv4 address 192.168.5.1. This is the default gateway router address

- The router is configured with a DHCP pool for the WLAN network
- AS hosts join the WLAN that is associated with the VLAN 5 interface, they will receive addressing information from this pool

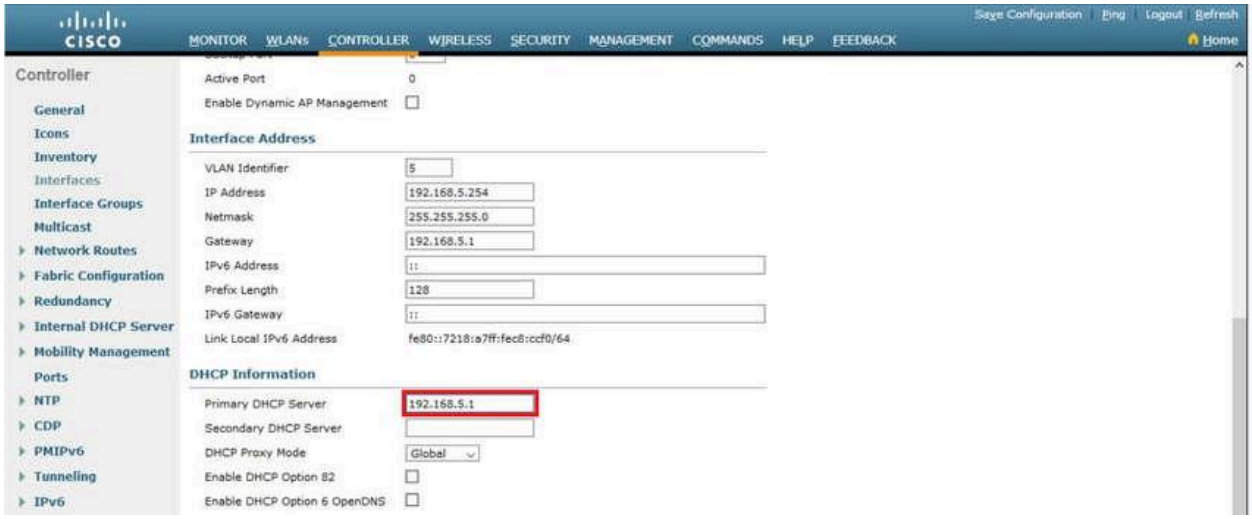

Apply and Confirm

Scroll to the top and click **Apply**. Click **OK** for the warning message

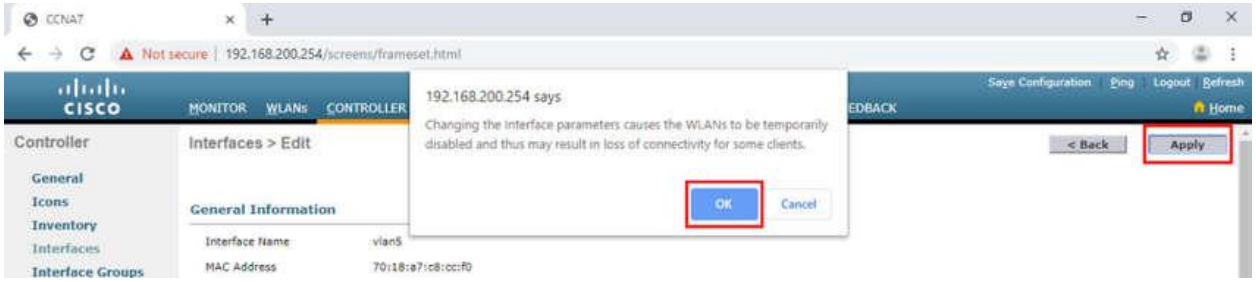

Verify Interfaces

Click **Interfaces**. The new **vlan5** interface is now shown in the list of interfaces with its IPv4 address

Configure a DHCP Scope

1. Create a new DHCP scope
2. Name the DHCP scope
3. Verify the new DHCP scope
4. Configure and enable the new DHCP scope
5. Verify the enable DHCP scope

Create a new DHCP scope

Very similar to a DHCP pool on a router

- Can include a variety of information including a pool of addresses to assign to DHCP clients, DNS server information, lease times, and more

To configure a new DHCP scope, click **Internal DHCP Server > DHCP Scope > New…**

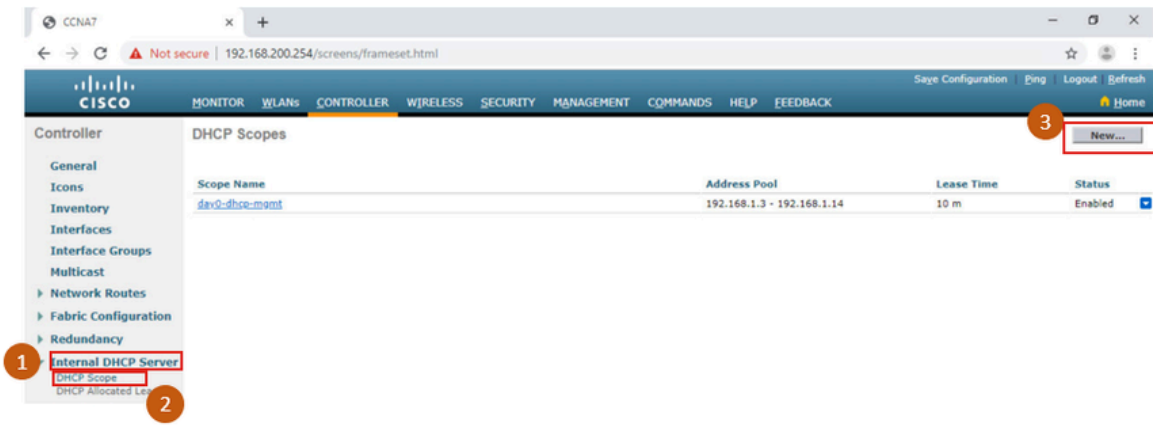

Name the DHCP scope

On the next screen, name the scope. Because this scope will apply to the wireless management network, the network administrator uses **Wireless_Management** as the Scope Name and clicks **Apply**

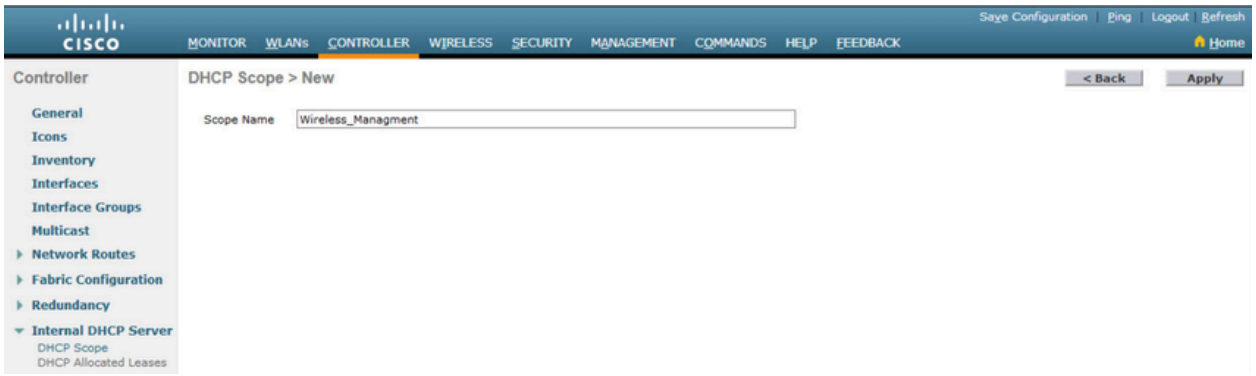

Verify the new DHCP scope

You are returned to the **DHCP Scopes** page and can verify the scope is ready to be configured. Click the new Scope Name to configure the DHCP scope

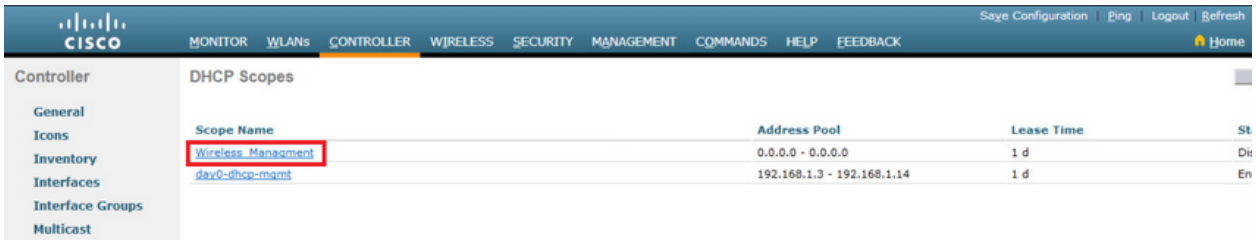

Configure and enable the new DHCP scope

On the edit screen for **Wireless_Management**, configure a pool of addresses for the 192.168.200.0/24 network starting at .240 and ending at .249. The network address and subnet mask are configured.

The default router IPv4 address is configured, which is the subinterface for R1 at 192.168.200.1

Example below: the rest of the scope is left unchanged. The network admin selects **Enabled** from the Status drop down and clicks **Apply**

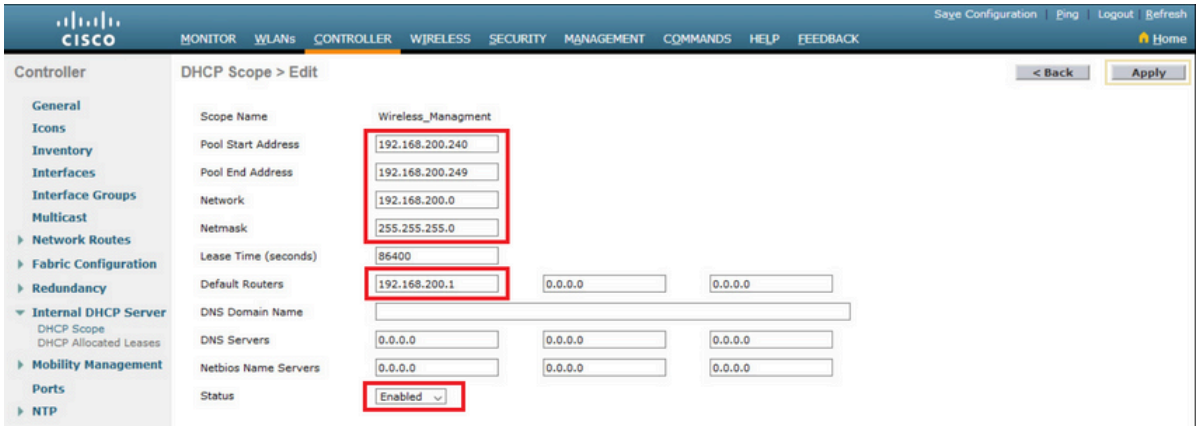

Verify the enable DHCP scope
The network admin is returned to the **DHCP Scopes** page and can verify the scope is ready to be allocated to a new WLAN

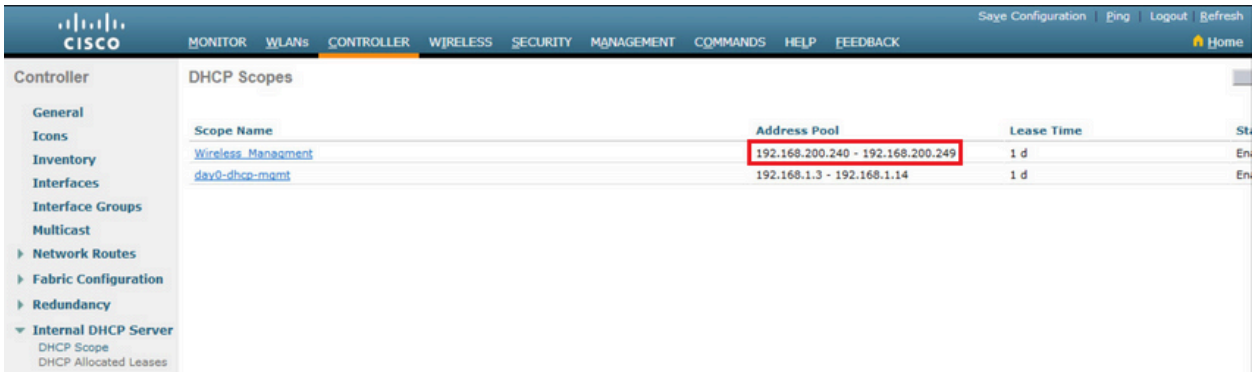

Configure a WPA2 Enterprise WLAN

By default, all newly created WLANs on the WLC will use WPA2 with Advanced Encryption System (AES). 802.1X is the default key management protocol used to communicate with the RADIUS server. Because the network admin already configured the WLC with the IPv4 address of the RADIUS server running on PC-A, the only configuration left to do is to create a new WLAN to use interface **vlan5**

Configuring a new WLAN on the WLC:

1. Create a new WLAN
2. Configure the WLAN name and SSID
3. Enable the WLAN for VLAN 5
4. Verify AES and 802.1X defaults
5. Configure WLAN security to use the RADIUS server
6. Verify the new WLAN is available

Create a new WLAN
Click the **WLANs** tab and then **Go** to create a new WLAN

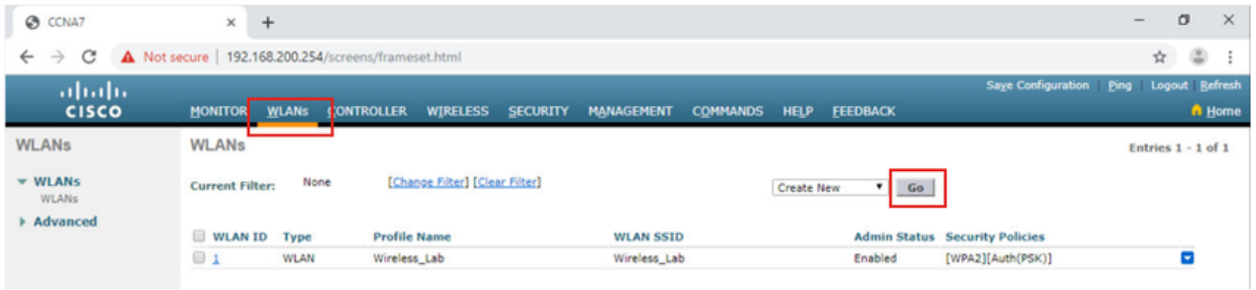

Configure the WLAN name and SSID

Fill in the profile name and SSID. In order to be consistent with the VLAN that was previously configured, choose an ID of **5**.

- Any available value can be used

Click **Apply** to create the WLAN

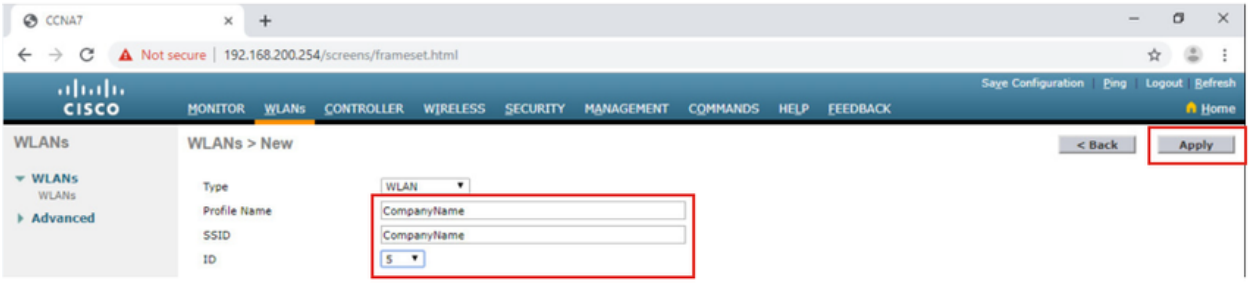

Enable the WLAN for VLAN 5

The WLAN is created but it still needs to be enabled and associated with the correct VLAN interface.

1. Change the status to **Enabled** and choose **vlan5** from the Interface/Interface Group(G) dropdown list
2. Click **Apply** and click **OK** to accept the popup message

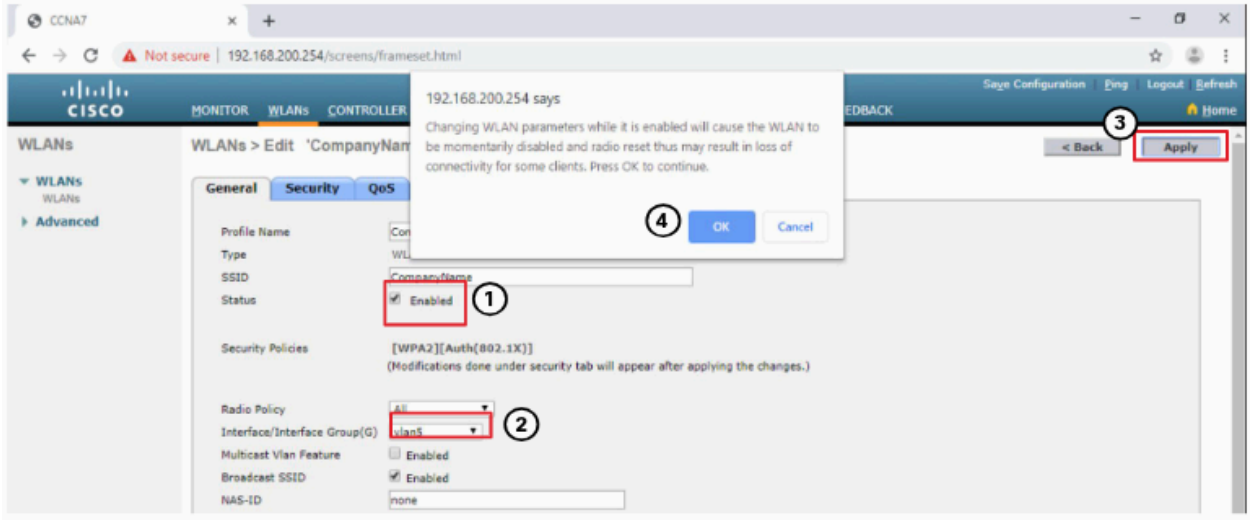

1. *Click **Enabled***
2. *Choose **vlan5***
3. *Click **Apply***
4. *Click **OK***

Verify AES and 802.1X defaults

Click the **Security** tab to view the default security configuration for the new WLAN.

- The WLAN will use WPA2 security with AES encryption.
- Authentication traffic is handled by 802.1X between the WLC and the RADIUS server

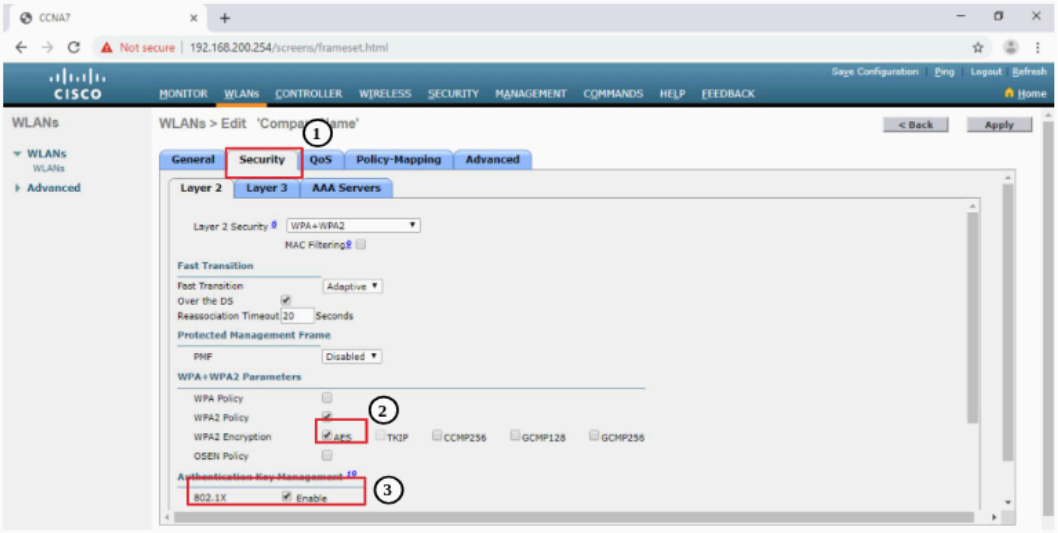

1. *Click **Security** tab*
2. *Verify **AES** is the encryption*
3. *Verify **802.1X** is the authentication key management protocol*

Configure the RADIUS server

We need to select the RADIUS server that will be used to authenticate users for this WLAN.

Click the **AAA Servers** tab

- In the dropdown box select the RADIUS server that was configured on the WLC previously.
- Apply your changes

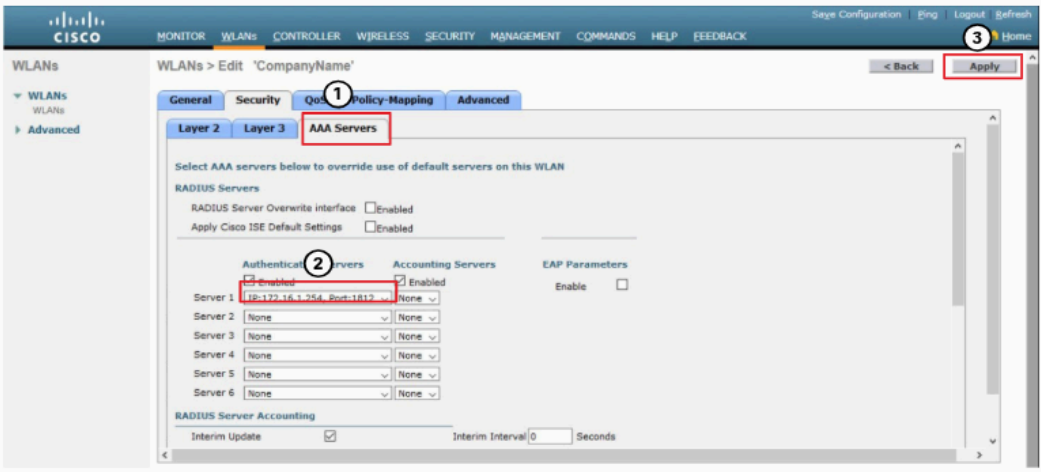

1. *Click the **AAA Servers** tab*
2. *Select the RADIUS server with the IP address **182.16.1.254** from the drop-down list next to Server 1*
3. ***Apply** your changes*

Verify that the new WLAN is available

To verify:

1. Click **Back** or the **WLANs** submenu on the left
    a. Both the **Wireless_LAN** WLAN and the **CompanyName** WLAN are listed.

In figure below: notice that both are enabled–

- **Wireless_LAN** is using WPA2 with PSK authentication
- **CompanyName** is using WPA2 security with 802.1X authentication

Troubleshoot WLAN Issues

Troubleshooting any sort of network problem should follow a systematic approach.

Step	Title	Description
1	Identify the Problem	The first step in the troubleshooting process is to identify the problem. While tools can be used in this step, a conversation with the user is often very helpful
2	Establish a Theory of Probable Causes	After you have talked to the user and identified the problem, you can try and establish a theory of probable causes. This step often yields more than a few probably causes to the problem
3	Test the Theory to Determine Cause	Based on the probable causes, test your theories to determine which one is the cause of the problem. A technician will often apply a quick procedure to test and see if it solves the problem. If a quick procedure does not correct the problem, you might need to research the problem further to establish the exact cause.
4	Establish a Plan of Action to Resolve the Problem and Implement the Solution	After you have determined the exactly cause of the problem, establish a plan of action to resolve the problem and implement the solution

5	Verify Full System Functionality and Implement Preventive Measures	After you have corrected the problem, verify full functionality and, if applicable, implement preventive measures
6	Document Findings, Actions, and Outcomes	In the final step of the troubleshoot process, document your findings, actions, and outcomes. This is very important for future reference

To assess the problem, determine how many devices on the network are experiencing the problem.

- If there is a problem with one device on the network, start the troubleshooting process at that device.
- If there is a problem with all devices on the network, start the troubleshooting process at the device where all other devices are connected

You should develop a logical and consistent method for diagnosing network problems by eliminating one problem at a time.

Wireless Client Not Connecting

When troubleshooting a WLAN, a process of elimination is recommended

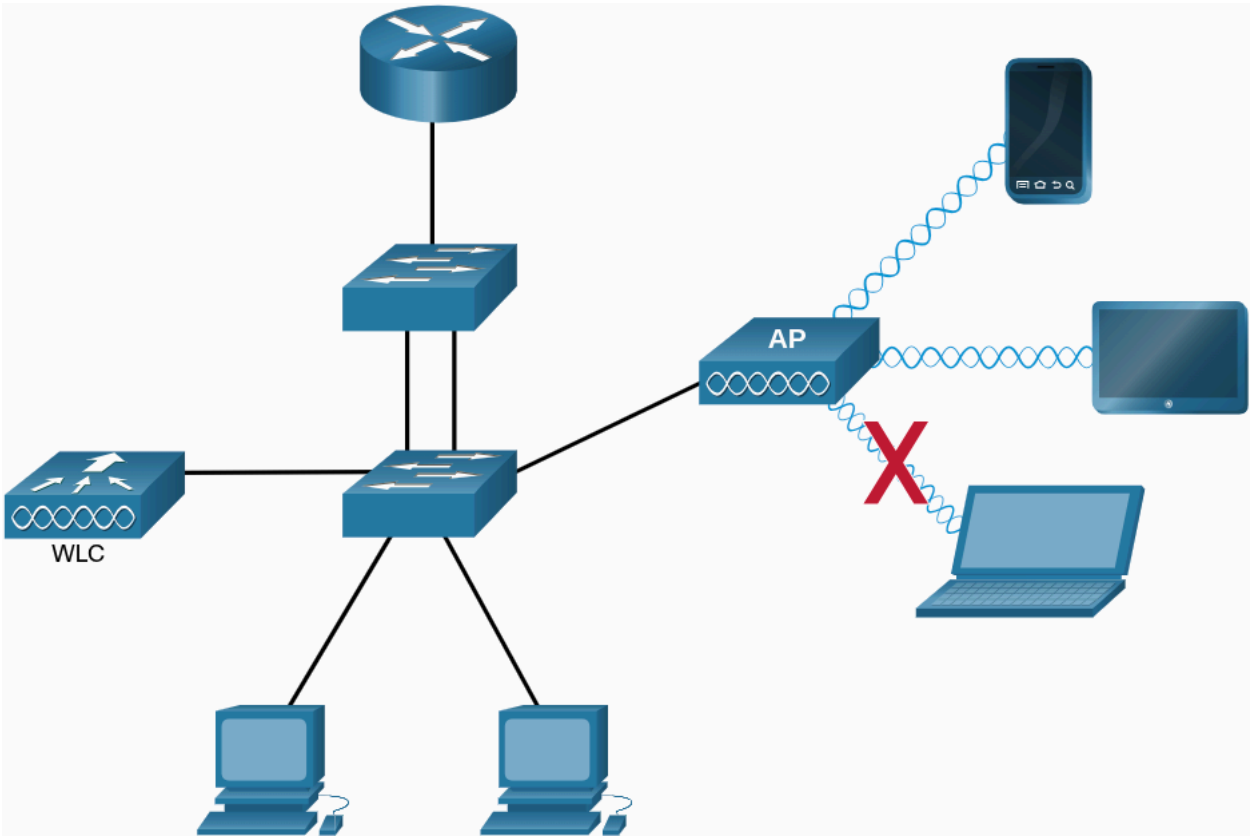

If there is no connectivity:

- Confirm the network configuration on the PC using the **ipconfig** command. Verify that the PC has received an IP address via DHCP or is configured with a static IP address
- Confirm that the device can connect to the wired network. Connect the device to the wired LAN and ping a known IP address
- If necessary, reload drivers as appropriate for the client. It may be necessary to try a different wireless NIC
- If the wireless NIC of the client is working, check the security mode and encryption settings on the client. If the security settings do not match, the client cannot gain access to the WLAN

If the PC is operational but the wireless connection is performing poorly:

- How far is the PC from the AP? Is the PC out of the planned cover area (BSA)?
- Check the channel settings on the wireless client. The client software should detect the appropriate channel as long as the SSID is correct
- Check for presence of other devices in the area that may be interfering with the 2.4Ghz band.
    - Examples: cordless phones, baby monitors, microwaves, wireless security systems, potentially rogue APs.
    - Data from these devices can cause interference in the WLAN and intermittent connection problems between a wireless client and AP

Next, ensure that all the devices are actually in place. Consider a possible physical security issue. Is there power to all devices and are they powered on?

Finally, inspect links between cabled devices looking for bad connectors or damaged or missing cables. If the physical plant is in place, verify the wired LAN by pinging devices, including the AP. If connectivity still fails at this point, perhaps something is wrong with the AP or its configuration.

When the user PC is eliminated as the source of the problem, and the physical status of devices is confirmed, begin investigating the performance of the AP. Check the power status of the AP.

Troubleshooting When the Network is Slow

To optimize and increase the bandwidth of 802.11 dual-band routers and APs

- **Upgrade your wireless clients** - Older 802.11b, 802.11g, and even 802.11n devices can slow the entire WLAN
    - For best performance, all wireless devices should supper the same highest acceptable standard
    - Although 802.11ax was released in 2019, 802.11ac is most likely the highest standard that enterprises can currently enforce
- **Split the traffic** - The easiest way to improve wireless performance is to split the wireless traffic between the 802.11n 2.4 Ghz band and the 5 Ghz band. Therefore,

802.11n (or better) can use the two bands as two separate wireless networks to help manage the traffic.

- For example:
  - Use 2.4 Ghz for basic internet tasks
  - Use 5 Ghz for streaming multimedia

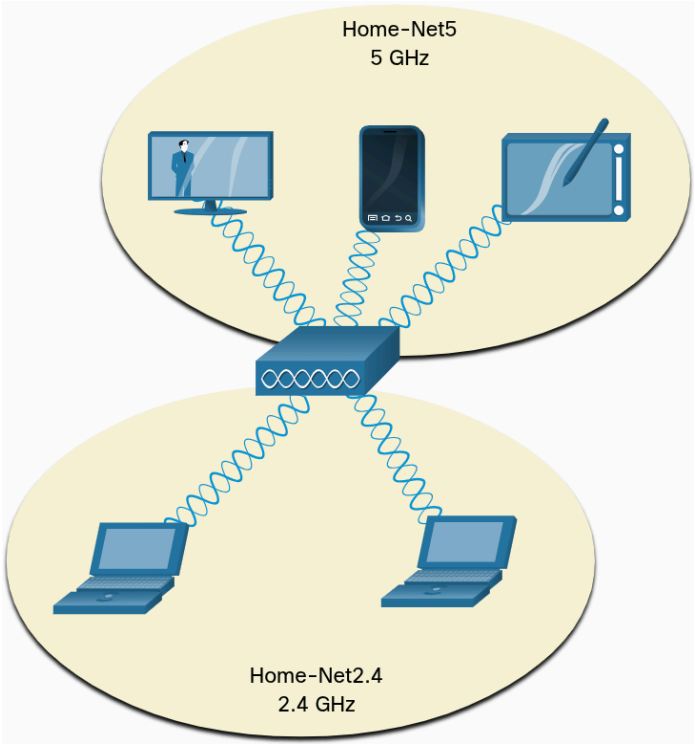

Reasons to use split-the-traffic approach

- The 2.4 Ghz band may be suitable for basic Internet traffic that is not time-sensitive
- The bandwidth may still be shared with other nearby WLANs
- The 5 Ghz band is much less crowded than the 2.4 Ghz band; ideal for streaming media
- The 5 Ghz band has more channels; therefore, the channel chosen is likely interference-free

By default, dual-band routers and APs use the same network name on both the 2.4 Ghz band and the 5 Ghz band. The simplest way to segment traffic is to rename one of the wireless networks. With a separate, descriptive name, it is easier to connect to the right network.

