CCNAX

Interconnecting Cisco Networking Devices: Accelerated

Student Guide Volume 2

Version 3.0

Part Number: 97-3637-01

Americas Headquarters	Asia Pacific Headquarters	Europe Headquarters
Cisco Systems, Inc. San Jose, CA	Cisco Systems (USA) Pte. Ltd. Singapore	Cisco Systems International BV Amsterdam, The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at www.cisco.com/go/offices.

Table of Contents

Module 2: Establishing Internet Connectivity

Lesson 9: Enabling Internet Connectivity

Introduction

Your boss dedicates you to the project with a customer who wants to connect to the Internet. The customer is going through the process of obtaining public IP addresses. They are asking about differences between manual IP address assignment and DHCP. Your focus is only to explain to them how DHCP can be used for address assignment by an ISP and what is needed on the customer side. The customer also wants to know about NAT and PAT in case they experience problems with public IP addressing. You should explain to them static and dynamic NAT configuration examples.

Demarcation Point

Although functions within the service provider network are not usually of concern to customers, there are some terms and concepts that you should be familiar with.

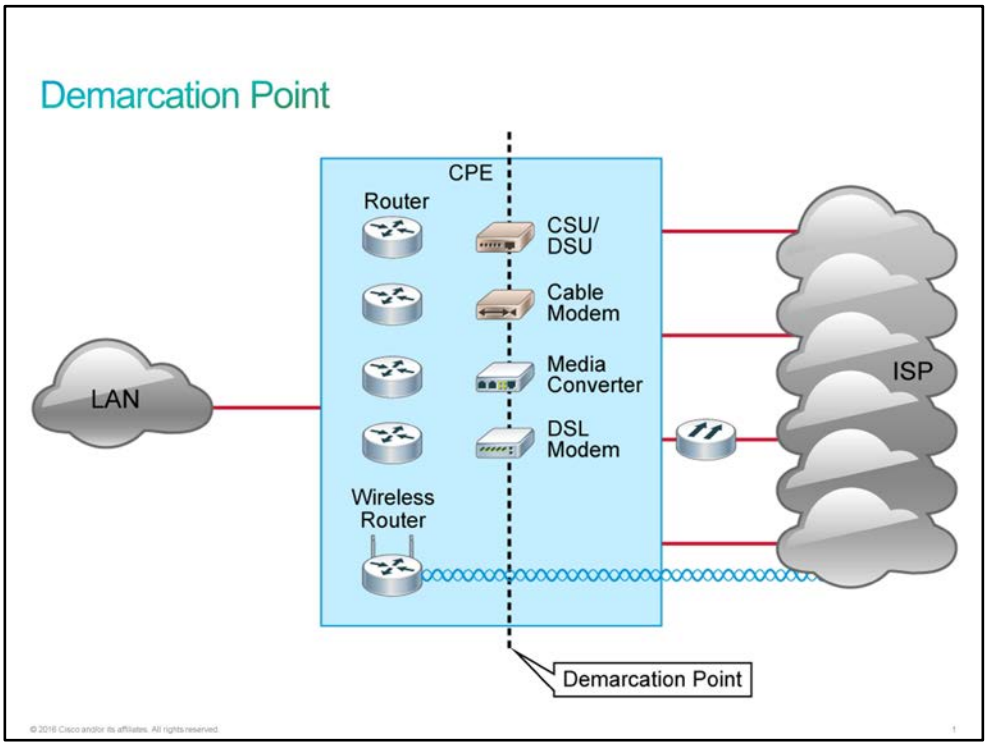

Service providers install a connection point (usually in the form of an RJ-45 jack) that physically connects a circuit to their nearest switching office. This link is known as the *demarcation point* and it represents the point at which the responsibility of the service provider is said to end. In other words, the service provider ensures that the link functions correctly up to that point. The other end of this link connects to the service provider network. These links are part of what is known as the local loop or last mile. The local loop may consist of various technologies, including DSL, cable, fiber optics, traditional twisted-pair wiring, and others.

The customer side of the demarcation point is the location of the CPE. The term CPE is often used quite loosely, but it traditionally refers to equipment that is owned and managed by the customer for connecting to the service provider network. However, many companies lease CPE from their service providers, and this equipment is still considered to be CPE. Before physically connecting to a service provider network, a company needs to determine the type of WAN service or connectivity that it requires.

Note	The exact demarcation point is different from country to country. The example that is described is for the United States.

Internet Connectivity

There are three common methods of connecting a small office to the Internet: by a copper medium, optical cable, or by wireless connection.

In wired connections, the medium is either copper, which carries electrical signals, or optical fiber, which carries signals in light. For wireless connections, the medium is the atmosphere of the earth, and radio frequencies carry the signals.

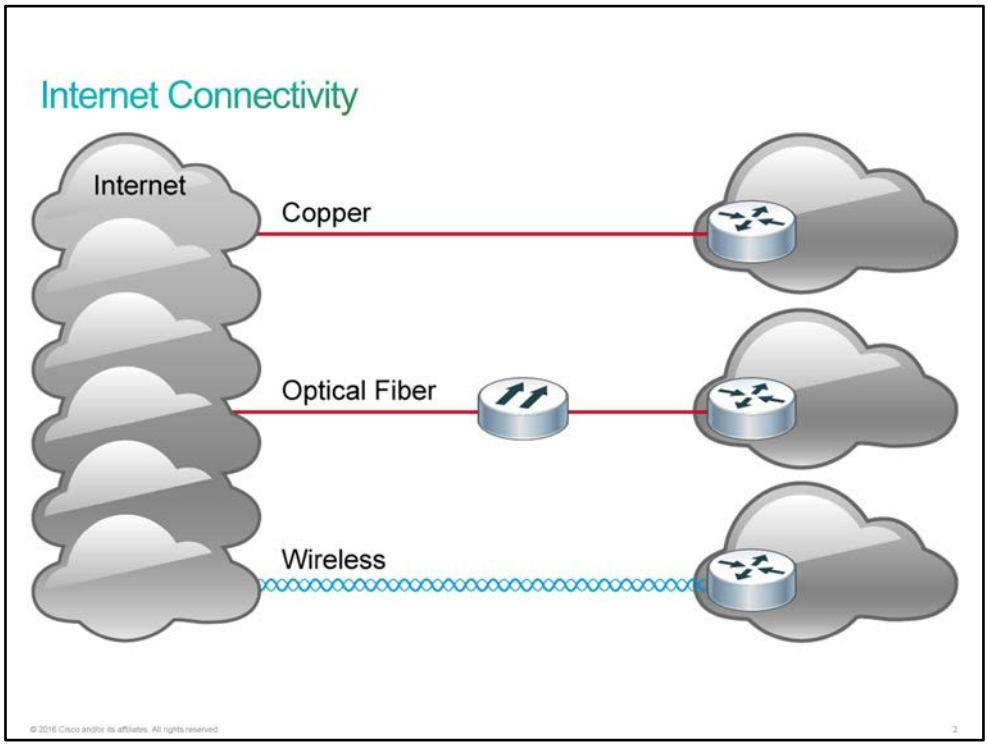

The copper medium includes cables, such as twisted-pair telephone wire, coaxial cable, or (most commonly), Category 6 UTP cable.

Copper medium is used by DSL, cable (DOCSIS), and serial connectivity methods. DSL sends signals across existing telephone lines, whereas cable Internet services leverage the CATV infrastructure.

Optical fibers are thin strands of glass or plastic that transmit digital signals by modulated pulses of light.

Due to its immunity to EMI and RFI, fiber-optic cabling is well suited for harsh environments. The fiber-optic cable medium has the added benefit of extending the distance of cable runs far beyond the capabilities of copper cable.

Wireless Internet carriers offer several different connectivity choices. One option is the home wireless connection between a wireless router and a computer with a wireless network card. Another option is the terrestrial wireless connection between two ground stations. Wireless Internet connectivity can also be achieved through the communication between ground receivers on earth and satellite communication between satellites in geostationary orbit.

The fourth generation (4G), is the fourth generation of mobile telecommunications technology, succeeding 3G and preceding 5G. Applications include amended mobile web access, IP telephony, gaming services, high-definition mobile TV, video conferencing, 3D television, and cloud computing.

WiMAX is a family of wireless communications standards that are designed to provide up to 1 Gbps data rates.

Copper media pros and cons:

- Better control
- Better security
- More reliability
- High speed
- Relatively cost-effective
- Costly and awkward to maintain

Wireless media pros and cons:

- Freedom to move around the office
- Neater, getting rid of all those unsightly cables
- Productivity benefits
- Harder to secure

Provider-Assigned IP Addresses

The DHCP service enables devices on a network to obtain IP addresses and other information from a DHCP server. This service automates assignment of IP addresses, subnet masks, gateways, and other IP networking parameters.

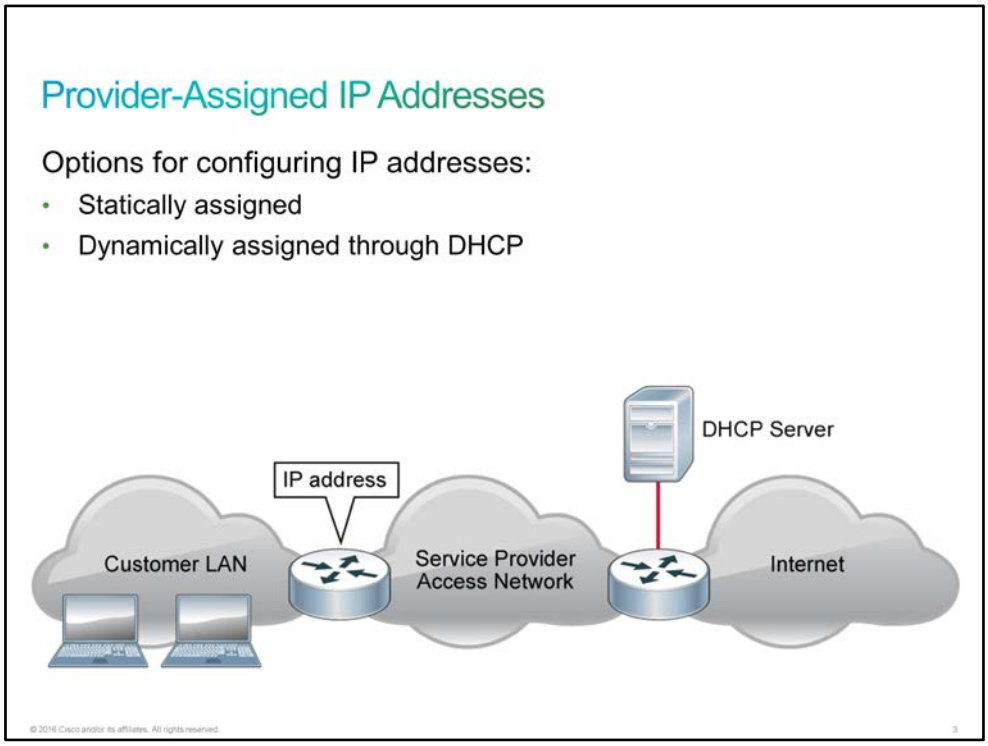

Configuring a Static Provider-Assigned IP Address

A service provider sometimes provides a static address for an interface that is connected to the Internet. In other cases, the address is provided using DHCP. On larger local networks or where the user population changes frequently, DHCP is preferred. New users may arrive with laptops and need a connection. Others have new workstations that need to be connected. Rather than have the network administrator assign an IP address for each workstation, it is more efficient to have IP addresses assigned automatically using DHCP.

If an ISP uses DHCP to provide interface addressing, no manual addresses can be configured. Instead, the interface is configured to operate as a DHCP client. This configuration means that when the router is connected to a cable modem, for example, it is a DHCP client and requests an IP address from the ISP.

Static provider-assigned IP addresses can be more useful in several respects than dynamic addresses. Static IP addresses can be linked to a domain name (such as http://www.cisco.com), and public or private servers can be run for access by outside users.

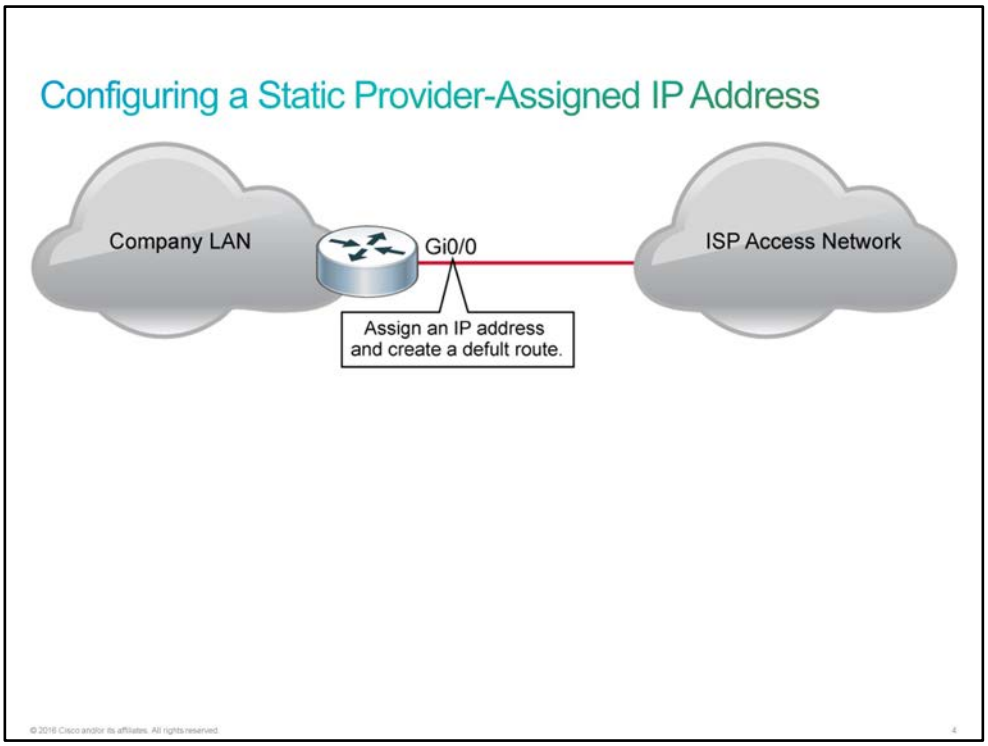

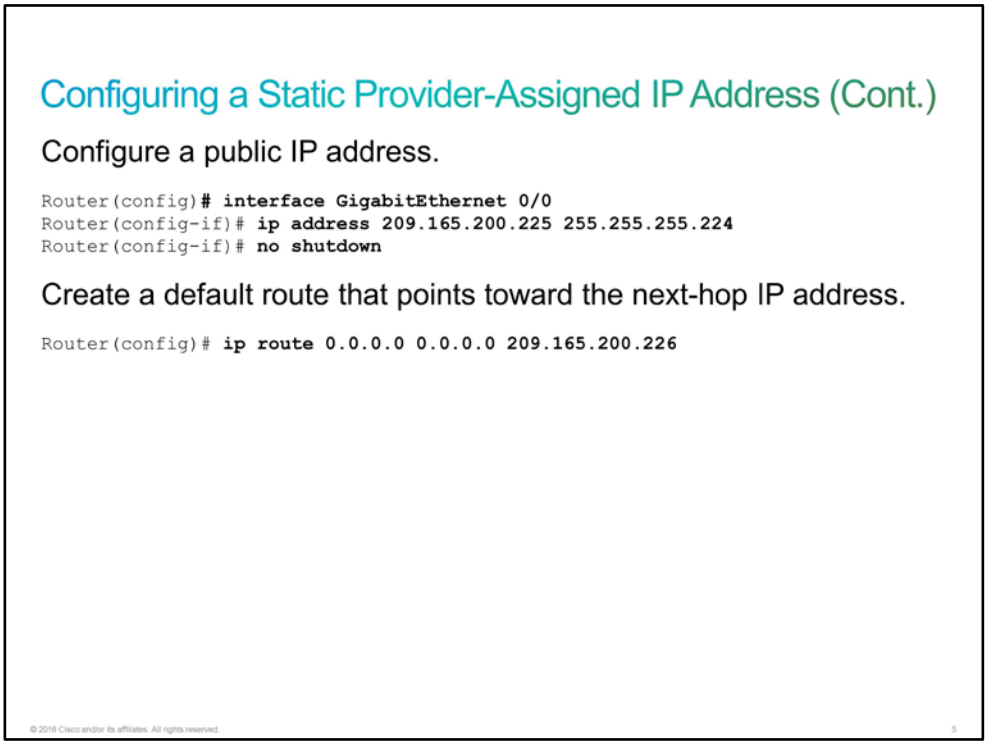

Static provider-assigned IP addresses can be more useful in several respects than dynamic addresses. Static IP addresses can be linked to a domain name (such as http://www.cisco.com), and public or private servers can be run for access by outside users.

For example, your ISP assigns you a static IP address of 209.165.200.225/27. You proceed with a two-step process. The first step is to configure the static IP address that you were assigned on the outside interface of the router. The second step is to configure a default route that forwards all traffic that is intended for the Internet to the outside interface.

Command	Description
interface *interface*	Specifies an interface and enters the interface configuration mode
ip address *address subnet_mask*	Sets the IP address and mask of the device
no shutdown	Enables an interface
ip route *net-prefix prefix-mask next_hop_ip_address*	Establishes a static route to destination

Dynamic Host Configuration Protocol

Managing a network can be very time-consuming. Network clients break, or are moved, and new clients are purchased that need network connectivity—these tasks are all part of the network administrator job. Depending on the number of IP hosts, manual configuration of IP addresses for every device on the network is virtually impossible.

DHCP can greatly decrease the workload of the network administrator. DHCP automatically assigns an IP address from an IP address pool that the administrator defines. However, DHCP is much more than just a mechanism that allocates IP addresses. This service automates the assignment of IP addresses, subnet masks, gateways, and other IP-required networking parameters.

DHCP is built on a client/server model. The DHCP server hosts allocate network addresses and deliver configuration parameters to dynamically configured hosts. The term "client" refers to a host that is requesting initialization parameters from a DHCP server. Several different devices can be DHCP clients, including Cisco IP phones, desktop PCs, laptops, printers, and even Blu-Ray players. Just about any device that you can configure to participate on a TCP/IP network has the option of using DHCP to obtain its IP configuration.

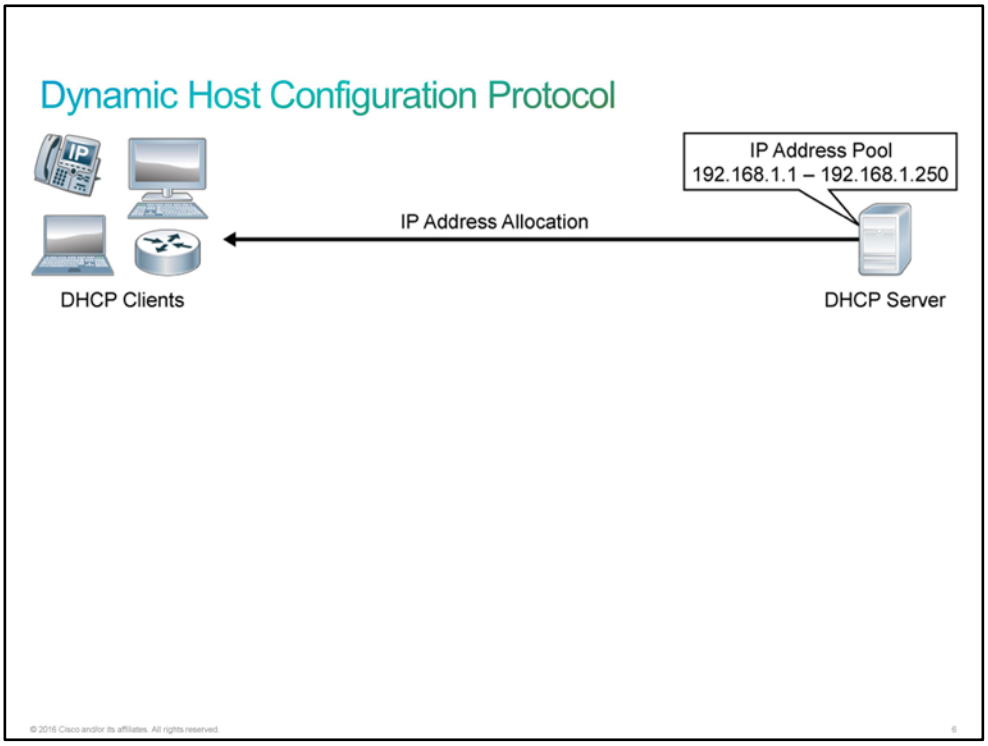

Depending on the actual DHCP server that is in use, there are three basic DHCP IP address allocation mechanisms:

- **Dynamic allocation:** Dynamic allocation of IP addresses is the most common type of address assignment. As devices boot and activate their Ethernet interfaces, the DHCP client service triggers a DHCP Discover broadcast that includes the MAC address of the DHCP client. If a DHCP server is listening on that IP subnet, it responds with a DHCP Offer message. As the name implies, the DHCP Offer message offers an unused IP address from the IP address pool that is on the DHCP server. If the IP address is acceptable, the DHCP client then sends a DHCP Request agreeing to the offered address. The DHCP server then marks the IP address as "in use" in its database and sends a final DHCP ACK to the DHCP client. The DHCP server also starts the countdown on a "lease timer." With a dynamic allocation, a DHCP client is given its IP configuration for a specified amount of time. When the lease time expires, the DHCP server can reclaim the address and return it to the address pool and lease it to another host.
- **Automatic allocation:** Automatic allocation of IP addresses is very similar to dynamic allocation, except that the lease time is set to never expire. This setting results in the DHCP client always being associated with the same IP address.
- **Static allocation:** Static allocation is an alternative that is generally used for devices such as servers and printers, where the device needs to remain at a given address more or less permanently. A static entry is made in the DHCP database that maps the MAC address to an IP address that is not part of the DHCP lease pool.

Note	The following examples show a simplified packet capture of a DHCP request.

```
Source              Destination         Protocol   Info
0.0.0.0             255.255.255.255     DHCP       DHCP Discover
78:ac:c0:52:e8:bd   ff:ff:ff:ff:ff:ff

Source              Destination          Protocol   Info
10.10.1.1           255.255.255.255      DHCP       DHCP Offer 10.10.1.241
00:1b:d5:9c:34:27   ff:ff:ff:ff:ff:ff

Source              Destination         Protocol   Info
0.0.0.0             255.255.255.255     DHCP       DHCP Request
78:ac:c0:52:e8:bd   ff:ff:ff:ff:ff:ff

Source              Destination         Protocol    Info
10.10.1.1           255.255.255.255     DHCP        DHCP ACK
00:1b:d5:9c:34:27   ff:ff:ff:ff:ff
```

Configuring a DHCP Client

An ISP sometimes provides a static address for an interface that is connected to the Internet. In other cases, an address is provided using DHCP. If the ISP uses DHCP to provide interface addressing, no manual address can be configured. Instead, the interface is configured to operate as a DHCP client.

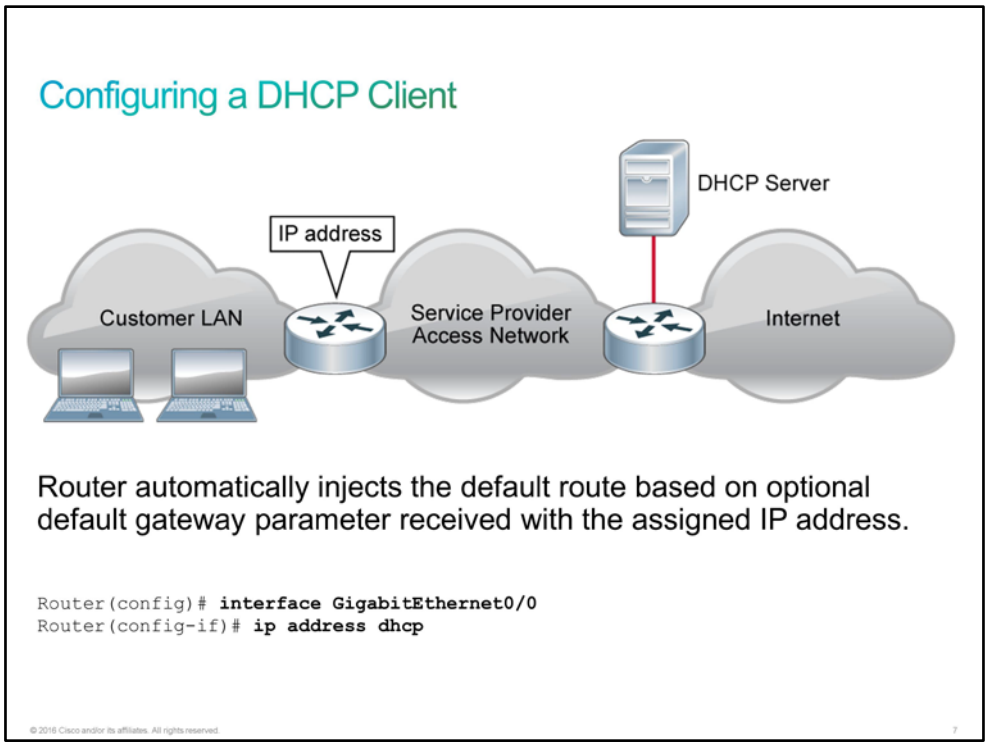

If the router receives an optional DHCP parameter that is called the default gateway with the assigned IP address, the default route will be injected into the routing table, pointing to the default gateway IP address.

Command	Description
interface *interface*	Specifies an interface and enters the interface configuration mode
ip address dhcp	Specifies that the interface acquires an IP address through DHCP

Public vs. Private IPv4 Addresses

Some networks connect to each other through the Internet, while others are private. For instance, the example addresses used in this course are private, which means that they are not assigned to public use. Public and private IP addresses are required for both of these network types.

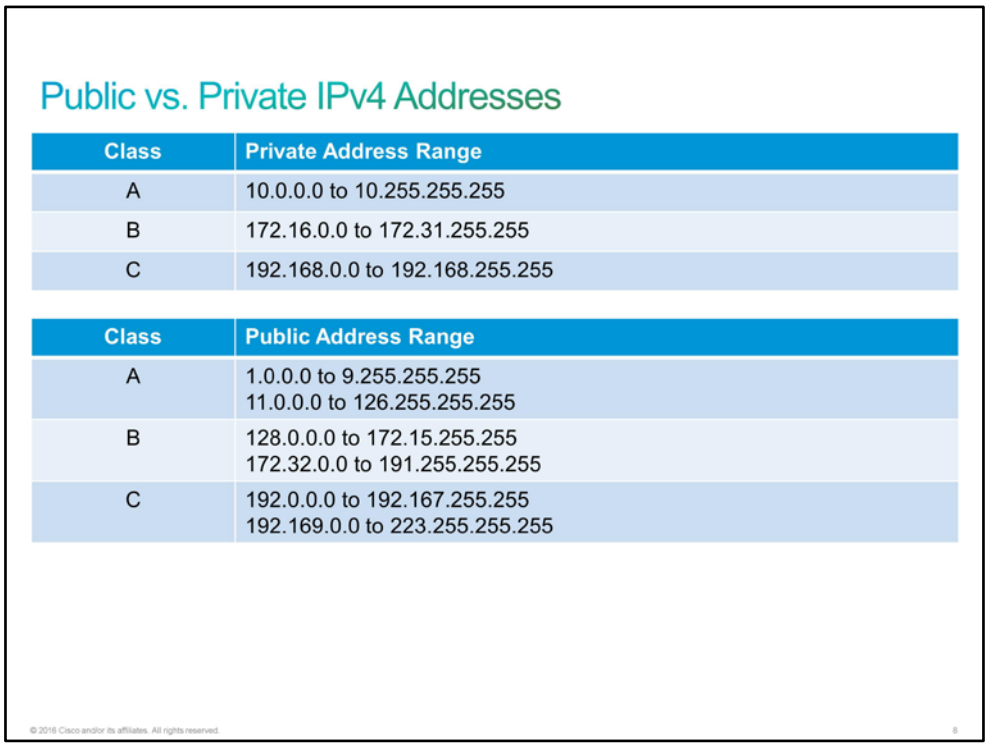

Private IP Addresses

Internet hosts require a globally unique IP address, but private hosts that are not connected to the Internet can use any valid address, as long as it is unique within the private network. However, because many private networks exist alongside public networks, grabbing "just any address" is strongly discouraged.

Three blocks of IP addresses (one Class A network, 16 Class B networks, and 256 Class C networks) are designated for private, internal use. The table shows the address ranges for each class. Addresses in these ranges are not routed on the Internet backbone. Internet routers are configured to discard private addresses.

In a private intranet, these private addresses can be used instead of globally unique addresses.

When a network that is using private addresses must connect to the Internet, it is necessary to translate the private addresses to public addresses. This translation process is called NAT. A router is often the network device that performs NAT.

Public IP Addresses

Public IP addresses are used for the hosts that are publicly accessible from the Internet. Internet stability depends directly on the uniqueness of publicly used network addresses. Therefore, a mechanism is needed to ensure that addresses are, in fact, unique. This mechanism was originally managed by the InterNIC. IANA succeeded InterNIC. IANA carefully manages the remaining supply of IP addresses to ensure that duplication of publicly used addresses does not occur. Duplication would cause instability in the Internet and would compromise its ability to deliver datagrams to networks using the duplicated addresses.

To obtain a provider-dependent IP address or block of addresses, you must contact an ISP. To obtain provider-independent IP addresses, you must contact an LIR. LIRs obtain IP address pools from their RIRs:

- AfriNIC
- APNIC
- ARIN
- LACNIC
- RIPE NCC

With the rapid growth of the Internet, public IP addresses began to run out. New mechanisms such as NAT, CIDR, VLSM, and IPv6 were developed to help solve the problem.

Discovery 12: Configure a Provider-Assigned IP Address

Introduction

This discovery lab will guide you through the aspects of connecting a small network to the Internet. You will implement the simplest of Internet connections where R1 will receive its IP address via DHCP.

Topology

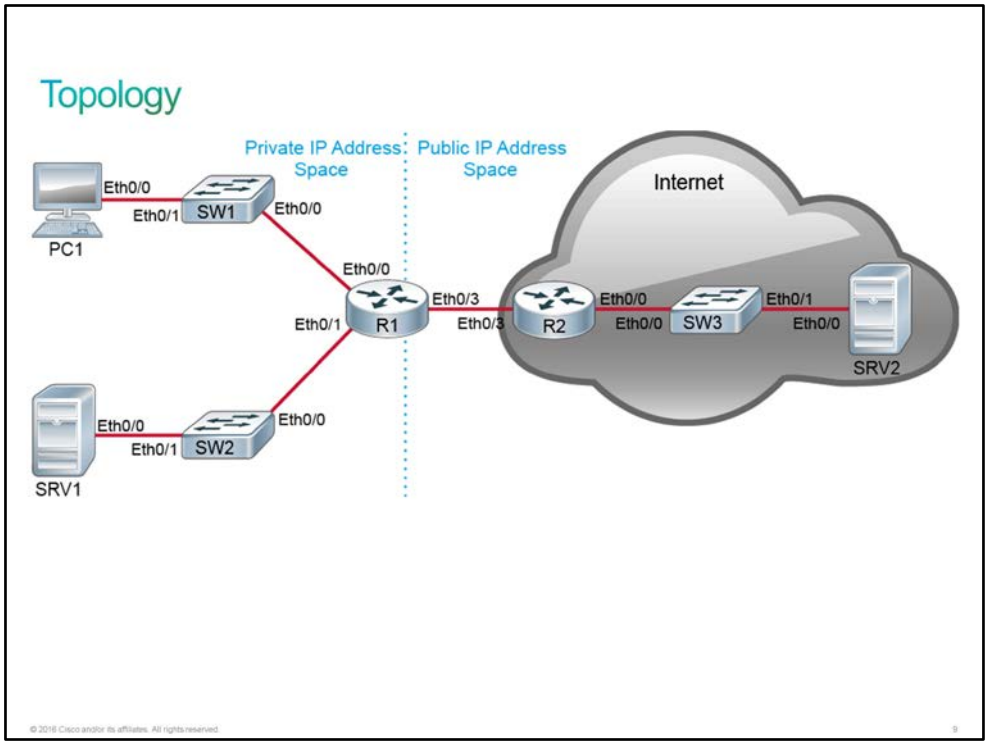

Job Aid

Device Information

Device Information Table

Device	Characteristic	Value
PC1	Hostname	PC1
PC1	IP address	10.10.1.10/24
PC1	Default gateway	10.10.1.1
SRV1	Hostname	SRV1

Device	Characteristic	Value
SRV1	IP address	10.10.2.20/24
SRV1	Default gateway	10.10.2.1
SRV2	Hostname	SRV2
SRV2	IP address	203.0.113.30/24
SRV2	Default gateway	203.0.113.1
SW1	Hostname	SW1
SW1	VLAN 1 IP address	10.10.1.4/24
SW1	Default gateway	10.10.1.1
SW1	Ethernet0/0 description	Link to R1
SW1	Ethernet0/1 description	Link to PC1
SW2	Hostname	SW2
SW2	VLAN 1 IP address	10.10.2.4/24
SW2	Default gateway	10.10.2.1
SW2	Ethernet0/0 description	Link to R2
SW2	Ethernet0/1 description	Link to PC2
SW3	Hostname	SW3
SW3	VLAN 1 IP address	203.0.113.4/24
SW3	Default gateway	203.0.113.1
SW3	Ethernet0/0 description	Link to R3
SW3	Ethernet0/1 description	Link to SRV2
R1	Hostname	R1
R1	Ethernet0/0 description	Link to SW1
R1	Ethernet0/0 IP address	10.10.1.1/24
R1	Ethernet0/1 description	Link to SW2
R1	Ethernet0/1 IP address	10.10.2.1/24

Device	Characteristic	Value
R1	Ethernet0/3 description	Link to R2
R1	Ethernet0/3 IP address	198.51.100.2/24
R2	Hostname	R2
R2	Ethernet0/0 description	Link to SW3
R2	Ethernet0/0 IP address	203.0.113.1/24
R2	Ethernet0/3 description	Link to R1
R2	Ethernet0/3 IP address	198.51.100.1/24

PC and SRVs in the virtual lab environment are simulated as routers, so you should use Cisco IOS commands to configure them or make verifications.

Global IP Address Networks

Address Block	Host Starting Address	Host Ending Address	Broadcast Address	Subnet Mask
192.0.2.0/24	192.0.2.1	192.0.2.254	192.0.2.255	255.255.255.0
198.51.100.0/24	198.51.100.1	198.51.100.254	198.51.100.255	255.255.255.0
203.0.113.0/24	203.0.113.1	203.0.113.254	203.0.113.255	255.255.255.0

Task 1: Configure a Provider-Assigned IP Address

Activity

Step 1 To provide some insight into the functioning of a DHCP server, you will configure DHCP server services on R2, which is acting as the ISP router. On R2, access the global configuration with the **configure terminal** command.

On R2, enter the following command:

```
R2# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
R2(config)#
```

The most commonly used commands are abbreviated in this guided discovery. For example, **en** is used for **enable** and **conf t** is used for **configure terminal**. If there is any confusion, you can perform tab completion of the commands to see the full commands during the discovery execution. For example, **en<tab>** would expand to **enable** and **conf<tab> t<tab>** would expand to **configure terminal**.

Step 2 Define a DHCP address pool for the subnet 198.51.100.0/24, specifying the interface (198.51.100.1) of R2 as the default gateway and an address lease length of 7 days. DHCP pools are identified with a name.

On R2, enter the following commands:

```
R2(config)# ip dhcp pool Clients
R2(dhcp-config)# network 198.51.100.0 /24
R2(dhcp-config)# default-router 198.51.100.1
R2(dhcp-config)# lease 7
R2(dhcp-config)# exit
R2(config)#
```

Step 3 By default, Cisco IOS DHCP server will serve all IP addresses within the subnet of a defined pool. Limit the assignable addresses from 198.51.100.101 to 198.51.100.254 by excluding the first 100 addresses in the subnet range. Then, leave the configuration mode on R2.

On R2, enter the following commands:

```
R2(config)# ip dhcp excluded-address 198.51.100.1 198.51.100.100
R2(config)# end
R2#
```

Step 4 R1 is currently configured with a static IP address on Ethernet0/3. Verify this fact.

Access the console of R1 and examine the current IP configuration on Ethernet0/3.

```
R1# sh ip int brie
Interface                  IP-Address      OK? Method Status
Protocol
Ethernet0/0                10.10.1.1       YES NVRAM  up                    up
Ethernet0/1                10.10.2.1       YES NVRAM  up                    up
Ethernet0/2                unassigned      YES NVRAM  administratively down
down
Ethernet0/3                198.51.100.2    YES NVRAM  up                    up
Serial1/0                  unassigned      YES NVRAM  administratively down
down
Serial1/1                  unassigned      YES NVRAM  administratively down
down
Serial1/2                  unassigned      YES NVRAM  administratively down
down
Serial1/3                  unassigned      YES NVRAM  administratively down
down
```

Step 5 Routing has not been configured on R1. Verify that only local and connected routes appear in the R1 routing table and that there is no default route configured.

On R1, enter the following command:


```
R1# sh 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
       + - replicated route, % - next hop override

Gateway of last resort is not set

      10.0.0.0/8 is variably subnetted, 4 subnets, 2 masks
C        10.10.1.0/24 is directly connected, Ethernet0/0
L        10.10.1.1/32 is directly connected, Ethernet0/0
C        10.10.2.0/24 is directly connected, Ethernet0/1
L        10.10.2.1/32 is directly connected, Ethernet0/1
      198.51.100.0/24 is variably subnetted, 2 subnets, 2 masks
C        198.51.100.0/24 is directly connected, Ethernet0/3
L        198.51.100.2/32 is directly connected, Ethernet0/3
```

Step 6 Reconfigure the interface Ethernet0/3 to obtain its IP address and default gateway via DHCP.

On R1, enter the following commands:

```
R1# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
R1(config)# int eth0/3
R1(config-if)# ip address dhcp
R1(config-if)# end
R1#
```

Step 7 Wait up to 30 seconds. Verify that R1 displays a syslog message indicating that it has been assigned an IP address via DHCP.

```
R1#
*Oct 20 14:47:21.312: %DHCP-6-ADDRESS_ASSIGN: Interface Ethernet0/3 assigned
DHCP address 198.51.100.101, mask 255.255.255.0, hostname R1
```

Step 8 Verify that the IP address is assigned to the interface Ethernet0/3.

On R1, enter the following command:

```
R1# show ip int brief
Interface                  IP-Address      OK? Method Status
Protocol
Ethernet0/0                10.10.1.1       YES NVRAM  up                    up
Ethernet0/1                10.10.2.1       YES NVRAM  up                    up
Ethernet0/2                unassigned      YES NVRAM  administratively down
down
Ethernet0/3                198.51.100.101  YES DHCP   up                    up
Serial1/0                  unassigned      YES NVRAM  administratively down
down
Serial1/1                  unassigned      YES NVRAM  administratively down
down
Serial1/2                  unassigned      YES NVRAM  administratively down
down
Serial1/3                  unassigned      YES NVRAM  administratively down
down
```

Step 9 Verify that there is now a default route on R1, using R2 (198.51.100.1) as the default route.

On R1, enter the following command:

```
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
       + - replicated route, % - next hop override

Gateway of last resort is 198.51.100.1 to network 0.0.0.0

S*    0.0.0.0/0 [254/0] via 198.51.100.1
      10.0.0.0/8 is variably subnetted, 4 subnets, 2 masks
C        10.10.1.0/24 is directly connected, Ethernet0/0
L        10.10.1.1/32 is directly connected, Ethernet0/0
C        10.10.2.0/24 is directly connected, Ethernet0/1
L        10.10.2.1/32 is directly connected, Ethernet0/1
      198.51.100.0/24 is variably subnetted, 2 subnets, 2 masks
C        198.51.100.0/24 is directly connected, Ethernet0/3
L        198.51.100.101/32 is directly connected, Ethernet0/3
```

Step 10 Verify connectivity from R1 to the public IP address side of the topology by pinging SRV2 (203.0.113.30).

On R1, enter the following command:

```
R1# ping 203.0.113.30
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 203.0.113.30, timeout is 2 seconds:
.!!!!
Success rate is 80 percent (4/5), round-trip min/avg/max = 1/1/1 ms
```

This is the end of the discovery lab.

Introducing NAT

Small networks are commonly implemented using private IP addressing as defined in RFC 1918. Private addressing gives enterprises considerable flexibility in a network design. This addressing enables operationally and administratively convenient addressing schemes and easier growth. However, you cannot route private addresses over the Internet, and there are not enough public addresses to allow all organizations to provide a private address to all their hosts. Therefore, network administrators need a mechanism to translate private addresses to public addresses (and back) at the edge of their network.

NAT provides this mechanism. Before NAT, a host with a private address could not access the Internet. Using NAT, companies can provide some or all their hosts with private addresses and provide address translation to allow access to the Internet.

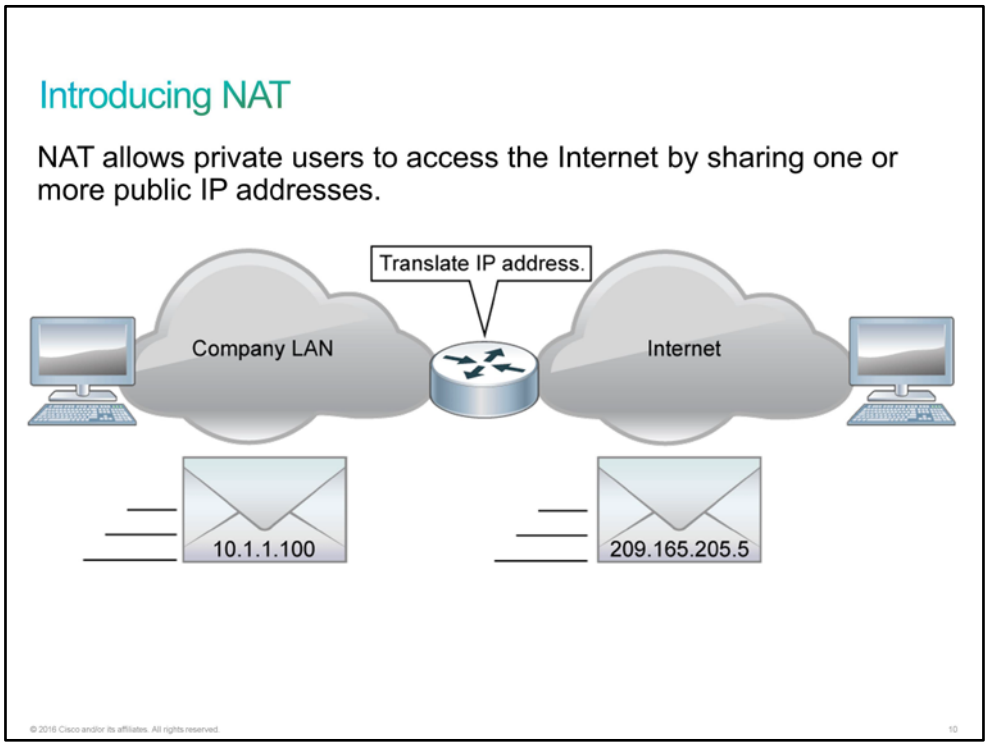

NAT is like the receptionist in a large office. Assume that you have left instructions with the receptionist not to forward any calls to you unless you request it. Later, you call a potential client and leave a message asking the client to call you back. You tell the receptionist that you are expecting a call from this client, and you ask the receptionist to put the call through to you. The client calls the main number to your office, which is the only number that the client knows. When the caller gives the receptionist your name, the receptionist checks a lookup table that matches your name to your extension. The receptionist knows that you requested this call and forwards the caller to your extension.

Usually, NAT connects two networks and translates the private (inside local) addresses in the internal network to public (inside global) addresses before packets are forwarded to another network. You can configure NAT to advertise only one address for the entire network to the outside world. Advertising only one address effectively hides the internal network, providing additional security as a side benefit.

The network address translation process of swapping one address for another is separate from the convention that is used to determine what is public and private, and devices must be configured to recognize which IP networks should be translated. This requirement is one of the reasons why NAT can also be deployed internally when there is a clash of private IP addresses, such as, for example, when two companies merge.

The benefits of NAT are the following:

- Eliminates the need to readdress all hosts that require external access, saving time and money.
- Conserves addresses through application port-level multiplexing. With PAT, multiple internal hosts can share a single registered IPv4 address for all external communication. In this type of configuration, relatively few external addresses are required to support many internal hosts. This characteristic conserves IPv4 addresses.
- Protects network security. Because private networks do not advertise their addresses or internal topology, they remain reasonably secure when they gain controlled external access with NAT.