To improve the range of a wireless network, ensure the wireless router or AP location is free of obstructions, such as furniture, fixtures, and tall appliances. These block the signal, which shortens the range of the WLAN.

- If this still does not solve the problem, then w Wi-Fi Range Extender or deploying the Powerline wireless technology may be used

Updating Firmware

Firmware releases may contain fixes for common problems reported by customers as well as security vulnerabilities.

- Periodically check the router and AP for updated firmware

Figure below: Network admin is verifying that the firmware is up to date on a Cisco Meraki AP

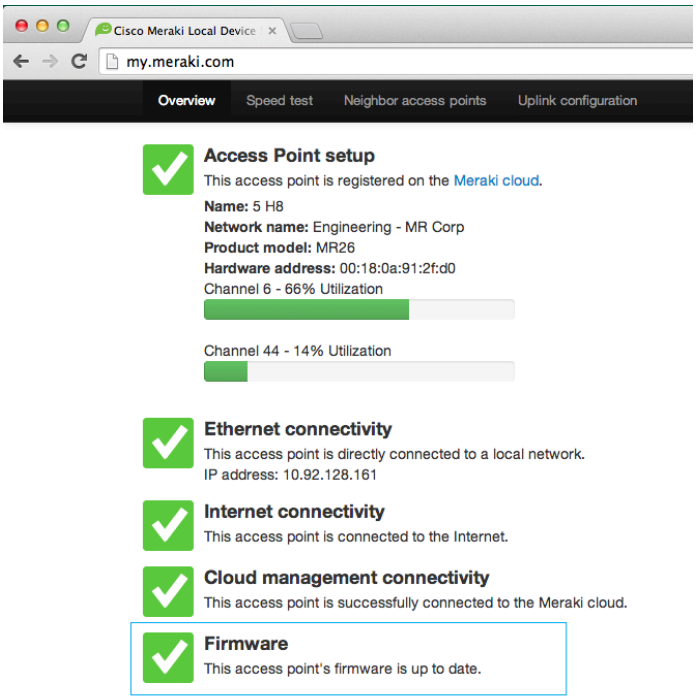

On the WLC, there will most likely be the ability to upgrade the firmware on all APs that the WLC controls.

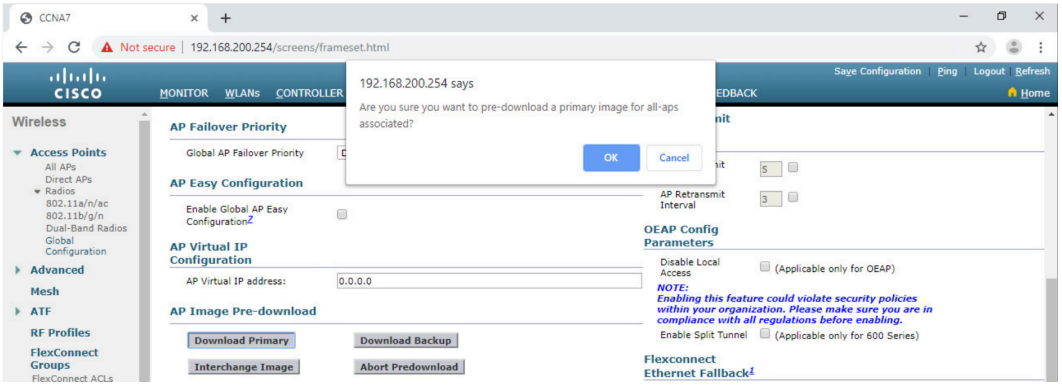

The network admin is downloading the firmware image that will be used to upgrade all the APs

On Cisco 3504 Wireless Controller, Click the **WIRELESS** tab > **Access Points** from the left menu > **Global Configuration** submenu. Then scroll to the bottom of the page for the AP Image Pre-download section

Users will be disconnected from the WLAN and the internet until the upgrade finishes. The wireless router may need to reboot several times before normal network operations are restored

13.5.3 What did I learn in this module?

Remote workers, small branch offices, and home networks often use a wireless router, which typically include a switch for wired clients, a port for an internet connection (sometimes labeled “WAN”), and wireless components for wireless client access. Most wireless routers are preconfigured to be connected to the network and provide services. The wireless router uses DHCP to automatically provide addressing information to connected devices. Your first priority should be to change the username and password of your wireless router. Use your router’s interface to complete basic network and wireless setup. If you want to extend the range beyond approximately 45 meters indoors and 90 meters outdoors, you can add wireless access points. The router will use a process called Network Address Translation (NAT) to convert private IPv4 addresses to Internet-routable IPv4 addresses. By configuring QoS, you can guarantee that certain traffic types, such as voice and video, are prioritized over traffic that is not as time-sensitive, such as email and web browsing.

Lightweight APs (LAPs) use the Lightweight Access Point Protocol (LWAPP) to communicate with a WLAN controller (WLC). Configuring a wireless LAN controller (WLC) is similar to configuring a wireless router except that a WLC controls APs and provides more services and management capabilities. Use the WLC interface to view an overall picture of the AP’s system information and performance, to access advanced settings and to configure a WLAN.

SNMP is used monitor the network. The WLC is set to forward all SNMP log messages, called traps, to the SNMP server. For WLAN user authentication, a RADIUS server is used for authentication, accounting, and auditing (AAA) services. Individual user access can be tracked and audited. Use the WLC interface to configure SNMP server and RADIUS server information, VLAN interfaces, DHCP scope, and a WPA2 Enterprise WLAN.

There are six steps to the troubleshooting process. When troubleshooting a WLAN, a process of elimination is recommended. Common problems are: no connectivity and poorly performing wireless connection when the PC is operational. To optimize and increase the bandwidth of 802.11 dual-band routers and APs, either: upgrade your wireless clients or split the traffic. Most wireless routers and APs offer upgradable firmware. Firmware releases may contain fixes for common problems reported by customers as well as security vulnerabilities. You should periodically check the router or AP for updated firmware.

Routing Concepts

Two functions of Router

When a router receives an IP packet on one interface, it determines which interface to use to forward the packet to the destination. This is known as routing. The interface that the router uses to forward the packet may be the final destination, or it may be a network connected to another router that is used to reach the destination network. Each network that a router connects to typically requires a separate interface, but not always the case.

- The primary functions of a router are to determine the ebay path to forward packets based on the information in its routing table, and to forward packets towards their destination.

Example:

- The router uses its IP routing table to determine which path (route) to use to forward a packet

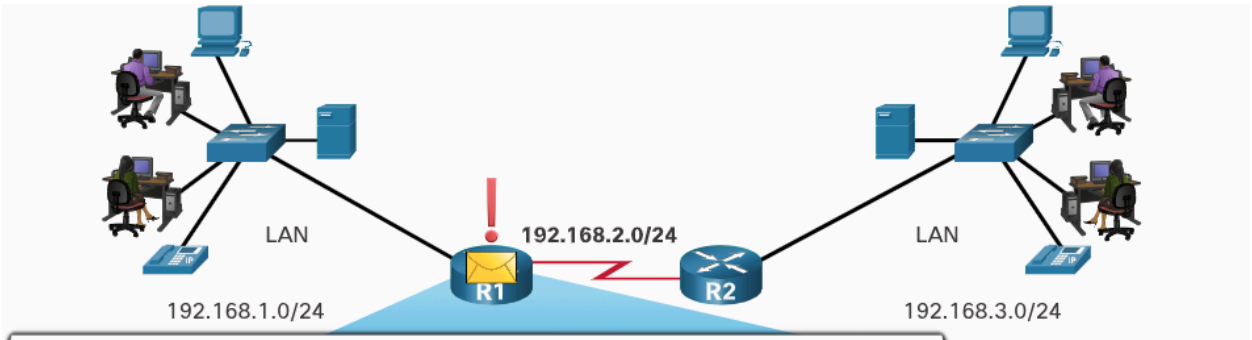


```
R1# show ip route
Codes:
C - connected, S - static, I - IGRP, R - RIP, M - mobile, B -
BGP
D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2
ia - IS-IS inter area, * - candidate default
U - per-user static route, o - ODR
P - periodic downloaded static route

Gateway of last resort is not set

C    192.168.1.0/24 is directly connected, FastEthernet0/0
C    192.168.2.0/24 is directly connected, Serial0/0/0
S    192.168.3.0/24 [1/0] via 192.168.2.2
```

Routers use the routing table like a map to discover the best path for a given network.

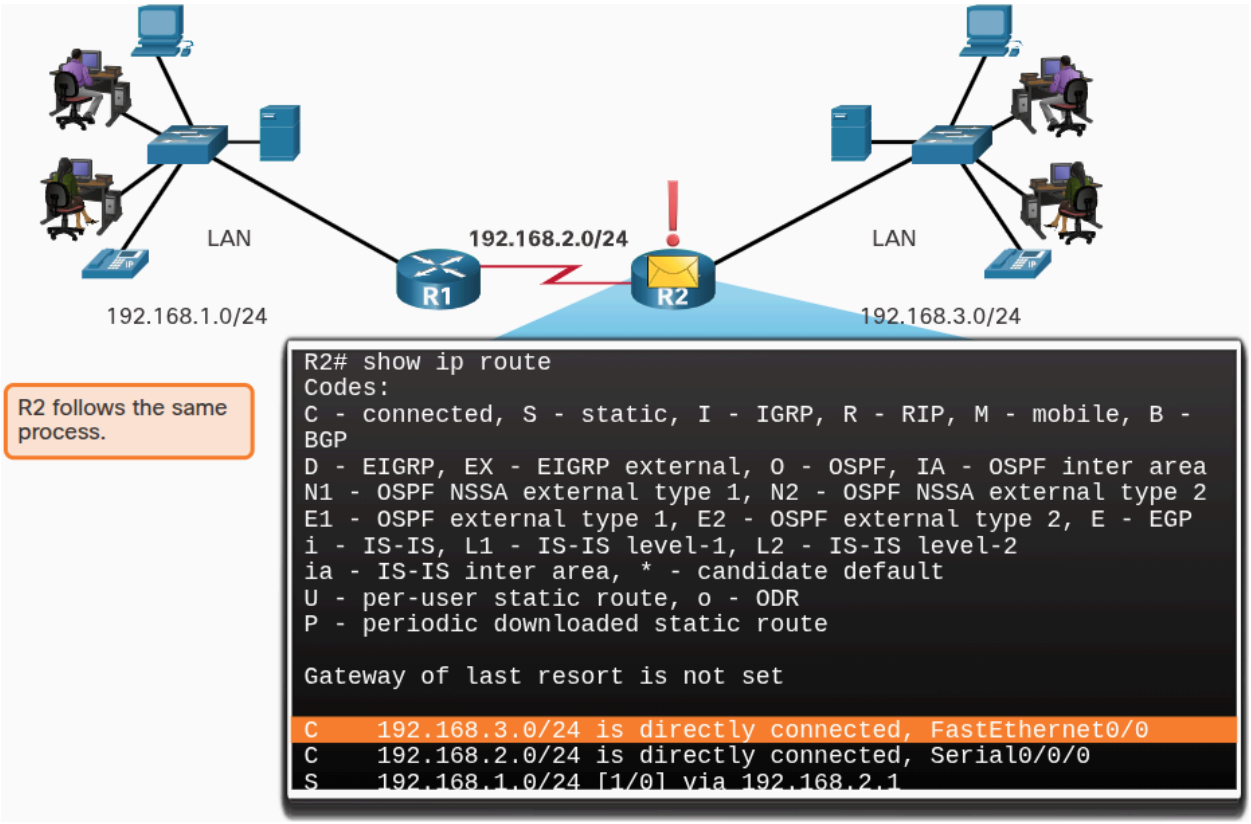

Best Path Equals Longest Match

The best path in the routing table is also known as the longest match

- The longest math is a process the router uses to find a match between the destination IP address of the packet and a routing entry in the routing table

The routing table contains route entries consisting of a prefix (network address) and prefix length of the route in the routing table is used to determine the minimum number of far-left bits that must match.

- Remember that an IP packet only contains the destination IP address and not the prefix length

The longest match is the route in the routing table that has the greatest number of far-left matching bits with the destination bits with the destination IP address of the packet. The route with the greatest number of equivalent far-left bits, or the longest match, is always the preferred route.

Note - The term prefix length will be used to refer to the network portion of both IPv4 and IPv6 address

IPv5 Address Longest Match Example

In table below: an IPv4 packet has the destination IPv4 address 172.16.0.10

- Router has three route entries that match this packet:
  - 172.16.0.0/12
  - 172.16.0.0/18
  - 172.16.0.0./26
  - 172.16.0.0/26 has the longest match and would be chosen to forward the packet
- Remember, for any of these routes to be considered a match, there must be at least the number of matching bits indicated by the subnet mask of the route

Destination IPv4 Address	**Address in Binary**
172.16.0.10	**10101100.00010000.00000000.00**001010

Route Entry	**Prefix/Prefix Length**	**Address in Binary**
1	172.16.0.0/**12**	**10101100.0001**0000.00000000.00001010
2	172.16.0.0/**18**	**10101100.00010000.00**000000.00001010
3	172.16.0.0/**26**	**10101100.00010000.00000000.00**001010

IPv6 Address Longest Match Example

In table below: An IPv6 packet has destination IPv6 address 2001:db8:c000::99

- Router has three entries, but only two are valid match, one being the longest match
  - The first entry with a prefix length of /40 matches the 40 far-left bits in the IPv6 address
  - The second entry has a prefix length of /58 and with all 48 bits matching the destination IPv6 address and is the longest match
  - The third entry is not a match because its /64 prefix requires 64 matching bits
- For prefix 2001:db8:c000:5555::/64 to be a match, the first 64 bits must be the destination IPv65 address of the packet.
  - For third entry, only the first 48 bits match, so this route entry is not considered a match

Route Entry	**Prefix/ Prefix Length**	Does it match?
1	**2001:db8:c0**00::/**40**	Match of 40 bits
2	**2001:db8:c000**::/**48**	Match of 48 bits (longest match)
3	**2001:db8:c000**:5555::/**64**	Does not match 64 bits

Build the Routing Table

- Routing table consists of prefixes and their prefix length
- How does the router learn about these networks?

Networks from the Perspective of R1

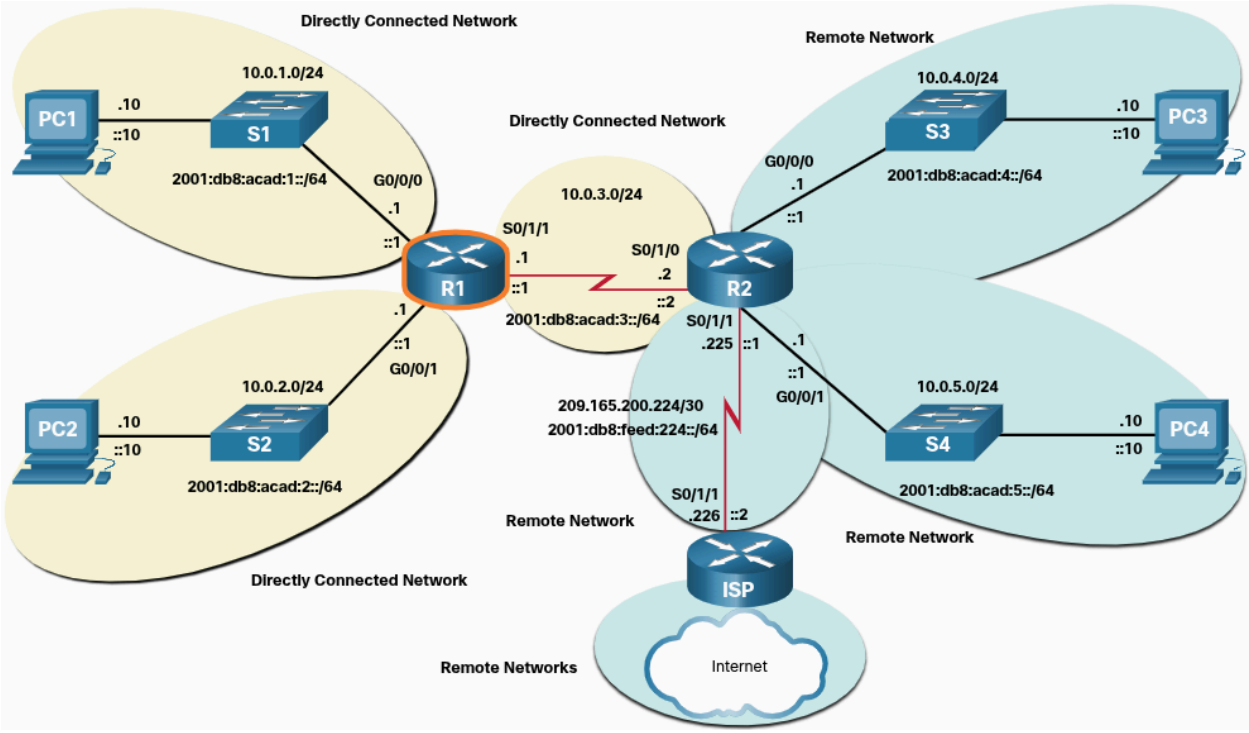

The networks in the topology are highlighted and labelled from the perspective of R1. All the IPv4 and IPv6 networks highlighted in yellow are directly connected networks. All the IPv4 and IPv6 networks highlighted in blue are remote networks

Directly Connected Networks

- Networks that are configured on the active interfaces of a router. A directly connected network is added to the routing table when an interface is configured with an IP address and subnet mask (prefix length) and is active (up and up)

Remote Networks

Networks that are not directly connected to the router. Routers learn in two ways:

- **Static routes** - Added to the routing table when a route is manually configured
- **Dynamic routing protocols** - Added to the routing table when routing protocols dynamically learn about the remote network. Dynamic routing protocols include Enhanced Interior Gate Routing Protocol (EIGRP), Open Shortest Path First (OSPF), etc

Default Route

- Specifies a next-hop router to use when the routing table does not contain a specific route that matches the destination IP address. The default route can be entered manually as a static route or learned automatically from dynamic routing protocol
- A default route over IPv4 has a route entry of 0.0.0.0/0 and a default route over IPv6 has a route entry of ::/0

- The /0 prefix length indicates that zero bits or no bits need to match the destination IP address for this route entry to be used.
    - If there are no routes with a longer match, more than 0 bits, then the default route is used to forward the packet
    - The default route is sometimes referred to as a gateway of last resort

Packet Forwarding

Decision Process

- Figure below demonstrates how a router first determines the best path, then forwards the packet

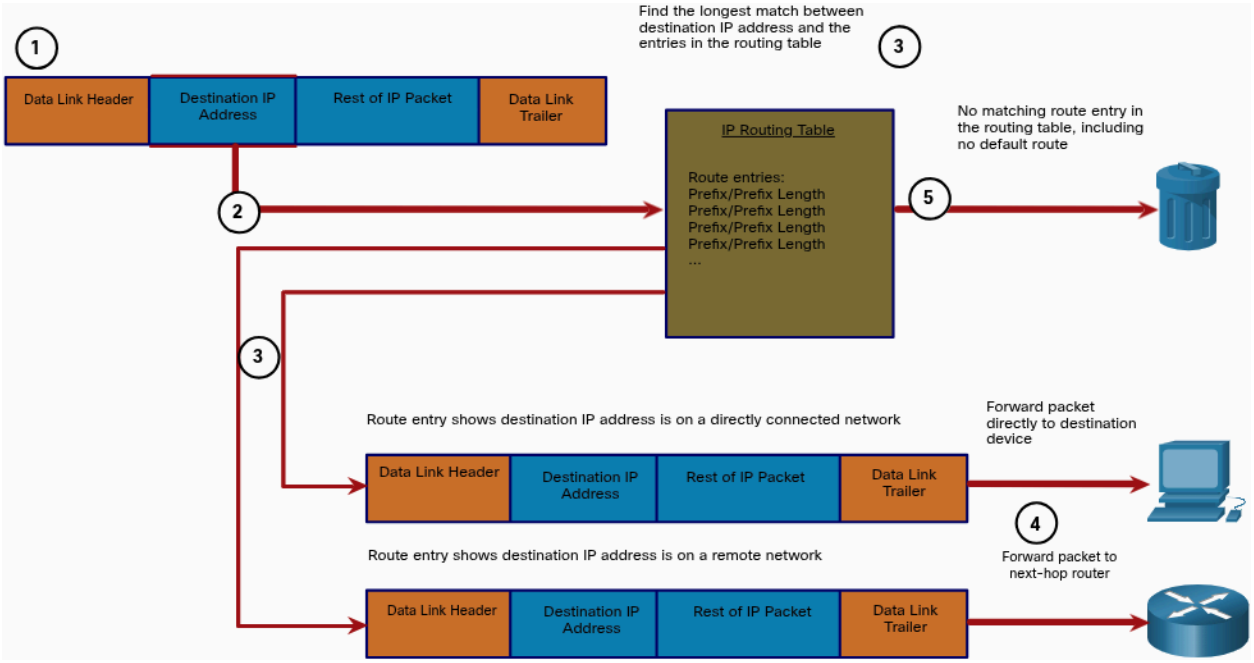

1. *The data link frame with an encapsulated IP packet arrives on the ingress interface*
2. *The router examines the destination IP address in the packet header and consults its IP routing table*
3. *The router finds the longest matching prefix in the routing table*
4. *The router encapsulates the packet in a data link frame and forwards it out the egress interface.*
    a. *The destination could be a device connected to the network or the next-hop router*
5. *However, if there is no matching route entry the packet is dropped*

Forwards the Packet to a Device on a Directly Connected Network

If the route entry indicates that the egress interface is a directly connected network, this means that the destination IP address of the packet belongs to a device on the directly connected network. Therefore, the packet can be forwarded directly to the destination device. The destination device is typically an end device on an Ethernet LAN, which means the packet must be encapsulated in an Ethernet frame.

To encapsulate the packet in the Ethernet frame, the router needs to determine the destination MAC address associated with the destination IP address of the packet. The process varies based on whether the packet is an IPv4 or IPv6 packet:

- **IPv4 packet** - The router checks its ARP table for the destination IPv4 address and an associated Ethernet MAC address
    - If there is no match, the router sends an ARP Request.
    - The destination device will return an ARP Reply with its MAC address
    - The router can now forward the IPv4 packet in an Ethernet frame with the proper destination MAC address
- **IPv6 packet** - The router checks its neighbor cache for the destination IPv6 address and an associated Ethernet MAC address.
    - If there is no match, the router send an ICMPv6 Neighbor Solicitation (NS) message
    - The destination device will return an ICMPv6 Neighbor Advertisement (NA) message with its MAC address
    - The router can now forward the IPv6 packet in an Ethernet frame with the proper destination MAC address

Forwards the Packet to a Next-Hop Router

If the route entry indicates that the destination IP address is on a remote network, this means the destination IP address of the packet belongs to a device on a network that is not directly connected. Therefore, the packet must be forwarded to another router, specifically a next-hop router. The next-hop address is indicated in the route entry.

If the forwarding router and the next-hop router are on an Ethernet network, a similar process (ARP and ICMPv6 Neighbor Discovery) will occur for determining the destination MAC address of the packet as described previously. The difference is that the router will search for the IP address of the next-hop router in its ARP table or neighbor cache, instead of the destination IP address of the packet.
Note - This process will vary for other types of Layer 2 networks

Drops the Packet - No Match in Routing Table

If there is no match between the destination IP address and a prefix in the routing table, and there is no default route, the packet will be dropped

End-to-End Packet Forwarding

The primary responsibility of the packet forwarding function is to encapsulate packets in the appropriate data link frame type for the outgoing interface

- Example: The data link frame format for a serial link could be Point-to-Point (PPP) protocol, High-Level Data Link Control (HDLC) protocol, or some other Layer 2 protocol

PC1 Sends Packet to PC2

In animation, PC1 sends a packet to PC2

- Since PC2 is on a different network, PC1 will forward the packet to its default gateway.
- PC1 will look in its ARP cache for the MAC address of the default gateway and add the indicated frame information

Note - If an ARp entry does not exist in the ARP table for the default gateway of 192.168.1.1, PC1 sends an ARP request. Router R1 would then return an ARP reply with its MAC address

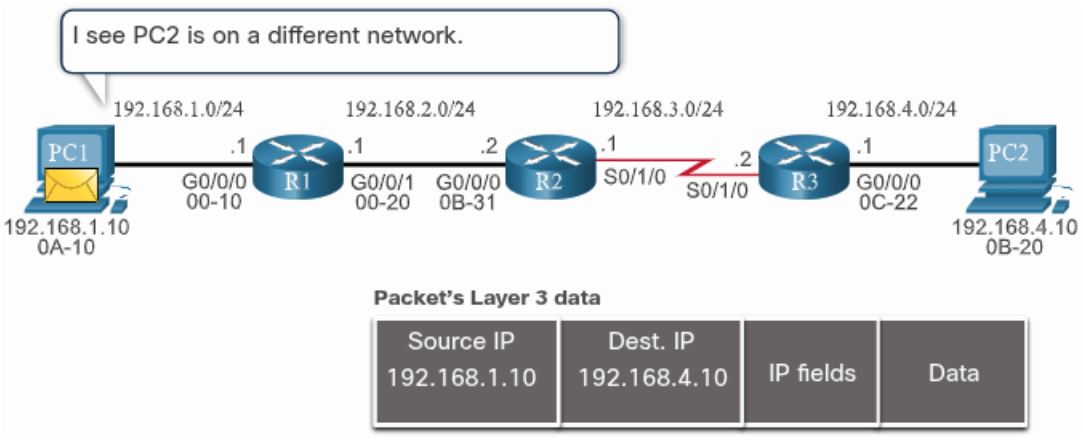

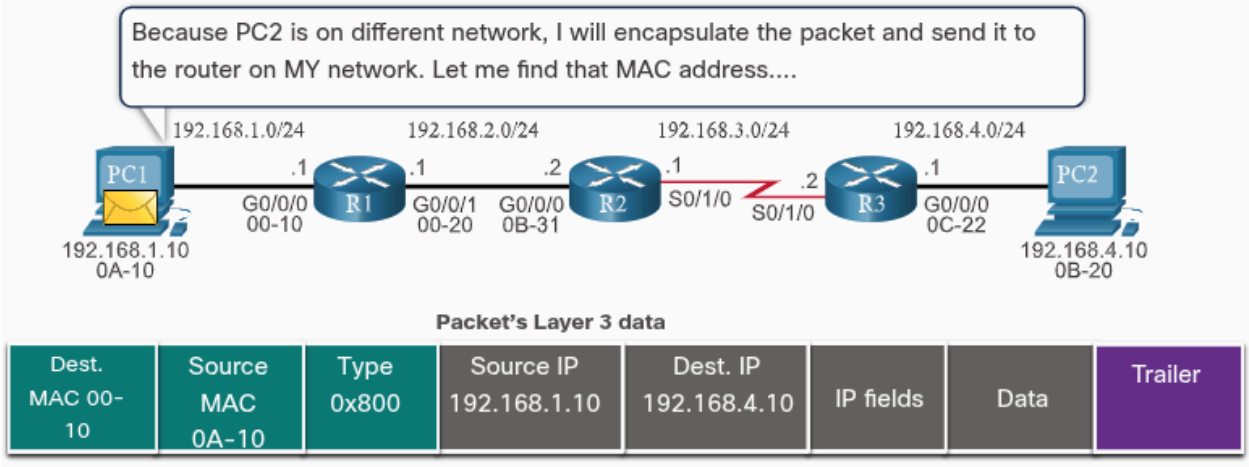

Layer 2 Data Link Frame

PC1's ARP Cache for R1	
IP Address	**MAC Address**
192.168.1.1	00-10

R1 Forwards the Packet to PC2

R1 now forwards the packet to PC2. Because the exit interface is on an Ethernet network, R1 must resolve the next-hope IPv4 address with a destination MAC address using its ARP table. If an ARP entry does not exist in the ARP table for the next-hop interface of 192.168.2.2, R1 sends an ARp request. R2 would then return an ARP Reply.

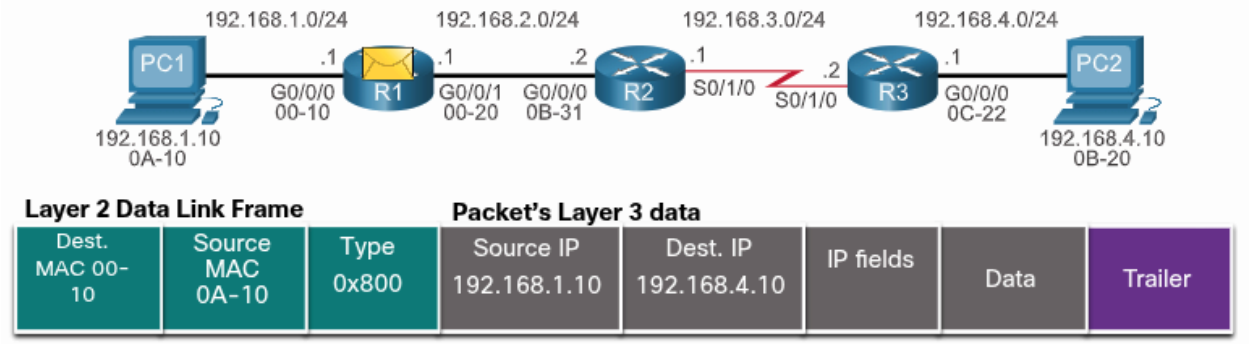

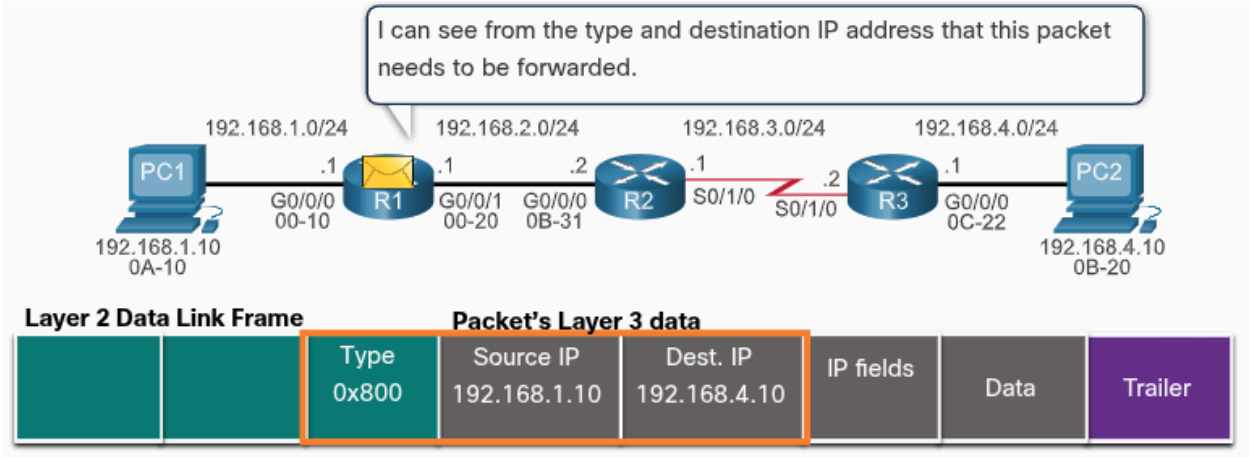

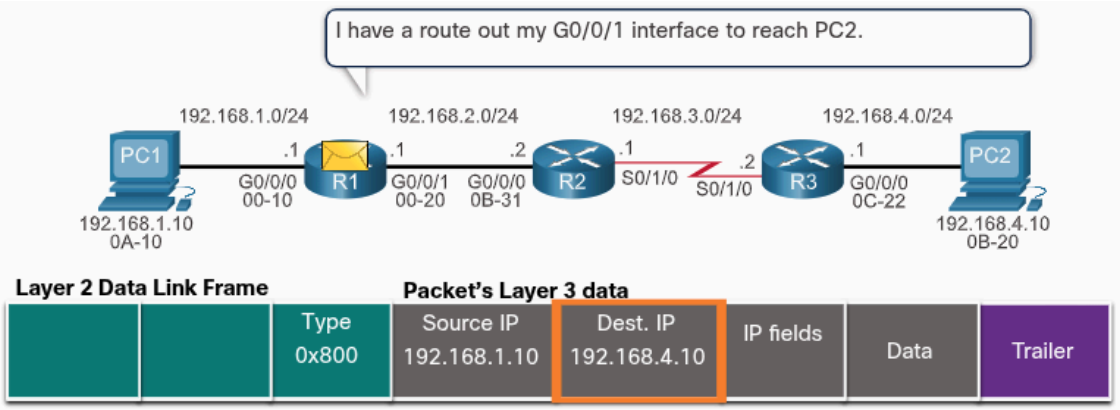

R1's Routing Table

Network	Hops	Next-hop-IP	Exit Interface
192.168.1.0/24	0	Dir. Connect.	G0/0/0
192.168.2.0/24	0	Dir. Connect.	G0/0/1
192.168.3.0/24	1	192.168.2.2	G0/0/1
192.168.4.0/24	2	192.168.2.2	G0/0/1

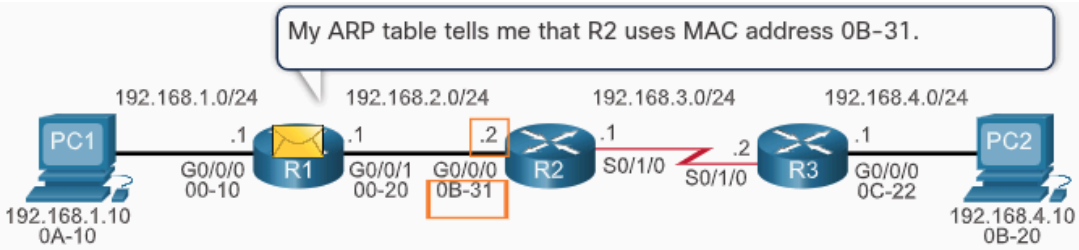

Layer 2 Data Link Frame | **Packet's Layer 3 data**

Dest. MAC 0B-31		Type 0x800	Source IP 192.168.1.10	Dest. IP 192.168.4.10	IP fields	Data	Trailer

R1's ARP Cache

IP Address	MAC Address
192.168.2.2	0B-31

R1's Routing Table

Network	Hops	Next-hop-IP	Exit Interface
192.168.1.0/24	0	Dir. Connect.	G0/0/0
192.168.2.0/24	0	Dir. Connect.	G0/0/1
192.168.3.0/24	1	192.168.2.2	G0/0/1
192.168.4.0/24	2	192.168.2.2	G0/0/1

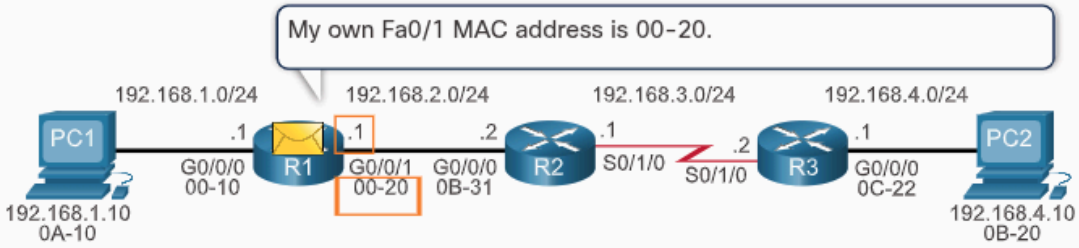

Layer 2 Data Link Frame | **Packet's Layer 3 data**

Dest. MAC 0B-31	Source MAC 00-20	Type 0x800	Source IP 192.168.1.10	Dest. IP 192.168.4.10	IP fields	Data	Trailer

R1's Routing Table

Network	Hops	Next-hop-IP	Exit Interface
192.168.1.0/24	0	Dir. Connect.	G0/0/0
192.168.2.0/24	0	Dir. Connect.	G0/0/1
192.168.3.0/24	1	192.168.2.2	G0/0/1
192.168.4.0/24	2	192.168.2.2	G0/0/1

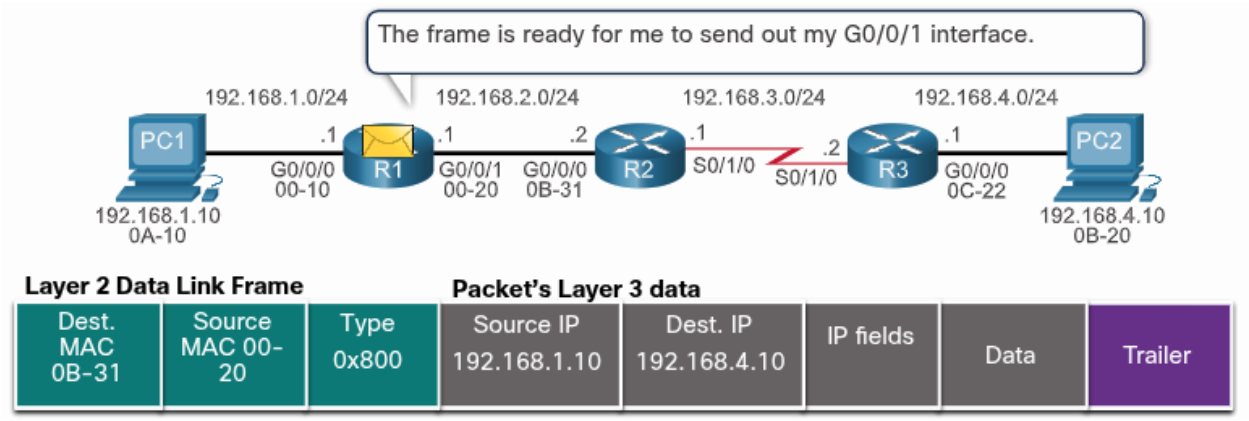

R2 Forwards the Packet to R3

R2 now forwards the packet to R3. Because the exit interface is not an Ethernet network, R2 does not have to resolve the next-hop IPv4 address with a destination MAC address. When the interface is a point-to-point (P2P) serial connection, the router encapsulates the IPv4 packet into the proper data link frame format used by the exit interface (HDLC, PPP, etc) Because there are no MAC addresses on serial interfaces, R2 sets the data link destination address to an equivalent of a broadcast

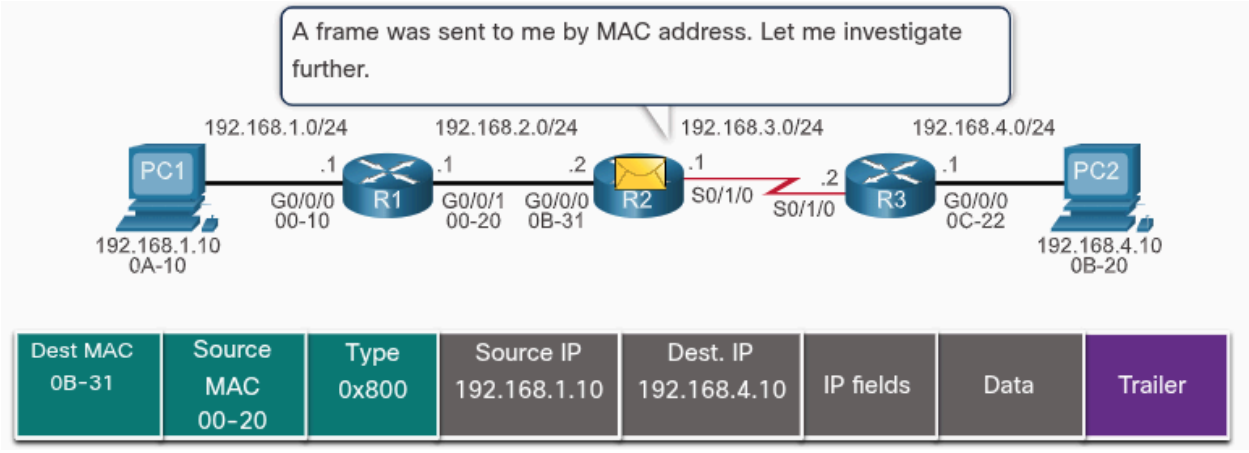

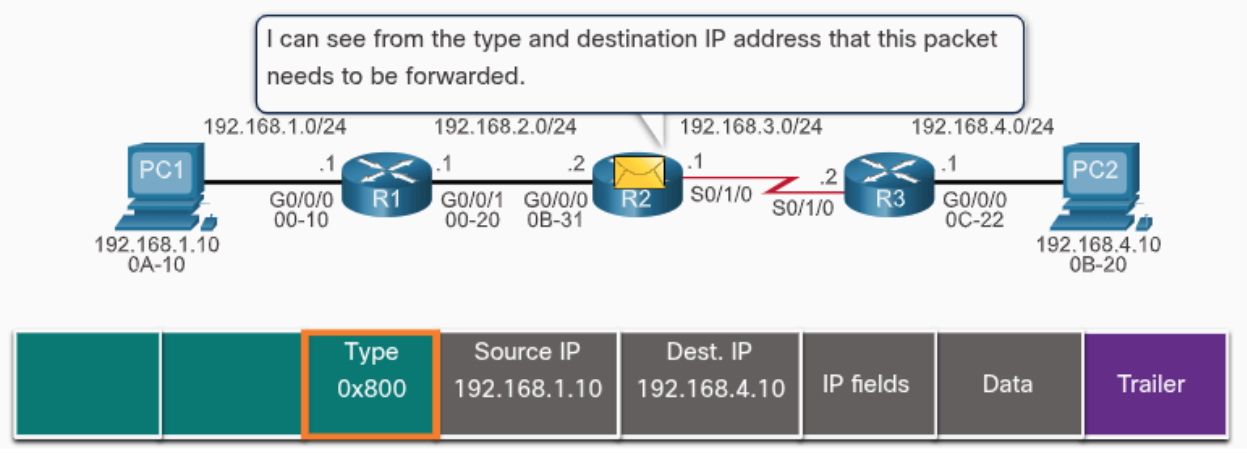

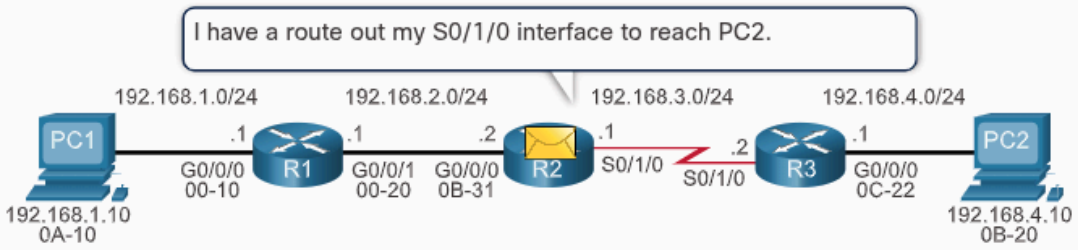

R2's Routing Table			
Network	**Hops**	**Next-hop-IP**	**Exit Interface**
192.168.1.0/24	1	192.168.2.1	G0/0/0
192.168.2.0/24	0	Dir. Connect.	G0/0/0
192.168.3.0/24	0	Dir. Connect.	S0/1/0
192.168.4.0/24	1	192.168.3.2	S0/1/0

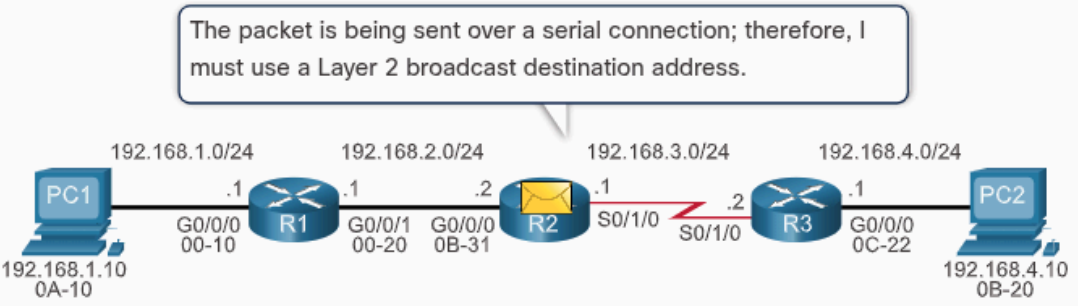

R2's Routing Table			
Network	**Hops**	**Next-hop-IP**	**Exit Interface**
192.168.1.0/24	1	192.168.2.1	G0/0/0
192.168.2.0/24	0	Dir. Connect.	G0/0/0
192.168.3.0/24	0	Dir. Connect.	S0/1/0
192.168.4.0/24	1	192.168.3.2	S0/1/0

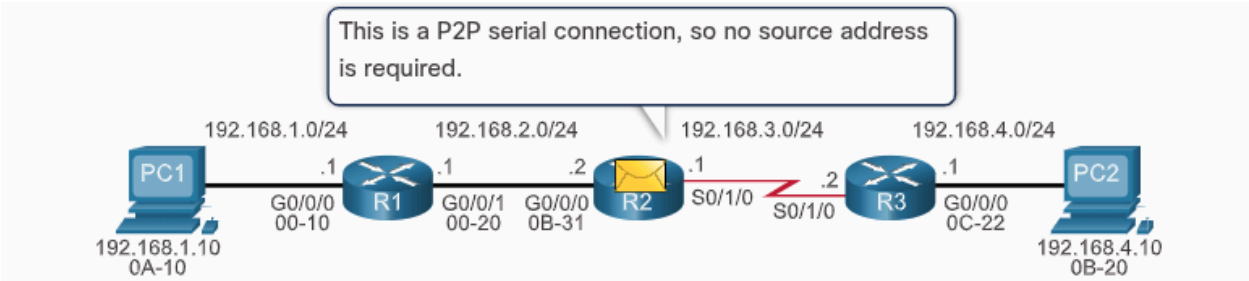

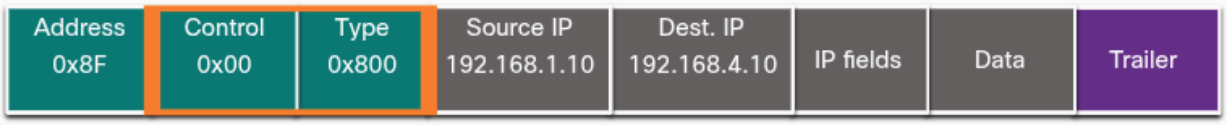

R2's Routing Table

Network	Hops	Next-hop-IP	Exit Interface
192.168.1.0/24	1	192.168.2.1	G0/0/0
192.168.2.0/24	0	Dir. Connect.	G0/0/0
192.168.3.0/24	0	Dir. Connect.	S0/1/0
192.168.4.0/24	1	192.168.3.2	S0/1/0

R3 Forwards the Packet to PC2

R3 now forwards the packet to PC2. Because the destination IPv4 address is on a directly connected Ethernet network, R3 must resolve the destination IPv4 address of the packet with its associated MAC address.

- If the entry is not the ARP table, R3 sends an ARP request out of its F0/0 interface.
- PC2 would then return an ARP reply with its MAC address

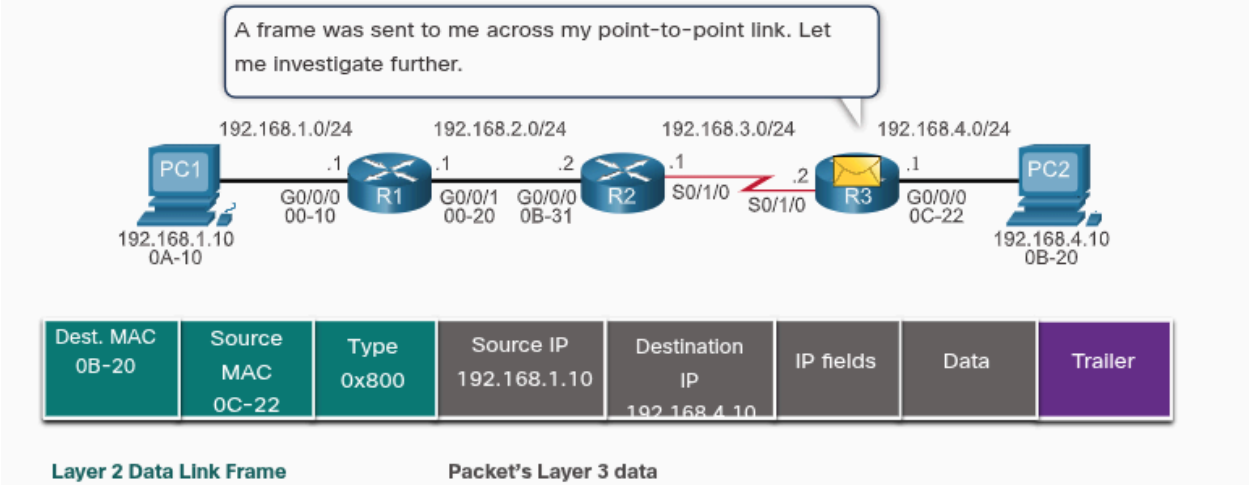

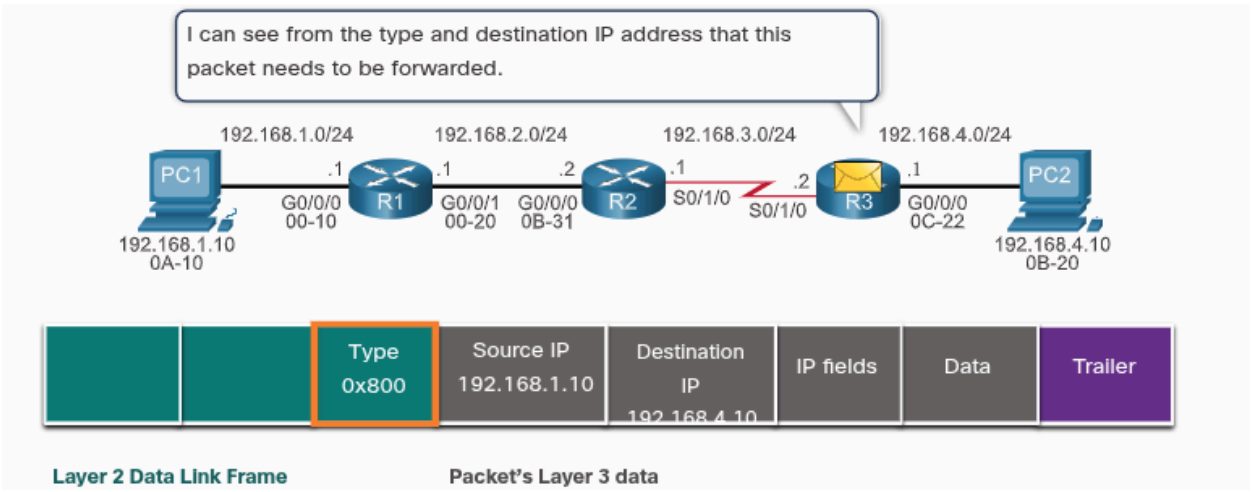

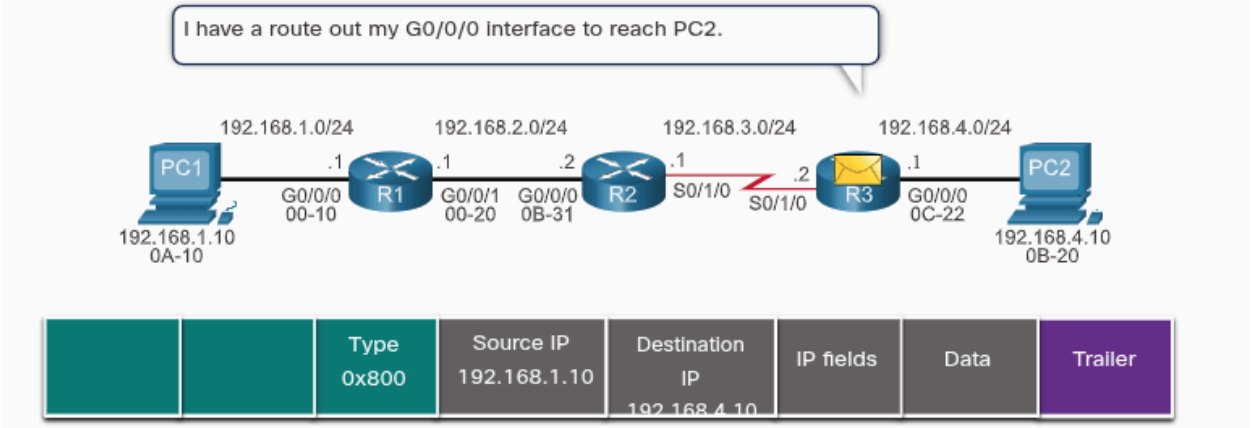

R3's Routing Table

Network	Hops	Next-Hop-IP	Exit Interface
192.168.1.0/24	2	192.168.3.1	S0/1/0
192.168.2.0/24	1	192.168.3.1	S0/1/0
192.168.3.0/24	0	Direct Connect	S0/1/0
192.168.4.0/24	0	Direct Connect	G0/0/0

My ARP table tells me that PC2 uses MAC address 0B-20.

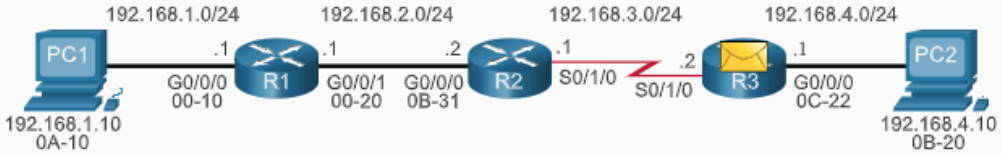

Dest. MAC 0B-20	Source MAC 0C-22	Type 0x800	Source IP 192.168.1.10	Destination IP 192.168.4.10	IP fields	Data	Trailer

R3's ARP Cache

IP Address	MAC Address
192.168.4.10	0B-20

R3's Routing Table

Network	Hops	Next-Hop-IP	Exit Interface
192.168.1.0/24	2	192.168.3.1	S0/1/0
192.168.2.0/24	1	192.168.3.1	S0/1/0
192.168.3.0/24	0	Direct Connect	S0/1/0
192.168.4.0/24	0	Direct Connect	G0/0/0

Oh look, a frame is being sent to my MAC address, let me process it. The packet also matches my IP address, so it MUST be mine.

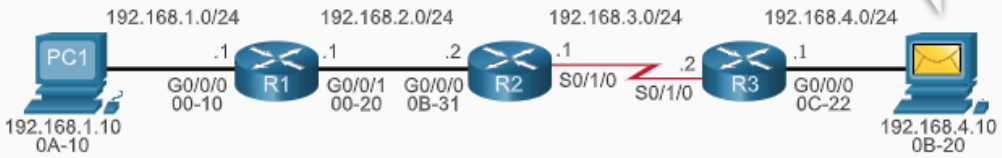

Dest. MAC 0B-20	Source MAC 0C-22	Type 0x800	Source IP 192.168.1.10	Destination IP 192.168.4.10	IP fields	Data	Trailer

Layer 2 Data Link Frame **Packet's Layer 3 data**

Packet Forwarding Mechanisms

The primary responsibility of the packet forwarding function is to encapsulate packets in the appropriate data link frame type for the outgoing interface. The more efficiently a router can perform this task, the faster packets can be forwarded by the router.

Routers support the following three packet forwarding mechanisms:

- Process switching
- Fast switching
- Cisco Express Forwarding (CEF)

Process Switching

An older packet forwarding mechanism still available for Cisco routers.
When a packet arrives on an interface, it is forwarded to the control plane where the CPU matches the destination address with an entry in its routing table, and then determines the exit interface and forwards the packet.

- It is important to understand that the router does this for every packet, even if the destination is the same for a stream of packets
- This process-switching mechanism is very slow and is rarely implemented in modern networks. Contrast this with fast switching.

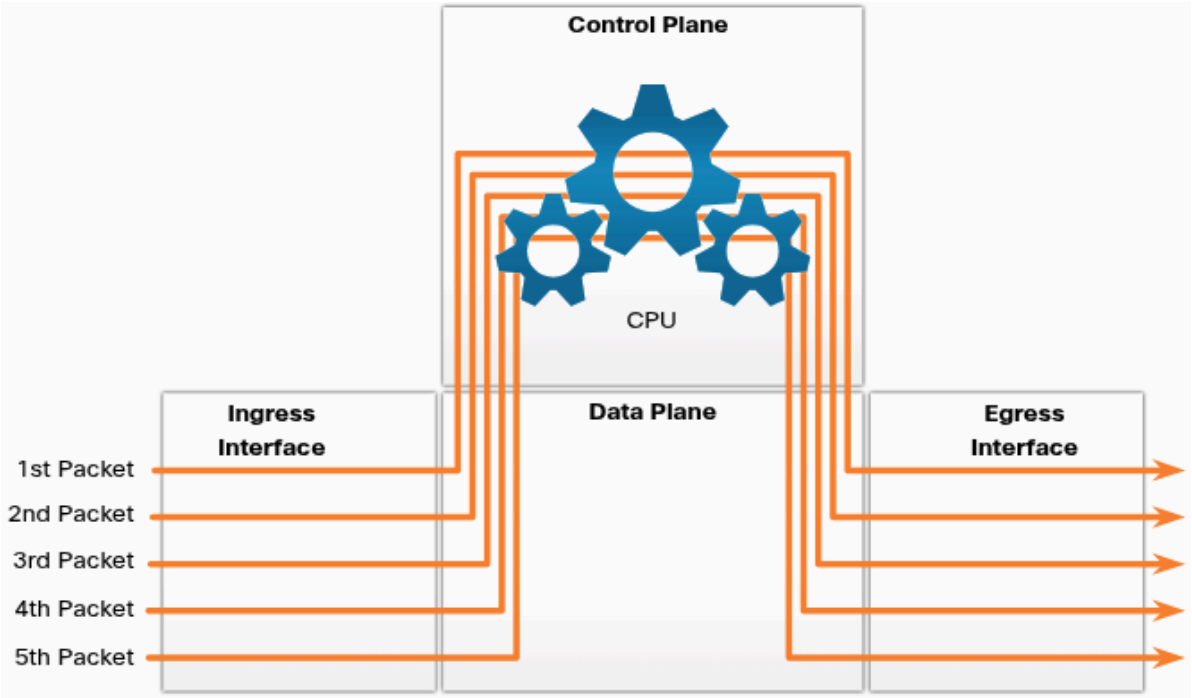

Fast Switching
Another, older packet forwarding mechanism which was the successor to process switching.

- Uses a fast-switching cache to store next-hop information.
- When a packet arrives on an interface, it is forwarded to the control plane where the CPU searches for a match in the fast-switching cache
    - If it is not there, it is process-switched and forwarded to the exit interface.
- The flow information for the packet is also stored in the fast-switching cache
    - If another packet going to the same destination arrives on an interface, the next-hop information in the cache is re-used without CPU intervention

With fast switching, notice how only the first packet of a flow is process-switched and added to the fast-switching cache. The next four packets are quickly processing based on the information in the fast-switching cache

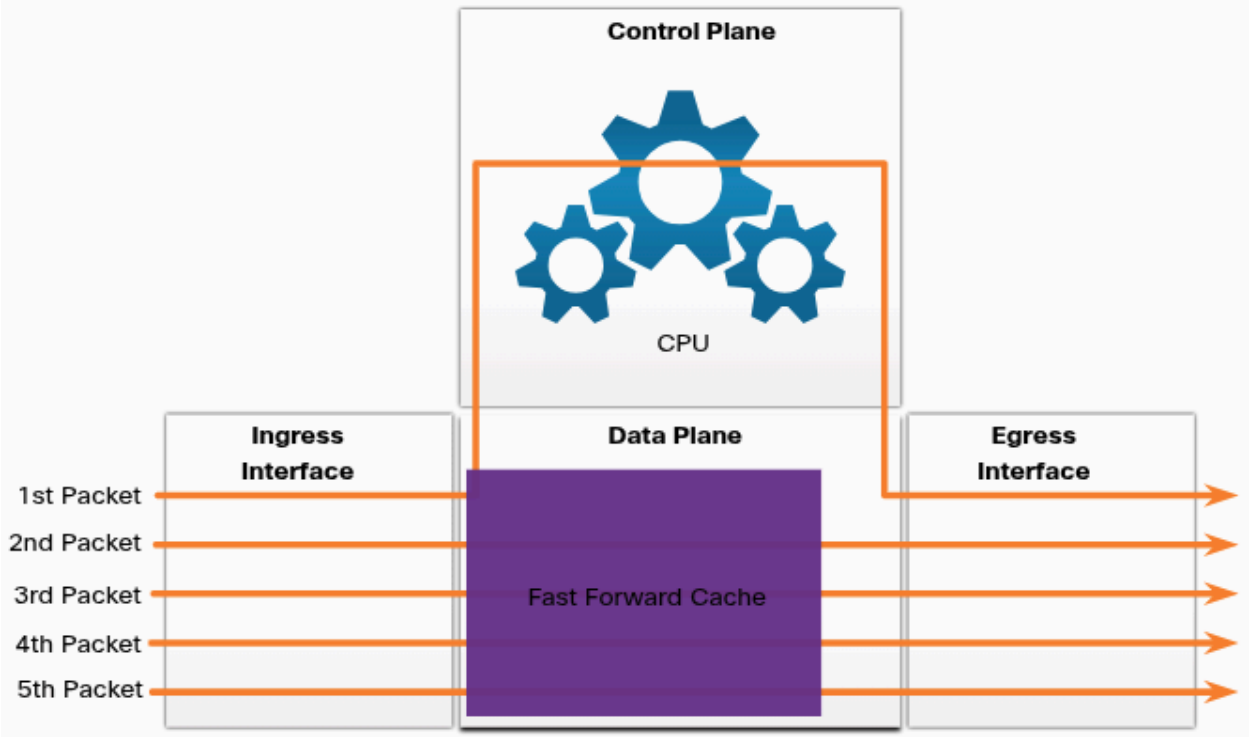

Cisco Express Forwarding (CEF)
The most recent and default Cisco IOS packet-forwarding mechanism.