The disadvantages of NAT are the following:

- Many IP addresses and applications depend on end-to-end functionality, with unmodified packets forwarded from the source to the destination. By changing end-to-end addresses, NAT blocks some applications that use IP addressing. For example, some security applications, such as digital signatures, fail because the source IP address changes. Applications that use physical addresses instead of a qualified domain name do not reach destinations that are translated across the NAT router. Sometimes, you can avoid this problem by implementing static NAT mappings.
- End-to-end IP traceability is also lost. It becomes much more difficult to trace packets that undergo numerous packet address changes over multiple NAT hops, so troubleshooting is challenging. On the other hand, hackers who want to determine the source of a packet find it difficult to trace or obtain the original source or destination address.
- Using NAT also complicates tunneling protocols, such as IPsec, because NAT modifies the values in the headers. This behavior interferes with the integrity checks that IPsec and other tunneling protocols perform.
- Services that require the initiation of TCP connections from the outside network, or stateless protocols such as those using UDP, can be disrupted. Unless the NAT router makes a specific effort to support such protocols, incoming packets cannot reach their destination. Some protocols can accommodate one instance of NAT between participating hosts (passive mode FTP, for example) but fail when NAT separates both systems from the Internet.
- The last disadvantage involves performance. NAT increases switching delays because translation of each IP address within the packet headers takes time. The first packet is process-switched, meaning that it always goes through the slower path. The router must look at each packet to decide whether it needs translation. The router needs to alter the IP header and possibly alter the TCP or UDP header. Remaining packets go through the fast-switched path if a cache entry exists; otherwise, they too are delayed.

Types of Addresses in NAT

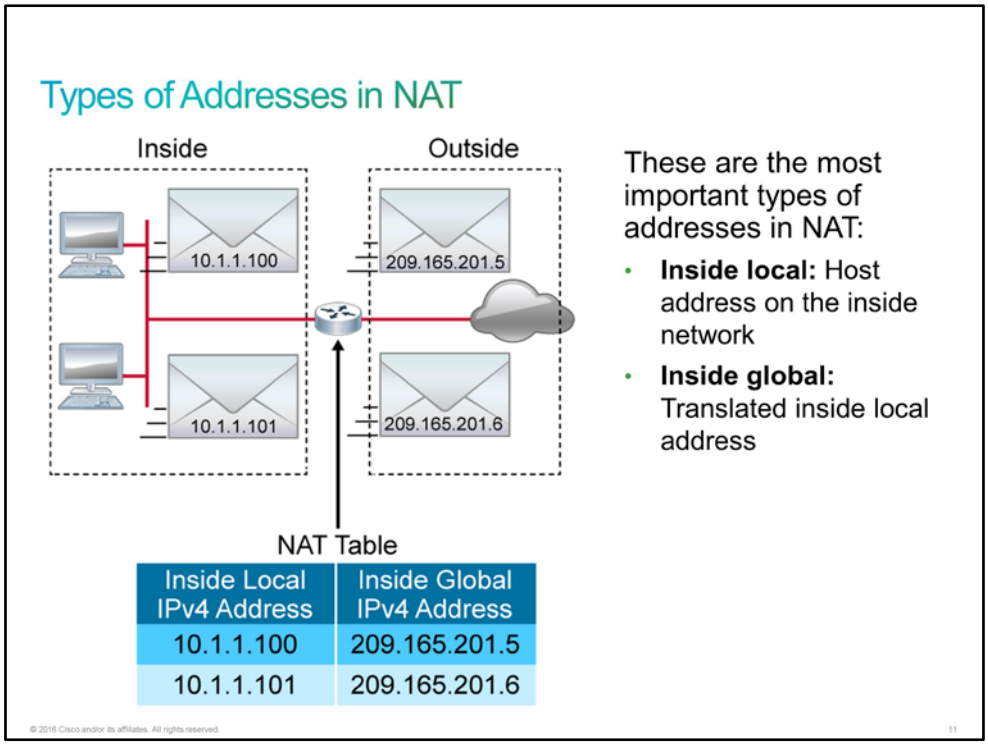

In NAT terminology, the *inside network* is the set of networks that are subject to translation. The *outside network* refers to all other addresses. Usually, these other addresses are valid addresses that are located on the Internet.

Cisco defines these NAT terms:

- **Inside local address:** The IPv4 address that is assigned to a host on the inside network. The inside local address is likely not an IPv4 address that the network information center or service provider assigns.
- **Inside global address:** The translated inside local address. It is typically a public IPv4 address.
- **Outside global address:** The IPv4 address that the host owner assigns to a host on the outside network. The outside global address is allocated from a globally routable address or network space.
- **Outside local address:** The IPv4 address of an outside host as it appears to the inside network. Not necessarily public, the outside local address is allocated from a routable address space on the inside.

A good way to remember what is local and what is global is to add the word *visible*. An address that is locally visible normally implies a private IP address, and an address that is globally visible normally implies a public IP address. The rest is simple. *Inside* means internal to your network and *outside* means external to your network. So, for example, an inside global address means that the device is physically inside your network and has an address that is visible from the Internet. It could be a web server, for instance.

Types of NAT

On a Cisco IOS router, NAT can be divided into three distinct categories, each having a clear use case.

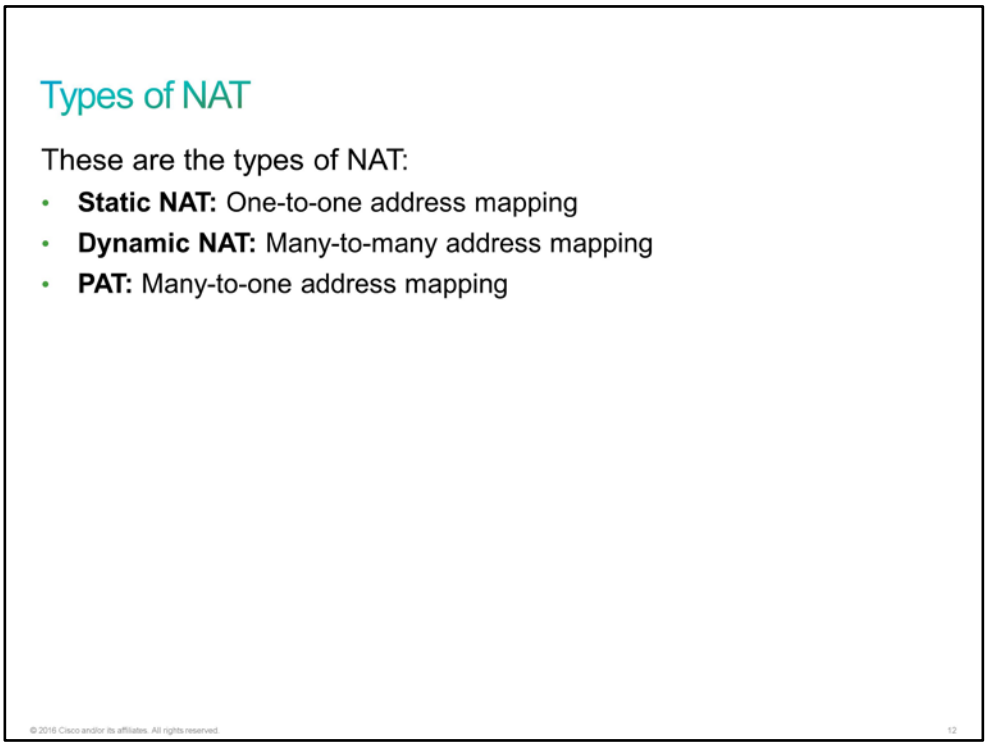

NAT can work in these ways:

- **Static NAT:** Maps a private IPv4 address to a public IPv4 address (one to one). Static NAT is particularly useful when a device must be accessible from outside the network. This type of NAT is used when a company has a server for which it needs a static IP address.
- **Dynamic NAT:** Maps a private IPv4 address to a public IPv4 address from a group of public IPv4 addresses. This type of NAT is used, for example, when two companies that are using the same private address space merge. With the use of dynamic NAT readdressing, using the entire address space is avoided or at least postponed.
- **PAT:** PAT maps multiple private IPv4 addresses to a single public IPv4 address (many to one) by using different ports. PAT is also known as NAT overloading. It is a form of dynamic NAT and is the most common use of NAT. It is used every day in your place of business or your home. Multiple users of PCs, tablets, and phones are able to access the Internet, even though only one public IP address is available for that LAN.

Understanding Static NAT

You can translate your own IPv4 addresses into globally unique IPv4 addresses when you are communicating outside your network.

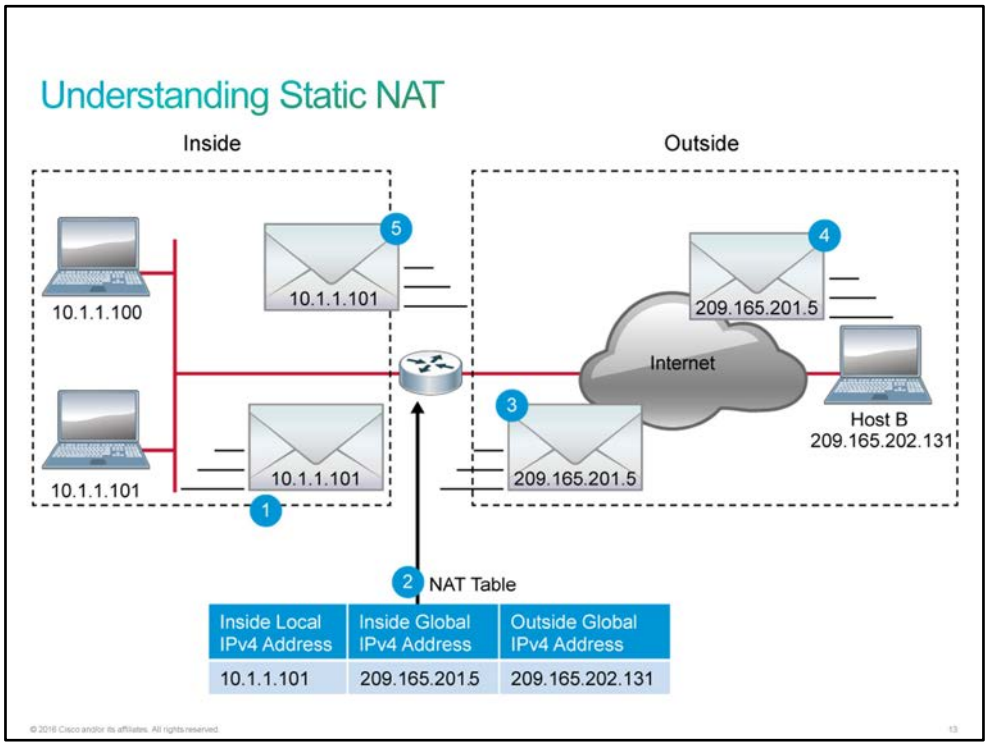

The figure illustrates a router that is translating a source address inside a network into a source address outside the network. The following are the steps for translating an inside source address:

1. The user at host 10.1.1.101 wants to open a connection to Host B (IP address 209.165.202.131).
2. The first packet that the router receives on its NAT inside-enabled interface from host 10.1.1.101 causes the router to check its NAT table.
3. The router replaces the inside local source address of host 10.1.1.101 with the translated inside global address (209.165.201.5) and forwards the packet.
4. Host B receives the packet and responds to host 10.1.1.101, using the inside global IPv4 destination address 209.165.201.5.
5. When the router receives the packet on its NAT outside-enabled interface with the inside global IPv4 address of 209.165.201.5, the router performs a NAT table lookup using the inside global address as a key. The router then translates the address back to the inside local address of host 10.1.1.101 and forwards the packet to host 10.1.1.101.
6. Host 10.1.1.101 receives the packet and continues the conversation. The router performs Steps 2 through 5 for each packet.

Configuring and Verifying Static NAT

Configuring Static NAT

Remember that static NAT is a one-to-one mapping between an inside address and an outside address. Static NAT allows external devices to initiate connections to internal devices. For instance, you may want to map an inside global address to a specific inside local address that is assigned to your web server. In the following example, you can see how to configure a static NAT.

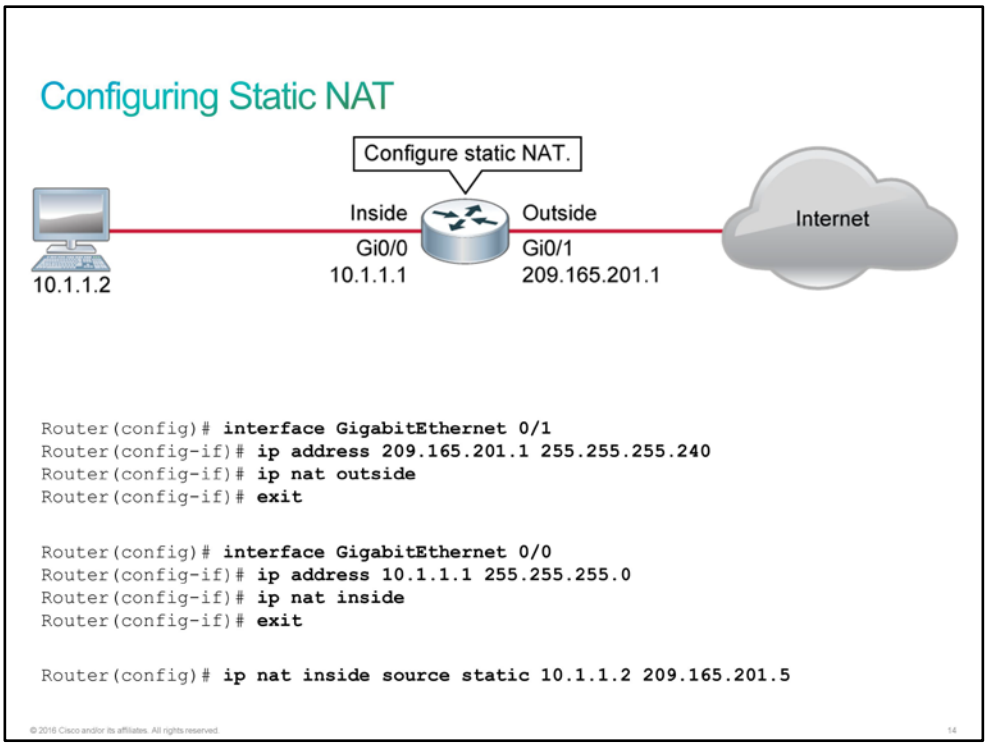

Configuring static NAT translations is a simple task. You need to define the addresses to translate and then configure NAT on the appropriate interfaces. Packets arriving on an inside interface from the identified IP address are subject to translation. Packets arriving on an outside interface that are addressed to the identified IP address are also subject to translation.

The figure shows examples of commands for the steps. You enter static translations directly into the configuration. Unlike dynamic translations, these translations are always in the NAT table.

Command	Description
interface *interface*	Specifies an interface and enters the interface configuration mode
ip address *address subnet_mask*	Sets the IP address and mask of the device
ip nat inside	Marks the interface as connected to the inside network

Command	Description
ip nat outside	Marks the interface as connected to the outside network
ip nat inside source static *inside_address outside_address*	Establishes a static translation between an inside local and inside global address

Verifying Static NAT Configuration

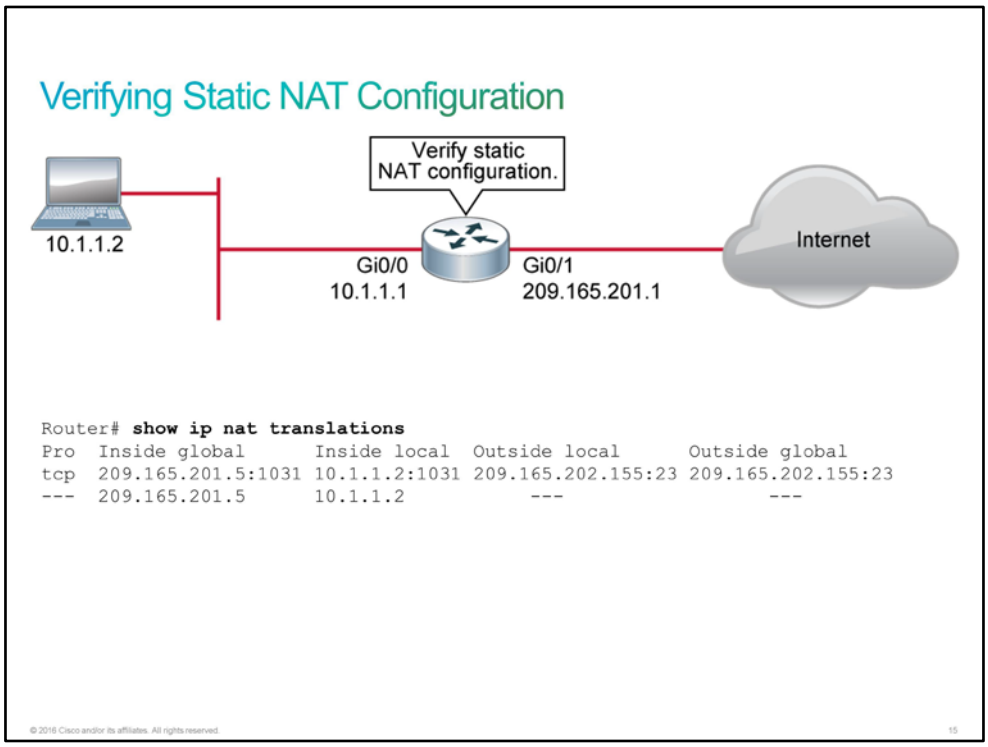

Command	Description
show ip nat translations	Displays active NAT translations

For more details about the **ip nat inside**, **ip nat pool**, **show ip nat translations**, and related commands, check the *Cisco IOS IP Addressing Services Command Reference* at http://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipaddr/command/ipaddr-cr-book.html.

Discovery 13: Configure Static NAT

Introduction

This discovery lab will guide you through the aspects of connecting a small network to the Internet. NAT is a very important concept for Internet connectivity. The private IP addresses that are used on most internal networks are not routable on the public Internet. Because they are not routable, the private IP addresses must be translated to assigned public IP addresses at the border to the Internet.

The lab is prepared with the devices that are represented in the topology diagram. All the devices have their basic configurations in place, including hostnames and IP addresses. Router R1 receives the default route from R2 via DHCP, but NAT has not been implemented. Implementing NAT will be your job during this discovery lab. You will implement a static NAT translation for SRV1. Static NAT, which can maintain persistent IP addresses for servers, facilitates inbound connectivity.

Topology

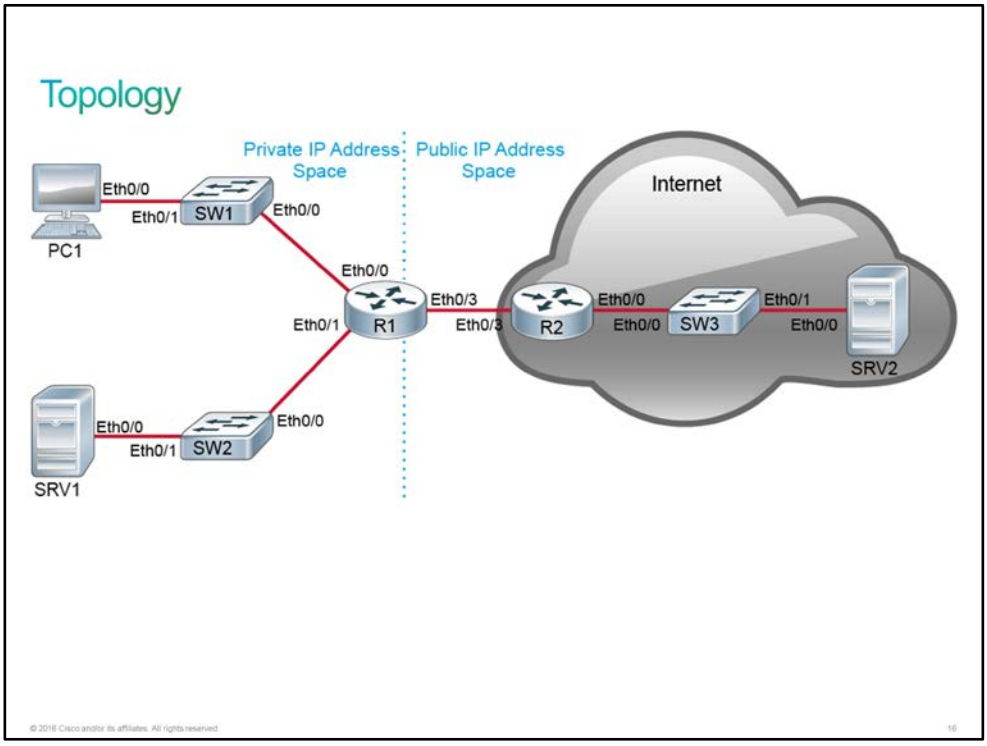

Job Aid

Device Information

Device Information Table

Device	Characteristic	Value
PC1	Hostname	PC1

Device	Characteristic	Value
PC1	IP address	10.10.1.10/24
PC1	Default gateway	10.10.1.1
SRV1	Hostname	SRV1
SRV1	IP address	10.10.2.20/24
SRV1	Default gateway	10.10.2.1
SRV2	Hostname	SRV2
SRV2	IP address	203.0.113.30/24
SRV2	Default gateway	203.0.113.1
SW1	Hostname	SW1
SW1	VLAN 1 IP address	10.10.1.4/24
SW1	Default gateway	10.10.1.1
SW1	Ethernet0/0 description	Link to R1
SW1	Ethernet0/1 description	Link to PC1
SW2	Hostname	SW2
SW2	VLAN 1 IP address	10.10.2.4/24
SW2	Default gateway	10.10.2.1
SW2	Ethernet0/0 description	Link to R2
SW2	Ethernet0/1 description	Link to PC2
SW3	Hostname	SW3
SW3	VLAN 1 IP address	203.0.113.4/24
SW3	Default gateway	203.0.113.1
SW3	Ethernet0/0 description	Link to R3
SW3	Ethernet0/1 description	Link to SRV2
R1	Hostname	R1
R1	Ethernet0/0 description	Link to SW1

Device	Characteristic	Value
R1	Ethernet0/0 IP address	10.10.1.1/24
R1	Ethernet0/1 description	Link to SW2
R1	Ethernet0/1 IP address	10.10.2.1/24
R1	Ethernet0/3 description	Link to R2
R1	Ethernet0/3 IP address	198.51.100.2/24
R2	Hostname	R2
R2	Ethernet0/0 description	Link to SW3
R2	Ethernet0/0 IP address	203.0.113.1/24
R2	Ethernet0/3 description	Link to R1
R2	Ethernet0/3 IP address	198.51.100.1/24

PC and SRVs in the virtual lab environment are simulated as routers, so you should use Cisco IOS commands to configure them or make verifications.

Global IP Address Networks

Address Block	Host Starting Address	Host Ending Address	Broadcast Address	Subnet Mask
192.0.2.0/24	192.0.2.1	192.0.2.254	192.0.2.255	255.255.255.0
198.51.100.0/24	198.51.100.1	198.51.100.254	198.51.100.255	255.255.255.0
203.0.113.0/24	203.0.113.1	203.0.113.254	203.0.113.255	255.255.255.0

Task 1: Configure Static NAT

Activity

Step 1 While R1 does have access to the public IP address space, systems within the private IP address space of the topology do not. Verify this fact. Access the console of PC1 and attempt to ping SRV2. This process should fail.

On PC1, enter the following command:


```
PC1# ping 203.0.113.30
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 203.0.113.30, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
```

For a ping operation to be successful, bidirectional connectivity must exist. In this case, the problem is not getting the echo requests from PC1 to SRV2. Instead, it is a failure in getting the echo replies from SRV2 back to PC1. Because NAT has not been configured, SRV2 is receiving IP packets from the private IP address of PC1 (10.10.1.10). Routers on the Internet are not aware of the private IP address space within the networks that connect to the Internet. When SRV2 generates an echo reply to 10.10.1.10 and sends that reply to R2 for forwarding, R2 does not have a route to use to forward the reply, so the reply is dropped.

Step 2 Configure R1 interfaces for NAT. Ethernet0/0 and Ethernet0/1 are on the inside, and Ethernet0/3 is on the outside.

On R1, enter the following commands:

```
R1# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
R1(config)# int e0/0
R1(config-if)# ip nat inside
*Dec  3 20:19:13.670: %LINEPROTO-5-UPDOWN: Line protocol on Interface NVI0,
changed state to up
R1(config-if)# int e0/1
R1(config-if)# ip nat inside
R1(config-if)# int e0/3
R1(config-if)# ip nat outside
```

There will be a significant pause in response to the first interface NAT command because R1 will have to initiate an internal NVI to support NAT.

Step 3 On R1 router, configure access list number 10, to identify addresses within 10.10.0.0/16 as NAT-eligible.

On R1, enter the following command:

```
R1(config)# access-list 10 permit 10.10.0.0 0.0.255.255
```

Step 4 To the R1 configuration, add a NAT statement that enables PAT. It translates the addresses that are permitted by access list 10 using the IP address that is assigned to the interface Ethernet0/3. Leave the configuration mode when you are done.

On R1, enter the following commands:

```
R1(config)# ip nat inside source list 10 interface e0/3 overload
R1(config)# end
R1#
```

Step 5 Return to the console of PC1 and reattempt the ping operation to SRV2 (203.0.113.30). This time, it should succeed.

On PC1, enter the following command:

```
PC1# ping 203.0.113.30
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 203.0.113.30, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/3 ms
```

Step 6 In this step and the next few steps, you will verify the status of the translation that is in place. Establish a Telnet session from PC1 to SRV2. Enter the username "admin" with the password "Cisco123."

On PC1, enter the following command:

```
PC1# telnet 203.0.113.30
Trying 203.0.113.30 ... Open


User Access Verification

Username: admin
Password: Cisco123
SRV2>
```

Step 7 You are now connected to a vty line of SRV2 from PC1. Using this interface, view the status of the IP sockets on SRV2, noting the foreign IP address that SRV2 sees.

On SRV2, enter the following command:

```
SRV2> show control-plane host open-ports
Active internet connections (servers and established)
Prot               Local Address             Foreign Address
Service    State
 tcp                        *:23                          *:0
Telnet   LISTEN
 tcp                        *:23        198.51.100.101:12959
Telnet ESTABLIS
```

SRV2 sees 198.51.100.101 as the source IP address of the connection that is coming in from PC1. This address is the IP address on Ethernet0/3 that R1 obtained via DHCP. PAT is in effect.

Step 8 Leave the connection to SRV2 from PC1 running. Access the console of SRV1 and establish a second connection to SRV2 from the private IP address space.

On SRV1, enter the following commands:

```
SRV1# telnet 203.0.113.30
Trying 203.0.113.30 ... Open

User Access Verification

Username: admin
Password: Cisco123
SRV2>
```

Step 9 Using the connection to SRV2 from SRV1, again review the IP socket status on SRV2.

On SRV2, enter the following command:


```
SRV2> show control-plane host open-ports
Active internet connections (servers and established)
Prot               Local Address             Foreign Address
Service    State
 tcp                        *:23                         *:0
Telnet   LISTEN
 tcp                        *:23        198.51.100.101:21299
Telnet ESTABLIS
 tcp                        *:23        198.51.100.101:34023
Telnet ESTABLIS
```

There are now two established Telnet sessions to SRV2. One is from PC1 and the other is from SRV1. But, from the perspective of SRV2, both connections are coming from 198.51.100.101. The two connections are uniquely identified by their source ports.

Step 10 Leaving both connections to SRV2 running, access the console of R1. Display the translation table on R1.

On R1, enter the following command:

```
R1# show ip nat translations
Pro Inside global      Inside local       Outside local      Outside global
tcp 198.51.100.101:21299 10.10.1.10:21299 203.0.113.30:23    203.0.113.30:23
tcp 198.51.100.101:34023 10.10.2.20:34023 203.0.113.30:23    203.0.113.30:23
```

R1 is also using the inside source port to uniquely identify the two translation sessions.

The source ports are dynamically generated so that the ports that are shown in the example will not match those ports that you see in the lab environment. But the source ports in the R1 translation table should match those in the open-ports status for SRV2.

Using the depicted example, when R1 receives a packet from 203.0.113.30 with a source port of 23 that is destined for 198.51.100.101 and a destination port of 21299, R1 knows to translate the destination address to 10.10.1.10 and forward the packet to PC1. On the other hand, if the destination port of a similar inbound packet is 34023, R1 will translate the destination address to 10.10.2.20 and forward the packet to SRV2.

Step 11 View the running translation statistics on R1.

On R1, enter the following command:

```
R1# show ip nat statistics
Total active translations: 2 (0 static, 2 dynamic; 2 extended)
Peak translations: 2, occurred 00:12:47 ago
Outside interfaces:
  Ethernet0/3
Inside interfaces:
  Ethernet0/0, Ethernet0/1
Hits: 389  Misses: 0
CEF Translated packets: 389, CEF Punted packets: 0
Expired translations: 4
Dynamic mappings:
-- Inside Source
[Id: 1] access-list 10 interface Ethernet0/3 refcount 2

Total doors: 0
Appl doors: 0
Normal doors: 0
Queued Packets: 0
```

The **show ip nat statistics** command displays information on the current configuration of NAT (interface assignment, ACL specification, and so on), active translation statistics, and historic translation statistics.

Step 12 One at a time, access the consoles of PC1 and SRV1. Terminate their Telnet sessions to SRV2.

Terminate PC1 Telnet session to SRV2:

```
SRV2> exit

[Connection to 203.0.113.30 closed by foreign host]
PC1#
```

Terminate SRV1 Telnet session to SRV2:

```
SRV2> exit

[Connection to 203.0.113.30 closed by foreign host]
SRV1#
```

Step 13 At this point, you will start to migrate from a DHCP and PAT-based configuration to a configuration that uses a static IP address and NAT. On R1, set the IP address of Ethernet0/3 to 198.51.100.2/24.

On R1, enter the following command:

```
R1# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
R1(config-if)# int e0/3
R1(config-if)# ip address 198.51.100.2 255.255.255.0
R1(config-if)# exit
R1(config)#
```

Step 14 From the configuration, use the **do** command to execute the **show ip interface brief** command to verify the configuration on Ethernet0/3.

On R1, enter the following command:


```
R1(config)# do show ip int brief
Interface                  IP-Address      OK? Method Status
Protocol
Ethernet0/0                10.10.1.1       YES NVRAM  up                    up
Ethernet0/1                10.10.2.1       YES NVRAM  up                    up
Ethernet0/2                unassigned      YES NVRAM  administratively down
down
Ethernet0/3                198.51.100.2    YES manual up                    up
Serial1/0                  unassigned      YES NVRAM  administratively down
down
Serial1/1                  unassigned      YES NVRAM  administratively down
down
Serial1/2                  unassigned      YES NVRAM  administratively down
down
Serial1/3                  unassigned      YES NVRAM  administratively down
down
NVI0                       10.10.1.1       YES unset  up                    up
```

Step 15 Statically configure R1 to use the R2 interface as its default route.

On R1, enter the following command:

```
R1(config)# ip route 0.0.0.0 0.0.0.0 198.51.100.1
```

Step 16 Remain in the configuration mode and verify that the default route is now in the routing table.

On R1, enter the following command:

```
R1(config)# do 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
       + - replicated route, % - next hop override

Gateway of last resort is 198.51.100.1 to network 0.0.0.0

S*    0.0.0.0/0 [1/0] via 198.51.100.1
      10.0.0.0/8 is variably subnetted, 4 subnets, 2 masks
C        10.10.1.0/24 is directly connected, Ethernet0/0
L        10.10.1.1/32 is directly connected, Ethernet0/0
C        10.10.2.0/24 is directly connected, Ethernet0/1
L        10.10.2.1/32 is directly connected, Ethernet0/1
      198.51.100.0/24 is variably subnetted, 2 subnets, 2 masks
C        198.51.100.0/24 is directly connected, Ethernet0/3
L        198.51.100.2/32 is directly connected, Ethernet0/3
```

Step 17 Leave the inside and outside NAT configurations on the R1 interfaces, but remove the PAT configuration statement from the running configuration of R1.

On R1, enter the following command:

```
R1(config)# no ip nat inside source list 10 interface Ethernet0/3 overload
```

Step 18 Add a static NAT configuration entry that translates the SRV1 IP address (10.10.2.20) to 198.51.100.20. Then leave the configuration mode.

On R1, enter the following command:

```
R1(config)# ip nat inside source static 10.10.2.20 198.51.100.20
R1(config)# end
R1#
```

Step 19 Display the translation table on R1.

On R1, enter the following command:

```
R1# show ip nat translations
Pro Inside global      Inside local       Outside local      Outside global
--- 198.51.100.20      10.10.2.20         ---                ---
```

Static translations continuously remain in the translation table, regardless of their use.

Step 20 Access the console of SRV1 and establish a Telnet session to SRV2.

On SRV1, enter the following command:

```
SRV1# telnet 203.0.113.30
Trying 203.0.113.30 ... Open

User Access Verification

Username: admin
Password: Cisco123
SRV2>
```

Step 21 Return to the console of R1 and display the translation table while the session from SRV1 to SRV2 is open.

On R1, enter the following command:

```
R1# show ip nat translations
Pro Inside global      Inside local       Outside local      Outside global
tcp 198.51.100.20:23024 10.10.2.20:23024  203.0.113.30:23    203.0.113.30:23
--- 198.51.100.20      10.10.2.20         ---                ---
```

This example shows two entries in the translation table.

The first entry is an extended entry because it embodies more details than just an IP address mapping to an IP address. In this case, it specifies the protocol (TCP) and also the ports in use on both systems.

The second entry is a simple entry. It simply maps one IP address to another.

The extended entry is due to the use of the static translation for the Telnet session from SRV1 to SRV2. It details the characteristics of that session.

The simple entry is the persistent entry that is associated with the configured static translation.

Step 22 The most common use for static NAT translations is to provide a persistent IP address that the systems in the public IP address space can use to communicate with specific systems in the private IP address space. Demonstrate this function. Access the console of SRV2 and establish a Telnet connection back to SRV1.

On SRV2, enter the following command:

```
SRV2# telnet 198.51.100.20
Trying 198.51.100.20 ... Open


User Access Verification

Username: admin
Password: Cisco123
SRV1>
```

Step 23 With the two Telnet connections running, return to the console of R1 and view the translation table.

On R1, enter the following command:

```
R1# show ip nat translations
Pro Inside global      Inside local       Outside local      Outside global
tcp 198.51.100.20:23   10.10.2.20:23      203.0.113.30:46401 203.0.113.30:46401
tcp 198.51.100.20:29158 10.10.2.20:29158  203.0.113.30:23    203.0.113.30:23
--- 198.51.100.20      10.10.2.20         ---                ---
```

There is the one simple entry that is associated with the configured static translation, and two extended entries, each associated with an active session.

Step 24 Access the console of SRV2 and terminate the Telnet session to SRV1.

On SRV2, enter the following command:

```
SRV1> exit

[Connection to 198.51.100.20 closed by foreign host]
SRV2#
```

Step 25 Access the console of SRV1 and terminate the Telnet session to SRV2.

On SRV1, enter the following command:

```
SRV2> exit

[Connection to 203.0.113.30 closed by foreign host]
SRV1#
```

This is the end of the discovery lab.

Understanding Dynamic NAT

While static NAT provides a permanent mapping between an internal address and a specific public address, dynamic NAT maps private IP addresses to public addresses. These public IP addresses come from a NAT pool. Dynamic NAT configuration differs from static NAT, but it also has some similarities. Like static NAT, it requires the configuration to identify each interface as an inside or outside interface. However, rather than creating a static map to a single IP address, a pool of inside global addresses is used.

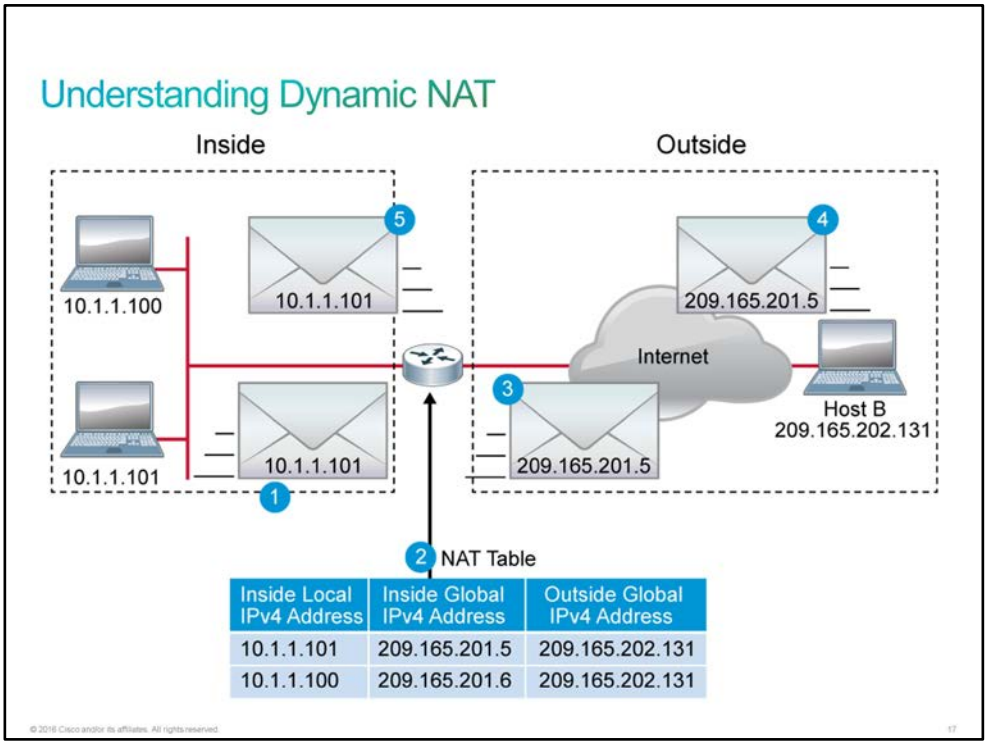

The figure illustrates a router that is translating a source address inside a network into a source address outside the network. The following are the steps for translating an inside source address:

1. The users at hosts 10.1.1.100 and 10.1.1.101 want to open a connection to Host B (IP address 209.165.202.131).
2. The first packet that the router receives from host 10.1.1.101 causes the router to check its NAT table. If no static translation entry exists, the router determines that the source address 10.1.1.101 must be translated dynamically. The router then selects a legal global address from the dynamic address pool and creates a translation entry (in this example, 209.165.201.5). This type of entry is called a *simple entry*. For the second host, 10.1.1.100, the router selects a legal global address from the dynamic address pool and creates a second translation entry (in this example, 209.165.201.6).
3. The router replaces the inside local source address of host 10.1.1.101 with the translated inside global address and forwards the packet.
4. Host B receives the packet and responds to host 209.165.201.5, using the inside global IPv4 destination address 209.165.201.5. When Host B receives the second packet, it responds to host 209.165.201.6, using the inside global IPv4 destination address 209.165.201.6.

5. When the router receives the packet with the inside global IPv4 address of 209.165.201.5, the router performs a NAT table lookup using the inside global address as a key. The router then translates the address back to the inside local address of host 10.1.1.101 and forwards the packet to host 10.1.1.101. When the router receives the packet with the inside global IPv4 address of 209.165.201.6, the router performs a NAT table lookup using the inside global address as a key. The router then translates the address back to the inside local address of host 10.1.1.100 and forwards the packet to host 10.1.1.100.
6. Hosts 10.1.1.100 and 10.1.1.101 receive the packets and continue the conversation. The router performs Steps 2 through 5 for each packet.

Configuring and Verifying Dynamic NAT

Configuring Dynamic NAT

Command	Description
interface *interface*	Specifies an interface and enters the interface configuration mode
ip nat pool *pool_name start_ip end_ip* **netmask** *netmask*	Defines an IP address pool
ip nat inside source list *acl_number* **pool** *pool_name*	Establishes a dynamic source translation by specifying the ACL and the address pool
ip address *address subnet_mask*	Sets the IP address and mask
ip nat inside	Marks the interface as connected to the inside network
ip nat outside	Marks the interface as connected to the outside network
access-list *acl_number* **permit** *ip_address netmask*	Creates an access list that defines the inside local addresses that are eligible to be translated

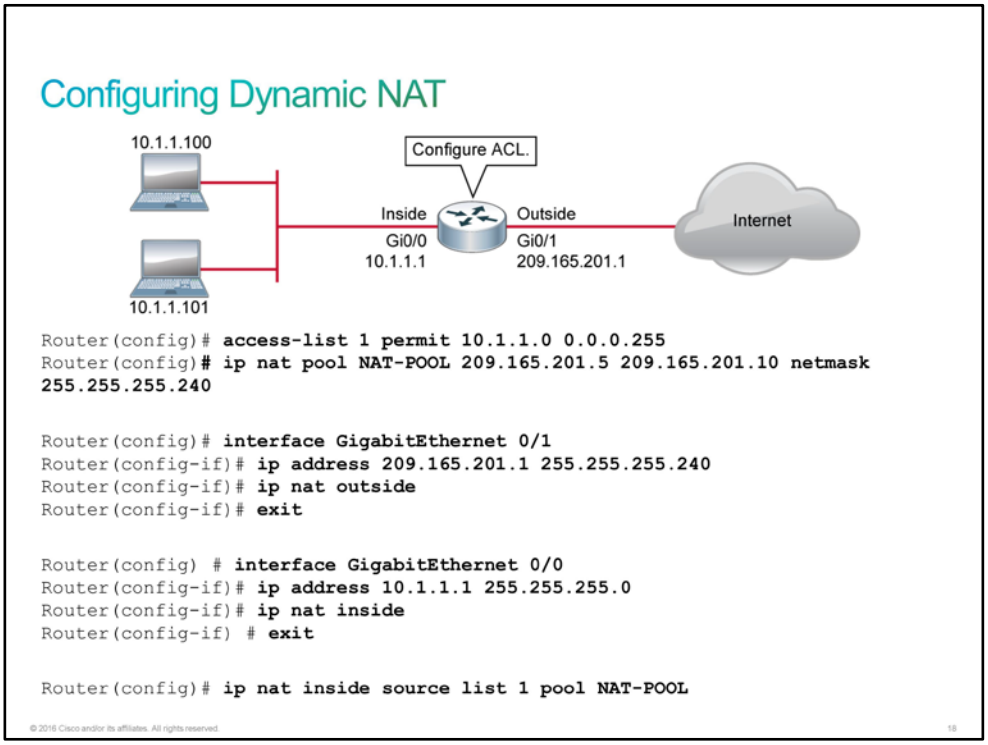

Note	The ACL must permit only those addresses that need to be translated. Remember that there is an implicit **deny any** statement at the end of each ACL. An ACL that is too permissive can lead to unpredictable results. Using **permit any** can result in NAT consuming too much router resources, which can cause network problems.

Verifying Dynamic NAT Configuration

Command	Description
show ip nat translations	Displays active NAT translations

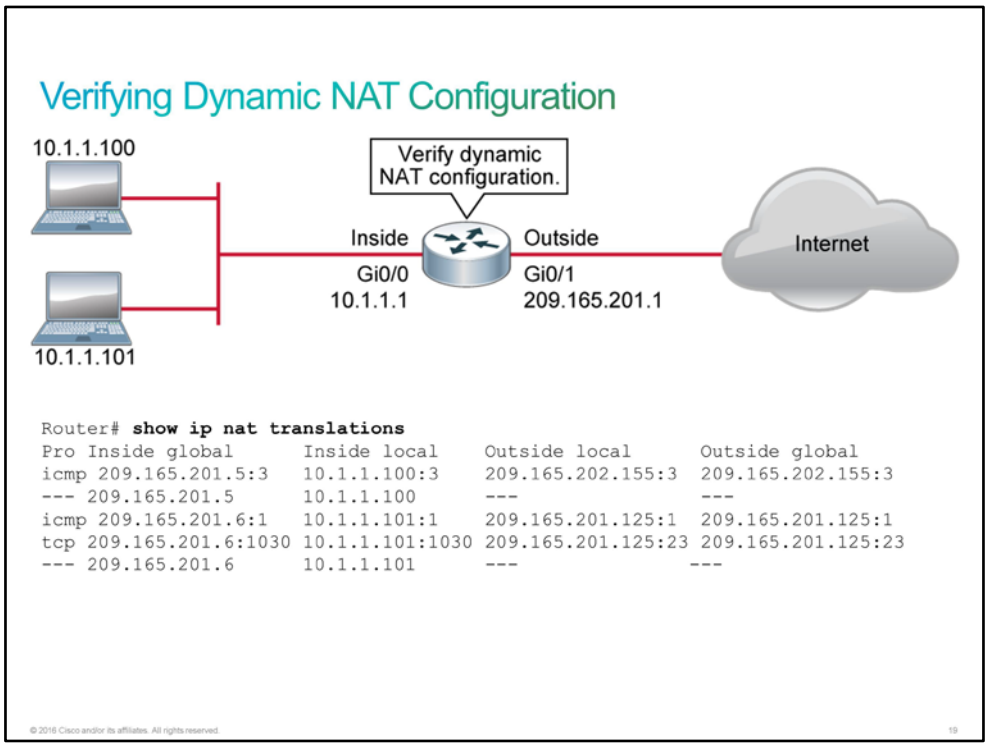

For more details about the **ip nat inside**, **ip nat pool**, **show ip nat translations** and related commands, check the *Cisco IOS IP Addressing Services Command Reference* at http://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipaddr/command/ipaddr-cr-book.html.