- Like fast switching, CEF builds a Forwarding Information Base (FIB), and an adjacent table

However, the table entries are not packet-triggered like fast switching but change-triggered, such as when something changes in the network topology. Therefore, when a network has converged, the FIB and adjacency tables contain all the information that a router would have to consider when forwarding a packet

Cisco Express Forwarding is the fastest forwarding mechanism and the default on Cisco routers and multilayer switches.

CEF builds the FIB and adjacency tables after the network has converged. All five packets are quickly processed in the data plane.

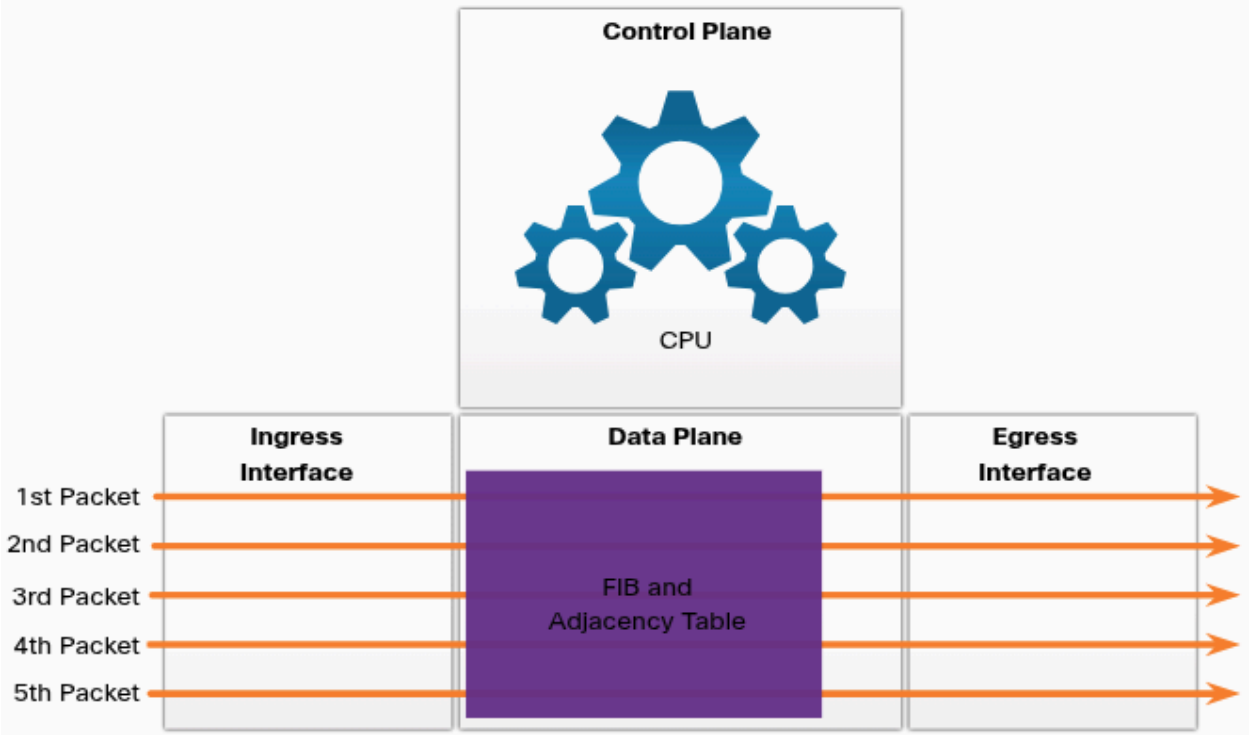

A common analogy used to describe these three packet-forwarding mechanisms:

- Process switching solves a problem by doing math long hand, even if it is the identical problem that was just solved
- Fast switching solves a problem by doing math long hand one time and remembering the answer for subsequent identical problems
- CEF solves every possible problem ahead of time in a spreadsheet

Basic Router Configuration Review

Topology

A router creates a routing table to help it determine whether to forward packets

- Topology in figure will be used for configuration and verification examples
- Will be used in next topic to discuss IP routing table

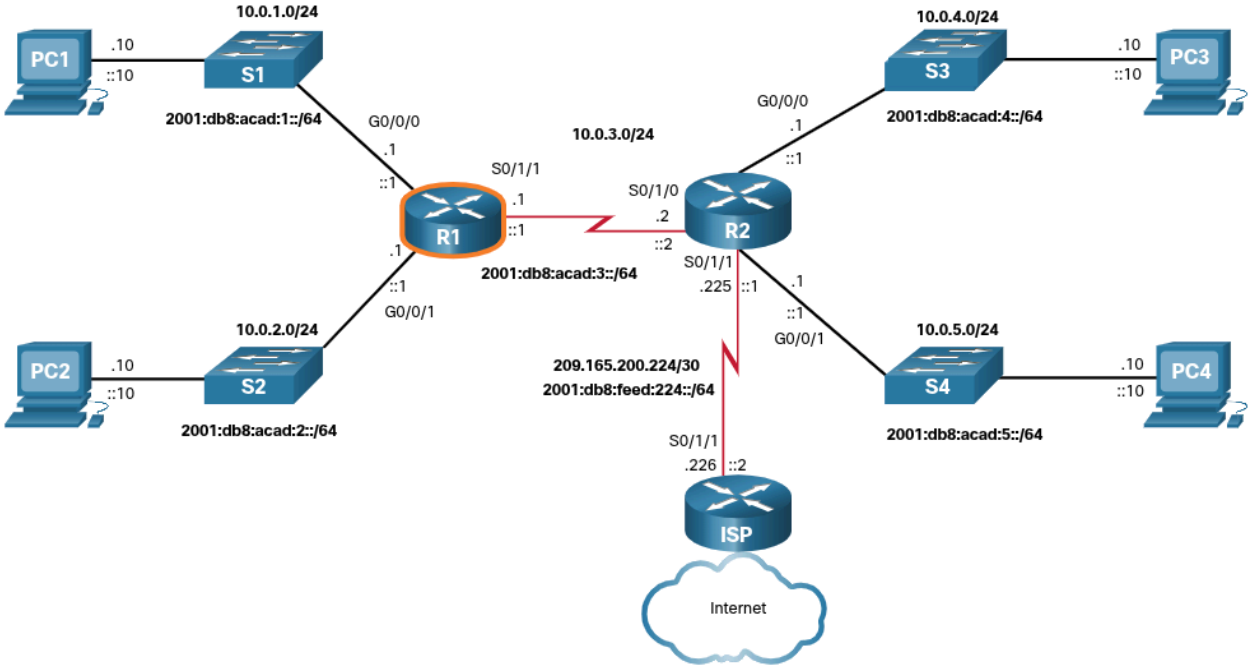

Configuration Commands

- Full configuration for R1

```
Router> enable
Router# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Router(config)# hostname R1
R1(config)# enable secret class
R1(config)# line console 0
R1(config-line)# logging synchronous
R1(config-line)# password cisco
R1(config-line)# login
R1(config-line)# exit
R1(config)# line vty 0 4
R1(config-line)# password cisco
R1(config-line)# login
R1(config-line)# transport input ssh telnet
R1(config-line)# exit
R1(config)# service password-encryption
```

```
R1(config)# banner motd #
Enter TEXT message. End with a new line and the #
*********************************************
WARNING: Unauthorized access is prohibited!
*********************************************
#
R1(config)# ipv6 unicast-routing
R1(config)# interface gigabitethernet 0/0/0
R1(config-if)# description Link to LAN 1
R1(config-if)# ip address 10.0.1.1 255.255.255.0
R1(config-if)# ipv6 address 2001:db8:acad:1::1/64
R1(config-if)# ipv6 address fe80::1:a link-local
R1(config-if)# no shutdown
R1(config-if)# exit
R1(config)# interface gigabitethernet 0/0/1
R1(config-if)# description Link to LAN 2
R1(config-if)# ip address 10.0.2.1 255.255.255.0
R1(config-if)# ipv6 address 2001:db8:acad:2::1/64
R1(config-if)# ipv6 address fe80::1:b link-local
R1(config-if)# no shutdown
R1(config-if)# exit
R1(config)# interface serial 0/1/1
R1(config-if)# description Link to R2
R1(config-if)# ip address 10.0.3.1 255.255.255.0
R1(config-if)# ipv6 address 2001:db8:acad:3::1/64
R1(config-if)# ipv6 address fe80::1:c link-local
R1(config-if)# no shutdown
R1(config-if)# exit
R1# copy running-config startup-config
Destination filename [startup-config]?
Building configuration...
[OK]
R1#
```

Verification Commands

- **show ip interface brief**
- **show running-config interface** *interface-type number*
- **show interfaces**
- **show ip interface**
- **show ip route**
- **ping**

In each case, replace **ip** with **ipv6** for IPv6 version of command

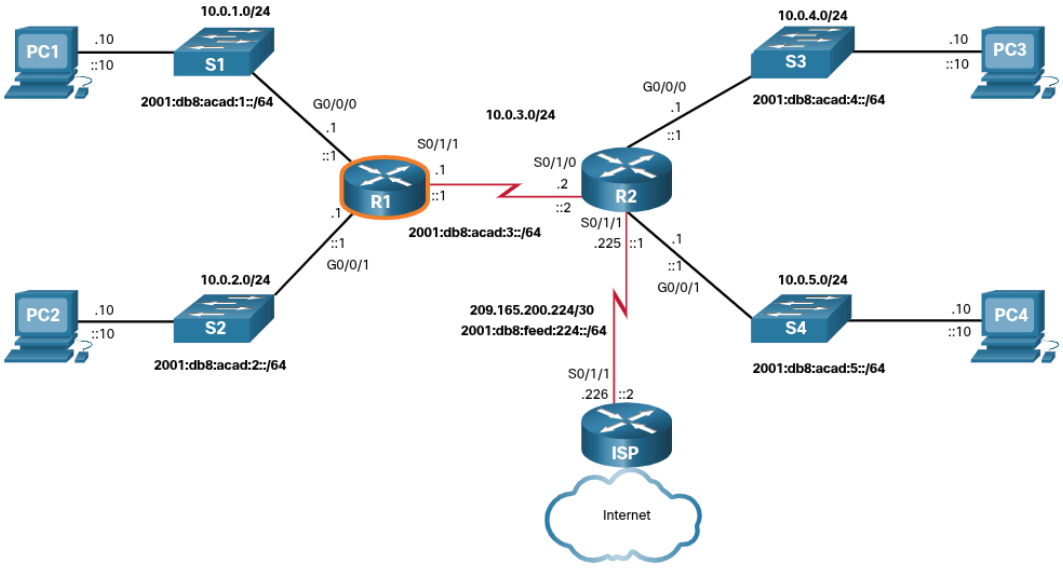

show ip interface brief

```
R1# show ip interface brief
Interface              IP-Address      OK? Method Status
Protocol
GigabitEthernet0/0/0   10.0.1.1        YES manual up
up
GigabitEthernet0/0/1   10.0.2.1        YES manual up
up
Serial0/1/0            unassigned      YES unset  administratively
down down
Serial0/1/1            10.0.3.1        YES manual up
up
GigabitEthernet0       unassigned      YES unset  down
down
R1#
```

show ipv6 interface brief

```
R1# show ipv6 interface brief
GigabitEthernet0/0/0   [up/up]
    FE80::1:A
    2001:DB8:ACAD:1::1
GigabitEthernet0/0/1   [up/up]
    FE80::1:B
    2001:DB8:ACAD:2::1
Serial0/1/0            [administratively down/down]
    unassigned
Serial0/1/1            [up/up]
    FE80::1:C
    2001:DB8:ACAD:3::1
GigabitEthernet0       [down/down]
    unassigned
R1#
```

show running-config interface

```
R1# show running-config interface gigabitethernet 0/0/0
Building configuration...
Current configuration : 189 bytes
!
interface GigabitEthernet0/0/0
 description Link to LAN 1
 ip address 10.0.1.1 255.255.255.0
 negotiation auto
 ipv6 address FE80::1:A link-local
 ipv6 address 2001:DB8:ACAD:1::1/64
end
R1#
```

show interfaces

```
R1# show interfaces gigabitEthernet 0/0/0
GigabitEthernet0/0/0 is up, line protocol is up
  Hardware is ISR4321-2x1GE, address is a0e0.af0d.e140 (bia a0e0.af0d.e140)
  Internet address is 10.0.1.1/24
  MTU 1500 bytes, BW 100000 Kbit/sec, DLY 100 usec,
     reliability 255/255, txload 1/255, rxload 1/255
  Encapsulation ARPA, loopback not set
  Keepalive not supported
  Full Duplex, 100Mbps, link type is auto, media type is RJ45
  output flow-control is off, input flow-control is off
  ARP type: ARPA, ARP Timeout 04:00:00
  Last input 00:00:00, output 00:00:06, output hang never
  Last clearing of "show interface" counters never
  Input queue: 0/375/0/0 (size/max/drops/flushes); Total output drops: 0
  Queueing strategy: fifo
  Output queue: 0/40 (size/max)
  5 minute input rate 2000 bits/sec, 1 packets/sec
  5 minute output rate 0 bits/sec, 0 packets/sec
     57793 packets input, 10528767 bytes, 0 no buffer
     Received 19711 broadcasts (0 IP multicasts)
     0 runts, 0 giants, 0 throttles
     0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored
     0 watchdog, 36766 multicast, 0 pause input
     10350 packets output, 1280030 bytes, 0 underruns
     0 output errors, 0 collisions, 1 interface resets
     0 unknown protocol drops
     0 babbles, 0 late collision, 0 deferred
     0 lost carrier, 0 no carrier, 0 pause output
     0 output buffer failures, 0 output buffers swapped out
R1#
```

show ip interface

```
R1# show ip interface gigabitethernet 0/0/0
GigabitEthernet0/0/0 is up, line protocol is up
  Internet address is 10.0.1.1/24
  Broadcast address is 255.255.255.255
  Address determined by setup command
  MTU is 1500 bytes
  Helper address is not set
  Directed broadcast forwarding is disabled
  Multicast reserved groups joined: 224.0.0.5 224.0.0.6
  Outgoing Common access list is not set
  Outgoing access list is not set
  Inbound Common access list is not set
  Inbound  access list is not set
  Proxy ARP is enabled
  Local Proxy ARP is disabled
  Security level is default
  Split horizon is enabled
  ICMP redirects are always sent
  ICMP unreachables are always sent
  ICMP mask replies are never sent
  IP fast switching is enabled
  IP Flow switching is disabled
  IP CEF switching is enabled
  IP CEF switching turbo vector
  IP Null turbo vector
  Associated unicast routing topologies:
        Topology "base", operation state is UP
  IP multicast fast switching is enabled
  IP multicast distributed fast switching is disabled
  IP route-cache flags are Fast, CEF
  Router Discovery is disabled
  IP output packet accounting is disabled
  IP access violation accounting is disabled
  TCP/IP header compression is disabled
  RTP/IP header compression is disabled
  Probe proxy name replies are disabled
  Policy routing is disabled
  Network address translation is disabled
  BGP Policy Mapping is disabled
  Input features: MCI Check
  IPv4 WCCP Redirect outbound is disabled
  IPv4 WCCP Redirect inbound is disabled
  IPv4 WCCP Redirect exclude is disabled
R1#
```

show ipv6 interface

```
R1# show ipv6 interface gigabitethernet 0/0/0
GigabitEthernet0/0/0 is up, line protocol is up
  IPv6 is enabled, link-local address is FE80::1:A
  No Virtual link-local address(es):
  Global unicast address(es):
    2001:DB8:ACAD:1::1, subnet is 2001:DB8:ACAD:1::/64
  Joined group address(es):
    FF02::1
    FF02::2
    FF02::5
    FF02::6
    FF02::1:FF00:1
    FF02::1:FF01:A
  MTU is 1500 bytes
  ICMP error messages limited to one every 100 milliseconds
  ICMP redirects are enabled
  ICMP unreachables are sent
  ND DAD is enabled, number of DAD attempts: 1
  ND reachable time is 30000 milliseconds (using 30000)
  ND advertised reachable time is 0 (unspecified)
  ND advertised retransmit interval is 0 (unspecified)
  ND router advertisements are sent every 200 seconds
  ND router advertisements live for 1800 seconds
  ND advertised default router preference is Medium
  Hosts use stateless autoconfig for addresses.
R1#
```

show ip route

```
R1# show ip route
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
(Output omitted)
Gateway of last resort is not set
      10.0.0.0/8 is variably subnetted, 6 subnets, 2 masks
C        10.0.1.0/24 is directly connected, GigabitEthernet0/0/0
L        10.0.1.1/32 is directly connected, GigabitEthernet0/0/0
C        10.0.2.0/24 is directly connected, GigabitEthernet0/0/1
L        10.0.2.1/32 is directly connected, GigabitEthernet0/0/1
C        10.0.3.0/24 is directly connected, Serial0/1/1
L        10.0.3.1/32 is directly connected, Serial0/1/1
R1#
```

show ipv6 route

```
R1# show ipv6 route
IPv6 Routing Table - default - 5 entries
Codes: C - Connected, L - Local, S - Static, U - Per-user Static route
(Output omitted)
C   2001:DB8:ACAD:1::/64 [0/0]
     via GigabitEthernet0/0/0, directly connected
L   2001:DB8:ACAD:1::1/128 [0/0]
     via GigabitEthernet0/0/0, receive
C   2001:DB8:ACAD:2::/64 [0/0]
     via GigabitEthernet0/0/1, directly connected
L   2001:DB8:ACAD:2::1/128 [0/0]
     via GigabitEthernet0/0/1, receive
C   2001:DB8:ACAD:3::/64 [0/0]
     via Serial0/1/1, directly connected
L   2001:DB8:ACAD:3::1/128 [0/0]
     via Serial0/1/1, receive
L   FF00::/8 [0/0]
     via Null0, receive
R1#
```

ping - Assumes s0/1/0 interface on R2 is configured and active

```
R1# ping 10.0.3.2
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.0.3.2, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 2/2/2 ms
R1# ping 2001:db8:acad:3::2
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2001:DB8:ACAD:3::2, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 2/2/2 ms
R1#
```

Filter Command Output

Can be used to display specific sections of output

- To enable the filtering command, enter a pipe (|) character after the **show** command then enter a filtering parameter and filtering expression

- **section** - Displays the entire section that starts with the filtering expression
- **include** - Includes all output lines that match the filtering expression
- **exclude** - Excludes all output lines that match the filtering expression
- **begin** - Displays all the output lines from a certain point, starting with the line that matches the filtering expression

Note - Output filters can be used in combination with any **show** command

```
R1# show running-config | section line vty
line vty 0 4
password 7 121A0C0411044C
login
transport input telnet ssh
R1#
R1# show ipv6 interface brief | include up
GigabitEthernet0/0/0   [up/up]
GigabitEthernet0/0/1   [up/up]
Serial0/1/1            [up/up]
R1#
R1# show ip interface brief | exclude unassigned
Interface              IP-Address      OK? Method Status
Protocol
GigabitEthernet0/0/0   192.168.10.1    YES manual up                    up
GigabitEthernet0/0/1   192.168.11.1    YES manual up                    up
Serial0/1/1            209.165.200.225 YES manual up                    up
R1#
R1# show ip route | begin Gateway
Gateway of last resort is not set
      192.168.10.0/24 is variably subnetted, 2 subnets, 2 masks
C        192.168.10.0/24 is directly connected, GigabitEthernet0/0/0
L        192.168.10.1/32 is directly connected, GigabitEthernet0/0/0
      192.168.11.0/24 is variably subnetted, 2 subnets, 2 masks
C        192.168.11.0/24 is directly connected, GigabitEthernet0/0/1
L        192.168.11.1/32 is directly connected, GigabitEthernet0/0/1
      209.165.200.0/24 is variably subnetted, 2 subnets, 2 masks
C        209.165.200.224/30 is directly connected, Serial0/1/1
L        209.165.200.225/32 is directly connected, Serial0/1/1
R1#
```

Route Sources

How does a router know where it can send packets? It creates a routing table that is based on the network in which it is located.

A routing table contains a list of routes to known networks (prefixes and prefix lengths). The source of this information is derived from:

- Directly connected networks
- Static routes
- Dynamic routing protocols

In figure below, R1 and R2 are using dynamic routing protocol OSPF to share routing information

- Additionally, R2 is configured with a default static route to ISP

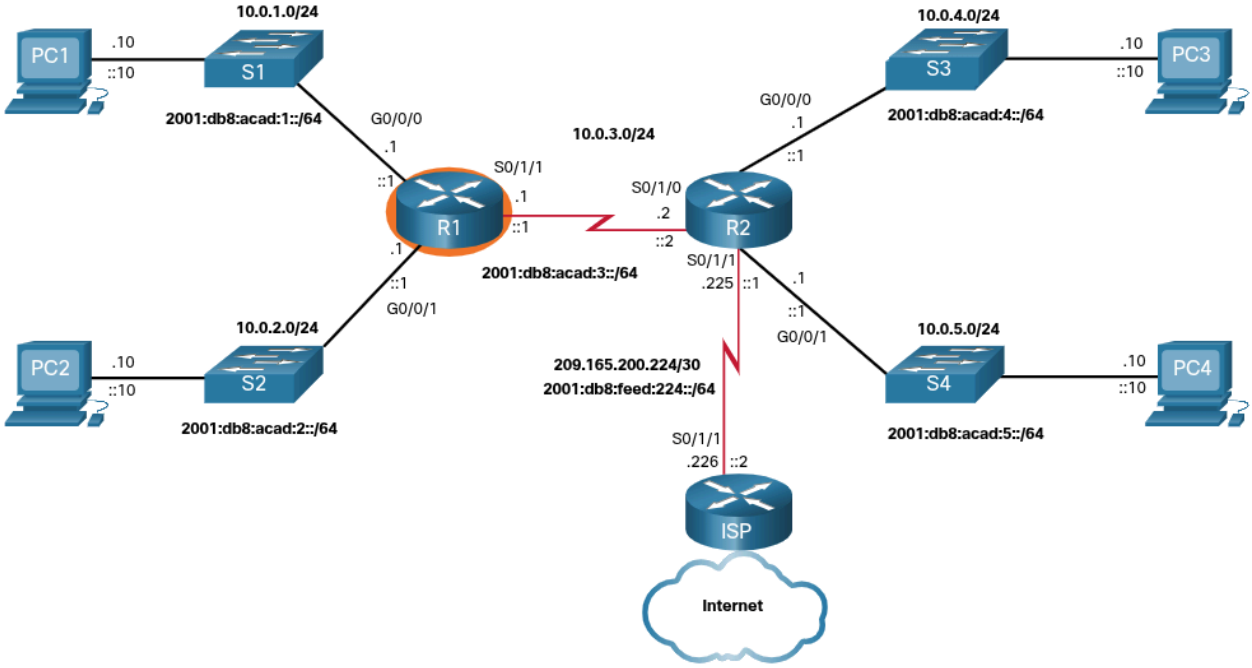

In the routing tables for R1 and R2, notice that the sources for each route are identified by a code. The code identifies how the route was learned.
Common Codes:

- **L** - Identifies the address assigned to a router interface. This allows the router to efficiently determine when it receives a packet for the interface instead of being forwarded
- **C** - Identifies a directly connected network
- **S** - Identifies a static route created to reach a specific network
- **O** - Identifies a dynamically learned network from another router using OSPF routing protocol
- ***** - This route is a candidate for a default route

R1 Routing Table

```
R1# show ip route
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP
       a - application route
       + - replicated route, % - next hop override, p - overrides from PfR
Gateway of last resort is 10.0.3.2 to network 0.0.0.0
O*E2  0.0.0.0/0 [110/1] via 10.0.3.2, 00:51:34, Serial0/1/1
      10.0.0.0/8 is variably subnetted, 8 subnets, 2 masks
C        10.0.1.0/24 is directly connected, GigabitEthernet0/0/0
L        10.0.1.1/32 is directly connected, GigabitEthernet0/0/0
C        10.0.2.0/24 is directly connected, GigabitEthernet0/0/1
L        10.0.2.1/32 is directly connected, GigabitEthernet0/0/1
C        10.0.3.0/24 is directly connected, Serial0/1/1
L        10.0.3.1/32 is directly connected, Serial0/1/1
O        10.0.4.0/24 [110/50] via 10.0.3.2, 00:24:22, Serial0/1/1
O        10.0.5.0/24 [110/50] via 10.0.3.2, 00:24:15, Serial0/1/1
R1#
```

R2 Routing Table

```
R2# show ip route
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP
       a - application route
       + - replicated route, % - next hop override, p - overrides from PfR
Gateway of last resort is 209.165.200.226 to network 0.0.0.0
S*   0.0.0.0/0 [1/0] via 209.165.200.226
     10.0.0.0/8 is variably subnetted, 8 subnets, 2 masks
O       10.0.1.0/24 [110/65] via 10.0.3.1, 00:31:38, Serial0/1/0
O       10.0.2.0/24 [110/65] via 10.0.3.1, 00:31:38, Serial0/1/0
C       10.0.3.0/24 is directly connected, Serial0/1/0
L       10.0.3.2/32 is directly connected, Serial0/1/0
C       10.0.4.0/24 is directly connected, GigabitEthernet0/0/0
L       10.0.4.1/32 is directly connected, GigabitEthernet0/0/0
C       10.0.5.0/24 is directly connected, GigabitEthernet0/0/1
L       10.0.5.1/32 is directly connected, GigabitEthernet0/0/1
     209.165.200.0/24 is variably subnetted, 2 subnets, 2 masks
C       209.165.200.224/30 is directly connected, Serial0/1/1
L       209.165.200.225/32 is directly connected, Serial0/1/1
R2#
```

Routing Table Principles

There are three routing table principles.

- These are issues addressed by the proper configuration of dynamic routing protocols or static routes on all the routers between the source and destination devices

Routing Table Principle	Example
Every router makes its decision alone, based on the information it has in its own routing table	• R1 can only forward packets using its own routing table • R1 does not know what routes are in the routing tables of other routers (IE: R2)
The information in a routing table of one router does not necessarily match the routing table of another router	Just because R1 has route in its routing table to a network in the internet via R2, that does not mean that R2 knows about that same network
Routing information about a path does not provide return routing information	R1 receives a packet with the destination IP address of PC1 and the source IP address of PC3. Just because R1 knows to forward the packet out its G0/0/0 interface, doesn't necessarily mean that it knows how to forward packets originating from PC1 back to the remote network of PC3

Routing Table Entries

- Imperative to know how to interpret the content of IPv4 and IPv5 routing tables
- Figure below displays IPv4 and IPv6 routing table entries on R1 for the route to remote network 10.0.4.0/24 and 2001:db8:acad:4::/64
    - Both routes were learned dynamically from the OSPF routing protocol

IPv4 Routing Table

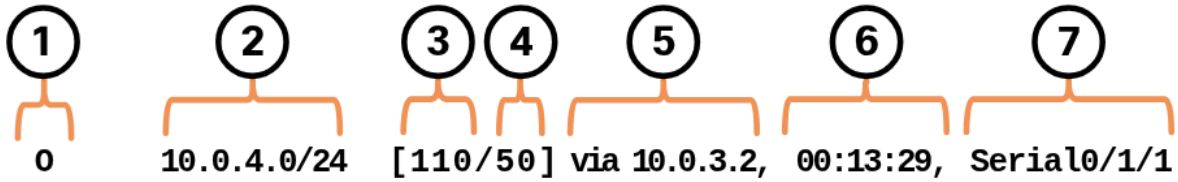

IPv6 Routing Table

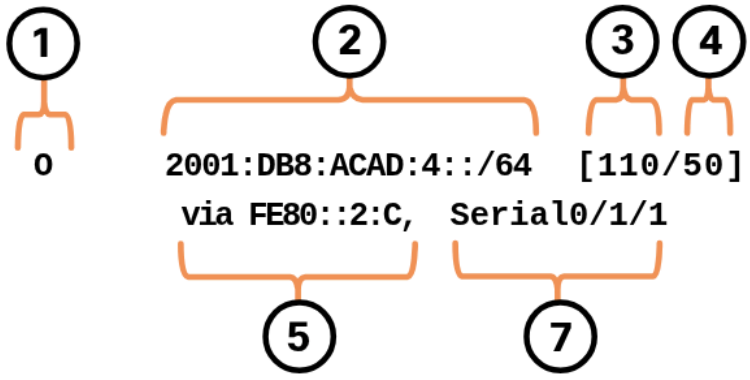

Directly Connected Networks

Before a router can learn about any remote networks, it must have at least one active interface configured with an IP address and subnet mask (prefix length)

- This is known as a directly connected network or directly connected route
- Routers add a directly connected route to its routing table when an interface is configured with an IP address and is activated

A directly connected network is denoted by a status code of **C** in the routing table. The route contains a network prefix and prefix length

The routing table also contains a local route for each of its directly connected networks, indicated by the status code of **L**

- This is the IP address that is assigned to the interface on that directly connected network

For IPv4 local routes

- Prefix length is /32

For IPv6 local routes

- Prefix length is /128

This means the destination IP address of the packet must match all the bits in the local route for this route to be a match.

- Purpose is to efficiently determine when it receives a packet for the interface instead of a packet that needs to be forwarded.

```
R1# show ip route
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
(Output omitted)
C        10.0.1.0/24 is directly connected, GigabitEthernet0/0/0
L        10.0.1.1/32 is directly connected, GigabitEthernet0/0/0
R1#
R1# show ipv6 route
IPv6 Routing Table - default - 10 entries
Codes: C - Connected, L - Local, S - Static, U - Per-user Static route
(Output omitted)

C   2001:DB8:ACAD:1::/64 [0/0]
     via GigabitEthernet0/0/0, directly connected
L   2001:DB8:ACAD:1::1/128 [0/0]
     via GigabitEthernet0/0/0, receive
R1#
```

Static Routes

After directly connected interfaces are configured and added to the routing table, static or dynamic routing can be implemented for accessing remote networks.