Understanding PAT

One of the main forms of NAT is PAT, which is also referred to as *overload* in Cisco IOS configuration. Several inside local addresses can be translated using NAT into just one or a few inside global addresses by using PAT. Most home routers operate in this manner. Your ISP assigns one address to your router, yet several members of your family can simultaneously surf the Internet.

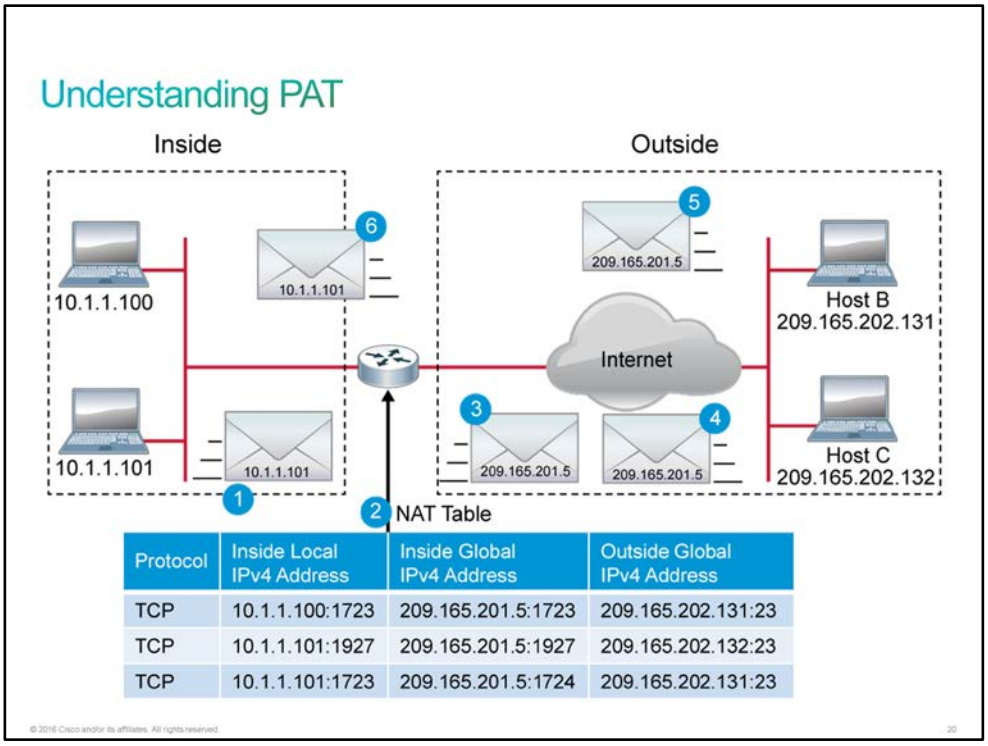

Protocol	Inside Local IPv4 Address	Inside Global IPv4 Address	Outside Global IPv4 Address
TCP	10.1.1.100:1723	209.165.201.5:1723	209.165.202.131:23
TCP	10.1.1.101:1927	209.165.201.5:1927	209.165.202.132:23
TCP	10.1.1.101:1723	209.165.201.5:1724	209.165.202.131:23

With NAT overload, multiple addresses can be mapped to one or a few addresses because a TCP or UDP port number tracks each private address. When a client opens a TCP/IP session, the NAT router assigns a port number to its source address. NAT overload ensures that clients use a different TCP or UDP port number for each client session with a server on the Internet. When a response comes back from the server, the source port number (which becomes the destination port number on the return trip) determines the client to which the router routes the packets. It also validates that the incoming packets were requested, which adds a degree of security to the session.

- PAT uses unique source port numbers on the inside global IPv4 address to distinguish between translations. Because the port number is encoded in 16 bits, the total number of internal addresses that NAT can translate into one external address is, theoretically, as many as 65,536.
- PAT attempts to preserve the original source port. If the source port is already allocated, PAT attempts to find the first available port number. It starts from the beginning of the appropriate port group, 0 to 511, 512 to 1023, or 1024 to 65535. If PAT does not find an available port from the appropriate port group and if more than one external IPv4 address is configured, PAT moves to the next IPv4 address and tries to allocate the original source port again. PAT continues trying to allocate the original source port until it runs out of available ports and external IPv4 addresses.

NAT generally translates IP addresses only as a 1:1 correspondence between publicly exposed IP addresses and privately held IP addresses. NAT overload modifies the private IP address and potentially the port number of the sender. NAT overload chooses the port numbers that hosts see on the public network.

NAT routes incoming packets to their inside destination by referring to the incoming destination IP address given by the host on the public network. With NAT overload, there is generally only one publicly exposed IP address (or a very few). Incoming packets from the public network are routed to their destinations on the private network by referring to a table in the NAT overload device that tracks public and private port pairs. This mechanism is called *connection tracking*.

The figure illustrates a PAT operation when one inside global address represents multiple inside local addresses. The TCP port numbers act as differentiators. Both Hosts B and C think that they are talking to a single host at the address 209.165.201.5. They are actually talking to different hosts, and the port number is the differentiator. In fact, many inside hosts could share the inside global IPv4 address by using many port numbers.

The router performs this process when it overloads inside global addresses:

1. The user at host 10.1.1.100 opens a connection to Host B. A second user at host 10.1.1.101 opens a connection to Hosts B and C.
2. The first packet that the router receives from host 10.1.1.100 causes the router to check its NAT table. If no translation entry exists, the router determines that address 10.1.1.100 must be translated and sets up a translation of the inside local address 10.1.1.100 into an inside global address. If overloading is enabled and another translation is active, the router reuses the inside global address from that translation and saves enough information to be able to translate back. This type of entry is called an extended entry.
3. The router replaces the inside local source address 10.1.1.100 with the selected inside global address 209.165.201.5 and forwards the packet.
4. Host B receives the packet and responds to host 10.1.1.100, using the inside global IPv4 address 209.165.201.5. Host C receives a packet with the same inside global address, even though the packet originated from host 10.1.1.101.
5. When the router receives the packet with the inside global IPv4 address, the router performs a NAT table lookup. Using the inside global address and port and outside global address and port as a key, the router translates the address back into the correct inside local address, 10.1.1.100, and forwards the packet to host 10.1.1.100.
6. Host 10.1.1.100 receives the packet and continues the conversation. The router performs Steps 2 through 5 for each packet.

Configuring and Verifying PAT

Configuring PAT

Command	Description
interface *interface*	Specifies an interface and enters the interface configuration mode
ip address *address subnet_mask*	Sets the IP address and mask
ip nat inside	Marks the interface as connected to the inside network
ip nat outside	Marks the interface as connected to the outside network
ip nat inside source list *access-list-number* **interface** *interface* **overload**	Establishes dynamic source translation, specifying the ACL
access-list *acl_number* **permit** *ip_address netmask*	Creates an ACL that defines the inside local addresses that are eligible to be translated

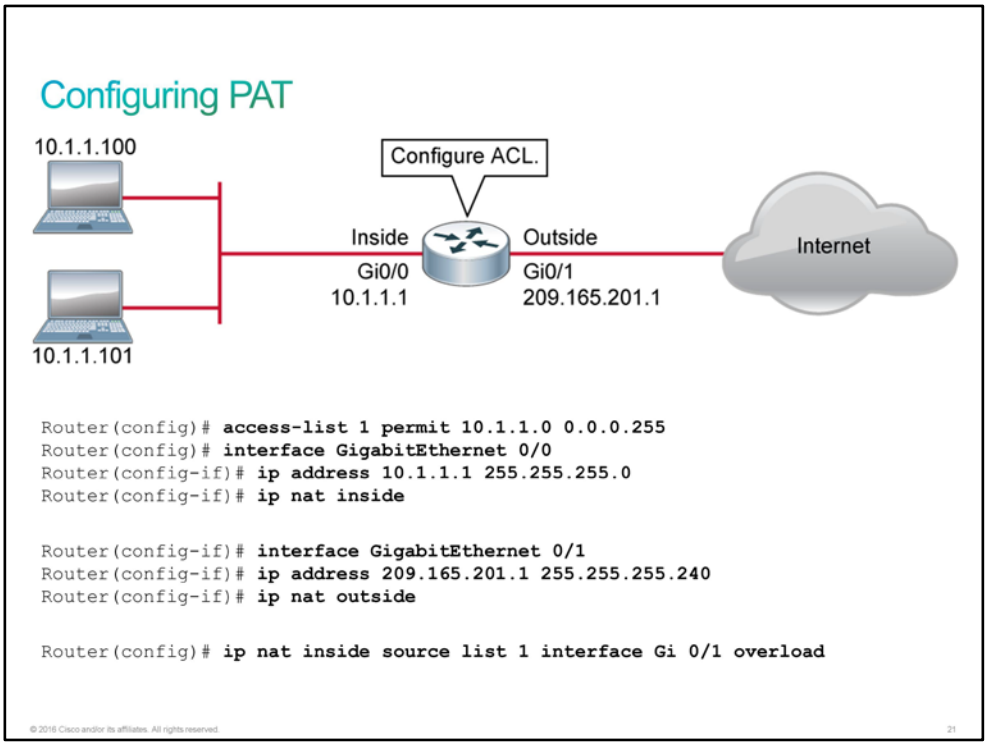

Verifying PAT Configuration

Command	Description
show ip nat translations	Displays active NAT translations

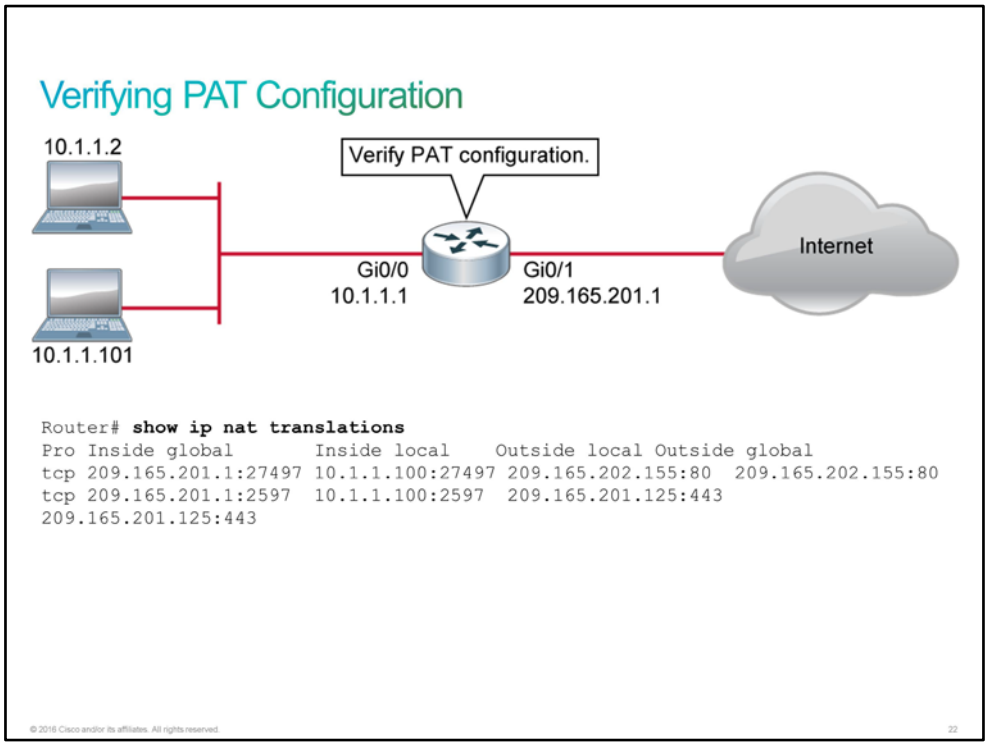

For more details about the **ip nat inside**, **ip nat pool**, and **show ip nat translations** and related commands, check the *Cisco IOS IP Addressing Services Command Reference* at http://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipaddr/command/ipaddr-cr-book.html.

Discovery 14: Configure Dynamic NAT and PAT

Introduction

In this discover lab, you will implement a dynamic NAT pool that other systems on the internal network can share for outbound connectivity.

Topology

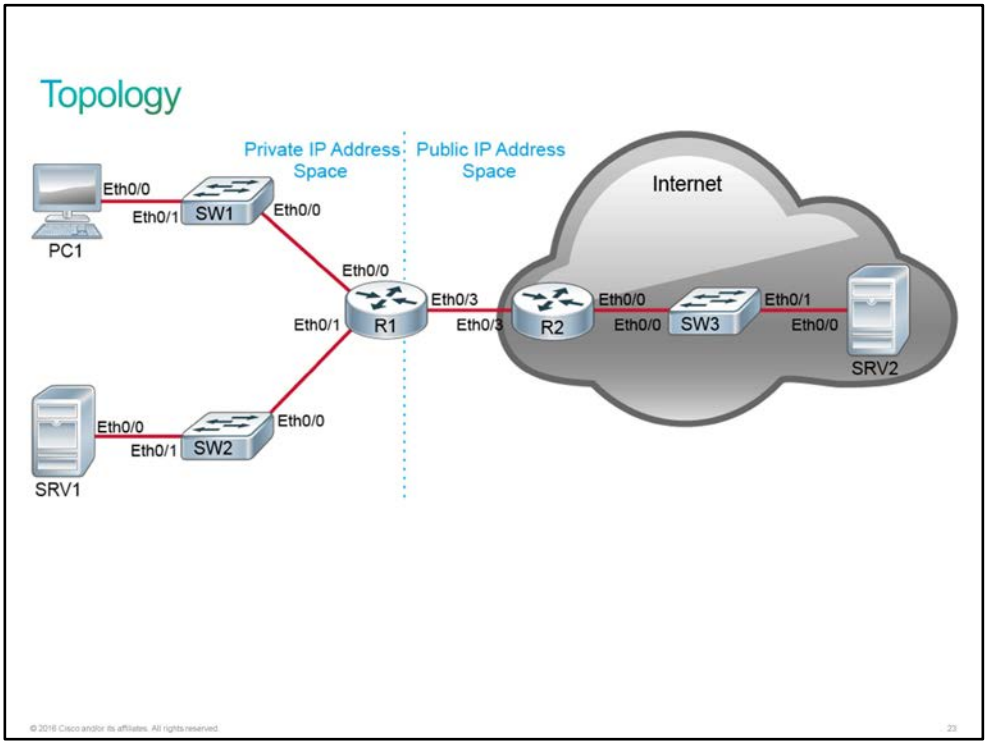

Job Aid

Device Information

Device Information Table

Device	Characteristic	Value
PC1	Hostname	PC1
PC1	IP address	10.10.1.10/24
PC1	Default gateway	10.10.1.1
SRV1	Hostname	SRV1
SRV1	IP address	10.10.2.20/24

Device	Characteristic	Value
SRV1	Default gateway	10.10.2.1
SRV2	Hostname	SRV2
SRV2	IP address	203.0.113.30/24
SRV2	Default gateway	203.0.113.1
SW1	Hostname	SW1
SW1	VLAN 1 IP address	10.10.1.4/24
SW1	Default gateway	10.10.1.1
SW1	Ethernet0/0 description	Link to R1
SW1	Ethernet0/1 description	Link to PC1
SW2	Hostname	SW2
SW2	VLAN 1 IP address	10.10.2.4/24
SW2	Default gateway	10.10.2.1
SW2	Ethernet0/0 description	Link to R2
SW2	Ethernet0/1 description	Link to PC2
SW3	Hostname	SW3
SW3	VLAN 1 IP address	203.0.113.4/24
SW3	Default gateway	203.0.113.1
SW3	Ethernet0/0 description	Link to R3
SW3	Ethernet0/1 description	Link to SRV2
R1	Hostname	R1
R1	Ethernet0/0 description	Link to SW1
R1	Ethernet0/0 IP address	10.10.1.1/24
R1	Ethernet0/1 description	Link to SW2
R1	Ethernet0/1 IP address	10.10.2.1/24
R1	Ethernet0/3 description	Link to R2

Device	Characteristic	Value
R1	Ethernet0/3 IP address	198.51.100.2/24
R2	Hostname	R2
R2	Ethernet0/0 description	Link to SW3
R2	Ethernet0/0 IP address	203.0.113.1/24
R2	Ethernet0/3 description	Link to R1
R2	Ethernet0/3 IP address	198.51.100.1/24

PC and SRVs in the virtual lab environment are simulated as routers, so you should use Cisco IOS commands to configure them or make verifications.

Global IP Address Networks

Address Block	Host Starting Address	Host Ending Address	Broadcast Address	Subnet Mask
192.0.2.0/24	192.0.2.1	192.0.2.254	192.0.2.255	255.255.255.0
198.51.100.0/24	198.51.100.1	198.51.100.254	198.51.100.255	255.255.255.0
203.0.113.0/24	203.0.113.1	203.0.113.254	203.0.113.255	255.255.255.0

Task 1: Configure Dynamic NAT and PAT

Activity

Step 1 The static NAT configuration is in place. It is now time to explore dynamic NAT translation by using a pool of IP addresses. First, access the console of PC1 and verify that it cannot ping SRV2.

On PC1, enter the following command:

```
PC1# ping 203.0.113.30
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 203.0.113.30, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
```

The failure of the ping process is not a failure on the delivery of the echo requests from PC1 to SRV2. It is a failure on the return of the echo replies from SRV2 to PC1. Because there is no translation that is configured for PC1, SRV2 receives the echo requests that come from 10.10.1.10. When SRV2 tries to reply to 10.10.1.10 from within the public IP address space, R2 does not have a route that it can to use to get to the private IP address space.

Step 2 On R1, define a pool of NAT addresses named "NatPool" by specifying the address range from 198.51.100.100 to 198.51.100.149.

On R1, enter the following command:

```
R1# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
R1(config)# ip nat pool NatPool 198.51.100.100 198.51.100.149 netmask
255.255.255.0
```

Step 3 Verify that access list 10 is still in place, permitting addresses from 10.10.0.0/16.

On R1, enter the following command:

```
R1(config)# do show access-list 10
Standard IP access list 10
    10 permit 10.10.0.0, wildcard bits 0.0.255.255
```

On R1, verify which interfaces are NAT inside and NAT outside.

```
R1(config)# do show ip nat statistics
Total active translations: 1 (1 static, 0 dynamic; 0 extended)
Peak translations: 1, occurred 00:04:10 ago
Outside interfaces:
  Ethernet0/3
Inside interfaces:
  Ethernet0/0, Ethernet0/1
Hits: 0  Misses: 0
CEF Translated packets: 0, CEF Punted packets: 0
Expired translations: 0
Dynamic mappings:

Total doors: 0
Appl doors: 0
Normal doors: 0
Queued Packets: 0
```

Step 4 Define a dynamic translation rule that specifies access list 10 as the source and that uses addresses from the pool NatPool. Then leave the configuration mode.

On R1, enter the following commands:

```
R1(config)# ip nat inside source list 10 pool NatPool
R1(config)# end
R1#
```

Step 5 Verify that there is now bidirectional connectivity between PC1 and SRV2. Access the console of PC1 and send a ping to SRV2.

On PC1, enter the following command:

```
PC1# ping 203.0.113.30
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 203.0.113.30, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms
```

Step 6 Access the console of R1 and view the current translation table.

On R1, enter the following command:

```
R1# show ip nat translations
Pro Inside global      Inside local       Outside local      Outside global
icmp 198.51.100.100:4  10.10.1.10:4       203.0.113.30:4     203.0.113.30:4
--- 198.51.100.100     10.10.1.10         ---                ---
--- 198.51.100.20      10.10.2.20         ---                ---
```

If you proceeded quickly enough, three translations will be in the table.

The extended translation that is associated with the ICMP session is short-lived and may have timed out. If it did, you can re-send the ping from PC1 and display the translation table again.

There is a simple entry in the table that is associated with the assignment of an address from the pool to PC1. By default, dynamic translations that are assigned from a NAT pool have a 24-hour inactivity timeout. So the translation for PC1 to 198.51.100.100 will persist as long as it is used at least once per day.

The third entry that is translating 10.10.2.20 to 198.51.100.20 is the static entry.

Step 7 One at a time, access the consoles of PC1, SW1, and SW2. From PC1, execute a Telnet session to SRV2. From SW1 and SW2, send pings to SRV2.

On PC1, enter the following command:

```
PC1# telnet 203.0.113.30
Trying 203.0.113.30 ... Open


User Access Verification

Username: admin
Password: Cisco123
SRV2>
```

On SW1, enter the following command:

```
SW1# ping 203.0.113.30
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 203.0.113.30, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/202/1008 ms
```

On SW2, enter the following command:


```
SW2# ping 203.0.113.30
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 203.0.113.30, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/201/1003 ms
```

Step 8 Return to the console of R1 and view the translation table.

On R1, enter the following command:

```
R1# show ip nat translations
Pro Inside global      Inside local       Outside local      Outside global
icmp 198.51.100.101:0  10.10.1.4:0        203.0.113.30:0     203.0.113.30:0
--- 198.51.100.101     10.10.1.4          ---                ---
tcp 198.51.100.100:24829 10.10.1.10:24829 203.0.113.30:23    203.0.113.30:23
--- 198.51.100.100     10.10.1.10         ---                ---
icmp 198.51.100.102:0  10.10.2.4:0        203.0.113.30:0     203.0.113.30:0
--- 198.51.100.102     10.10.2.4          ---                ---
--- 198.51.100.20      10.10.2.20         ---                ---
```

The extended ICMP entries that are associated with the ping activity are short-lived. You can always try to resend the ping and re-display the translation table.

SW1 (10.10.1.4) and SW2 (10.10.2.4) have been assigned IP addresses from the NAT pool. Again, there is a 24-hour inactivity timeout on these dynamic entries by default.

Step 9 Display the translation statistics on R1.

On R1, enter the following command:

```
R1# sh ip nat statistics
Total active translations: 5 (1 static, 4 dynamic; 1 extended)
Peak translations: 7, occurred 00:03:10 ago
Outside interfaces:
  Ethernet0/3
Inside interfaces:
  Ethernet0/0, Ethernet0/1
Hits: 112  Misses: 0
CEF Translated packets: 112, CEF Punted packets: 0
Expired translations: 7
Dynamic mappings:
-- Inside Source
[Id: 1] access-list 10 pool NatPool refcount 4
 pool NatPool: netmask 255.255.255.0
        start 198.51.100.100 end 198.51.100.149
        type generic, total addresses 50, allocated 3 (6%), misses 0

Total doors: 0
Appl doors: 0
Normal doors: 0
Queued Packets: 0
```

The statistics that are displayed in the lab environment will likely differ from the example. But, in any case, statistics include the current status such as the current active translation count, historical statistics such as the largest number of translations seen on R1, and configuration information such as the details of the NAT pools.

Step 10 Return to the console of PC1 and terminate the Telnet session to SRV2.

On PC1, enter the following command:

```
SRV2> exit

[Connection to 203.0.113.30 closed by foreign host]
PC1#
```

Step 11 Return to the console of R1 and clear all dynamic translations from the translation table.

On R1, enter the following command:

```
R1# clear ip nat translation *
```

Step 12 Display the translation table, verifying the removal of the dynamic entries.

On R1, enter the following command:

```
R1# show ip nat translation
Pro Inside global      Inside local       Outside local      Outside global
--- 198.51.100.20      10.10.2.20         ---                ---
```

The dynamic entries have been removed, but the statically configured entry for SRV1 remains.

Feel free to continue with independent exploration of NAT concepts within the lab environment.

This is the end of the discovery lab.

Troubleshooting NAT

When you have IPv4 connectivity problems in a NAT environment, it is often difficult to determine the cause of the problem. NAT is often blamed, when in reality there is an underlying problem. When you are trying to determine the cause of an IPv4 connectivity problem, it helps to eliminate NAT as the problem.

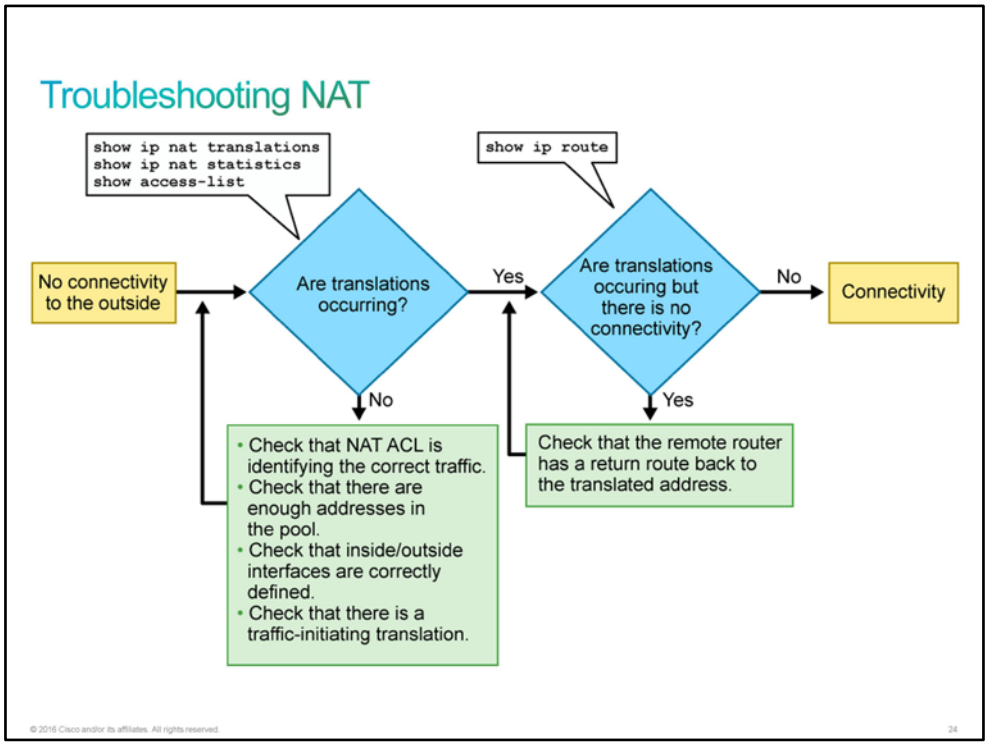

Follow these steps to verify that NAT is operating as expected:

1. Verify that translations are occurring:
   - Use the **show ip nat translations** command to determine if translations exist in the translation table.
   - Verify that the translation is actually occurring by using the **show ip nat statistics** and **debug ip nat** commands.
   - Use the **show access-list** command to verify that the ACL that is associated with the NAT command is permitting all necessary networks.
   - Use the **show ip nat statistics** command to verify that the router interfaces are appropriately defined as NAT inside or NAT outside.
   - If some devices have connectivity but others do not, the NAT pool might be out of addresses.
2. If translations are occurring but there is no connectivity, use the **show ip route** command to verify that there is a return route to the translated address.

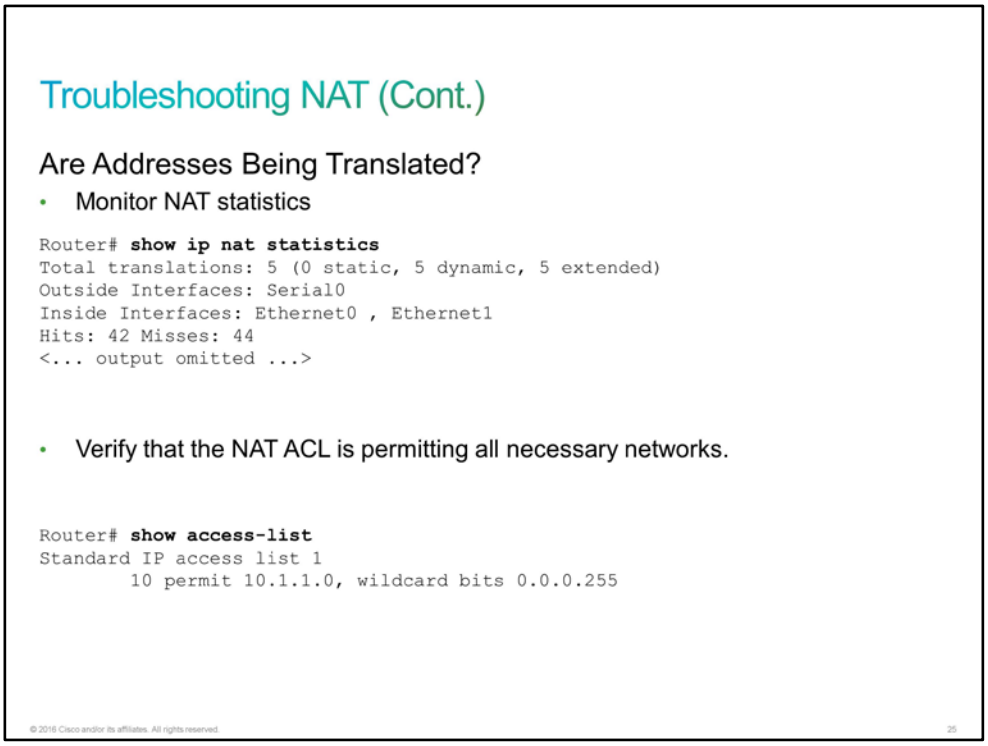

In a simple network environment, it is useful to monitor NAT statistics with the **show ip nat statistics** command. The **show ip nat statistics** command displays information about the total number of active translations, NAT configuration parameters, the number of addresses in the pool, and the numbers that have been allocated. However, in a more complex NAT environment, with several translations taking place, this **show** command may not clearly identify the issue. It may be necessary to run **debug** commands on the router.

Note	You can use the **clear ip nat translation *** command to clear all dynamic address translation entries. By default, translation entries time out after 24 hours. When testing the NAT configuration, it can be useful to clear translations.

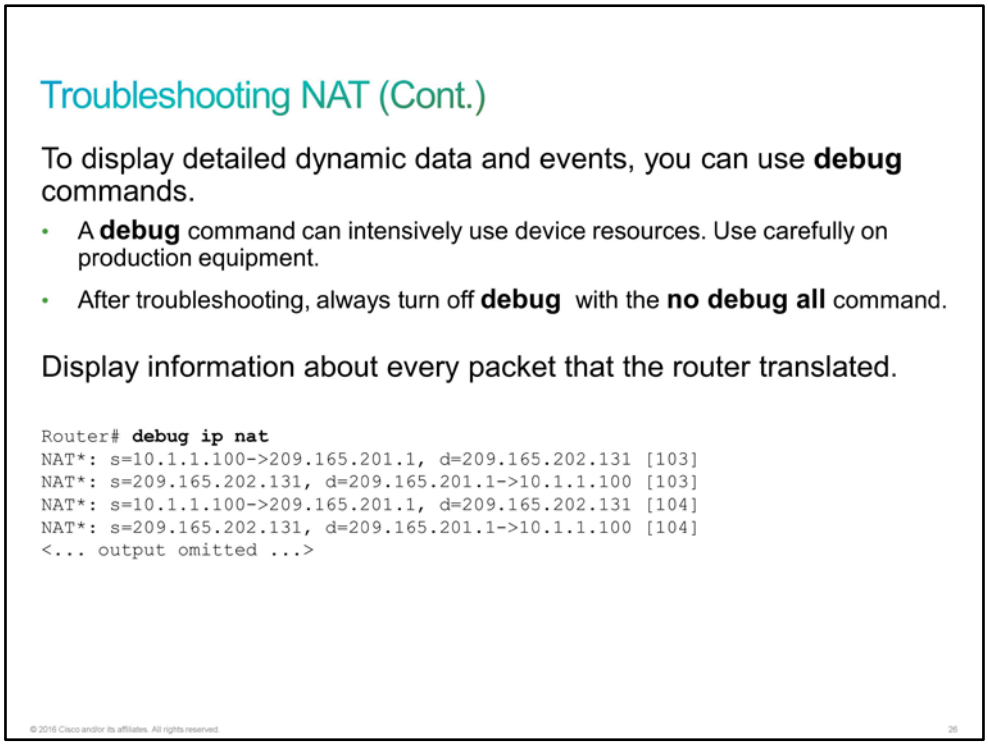

Note	The **debug** command, especially the **debug all** command, should be used sparingly. These commands can disrupt router operations. The **debug** commands are useful when configuring or troubleshooting a network. However, they can make intensive use of CPU and memory resources. It is recommended that you run as few debug processes as necessary and disable them immediately when they are no longer needed. On production networks, you should use the **debug** commands with caution because they can affect the performance of the device.

The **debug ip nat** command displays information about every packet that the router translates, which helps you to verify NAT operation. The **debug ip nat detailed** command generates a description of each packet that is considered for translation. This command also provides information about certain errors or exception conditions, such as the failure to allocate a global address. The **debug ip nat detailed** command generates more overhead than the **debug ip nat** command, but it can provide the detail that you need to troubleshoot the NAT problem. Always remember to turn off debugging when finished.

The example shows a sample **debug ip nat** output. In the output, you can see that the inside host 10.1.1.100 initiated traffic to the outside host 209.165.202.131 and has been translated to the address 209.165.201.1.

For decoding the **debug** output, note what the following symbols and values indicate:

- *****: The asterisk next to "NAT" indicates that the translation is occurring in the fast-switched path. The first packet in a conversation is always process-switched, which is slower. The remaining packets go through the fast-switched path if a cache entry exists.
- **s=**: Refers **t**o the source IP address.

- **a.b.c.d->w.x.y.z**: Indicates that source address a.b.c.d is translated to w.x.y.z.
- **d=**: Refers to the destination IP address.
- **[xxxx]**: The value in brackets is the IP identification number. This information may be useful for debugging because it enables correlation with other packet traces from protocol analyzers.

Finally, you should make sure that the ACL that the NAT command references is permitting all the necessary networks. Notice that ACLs use wildcard masks and not subnet masks.

If translations are occurring, but there is no connectivity, verify that the remote router has a route to the translated address.

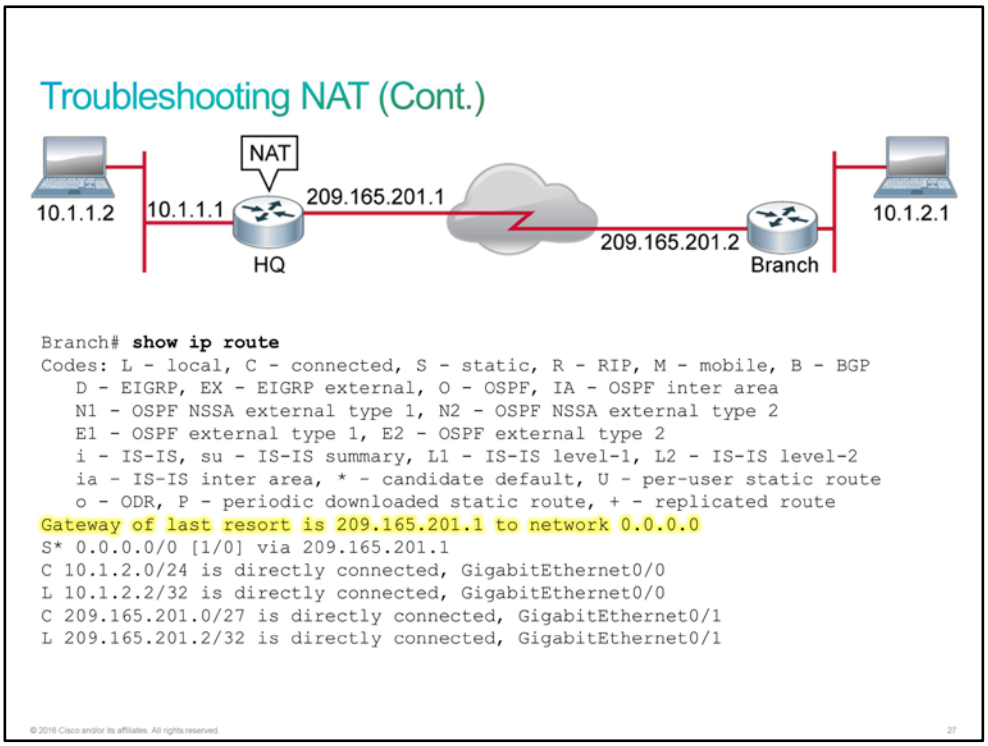

If translations are occurring but a ping to the remote network still fails, the issue might be a missing route back to the translated address. This problem can arise in NAT between a headquarters and branch office. It is usually not an issue when connecting to an ISP, because the service provider takes care of routing all the necessary traffic back to the customer.

Discovery 15: Troubleshoot NAT

Introduction

In this discovery lab, you will use different **show** commands to troubleshoot common NAT-related issues.

Topology

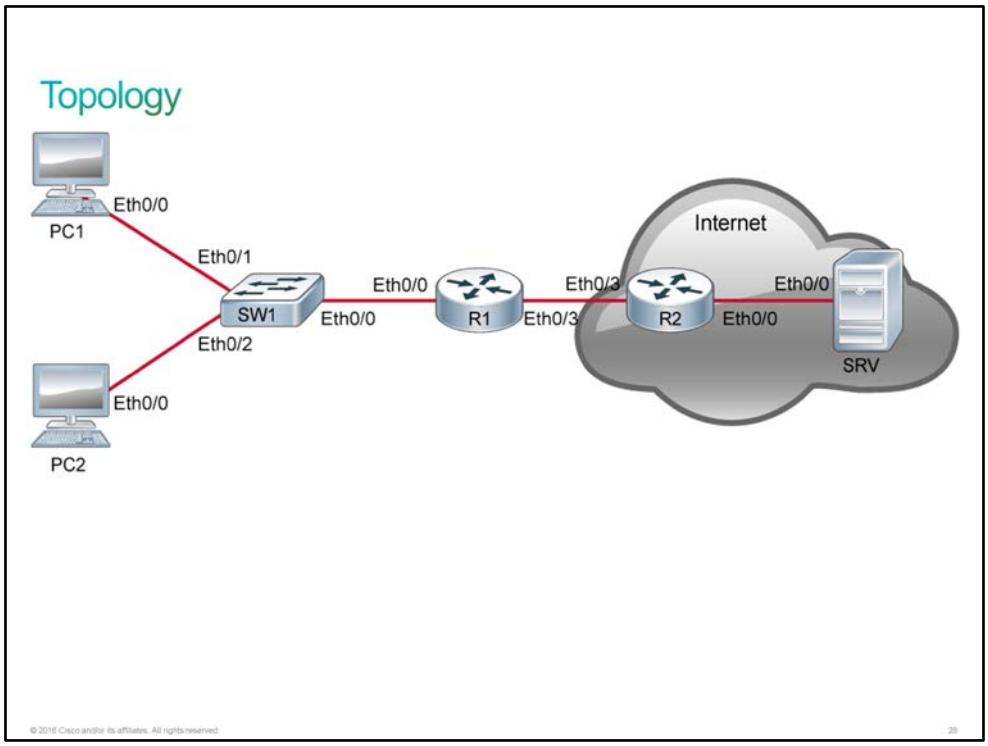

Job Aid

Device Information

Device Information Table

Device	Characteristic	Value
PC1	Hostname	PC1
PC1	IP address	10.10.1.10/24
PC1	Default gateway	10.10.1.1
PC2	Hostname	PC2
PC2	IP address	10.10.1.20/24
PC2	Default gateway	10.10.1.1

Device	Characteristic	Value
SRV	Hostname	SRV
SRV	IP address	203.0.113.30/24
SRV	Default gateway	203.0.113.1
SW1	Hostname	SW1
SW1	VLAN 1 IP address	10.10.1.4/24
SW1	Default gateway	10.10.1.1
SW1	Ethernet0/0 description	Link to R1
SW1	Ethernet0/1 description	Link to PC1
SW1	Ethernet0/2 description	Link to PC2
R1	Hostname	R1
R1	Ethernet0/0 description	Link to SW1
R1	Ethernet0/0 IP address	10.10.1.1/24
R1	Ethernet0/3 description	Link to R2
R1	Ethernet0/3 IP address	198.51.100.2/24
R2	Hostname	R2
R2	Ethernet0/0 description	Link to SRV
R2	Ethernet0/0 IP address	203.0.113.1/24
R2	Ethernet0/3 description	Link to R1
R2	Ethernet0/3 IP address	198.51.100.1/24

PC and SRV in the virtual lab environment are simulated as routers, so you should use Cisco IOS commands to configure them or make verifications.

Task 1: Troubleshoot NAT

Activity

Step 1 PC1 and SRV are unable to ping after a new NAT configuration is put in place.

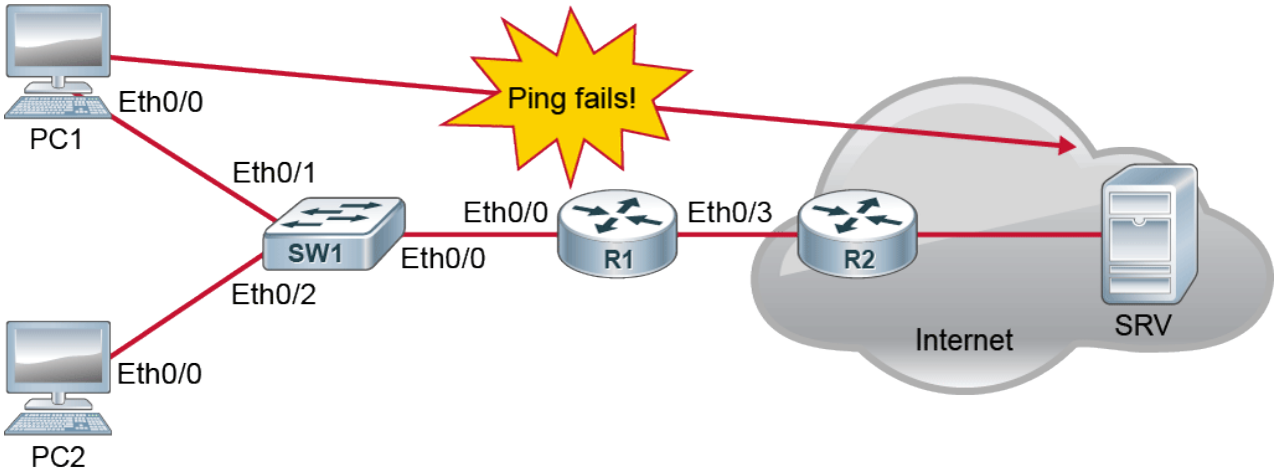

The figure shows that PC1 (10.10.1.10) cannot ping SRV (203.0.113.30). R1 router has a default gateway set to 198.51.100.1.

Ping from PC1 to SRV will fail.

```
PC1# ping 203.0.113.30
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 203.0.113.30, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
```

Verify relevant part of configuration on the R1 router.

```
R1# show running-config
<... output omitted ...>
interface Ethernet0/0
 description Link to SW1
 ip address 10.10.1.1 255.255.255.0
 ip nat outside
<... output omitted ...>
!
interface Ethernet0/3
 description Link to R2
 ip address 198.51.100.2 255.255.255.0
 ip nat inside
!
ip nat inside source list 20 interface Ethernet0/3 overload
ip route 0.0.0.0 0.0.0.0 198.51.100.1
!
access-list 20 permit 0.0.0.0 255.255.255.0
<... output omitted ...>
```

Step 2 To troubleshoot the problem, use the **show ip nat translation** command to see if any translations are currently in the table.

On R1, enter the following command:

```
R1# show ip nat translations
R1#
```

Translations are not occurring.

Step 3 Next, you must determine whether any translations have ever taken place and identify the interfaces between which translation should be occurring. You use the **show ip nat statistics** command.

On R1, enter the following command:

```
R1# show ip nat statistics
Total active translations: 0 (0 static, 0 dynamic; 0 extended)
Peak translations: 0
Outside interfaces:
  Ethernet0/0
Inside interfaces:
  Ethernet0/3
Hits: 0  Misses: 0
CEF Translated packets: 0, CEF Punted packets: 0
Expired translations: 0
Dynamic mappings:
-- Inside Source
[Id: 1] access-list 20 interface Ethernet0/3 refcount 0

Total doors: 0
Appl doors: 0
Normal doors: 0
Queued Packets: 0
```

Step 4 The NAT counters are at 0, verifying that no translation has occurred. The R1 router interfaces are incorrectly defined as NAT inside and NAT outside. Fix the R1 router configuration.

On R1, enter the following command:

```
R1# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
R1(config)# inter Eth0/0
R1(config-if)# no ip nat outside
R1(config-if)# ip nat inside
R1(config-if)# inter Eth0/3
R1(config-if)# no ip nat inside
R1(config-if)# ip nat outside
```

Verify connectivity between PC1 and SRV. Ping from PC1 to SRV will fail again.

```
PC1# ping 203.0.113.30
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 203.0.113.30, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
```

Step 5 Verify that the access list is correct.

On R1, enter the following command:

```
R1# show access-list
Standard IP access list 20
    10 permit 0.0.0.0, wildcard bits 255.255.255.0
```

Access list has wrong wildcard mask. Wildcard mask is matching on the fourth octet only. You will need to invert wildcard mask and define the correct network part of the access list.

Step 6 On the R1 router fix access list.