Static routes are manually configured

- Defines an explicit path between two networking devices
- Unlike a dynamic routing protocol, static routes are not automatically updated and must be manually reconfigured if the network topology changes

Benefits:

- Improved security
- Resource efficiency
- Uses less bandwidth than dynamic routing protocols
- No CPU cycles are used to calculate and communicate routes

Disadvantage

- Lack of automatic reconfiguration if the network topology changes

Three primary uses:

- It provides ease of routing table maintenance in smaller networks that are not expected to grow significantly
- It uses a single default route to represent a path to any network that does not have a more specific match with another route in the routing table
    - Default routes are used to send traffic to any destination beyond the next upstream router
- It routes to and from stub networks.
    - A stub network is network accessed by a single route, and the router has only one neighbor

The figure below shows an example of stub networks. Notice that any network attached to R1 would only have one way to reach other destinations, whether to networks attached to R2, or to destinations beyond R2. This means that networks 10.0.1.0/24 and 10.0.2.0/24 are stub networks and R1 is a stub router

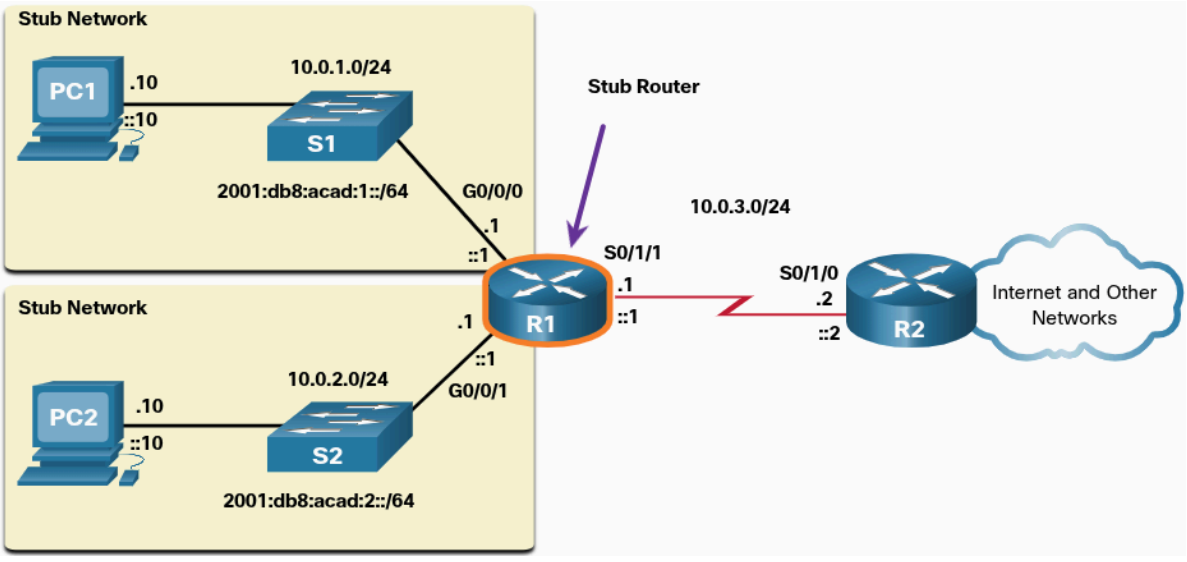

Static Routes in the IP Routing Table

For demonstration, topology in figure is simplified to show only one LAN attached to each router

- Figure shows IPv4 and IPv6 static routes configured on R1 to reach the 10.0.4.0/24 and 2001:db8:acad:4::/64 networks on R2

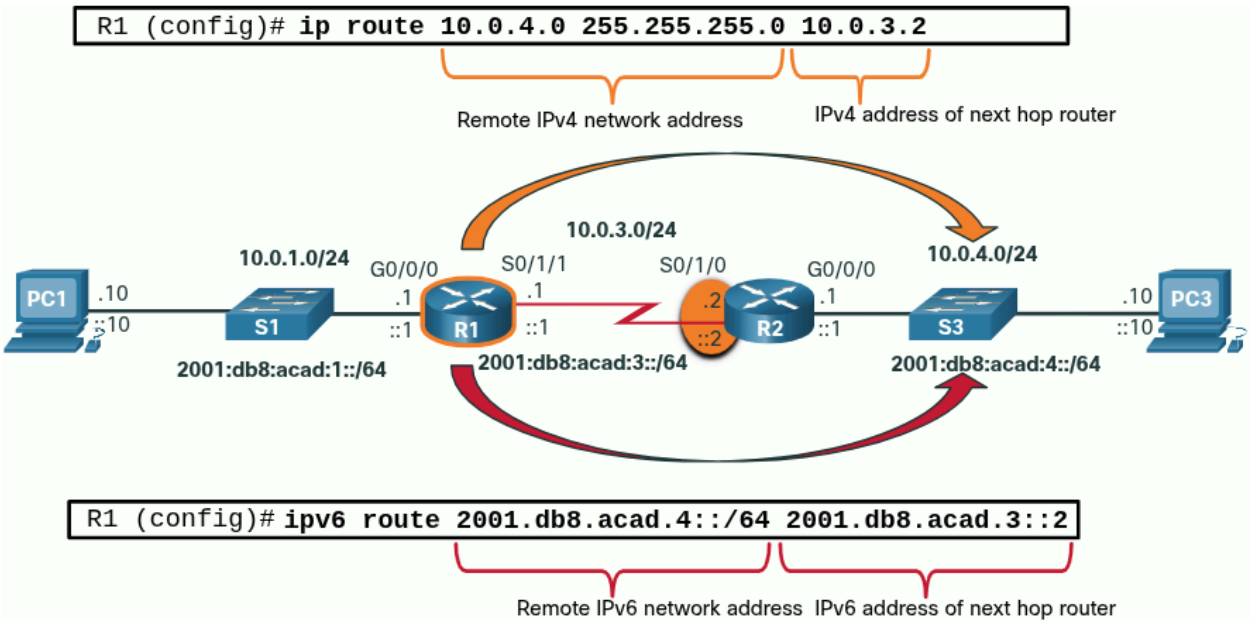

The output shows the IPv4 and IPv6 static routing entries on R1 that can reach the 10.0.4.0/24 and 2001:db8:acad:4::/64 networks on R2

- Notice that both routing entries use the status code of **S** indicating that the route was learned by a static route
- Both entries also include the IP address of the next hop router, via *ip-address*
- The **static** parameter at the end of the command displays only static routes

```
R1# show ip route static
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
(output omitted)

      10.0.0.0/8 is variably subnetted, 8 subnets, 2 masks
S        10.0.4.0/24 [1/0] via 10.0.3.2
R1# show ipv6 route static
IPv6 Routing Table - default - 8 entries
Codes: C - Connected, L - Local, S - Static, U - Per-user Static route
(output omitted)

S   2001:DB8:ACAD:4::/64 [1/0]
     via 2001:DB8:ACAD:3::2
```

Dynamic Routing Protocols

Used by routers to automatically share information about the reachability and status of remote networks.

- Perform several activities including network discovery and maintaining routing tables

Important advantages:

- Ability to select a best path
- Ability to automatically discover a new best path when there is a change in topology

Network Discovery

Ability of a routing protocol to share information about the networks that it knows about with other routers that are also using the same routing protocol

- Instead of depending on manually configured static routes to remote networks on every router, a dynamic routing protocol allows the routers to automatically learn about these networks from other routers
    - These networks, and the best path to each, are added to the routing table of the router and identified as a network learned by a specific dynamic routing protocol

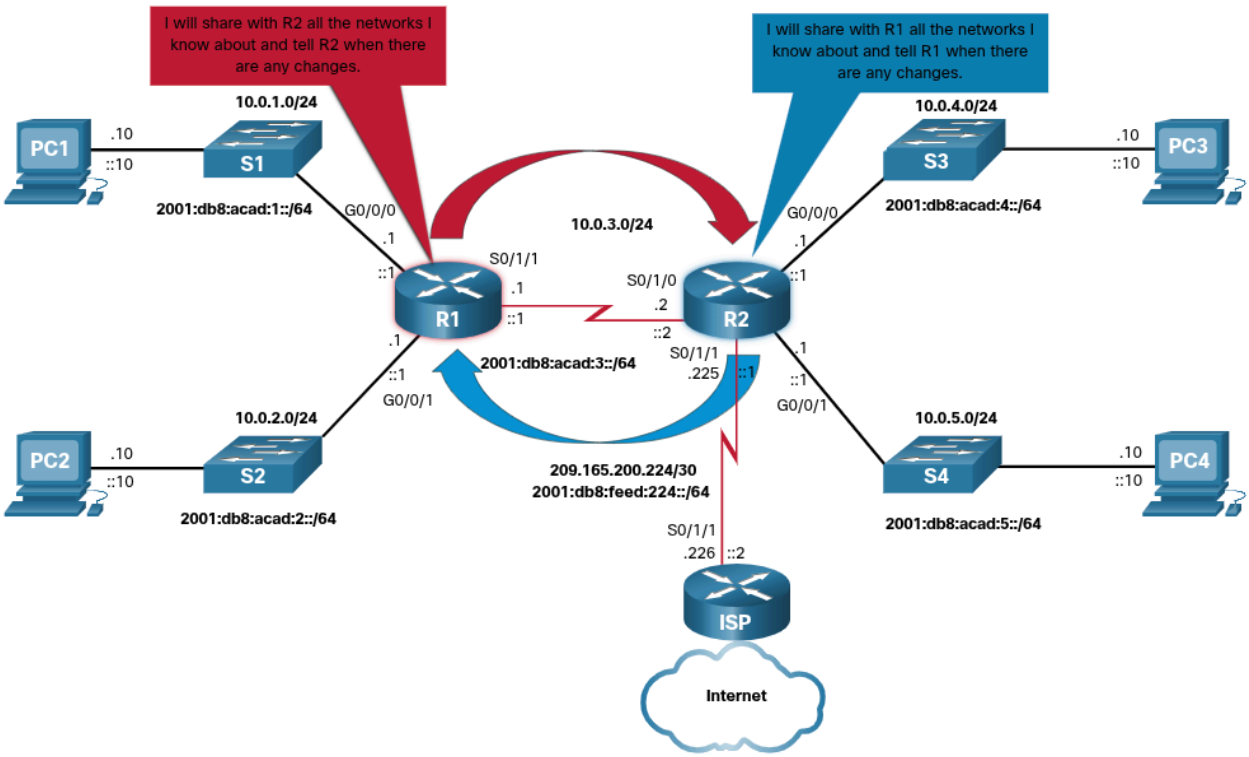

Dynamic Routes in the IP Routing Table

- Previous example used static routes to 10.0.4.0/24 and 2001:db8:acad:4::/64
    - No longer configured and OSPF is now being used to dynamically learn all the networks connected to R1 and R2

Following examples show the IPv4 and IPv6 OSPF routing entries on R1 that can reach these networks on R2

- Notice that both routing entries use the status code of **O** to indicate the route was learned by the OSPF routing protocol. Both entries also include the IP address of the next-hop router, via *ip-address*

Note - IPv6 routing protocols use link-local address of the next-hop router

```
R1# show ip route
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
(output omitted for brevity)
O         10.0.4.0/24 [110/50] via 10.0.3.2, 00:24:22, Serial0/1/1
O         10.0.5.0/24 [110/50] via 10.0.3.2, 00:24:15, Serial0/1/1
R1# show ipv6 route
IPv6 Routing Table - default - 10 entries
(Output omitted)
       NDr - Redirect, RL - RPL, O - OSPF Intra, OI - OSPF Inter
O   2001:DB8:ACAD:4::/64 [110/50]
       via FE80::2:C, Serial0/1/1
O   2001:DB8:ACAD:5::/64 [110/50]
       via FE80::2:C, Serial0/1/1
```

Default Route

Similar to a default gateway on a host

- Specifies a next-hop router to use when the routing table does not contain a specific route that matches the destination IP address
- Can be either a static route or learned automatically from a dynamic routing protocol
- Has an IPv4 route entry of 0.0.0.0/0 or an IPv6 entry of ::/0
    - This means that zero or no bits need to match between the destination IP address and the default route

Most enterprise routers have a default route in their routing table. This is to reduce the number of routes in a routing table.

A router, such as a home or small office router that only has one LAN, may reach all its remote networks through a default route. This is useful when the router has only directly connected networks and one exit point to a service provider router.

In figure below, routers R1 and R2 are using OSPF to share routing information about their own networks (10.0.x.x/24 and 2001:db8:acad:x::/64)

- R2 has a static default route to the ISP router
- R2 will now forward any packet with a destination IP address that does not specifically match one of the networks in its routing table to the ISP router
    - This would include all packets destined for the internet

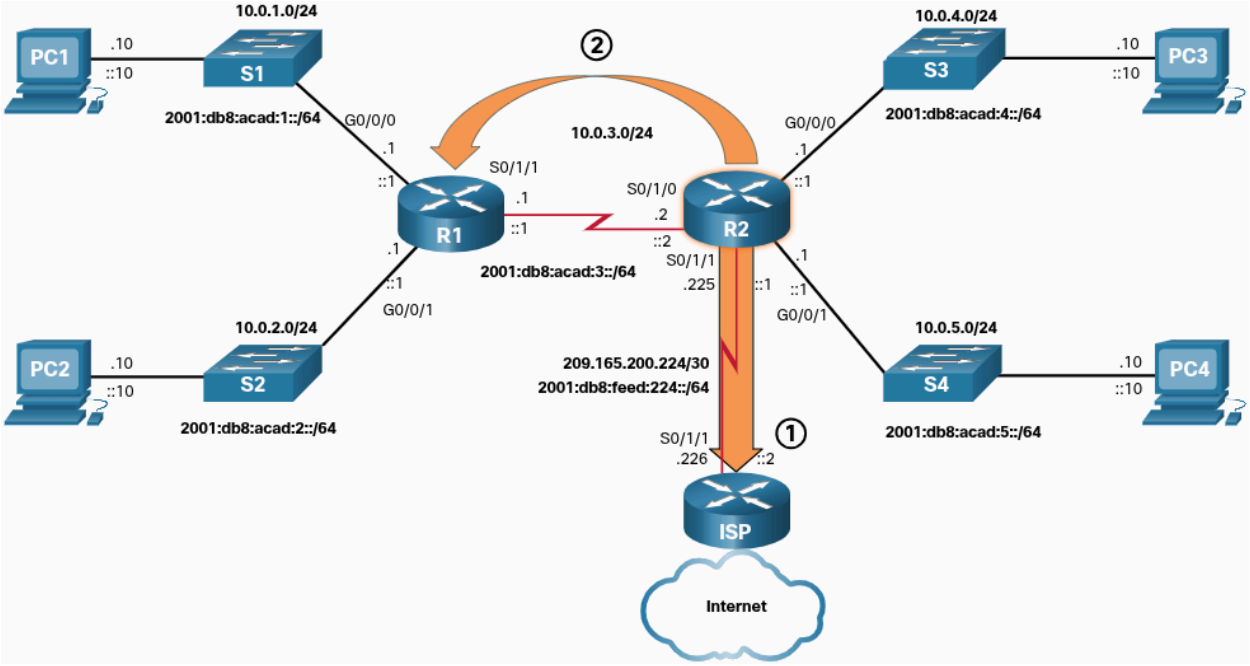

1. *R2 has a static default route to the ISP router*
2. *The default route is advertised by R2 to R1 using the dynamic routing protocol OSPF*

R2 has shared its default route to R1 using OSPF.

- R1 will not have a default route in its routing table that it learned dynamically from OSPF.
- R1 will also forward any packets with a destination IP address that does not specifically match one of the networks in its routing table to R2

Following example shows the IPv4 and IPv6 routing table entries for the static default routes configured on R2

```
R2# show ip route
(Output omitted)
S*    0.0.0.0/0 [1/0] via 209.165.200.226
R2#
R2# show ipv6 route
(Output omitted)
S   ::/0 [1/0]
     via 2001:DB8:FEED:224::2
R2#
```

Structure of an IPv4 Routing Table

IPv4 was standardized in the early 1980s using the now obsolete classful addressing architecture.
The IPv4 routing table is organized using this same classful structure.

In the **show ip route** output, notice that some route entries are left justified while others are indented

- This is based on how the routing process searches the IPv4 routing table for the longest match
- This was all because of classful addressing
    - Although the lookup process no longer uses classes, the structure of the IPv4 routing table still retains in this format

```
Router# show ip route
(Output omitted)
      192.168.1.0/24 is variably subnetted, 2 subnets, 2 masks
C     192.168.1.0/24 is directly connected, GigabitEthernet0/0
L     192.168.1.1/32 is directly connected, GigabitEthernet0/0
O    192.168.2.0/24 [110/65] via 192.168.12.2, 00:32:33, Serial0/0/0
O    192.168.3.0/24 [110/65] via 192.168.13.2, 00:31:48, Serial0/0/1
      192.168.12.0/24 is variably subnetted, 2 subnets, 2 masks
C     192.168.12.0/30 is directly connected, Serial0/0/0
L     192.168.12.1/32 is directly connected, Serial0/0/0
      192.168.13.0/24 is variably subnetted, 2 subnets, 2 masks
C     192.168.13.0/30 is directly connected, Serial0/0/1
L     192.168.13.1/32 is directly connected, Serial0/0/1
      192.168.23.0/30 is subnetted, 1 subnets
O     192.168.23.0/30 [110/128] via 192.168.12.2, 00:31:38,
Serial0/0/0
Router#
```

An indented entry is known as a child route.

- A route entry is indented if it is the subnet of a classful address (Class A, B, or C network)

Directly connected networks will always be indented (child routes) because the local address of the interface is always entered in the routing table as a /32.

- The child route will include the route source and all the forwarding information such as the next-hop address
- The classful network address of this subnet will be shown above the route entry, less indented, and without the source code
    - That route is known as a parent route

Next example shows the IPv4 routing table of R1 in the topology

- Notice that all of the networks in the topology are subnets, which are child routes, of the class A network and parent route 10.0.0.0/8

```
R1# show ip route
(output omitted for brevity)
O*E2  0.0.0.0/0 [110/1] via 10.0.3.2, 00:51:34, Serial0/1/1
      10.0.0.0/8 is variably subnetted, 8 subnets, 2 masks
C        10.0.1.0/24 is directly connected, GigabitEthernet0/0/0
L        10.0.1.1/32 is directly connected, GigabitEthernet0/0/0
C        10.0.2.0/24 is directly connected, GigabitEthernet0/0/1
L        10.0.2.1/32 is directly connected, GigabitEthernet0/0/1
C        10.0.3.0/24 is directly connected, Serial0/1/1
L        10.0.3.1/32 is directly connected, Serial0/1/1
O        10.0.4.0/24 [110/50] via 10.0.3.2, 00:24:22, Serial0/1/1
O        10.0.5.0/24 [110/50] via 10.0.3.2, 00:24:15, Serial0/1/1
R1#
```

Structure of an IPv6 Routing Table

The concept of classful address was never part of IPv6, so the structure of an IPv6 routing table is very straight forward. Every IPv6 route entry is formatted and aligned the same way

```
R1# show ipv6 route
(output omitted for brevity)
OE2 ::/0 [110/1], tag 2
      via FE80::2:C, Serial0/0/1
C   2001:DB8:ACAD:1::/64 [0/0]
      via GigabitEthernet0/0/0, directly connected
L   2001:DB8:ACAD:1::1/128 [0/0]
      via GigabitEthernet0/0/0, receive
C   2001:DB8:ACAD:2::/64 [0/0]
      via GigabitEthernet0/0/1, directly connected
L   2001:DB8:ACAD:2::1/128 [0/0]
      via GigabitEthernet0/0/1, receive
C   2001:DB8:ACAD:3::/64 [0/0]
      via Serial0/1/1, directly connected
L   2001:DB8:ACAD:3::1/128 [0/0]
      via Serial0/1/1, receive
O   2001:DB8:ACAD:4::/64 [110/50]
      via FE80::2:C, Serial0/1/1
O   2001:DB8:ACAD:5::/64 [110/50]
      via FE80::2:C, Serial0/1/1
L   FF00::/8 [0/0]
      via Null0, receive
R1#
```

Administrative Distance

A route entry for a specific network address (prefix and prefix length) can only appear once in the routing table

- However, it is possible that the routing table learns about the same network address from more than one routing source

Except for very specific circumstances, only one dynamic routing protocol should be implemented on a router

- However, it is possible to configure both OSPF and EIGRP on a router, and both routing protocols may learn of the same destination network
    - Each routing protocol may device on a different path to reach the destination based on the metric of that routing protocol

Questions:

- How does the router know which source to use?
- Which route should it install in the routing table? The route learned from OSPF, or the route learned from EIGRP?

* Cisco IOS uses Administrative Distance (AD) to determine the route to install into the IP routing table.

- The AD represents the “trustworthiness” of the route
    - The lower the AD, the more trustworthy the route source
    - Because EIGRP has an AD of 90 and OSPF has an AD of 110, the EIGRP route entry would be installed in the routing table

Note - The AD does not necessarily represent which dynamic routing protocol is best

A common example:

- A router learning the same network address from a static route and a dynamic routing protocol, such as OSPF. A static route has an AD of 1, whereas an OSPF-discovered route has an AD of 110
- Given two separate route sources to the same destination, the router chooses to install the route with the lowest AD.
- When a router has the choice of a static route and an OSPF route, the static route takes precedence.

Note - Directly connected networks have the lowest AD of 0. Only a directly connected network can have an AD of 0.

* Format example:
S 10.2.0.0 [1/0] via 172.16.2.2
[1/0] = [Administrative Distance / Metric]

Route Source	Administrative Distance
Directly connected	0
Static route	1
EIGRP summary route	5
External BGP	20
Internal EIGRP	90
OSPF	110
IS-IS	115
RIP	120
External EIGRP	170
Internal BGP	200

Static and Dynamic Routing

Static or Dynamic?

- The answer is both
- Static and dynamic are not mutually exclusive.
    - Most networks use a combination of both

Static Routes
Commonly used in scenarios:

- As a default route forwarding packets to a service provider
- For routes outside the routing domain and not learned by the dynamic routing protocol
- When the network administrator wants to explicitly define the path for a specified network
- For routing between stub networks

Static routes are useful for smaller networks with only one path to an outside network. They also provide security in a larger network for certain types of traffic, or links to other networks that need more control

Dynamic Routing Protocols
Helps the network administrator manage the time-consuming and exacting process of configuring and maintaining static routes.

- Are implemented in any type of network consisting of more than just a few routes
- Are scalable and automatically determine better routes if there is a change in the topology

Commonly used in scenarios:

- In networks consisting of more than just a few routers
- When a change in the network topology requires the network to automatically determine another path
- For scalability. As the network grows, the dynamic routing protocol automatically learns about any new networks

Comparison differences between dynamic and static routing

Feature	**Dynamic Routing**	**Static Routing**
Configuration complexity	Independent of network size	Increases with network size
Topology changes	Automatically adapts to topology changes	Administer intervention required
Scalability	Suitable for simple to complex network topologies	Suitable for simple topologies
Security	Security must be configured	Security is inherent
Resource Usage	Uses CPU, memory, and link bandwidth	No additional resources needed
Path Predictability	Route depends on topology and routing protocol used	Explicitly defined by the administrator

Dynamic Routing Evolution

Dynamic routing protocols have been used in networks since the late 1980s.

- One of the first routing protocols was RIP
    - RIPv1 was released in 1988, but some of the basic algorithms within the protocol were used on the Advanced Research Projects Agency Network (ARPANET) as early as 1969

The RIP protocol was updated to RIPv2 to accommodate growth in the network environment

- However, RIPv2 still does not scale to larger network implementations

To address the needs of larger networks, two advanced routing protocols were developed:

- OSPF

- Intermediate System-to-Intermediate System (IS-IS)

Cisco developed the interior Gateway Routing Protocol (IGRP), which was later replaced by Enhanced IGRP (EIGRP), which also scales well in larger network implementations.

Additionally, there was the need to connect different routing domains of different organizations and provide routing between them.
The Border Gateway Protocol (BGP), successor of Exterior Gateway Protocol (EGP) is used between Internet Service Providers (ISPs)

- BGP is also used between ISPs and some private organizations to exchange routing information

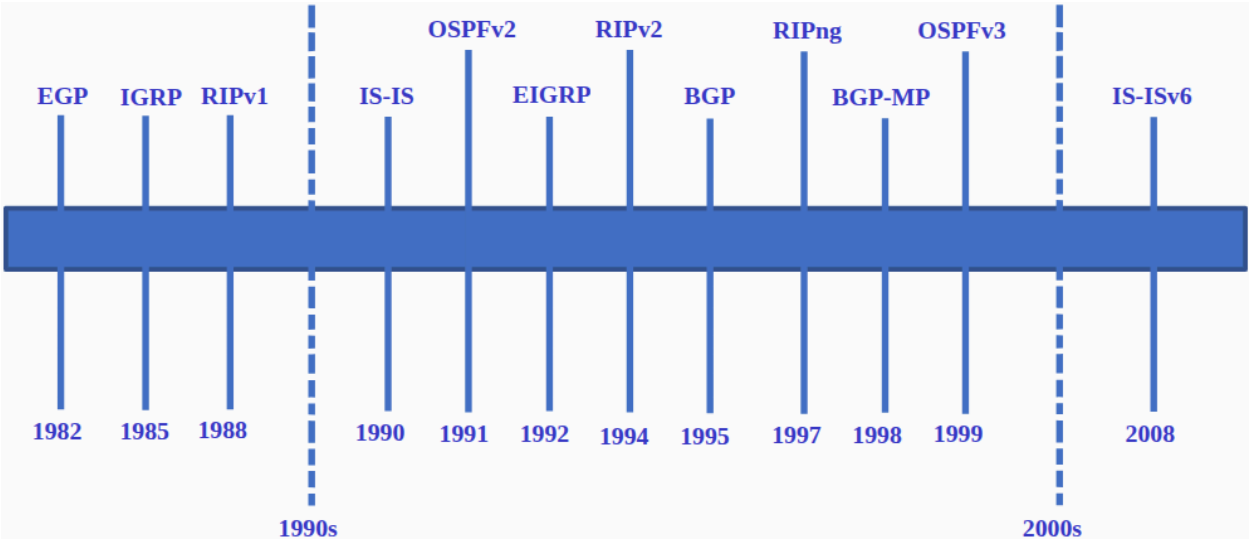

To support IPv6 communication, newer versions of the IP routing protocols have been developed

The table classifies the current routing protocols

- Internal Gateway Protocols (IGPs) are routing protocols used to exchange routing information within a routing domain administered by a single organization

There is only one EGP and it is BGP

- BGP is used to exchange routing information between different organizations, known as autonomous systems (AS)
- * BGP is used by ISPs to route packets over the internet.

	Interior Gateway Protocols				**Exterior Gateway Protocols**
	Distance Vector		**Link-State**		**Path Vector**
IPv4	RIPv2	EIGRP	OSPFv2	IS-IS	BGP-4

IPv6	RIPng	EIGRP for IPv6	OSPFv3	IS-IS for IPv6	BGP-MP

Dynamic Routing Protocol Concepts

A routing protocol is a set of processes, algorithms, and messages that are used to exchange routing information and populate the routing table with the choice of best paths.

The purpose of dynamic routing protocol:

- Discovery of remote networks
- Maintaining up-to-date routing information
- Choosing the best path to destination networks
- Ability to find a new best path if the current path is no longer available

The main components of dynamic routing protocols:

- **Data structures** - Routing protocols typically use tables or databases for their operations. This information is kept in RAM
- **Routing protocol messages** - Routing protocols use various types of messages to discover neighboring routers, exchange routing information, and other tasks to learn and maintain accurate information about the network
- **Algorithm** - An algorithm is a finite list of steps used to accomplish a task. Routing protocols use algorithms for facilitating routing information and for the best path determination

Routing protocols allow routers to dynamically share information about remote networks and automatically offer this information on their own routing tables.

Routing protocols determine the best path, or route, to each network. That route is then offered to the routing table. The route will still be installed in the routing table if there is not another routing source with a lower AD.

A primary benefit of dynamic routing protocols is that routers exchange routing information when there is a topology change.

- This exchange allows routers to automatically learn about new networks and to find alternate paths when there is a link failure to a current network.

Best Path

Determining the best path may involve the evaluation of multiple paths to the same destination network and selecting the optimum or shortest path to reach that network.

The best path is selected by a routing protocol based on the value or metric it uses to determine the distance to reach a network

- A metric is the quantitative value to measure the distance to a given network. The best path to a network is the path with the lowest metric.

Dynamic routing protocols typically use their own rules and metrics to build and update routing tables.
The routing algorithm generates a value, or a metric, for each path through the network.

- Metrics can be based on either a single characteristic or several characteristics of a path.
- Some routing protocols can base route selection on multiple metrics, combining them into a single metric

Routing Protocol	Metric
Routing Information Protocol (RIP)	• The metric is “hop count” • Each router along a path adds a hop to the hop count • A maximum of 15 hops allowed
Open Shortest Path First (OSPF)	• The metric is “cost” which is the based on the cumulative bandwidth from source to destination • Faster links are assigned lower costs compared to slower (higher cost) links
Enhanced Interior Gateway Routing Protocol (EIGRP)	• It calculates a metric based on the slowest bandwidth and delay values • It could also include load and reliability into the metric calculation

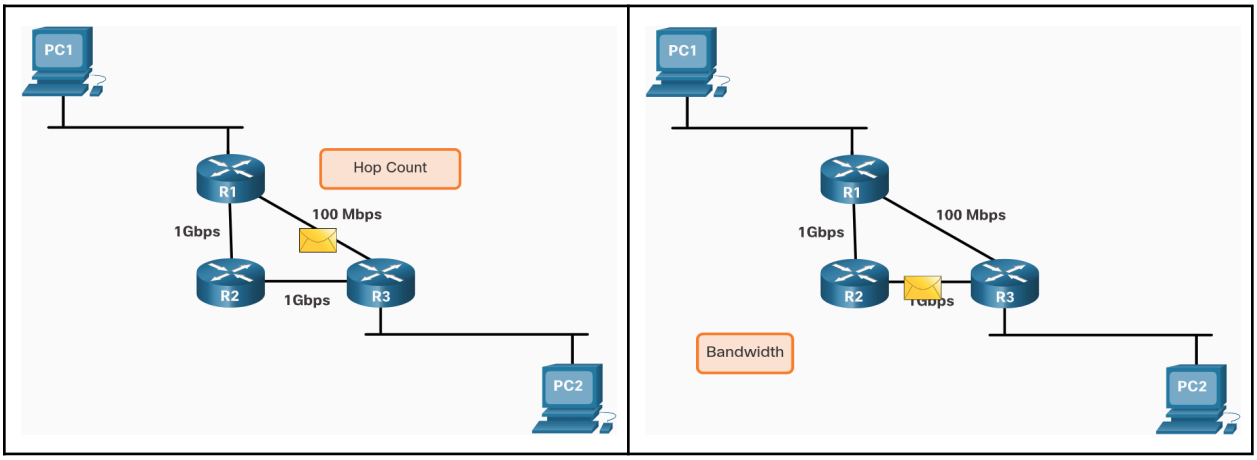

Load Balancing

When a router has two or more paths to a destination with equal cost metrics, then the router forwards the packets using both paths equally.

- This is called equal cost load balancing

The routing table contains the single destination network, but has multiple exit interfaces, one for each equal cost path. The router forwards packets using the multiple exit interfaces lists in the routing table.

If configured correctly, load balancing can increase the effectiveness and performance of the network.

Equal cost load balancing is implemented automatically by dynamically routing protocols. It is enabled with static routes when there are multiple static routes to the same destination network using different next-hop routers.