On R1, enter the following commands:

```
R1# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
R1(config)# no access-list 20
R1(config)# access-list 20 permit 10.10.1.0 0.0.0.255
```

Step 7 After you have corrected the wildcard mask, you generate another ping from PC1 to SRV. The connectivity test is now a success. Verify that translations are occurring and you have connectivity to the remote network.

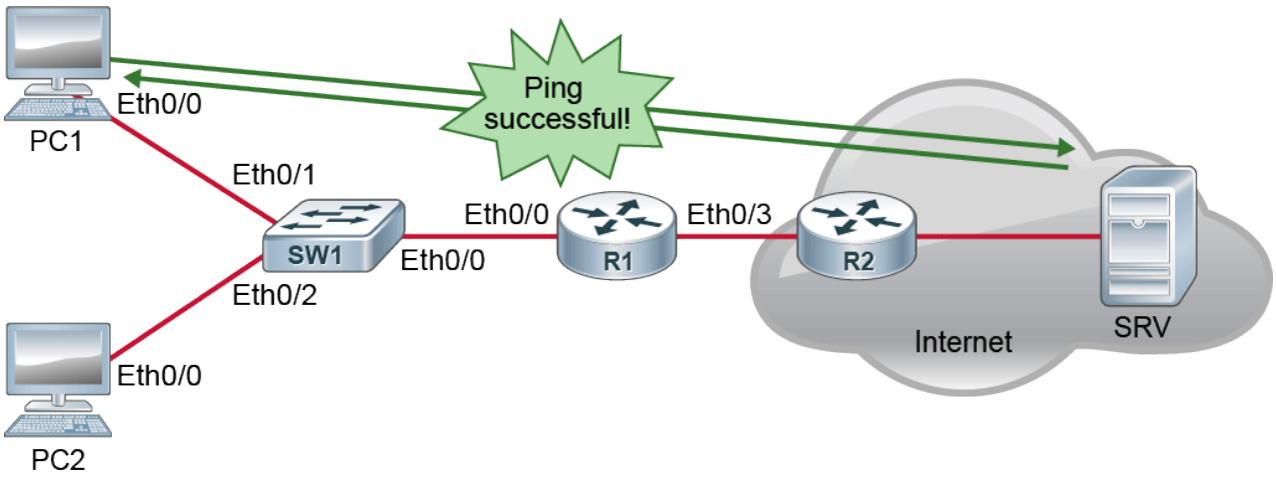

On PC1, enter the following command:

```
PC1# ping 203.0.113.30
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 203.0.113.30, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/2 ms
```

On R1, enter the following command:

```
R1# sh ip nat translations
Pro Inside global      Inside local       Outside local      Outside global
icmp 198.51.100.2:12   10.10.1.10:12      203.0.113.30:12    203.0.113.30:12
```

This is the end of the discovery lab.

Challenge

1. Which of the following is the customer side of the demarcation point?

   A. CPE
   B. CTE

2. The following are methods that are used to connect small offices to the internet. Which of them would you use if you had an environment filled with EMI and RFI?

   A. Copper
   B. Optic Fiber
   C. Wifi

3. **show ip nat translations** command displays which interfaces are enabled for NAT configuration on a router. True or False?

   A. True
   B. False

4. Which of the following is eliminated with the use of NAT? (Choose two.)

   A. Need to readdress all hosts that require external access
   B. IP Address conservation
   C. Revealing of private addresses outside of the network
   D. Performance problems and switching delays

5. Which of the following is the IPv4 address of an outside host as it appears to the inside network?

   A. Inside local address
   B. Inside global address
   C. Outside global address
   D. Outside local address

6. What is the difference between static NAT and dynamic NAT?

   A. Static NAT maps one-to-one and dynamic NAT maps one-to-many
   B. Static NAT maps one-to-many and dynamic NAT maps many-to-one
   C. Static NAT maps one-to-one and dynamic NAT maps many-to-many

7. Which Translation technology would most likely be used at home, especially for connecting devices such as tablets, phones, and PC's through the DSL internet connection?

   A. static NAT
   B. dynamic NAT
   C. PAT

Answer Key

Challenge

1. A
2. B
3. B
4. A, C
5. D
6. C
7. C

Module 3: Summary Challenge

Introduction

This module challenges you to use the knowledge and skill that you have obtained related in the *Building a Simple Network* and *Establishing Internet Connectivity* modules.

Lesson 1: Establish Internet Connectivity

Introduction

In this lesson, you are required to implement and establish internet connectivity.

Challenge

1. A router needs which of the following to allow the users inside it's network to connect to the internet? (Choose four.)

   A. A Default route to the ISP router
   B. A Static route to the ISP router
   C. A static NAT statement
   D. A PAT statement using an ACL
   E. The interface facing the internet to be designated as 'NAT Outside'
   F. The interface facing the internet to be designated as 'NAT Inside'
   G. An ACL permitting the users that need to be connected to the internet
   H. An ACL permitting the IP Addresses on the internet that need to be accessed

2. If you want to allow specific servers on the network to be accessible from the internet, which of the following would you need?

   A. A static NAT
   B. A dynamic NAT
   C. PAT

3. You need to create an ACL that will allow only users from the marketing VLAN to access certain servers on the internet. If you had to use a numbered ACL, which of the following choices would you use?

   A. 1
   B. 45
   C. 99
   D. 199
   E. 1399

4. When implementing static routing to enable internet access, which is the most suitable solution?

   A. All devices have a default route to the border router, the border router has a default route to the ISP Router
   B. All devices have a static route to the border router, the border router has a default route to the ISP Router
   C. All devices have a default route to the border router, the border router has a static route to the ISP Router
   D. All devices have a static route to the border router, the border router has a static route to the ISP Router

5. When specifying the next hop in a default route to the ISP network, which of the following can you use?

   A. Only IP Address of the ISP router
   B. Only the interface connected to the ISP Router
   C. Only the IP Address of the interface connected to the ISP Router
   D. Either the IP Address of the ISP router or the Interface connected to ISP Router

6. You have created an ACL named "Public_ACL" created to deny a set of public IP Addresses, and another ACL named "Internet_ACL" created to permit the user network to go out to the internet. Your manager has asked you to apply both ACLs on the interface that is facing the user network. How would you do it?

   A. **interface F0/0**
      ip access-group Public_ACL in
      ip access-group Internet_ACL out
   B. **interface F0/0**
      ip access-group Public_ACL out
      ip access-group Internet_ACL in
   C. interface F0/0
      ip access-group Public_ACL in
      ip access-group Internet_ACL in
   D. **interface F0/0**
      ip access-group Public_ACL out
      ip access-group Internet_ACL out

7. Which of the following commands ensures the IP Address that will be received by a DHCP client? (Choose two.)

   A. **ip dhcp pool *name***
   B. **utilization mark high *percentage-number***
   C. **network *network-number {[mask] | /prefix-length}***
   D. **domain-name *domain***
   E. **dns-server *address [address 2... address8]***
   F. **default-router *address [address2...address 8]***
   G. **lease {*[specific time]* | infinite}**

Answer Key

Challenge

1. A, D, E, G
2. A
3. D
4. A
5. D
6. B
7. A, C

Lesson 2: Troubleshoot Internet Connectivity

Introduction

In this lesson, you are required to Troubleshoot Internet Connectivity.

Challenge

1. You are troubleshooting the following ACL which is supposed to permit 192.168.123.1 and 192.168.123.2 but not the rest of the 192.168.123.0/24 subnet. Which of the options would be a solution to making the ACL work?

   ip access-list 12
   10 permit host 192.168.123.1
   20 permit host 192.168.123.2
   30 permit 192.168.123.0 0.0.0.255
   40 permit any

   A. Change ACL line 40 to deny any
   B. Change ACL line 30 to permit 192.168.123.0 255.255.255.0
   C. Change ACL line 30 to permit 192.168.123.0 255.255.0.0
   D. Change ACL line 30 to deny 192.168.123.0 0.0.0.255

2. Refer to the static route configuration on Router A. Which statement about interface serial0/0/0 is correct?

   RouterA(config)# ip route 172.16.1.0 255.255.255.0 serial0/0/0

   A. The interface serial0/0/0 configured in the static route configuration is the outbound interface of local router Router A.
   B. The interface serial0/0/0 configured in the static route configuration is the inbound interface on the remote router connected to Router A.
   C. It doesn't matter you can use either outbound interface of the local router Router A or inbound interface of the remote router connected to Router A.
   D. The static route configuration is incorrect, usage of exit interface is not accepted in IPv4 static route configuration.

3. Which of the following commands would you issue on a Cisco Router if you are looking for a device based on its MAC Address?

   A. **traceroute**
   B. **tracert**
   C. **show ip arp**
   D. **arp -a**

4. You see documentation of commands to be applied to a router to allow traffic between a PC with IP Address 192.168.1.1/24 and a Server with IP Address 172.16.1.1 /24.Inspect the ACL configuration. What needs to be done to make it work?

   access-list 101 permit ip 192.168.1.1 172.16.1.1

   A. Subnet masks of 255.255.255.0 need to be added.
   B. Wildcard masks of 0.0.0.255 need to be added.
   C. The keyword 'host' needs to be added before each IP Address.
   D. The ACL is already correct. It needs to be applied on the appropriate interface correctly.

5. You see the following config on a router. What would you do to fix the working of the NAT?

interface GigabitEthernet 0/0
ip address 192.168.1.1 255.255.255.0
ip nat inside

interface GigabitEthernet 0/1
ip address 209.165.200.225 255.255.255.224
ip nat inside

ip nat inside source static 192.168.1.2 interface GigabitEthernet 0/1

A. Interface GigabitEthernet 0/1 needs to be changed from **ip nat inside** to **ip nat outside**
B. Interface GigabitEthernet 0/0 needs to be changed from **ip nat outside** to **ip nat inside**
C. The NAT Statement needs to be changed so that it uses an ACL to identify the source IP Address that will be translated.
D. The NAT Statement needs to be changed so that it is **ip nat outside source static 192.168.1.2 interface gigabitethernet 0/1**

6. Inspect the following configuration. Which of the following in correct statement?

RouterA# Show cdp neighbors
Capability Codes: R - Router, T - TransBridge, B - Source Route Bridge
S - Switch, H - Host, I - IGMP, r - repeater

Device ID	Local Interface	Holdtime	Capability	Platform
SwitchA	f0/0	122	S I	WS-2960-fa0/2
RouterB	s0/0/0	177	R S I	2811

RouterA# show run interface f0/0
interface FastEthernet 0/0
ip access-group 10 in

RouterA# show run interface f0/1
interface FastEthernet 0/1
ip address 192.168.1.1 255.255.255.0

RouterA# show run interface S0/0/0
interface Serial 0/0/0
ip address 172.16.1.1 255.255.255.0

RouterA# show access-list
access-list 10 permit any

A. From the CDP table, it is clear from the 'platform' information that the device with hostname SwitchA is actually a router.
B. The connection to SwitchA is incorrectly appearing on the CDP neighbor table.
C. Access-list 10 is preventing any connectivity to SwitchA
D. The connection to SwitchA is through F0/0. However it is F0/1 that is configured with an IP Address and not F0/0.

7. Inspect the config below and explain why the interface on the router is not able to receive a DHCP assigned IP Address from the neighbor router.

RouterA# show run interface Ethernet 0/0
Ethernet 0/0
no ip address

A. The interface already has an IP Address statically assigned.
B. The interface is probably going to get an IP Address from a dynamic routing protocol
C. The interface has a **no ip address** command configured. It needs the **ip address dhcp** command instead.
D. The interface is shut down. It needs to be issued with the **no shutdown** command.

Answer Key

Challenge

1. D
2. A
3. C
4. C
5. A
6. D
7. C

Module 4: Implementing Scalable Medium-Sized Networks

Introduction

When you understand how a switch and router operate, and how they communicate, you can move on to understanding an expanded network. This module shows how to "virtualize" your LAN using VLANs and how to configure Layer 3 connectivity between these VLANs. Then it describes how to decrease the administrative burden of assigning IP addresses by using DHCP. You will also learn how to configure and troubleshoot RIPv2.

Lesson 1: Implementing and Troubleshooting VLANs and Trunks

Introduction

Your boss sends you to your customer to add a VLAN into their network for their IT department. You need to understand the common issues in a poorly designed local network, such as large broadcast, failure domain, limited security control, and so on. You need to understand the VLAN operation along with trunk encapsulation.

Before going to the customer premise, you will need to design VLANs, IP addressing, VLANs for special traffic types, and VLAN security practice. You will also explain the configuration steps to the customer IT department and inform them about the role of the DTP and VTP.

Enterprise Network Design

Each layer—access, distribution, and core—provides different functionality and capability to the network.

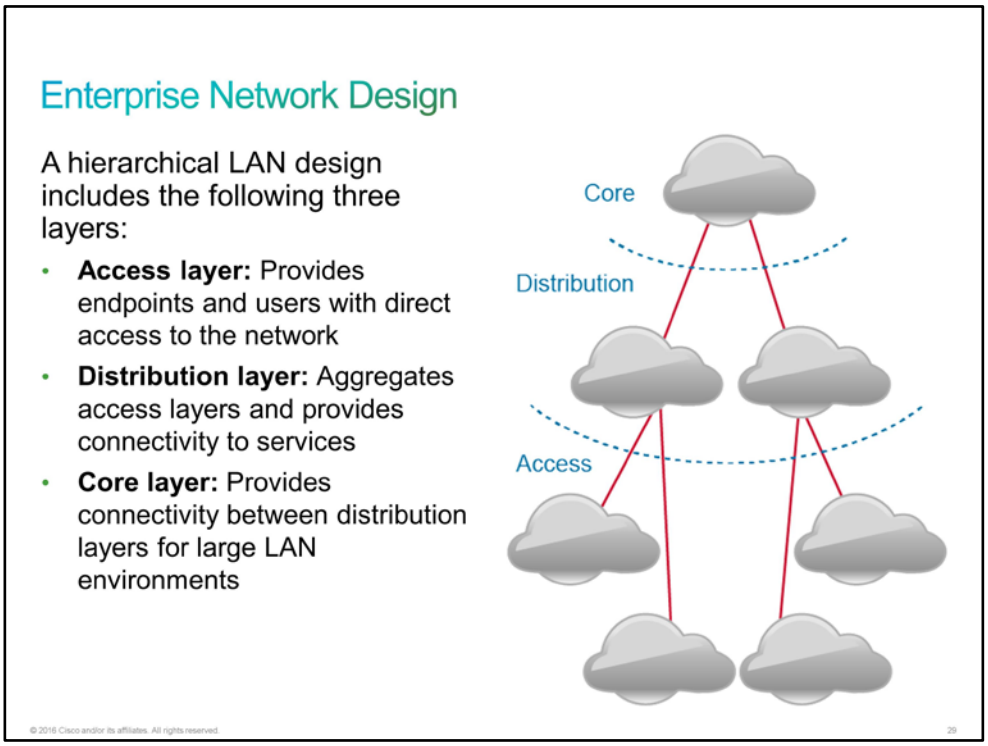

Depending on the characteristics of the deployment site, you might need one, two, or all three layers. For example, a site that occupies a single building might only require the access and distribution layers, while a campus of multiple buildings will most likely require all three layers.

Access Layer

The access layer is where user-controlled devices, user-accessible devices, and other endpoint devices are connected to the network. The access layer provides both wired and wireless connectivity and contains features and services that ensure security and resiliency for the entire network.

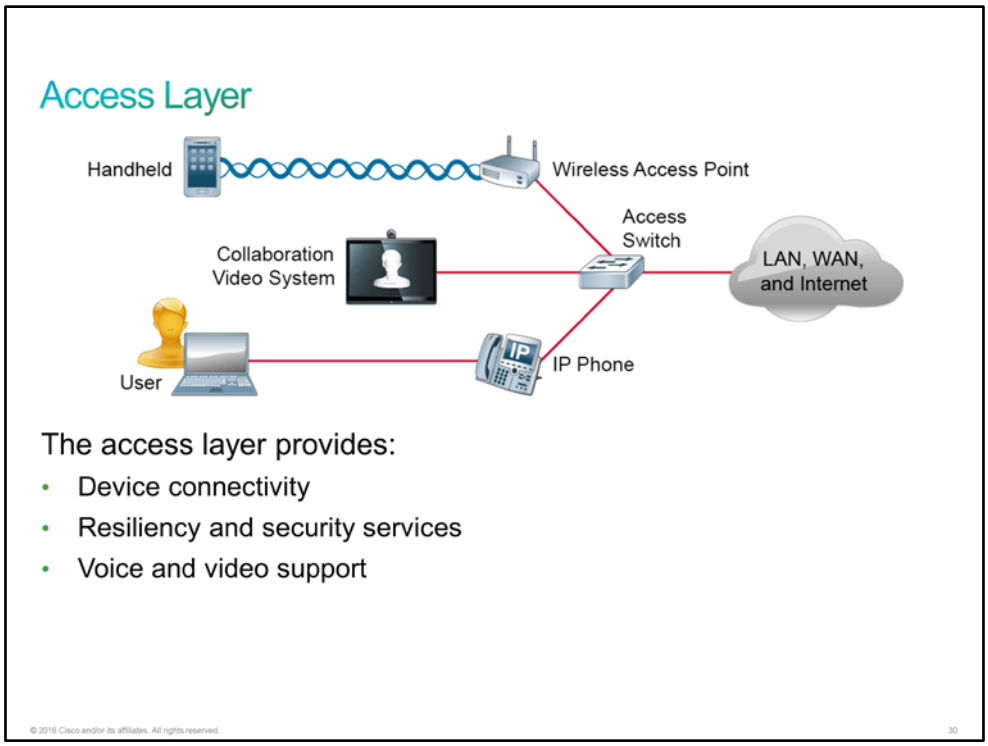

- **Device connectivity:** The access layer provides high-bandwidth device connectivity. To help make the network a transparent part of an end-user day-to-day job, the access layer must support bursts of high-bandwidth traffic when users perform routine tasks. Common routine tasks include sending large emails or opening a file from an internal web page. Because many types of end-user devices connect at the access layer—personal computers, IP phones, wireless access points, and IP video surveillance cameras—the access layer can support many logical networks, delivering benefits for performance, management, and security.
- **Resiliency and security services:** The access layer design must ensure that the network is available for all users who need it, whenever they need it. As the connection point between the network and client devices, the access layer must help protect the network from human errors and from malicious attacks. This protection includes ensuring that users have access only to authorized services and preventing end-user devices from taking over the role of other devices on the network. When possible, this protection mechanism should also verify that each end-user device is allowed on the network.
- **Advanced technology capabilities:** The access layer provides a set of network services that support advanced technologies, such as voice and video. The access layer must provide specialized access for devices using advanced technologies to ensure that traffic from these devices is not impaired by traffic from other devices. The access layer must also ensure efficient delivery of traffic that many devices in the network need.

Distribution Layer

The distribution layer supports many important services. In a network where connectivity needs to traverse the LAN end-to-end, whether between different access layer devices or from an access layer device to the WAN, the distribution layer facilitates this connectivity.

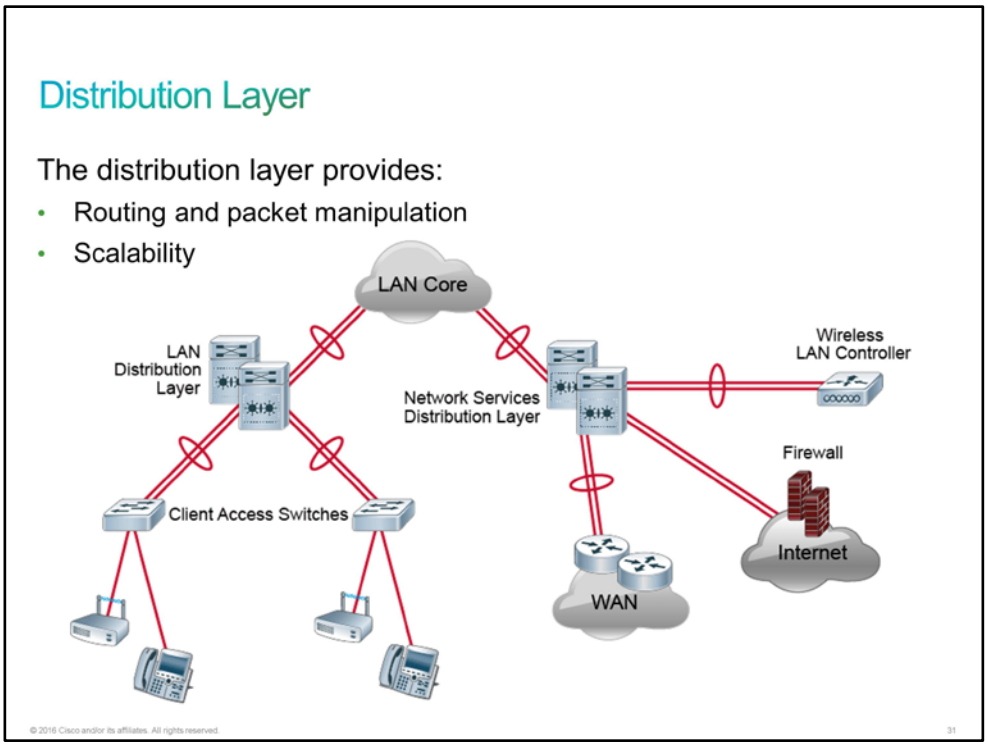

- **Routing and packet manipulation:** The distribution layer is the layer that provides policy-based connectivity. In terms of IP routing, the distribution layer represents a redistribution point between routing domains or the demarcation between the static and dynamic routing protocols. The distribution layer can also be the point at which tasks such as controlled routing decisions and filtering occur.
- **Scalability:** At any site with more than two or three access layer devices, it is impractical to interconnect all access switches. The distribution layer serves as an aggregation point for multiple access layer switches. The distribution layer can lower operating costs by making the network more efficient. Efficiency can be accomplished by requiring less memory, creating fault domains that compartmentalize failures or network changes, and by processing resources for devices elsewhere in the network. The distribution layer also increases network availability by containing failures to smaller domains.

Core Layer

In a large LAN environment, you often need to have multiple distribution layer switches. One reason for this is that when access layer switches are located in multiple geographically dispersed buildings, you can save potentially costly fiber-optic runs between buildings by locating a distribution layer switch in each of those buildings. As networks grow beyond three distribution layers in a single location, you should use a core layer to optimize the design.

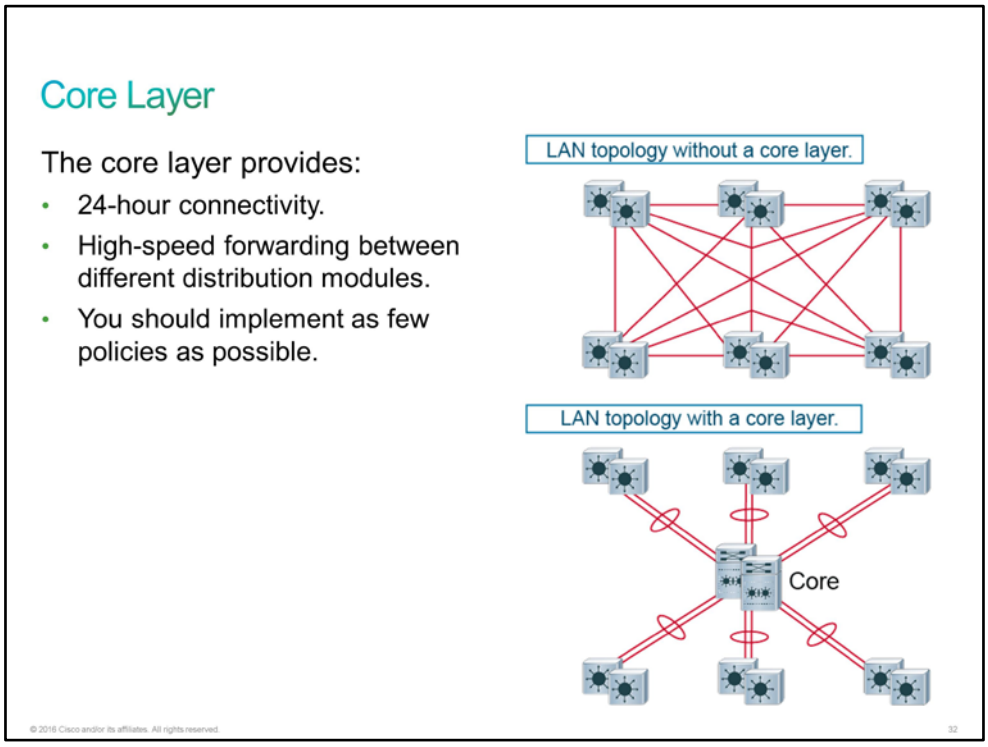

Another reason to use multiple distribution layer switches is when the number of access layer switches connecting to a single distribution layer exceeds the performance limits of the distribution switches. In a modular and scalable design, you can co-locate distribution layers for the data center, WAN connectivity, or Internet edge services.

In environments where multiple distribution layer switches exist in proximity and where fiber optics provide the ability for high-bandwidth interconnect, a core layer reduces the network as the example shows.

The core layer of the LAN is a critical part of a scalable network, and yet it is one of the simplest by design. The distribution layer provides fault domains, and the core represents 24-hour connectivity between them, which organizations must have in the modern business environment where connectivity to resources is critical.

Issues in a Poorly Designed Network

A poorly designed network has increased support costs, reduced service availability, and limited support for new applications and solutions. A less-than-optimal performance directly affects end users and their access to central resources.

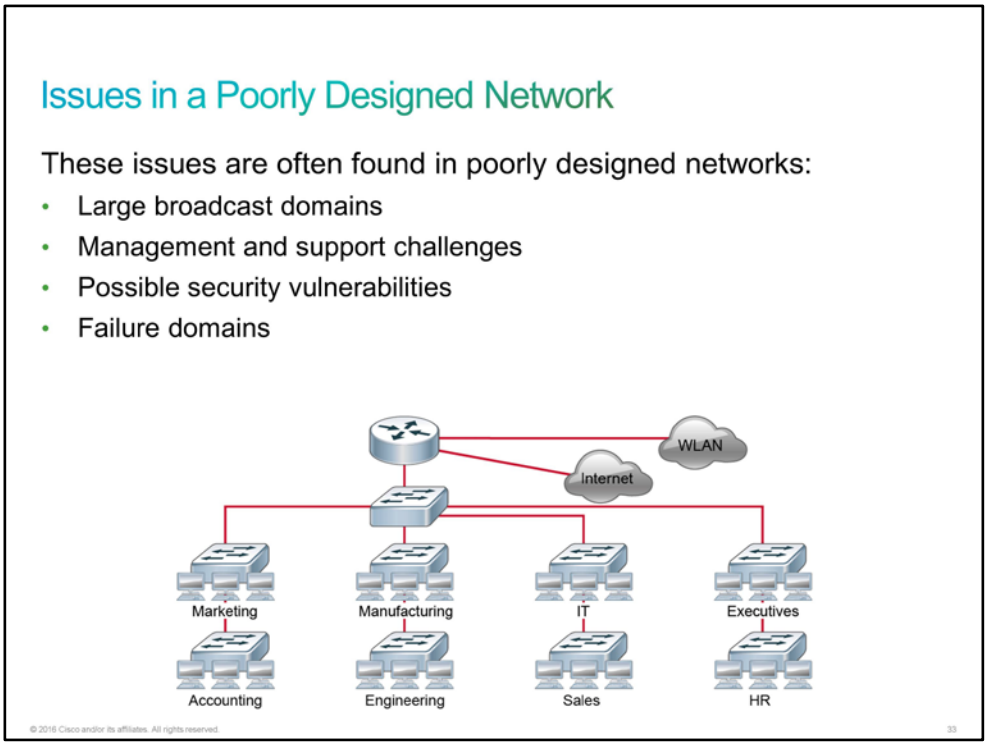

- **Large broadcast domains:** Broadcasts exist in every network. Many applications and network operations use broadcasts to function properly. Therefore, you cannot eliminate them completely. In the same way that avoiding failure domains involves clearly defining boundaries, broadcast domains should also have clear boundaries. They should also include an optimal number of devices to minimize the negative effect of broadcasts.
- **Management and support difficulties:** A poorly designed network may be disorganized, poorly documented, and lack easily identified traffic flows. These issues can make support, maintenance, and problem resolution time-consuming and difficult.
- **Possible security vulnerabilities:** A switched network that has been designed with little attention to security requirements at the access layer can compromise the integrity of the entire network.
- **Failure domains:** One of the reasons to implement an effective network design is to minimize the extent of problems when they occur. When you don't clearly define Layer 2 and Layer 3 boundaries, a failure in one network area can have a far-reaching effect.

A poorly designed network always has a negative effect. It becomes a support burden and a cost burden for any organization.

VLAN Introduction

To understand VLANs, it is important that you have a solid understanding of LANs. A LAN is a group of devices that share a common broadcast domain. When a device on the LAN sends broadcast messages, all the other devices on the LAN receive them. You can think of a LAN and a broadcast domain as being basically the same thing. Without VLANs, a switch considers all its interfaces to be in the same broadcast domain. In others words, all connected devices are in the same LAN. With VLANs, a switch can put some interfaces into one broadcast domain and some into another. The individual broadcast domains that are created by the switch are called *virtual LANs*, or VLANs.

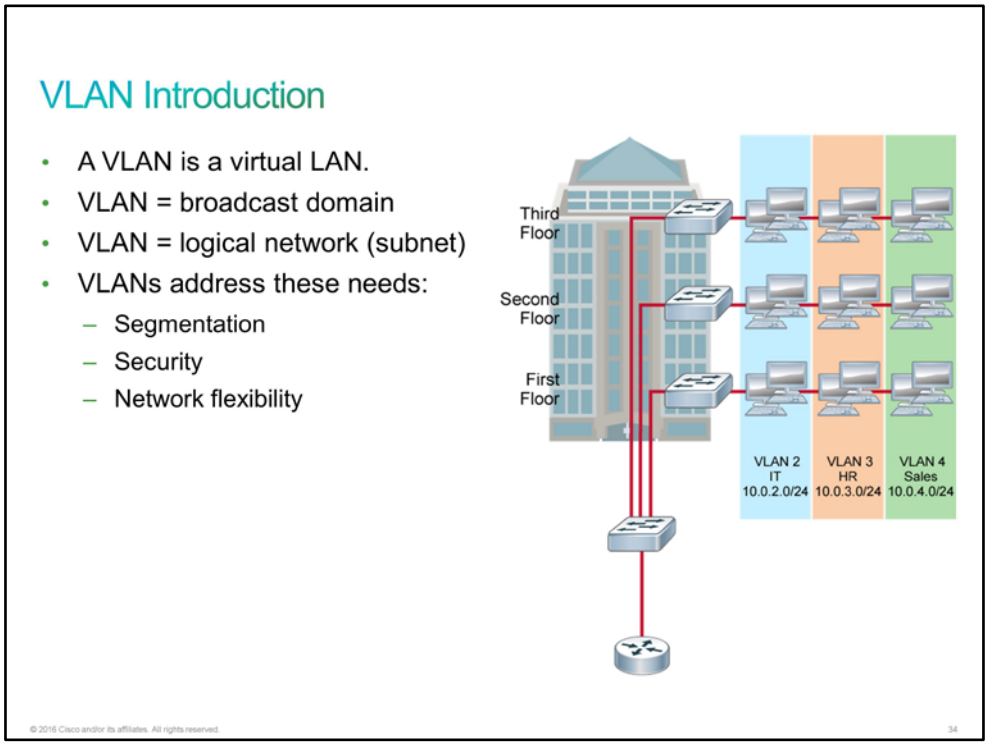

VLANs improve network performance by separating large broadcast domains into smaller segments. A VLAN allows a network administrator to create logical groups of network devices. These devices act like they are in their own independent network, even if they share a common infrastructure with other VLANs. A VLAN is a logical broadcast domain that can span multiple physical LAN segments. Within the switched internetwork, VLANs provide segmentation and organizational flexibility. You can design a VLAN structure that lets you group stations that are segmented logically by functions, project teams, and applications, without regard to the physical location of the users. VLANs allow you to implement access and security policies to particular groups of users. You can assign each switch port to only one VLAN, which adds a layer of security (if the port is operating as an access port). Ports in the same VLAN share broadcasts. Ports in different VLANs do not share broadcasts. Containing broadcasts within a VLAN improves the overall performance of the network.

A VLAN can exist on a single switch or span multiple switches. VLANs can include stations in a single building or multiple buildings. VLANs can also connect across WANs. The process of forwarding network traffic from one VLAN to another VLAN using a router is called inter-VLAN routing. VLANs are associated with unique IP subnets on the network. This subnet configuration facilitates the routing process in a multi-VLAN environment. When you are using a router to facilitate inter-VLAN routing, you can connect the router interfaces to separate VLANs. The devices on those VLANs send traffic through the router to reach other VLANs.

Usually, the subnets are chosen according to which VLANs they are associated with. The figure shows that VLAN 2 uses the subnet 10.0.2.0/24, VLAN 3 uses 10.0.3.0/24, and VLAN 4 uses 10.0.4.0/24. In this example, the third octet clearly identifies the VLAN that the device belongs to.

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

Creating a VLAN

For many Cisco Catalyst switches, you can use the **vlan** global configuration command to create a VLAN and enter the VLAN configuration mode. Use the **no** form of this command to delete the VLAN. The example shows how to add VLAN 2 to the VLAN database and how to name it "Sales."

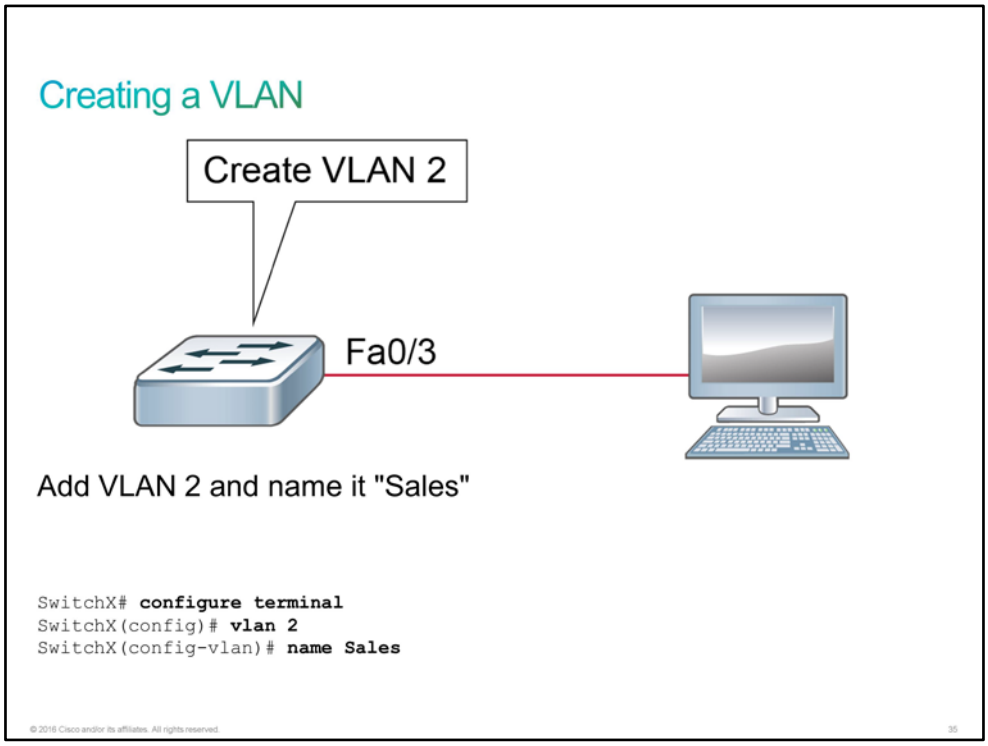

The table lists the commands to use when adding a VLAN.

Command and Variable	Description
vlan *vlan-id*	The ID of the VLAN that you want to add and configure. Do not enter leading zeros. You can enter a single VID, a series of VIDs that are separated by commas, or a range of VIDs that are separated by hyphens.
name *vlan-name*	(Optional) Specifies the VLAN name, which is an ASCII string from 1 to 32 characters that must be unique within the administrative domain.

To add a VLAN to the VLAN database, assign a number and name to the VLAN. VLAN 1 is the factory default VLAN. Normal-range VLANs are identified with a number between 1 and 1001. The VLAN numbers 1002 through 1005 are reserved for Token Ring and FDDI VLANs. VIDs 1 and 1002 to 1005 are automatically created, and you cannot remove them.

The configurations for VIDs 1 to 1005 are written to the vlan.dat file (VLAN database). You can display the VLANs by entering the **show vlan** privileged EXEC command. The *vlan.dat* file is stored in flash memory.

To add an Ethernet VLAN, you must specify at least a VLAN number. If you don't enter a name for the VLAN, the default is to append the VLAN number to the **vlan** command. For example, VLAN0004 would be the default name for VLAN 4 if you don't specify a name.

For more details about the **vlan** (VLAN configuration mode) command, see the *Cisco IOS LAN Switching Command Reference* at http://www.cisco.com/c/en/us/td/docs/ios/lanswitch/command/reference/lsw_book.html.

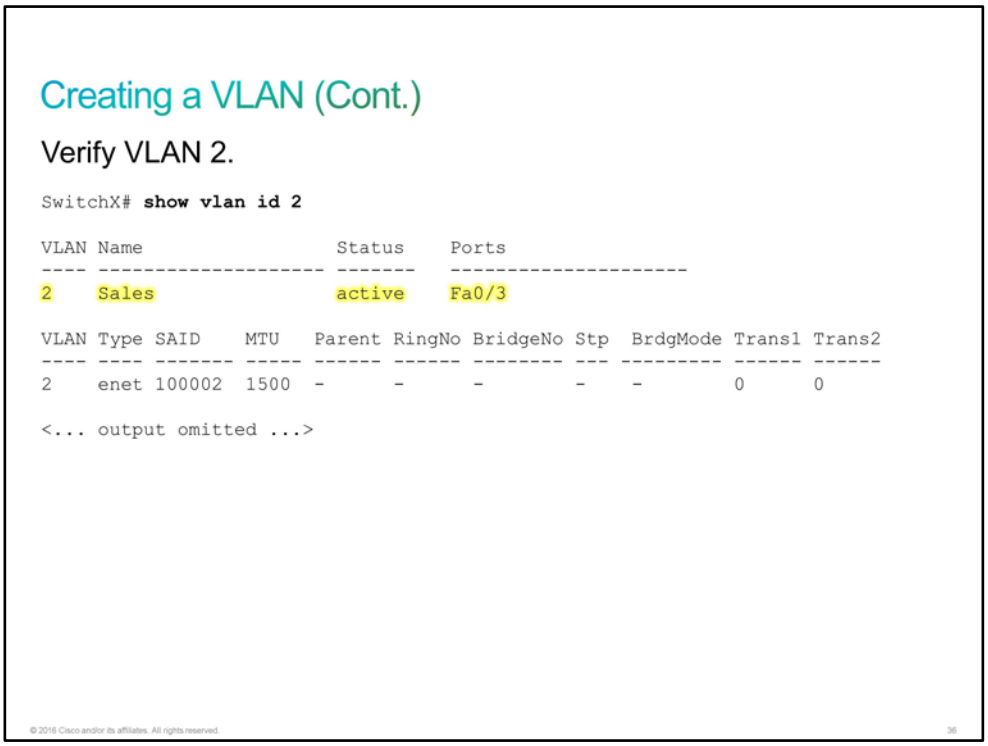

After you configure the VLAN, validate the parameters for this VLAN.

Use the **show vlan id** *vlan_number* or the **show vlan name** *vlan-name* command to display information about a particular VLAN. The figure shows an example of using the **show vlan** command to display the contents of the vlan.dat file. The "Sales" VLAN, which is VLAN 2, is highlighted in the example.

Use the **show vlan** command to display information on all configured VLANs. The **show vlan** command displays the switch ports that are assigned to each VLAN.

For more details about the **show vlan** command, see the *Cisco IOS LAN Switching Command Reference* at http://www.cisco.com/c/en/us/td/docs/ios/lanswitch/command/reference/lsw_book.html.

Assigning a Port to a VLAN

When you connect an end system to a switch port, you should associate it with a VLAN, in accordance with the network design. To associate a device with a VLAN, assign the switch port to which the device connects to a single-data VLAN. The switch port, therefore, becomes an access port.

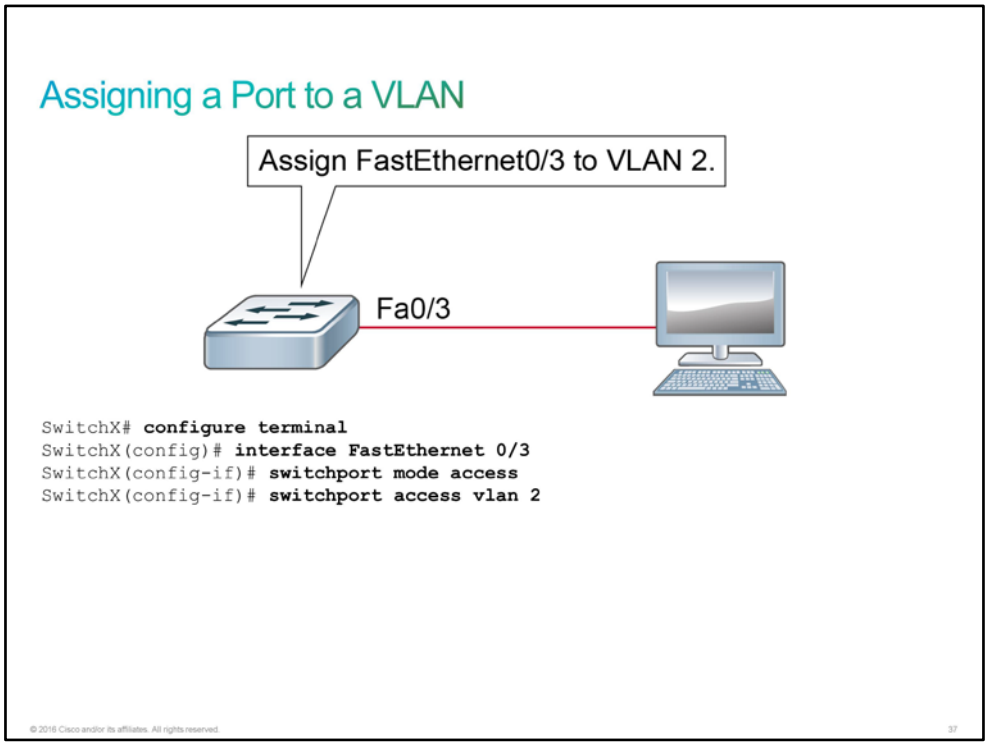

The table lists the commands to use when assigning a port to a VLAN.

Command and Variable	Description
interface *interface*	Enters the interface configuration mode
switchport access vlan *vlan_number*	Sets a nontrunking, untagged, single VLAN Layer 2 interface

After creating a VLAN, you can manually assign a port or many ports to this VLAN. A port can belong to only one VLAN at a time.

Note By default, all ports are members of VLAN 1.

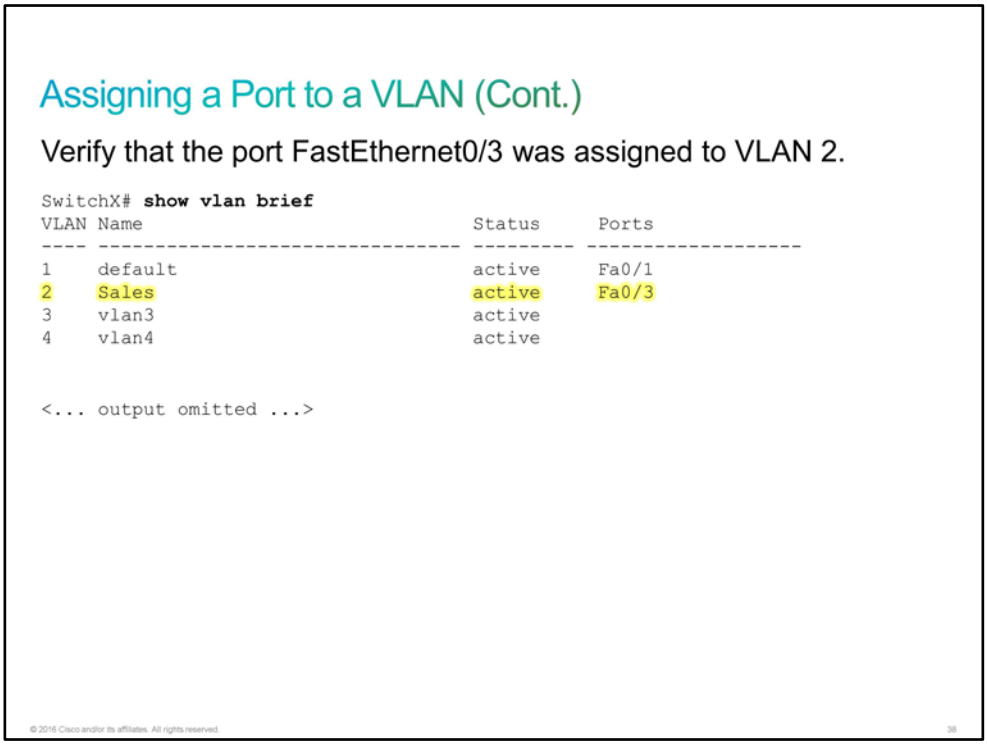

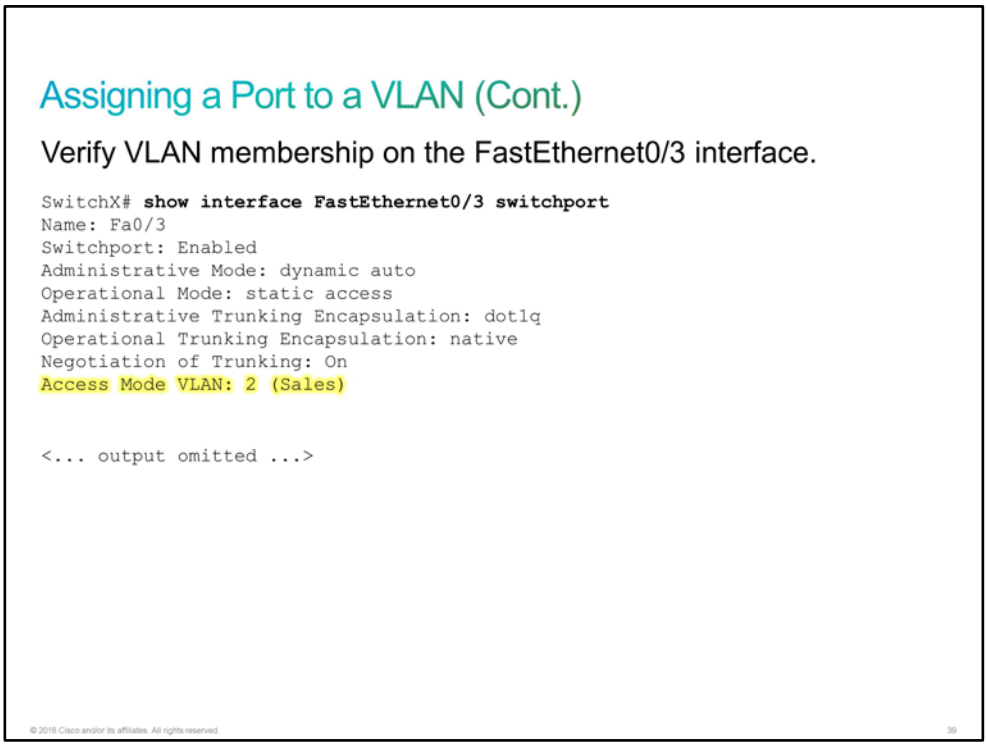

Use the **show vlan** privileged EXEC command to display the VLAN assignment and membership type for all switch ports. The **show vlan** command displays one line for each VLAN. The output for each VLAN includes the VLAN name, status, and switch ports.

For more details about the **show vlan** command, see the *Cisco IOS LAN Switching Command Reference* at http://www.cisco.com/c/en/us/td/docs/ios/lanswitch/command/reference/lsw_book.html.

Alternatively, use the **show interfaces switchport** privileged EXEC command to display the VLAN information for a particular interface. The output in the example shows the information about the FastEthernet0/3 interface, where VLAN 2, which is named "Sales," is assigned.

For more details about the **show interfaces switchport** command, see the *Cisco IOS LAN Switching Command Reference* at http://www.cisco.com/c/en/us/td/docs/ios/lanswitch/command/reference/lsw_book.html.

Trunking with 802.1Q

Running many VLANs between switches would require the same number of interconnecting links.

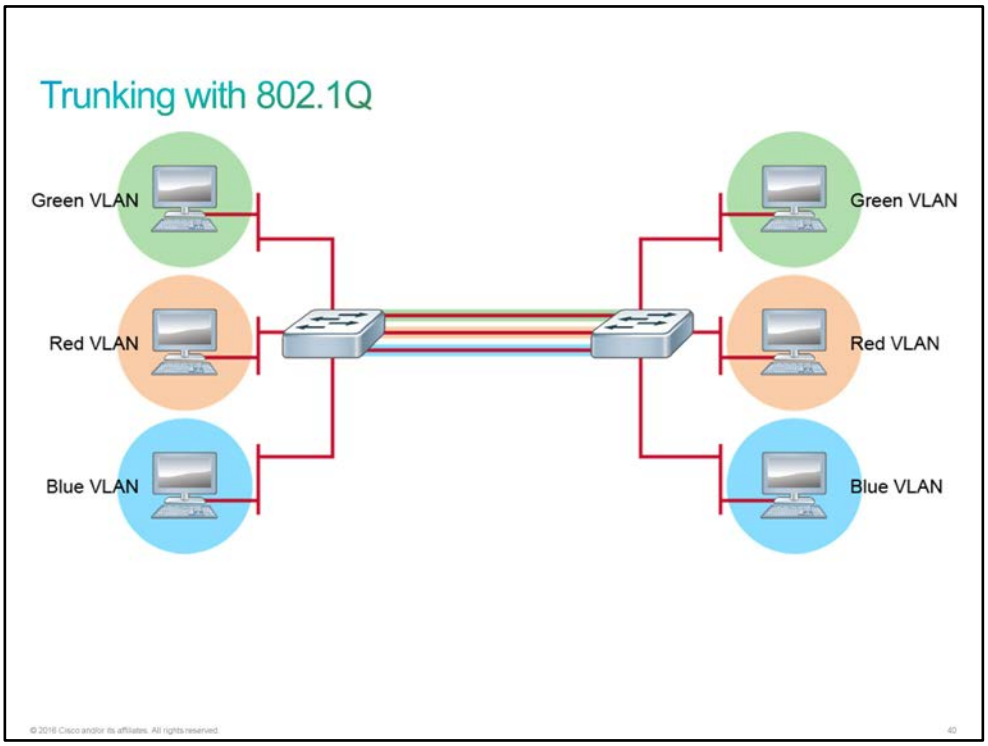

If every port belongs to one VLAN and you have several VLANs that are configured on switches, interconnecting these VLANs would require one physical cable per VLAN. When the number of VLANs increases, so does the number of required interconnecting links. Ports are then used for interswitch connectivity instead of attaching end devices.

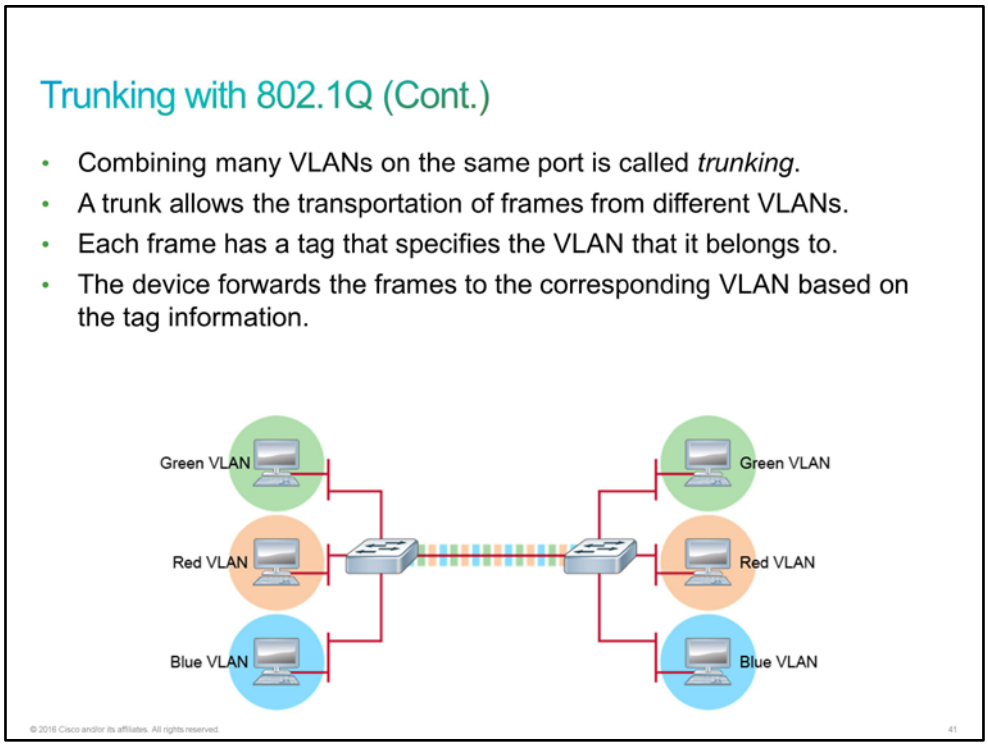

VLAN Tagging

For VLANs in networks that have multiple interconnected switches, the switches must use VLAN trunking on the segments between the switches. VLAN trunking causes the switches to use a process that is called *VLAN tagging*, so that the sending switch adds another header to the frame before sending it over the trunk. This extra VLAN header includes a VID (VID field) so that the sending switch can list the VLAN ID and the receiving switch can identify the VLAN that each frame belongs to. The figure shows the basic idea.

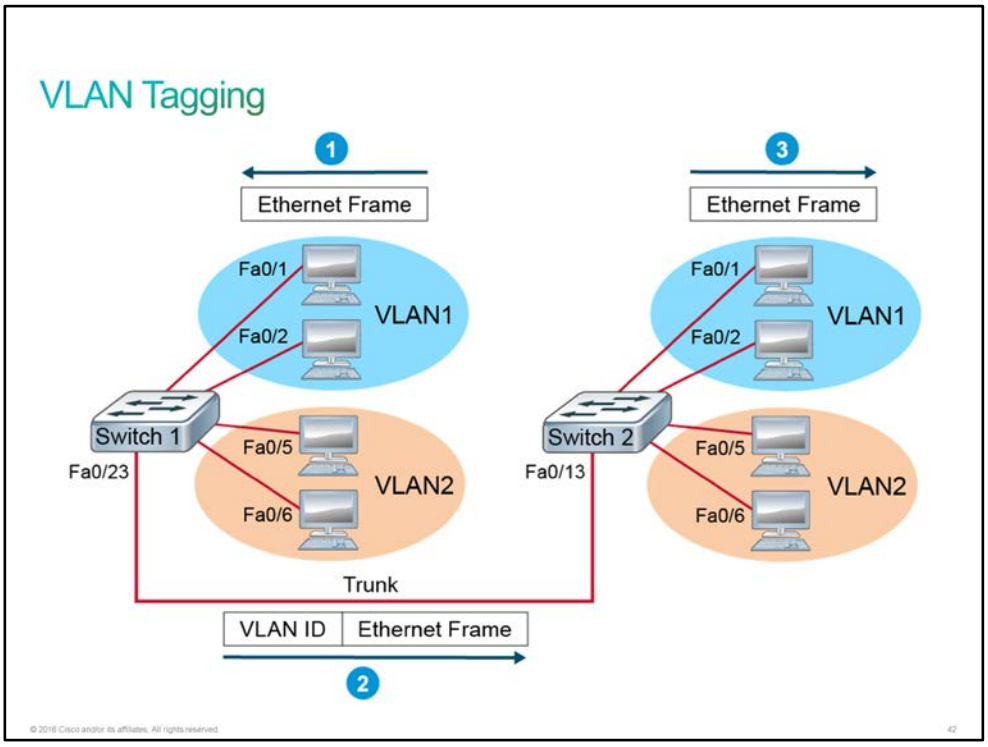

Trunking allows switches to pass frames from multiple VLANs over a single physical connection. For example, the figure shows Switch 1 receiving a broadcast frame on the interface Fa0/1, which is a member of VLAN1. In a broadcast, the frame must be forwarded to all ports in VLAN1. Because there are ports on Switch 2 that are members of the VLAN1 switch, the frame must be forwarded to Switch 2. Before forwarding the frame, Switch 1 adds a header that identifies the frame as belonging to VLAN1. This header tells Switch 2 that the frame should be forwarded to the VLAN1 ports. Switch 2 removes the header and then forwards the frame for all ports that are part of VLAN1.

As another example, the device on the Switch 1 Fa0/5 interface sends a broadcast. Switch 1 sends the broadcast out of port Fa0/6 (because that port is in VLAN 2) and out Fa0/23 (because it is a trunk, meaning that it supports multiple VLANs). Switch 1 adds a trunking header to the frame, listing a VLAN ID of 2. Switch 2 strips off the trunking header because the frame is part of VLAN 2, so Switch 2 knows to forward the frame out of only ports Fa0/5 and Fa0/6, and not ports Fa0/1 and Fa0/2.

IEEE 802.1Q

The 802.1Q inserts an extra 4-byte VLAN header into the Ethernet header of the original frame. As a result, the frame still has the original source and destination MAC addresses. Also, because the original header has been expanded, 802.1Q encapsulation forces a recalculation of the original FCS field in the Ethernet trailer, because the FCS is based on the content of the entire frame. The figure shows the 802.1Q header and framing of the revised Ethernet header.

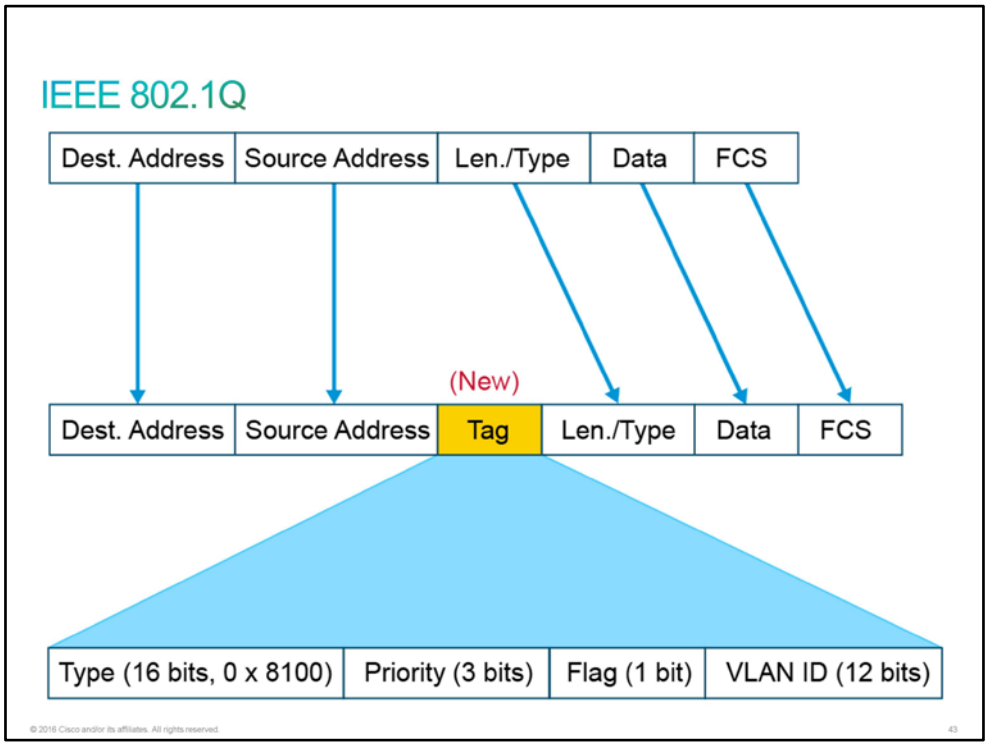

These are Tag fields:

- **Type** or tag protocol identifier is set to a value of 0x8100 in order to identify the frame as an IEEE 802.1Q-tagged frame
- **Priority** indicates the frame priority level which can be used for the prioritization of traffic
- If **flag** is 1, the MAC address is in noncanonical format. If flag is 0, the MAC address is in canonical format.
- **VLAN ID** uniquely identifies the VLAN to which the frame belongs

Configuring an 802.1Q Trunk

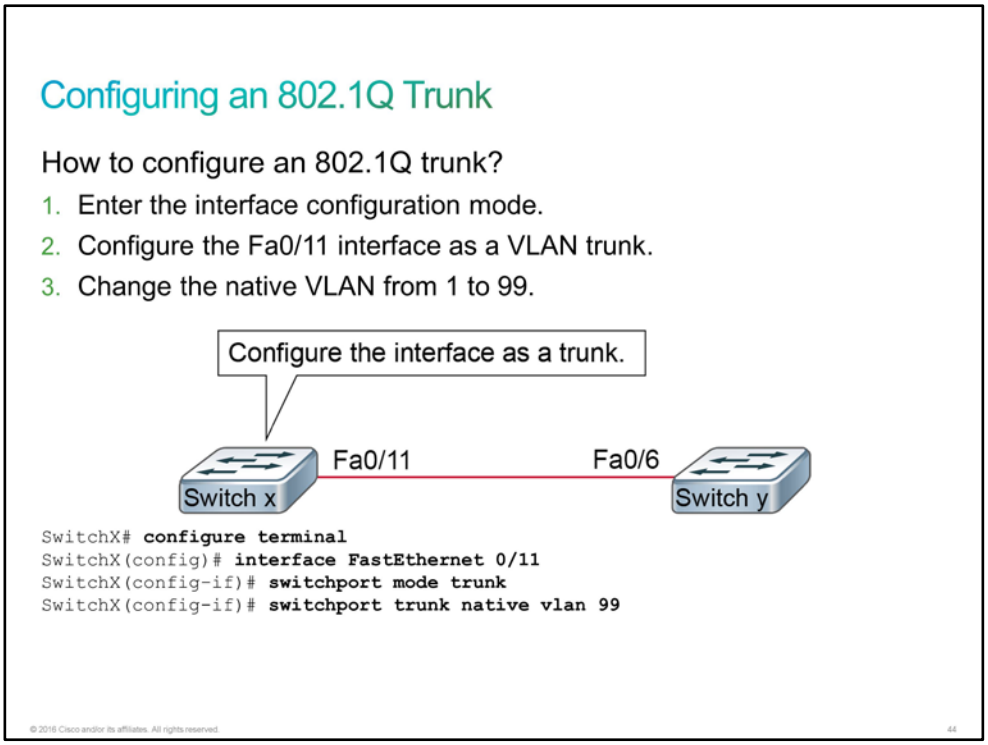

Command and Variable	Description
interface *interface*	Enters interface configuration mode.
switchport mode trunk	Sets the interface type. The keyword **trunk** specifies a trunking VLAN Layer 2 interface.
switchport trunk native vlan *vlan_number*	Sets the native VLAN for the trunk in the 802.1Q trunking mode.

The example configures the FastEthernet0/11 port on Switch X as a trunk port. Use the **switchport mode** interface configuration command to set a Fast Ethernet port to trunk mode. Many Cisco Catalyst switches support DTP, which manages automatic trunk negotiation. DTP is a Cisco proprietary protocol. Switches from other vendors do not support DTP. DTP is automatically enabled on a switch port when certain trunking modes are configured on the switch port. DTP manages trunk negotiation only if the port on the other switch is configured in a trunk mode that supports DTP.

The example shows the configuration of interface FastEthernet0/11. The **switchport trunk mode** command sets the FastEthernet0/11 port to the trunk mode. The example shows the reconfiguration of the native VLAN. VLAN 99 is configured as the native VLAN. Therefore, the device will send the traffic from VLAN 99 untagged.

Make sure that the other end of the trunk link (Switch Y) is configured for trunking and with the native VLAN that is changed to 99.

Note	For details on all the parameters that are associated with the **switchport mode** interface command, visithttp://www.cisco.com/en/US/docs/ios/mcl/allreleasemcl/all_15.htmlhttp://www.cisco.com/en/US/docs/ios/mcl/allreleasemcl/all_15.html.

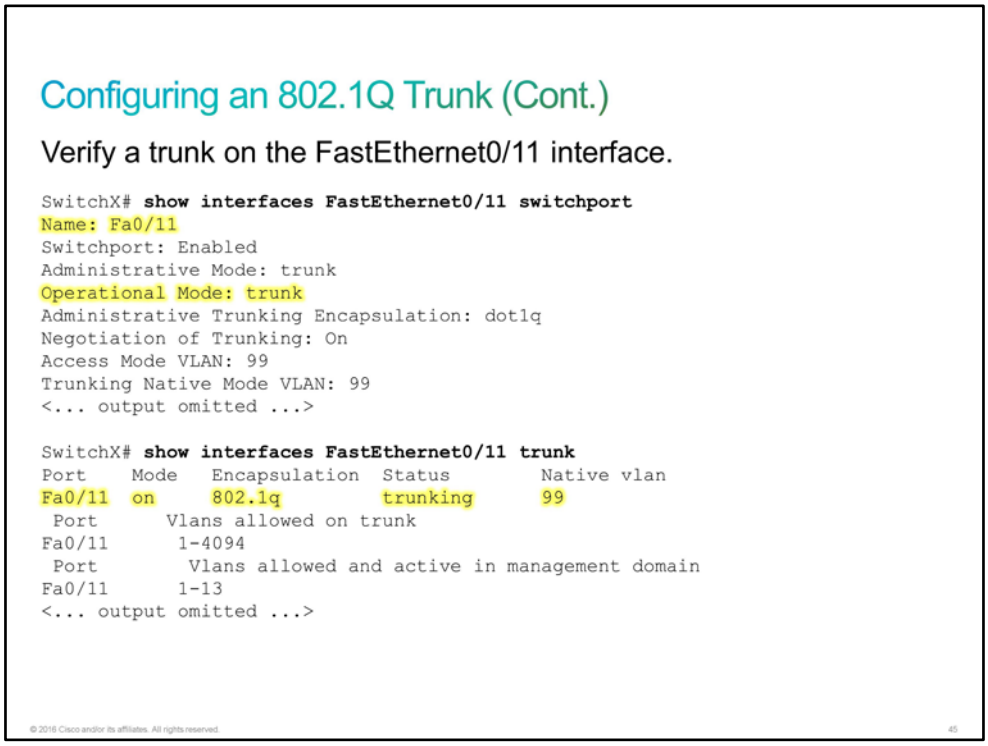

To verify a trunk configuration on a switch, use the **show interfaces switchport** and **show interfaces trunk** commands. These two commands display the trunk parameters and VLAN information of the port.

For more details about the **show interfaces switchport** and **show interfaces trunk** commands, see the *Cisco IOS Interface and Hardware Component Command Reference* at http://www.cisco.com/en/US/docs/ios-xml/ios/interface/command/ir-s5.html.

Discovery 16: Configure VLAN and Trunk

Introduction

This discovery lab will guide you through several aspects of VLAN operations, including the management of VLANs, and using trunks to carry multiple VLANs across a single physical link. The devices are configured as pictured in the topology diagram. Currently, all devices have IP addresses in the 10.10.1.0/24 subnet. Only the default VLAN, VLAN 1, exists initially. You will start by migrating this configuration to one that uses two VLANs.

Topology

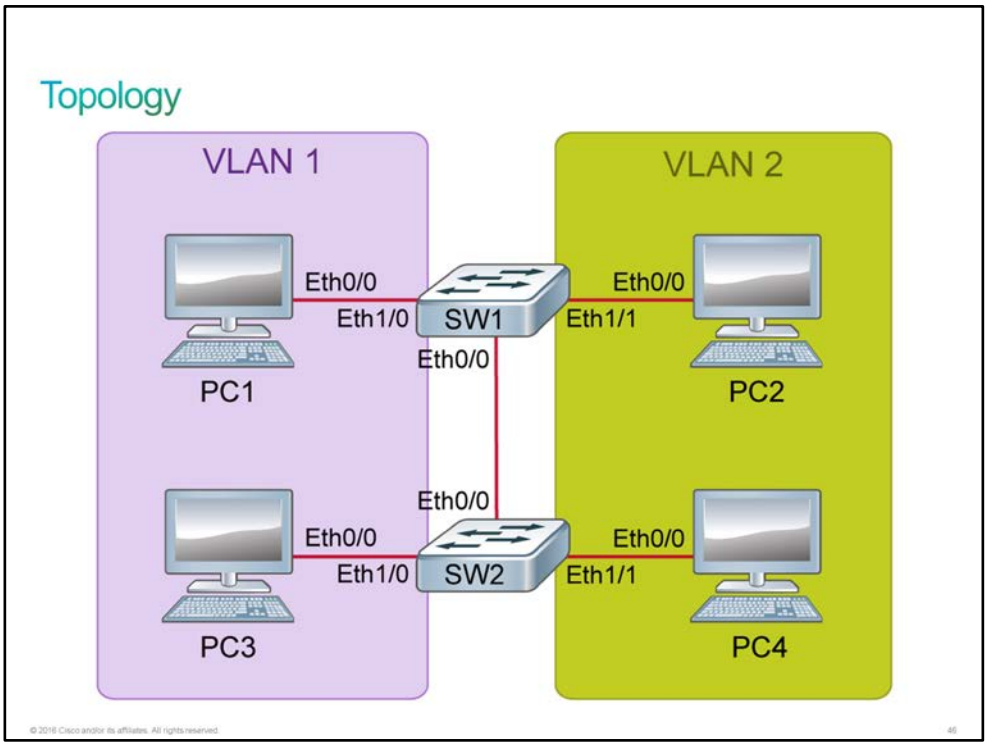

Job Aid

Device Information

Device Information Table

Device	Characteristic	Value
PC1	Hostname	PC1
PC1	IP address	10.10.1.10/24
PC2	Hostname	PC2
PC2	IP address	10.10.1.20/24

Device	Characteristic	Value
PC3	Hostname	PC3
PC3	IP address	10.10.1.30/24
PC4	Hostname	PC4
PC4	IP address	10.10.1.40/24
SW1	Hostname	SW1
SW1	VLAN 1 IP Address	10.10.1.4/24
SW1	Ethernet0/0 description	Link to SW2
SW1	Ethernet1/0 description	Link to PC1
SW1	Ethernet1/1 description	Link to PC2
SW2	Hostname	SW2
SW2	VLAN 1 IP address	10.10.1.5/24
SW2	Ethernet0/0 description	Link to SW1
SW2	Ethernet1/0 description	Link to PC3
SW2	Ethernet1/1 description	Link to PC4

PCs in the virtual lab environment are simulated as routers, so you should use Cisco IOS commands to configure them or make verifications.

Device Information Table (Changes)

Device	Characteristic	Value
PC2	VLAN	2
PC2	IP address	10.10.2.20/24
PC4	VLAN	2
PC4	IP address	10.10.2.40/24

Task 1: Configure VLAN and Trunk

Activity

Step 1 Start by demonstrating that there is full connectivity between the devices in VLAN 1 on the 10.10.1.0/24 subnet. Access the console of PC1 and ping the IP addresses of the other devices.

Enter following commands to PC1:

```
PC1# ping 10.10.1.20
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.10.1.20, timeout is 2 seconds:
.!!!!
Success rate is 80 percent (4/5), round-trip min/avg/max = 1/1/1 ms
PC1# ping 10.10.1.30
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.10.1.30, timeout is 2 seconds:
.!!!!
Success rate is 80 percent (4/5), round-trip min/avg/max = 1/1/1 ms
PC1# ping 10.10.1.40
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.10.1.40, timeout is 2 seconds:
.!!!!
Success rate is 80 percent (4/5), round-trip min/avg/max = 1/1/1 ms
PC1# ping 10.10.1.4
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.10.1.4, timeout is 2 seconds:
.!!!!
Success rate is 80 percent (4/5), round-trip min/avg/max = 1/1/1 ms
PC1# ping 10.10.1.5
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.10.1.5, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
```

Step 2 Now access the console of PC2 and change its IP address to 10.10.2.20 on the 10.10.2.0/24 subnet.

The most commonly used commands are abbreviated in this guided discovery. For example, **conf t** is used for **configure terminal**. If there is any confusion, you can attempt tab completion to expand the full command syntax. For example, **conf<tab> t<tab>** would expand to **configure terminal**.

Enter the following commands to PC2:

```
PC2# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
PC2(config)# int e0/0
PC2(config-if)# ip address 10.10.2.20 255.255.255.0
PC2(config-if)# end
PC2#
```

At this point, PC2 is still in VLAN 1, so it is in the same broadcast domain as all the other hosts. But its IP address is configured for a different IP subnet. PC2 will not attempt ARP resolution for hosts on the 10.10.1.0/24 subnet. It must use a gateway to reach the 10.10.1.0/24 subnet; however, that gateway does not even exist. PC2 is currently isolated by the IP configuration.

Step 3 Access the console of PC4 and reconfigure its IP address to be 10.10.2.40 on the 10.10.2.0/24 subnet.

Enter the following commands to the PC4:

```
PC4# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
PC4(config)# int e0/0
PC4(config-if)# ip address 10.10.2.40 255.255.255.0
PC4(config-if)# end
PC4#
```

Now both PC2 and PC4 are configured for the 10.10.2.0/24 subnet, while the rest of the hosts are configured for the 10.10.1.0/24 subnets. They are all in the same broadcast domain (VLAN 1), but they are isolated by the IP configuration.

Step 4 Verify that PC4 can communicate with PC2, because they are both configured for the 10.10.2.0/24 subnet. Attempt to ping 10.10.2.20. The ping should succeed.

Enter the following commands to PC4:

```
PC4# ping 10.10.2.20
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.10.2.20, timeout is 2 seconds:
.!!!!
Success rate is 80 percent (4/5), round-trip min/avg/max = 1/1/1 ms
```

Step 5 Access the console of SW1 and verify that the only Ethernet VLAN is the default VLAN—VLAN 1.

Enter the following command to the SW1 switch:

```
SW1# show vlan

VLAN Name                              Status    Ports
---- --------------------------------- --------- -------------------------------
1    default                           active    Et0/1, Et0/2, Et0/3, Et1/0
                                                 Et1/1, Et1/2, Et1/3
1002 fddi-default                      act/unsup
1003 token-ring-default                act/unsup
1004 fddinet-default                   act/unsup
1005 trnet-default                     act/unsup

VLAN Type  SAID       MTU   Parent RingNo BridgeNo Stp  BrdgMode Trans1 Trans2
---- ----- ---------- ----- ------ ------ -------- ---- -------- ------ ------
1    enet  100001     1500  -      -      -        -    -        0      0
1002 fddi  101002     1500  -      -      -        -    -        0      0
1003 tr    101003     1500  -      -      -        -    -        0      0
1004 fdnet 101004     1500  -      -      -        ieee -        0      0
1005 trnet 101005     1500  -      -      -        ibm  -        0      0

Primary Secondary Type              Ports
------- --------- ----------------- ------------------------------------------
```

Besides VLAN 1, which is the default Ethernet VLAN, there are four other VLANs that exist by default. VLANs 1002 to 1005 exist to support the legacy Token Ring and FDDI technology. They are very rarely used in networks today.

Step 6 Create VLAN 2 and assign "Engineering" as its name.

Enter the following commands to the SW1 switch:

```
SW1# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
SW1(config)# vlan 2
SW1(config-vlan)# name Engineering
SW1(config-vlan)# end
SW1#
```

Step 7 Verify that the VLAN has been created and is active.

Enter the following command to the SW1 switch:

```
SW1#show vlan brief

VLAN Name                              Status    Ports
---- --------------------------------- --------- -------------------------------
1    default                           active    Et0/1, Et0/2, Et0/3, Et1/0
                                                 Et1/1, Et1/2, Et1/3
2    Engineering                       active
1002 fddi-default                      act/unsup
1003 token-ring-default                act/unsup
1004 fddinet-default                   act/unsup
1005 trnet-default                     act/unsup
```

You can compare the output of the **show vlan brief** command to the output of the **show vlan** command that was used previously. With the brief argument, the characteristics that are only appropriate to Token Ring and FDDI networks (such as parent and ring number) are hidden from the display.

Although VLAN 2 is active, no active ports appear to be using VLAN 2.

Step 8 Take a closer look at the status of VLAN 2 by specifying its ID with the **show vlan** command.

Enter the following command to the SW1 switch:

```
SW1# show vlan id 2

VLAN Name                              Status    Ports
---- --------------------------------- --------- -------------------------------
2    Engineering                       active    Et0/0

VLAN Type  SAID       MTU   Parent RingNo BridgeNo Stp  BrdgMode Trans1 Trans2
---- ----- ---------- ----- ------ ------ -------- ---- -------- ------ ------
2    enet  100002     1500  -      -      -        -    -        0      0

Primary Secondary Type              Ports
------- --------- ----------------- ---------------------------------------
```

When you show all VLANs, only the access mode ports are displayed. When you show a particular VLAN, the trunk ports that carry the VLAN are also displayed. Ethernet0/0 is the trunk port connecting SW1 and SW2.

Step 9 View the switch port status of the interface Ethernet0/0.

Enter the following command to the SW1 switch:


```
SW1# show int e0/0 switchport
Name: Et0/0
Switchport: Enabled
Administrative Mode: dynamic desirable
Operational Mode: trunk
Administrative Trunking Encapsulation: negotiate
Operational Trunking Encapsulation: isl
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
Capture Mode Disabled
Capture VLANs Allowed: ALL

Appliance trust: none
```

The default administrative trunking mode on a switch port is "dynamic desirable." If this default is maintained, the connection between two switches will automatically become an operational trunk.

The default administrative trunking mode varies between Switch models.

The SW1 default trunking encapsulation method is "ISL." This characteristic is model-dependent. ISL is an older, Cisco proprietary trunking protocol. 802.1Q is much more common in networks today, and some switch models no longer support ISL.

Step 10 While the trunking status was automatically negotiated between the switches, the best practice is to explicitly configure the trunking status on switch ports. Also, it is best practice to assign a native VLAN to 802.1Q trunks that is not used by any endpoint hosts on the network. Begin this explicit configuration by defining VLAN 256 and assigning it the "NoHosts" name.

Enter the following commands to the SW1 switch:

```
SW1# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
SW1(config)# vlan 256
SW1(config-vlan)# name NoHosts
SW1(config-vlan)# exit
SW1(config)#
```

While it makes no difference to the switch which IP subnet you implement on which VLAN, for ease of network management, it is common to use the value of the third octet of the IP network as the VLAN ID, when possible. For example, you would pair VLAN 1 with 10.10.1.0/24 and pair VLAN 2 with 10.10.2.0/24.

The number 256 is not a valid IPv4 octet. The X.Y.256.Z addresses are invalid IPv4 addresses. Therefore, 256 can be an effective VID to use for a VLAN that intentionally services no hosts and is used for the native VLAN on 802.1Q trunks.

Step 11 Now, explicitly configure Ethernet0/0 as 802.1Q trunks using the VLAN 256 as the native VLAN.

Enter the following commands to the SW1 switch:

```
SW1(config)# int eth0/0
SW1(config-if)# switchport trunk encapsulation dot1q
SW1(config-if)# switchport trunk native vlan 256
SW1(config-if)# switchport mode trunk
SW1(config-if)# end
SW1#
```

Before changing native VLAN on SW2, you will see on SW1 console the *%CDP-4-NATIVE_VLAN_MISMATCH* message every 60 seconds.

```
*Feb  2 12:34:09.712: %CDP-4-NATIVE_VLAN_MISMATCH: Native VLAN mismatch
discovered on Ethernet0/0 (256), with SW2 Ethernet0/0 (1).
```

Step 12 You must configure SW2 to be synchronized with the configuration operations that you just performed on SW1. Access the console of SW2, configure VLAN 2 and VLAN 256, and configure Ethernet0/0 explicitly as 802.1Q trunk, with VLAN 256 as the native VLAN.

Enter following commands to the SW2 switch:

```
SW2# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
SW2(config)# vlan 2
SW2(config-vlan)# name Engineering
SW2(config-vlan)# vlan 256
SW2(config-vlan)# name NoHosts
SW2(config-vlan)# exit
SW2(config)# int eth0/0
SW2(config-if)# switchport trunk encapsulation dot1q
SW2(config-if)# switchport trunk native vlan 256
SW2(config-if)# switchport mode trunk
SW2(config-if)# end
SW2#
```

Step 13 Verify the trunk status of Ethernet0/0 on SW2.

Enter the following commands to the SW2 switch:


```
SW2# show int e0/0 switchport
Name: Et0/0
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: 256 (NoHosts)
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
Capture Mode Disabled
Capture VLANs Allowed: ALL

Appliance trust: none
```

Both administrative and operational modes are 802.1Q trunk. Optionally, you may repeat this verification on SW1.

Step 14 VLAN 2 is now ready on both switches, and the trunk link is configured between the two switches. Explicitly define the PC4 switch port as an access port that is assigned to VLAN 2.

Enter the following commands to the SW2 switch:

```
SW2# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
SW2(config)# int e1/1
SW2(config-if)# switchport access vlan 2
SW2(config-if)# switchport mode access
SW2(config-if)# end
SW2#
```

Step 15 Verify the status of the Ethernet1/1 switch port configuration.

Enter the following command to the SW2 switch:

```
SW2# show int e1/1 switchport
Name: Et1/1
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: 2 (Engineering)
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
Capture Mode Disabled
Capture VLANs Allowed: ALL

Appliance trust: none
```

Step 16 Verify the interface status of the trunk link (Ethernet0/0) and the ports that are supporting PC3 and PC4.

Enter the following command to the SW2 switch:

```
SW2# sh int status

Port      Name               Status       Vlan       Duplex  Speed Type
Et0/0     Link to SW1        connected    trunk        auto   auto unknown
Et0/1                        connected    1            auto   auto unknown
Et0/2                        connected    1            auto   auto unknown
Et0/3                        connected    1            auto   auto unknown
Et1/0     Link to PC3        connected    1            auto   auto unknown
Et1/1     Link to PC4        connected    2            auto   auto unknown
Et1/2                        connected    1            auto   auto unknown
Et1/3                        connected    1            auto   auto unknown
```

Step 17 Now PC4 and PC2 are on different VLANs, so even though they are configured for the same IP subnet, they should no longer be able to communicate. Verify this status by attempting to ping 10.10.2.20 from PC4. This ping should fail.

Enter the following command to the PC4:

```
PC4# ping 10.10.2.20
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.10.2.20, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
```


Step 18 Access the console of SW1. Configure the PC2 switch port to be an access port that is assigned to VLAN 2.

Enter the following commands to the SW1 switch:

```
SW1# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
SW1(config)# int e 1/1
SW1(config-if)# switchport access vlan 2
SW1(config-if)# switchport mode access
SW1(config-if)# end
SW1#
```

Step 19 Verify the switch port status of Ethernet1/1.

Enter the following command to the SW1 switch:

```
SW1# show int e1/1 switchport
Name: Et1/1
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: 2 (Engineering)
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
Capture Mode Disabled
Capture VLANs Allowed: ALL

Appliance trust: none
```

Step 20 Verify the interface of the SW1 trunk ports and the links to the PCs.

Enter the following command to the SW1 switch:

```
SW1# show int status

Port      Name               Status       Vlan       Duplex  Speed Type
Et0/0     Link to SW2        connected    trunk        auto   auto unknown
Et0/1                        connected    1            auto   auto unknown
Et0/2                        connected    1            auto   auto unknown
Et0/3                        connected    1            auto   auto unknown
Et1/0     Link to PC1        connected    1            auto   auto unknown
Et1/1     Link to PC2        connected    2            auto   auto unknown
Et1/2                        connected    1            auto   auto unknown
Et1/3                        connected    1            auto   auto unknown
```

Step 21 PC2 and PC4 are now both configured for the 10.10.2.0/24 subnet and are in the same broadcast domain (VLAN 2). Access the console of PC4 and verify that it can once again ping PC2.

Enter the following command to the PC4:

```
PC4# ping 10.10.2.20
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.10.2.20, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
```

Step 22 At this point, there is no routing configured. PC2 and PC4 are isolated from the other hosts that are in VLAN 1. Demonstrate that PC4 cannot ping PC1.

Enter the following command to the PC4:

```
PC4# ping 10.10.1.10
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.10.1.10, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
```

This is the end of the discovery lab.

Dynamic Trunking Protocol

Many Cisco Catalyst switches support DTP, which manages automatic trunk negotiation. DTP is a Cisco proprietary protocol. Switches from other vendors do not support DTP.

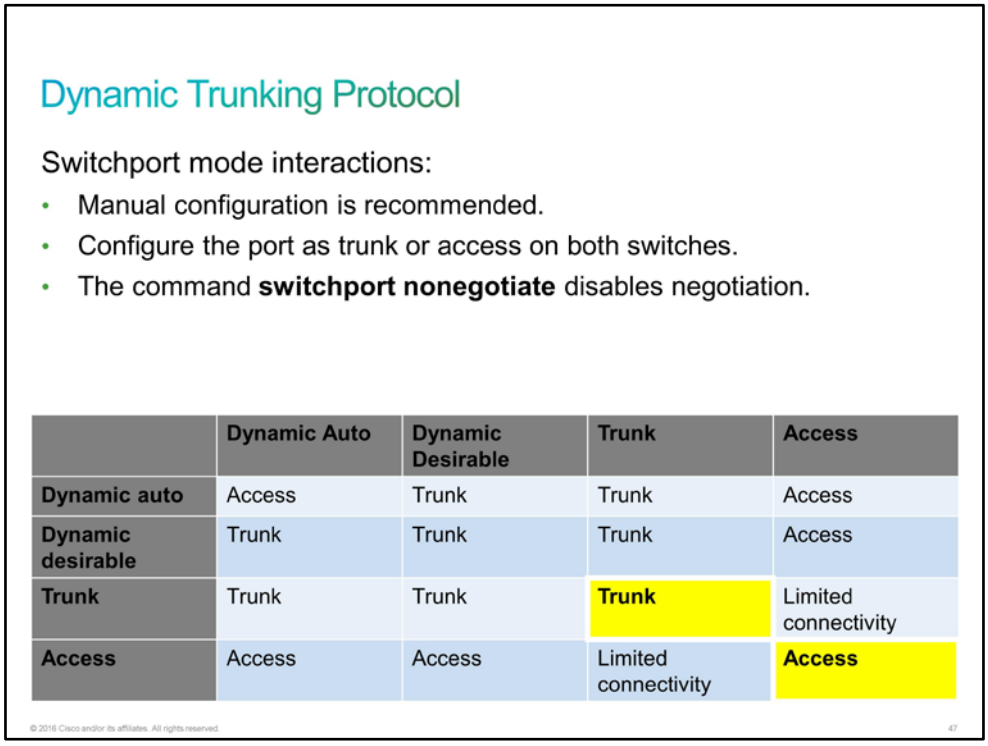

DTP is automatically enabled on a switch port when certain trunking modes are configured on the switch port. DTP manages trunk negotiation only if the port on the other switch is configured in a mode that supports DTP.

You should configure trunk links statically whenever possible. However, Cisco switch ports can run DTP, which can automatically negotiate a trunk link. This protocol can determine an operational trunking mode and protocol on a switch port when it is connected to another device that is also capable of dynamic trunk negotiation.

The default DTP mode depends on the Cisco IOS Software version and on the platform. To determine the current DTP mode, issue the **show dtp interface** command.

```
SW1# show dtp interface FastEthernet0/1
DTP information for FastEthernet0/1:
    TOS/TAS/TNS:                      TRUNK/DESIRABLE/TRUNK
    TOT/TAT/TNT:                      802.1Q/802.1Q/802.1Q
    Neighbor address 1:              001646FA9B01
    Neighbor address 2:               000000000000
    Hello timer expiration (sec/state):  17/RUNNING
    Access timer expiration (sec/state) 287/RUNNING
<... output omitted ...>
```

You can configure the DTP mode to turn off the protocol or to instruct it to negotiate a trunk link only under certain conditions, as described in the table.

Command	Function
switchport mode dynamic auto	Creates the trunk based on the DTP request from the neighboring switch.
switchport mode dynamic desirable	Communicates to the neighboring switch via DTP that the interface is attempting to become a trunk if the neighboring switch interface is able to become a trunk.
switchport mode trunk	Automatically enables trunking regardless of the state of the neighboring switch and regardless of any DTP requests that the neighboring switch sends.
switchport mode access	Trunking not allowed on this port regardless of the state of the neighboring switch interface and regardless of any DTP requests that the neighboring switch sends.
switchport nonegotiate	Prevents the interface from generating DTP frames. This command can be used only when the interface switch port mode is access or trunk. You must manually configure the neighboring interface as a trunk interface to establish a trunk link.

The **switchport nonegotiate** interface command specifies that DTP negotiation packets are not sent. The switch does not engage in DTP negotiation on this interface. This command is valid only when the interface switchport mode is access or trunk (configured by using the **switchport mode access** or the **switchport mode trunk** interface configuration commands). This command returns an error if you attempt to execute it in dynamic (auto or desirable) mode. Use the **no** form of this command to return to the default setting. When you configure a port with the **switchport nonegotiate** command, the port trunks only if the other end of the link is specifically set to trunk. The **switchport nonegotiate** command does not form a trunk link with ports in either dynamic desirable or dynamic auto mode.

Note	A general best practice is to set the interface to **trunk** and **nonegotiate** when a trunk link is required. On links where trunking is not intended, you should turn off DTP. Ideally, links that are not indented to be trunks should be set to access mode and placed in an unused VLAN.

VLAN Trunking Protocol

To minimize misconfiguration and configuration inconsistencies of VLANs in your network, use VTP. VTP is a data link layer (Layer 2) protocol that facilitates the management of VLANs across several switches in a network.