Note - Only EIGRP supports unequal cost load balancing

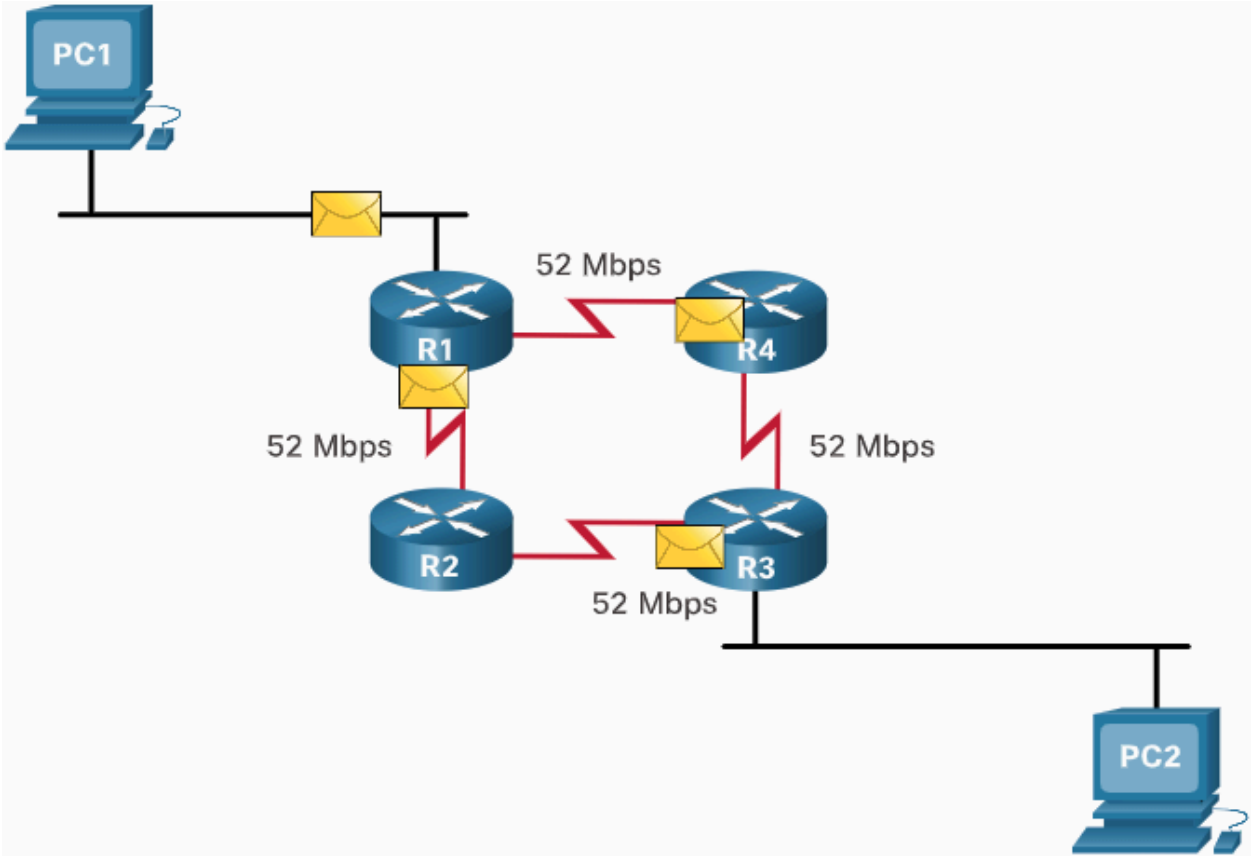

14.6.1 What did I learn in this module?

Path Determination

The primary functions of a router are to determine the best path to forward packets based on the information in its routing table, and to forward packets toward their destination. The best path in the routing table is also known as the longest match. The longest match is the route in the routing table that has the greatest number of far-left matching bits with the destination IP address of the packet. Directly connected networks are networks that are configured on the active interfaces of a router. A directly connected network is added to the routing table when an interface is configured with an IP address and subnet mask (prefix length) and is active (up and up). Routers learn about remote networks in two ways: static routes are added to the routing table when a route is manually configured, and with dynamic routing protocols. Using dynamic routing protocols such as EIGRP and OSPF, routes are added to the routing table when routing protocols dynamically learn about the remote network.

Packet Forwarding

After a router determines the correct path, it can forward the packet on a directly connected network, it can forward the packet to a next-hop router, or it can drop the packet. The primary responsibility of the packet forwarding function is to encapsulate packets in the appropriate data link frame type for the outgoing interface. Routers support three packet forwarding mechanisms: process switching, fast switching, and CEF. The following steps describe the packet forwarding process:

1. The data link frame with an encapsulated IP packet arrives on the ingress interface.
2. The router examines the destination IP address in the packet header and consults its IP routing table.
3. The router finds the longest matching prefix in the routing table.
4. The router encapsulates the packet in a data link frame and forwards it out the egress interface. The destination could be a device connected to the network or a next-hop router.
5. However, if there is no matching route entry the packet is dropped.

Basic Router Configuration Review

There are several configuration and verification commands for routers, including show ip route, show ip interface, show ip interface brief and show running-config. To reduce the amount of command output, use a filter. Filtering commands can be used to display specific sections of output. To enable the filtering

command, enter a pipe (|) character after the show command and then enter a filtering parameter and a filtering expression. The filtering parameters that can be configured after the pipe include the following:

- section - Shows entire section that starts with the filtering expression
- include - Includes all output lines that match the filtering expression
- exclude - Excludes all output lines that match the filtering expression
- begin - Shows all the output lines from a certain point, starting with the line that matches the filtering expression

IP Routing Table

A routing table contains a list of routes known networks (prefixes and prefix lengths). The source of this information is derived from directly connected networks, static routes, and dynamic routing protocols. Common routing table codes include:

- L - Identifies the address assigned to a router interface. This allows the router to efficiently determine when it receives a packet for the interface instead of being forwarded.
- C - Identifies a directly connected network.
- S - Identifies a static route created to reach a specific network.
- O - Identifies a dynamically learned network from another router using the OSPF routing protocol.
- * - This route is a candidate for a default route.

Every router makes its decision alone, based on the information it has in its own routing table. The information in a routing table of one router does not necessarily match the routing table of another router. Routing information about a path does not provide return routing information. Routing table entries include the route source, destination network, AD, metric, next-hop, route timestamp, and exit interface. To learn about remote networks, a router must have at least one active interface configured with an IP address and subnet mask (prefix length), called a directly connected network. Static routes are manually configured and define an explicit path between two networking devices. Dynamic routing protocols can discover a network, maintain routing tables, select a best path, and automatically discover a new best path if the topology changes. The default route specifies a next-hop router to use when the routing table does not contain a specific route that matches the destination IP address. A default route can be either a static route or learned automatically from a dynamic routing protocol. A default route has an IPv4 route entry of 0.0.0.0/0 or an IPv6 route entry of ::/0. IPv4 routing tables still have a structure based on classful addressing represented by levels of indentation. IPv6 routing tables do not use the IPv4 routing table structure. Cisco IOS uses what is known as the administrative distance (AD) to determine the route to install into the IP routing table. The AD represents the "trustworthiness" of the route. The lower the AD, the more trustworthy the route source.

Static and Dynamic Routing

Static routes are commonly used:

- As a default route forwarding packets to a service provider.
- For routes outside the routing domain and not learned by the dynamic routing protocol.
- When the network administrator wants to explicitly define the path for a specific network.
- For routing between stub networks.

Dynamic routing protocol are commonly used:

- In networks consisting of more than just a few routers
- When a change in the network topology requires the network to automatically determine another path
- For scalability. As the network grows, the dynamic routing protocol automatically learns about any new networks.

Current routing protocols include IGPs and EGPs. IGPs exchange routing information within a routing domain administered by a single organization. The only EGP is BGP. BGP exchanges routing information between different organizations. BGP is used by ISPs to route packets over the internet. Distance vector, link-state, and path vector routing protocols refer to the type of routing algorithm used to determine best path. The main components of dynamic routing protocols are data structures, routing protocol messages, and algorithms. The best path is selected by a routing protocol based on the value or metric it uses to determine the distance to reach a network. A metric is the quantitative value used to measure the distance to a given network. The best path to a network is the path with the lowest metric. When a router has two or more paths to a destination with equal cost metrics, then the router forwards the packets using both paths equally. This is called equal cost load balancing.

Types of Static Routes

Static routes are commonly implemented on a network. This is true even when there is a dynamic routing protocol configured.

- For instance, an organization could configure a default static route to the service provider and advertise this route to other corporate routers using the dynamic routing protocol

Static routes can be configured for IPv4 and IPv6. Both protocols support these types:

- standard static route
- Default static route
- Floating static route
- Summary static route

Static routes are configured using the **ip route** and **ipv6 route** global configuration commands

Next-Hop Options

When configuring a static route, the next hop can be identified by an IP address, exit interface, or both. How the destination is specified creates one of three following types:

- **Next-hop route** - Only the next-hop IP address is specified
- **Directly connected static route** - Only the router exit interface is specified
- **Full specified static route** - The next-hop IP address and exit interface are specified

IPv4 Static Route Command

Configured using following global configuration command:

```
Router(config)# ip route network-address subnet-mask { ip-address | exit-intf [ip-address]} [distance]
```

Note - Either the *ip-address, exit-intf*, or the *ip-address* and *exit-intf* parameters must be configured

Table below describes **ip route** command parameters

Parameter	Description
network-address	Identifies the destination IPv4 network address of the remote network to add to the routing table
subnet-mask	● Identifies the subnet mask of the remote network ● The subnet mask can be modified to summarize a group of

	networks and create a summary static route
ip-address	• Identifies the next-hop router IPv4 address • Typically used with broadcast networks (IE: Ethernet) • Could create a recursive static route where the router performs an additional lookup to find the exit interface
exit-intf	• Identifies the exit interface to forward packets • Creates a directly connected static route • Typically used in a point-to-point configuration
exit-intf ip-address	Creates a fully specified static route because it specifies the exit interface and next-hop IPv4 address
distance	• Optional command that can be used to assign an administrative distance value between 1 and 255 • Typically used to configure a floating static route by setting an administrative distance that is higher than a dynamically learned route

IPv6 Static Route Command

Configured using the following global configuration command:

```
Router(config)# ipv6 route ipv6-prefix/prefix-length
{ipv6-address | exit-intf [ipv6-address]} [distance]
```

Most parameters are identical to IPv4 version

Table below shows various **ipv6 route** command parameters and their descriptions

Parameter	Description
ipv6-prefix	Identifies the destination IPv6 network address of the remote network to add to the routing table
/prefix-length	Identifies the prefix length of the remote network
ipv6-address	• Identifies the next-hop router IPv6 address • Typically used with broadcast networks (IE: Ethernet) • Could create a recursive static route where the router performs an additional lookup to find the exit interface
exit-intf	• Identifies the exit interface to forward packets • Creates a directly connected static

	route • Typically used in a point-to-point configuration
exit-intf ipv6-address	Creates a fully specified static route because it specifies the exit interface and next-hop IPv6 address
distance	• Optional command that can be used to assign and administrative distance value between 1 and 255 • Typically used to configure a floating static route by setting an administrative distance that is higher than a dynamically learned route

Note - The **ipv6 unicast-routing** global configuration command must be configured to enable the router to forward IPv6 packets

Dual-Stack Topology

Figure below– currently no static routes are configured for either IPv4 or IPv6

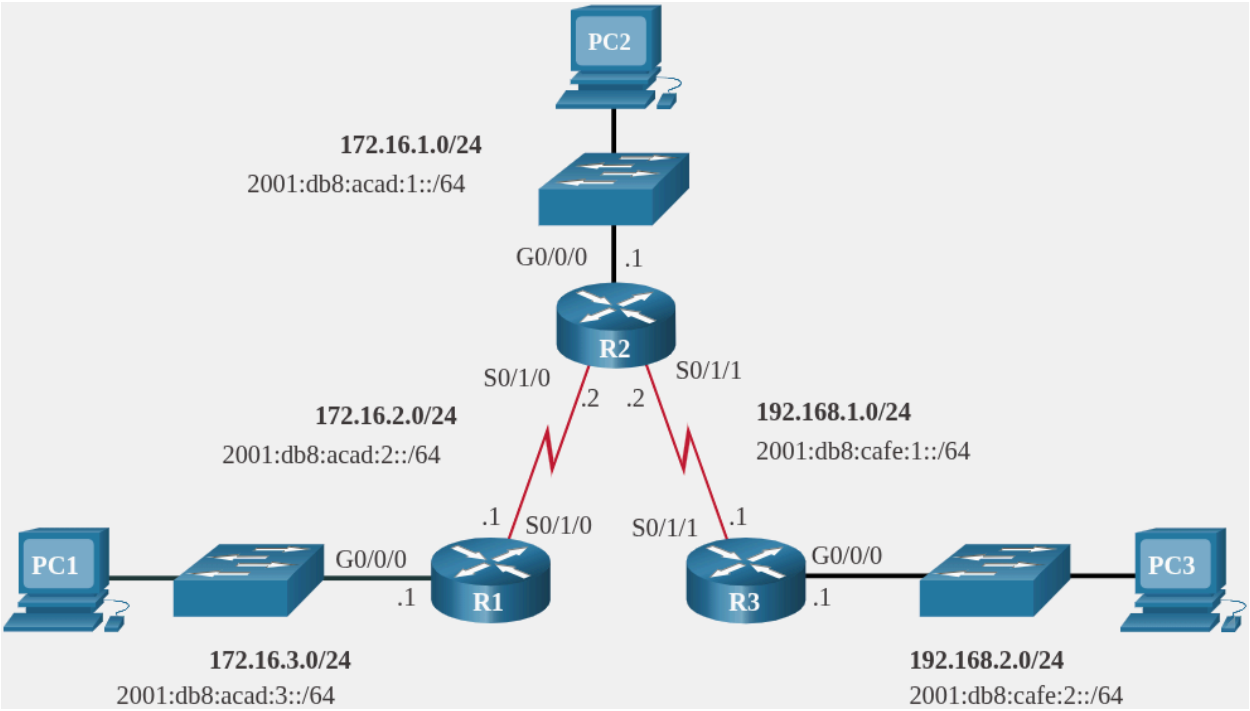

IPv4 Starting Routing Tables

R1 IPv3 Routing Table

```
R1# show ip route | begin Gateway
```

```
Gateway of last resort is not set
      172.16.0.0/16 is variably subnetted, 4 subnets, 2 masks
C        172.16.2.0/24 is directly connected, Serial0/1/0
L        172.16.2.1/32 is directly connected, Serial0/1/0
C        172.16.3.0/24 is directly connected, GigabitEthernet0/0/0
L        172.16.3.1/32 is directly connected, GigabitEthernet0/0/0
R1#
```

R2 IPv4 Routing Table

```
R2# show ip route | begin Gateway
Gateway of last resort is not set
      172.16.0.0/16 is variably subnetted, 4 subnets, 2 masks
C        172.16.1.0/24 is directly connected, GigabitEthernet0/0/0
L        172.16.1.1/32 is directly connected, GigabitEthernet0/0/0
C        172.16.2.0/24 is directly connected, Serial0/1/0
L        172.16.2.2/32 is directly connected, Serial0/1/0
      192.168.1.0/24 is variably subnetted, 2 subnets, 2 masks
C        192.168.1.0/24 is directly connected, Serial0/1/1
L        192.168.1.2/32 is directly connected, Serial0/1/1
R2#
```

R3 IPv4 Routing Table

```
R3# show ip route | begin Gateway
Gateway of last resort is not set
      192.168.1.0/24 is variably subnetted, 2 subnets, 2 masks
C        192.168.1.0/24 is directly connected, Serial0/1/1
L        192.168.1.1/32 is directly connected, Serial0/1/1
      192.168.2.0/24 is variably subnetted, 2 subnets, 2 masks
C        192.168.2.0/24 is directly connected, GigabitEthernet0/0/0
L        192.168.2.1/32 is directly connected, GigabitEthernet0/0/0
R3#
```

R1 Can Ping R2

None of the routers have knowledge of any networks beyond the directly connected interfaces. This means each router can only reach directly connected networks.

A **ping** from R1 to the serial 0/1/0 interface of R2 should be successful because it is a directly-connected network

```
R1# ping 172.16.2.2
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.16.2.2, timeout is 2 seconds:
!!!!!
```

R1 Cannot Ping R3 LAN

However, a **ping** from R1 to the R3 LAN should fail because R1 does not have any entry in its routing table for the R3 LAN network

```
R1# ping 192.168.2.1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.168.2.1, timeout is 2
seconds:
.....
Success rate is 0 percent (0/5)
```

IPv6 Starting Routing Tables

R1 IPv6 Routing Table

```
R1# show ipv6 route | begin C
C   2001:DB8:ACAD:2::/64 [0/0]
     via Serial0/1/0, directly connected
L   2001:DB8:ACAD:2::1/128 [0/0]
     via Serial0/1/0, receive
C   2001:DB8:ACAD:3::/64 [0/0]
     via GigabitEthernet0/0/0, directly connected
L   2001:DB8:ACAD:3::1/128 [0/0]
     via GigabitEthernet0/0/0, receive
L   FF00::/8 [0/0]
     via Null0, receive
R1#
```

R2 IPv6 Routing Table

```
R2# show ipv6 route | begin C
C   2001:DB8:ACAD:1::/64 [0/0]
     via GigabitEthernet0/0/0, directly connected
L   2001:DB8:ACAD:1::1/128 [0/0]
     via GigabitEthernet0/0/0, receive
C   2001:DB8:ACAD:2::/64 [0/0]
     via Serial0/1/0, directly connected
L   2001:DB8:ACAD:2::2/128 [0/0]
     via Serial0/1/0, receive
C   2001:DB8:CAFE:1::/64 [0/0]
     via Serial0/1/1, directly connected
L   2001:DB8:CAFE:1::2/128 [0/0]
     via Serial0/1/1, receive
L   FF00::/8 [0/0]
     via Null0, receive
R2#
```

R3 IPv6 Routing Table

```
R3# show ipv6 route | begin C
C   2001:DB8:CAFE:1::/64 [0/0]
     via Serial0/1/1, directly connected
L   2001:DB8:CAFE:1::1/128 [0/0]
     via Serial0/1/1, receive
C   2001:DB8:CAFE:2::/64 [0/0]
     via GigabitEthernet0/0/0, directly connected
L   2001:DB8:CAFE:2::1/128 [0/0]
     via GigabitEthernet0/0/0, receive
L   FF00::/8 [0/0]
     via Null0, receive
R3#
```

R1 Can Ping R2

None of the routers have knowledge of any networks beyond the directly connected interfaces

A **ping** from R1 to the serial 0/1/0 interface on R2 should be successful

```
R1# ping 2001:db8:acad:2::2
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2001:DB8:ACAD:2::2, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 2/2/3 ms
```

R1 Cannot Ping R3 LAN

However, a **ping** to the R3 LAN is unsuccessful. This is because R1 does not have any entry in its routing table for that network

```
R1# ping 2001:DB8:cafe:2::1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2001:DB8:CAFE:2::1, timeout is 2 seconds:
% No valid route for destination
Success rate is 0 percent (0/1)
```

Configure IP Static Route

IPv4 Next-Hop Static Route

The commands to configure standard static routes vary slightly between IPv3 and Ipv6.

In a next-hop static route, only the next-hop IP address is specified. The exit interface is derived from the next hop

- For example: three next-hop IPv4 static routes are configured on R1 using the IP address of the next hop, R2

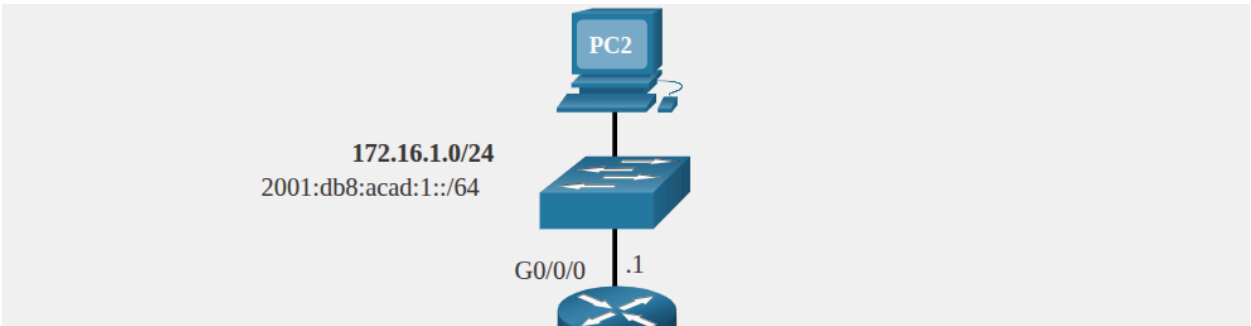

The commands to configure R1 with the IPv4 static routes to the three remote networks are:

```
R1(config)# ip route 172.16.1.0 255.255.255.0 172.16.2.2
R1(config)# ip route 192.168.1.0 255.255.255.0 172.16.2.2
R1(config)# ip route 192.168.2.0 255.255.255.0 172.16.2.2
```

The routing table for R1 now has routes to the three remote IPv4 networks

```
R1# show ip route | begin Gateway
Gateway of last resort is not set
      172.16.0.0/16 is variably subnetted, 5 subnets, 2 masks
S        172.16.1.0/24 [1/0] via 172.16.2.2
C        172.16.2.0/24 is directly connected, Serial0/1/0
L        172.16.2.1/32 is directly connected, Serial0/1/0
C        172.16.3.0/24 is directly connected,
GigabitEthernet0/0/0
L        172.16.3.1/32 is directly connected,
GigabitEthernet0/0/0
S    192.168.1.0/24 [1/0] via 172.16.2.2
S    192.168.2.0/24 [1/0] via 172.16.2.2

R1#
```

IPv6 Next-Hop Static Route

The commands to configure R1 with the IPv6 static routes to the three remote networks are:

```
R1(config)# ipv6 unicast-routing
R1(config)# ipv6 route 2001:db8:acad:1::/64 2001:db8:acad:2::2
R1(config)# ipv6 route 2001:db8:cafe:1::/64 2001:db8:acad:2::2
R1(config)# ipv6 route 2001:db8:cafe:2::/64 2001:db8:acad:2::2
```

The routing table for R1 now has routes to the three remote IPv6 networks

```
R1# show ipv6 route
IPv6 Routing Table - default - 8 entries
Codes: C - Connected, L - Local, S - Static, U - Per-user Static route
      B - BGP, R - RIP, H - NHRP, I1 - ISIS L1
      I2 - ISIS L2, IA - ISIS interarea, IS - ISIS summary, D - EIGRP
      EX - EIGRP external, ND - ND Default, NDp - ND Prefix, DCE - Destination
      NDr - Redirect, RL - RPL, O - OSPF Intra, OI - OSPF Inter
      OE1 - OSPF ext 1, OE2 - OSPF ext 2, ON1 - OSPF NSSA ext 1
      ON2 - OSPF NSSA ext 2, la - LISP alt, lr - LISP site-registrations
      ld - LISP dyn-eid, lA - LISP away, le - LISP extranet-policy
      a - Application
S   2001:DB8:ACAD:1::/64 [1/0]
      via 2001:DB8:ACAD:2::2
C   2001:DB8:ACAD:2::/64 [0/0]
```

```
      via Serial0/1/0, directly connected
L   2001:DB8:ACAD:2::1/128 [0/0]
      via Serial0/1/0, receive
C   2001:DB8:ACAD:3::/64 [0/0]
      via GigabitEthernet0/0/0, directly connected
L   2001:DB8:ACAD:3::1/128 [0/0]
      via GigabitEthernet0/0/0, receive
S   2001:DB8:CAFE:1::/64 [1/0]
      via 2001:DB8:ACAD:2::2
S   2001:DB8:CAFE:2::/64 [1/0]
      via 2001:DB8:ACAD:2::2
L   FF00::/8 [0/0]
      via Null0, receive
```

IPv4 Directly Connected Static Route

When configuring a static route, another option is to use the next interface to specify the next-hop address

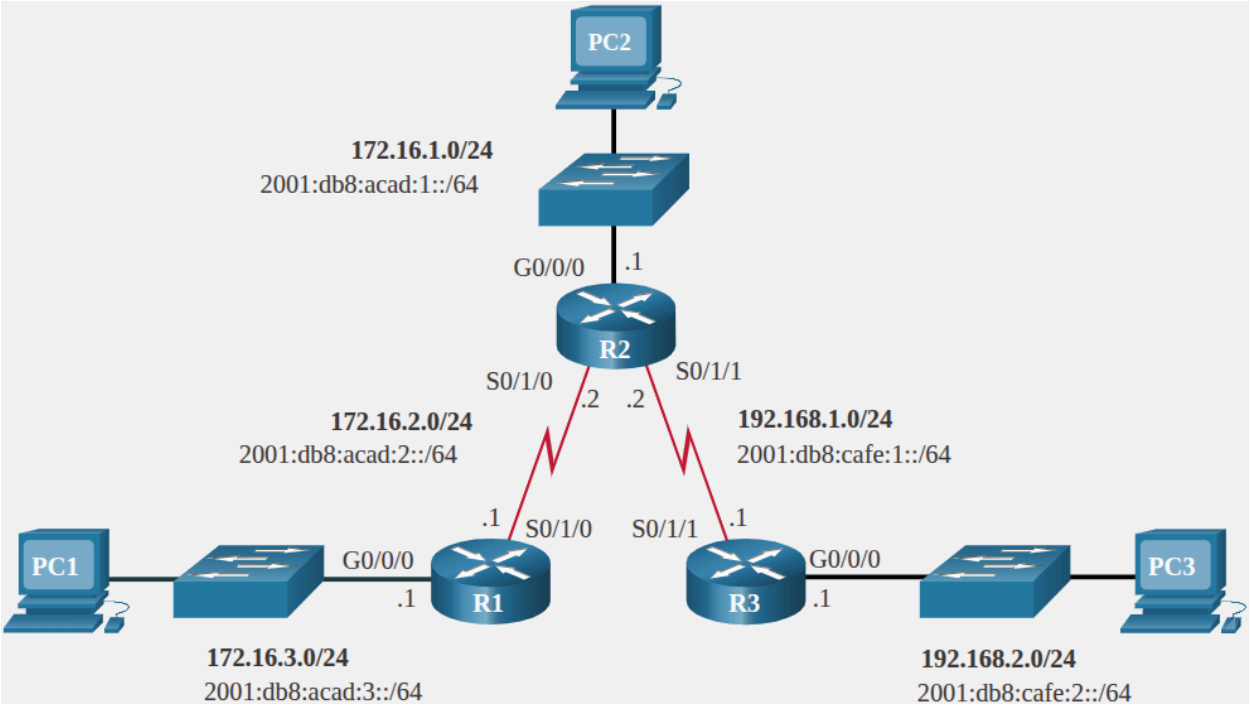

Three directly connected IPv3 static routes are configured on R1 using the exit interface

```
R1(config)# ip route 172.16.1.0 255.255.255.0 s0/1/0
R1(config)# ip route 192.168.1.0 255.255.255.0 s0/1/0
R1(config)# ip route 192.168.2.0 255.255.255.0 s0/1/0
```

The IPv4 routing table for R1 shows that when a packet is destined for the 192.168.2.0/24 network, R1 looks for a match in the routing table, and finds that it can forward the packet out of its Serial 0/1/0 interface

Note - Using a next-hop address is generally recommended. Directly connected static routes should only be used with point-to-point serial interfaces

```
R1# show ip route | begin Gateway
Gateway of last resort is not set
      172.16.0.0/16 is variably subnetted, 5 subnets, 2 masks
S        172.16.1.0/24 is directly connected, Serial0/1/0
C        172.16.2.0/24 is directly connected, Serial0/1/0
L        172.16.2.1/32 is directly connected, Serial0/1/0
C        172.16.3.0/24 is directly connected,
GigabitEthernet0/0/0
L        172.16.3.1/32 is directly connected,
GigabitEthernet0/0/0
S    192.168.1.0/24 is directly connected, Serial0/1/0
S    192.168.2.0/24 is directly connected, Serial0/1/0
```

IPv6 Directly Connected Static Route

In the example, three directly connected IPv6 static routes are configured on R1 using the exit interface

```
R1(config)# ipv6 route 2001:db8:acad:1::/64 s0/1/0
R1(config)# ipv6 route 2001:db8:cafe:1::/64 s0/1/0
R1(config)# ipv6 route 2001:db8:cafe:2::/64 s0/1/0
```

The IPv6 routing table for R1 in the example output shows that when a packet is destined for the 2001:db8:cafe:2::/64 network, R1 looks for a match in the routing table and finds that it can forward the packet out of its Serial 0/1/0 interface

Note - Using a next-hop address is generally recommended. Directly connected static routes should only be used with point-to-point serial interfaces

```
R1# show ipv6 route
IPv6 Routing Table - default - 8 entries
Codes: C - Connected, L - Local, S - Static, U - Per-user Static route
       B - BGP, R - RIP, H - NHRP, I1 - ISIS L1
       I2 - ISIS L2, IA - ISIS interarea, IS - ISIS summary, D - EIGRP
       EX - EIGRP external, ND - ND Default, NDp - ND Prefix, DCE - Destination
       NDr - Redirect, RL - RPL, O - OSPF Intra, OI - OSPF Inter
       OE1 - OSPF ext 1, OE2 - OSPF ext 2, ON1 - OSPF NSSA ext 1
       ON2 - OSPF NSSA ext 2, la - LISP alt, lr - LISP site-registrations
       ld - LISP dyn-eid, lA - LISP away, le - LISP extranet-policy
       a - Application
S   2001:DB8:ACAD:1::/64 [1/0]
       via Serial0/1/0, directly connected
```

```
C   2001:DB8:ACAD:2::/64 [0/0]
      via Serial0/1/0, directly connected
L   2001:DB8:ACAD:2::1/128 [0/0]
      via Serial0/1/0, receive
C   2001:DB8:ACAD:3::/64 [0/0]
      via GigabitEthernet0/0/0, directly connected
L   2001:DB8:ACAD:3::1/128 [0/0]
      via GigabitEthernet0/0/0, receive
S   2001:DB8:CAFE:1::/64 [1/0]
      via Serial0/1/0, directly connected
S   2001:DB8:CAFE:2::/64 [1/0]
      via Serial0/1/0, directly connected
L   FF00::/8 [0/0]
      via Null0, receive
R1#
```

IPv4 Fully Specified Static Route

in a fully specified static route, both the exit interface and the next-hop IP address are specified.