Using VTP, you do not need to log into each switch to create and name each VLAN manually. Managing VLANs manually on each switch in your network works well for a few switches, but VTP is a better solution in large networks.

Note	You still need to assign ports to each VLAN either manually or automatically.

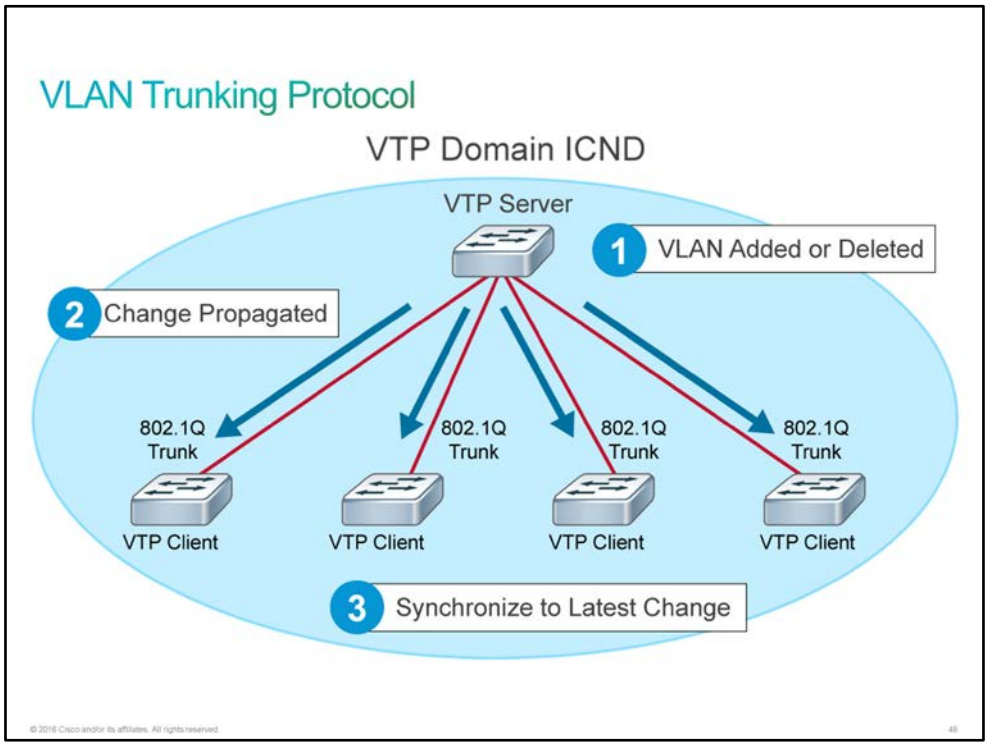

A VTP domain consists of one switch or several interconnected switches sharing the same VTP environment. A switch can belong to only one domain.

By default, a Cisco Catalyst switch is in the no-management-domain state until it receives an advertisement for a domain over a trunk link or until you configure a management domain. The configurations that you make to a VTP server are propagated across trunk links to all the connected switches in the network.

Note	VTP advertisements are flooded throughout the management domain. VTP advertisements are sent every 5 minutes or whenever there is a change in VLAN configurations.

The default VTP version that is enabled on a Cisco switch is version 1. However, three different VTP versions exist: 1, 2, and 3. You can change the switch to run VTP version 2 or 3, but these versions are not compatible. You need to configure the same VTP version on every switch in the domain.

Note	Version 1 and version 2 do not propagate configuration information for extended-range VLANs so you must configure extended-range VLANs manually.

VTP Modes

VTP operates in one of three modes: server, transparent, or client. You can complete various tasks depending on the VTP operation mode.

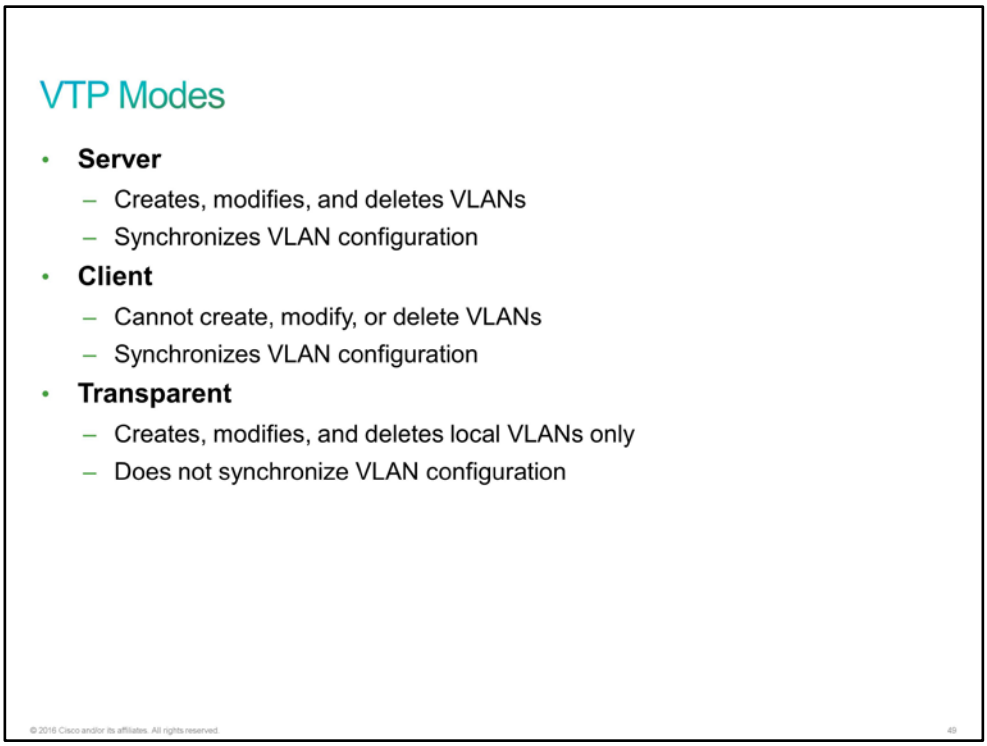

The following are the characteristics of the three VTP modes:

- **Server:** The default VTP mode is server mode. However, VLANs are not propagated over the network until a management domain name is specified or learned. When you change (create, modify, or delete) the VLAN configuration on a VTP server, the change is propagated to all switches in the VTP domain. VTP messages are transmitted out all the trunk connections. A VTP server synchronizes its VLAN database file with other VTP servers and clients. Use the **vtp mode server** Cisco IOS command to configure a switch to be a VTP server.
- **Transparent:** When you change the VLAN configuration in VTP transparent mode, the change affects only the local switch and does not propagate to other switches in the VTP domain. VTP transparent mode forwards VTP advertisements that it receives within the domain. A VTP transparent device does not synchronize its database with any other device. Use the **vtp mode transparent** Cisco IOS command to configure a switch to be transparent.
- **Client:** You cannot change the VLAN configuration when in VTP client mode. However, a VTP client can send any VLANs that are currently listed in its database to other VTP switches. VTP advertisements are forwarded in VTP client mode. A VTP client synchronizes its database with other VTP servers and clients. You can use the **vtp mode client** Cisco IOS command to configure a switch to be a VTP client.

VTP Configuration

When creating VLANs, you must decide whether to use VTP in your network. With VTP, you can make configuration changes on one or more switches, and those changes are automatically communicated to all other switches in the same VTP domain.

Default VTP configuration values depend on the switch model and the software version. The following are the default values for Cisco Catalyst switches:

- **VTP domain name:** Null
- **VTP mode:** Server
- **VTP password:** None
- **VTP pruning:** Enabled or disabled (operating system version-specific)
- **VTP version:** Version 1

Note	When the VTP pruning option is enabled in a VTP domain, VTP client switches receive VTP update frames only for VLANs that are enabled on each switch. Thus, VTP pruning saves some bandwidth on trunk ports and on switches by limiting the number of VTP update transmissions. You should always prune the VLANs from switches where the VLANs are not used.

The VTP domain name can be specified or learned. By default, the domain name is not set. You can set a password for the VTP management domain. However, if you do not assign the same password for each switch in the domain, VTP does not function properly.

VTP pruning eligibility is one VLAN parameter that the VTP protocol advertises. Enabling or disabling VTP pruning on a VTP server propagates the change throughout the management domain.

Use the **vtp** global configuration command to modify the VTP configuration, domain name, interface, and mode:

```
Switch# configure terminal
Switch(config)# vtp mode [server | client | transparent]
Switch(config)# vtp domain domain-name
Switch(config)# vtp password password
Switch(config)# vtp pruning
```

Use the **no** form of this command to remove the filename or to return to the default settings. When the VTP mode is transparent, you can save the VTP configuration in the switch configuration file by entering the **copy running-config startup-config** privileged EXEC command.

The following example demonstrates how to configure VTP and display VTP status.

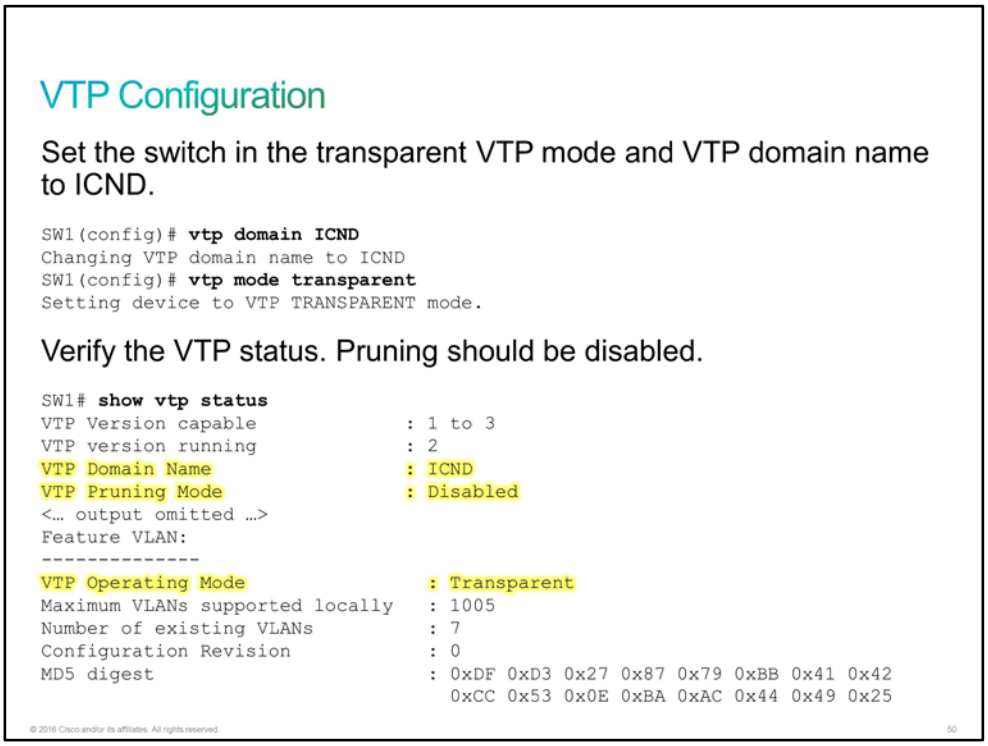

Note	In the output of the **show vtp status** command, "VTP Version capable" identifies the version of VTP that the switch is capable of running. "VTP version running" indicates which VTP version is being used.

On switches that are configured in VTP client or VTP server mode, you cannot see any configuration related to VLANs or VTP in the running configuration. To verify VTP configuration, you have to use **show vtp status** and **show vtp password** commands. To verify configured VLANs, you should use **show vlan** command.

Discovery 17: Troubleshoot VLANs and Trunks

Introduction

This discovery will guide you through a scenario involving VLAN configuration, Layer 2 connectivity, and IP connectivity. The topology diagram is intentionally vague and there is no connectivity table. So, you are on your first day at a new job as a network engineer. You are not yet familiar with the network of your organization. A member of the security team comes to you because the intrusion prevention system has flagged malicious traffic from the IP address 10.10.10.182. You are asked to help in isolating this system and removing it from the network.

This discovery will also guide you through the IP connectivity issue between two hosts.

Topology

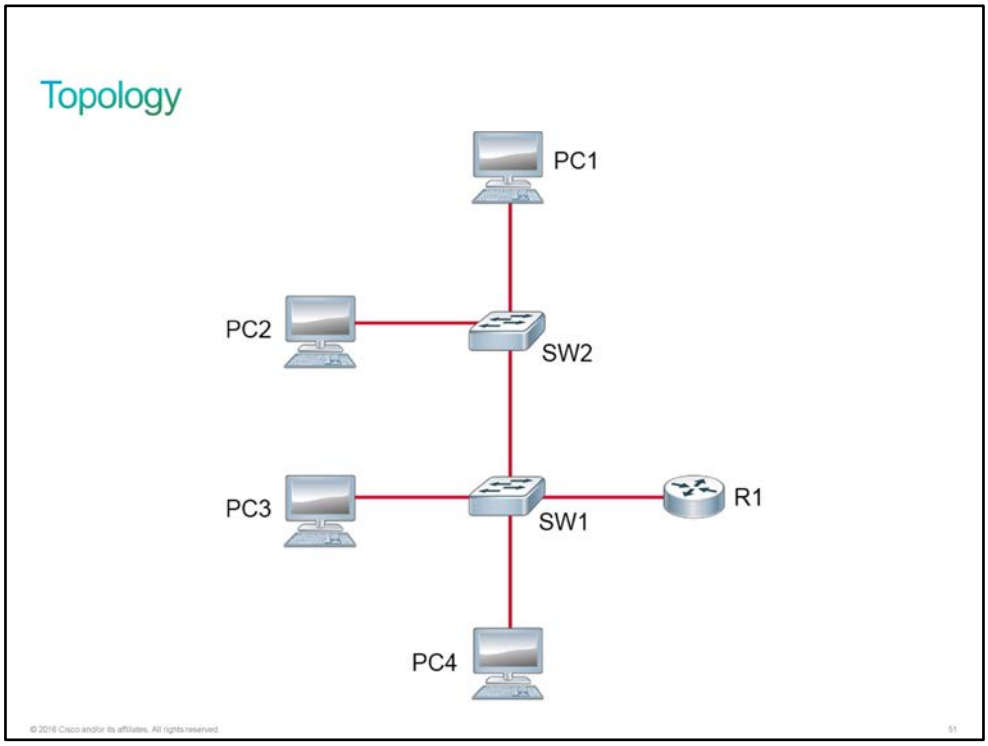

Job Aid

There is no Job Aid available for this lab exercise, because one of the objectives of the lab is to map the connectivity within an unfamiliar network.

Note	PCs in the virtual lab environment are simulated as routers, so you should use Cisco IOS commands to configure them or make verifications.

Task 1: Troubleshoot VLAN Issues

The following figure shows the flow for troubleshooting VLANs.

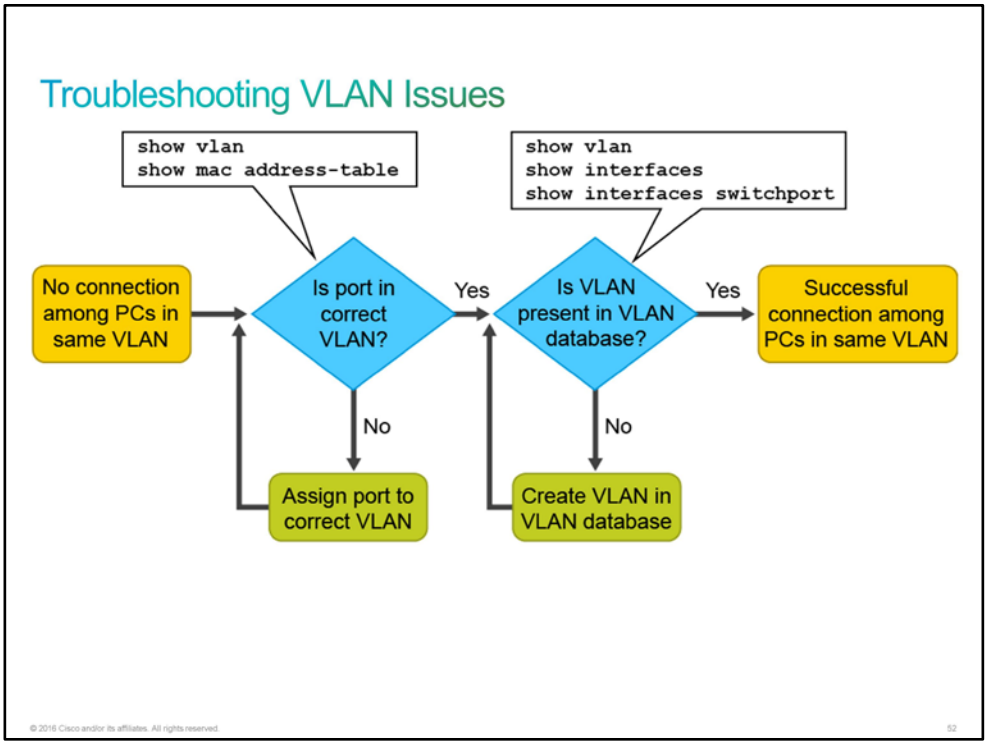

To troubleshoot VLAN issues when you have no connection between PCs that belong to the same VLAN, follow these high-level steps:

1. Use the **show vlan** command to check whether the port belongs to the expected VLAN. If the port is assigned to the wrong VLAN, use the **switchport access vlan** command to correct the VLAN membership. Use the **show mac address-table** command to check which addresses were learned on a particular port of the switch and to which VLAN that port is assigned.
2. If the VLAN to which the port is assigned is deleted, the port becomes inactive. Use the **show vlan** or **show interfaces switchport** command to verify that the VLAN is present in the VLAN database.
   - Also note, that you can shut the VLAN using **shutdown** command, so you may need to verify that the VLAN is not disabled using the **show vlan** command.

MAC Address Table Verification

To display the MAC address table, use the **show mac address-table** command in privileged EXEC mode as shown in the following example. This command displays the MAC address table for the switch. You can define specific views by using the optional keywords and arguments. The example shows MAC addresses that were learned on the FastEthernet0/1 interface. As you can see, MAC address 000c.296a.a21c was learned on the interface FastEthernet0/1 in VLAN 10. If this number is not the expected VLAN number, change the port VLAN membership using the **switchport access vlan** command.

```
SW1# show mac address-table interface Ethernet0/1
          Mac Address Table
-------------------------------------------

Vlan    Mac Address       Type        Ports
----    -----------       --------    -----
  10    000c.296a.a21c    DYNAMIC     Fa0/1
  10    000f.34f9.9181    DYNAMIC     Fa0/1
Total Mac Addresses for this criterion: 2
```

Troubleshooting Missing VLANs

Each port on a switch belongs to a VLAN. If the VLAN to which the port belongs is deleted, the port becomes inactive. All ports belonging to the VLAN that was deleted are unable to communicate with the rest of the network.

As shown in the following example, use the command **show interface***interface* **switchport** to check whether the port is inactive. If the port is inactive, it will not be functional until you create the missing VLAN using the **vlan** *vlan_id* command or until you assign the port to a valid VLAN.

```
SW1# show interfaces Ethernet0/1 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: 10 (Inactive)
Trunking Native Mode VLAN: 1 (default)
Administrative Native VLAN tagging: enabled
Voice VLAN: none
```

Activity

Complete the following steps:

Step 1 Your task in this discovery is to find the system using the IP address 10.10.10.182 and to disconnect it from the network. You might assume that VLANs were configured by a logical pattern.

Access the console of SW1 and display the VLAN configuration to show how incorrect that assumption is.

```
SW1# show vlan

VLAN Name                              Status    Ports
---- --------------------------------- --------- -------------------------------
1    default                           active
62   SixtyTwo                          active    Et0/2, Et0/3
1002 fddi-default                      act/unsup
1003 token-ring-default                act/unsup
1004 fddinet-default                   act/unsup
1005 trnet-default                     act/unsup
<... output omitted ...>
```

This disorganized set of VLANs demonstrates why it can be beneficial to set a standard. For example, you can have the VLAN ID match the third octet of the IP network running on that VLAN.

Step 2 To determine which VLAN supports the network to which 10.10.10.182 belongs, access the console of R1 and display the brief summary status of its IP interfaces.

When the display output pauses with the --More-- prompt, you can use the space bar to display the next page of the output.

```
R1# show ip interface brief
Interface                  IP-Address      OK? Method Status
Protocol
Ethernet0/0                unassigned      YES manual up                    up
Ethernet0/0.21             10.10.1.1       YES manual up                    up
Ethernet0/0.134            10.10.10.1      YES manual up                    up
Ethernet0/1                unassigned      YES NVRAM  administratively down
down
Ethernet0/2                unassigned      YES NVRAM  administratively down
down
Ethernet0/3                unassigned      YES NVRAM  administratively down
down
Serial1/0                  unassigned      YES NVRAM  administratively down
down
Serial1/1                  unassigned      YES NVRAM  administratively down
down
Serial1/2                  unassigned      YES NVRAM  administratively down
down
Serial1/3                  unassigned      YES NVRAM  administratively down
down
Loopback0                  10.10.99.1      YES manual up                    up
```

Step 3 The IP address of Ethernet0/0.134 is 10.10.10.1. If you configure it with a 24-bit subnet mask, it would be on the same subnet as 10.10.10.182. If its subinterface ID matches the VLAN ID, the VLAN would be 134. Display the running configuration that is associated with this interface to determine if either of these values are true.

Verify the running configuration on the R1 router:


```
R1# show run interface Ethernet0/0.134
Building configuration...

Current configuration : 94 bytes
!
interface Ethernet0/0.134
 encapsulation dot1Q 62
 ip address 10.10.10.1 255.255.255.0
end
```

The mask is indeed 24 bits. This interface is on the same subnet as 10.10.10.182.

The VLAN, as set by the encapsulation command, is actually 62, not 134.

Step 4 The security team member gave you the IP address. Determine the system MAC address by first pinging it from R1 and then finding the entry in the R1 ARP cache.

```
R1# ping 10.10.10.182
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.10.10.182, timeout is 2 seconds:
.!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
R1# show ip arp 10.10.10.182
Protocol  Address          Age (min)  Hardware Addr   Type   Interface
Internet  10.10.10.182           23   aabb.cc00.5300  ARPA   Ethernet0/0.134
```

The system that you are looking for has the MAC address aabb.cc00.5300.

Note: The MAC address that you will see in your output can be different. Further in the lab, refer to the MAC address determined in your output.

Step 5 Access the console of SW1 and view its MAC address table to find the port that is connecting to aabb.cc00.5300, or whatever your MAC address is.

You have to search for the MAC address that you discovered in the previous step.

```
SW1# show mac address-table
          Mac Address Table
-------------------------------------------

Vlan    Mac Address       Type        Ports
----    -----------       --------    -----
   1    aabb.cc00.5000    DYNAMIC     Et0/0
  62    aabb.cc00.5000    DYNAMIC     Et0/0
  62    aabb.cc00.5200    DYNAMIC     Et0/1
  62    aabb.cc00.5300    DYNAMIC     Et0/2
  62    aabb.cc00.5400    DYNAMIC     Et0/3
Total Mac Addresses for this criterion: 5
```

Interface Ethernet0/2 is where the offending system is connected.

Step 6 Since there were few addresses in the MAC address table, it was pretty easy to pick out the appropriate entry. If there are thousands of entries in the table, you would want to filter down the output. Try displaying the MAC address table using the include filter to only include addresses that have 5300, or whatever the last 4 digits of your MAC address are, as part of their address.

In a larger environment, you might find that the port with the offending MAC address is actually a link to another switch. In this case, you would have to go to that switch and view its MAC address table. It might again be on a link to a third switch. You would have to continue the process until you reached a switch with the address on an end-host port.

```
SW1# show mac address-table | include 5300
  62    aabb.cc00.5300    DYNAMIC     Et0/2
```

Step 7 Display the interface status summary on SW1 to observe the status of Ethernet0/2.

One thing that was sensibly configured in this environment is the description on the switch ports. PC3 is the offending system.

```
SW1# show interface status

Port      Name               Status       Vlan       Duplex  Speed Type
Et0/0     Link to R1         connected    trunk        auto   auto unknown
Et0/1     Link to SW2        connected    trunk        auto   auto unknown
Et0/2     Link to PC3        connected    62           auto   auto unknown
Et0/3     Link to PC4        connected    62           auto   auto unknown
```

Step 8 Verify that the offending system, PC3, has access to the network. Attempt to ping R1 (10.10.10.1) from PC3.

Ping should be successful.

```
PC3# ping 10.10.10.1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.10.10.1, timeout is 2 seconds:
.!!!!
Success rate is 80 percent (4/5), round-trip min/avg/max = 1/1/1 ms
```

Step 9 Disable interface Ethernet0/2 on SW1.

On SW1, enter the following commands:

```
SW1# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
SW1(config)# interface Ethernet0/2
SW1(config-if)# shutdown
SW1(config-if)#
*Sep 17 07:22:54.192: %LINK-5-CHANGED: Interface Ethernet0/2, changed state to
administratively down
*Sep 17 07:22:55.196: %LINEPROTO-5-UPDOWN: Line protocol on Interface
Ethernet0/2, changed state to down
SW1(config-if)# end
SW1#
*Sep 17 07:22:57.180: %SYS-5-CONFIG_I: Configured from console by console
SW1#
```

Step 10 The offending system is PC3. Access the console of PC3 and verify that it has been isolated from the network. Attempt to ping R1 (10.10.10.1).

The attempt should fail.


```
PC3# ping 10.10.10.1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.10.10.1, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
```

Task 2: Troubleshoot Trunk Issues

The figure shows the flow for troubleshooting trunks.

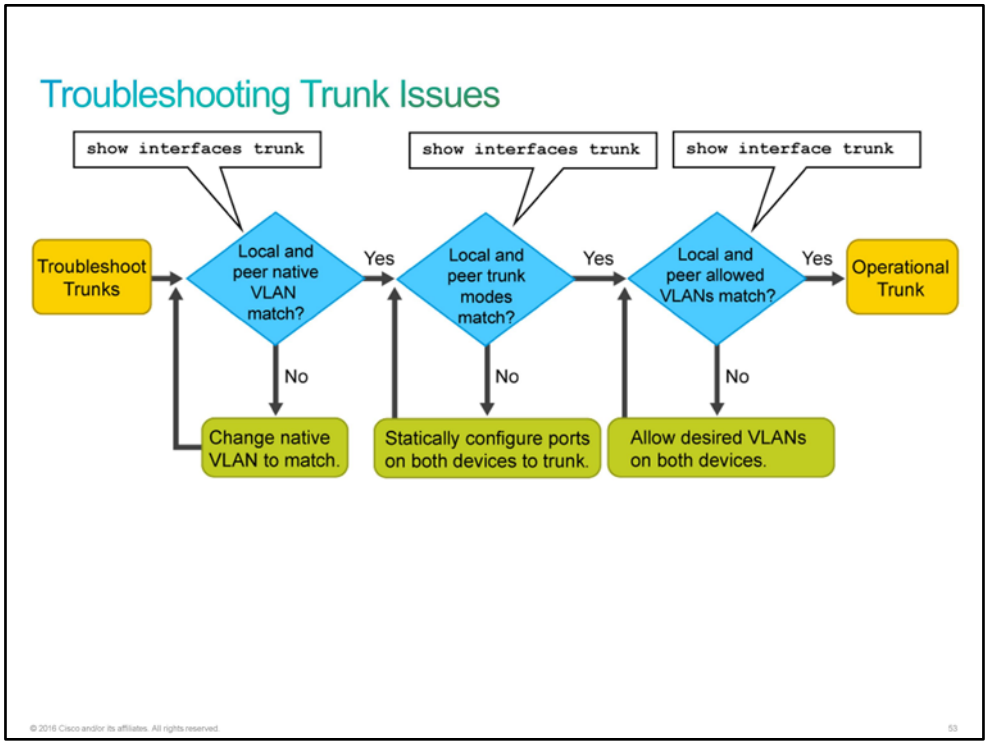

To troubleshoot trunk issues when the trunk is not established, follow these high-level steps:

1. Use the **show interfaces trunk** command to check whether the local and peer native VLANs match. If the native VLAN does not match on both sides, VLAN leaking occurs.
2. Use the **show interfaces trunk** command to check whether a trunk has been established between switches. You should statically configure trunk links whenever possible. However, Cisco Catalyst switch ports by default run DTP, which tries to negotiate a trunk link.
3. Use the **show interface trunk** command to check whether the desired VLANs have been allowed on both the sides of the trunk link.

Verifying Trunk Establishment

To display the status of the trunk and native VLAN that is used on a trunk link and to verify trunk establishment, use the **show interface trun**k command in privileged EXEC mode. The following example shows that the native VLAN on one side of the trunk link was changed to VLAN 2. If one end of the trunk is configured as native VLAN 1 and the other end is configured as native VLAN 2, a frame that is sent from VLAN 1 on one side is received on VLAN 2 on the other. VLAN 1 "leaks" into the VLAN 2 segment. This behavior would never be required, and connectivity issues occur in the network if a native VLAN mismatch exists. Change the native VLAN to the same VLAN on both sides of the VLAN to avoid this behavior.

```
SW1# show interfaces Ethernet 0/3 trunk

Port        Mode             Encapsulation  Status        Native vlan
Et0/3       auto             802.1q         not-trunking  2

<...output omitted...>
```

Cisco Discovery Protocol notifies you of a native VLAN mismatch on a trunk link with this message:

```
Aug 31 08:34:48.714: %CDP-4-NATIVE_VLAN_MISMATCH: Native VLAN mismatch discovered on
FastEthernet0/3 (2),
with SW2 FastEthernet0/3 (1).
```

You should statically configure trunk links whenever possible. Cisco Catalyst switch ports by default run DTP. DTP can determine the operational trunking mode and protocol on a switch port when it is connected to another device that is also capable of dynamic trunk negotiation. Remember that if both ends of a trunk are set to dynamic auto trunk mode, a trunk will not be established. The example shows the status of the link as "not-trunking."

Activity

Complete the following steps:

Step 1 User that is using PC1 is reporting that PC1 can reach PC2 (10.10.10.20), but cannot reach PC4 (10.10.10.40). Help the user find the issue and resolve it.

Access PC1 and verify IP connectivity to PC2 and PC4 to exclude an IP connectivity issue.

```
PC1# ping 10.10.10.20
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.10.10.20, timeout is 2 seconds:
.!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
PC1# ping 10.10.10.40
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.10.10.40, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
```

You should find out that there is an IP connectivity issue between PC1 and PC4.

Step 2 Access the SW2 switch and check which VLAN is set on the interface that PC1 is connected to.

First, you need to use Cisco Discovery Protocol to verify to which port PC1 is connected.

Note: With real PCs, PC would not be seen as CDP neighbor, so you would need to use the same approach that you used in the first procedure of this discovery.

```
SW2# 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
PC2              Eth 0/2           133              R    Linux Uni Eth 0/0
PC1              Eth 0/1           177              R    Linux Uni Eth 0/0
SW1              Eth 0/0           170            S I    Linux Uni Eth 0/1
SW2# show vlan

VLAN Name                              Status    Ports
---- --------------------------------- --------- -------------------------------
1    default                           active    Et0/3
62   SixtyTwo                          active    Et0/1, Et0/2
1002 fddi-default                      act/unsup
1003 token-ring-default                act/unsup
1004 fddinet-default                   act/unsup
1005 trnet-default                     act/unsup
<... output omitted ...>
```

You will find out that PC1 is connected to Ethernet0/1 and that it is placed into active VLAN 62.

Step 3 Access the SW1 switch and check which VLAN is set on the interface that PC4 is connected to.

First, you need to use Cisco Discovery Protocol to verify to which port PC4 is connected.

Note: With real PCs, PC would not be seen as CDP neighbor, so you would need to use the same approach that you used in the first procedure of this discovery.

```
SW1# 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
PC4              Eth 0/3           129              R    Linux Uni Eth 0/0
SW2              Eth 0/1           170            S I    Linux Uni Eth 0/0
R1               Eth 0/0           163              R    Linux Uni Eth 0/0.21
SW1# show vlan

VLAN Name                              Status    Ports
---- --------------------------------- --------- -------------------------------
1    default                           active
62   SixtyTwo                          active    Et0/2, Et0/3
1002 fddi-default                      act/unsup
1003 token-ring-default                act/unsup
1004 fddinet-default                   act/unsup
1005 trnet-default                     act/unsup
<... output omitted ...>
```

You will find out that both PC1 and PC4 are in the same VLAN.

Step 4 While troubleshooting, you first noticed the following message on the SW1 console:

```
*Sep 17 09:09:21.594: %CDP-4-NATIVE_VLAN_MISMATCH: Native VLAN mismatch
discovered on Ethernet0/1 (1), with SW2 Ethernet0/0 (2).
```

This message indicates that SW1 and SW2 have different native VLANs configured.

On SW1, check which VLAN is used as native on Ethernet0/1:

```
SW1# show interfaces trunk

Port        Mode             Encapsulation  Status        Native vlan
Et0/0       on               802.1q         trunking      1
Et0/1       on               802.1q         trunking      1
<... output omitted ...>
```

On SW2, check which VLAN is used as native on Ethernet0/1:

```
SW2# show interfaces trunk

Port        Mode             Encapsulation  Status        Native vlan
Et0/0       desirable        802.1q         trunking      2
<... output omitted ...>
```

Step 5 Change the native VLAN configuration on the SW2 switch.

On SW2, enter the following commands:

```
SW2# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
SW2(config)# interface Ethernet0/0
SW2(config-if)# switchport trunk native vlan 1
```

Messages to the console stopped.

Step 6 Verify if native VLAN was the reason for broken connectivity between PC1 and PC4.

Access PC1 and verify IP connectivity to PC4.

```
PC1# ping 10.10.10.40
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.10.10.40, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
```

PC1 still has no connectivity to PC4, so you need to investigate further.

Step 7 You have determined that PC1 and PC4 are both in VLAN 62. Now, you will verify trunk link between SW1 and SW2.


```
SW2# show interfaces trunk

Port        Mode             Encapsulation  Status        Native vlan
Et0/0       desirable        802.1q         trunking      1

Port        Vlans allowed on trunk
Et0/0       1-1000

Port        Vlans allowed and active in management domain
Et0/0       1,62

Port        Vlans in spanning tree forwarding state and not pruned
Et0/0       1,62
```

VLAN 62 is correctly allowed on the link to SW1.

```
SW1# show interfaces trunk

Port        Mode             Encapsulation  Status        Native vlan
Et0/0       on               802.1q         trunking      1
Et0/1       on               802.1q         trunking      1

Port        Vlans allowed on trunk
Et0/0       1-4094
Et0/1       1-61,63-1000

Port        Vlans allowed and active in management domain
Et0/0       1,62
Et0/1       1

Port        Vlans in spanning tree forwarding state and not pruned
Et0/0       1,62
Et0/1       1
```

VLAN 62 is missing from the allowed VLANs on the link toward SW2.

Step 8 On SW1, verify the interface Ethernet0/1 configuration.

Here, you can confirm that VLAN 62 is excluded from the allowed VLAN list:

```
SW1# show run interface Ethernet0/1
Building configuration...

Current configuration : 172 bytes
!
interface Ethernet0/1
 description Link to SW2
 switchport trunk encapsulation dot1q
 switchport trunk allowed vlan 1-61,63-1000
 switchport mode trunk
 duplex auto
end
```

Step 9 On the SW1 interface Ethernet0/1, add VLAN 62 into trunk.

On SW1, enter the following commands:

```
SW1# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
SW1(config)# interface Ethernet0/1
SW1(config-if)# switchport trunk allowed vlan add 62
```

Step 10 From PC1, verify that the IP connectivity issue to the PC4 is resolved.

Ping should be successful:

```
PC1# ping 10.10.10.40
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.10.10.40, timeout is 2 seconds:
.!!!!
Success rate is 80 percent (4/5), round-trip min/avg/max = 1/1/1 ms
```

Note: It may take a while for the ping to work.

This is the end of the discovery lab.

VLAN Design Consideration

VLANs create boundaries that can isolate nodes or traffic, so you should design a multi-VLAN topology thoughtfully. The general question that you should ask yourself is the following: "Who is talking to whom and what are they trying to get done?" Here are some considerations that you need to take into account before implementing VLANs.

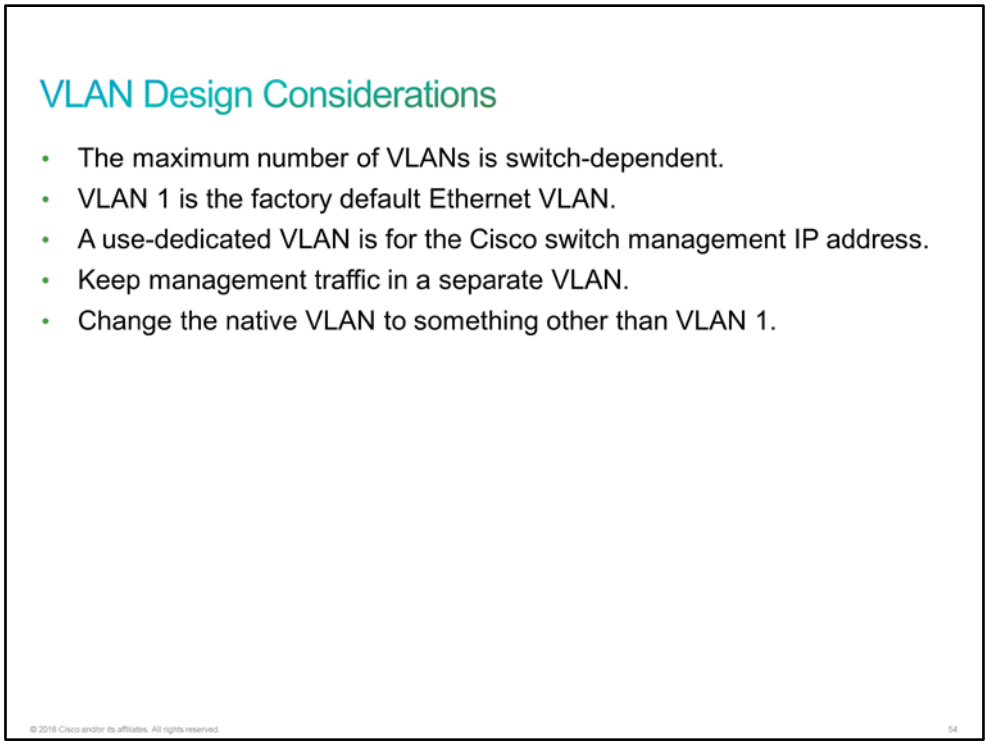

Typically, access layer Cisco switches support up to 64, 256, or 1024 VLANs. The maximum number of VLANs is switch-dependent.

Cisco switches have a factory default configuration in which various default VLANs are preconfigured to support various media and protocol types. The default Ethernet VLAN is VLAN 1. Cisco Discovery Protocol advertisements are sent on VLAN 1. A good security practice is to separate management and user data traffic because you do not want users to be able to establish Telnet sessions to the switch.

If you want to communicate with a Cisco switch remotely for management purposes, the switch must have an IP address. This IP address must be in the management VLAN, which is VLAN 1 by default.

A good security practice is to change the native VLAN to something other than VLAN 1 (for example, VLAN 90) and therefore tag the VLAN 1 traffic.

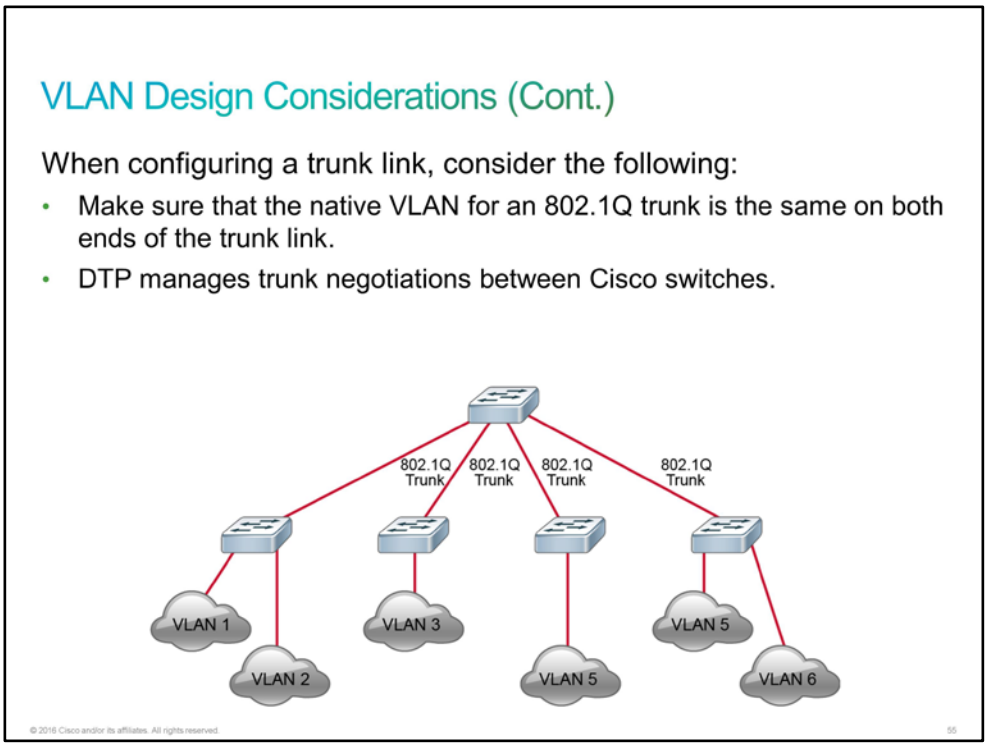

Make sure that the native VLAN for an 802.1Q trunk is the same on both ends of the trunk link. If the ends are different, spanning-tree loops might result. If IEEE 802.1Q trunk configuration is not the same on both ends, Cisco IOS Software will report error messages. Also make sure that native VLAN frames are untagged.

The DTP helps to automatically negotiate whether the port should be put into the access or trunk mode and which trunking protocol (802.1Q or ISL) should be used. The individual DTP modes are *dynamic auto* (the port will negotiate the mode automatically; however, it prefers to be an access port) and *dynamic desirable* (the port will negotiate the mode automatically; however, it prefers to be a trunk port). If you do not want the switch to negotiate, use the **switchport nonegotiate** command. For details on all the parameters that are associated with the **switchport mode** interface command, go to http://www.cisco.com/en/US/docs/ios/mcl/allreleasemcl/all_15.html.

Challenge

1. Which of the three layers in hierarchical LAN design will you implement Routing and packet manipulation?

   A. Access
   B. Distribution
   C. Core

2. A poorly designed network includes (or is associated with) which of the following? (Choose two.)

   A. Large broadcast domains
   B. Small broadcast domains
   C. Several Security Vulnerabilities
   D. Proper documentation
   E. Ease of Management and support

3. VLANs improve network performance by doing which of the following:

   A. Separating large broadcast domains into smaller segments
   B. Creating large broadcast domains out of smaller segments
   C. Creating one large virtual switch out of many physical switches
   D. Allowing users to connect over radio frequency

4. Traffic from one VLAN can reach another VLAN on a Layer 2 switch. True or False?

   A. True
   B. False

5. How does 802.1Q incorporate VLAN information onto a packet?

   A. It creates a 4 byte header and a 26 byte tag
   B. It creates a 4 byte header only
   C. It changes the frame body to mention the VLAN information
   D. It re-routes the frame through the VLAN interface, causing a different source Address.

6. Which of the following must you ensure when configuring two ends of a 802.1Q trunk?

   A. Native VLAN must be tagged
   B. Native VLAN must be the same
   C. DTP must be disabled
   D. DTP modes on both ends must be the same

7. Which one of the following is correct?

   A. STP blocks certain ports to increase efficiency and only allow the ports that are utilized to be 'Up'.
   B. STP blocks certain ports to ensure that loops do not occur.
   C. STP is disabled by default on Cisco Switches.
   D. If there is a problem with connectivity, STP alerts the administrator so the issue can be rectified.

Answer Key

Challenge

1. B
2. A, C
3. A
4. B
5. B
6. B
7. B

Lesson 2: Building Redundant Switched Topologies

Introduction

The law firm's client calls CCS complaining that employees in its international and constitutional law departments are unable to communicate digitally or share resources on the intranet. Bob has already determined that the cause of the problem is the failure of a single switch. The law firm has agreed to have CCS implement and troubleshoot a redundant switched topology and optimize network reliability by implementing PVST+. Bob wants to know if you are ready to go to the law firm to implement and troubleshoot the redundant switched topology, or if you need some time to prepare.

Physical Redundancy in a LAN

Loops can affect performance in a switched LAN, and STP is a solution.

Loops may occur in the network as part of a design strategy for redundancy. Adding switches to LANs can add the benefit of redundancy. Connecting two switches to the same network segments ensures continuous operation if there are problems with one of the segments. Redundancy can ensure the constant availability of the network.

However, when switches are used for redundancy in a network, loops are a potential problem. When a host on one network segment transmits data to a host on another network segment, and the two are connected by two or more switches, each switch receives the data frames, looks up the location of the receiving device, and forwards the frame. Because each switch forwards the frame, each frame is duplicated. As a result, a loop occurs, and the frame circulates between the two paths without being removed from the network. The MAC address tables may also be updated with incorrect address information, resulting in inaccurate forwarding.

In the topology that is shown in the figure, suppose that host A sends a frame to host B. Host A resides on network segment A, and host B resides on network segment B. Redundant connections between hosts ensure continuous operation if a segment fails. For this example, it is assumed that none of the switches have learned the address of host B.

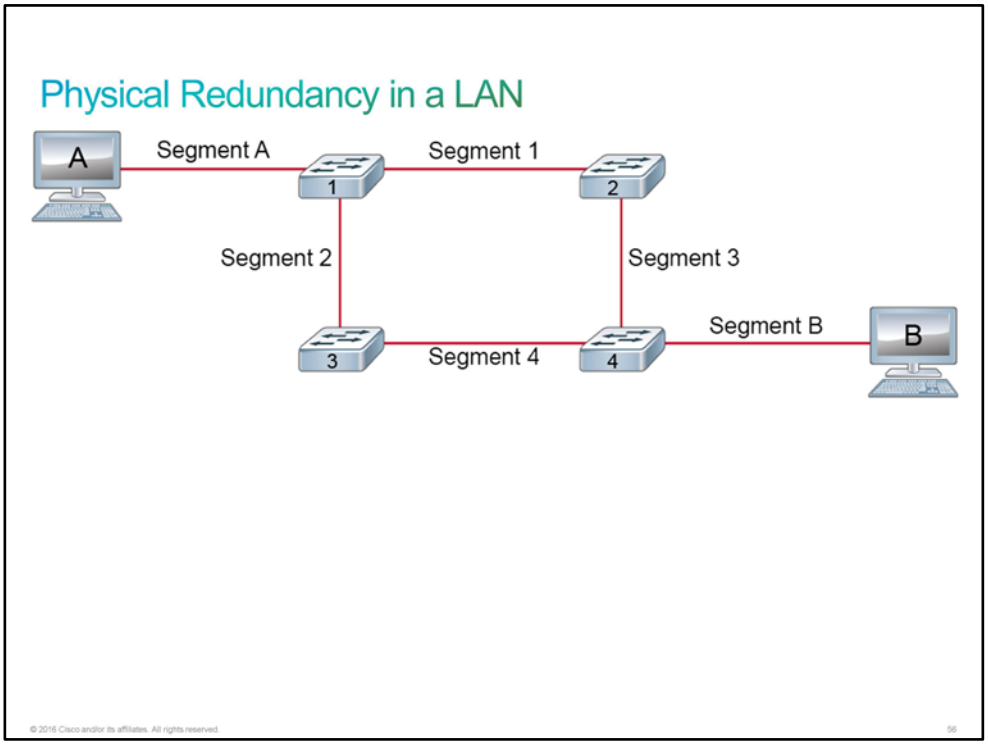

Switch 1 receives the frame that is destined for host B and floods it out to switches 2 and 3. Switch 2 and switch 3 both receive the frame from host A (via switch 1) and correctly learn that host A is on segment 1 and 2. Each switch forwards the frame to switch 4.

Switch 4 receives two copies of the frame from host A, one copy through switch 2 and one copy through switch 3. Assume that the frame from switch 2 arrives first. Switch 4 learns that host A resides on segment 3. Because switch 4 does not know where host B is connected, it forwards the frame to all its ports (except the incoming port) and therefore to host B and switch 3. When the frame from switch 3 arrives at switch 4, switch 4 updates its table to indicate that host A resides on segment 4. It then forwards the frame to host B and switch 2.

Switches 2 and 3 now change their internal tables to indicate that host A is on segment 3 and 4. If the initial frame from host A was a broadcast frame, both switches forward the frames endlessly. They would use all available network bandwidth and block transmission of other packets on both segments. This situation is called a *broadcast storm*.

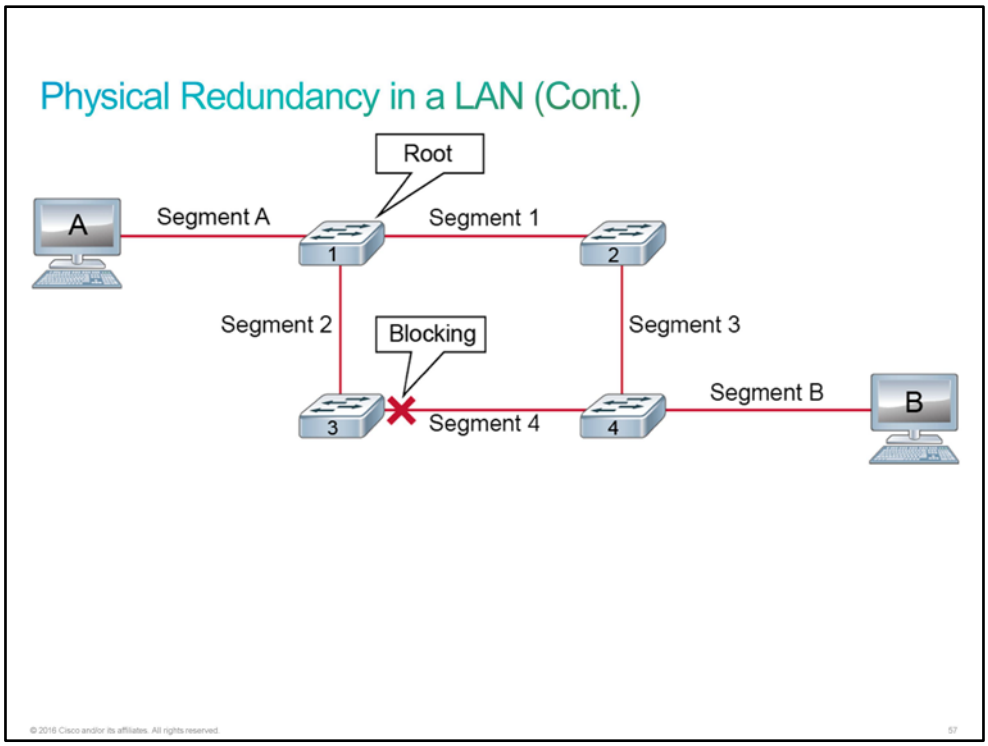

The solution to loops is STP, which manages the physical paths to given network segments. STP provides physical path redundancy while preventing the undesirable effects of active loops in the network. By default, STP is turned on in Cisco Catalyst switches.

STP behaves as follows:

- STP forces certain ports into a standby state so that they do not listen to, forward, or flood data frames. The overall effect is that there is only one path to each network segment that is active at any time.
- If there is a problem with connectivity to any of the segments within the network, STP re-establishes connectivity by automatically activating a previously inactive path, if one exists.

Issues in Redundant Topologies

Enterprise voice and data networks are designed with physical component redundancy to eliminate the possibility of any single point of failure causing a loss of function for an entire switched network. However, redundant OSI Layer 2 switch topologies require planning and configuration to operate without introducing loops.

OSI Layer 2 LAN protocols, such as Ethernet, lack a mechanism to recognize and eliminate endlessly looping frames, as illustrated in the figure.

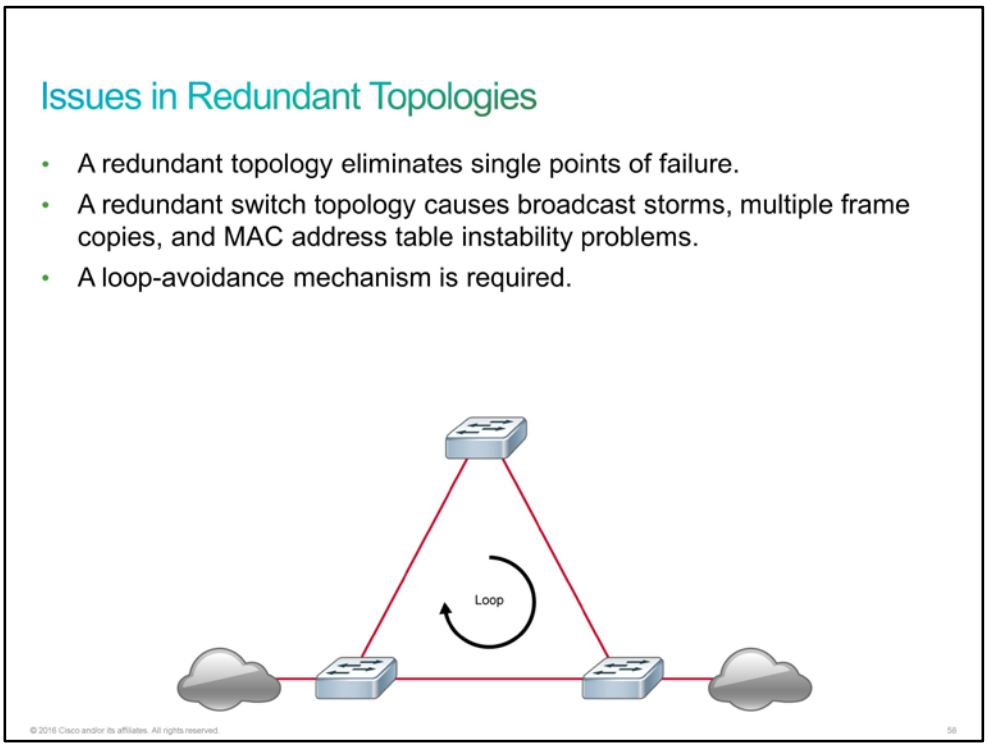

In the absence of a protocol to monitor link forwarding states, a redundant switch topology is vulnerable to the following conditions:

- **Broadcast storms:** Without some loop-avoidance process, each switch floods broadcasts endlessly. This situation is commonly called a broadcast storm.
- **Multiple frame transmission:** Multiple copies of unicast frames may be delivered to destination stations. Many protocols expect to receive only a single copy of each transmission. Multiple copies of the same frame can cause unrecoverable errors.
- **MAC database instability:** Instability in the content of the MAC address table results from the fact that different ports of the switch receive copies of the same frame. Data forwarding can be impaired when the switch consumes the resources that are coping with instability in the MAC address table.

Layer 2 LAN protocols, such as Ethernet, do not have a mechanism to recognize and eliminate endlessly looping frames. Some Layer 3 protocols implement a TTL mechanism that limits the number of times that a Layer 3 networking device can retransmit a packet. Lacking such a mechanism, Layer 2 devices continue to retransmit looping traffic indefinitely.

A loop-avoidance mechanism solves these problems. STP was developed to address them.

Loop Resolution with STP

STP provides loop resolution by managing the physical paths to given network segments. STP allows physical path redundancy while preventing the undesirable effects of active loops in the network. STP is an IEEE committee standard, which is defined as 802.1D.

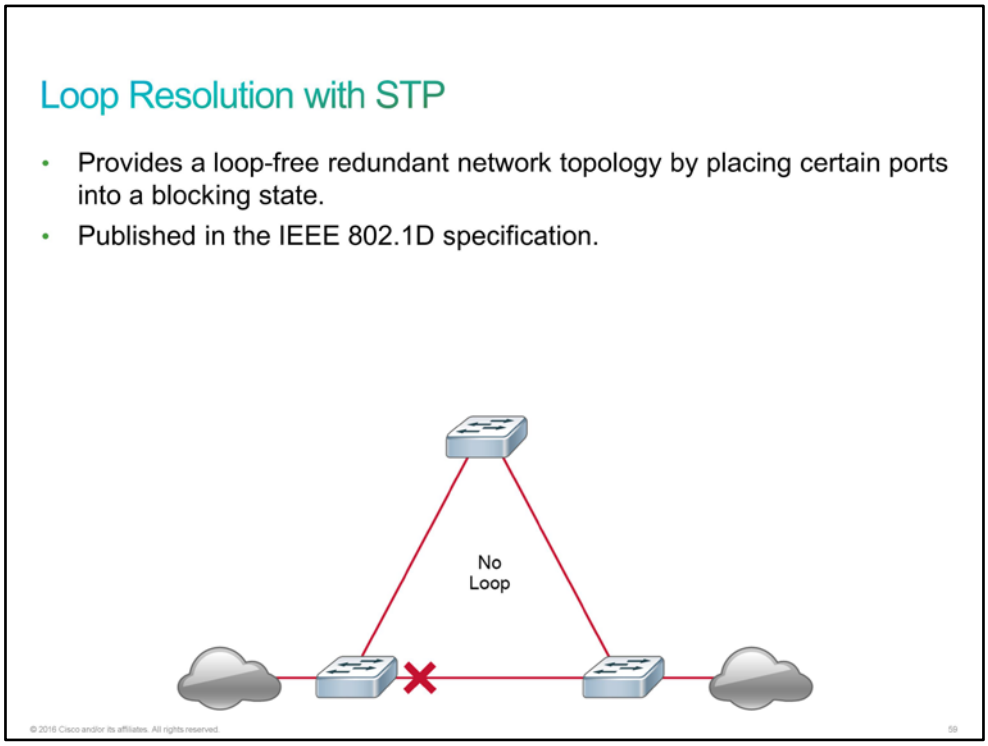

STP behaves as follows:

- STP uses BPDUs for communication between switches.
- STP forces certain ports into a blocked state so that they do not listen to, forward, or flood data frames. The overall effect is that only one path to each network segment is active at any time.
- If there is a problem with connectivity to any of the segments within the network, STP re-establishes connectivity by automatically activating a previously inactive path, if one exists (changing blocked port to forwarding state).

Spanning-Tree Operation

STP and its successor protocols provide loop resolution by managing the physical paths to given network segments. STP allows physical path redundancy while preventing the undesirable effects of active loops in the network. STP forces certain ports into a blocking state. These blocking ports do not forward data frames.

The overall effect is that only one path to each network segment is active at any time. If there is a problem with connectivity to any of the segments within the network, STP re-establishes connectivity by automatically activating a previously inactive path, if one exists.

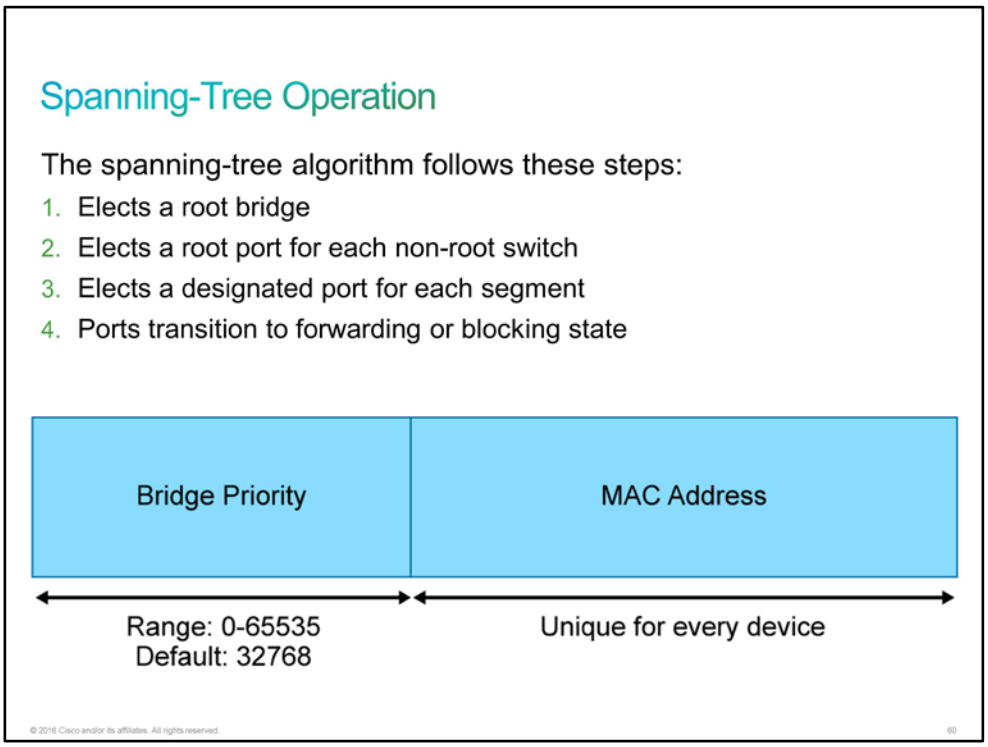

The following are the steps of the spanning-tree algorithm:

1. Elects a root bridge. The root bridge becomes the switch with the lowest BID. You can have only one root bridge per network. Bridge ID is a combination of bridge priority and the MAC address of the switch. Bridge priority is a number between 0 and 65535 in increments of 4096, and the default is 32768. If one or more bridges have equally lowest bridge priorities, then the bridge with the lowest MAC address will be elected the root bridge.
2. Elects a root port for each non-root switch based on the lowest root path cost. The root bridge does not have root ports. Each non-root switch has one root port. The root port shows the direction of the best path to the root bridge.
3. Elects a designated port for each segment based on the lowest root path cost. Each link will have one designated port.
4. The root ports and designated ports transition to the forwarding state, and the other ports stay in the blocking state.

STP path cost depends on the speed of the link. The table shows STP link costs.

Data rate	STP Cost (802.1D-1998)	STP Cost (802.1D-2004)
4 Mbps	250	5,000,000
10 Mbps	100	2,000,000
16 Mbps	62	1,250,000
100 Mbps	19	200,000
1 Gbps	4	20,000
2 Gbps	3	10,000
10 Gbps	2	2000

STP Port Roles

Port Role	Description
Root port	This port exists on non-root bridges. It is the switch port with the best path to the root bridge. Root ports forward traffic toward the root bridge and the source MAC address of the frames received on the root port that is capable of populating the MAC table. Only one root port is allowed per bridge.
Designated port	This port exists on root and non-root bridges. For root bridges, all switch ports are designated ports. For non-root bridges, a designated port is the switch port that will receive and forward frames toward the root bridge as needed. Only one designated port is allowed per segment. If multiple switches exist on the same segment, an election process determines the designated switch, and the corresponding switch port begins forwarding frames for the segment. Designated ports are capable of populating the MAC table.
Nondesignated port	The nondesignated port is a switch port that is not forwarding (blocking) data frames and is not populating the MAC address table with the source addresses of frames that are seen on that segment.
Disabled port	The disabled port is a switch port that is shut down.

Spanning-Tree Operation Example

The first step in the spanning-tree algorithm is the election of a root bridge. Initially, all switches assume that they are the root. They start transmitting BPDUs with the Root ID field containing the same value as the Bridge ID field. Thus, each switch essentially claims that it is the root bridge on the network.

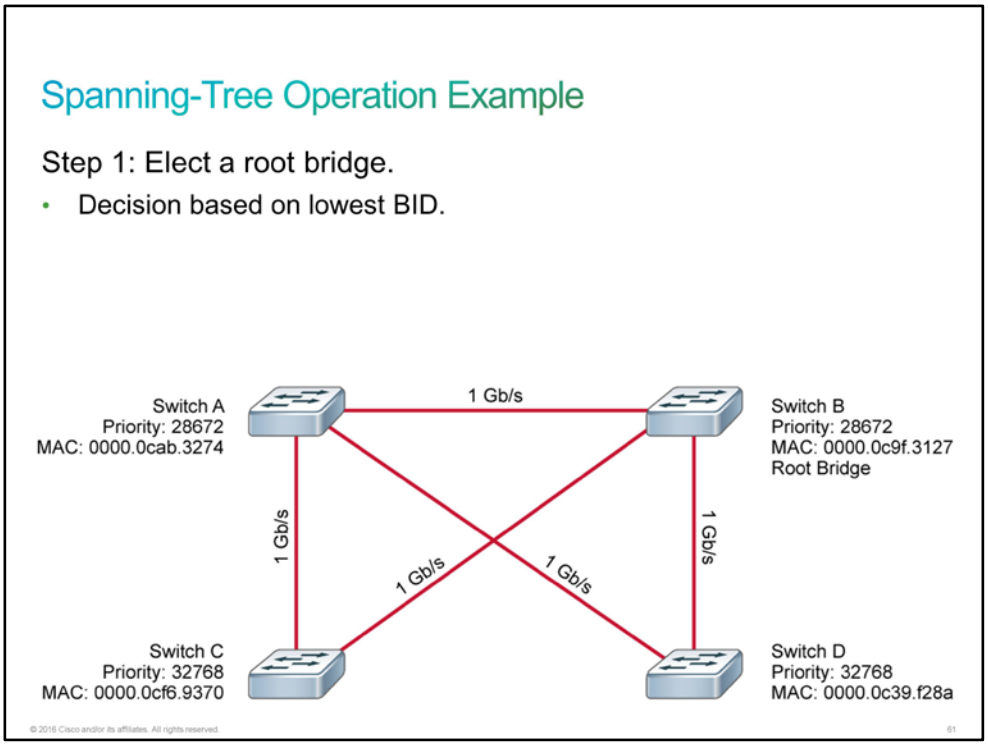

When the switches start receiving BPDUs from the other switches, each switch compares the root ID in the received BPDUs against the value that it currently has recorded as the root ID. If the received value is lower than the recorded value (which was originally the BID of that switch), the switch replaces the recorded value with the received value and starts transmitting this value in the Root ID field in its own BPDUs.

Eventually, all switches learn and record the BID of the switch that has the lowest BID. The switches all transmit this ID in the Root ID field of their BPDUs.

In the example, Switch B becomes the root bridge because it has the lowest BID. Switch A and switch B have the same priority, but switch B has a lower MAC value.

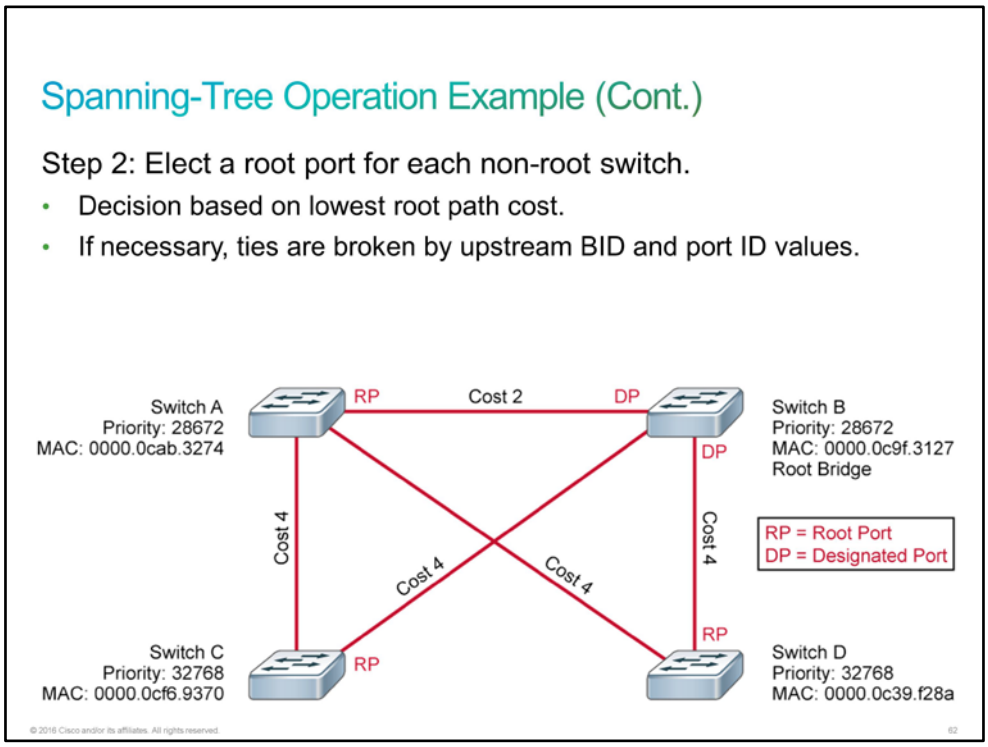

When a switch recognizes that it is not the root (because it is receiving BPDUs that have a root ID value that is lower than its own BID), it marks the port on which it is receiving those BPDUs as its root port.

A switch could receive BPDUs on multiple ports. In this case, the switch elects the port that has the lowest-cost path to the root as its root port. If two ports have an equal path cost to the root, the switch looks at the BID values in the received BPDUs to make a decision (where the lowest BID is considered best, similar to root bridge election). If the root path cost and the BID in both BPDUs are the same because both ports are connected to the same upstream switch, the switch looks at the Port ID field in the BPDUs and selects its root port based on the lowest value in that field.

By default, the cost that is associated with each port is related to its speed (the higher the interface bandwidth, the lower the cost), but the cost can be manually changed.

Switches A, C, and D mark the ports that are directly connected to switch B (which is the root bridge) as the root port. These directly connected ports on switches A, C, and D have the lowest cost to the root bridge.

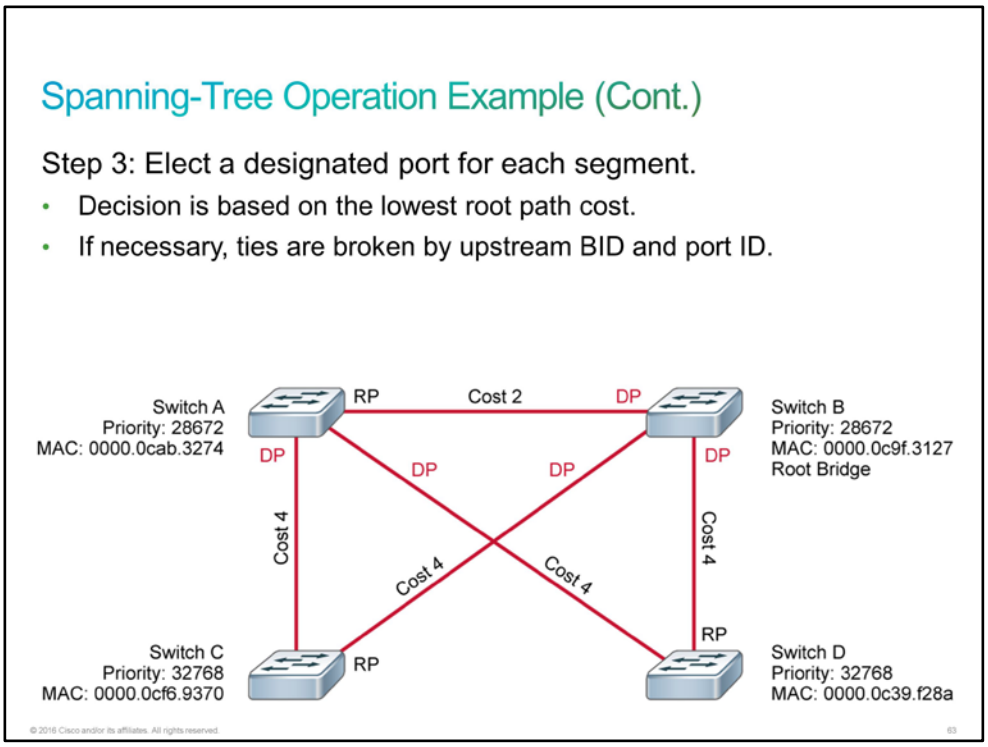

After electing the root bridge and root ports, the switches determine which switch will become the designated bridge for each Ethernet segment. This process is similar to the root bridge and root port elections. Each switch that is connected to a segment sends BPDUs out of the port that is connected to that segment, claiming to be the designated bridge for that segment. At this point, it considers its port to be a designated port.

When a switch starts receiving BPDUs from other switches on that segment, it compares the received values of the root path cost, BID, and port ID fields (in that order) against the values in the BPDUs that it is sending out its own port. The switch stops transmitting BPDUs on the port and marks it as a nondesignated port if the other switch has lower values.

In the example, all ports on the root bridge (switch B) are designated ports. The ports on switch A that are connecting to switch C and switch D become designated ports, because they have lower root path costs on each segment.

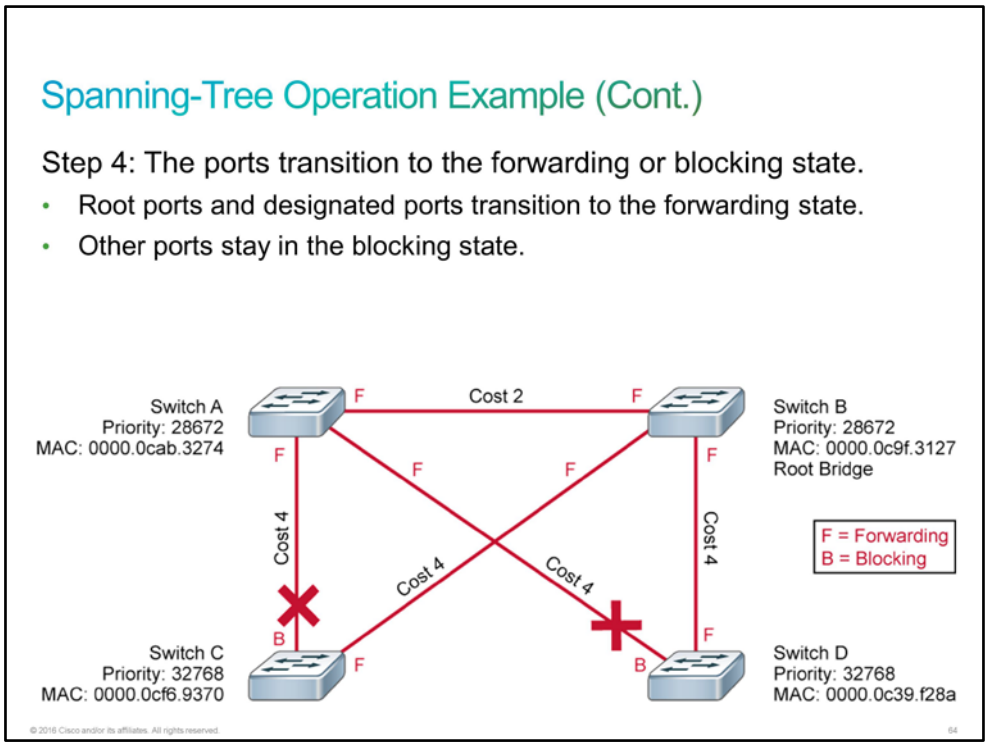

To prevent bridging loops while STP needs to execute its algorithm, all ports start out in the blocking state. When STP marks a port as either a root port or a designated port, the algorithm starts to transition this port to the forwarding state.

Classic (802.1D-1998) and rapid (802.1w and 802.1D-2004) versions of STP both execute the same algorithm in the decision-making process. However, in the transition of a port from the blocking (or discarding, in rapid spanning-tree terms) to the forwarding state, there is a big difference between those two spanning-tree versions. Classic 802.1D would simply take 30 seconds to transition the port to forwarding. The rapid spanning tree algorithm can leverage additional mechanisms to transition the port to forwarding in less than a second.

Although the order of the steps that are listed in the diagrams suggests that STP goes through them in a coordinated, sequential manner, that is not actually the case. If you look back at the description of each step in the process, you see that each switch is going through these steps in a parallel line. Also, each switch might adapt its selection of root bridge, root ports, and designated ports as it receives new BPDUs. As the BPDUs are propagated through the network, all switches eventually have a consistent view of the topology of the network. When this stable state is reached, BPDUs are transmitted only by designated ports.

There are two loops in the sample topology, meaning that two ports should be in the blocking state to break both loops. The port on Switch C that is not directly connected to Switch B (root bridge) is blocked, because it is a nondesignated port. The port on Switch D that is not directly connected to Switch B (root bridge) is also blocked, because it is a nondesignated port.

Types of Spanning-Tree Protocols

The STP is a network protocol that ensures a loop-free topology. Several varieties of spanning-tree protocols exist.

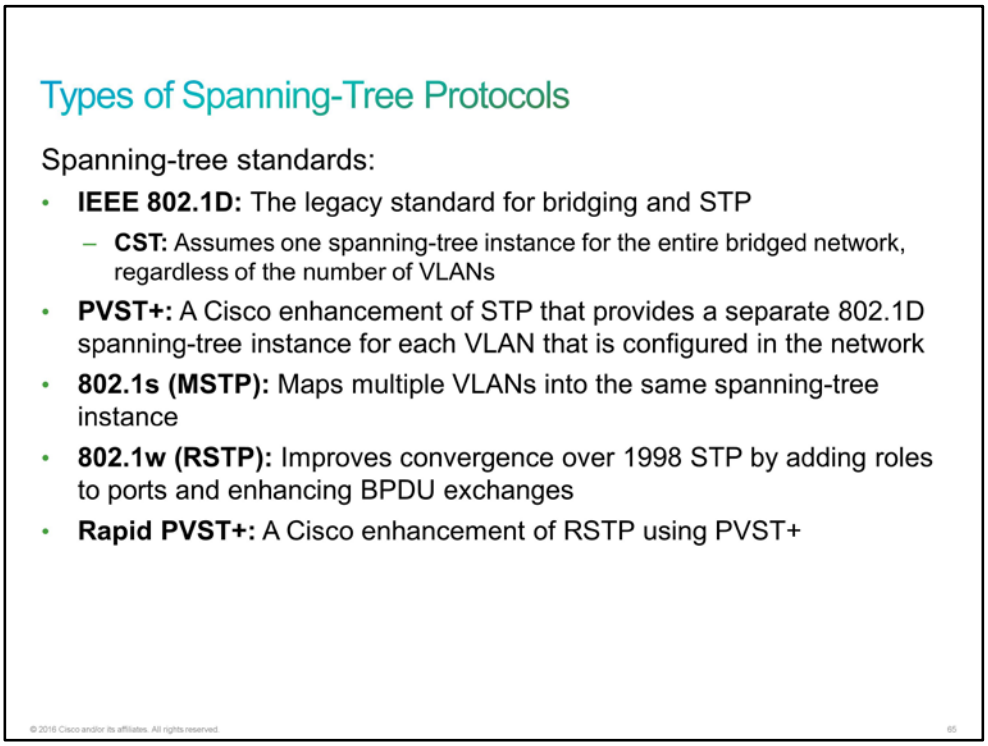

- STP (IEEE 802.1D) provides a loop-free topology in a network with redundant links.
  - CST assumes one spanning-tree instance for the entire bridged network, regardless of the number of VLANs.
- PVST+ is a Cisco enhancement of STP that provides a separate 802.1D spanning-tree instance for each VLAN that is configured in the network.
- MSTP, or IEEE 802.1s, is an IEEE standard that is inspired by the earlier Cisco proprietary MISTP implementation. MSTP maps multiple VLANs into the same spanning-tree instance.
- RSTP, or IEEE 802.1w, is evolution of STP that provides faster convergence of STP. It redefines port roles and link costs.
- Rapid PVST+ is a Cisco enhancement of RSTP that uses PVST+. Rapid PVST+ provides a separate instance of 802.1w per VLAN.

Note	When Cisco documentation and this course refer to implementing RSTP, they are referring to the Cisco RSTP implementation—Rapid PVST+.

Comparison of Spanning-Tree Protocols

The following are characteristics of various spanning-tree protocols:

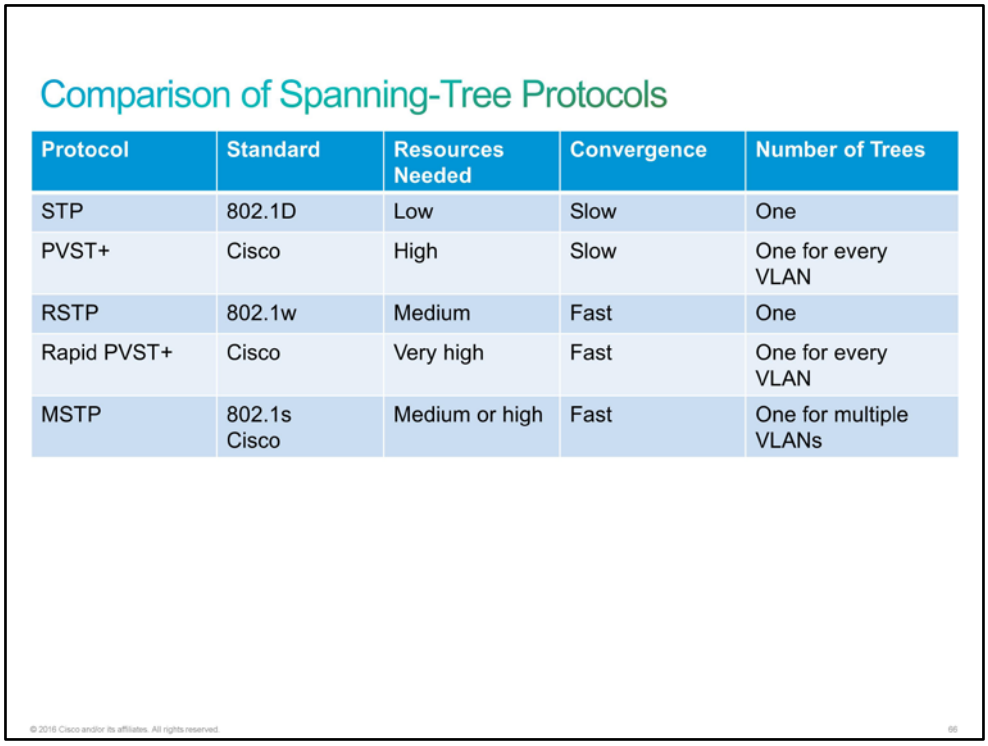

- STP assumes one 802.1D spanning-tree instance for the entire bridged network, regardless of the number of VLANs. Because only one instance exists, the CPU and memory requirements for this version are lower than for the other protocols. However, because of only one instance, there is only one root bridge and one tree. Traffic for all VLANs flows over the same patch, which can lead to suboptimal traffic flows. Because of the limitations of 802.1D, this version is slow to converge.
- PVST+ is a Cisco enhancement of STP that provides a separate 802.1D spanning-tree instance for each VLAN that is configured in the network. The separate instance supports PortFast, UplinkFast, BackboneFast, BPDU guard, BPDU filter, root guard, and loop guard. Creating an instance for each VLAN increases the CPU and memory requirements but allows for per-VLAN root bridges. This design allows the STP tree to be optimized for the traffic of each VLAN. Convergence of this version is similar to the convergence of 802.1D. However, convergence is per-VLAN.
- RSTP, or IEEE 802.1w, is evolution of STP that provides faster STP convergence. This version addresses many convergence issues, but because it still provides a single instance of STP, it does not address the suboptimal traffic flow issues. To support that faster convergence, the CPU usage and memory requirements of this version are slightly higher than the requirements of CST but lower than those of RSTP+.

- Rapid PVST+ is a Cisco enhancement of RSTP that uses PVST+. It provides a separate instance of 802.1w per VLAN. This version addresses both the convergence issues and the suboptimal traffic flow issues. However, this version has the largest CPU and memory requirements.
- MSTP is an IEEE standard that is inspired by the earlier Cisco proprietary MISTP implementation. To reduce the number of required STP instances, MSTP maps multiple VLANs that have the same traffic flow requirements into the same spanning-tree instance. The Cisco implementation of MSTP is MST. MST provides up to 16 instances of RSTP (802.1w) 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. The CPU and memory requirements of this version are lower than the requirements of Rapid PVST+ but are higher than those of RSTP.

Default Spanning-Tree Configuration

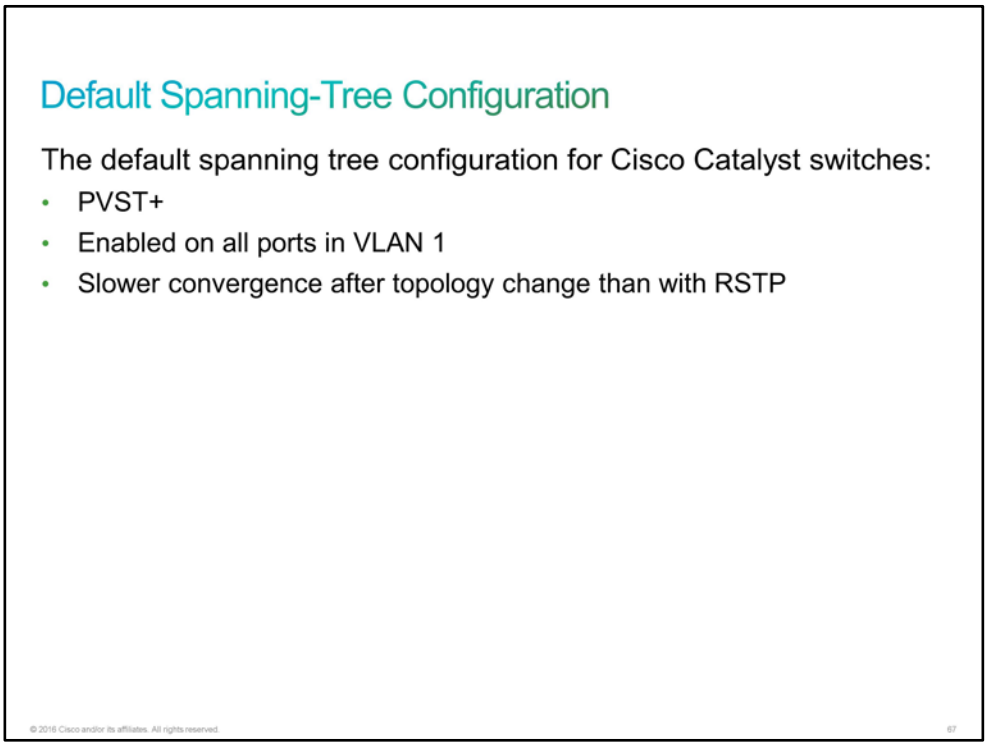

The default spanning-tree mode for Cisco Catalyst switches is PVST+, which is enabled on all ports. PVST+ has much slower convergence after a topology change than the Rapid PVST but requires less control plane CPU and memory resources to compute the shortest path tree upon topology changes.

Per VLAN Spanning Tree Plus

The 802.1D standard defines a CST that assumes only one spanning-tree instance for the entire switched network, regardless of the number of VLANs. A network that is running CST has these characteristics:

- No load sharing is possible. One uplink must block for all VLANs.
- The CPU is spared. Only one instance of spanning tree must be computed.

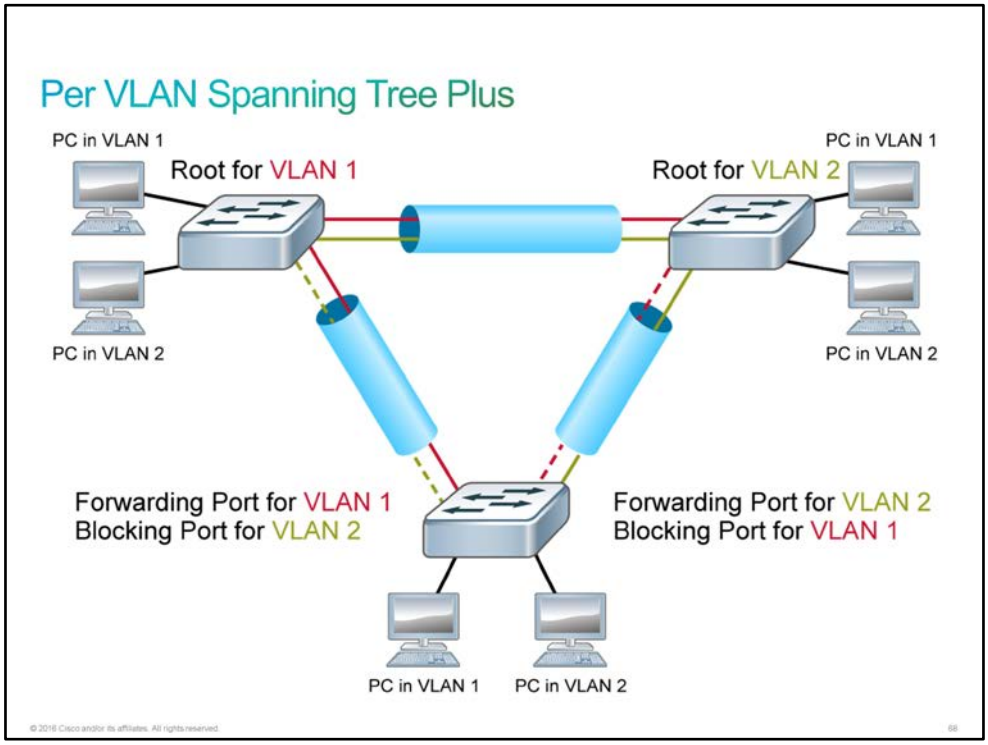

PVST+ defines a spanning-tree protocol that has several spanning-tree instances running for the network (one instance of STP per VLAN). Networks that are running several spanning-tree instances have these characteristics:

- Optimum load sharing can occur. In a Cisco PVST+ environment, you can tune the spanning-tree parameters so that half the VLANs forward on each uplink trunk. The configuration must define a different root bridge for each half of the VLANs. Providing different STP root switches per VLAN creates a more redundant network.
- One spanning-tree instance for each VLAN maintained can mean a considerable waste of CPU cycles for all the switches in the network (in addition to the bandwidth that is used for each instance to send its own BPDUs). This situation would only be problematic if many VLANs are configured.

Rapid PVST+ is Cisco proprietary version of the RSTP. It creates a spanning tree for each VLAN, just like PVST.