- This form of static route is used when the exit interface is a multi-access interface and it is necessary to explicitly identify the next hop
- The next hop must be directly connected to the specified exit interface
    - Using an exit interface is optional, however it is necessary to use a next-hop address

Suppose that the network link between R1 and R2 is an Ethernet link and that the GigabitEthernet 0/0/1 interface of R1 is connected to that network

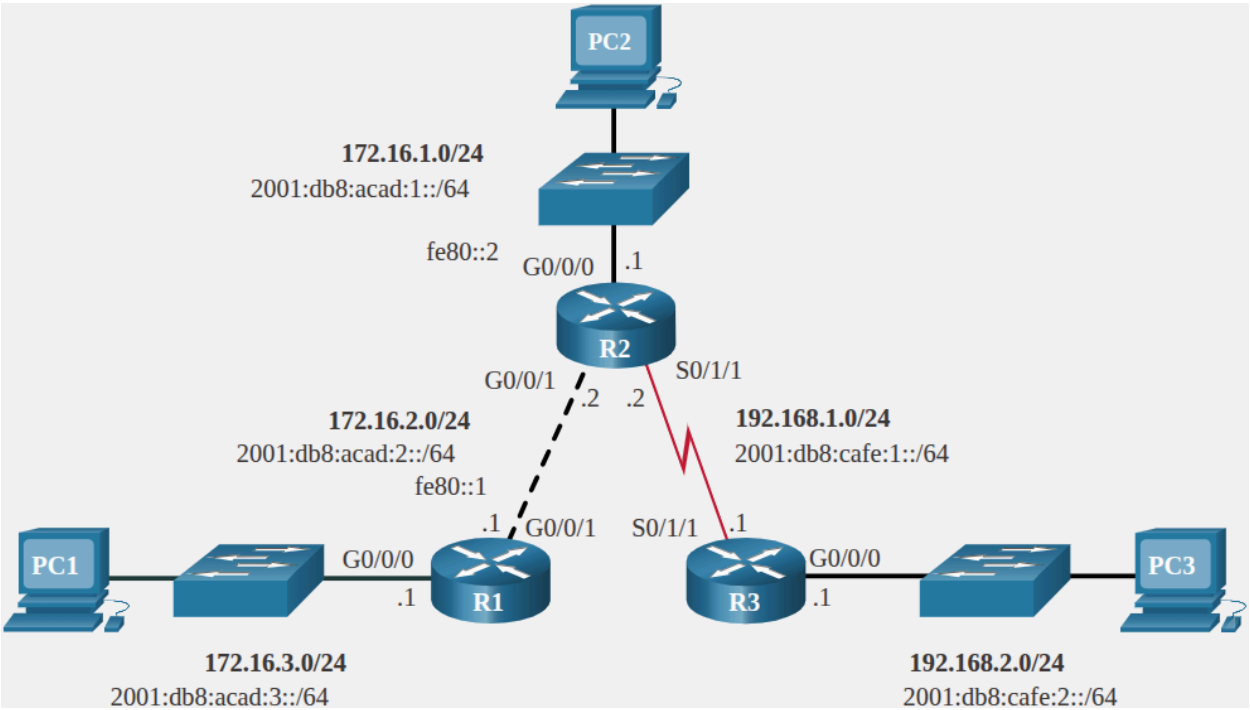

The difference between an Ethernet multi-access network and a point-to-point serial network:

- A point-to-point serial network has only one other device on that network, the router at the other end of the link
- With Ethernet networks, there may be many different devices sharing the same multi-access network, including hosts and even multiple routers

It is recommended that when the exit interface is an Ethernet network, that the static route includes a next-hop address.

- Can also use a fully specified static route that includes both the exit interface and the next-hop address

```
R1(config)# ip route 172.16.1.0 255.255.255.0 GigabitEthernet 0/0/1 172.16.2.2
R1(config)# ip route 192.168.1.0 255.255.255.0 GigabitEthernet 0/0/1 172.16.2.2
R1(config)# ip route 192.168.2.0 255.255.255.0 GigabitEthernet 0/0/1 172.16.2.2
```

When forwarding packets to R2, the exit interface is GigabitEthernet 0/0/1 and the next-hop IPv4 address is 172.16.2.2 as shown in the **show ip route** output from R1

```
R1# show ip route | begin Gateway
Gateway of last resort is not set
      172.16.0.0/16 is variably subnetted, 5 subnets, 2 masks
S        172.16.1.0/24 [1/0] via 172.16.2.2, GigabitEthernet0/0/1
C        172.16.2.0/24 is directly connected, GigabitEthernet0/0/1
L        172.16.2.1/32 is directly connected, GigabitEthernet0/0/1
C        172.16.3.0/24 is directly connected, GigabitEthernet0/0/0
L        172.16.3.1/32 is directly connected, GigabitEthernet0/0/0
S     192.168.1.0/24 [1/0] via 172.16.2.2, GigabitEthernet0/0/1
S     192.168.2.0/24 [1/0] via 172.16.2.2, GigabitEthernet0/0/1
```

IPv6 Fully Specified Static Route

In a fully specified IPv6 static route, both the exit interface and the next-hop IPv6 address are specified.

There is a situation in IPv6 when a fully specified static route must be used.

- If the IPv6 static route uses an IPv6 link-local address as the next-hop address, use a fully specified static route

Example of fully specified IPv6 static route using an IPv6 link-local address as next-hop address

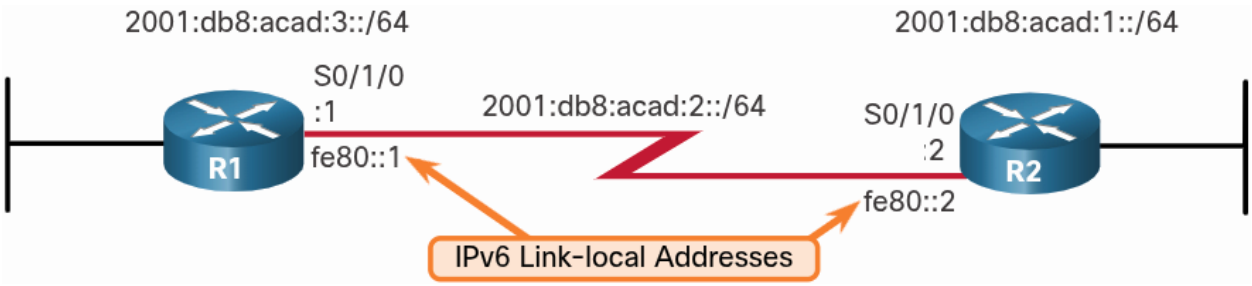

```
R1(config)# ipv6 route 2001:db8:acad:1::/64 fe80::2
%Interface has to be specified for a link-local nexthop
R1(config)# ipv6 route 2001:db8:acad:1::/64 s0/1/0 fe80::2
```

In example, a fully specified static route is configured using the link-local address of R2 as the next-hop address. Notice that IOS required that an exit interface be specified.

The reason a fully specified static route must be used is because IPv6 link-local addresses are not contained in the IPv6 routing table.

- Link-local addresses are only unique on a given link or network.
- The next-hop link-local address may be a valid address on multiple networks connected to the router
    - Therefore, it is necessary that the exit interface be included

The following shows the IPv6 routing table entry for this route. Notice that both the next-hop link-local address and the exit interface are included

```
R1# show ipv6 route static | begin 2001:db8:acad:1::/64
S   2001:DB8:ACAD:1::/64 [1/0]
    via FE80::2, Seria0/1/0
```

Verify a Static Route

Along with **show ip route**, **show ipv6 route**, **ping**, and **traceroute,** other commands to verify:

- **show ip route static**
- **show ip route** *network*
- **show running-config | section ip route**

Replace **ip** with **ipv6** for the IPv6 version

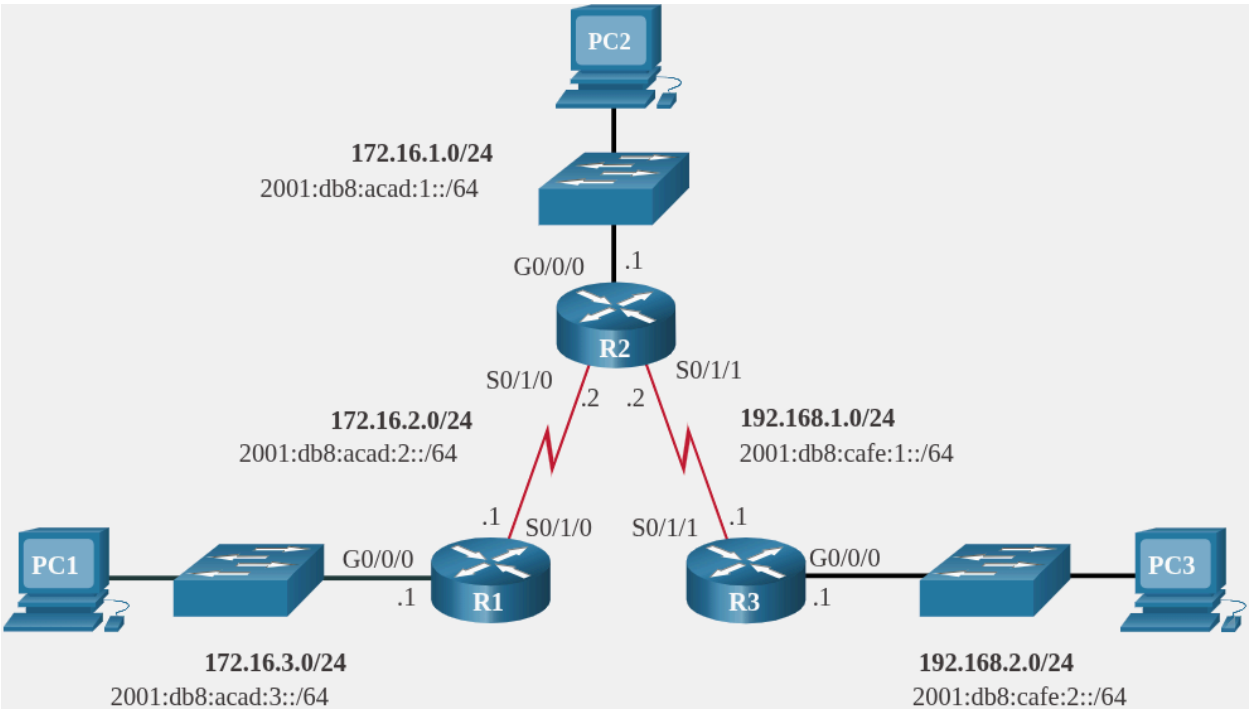

Display Only IPv4 Static Routes

- Note where the filter begins the output, excluding all the codes

```
R1# show ip route static | begin Gateway
Gateway of last resort is not set
      172.16.0.0/16 is variably subnetted, 5 subnets, 2 masks
S        172.16.1.0/24 [1/0] via 172.16.2.2
S     192.168.1.0/24 [1/0] via 172.16.2.2
S     192.168.2.0/24 [1/0] via 172.16.2.2
R1#
```

Display a Specific IPv4 Network

```
R1# show ip route 192.168.2.1
Routing entry for 192.168.2.0/24
  Known via "static", distance 1, metric 0
  Routing Descriptor Blocks:
  * 172.16.2.2
      Route metric is 0, traffic share count is 1
R1#
```

Display the IPv4 Static Route Configuration

```
R1# show running-config | section ip route
ip route 172.16.1.0 255.255.255.0 172.16.2.2
ip route 192.168.1.0 255.255.255.0 172.16.2.2
ip route 192.168.2.0 255.255.255.0 172.16.2.2
R1#
```

Display Only IPv6 Static Routes

- Note where the filter begins the output, excluding all the codes

```
R1# show ipv6 route static
IPv6 Routing Table - default - 8 entries
Codes: C - Connected, L - Local, S - Static, U - Per-user Static route
       B - BGP, R - RIP, H - NHRP, I1 - ISIS L1
       I2 - ISIS L2, IA - ISIS interarea, IS - ISIS summary, D - EIGRP
       EX - EIGRP external, ND - ND Default, NDp - ND Prefix, DCE -
Destination
       NDr - Redirect, RL - RPL, O - OSPF Intra, OI - OSPF Inter
       OE1 - OSPF ext 1, OE2 - OSPF ext 2, ON1 - OSPF NSSA ext 1
       ON2 - OSPF NSSA ext 2, la - LISP alt, lr - LISP site-registrations
       ld - LISP dyn-eid, lA - LISP away, le - LISP extranet-policy
       a - Application
S   2001:DB8:ACAD:1::/64 [1/0]
     via 2001:DB8:ACAD:2::2
S   2001:DB8:CAFE:1::/64 [1/0]
     via 2001:DB8:ACAD:2::2
S   2001:DB8:CAFE:2::/64 [1/0]
     via 2001:DB8:ACAD:2::2
R1#
```

Display a Specific IPv6 Network

```
R1# show ipv6 route 2001:db8:cafe:2::
Routing entry for 2001:DB8:CAFE:2::/64
  Known via "static", distance 1, metric 0
  Route count is 1/1, share count 0
  Routing paths:
    2001:DB8:ACAD:2::2
      Last updated 00:23:55 ago

R1#
```

Display the IPv6 Static Route Configuration

```
R1# show running-config | section ipv6 route
ipv6 route 2001:DB8:ACAD:1::/64 2001:DB8:ACAD:2::2
ipv6 route 2001:DB8:CAFE:1::/64 2001:DB8:ACAD:2::2
ipv6 route 2001:DB8:CAFE:2::/64 2001:DB8:ACAD:2::2
R1#
```

Default Static Route

A default route is a static route that matches all packets

- Instead of routers storing routes for all of the networks in the internet, they can store a single default route to represent any network that is not in the routing table

Commonly used when connecting an edge router to a service provider network, or a stub router (a router with only one upstream neighbor router)

* Routers commonly use default routes that are either configured locally or learned from another router, using a dynamic routing protocol

- A default route does not require any far-left bits to match between the default route and the destination IP address
- A default route is used when no other routes in the routing table match the destination IP address of the packet
    - In other words, if a more specific match does not exist, then the default route is used as the Gateway of Last Resort

R1 only needs to know about directly connected networks,

For all other networks it can use a default static route pointing to R2

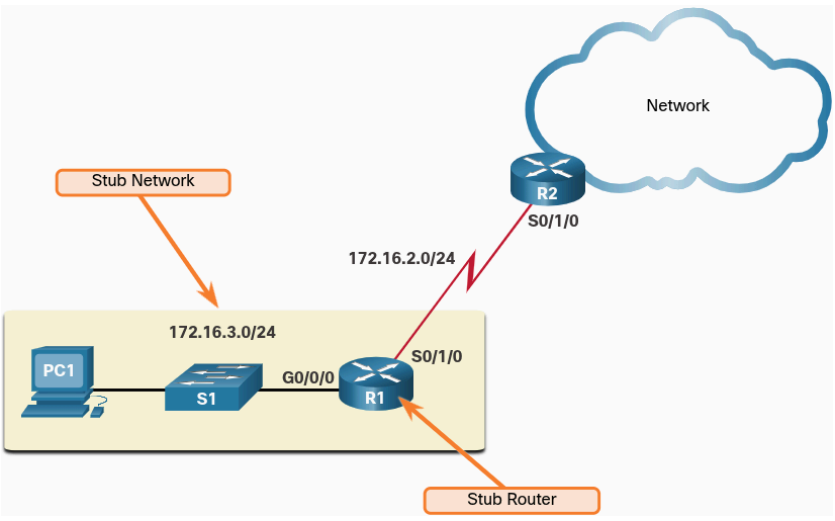

IPv4 Default Static Route
The command syntax for an IPv4 default static route is similar to any other IPv4 static route, except that the network address is **0.0.0.0** and the subnet mask is **0.0.0.0**. The 0.0.0.0.0.0.0.0 in the route will match any network address.

Note - An IPv4 default static route is commonly referred to as a quad-zero route

The basic command syntax for IPv4 default route:

```
Router(config)# ip route 0.0.0.0 0.0.0.0 {ip-address | exit-intf}
```

IPv6 Default Static Route
The command syntax for an IPv4 default static route is similar to any other IPv6 static route, except that the IPv6-prefix/prefix-length is ::**/0**, which matches all routes

The basic syntax for an IPv6 default static route:

```
Router(config)# ipv6 route ::/0 {ipv6-address | exit-intf}
```

Configure a Default Static Route

In figure, R1 could be configured with three static routes, one to reach out of the remote networks in the example topology

- However, R1 is a stub router because it is only connected to R2
- Therefore, it would be more efficient to configure a single default static route

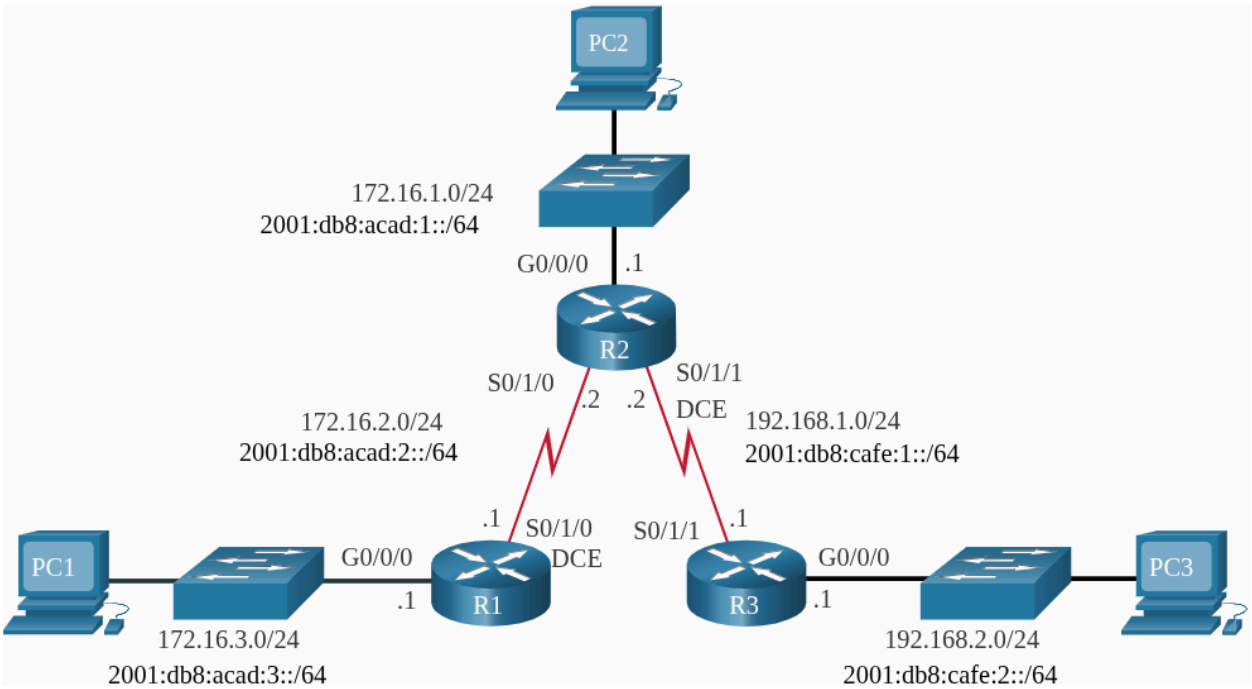

Example shows an IPv4 default route configured on R1. With configuration in example, any packets not matching more specific route entries are forwarded to R2 at 172.16.2.2

```
R1(config)# ip route 0.0.0.0 0.0.0.0 172.16.2.2
```

An IPv6 default static route is configured in a similar fashion. With this configuration any packets not matching more specific IPv6 route entries are forwarded to R2 at 2001:db8:acad:2::2

```
R1(config)# ipv6 route ::/0 2001:db8:acad:2::2
```

Verify a Default Static Route

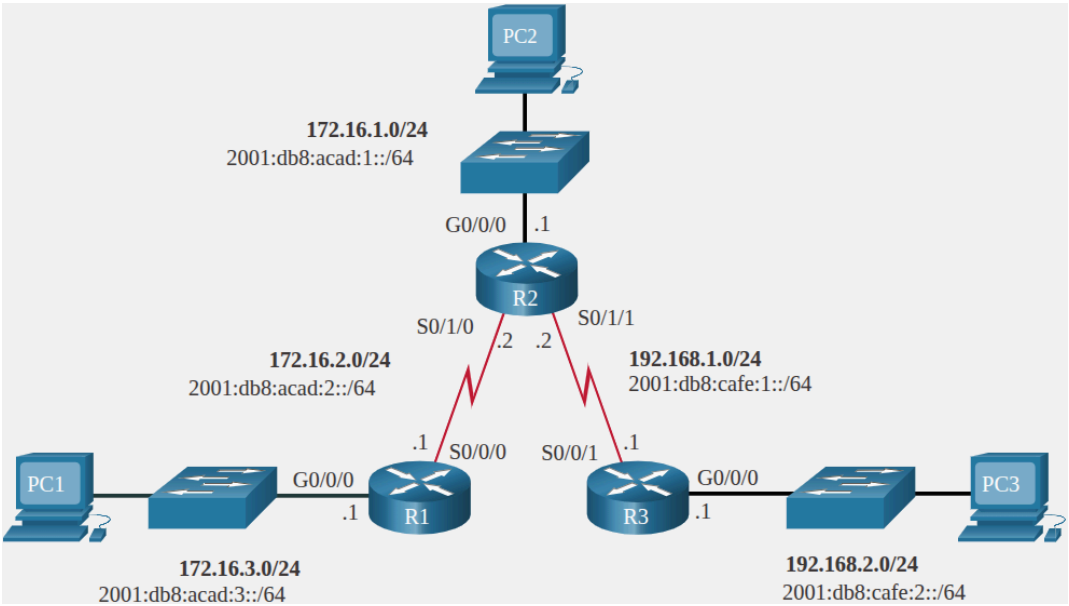

Verify IPv4 Default Static Route

The **show ip route static** command output from R1 displays the contents of the static routes in the routing table.

Note the asterisk (*) next to the route with code 'S'.

Displayed in the codes table in the **show ip route** output, the asterisk indicates that this static route is a candidate default route, which is why it's selected as the Gateway of Last Resort.

```
R1# show ip route static
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
E1 - OSPF external type 1, E2 - OSPF external type 2
i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
ia - IS-IS inter area, * - candidate default, U - per-user static route
o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP
+ - replicated route, % - next hop override

Gateway of last resort is 172.16.2.2 to network 0.0.0.0

S* 0.0.0.0/0 [1/0] via 172.16.2.2
R1#
```

Verify IPv6 Default Static Route

Example: the **show ipv6 route static** command output to display contents of routing table

```
R1# show ipv6 route static
IPv6 Routing Table - default - 8 entries
Codes: C - Connected, L - Local, S - Static, U - Per-user Static route
       B - BGP, R - RIP, H - NHRP, I1 - ISIS L1
       I2 - ISIS L2, IA - ISIS interarea, IS - ISIS summary, D - EIGRP
       EX - EIGRP external, ND - ND Default, NDp - ND Prefix, DCE - Destination
       NDr - Redirect, RL - RPL, O - OSPF Intra, OI - OSPF Inter
       OE1 - OSPF ext 1, OE2 - OSPF ext 2, ON1 - OSPF NSSA ext 1
       ON2 - OSPF NSSA ext 2, la - LISP alt, lr - LISP site-registrations
       ld - LISP dyn-eid, lA - LISP away, le - LISP extranet-policy
       a - Application
S   ::/0 [1/0]
       via 2001:DB8:ACAD:2::2

R1#
```

Notice that the static default route configuration uses the /0 mask for IPv4 default routes and the ::/0 prefix for IPv5 default routes.

Remember that the IPv4 subnet mask and IPv6 prefix-length in a routing table determines how many bits must match between the destination IP address of the packet and the route in the routing table.

A /0 mask or ::/0 prefix indicates that none of the bits are required to match. As long as a more specific match does not exist, the default static route matches all packets.

Floating Static Routes

- Static routes that are used to provide a backup path to a primary static or dynamic route, in the event of a link failure.
- Only used when the primary route is not available

To accomplish this, the floating static route is configured with a higher administrative distance than the primary route. The administrative distance represents the trustworthiness of a route

- If multiple paths to the destination exist, the router will choose the path with the lowest administrative distance

Example: Assume that an admin wants to create a floating static route as a backup to an EIGRP-learned route.

- The floating static route must be configured with a higher administrative distance than EIGRP. EIGRP has an administrative distance of 90
- If the floating static route is configured with an administrative distance of 95, the dynamic route learned through EIGRP is preferred to the floating static route
- If the EIGRP-learned route is lost, the floating static route is used in its place

In figure, the branch router typically forwards all traffic to the HQ router over the private WAN link.

- In example, the routers exchange route information using EIGRP.
- A floating static route, with an administrative distance of 91 or higher, could be configured to serve as a backup route.
- If the private WAN link fails and the EIGRP route disappears from the routing table, the router selects the floating static route as the best path to reach the HQ LAN.

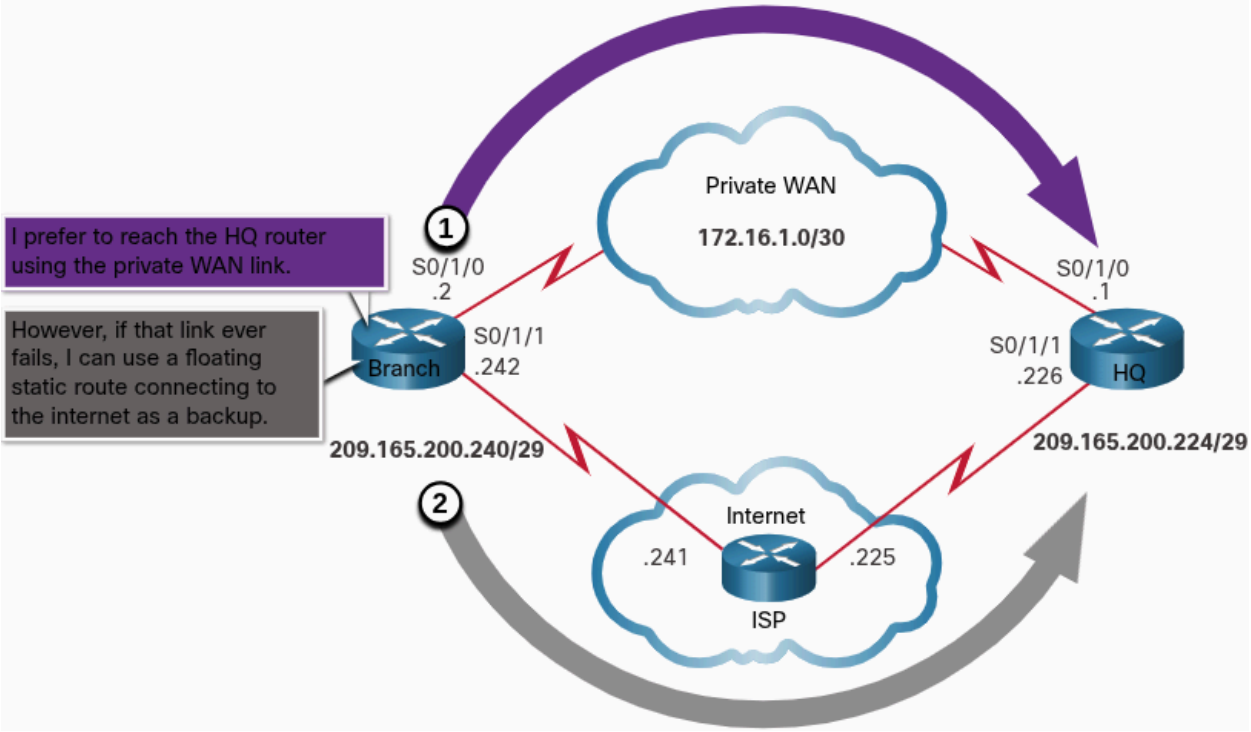

1. *A route learned through dynamic routing is preferred*
2. *If a dynamic route is lost, the floating static route will be used*

By default, static routes have an administrative distance of 1, making them preferable to routes learned from dynamic routing protocols.

- For example, the administrative distances of some common interior gateway dynamic protocols are:
    - EIGRP = 90
    - OSPF = 110
    - IS-IS = 115

The administrative distance of a static route can be increased to make the route less desirable than that of another static route or a route learned through a dynamic routing protocol

In this way, the static route “floats” is not used when the route with the better administrative distance is active. However, if the preferred route is lost, the floating static route can take over, and traffic can be sent through this alternate route.

Configure IPv4 and IPv6 Floating Static Routes

IP floating static routes are configured by using the **distance** argument to specify an administrative distance. If no administrative distance is configured, the default value (1) is used.

In this scenario, the preferred default route from R1 is to R2. The connection to R3 should be used for backup only

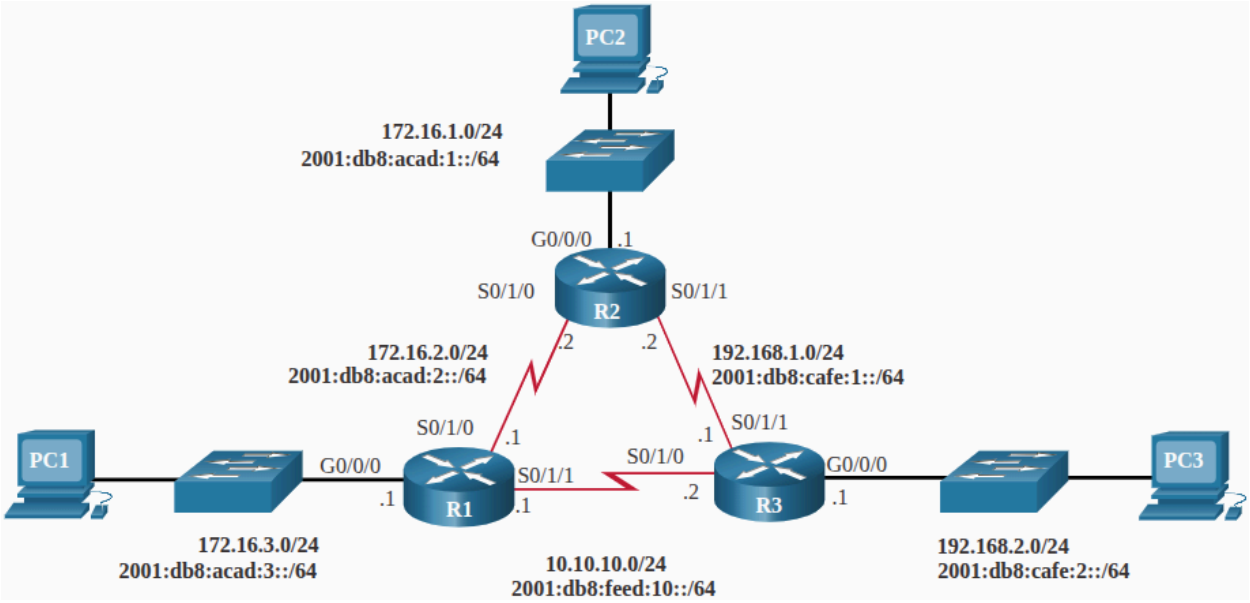


```
R1(config)# ip route 0.0.0.0 0.0.0.0 172.16.2.2
R1(config)# ip route 0.0.0.0 0.0.0.0 10.10.10.2 5
R1(config)# ipv6 route ::/0 2001:db8:acad:2::2
R1(config)# ipv6 route ::/0 2001:db8:feed:10::2 5
```

R1 is configured with IPv4 and IPv6 default static routes pointing to R2. Because no administrative distance is configured, the default value (1) is used for these static routes.

R1 is also configured with IPv4 and IPv6 floating static routes pointing to R3 with an administrative distance of 5

- This value is greater than a default value of 1 and therefore; this route floats and is not present in the routing table unless the preferred route fails

The **show ip route** and **show ipv6 route** output verifies that the default routes to R2 are installed in the routing table. Not that the IPv4 floating static route to R3 is not present

```
R1# show ip route static | begin Gateway
Gateway of last resort is 172.16.2.2 to network 0.0.0.0

S*   0.0.0.0/0 [1/0] via 172.16.2.2
R1# show ipv6 route static | begin S :
S   ::/0 [1/0]
      via 2001:DB8:ACAD:2::2
R1#
```

Use the **show run** command to verify that floating static routes are in the configuration

- Example: the following command output verifies that both IPv6 static default routes are in the running configuration

```
R1# show run | include ipv6 route
ipv6 route ::/0 2001:db8:feed:10::2 5
ipv6 route ::/0 2001:db8:acad:2::2
R1#
```

Test the Floating Static Route

In figure, what would happen if R2 failed?

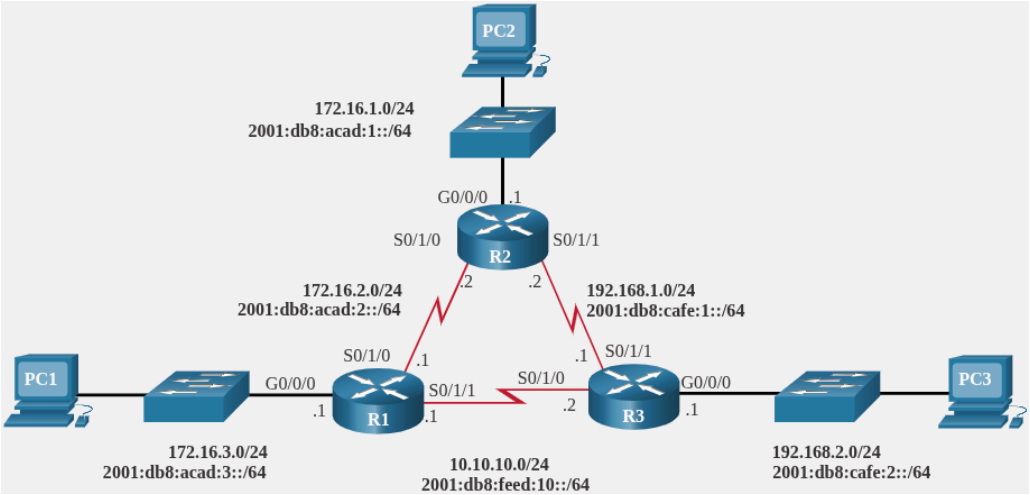

To simulate this failure, both serial interfaces of R2 are shut down

```
R2(config)# interface s0/1/0
R2(config-if)# shut
*Sep 18 23:36:27.000: %LINK-5-CHANGED: Interface Serial0/1/0, changed state to
administratively down
*Sep 18 23:36:28.000: %LINEPROTO-5-UPDOWN: Line protocol on Interface Serial0/1/0,
changed state to down
R2(config-if)# interface s0/1/1
R2(config-if)# shut
*Sep 18 23:36:41.598: %LINK-5-CHANGED: Interface Serial0/1/1, changed state to
administratively down
*Sep 18 23:36:42.598: %LINEPROTO-5-UPDOWN: Line protocol on Interface Serial0/1/1,
changed state to down
```

Notice R1 automatically generates messages indicating that the serial interface to R2 is down

```
R1#
*Sep 18 23:35:48.810: %LINK-3-UPDOWN: Interface Serial0/1/0, changed state to
down
R1#
*Sep 18 23:35:49.811: %LINEPROTO-5-UPDOWN: Line protocol on Interface
Serial0/1/0, changed state to down
R1#
```

A look at the IP routing tables of R1 verifies that the floating static default routes are now installed as the default routes and are pointing to R3 as the next-hop router

```
R1# show ip route static | begin Gateway
Gateway of last resort is 10.10.10.2 to network 0.0.0.0
S*    0.0.0.0/0 [5/0] via 10.10.10.2
R1# show ipv6 route static | begin ::
S   ::/0 [5/0]
     via 2001:DB8:FEED:10::2
R1#
```

Host Routes

A host route is an IPv4 address with a 32-bit mask, or an IPv6 address with a 128-bit mask. The following shows three ways a host route can be added to the routing table:

- Automatically installed when an IP address is configured on the router
- Configured as a static host route
- Host route automatically obtained through other methods

Automatically Installed Host Routes

Cisco IOS automatically installs a host route, also known as a local host route, when an interface address is configured on the router.

A host route allows for a more efficient process for packets that are directed to the router itself, rather than packet forwarding

- This is an addition to the connected route, designated with a **C** in the routing table for the network address of the interface

When an active interface on a router is configured with an IP address, a local host route is automatically added to the routing table. The local routes are marked with **L** in the output of the routing table

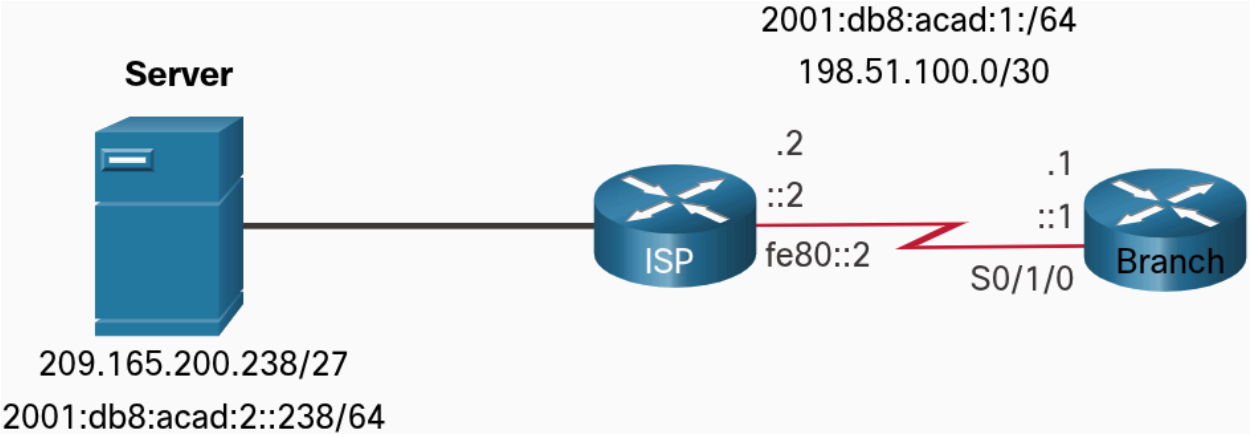

The IP addresses assigned to the Branch Serial 0/1/0 interface are 198.51.100.1/30 and 2001:db8:acad:1::1/64. The local routes for the interface are installed by the IOS in the IPv4 and IPv6 routing tables

```
Branch# show ip route | begin Gateway
Gateway of last resort is not set
      198.51.100.0/24 is variably subnetted, 2 subnets, 2 masks
C        198.51.100.0/30 is directly connected, Serial0/1/0
L        198.51.100.1/32 is directly connected, Serial0/1/0
Branch# show ipv6 route | begin ::
C   2001:DB8:ACAD:1::/64 [0/0]
     via Serial0/1/0, directly connected
L   2001:DB8:ACAD:1::1/128 [0/0]
     via Serial0/1/0, receive
L   FF00::/8 [0/0]
     via Null0, receive
```

Static Host Routes

A host route can be a manually static route to direct traffic to a specified destination device. The static route uses a destination IP address and a 255.255.255.255 (/32) mask for IPv4 host routes, and a /128 prefix length for IPv6 host routes

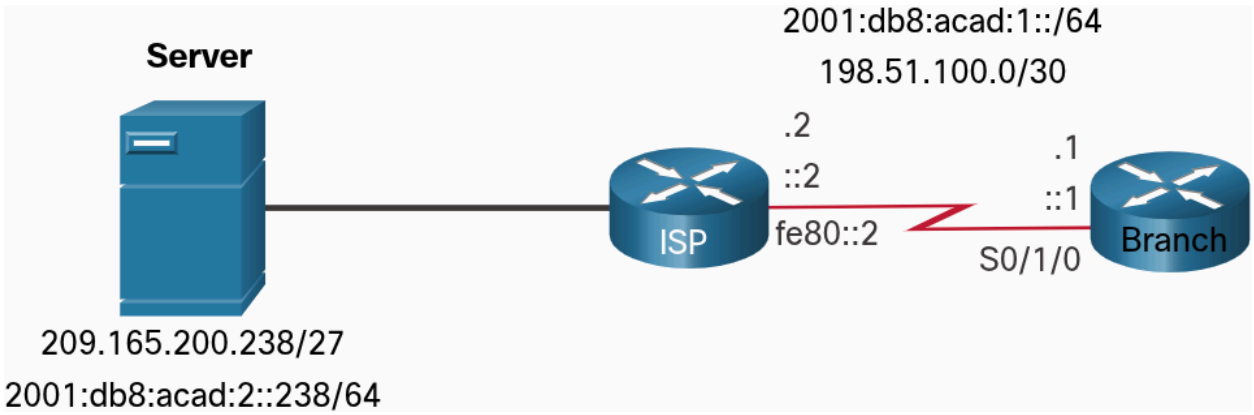

Configure Static Host Routes

Example shows the IPv4 and IPv6 static host route configuration on the Branch router to access the server

```
Branch(config)# ip route 209.165.200.238 255.255.255.255 198.51.100.2
Branch(config)# ipv6 route 2001:db8:acad:2::238/128
2001:db8:acad:1::2
Branch(config)# exit
Branch#
```

Verify Static Host Routes

A review of both the IPv4 and IPv6 route tables verifies that routes are active

```
Branch# show ip route | begin Gateway
Gateway of last resort is not set
      198.51.100.0/24 is variably subnetted, 2 subnets, 2 masks
C        198.51.100.0/30 is directly connected, Serial0/1/0
L        198.51.100.1/32 is directly connected, Serial0/1/0
      209.165.200.0/32 is subnetted, 1 subnets
S        209.165.200.238 [1/0] via 198.51.100.2
Branch# show ipv6 route
(Output omitted)
C   2001:DB8:ACAD:1::/64 [0/0]
     via Serial0/1/0, directly connected
L   2001:DB8:ACAD:1::1/128 [0/0]
     via Serial0/1/0, receive
S   2001:DB8:ACAD:2::238/128 [1/0]
     via 2001:DB8:ACAD:1::2
Branch#
```

Configure IPv6 Static Host Route with Link-Local Next-Hop

For IPv6 static routes, the next-hop address can be the link-local address of the adjacent router

However, you must specify an interface type and an interface number when using a link-local address as the next hop, as shown in example.

First, the original IPv6 static host route is removed, then a fully specified route configured with the IPv6 address of the server and the IPv6 link-local address of the ISP router

```
Branch(config)# no ipv6 route 2001:db8:acad:2::238/128
2001:db8:acad:1::2
Branch(config)# ipv6 route 2001:db8:acad:2::238/128 serial 0/1/0
fe80::2
Branch# show ipv6 route | begin ::
C   2001:DB8:ACAD:1::/64 [0/0]
     via Serial0/1/0, directly connected
L   2001:DB8:ACAD:1::1/128 [0/0]
     via Serial0/1/0, receive
S   2001:DB8:ACAD:2::238/128 [1/0]
     via FE80::2, Serial0/1/0
Branch#
```

15.6.3 What did I learn in this module?

Static Routes

Static routes can be configured for IPv4 and IPv6. Both protocols support the following types of static routes: standard static route, default static route, floating static route, and summary static route. Static routes are configured using the ip route and ipv6 route global configuration commands. When configuring a static route, the next hop can be identified by an IP address, exit interface, or both. How the destination is specified creates one of the three following types of static route: next-hop, directly connected, and fully specified. IPv4 static routes are configured using the following global configuration command: ip route *network-address subnet-mask* { *ip-address* | *exit-intf* [*ip=address*] } [distance]. IPv6 static routes are configured using the following global configuration command: ipv6 route *ipv6-prefix/prefix-length* { *ipv6-address* | *exit-intf* [*ipv6-address*]} [distance]. The command to start an IPv4 routing table is show ip route | begin Gateway. The command to start an IPv6 routing table is show ipv6 route | begin C.

Configure IP Static Routes

In a next-hop static route, only the next-hop IP address is specified. The exit interface is derived from the next hop. When configuring a static route, another option is to use the exit interface to specify the next-hop address. Directly connected static routes should only be used with point-to-point serial interfaces. In a fully specified static route, both the exit interface and the next-hop IP address are specified. This form of static route is used when the exit interface is a multi-access interface and it is necessary to explicitly identify the next hop. The next hop must be directly connected to the specified exit interface. In a fully specified IPv6 static route, both the exit interface and the next-hop IPv6 address are specified. Along with show ip route, show ipv6 route, ping and traceroute, other useful commands to verify static routes include: show ip route static, show ip route *network*, and show running-config | section ip route. Replace ip with ipv6 for the IPv6 versions of the command.

Configure IP Default Static Routes

A default route is a static route that matches all packets. A default route does not require any far-left bits to match between the default route and the destination IP address. Default static routes are commonly used when connecting an edge router to a service provider network, and a stub router. The command syntax for an IPv4 default static route is similar to any other IPv4 static route, except that the network address is 0.0.0.0 and the subnet mask is 0.0.0.0. The 0.0.0.0 0.0.0.0 in the route will match any network

address. The command syntax for an IPv6 default static route is similar to any other IPv6 static route, except that the ipv6-prefix/prefix-length is ::/0, which matches all routes. To verify an IPv4 default static route, use the show ip route static command. For IPV6 use the show ipv6 route static command.

Configure Floating Static Routes

Floating static routes are static routes that are used to provide a backup path to a primary static or dynamic route in the event of a link failure. The floating static route is configured with a higher administrative distance than the primary route. By default, static routes have an administrative distance of 1, making them preferable to routes learned from dynamic routing protocols. The administrative distances of some common interior gateway dynamic routing protocols are EIGRP = 90, OSPF = 110, and IS-IS = 115. IP floating static routes are configured by using the distance argument to specify an administrative distance. If no administrative distance is configured, the default value (1) is used. The show ip route and show ipv6 route output verifies that the default routes to a router are installed in the routing table.

Configure Static Host Routes

A host route is an IPv4 address with a 32-bit mask or an IPv6 address with a 128-bit mask. There are three ways a host route can be added to the routing table: automatically installed when an IP address is configured on the router, configured as a static host route, or automatically obtained through other methods not covered in this module. Cisco IOS automatically installs a host route, also known as a local host route, when an interface address is configured on the router. A host route can be a manually configured static route to direct traffic to a specific destination device. For IPv6 static routes, the next-hop address can be the link-local address of the adjacent router; however, you must specify an interface type and an interface number when using a link-local address as the next hop. To do this, the original IPv6 static host route is removed, then a fully specified route is configured with the IPv6 address of the server and the IPv6 link-local address of the ISP router.

Packet Processing with Static Routes

Static Routes and Packet Forwarding

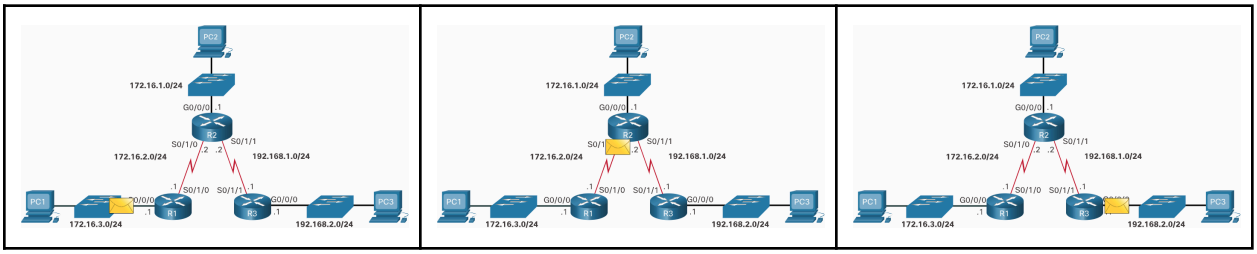

Demonstration Explanation:

1. The packet arrives on the GigabitEthernet 0/0/0 interface of R1
2. R1 does not have a specific route to the destination network, 192.168.2.0/24. Therefore, R1 uses the default static route
3. R1 encapsulates the packet in a new frame. Because the link to R2 is a point-to-point link, R1 adds an “all 1s” address for the Layer 2 destination address
4. The frame is forwarded out of the Serial 0/1/0 interface. The packet arrives on the Serial 0/1/0 interface on R2
5. R2 de-encapsulates the frame and looks for a route to the destination. R2 has a static route to 192.168.2.0/24 out of the Serial 0/1/1 interface
6. R2 encapsulates the packet in a new frame. Because the link to R3 is a point-to-point link, R2 adds an “all 1s” address for the Layer 2 destination address
7. The frame is forwarded out of the Serial 0/1/1 interface. The packet arrives on the SErial 0/1/1 interface on R3
8. R3 de-encapsulates the frame and looks for a route to the destination. R3 has a connected route to 192.168.2.0/24 out of the GigabitEthernet 0/0/0 interface
9. R3 looks up the ARP table entry for 192.168.2.10 to find the Layer 2 Media Access Control (MAC) address for PC3.
    a. If no entry exists, R3 sends an Address Resolution Protocol (ARP) request out of the GigabitEthernet 0/0/0 interface, and PC3 responds with an ARP reply, which includes the PC3 MAC address
10. R3 encapsulates the packet in a new frame with the MAC address of the GigabitEthernet 0/0/0 interface as the source Layer 2 address, and the MAC address of PC3 as the destination MAC address
11. The frame is forwarded out of GigabitEthernet 0/0/0 interface. The packet arrives on the network interface card (NIC) interface of PC3

Network Changes

Troubleshoot IPv4 Static and Default Route Configuration

Common Troubleshooting Commands

- ping
- traceroute
- show ip route
- show ip interface brief
- show cdp neighbors detail

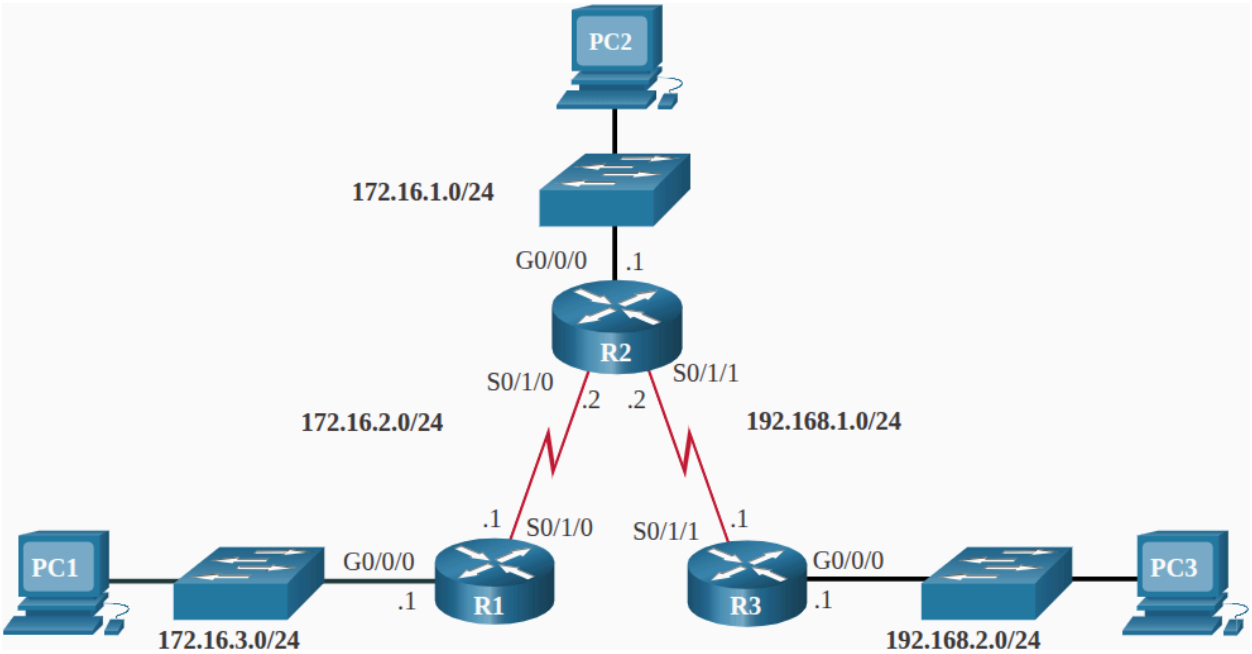

ping

Example below displays result of an extended ping from the source interface of R1 to the LAN interface of R3.

- An extended ping is an enhanced version of the ping utility.
- Enables you to specify the source IP address for the ping packets

```
R1# ping 192.168.2.1 source 172.16.3.1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.168.2.1, timeout is 2
seconds:
Packet sent with a source address of 172.16.3.1
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 3/3/5
ms
R1#
```

traceroute

Example displays result of traceroute from R1 to R3 LAN

- Not that each hop route returns an ICMP reply

```
R1# traceroute 192.168.2.1
Type escape sequence to abort.
Tracing the route to 192.168.2.1
VRF info: (vrf in name/id, vrf out name/id)
  1 172.16.2.2 1 msec 2 msec 1 msec
  2 192.168.1.1 2 msec 3 msec *
R1#
```

show ip route

Displays the route table of R1

```
R1# show ip route | begin Gateway
Gateway of last resort is not set
      172.16.0.0/16 is variably subnetted, 5 subnets, 2 masks
S        172.16.1.0/24 [1/0] via 172.16.2.2
C        172.16.2.0/24 is directly connected, Serial0/1/0
L        172.16.2.1/32 is directly connected, Serial0/1/0
C        172.16.3.0/24 is directly connected,
GigabitEthernet0/0/0
L        172.16.3.1/32 is directly connected,
GigabitEthernet0/0/0
S     192.168.1.0/24 [1/0] via 172.16.2.2
S     192.168.2.0/24 [1/0] via 172.16.2.2
R1#
```

show ip interface brief

A quick status of all interfaces on the router

```
R1# show ip interface brief
Interface              IP-Address      OK? Method Status
Protocol
GigabitEthernet0/0/0   172.16.3.1      YES manual up
up
GigabitEthernet0/0/1   unassigned      YES unset  up
up
Serial0/1/0            172.16.2.1      YES manual up
up
Serial0/1/1            unassigned      YES unset  up
```

```
up
R1#
```

show cdp neighbors

Provides a list of directly connected Cisco devices.

- Validates layer 2 (and therefore Layer 1) connectivity
- Example: If a neighbor device is listed in the command output, but cannot be pinged, then Layer 3 addressing should be investigated

```
R1# show cdp neighbors
Capability Codes: R - Router, T - Trans Bridge, B - Source Route Bridge
                  S - Switch, H - Host, I - IGMP, r - Repeater, P - Phone,
                  D - Remote, C - CVTA, M - Two-port Mac Relay
Device ID        Local Intrfce     Holdtme    Capability  Platform  Port ID
Switch           Gig 0/0/0         129              S I   WS-C3560- Fas 0/5
R2               Ser 0/1/0         156            R S I   ISR4221/K Ser 0/1/0
Total cdp entries displayed : 3
R1#
```

Solve a Connectivity Problem

Example: The user at PC1 reports that he cannot access resources on the R3 LAN

- This can be confirmed by pinging the LAN interface of R3 using the LAN interface of R1 as the source

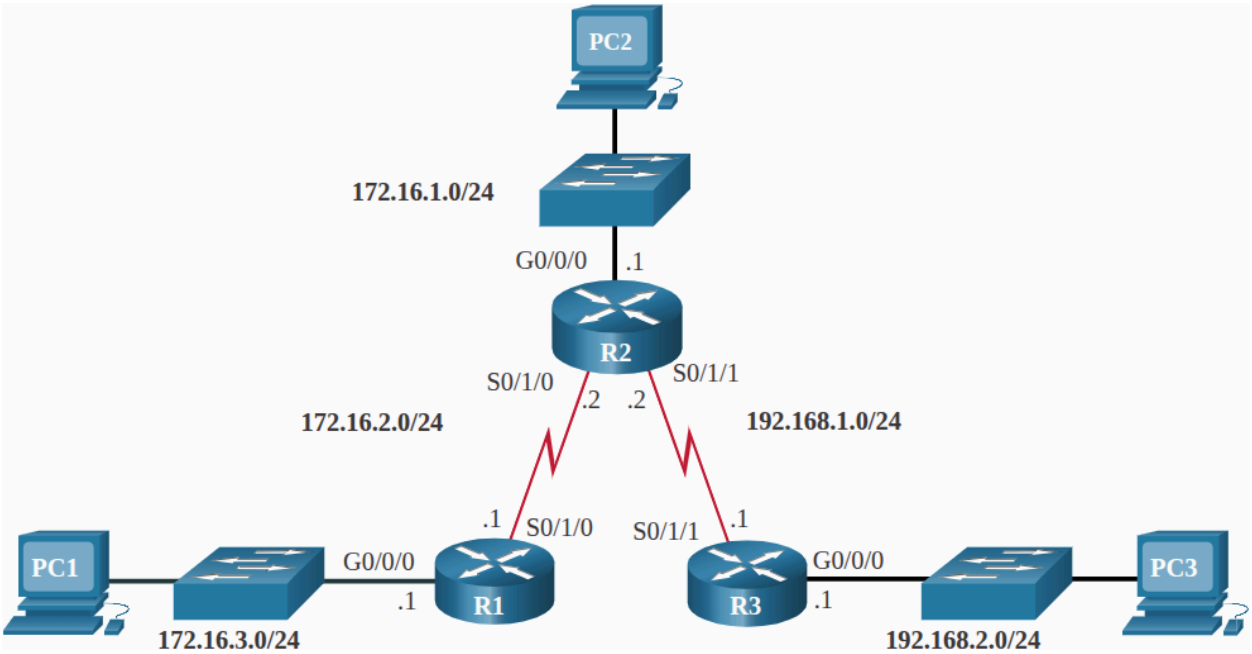

Ping the Remote LAN

Test connectivity between the two LANs from R1 instead of PC1

- Can be done by sourcing the ping from the G0/0/0 interface on R1 to the G0/0/0 interface on R3

```
R1# ping 192.168.2.1 source g0/0/0
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.168.2.1, timeout is 2
seconds:
Packet sent with a source address of 172.16.3.1
.....
Success rate is 0 percent (0/5)
```

Ping the Next-Hop Router

Next, a ping to the S0/1/0 interface on R2 is successful

- Ping is sourced from S0/1/0 interface of R1, therefore, issue is not loss of connectivity between R1 and R2

```
R1# ping 172.16.2.2
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.16.2.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 3/3/4
ms
```

Ping R3 LAN from S0/1/0

A ping from R1 to the R3 interface 192.168.2.1 is successful as well.

- Ping is sourced from the S0/1/0 interface on R1

R3 has a route back to the network between R1 and R2, 172.16.2.0/24.

- This confirms that R1 can reach the remote LAN on R3, however, packets sourced from the LAN on R1 cannot
    - This indicates that either R2 or R3 may have an incorrect or missing route to the LAN on R1

```
R1# ping 192.168.2.1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.168.2.1, timeout is 2
seconds:
!!!!!
```

```
Success rate is 100 percent (5/5), round-trip min/avg/max = 3/3/4
ms
```

Verify the R2 Routing Table

The next step is to investigate the routing tables of R2 and R3.
Notice in example that the 172.16.3.0/24 network is configured incorrectly.

- The static route to 172.16.3.0/24 network has been configured using the next-hop address 192.168.1.1
    - Therefore, packets destined for the 172.16.3.0/24 network are sent back to R3 instead of to R1

```
R2# show ip route | begin Gateway
Gateway of last resort is not set
      172.16.0.0/16 is variably subnetted, 5 subnets, 2 masks
C        172.16.1.0/24 is directly connected, GigabitEthernet0/0/0
L        172.16.1.1/32 is directly connected, GigabitEthernet0/0/0
C        172.16.2.0/24 is directly connected, Serial0/1/0
L        172.16.2.2/32 is directly connected, Serial0/1/0
S        172.16.3.0/24 [1/0] via 192.168.1.1
      192.168.1.0/24 is variably subnetted, 2 subnets, 2 masks
C        192.168.1.0/24 is directly connected, Serial0/1/1
L        192.168.1.2/32 is directly connected, Serial0/1/1
S     192.168.2.0/24 [1/0] via 192.168.1.1
R2#
```

Connect the R2 Static Route Configuration

Next, the running configuration does reveal the incorrect **ip route** statement

- The incorrect route is removed and the correct route is then entered

```
R2# show running-config | include ip route
ip route 172.16.3.0 255.255.255.0 192.168.1.1
ip route 192.168.2.0 255.255.255.0 192.168.1.1
R2#
R2# configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.
R2(config)# no ip route 172.16.3.0 255.255.255.0 192.168.1.1
R2(config)# ip route 172.16.3.0 255.255.255.0 172.16.2.1
R2(config)#
```

Verify New Static Route is installed

The routing table on R2 is checked once again to confirm the route entry to the LAN on R1, 172.16.3.0, is correct and pointing toward R1

```
R2(config) # exit
R2#
*Sep 20 02:21:51.812: %SYS-5-CONFIG_I: Configured from console by console
R2# show ip route | begin Gateway
Gateway of last resort is not set
      172.16.0.0/16 is variably subnetted, 5 subnets, 2 masks
C        172.16.1.0/24 is directly connected, GigabitEthernet0/0/0
L        172.16.1.1/32 is directly connected, GigabitEthernet0/0/0
C        172.16.2.0/24 is directly connected, Serial0/1/0
L        172.16.2.2/32 is directly connected, Serial0/1/0
S        172.16.3.0/24 [1/0] via 172.16.2.1
      192.168.1.0/24 is variably subnetted, 2 subnets, 2 masks
C        192.168.1.0/24 is directly connected, Serial0/1/1
L        192.168.1.2/32 is directly connected, Serial0/1/1
S     192.168.2.0/24 [1/0] via 192.168.1.1
R2#
```

Ping the Remote LAN Again

Next, a ping from R1 sourced from G0/0/0 is used to verify that R1 can now reach the LAN interface of R3.

- As last step in confirmation, user on PC1 should also test connectivity to the 192.168.2.0/24 LAN

```
R1# ping 192.168.2.1 source g0/0/0
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.168.2.1, timeout is 2
seconds:
Packet sent with a source address of 172.16.3.1
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 4/4/4
ms
```

16.3.3 What did I learn in this module?

Packet Processing with Static Routes

1. The packet arrives on the interface of R1.
2. R1 does not have a specific route to the destination network; therefore, R1 uses the default static route.
3. R1 encapsulates the packet in a new frame. Because the link to R2 is a point-to-point link, R1 adds an "all 1s" address for the Layer 2 destination address.
4. The frame is forwarded out of the appropriate interface. The packet arrives on the interface on R2.
5. R2 de-encapsulates the frame and looks for a route to the destination. R2 has a static route to the destination network out of one of its interfaces.
6. R2 encapsulates the packet in a new frame. Because the link to R3 is a point-to-point link, R2 adds an "all 1s" address for the Layer 2 destination address.
7. The frame is forwarded out of the appropriate interface. The packet arrives on the interface on R3.
8. R3 de-encapsulates the frame and looks for a route to the destination. R3 has a connected route to the destination network out of one of its interfaces.
9. R3 looks up the ARP table entry for the destination network to find the Layer 2 MAC address for PC3. If no entry exists, R3 sends an ARP request out of one of its interfaces, and PC3 responds with an ARP reply, which includes the PC3 MAC address.
10. R3 encapsulates the packet in a new frame with the MAC address of the appropriate interface as the source Layer 2 address and the MAC address of PC3 as the destination MAC address.
11. The frame is forwarded out of the appropriate interface. The packet arrives on the network interface card (NIC) interface of PC3.

Troubleshoot IPv4 Static and Default Route Configuration

Networks are frequently subject to events that can cause their status to change. An interface can fail, or a service provider drops a connection. Links can become oversaturated, or an

administrator may enter a wrong configuration. Common IOS troubleshooting commands include the following:

- ping
- traceroute
- show ip route
- show ip interface brief
- show cdp neighbors detail