PVST+ Extended Bridge ID

Spanning-tree operation requires that each switch has a unique BID. In the original 802.1D standard, the BID consisted of the bridge priority and the MAC address of the switch, and a CST represented all VLANs. PVST+ requires that a separate instance of spanning tree that is run for each VLAN and the BID field must carry VID information. This functionality is accomplished by reusing a portion of the Priority field as the extended system ID to carry a VID.

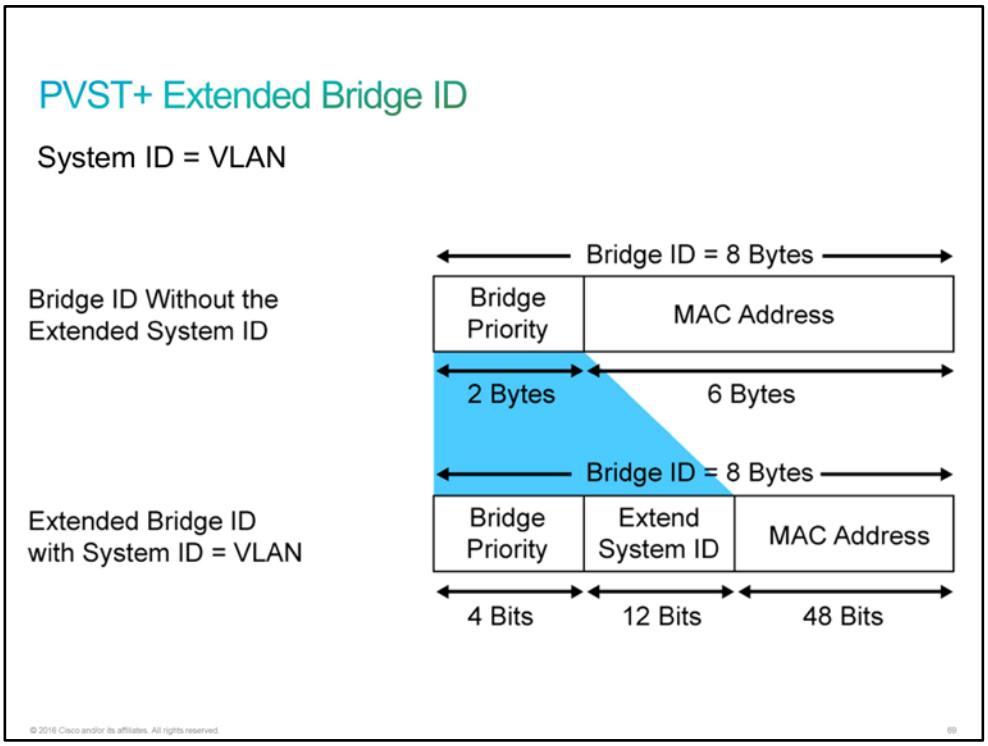

To accommodate the extended system ID, the original 802.1D 16-bit bridge priority field is split into two fields. The BID includes the following fields:

- **Bridge priority:** A 4-bit field that is still used to carry bridge priority. The priority is conveyed in discrete values in increments of 4096 rather than discrete values in increments of 1, because only the four most significant bits are available from the 16-bit field. In other words, in binary, the following applies: priority 0 = [0000|<sys-id-ext #>], priority 4096 = [0001|<sys-id-ext #>], and so on. Increments of 1 would be used if the complete 16-bit field was available. The default priority, in accordance with IEEE 802.1D, is 32768, which is the midrange value.
- **Extended system ID:** A 12-bit field carrying, in this case, the VID for PVST+. This value is expressed as the **sys-id-ext** in Cisco IOS software and elsewhere in this course.
- **MAC address:** A 6-byte field with the MAC address of a single switch.

By virtue of the MAC address, a BID is always unique. When the priority and extended system ID are prepended to the switch MAC address, each VLAN on the switch can be represented by a unique BID.

For example, the VLAN 2 default BID would be 32770 (priority 32768 plus the extended system ID of 2).

If no priority is configured, every switch will have the same default priority. In this case, the election of the root for each VLAN is based on the MAC address. This method is a random means of selecting the ideal root bridge. For this reason, it is recommended that you assign a lower priority to the switch that should serve as the root bridge.

Note	In the Cisco PVST+ environment, you can tune the spanning-tree parameters so that half the VLANs forward on each uplink trunk. The network must be correctly configured. The configuration must define a different root bridge for each half of the VLANs. Providing different STP root switches per VLAN creates a more redundant network.

Discovery 18: Configure Root Bridge and Analyze STP Topology

Introduction

The purpose of this discovery is to demonstrate how to determine the map of a spanning tree across a topology. The live virtual lab is prepared with the devices that are represented in the topology diagram and the connectivity table. All devices have their basic configurations in place, including hostnames and IP addresses.

During the discovery, you will map out the spanning tree for VLAN 20. SRV2 is the server on VLAN 20 and it is connected to SW4. You will observe that the spanning tree does not currently provide optimized paths from the clients to SRV2. You will then modify the spanning tree and verify the results.

Topology

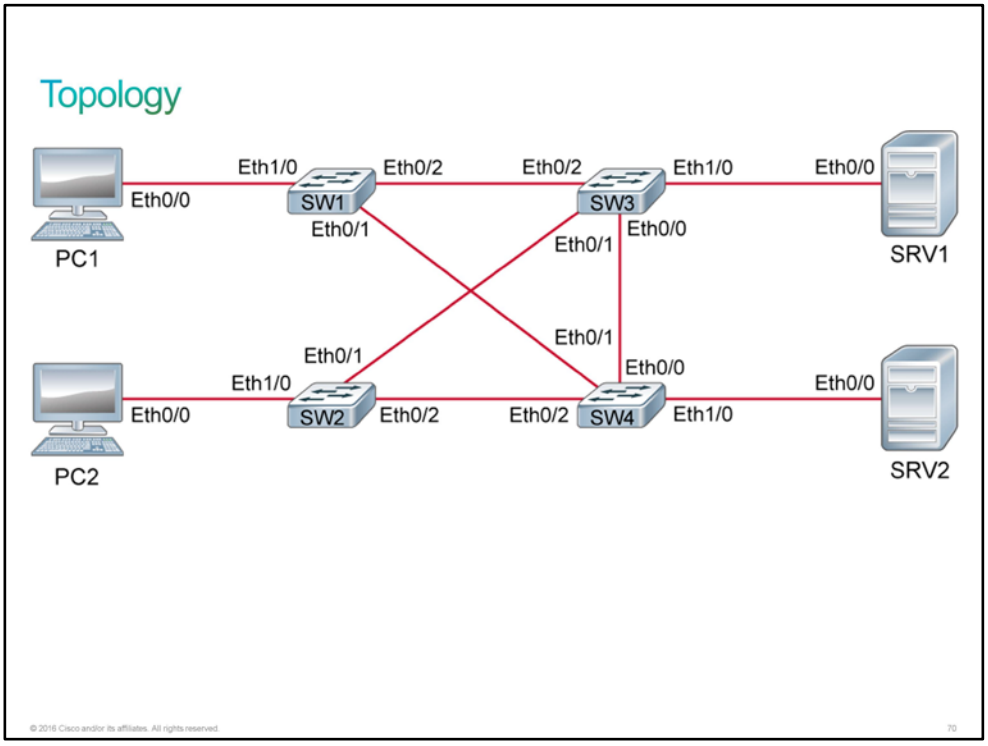

Job Aid

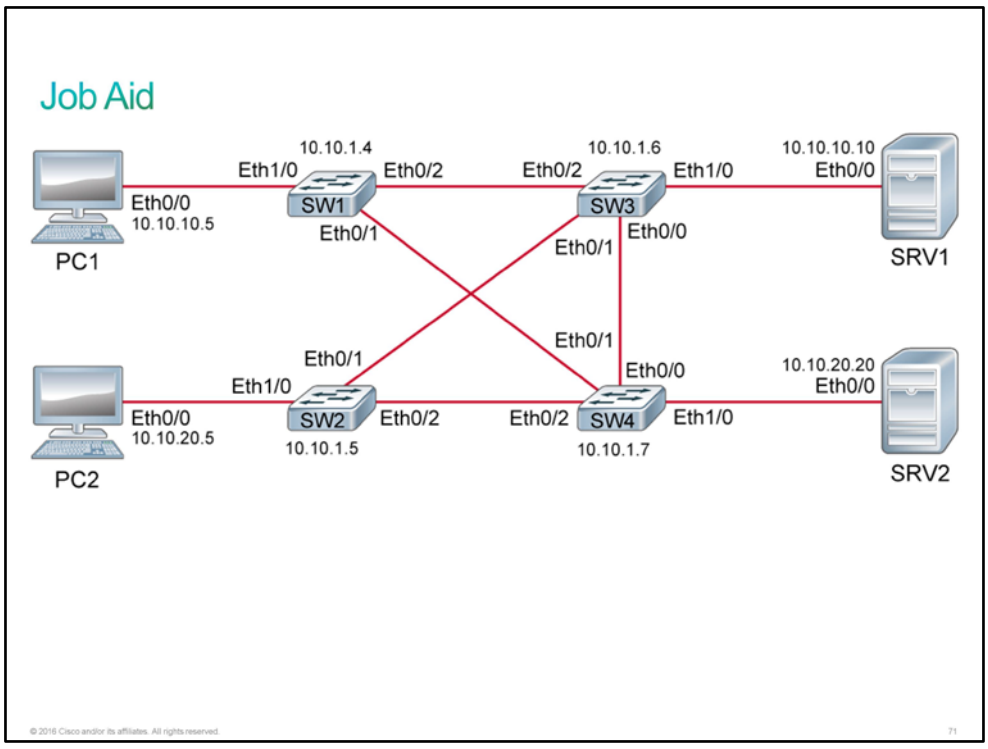

The configuration is as follows:

- All devices have their basic configurations in place, including hostnames and IP addresses.

Device Information

Device Details

Device	Interface	Neighbor	IP Address
PC1	Ethernet0/0	SW1	10.10.10.5/24
PC2	Ethernet0/0	SW2	10.10.20.5/24
SW1	VLAN 1	—	10.10.1.4/24
SW2	VLAN 1	—	10.10.1.5/24
SW3	VLAN 1	—	10.10.1.6/24
SW4	VLAN 1	—	10.10.1.7/24
SVR1	Ethernet0/0	SW3	10.10.10.10/24
SVR2	Ethernet0/0	SW4	10.10.20.20/24

Device Cabling Details

Switch	Port	Switch	Port
SW1	Ethernet0/1	SW4	Ethernet0/1
SW1	Ethernet0/2	SW3	Ethernet0/2
SW2	Ethernet0/1	SW3	Ethernet0/1
SW2	Ethernet0/2	SW4	Ethernet0/2
SW3	Ethernet0/0	SW4	Ethernet0/0

Note PCs and SRVs in the virtual lab environment are simulated as routers, so you should use Cisco IOS commands to configure them or make verifications.

Task 1: Modify the Bridge ID

By modifying the BID of a switch, you can influence the root bridge election.

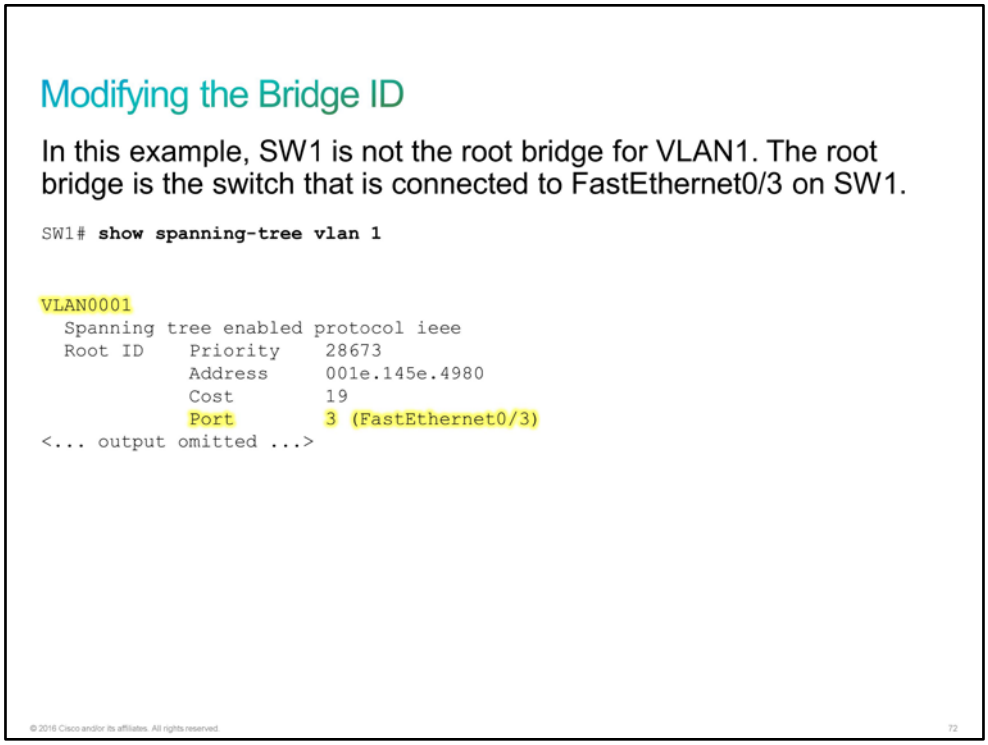

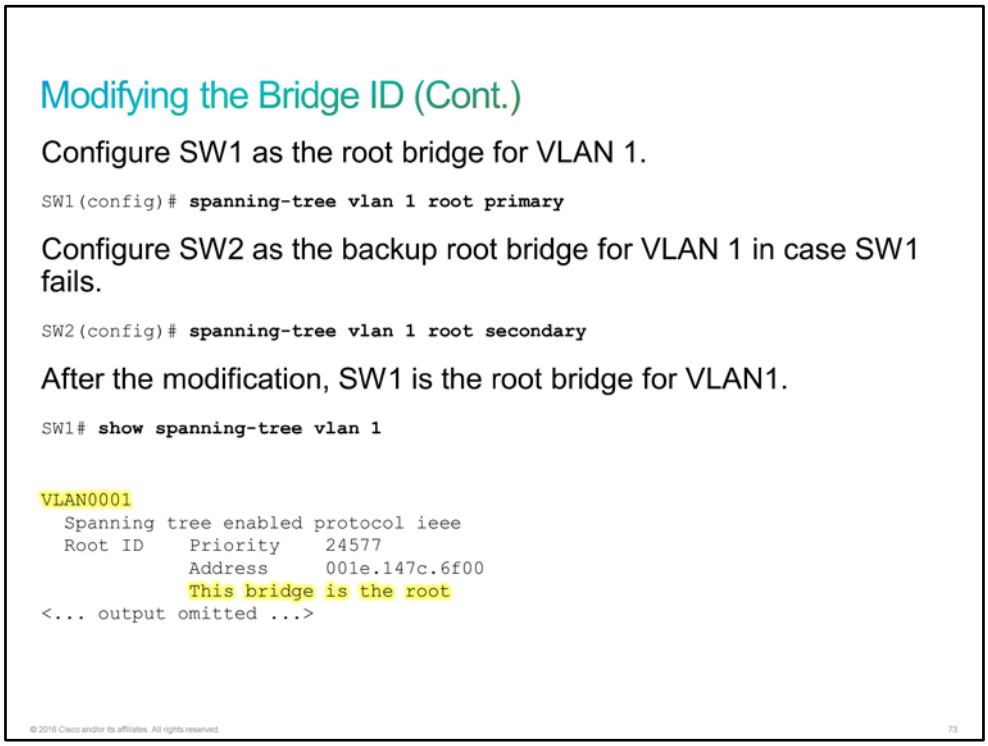

Command	Description
spanning-tree vlan *vlan_number* **root primary**	Forces this switch to be the root bridge for the specified VLAN.
spanning-tree vlan *vlan_number* **root secondary**	Configures the backup root bridge for the specified VLAN.

The root bridge is elected based on the BID. Since by default the priority part of the BID is the same for all switches (32768), the root bridge will be the switch with the lowest MAC address. For load balancing between switches (for example, if you want one switch to be the root bridge for VLAN 1 and the other switch to be the root bridge for VLAN 2), you can modify the priority of the bridge. The easiest way that you can make a switch the root bridge for a VLAN is if you use the **spanning-tree vlan** *vlan_number* **root primary** command. If the primary root bridge fails, you do not want the slowest, oldest access-layer switch becoming the root bridge. For this reason, you can configure the backup, secondary root bridge for a VLAN, if you use the **spanning-tree vlan** *vlan_number* **root secondary** command.

Complete the following step:

Step 1 On all four switches verify, if there is any spanning-tree preconfiguration for the root bridge.

```
SW1# show running-config | include root
SW1#

SW2# show running-config | include root
SW2#

SW3# show running-config | include root
SW3#

SW4# show running-config | include root
SW4#
```

There is no configuration for the root bridge on any of the switches, so the root bridge has been elected automatically.

Task 2: Analyze STP Topology

To analyze the STP topology, follow these steps:

1. Discover the physical Layer 2 topology. You could use network documentation, if it exists, or use the **show cdp neighbors** command to discover the physical topology.
2. After you have discovered the physical topology, use your knowledge of STP to determine the expected Layer 2 path. You will need to know which switch is the root bridge.
3. Use the **show spanning-tree vlan** command to determine which switch is the root bridge.
4. Use the **show spanning-tree vlan** command on all switches to find out which ports are in the blocking or forwarding state, and thus confirm your expected Layer 2 path.

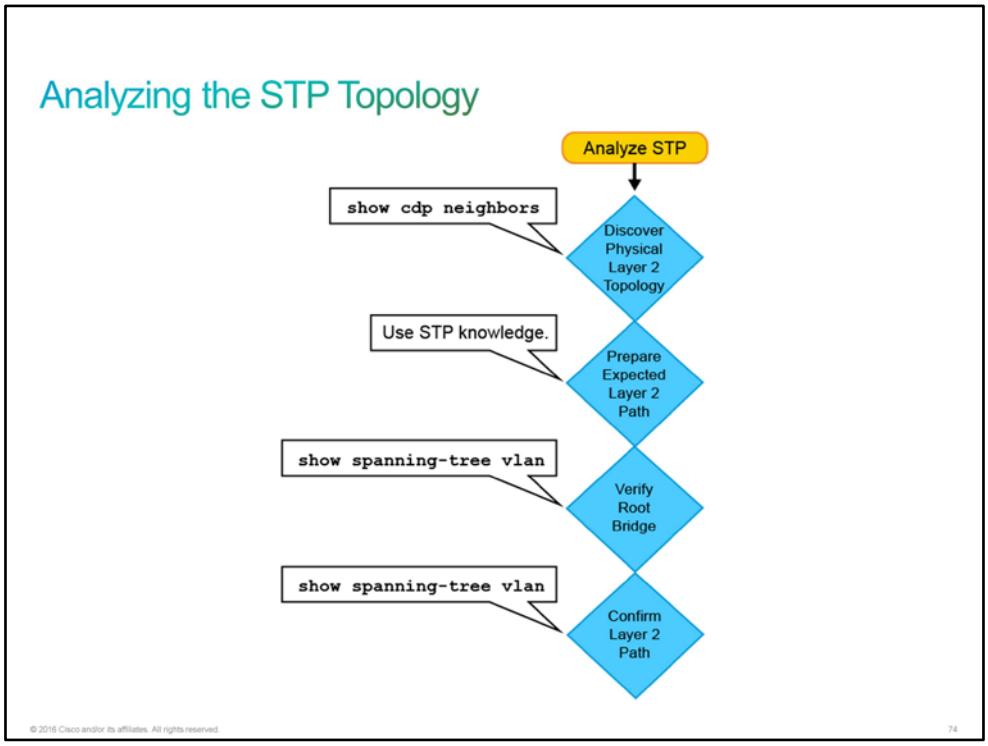

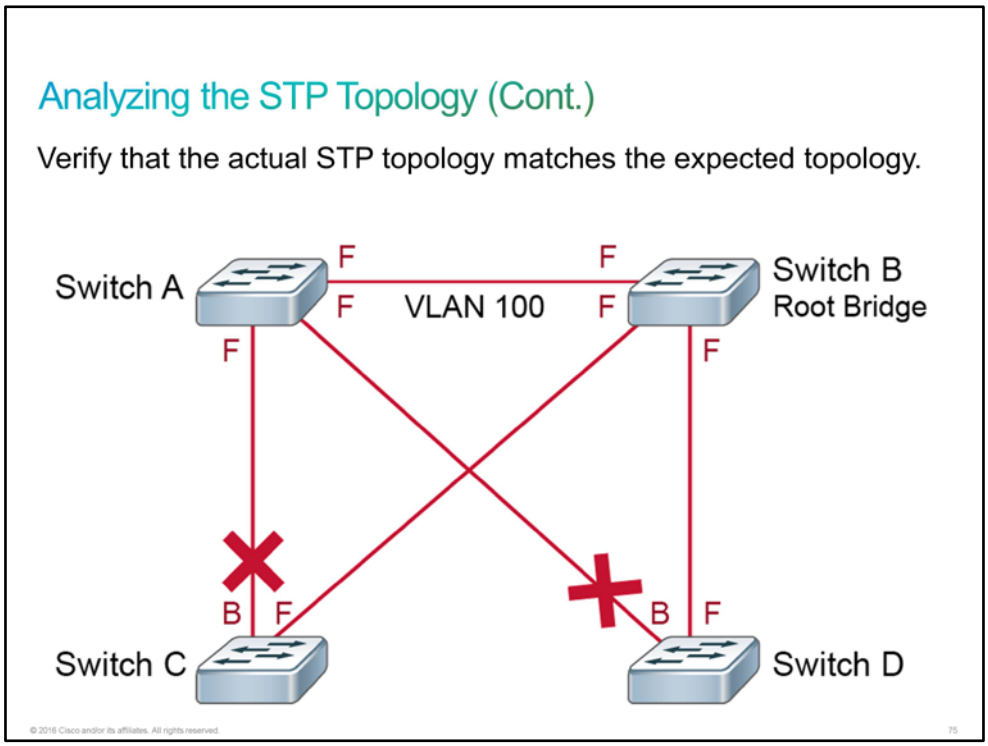

In many networks, the optimal STP topology is determined as part of the network design and then implemented through manipulation of STP priority and cost values. You might run into situations where STP was not considered in the design and implementation, or where it was considered initially, before the network underwent significant growth and change. In such situations, it is important that you to know how to analyze the actual STP topology in the operational network.

In addition, a part of troubleshooting also consists of comparing the actual state of the network against the expected state of the network. This way, you can spot the differences to gather clues about the problem that you are troubleshooting. You should be able to examine the switches and determine the actual topology, in addition to knowing what the spanning-tree topology is supposed to be.

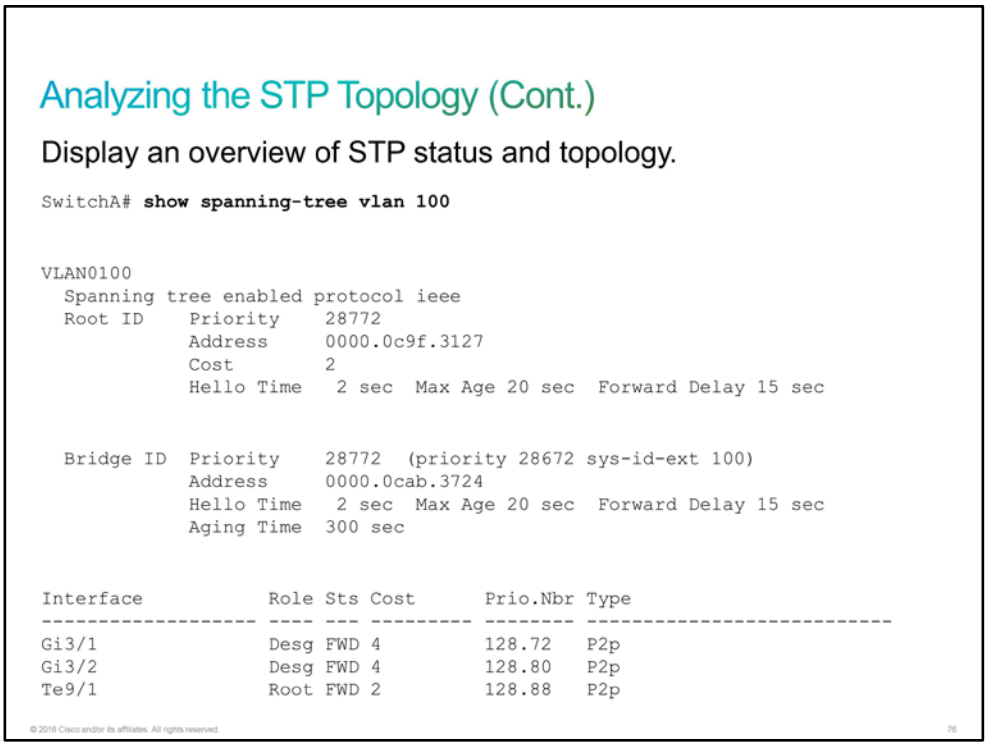

Using the **show spanning-tree** command without specifying any additional options is a good way to get a quick overview of the status of STP for all VLANs that are defined on a switch. If you are interested only in a particular VLAN, you can limit the scope of this command by specifying that VLAN as an option.

Use the **show spanning-tree vlan** *vlan_id* command to obtain STP information for a particular VLAN. Use this command to get information about the role and status of each port on the switch. The example output on Switch A shows all three ports in the forwarding state (FWD) and the role of the three ports as either designated ports or root ports. Any ports that are being blocked have the status "BLK" in the output.

The output also gives information about the BID of the local switch and the root ID. If Switch A is the root bridge, the root ID and bridge ID MAC addresses listed would be the same.

Activity

Complete the following steps:

Step 1 Begin to map out the spanning tree for VLAN 20. Start by accessing the console of SW1 and displaying the spanning-tree status for VLAN 20.

```
SW1# show spanning-tree vlan 20

VLAN0020
  Spanning tree enabled protocol ieee
  Root ID    Priority    32788
             Address     aabb.cc00.5400
             This bridge is the root
             Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec

  Bridge ID  Priority    32788  (priority 32768 sys-id-ext 20)
             Address     aabb.cc00.5400
             Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec
             Aging Time  300 sec

Interface           Role Sts Cost      Prio.Nbr Type
------------------- ---- --- --------- -------- --------------------------------
-
Et0/1               Desg FWD 100       128.2    Shr
Et0/2               Desg FWD 100       128.3    Shr
```

SW1 is the root switch for VLAN 20. The reason for this is the lowest Bridge ID. All switches have default priority set, so the MAC address decides which switch becomes the root bridge. SW1 has the lowest MAC address, hence resulting in SW1 becoming the root bridge.

Both Ethernet0/1 and 0/2 are designated (the port on the link that is closest to the root). Designated ports always forward.

Note: The MAC address might differ in your output.

Step 2 The topology, as you currently understand it, looks like the following example:

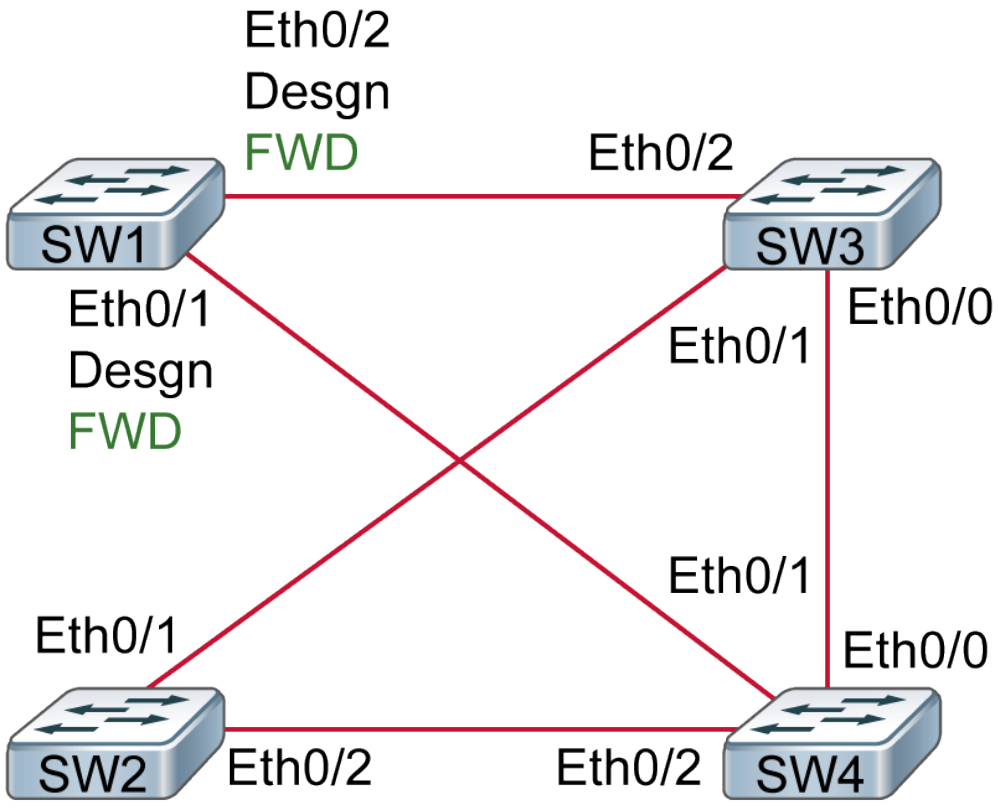

Step 3 Access the console of SW2 and display the spanning-tree status for VLAN 20.

Ethernet0/1 is the root port for SW2. That is the port that provides the lowest-cost path back to the root bridge. Root ports always forward.

```
SW2# show spanning-tree vlan 20

VLAN0020
  Spanning tree enabled protocol ieee
  Root ID    Priority    32788
             Address     aabb.cc00.5400
             Cost        200
             Port        2 (Ethernet0/1)
             Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec

  Bridge ID  Priority    32788  (priority 32768 sys-id-ext 20)
             Address     aabb.cc00.5500
             Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec
             Aging Time  300 sec

Interface           Role Sts Cost      Prio.Nbr Type
------------------- ---- --- --------- -------- --------------------------------
-
Et0/1               Root FWD 100       128.2    Shr
Et0/2               Altn BLK 100       128.3    Shr
Et1/0               Desg FWD 100       128.5    Shr
```

Ethernet0/2 is an alternate port for the link. There is another switch that provides forwarding back to the root for this link. Alternate ports are always in a blocking status. If they forwarded, it would cause a loop.

Ethernet1/0 is designated and forwarding. This is the port to which PC2 connects.

Note: The MAC address might differ in your output.

Step 4 The topology, as you currently understand it, looks like the following example:

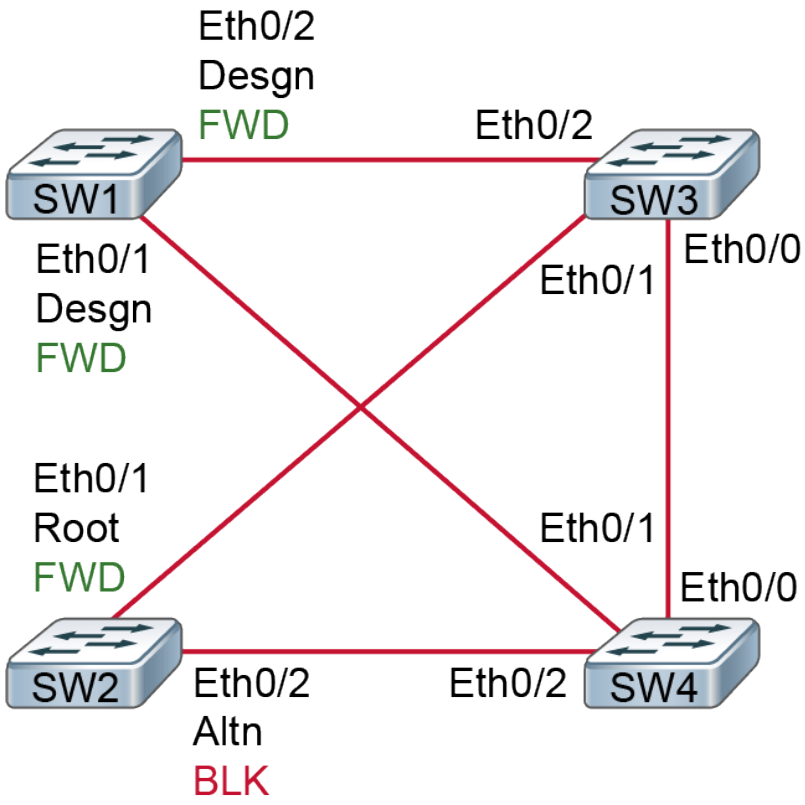

Step 5 Access the console of SW3 and display the spanning-tree status for VLAN 20.

Ethernet0/0 and 0/1 are designated and forwarding. Ethernet0/2 is the SW3 root port and is forwarding.

```
SW3# show spanning-tree vlan 20

VLAN0020
  Spanning tree enabled protocol ieee
  Root ID     Priority    32788
              Address     aabb.cc00.5400
              Cost        100
              Port        3 (Ethernet0/2)
              Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec

  Bridge ID   Priority    32788  (priority 32768 sys-id-ext 20)
              Address     aabb.cc00.5600
              Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec
              Aging Time  300 sec

Interface           Role Sts Cost      Prio.Nbr Type
------------------- ---- --- --------- -------- --------------------------------
-
Et0/0               Desg FWD 100       128.1    Shr
Et0/1               Desg FWD 100       128.2    Shr
Et0/2               Root FWD 100       128.3    Shr
```

You will not see the interface Ethernet1/0, the one connecting to the server, listed because it is on a different VLAN.

Note: The MAC address might differ in your output.

Step 6 The topology, as you currently understand it, looks like the following example:

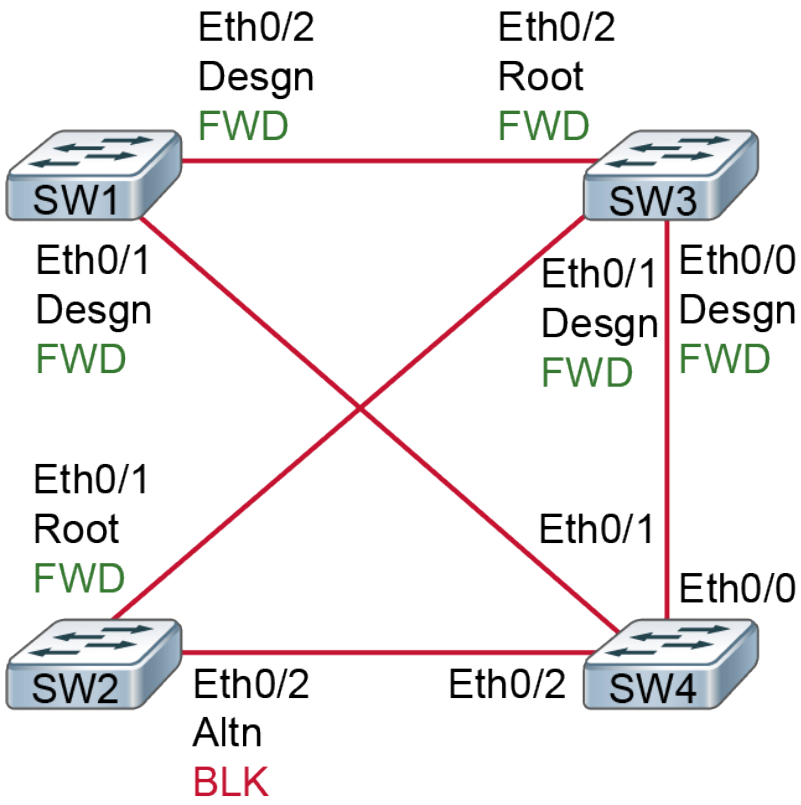

Step 7 Access the console of SW4 and display the spanning-tree status for VLAN 20.

Ethernet0/1 provides the best path back to the root switch. Therefore, it is the root port and is forwarding.

```
SW4# show spanning-tree vlan 20

VLAN0020
  Spanning tree enabled protocol ieee
  Root ID     Priority    32788
              Address     aabb.cc00.5400
              Cost        100
              Port        2 (Ethernet0/1)
              Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec

  Bridge ID   Priority    32788  (priority 32768 sys-id-ext 20)
              Address     aabb.cc00.5a00
              Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec
              Aging Time  300 sec

Interface           Role Sts Cost      Prio.Nbr Type
------------------- ---- --- --------- -------- --------------------------------
-
Et0/0               Altn BLK 100       128.1    Shr
Et0/1               Root FWD 100       128.2    Shr
Et0/2               Desg FWD 100       128.3    Shr
Et1/0               Desg FWD 100       128.5    Shr
```


Ethernet0/2 and 1/0 are both designated and forwarding. Ethernet1/0 is the port to which SRV2 connects.

Ethernet0/0 is alternate and blocking.

Note: The MAC address might differ in your output.

Step 8 The topology, as you currently understand it, looks like the following example:

With SW1 as the root switch, the spanning-tree topology is optimized for all switches to provide a best path back to SW1.

PC2 is attached to SW2 and SRV2 is attached to SW4. With the spanning tree in this state, the link between SW2 and SW4 is blocked. Traffic from PC2 to SRV2 must travel through SW2, then SW3, then SW1, and then finally to SW4.

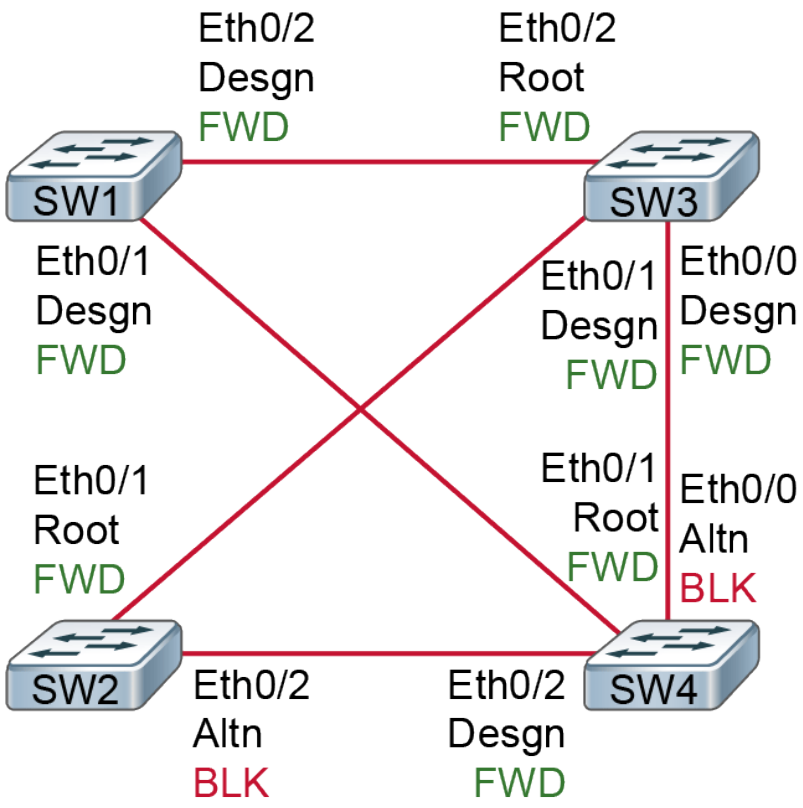

Step 9 The single most important spanning-tree tuning operation that you should do is to set the root switch for a VLAN to be the switch where most of the traffic on that VLAN is destined. Usually, this is a switch to which routers or servers are connected. Configure SW4 to be the root switch for VLAN 20.

Enter the following commands to the SW4 switch:

```
SW4# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
SW4(config)# spanning-tree vlan 20 root primary
SW4(config)# end
SW4#
*Sep 18 09:01:21.569: %SYS-5-CONFIG_I: Configured from console by console
SW4#
```

Step 10 Display the spanning-tree status for VLAN 20 on SW4.

The Forward Delay is 15 seconds, by default. A blocking port must transition through listening for 15 seconds and learning for 15 seconds before proceeding to forwarding.

```
SW4# show spanning-tree vlan 20

VLAN0020
  Spanning tree enabled protocol ieee
  Root ID    Priority    24596
             Address     aabb.cc00.5a00
             This bridge is the root
             Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec

  Bridge ID  Priority    24596  (priority 24576 sys-id-ext 20)
             Address     aabb.cc00.5a00
             Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec
             Aging Time  300 sec

Interface           Role Sts Cost      Prio.Nbr Type
------------------- ---- --- --------- -------- --------------------------------
-
Et0/0               Desg LRN 100       128.1    Shr
Et0/1               Desg FWD 100       128.2    Shr
Et0/2               Desg FWD 100       128.3    Shr
Et1/0               Desg FWD 100       128.5    Shr
```

Step 11 If all four ports are not yet in the forwarding state, continue to execute the **show spanning-tree** command until they are forwarding.

```
SW4# show spanning-tree vlan 20

VLAN0020
  Spanning tree enabled protocol ieee
  Root ID    Priority    24596
             Address     aabb.cc00.5a00
             This bridge is the root
             Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec

  Bridge ID  Priority    24596  (priority 24576 sys-id-ext 20)
             Address     aabb.cc00.5a00
             Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec
             Aging Time  300 sec

Interface           Role Sts Cost      Prio.Nbr Type
------------------- ---- --- --------- -------- --------------------------------
-
Et0/0               Desg FWD 100       128.1    Shr
Et0/1               Desg FWD 100       128.2    Shr
Et0/2               Desg FWD 100       128.3    Shr
Et1/0               Desg FWD 100       128.5    Shr
```

Step 12 For the backup purposes, also configure SW3 as the secondary root bridge for VLAN 20.


```
SW3# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
SW3(config)# spanning-tree vlan 20 root secondary
SW3(config)# end
SW3#
*Sep 18 09:01:21.569: %SYS-5-CONFIG_I: Configured from console by console
SW3#
```

Step 13 Optionally, you can view the spanning tree details for VLAN 20 on the other three switches. The topology that would emerge would look like the following.

The spanning tree for VLAN 20 is now optimized to provide an optimal path from all switches to SW4, where SRV2 on VLAN 20 is connected.

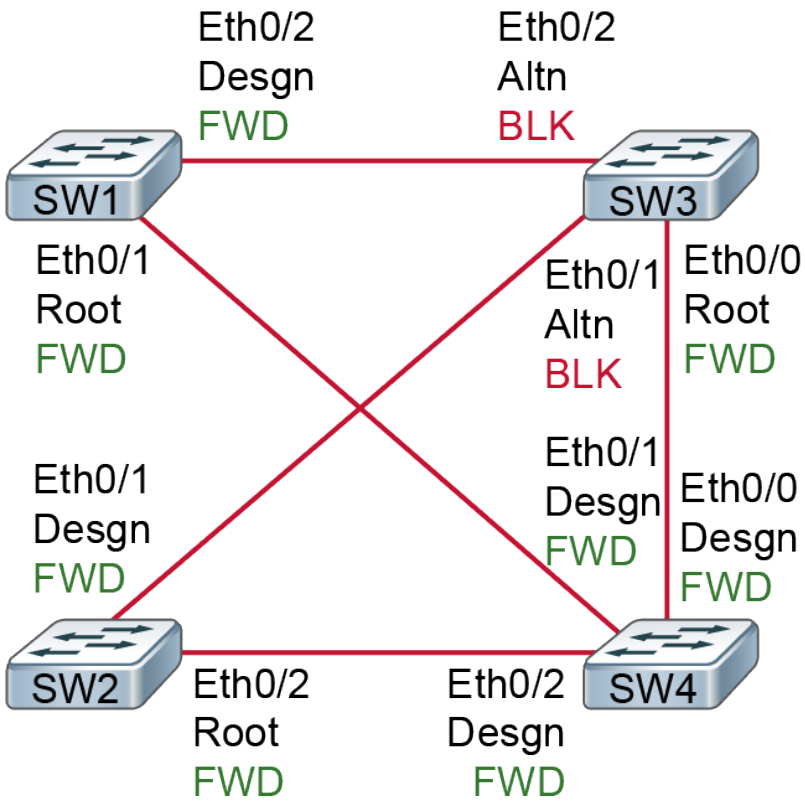

This is the end of the discovery lab.

PortFast and BPDU Guard

You will explore two features that called PortFast and BPDU guard. Before you can fully appreciate the benefits of these features, you will review the STP initialization process that a switch port transitions through when it is enabled.

Because STP is responsible for maintaining a loop-free topology, precautions are required each time that you enable a switch port. If the port is connected to another switch, BPDUs are exchanged to ensure that a loop is not introduced into the topology. The following are the stages that a port goes through when it is enabled.

1. **Blocking:** For up to 20 seconds, the port remains in the blocking state.
2. **Listening:** For 15 seconds, the port listens to BPDUs that it received and listens for new topology information that would cause it to transition back to the blocking state. It does not populate the MAC address table with the addresses it learns and it does not forward any frames.
3. **Learning:** For up to 15 seconds, the port updates the MAC address forwarding table, but it does not begin forwarding.
4. **Forwarding:** Once the switch port is certain it will not form a loop by forwarding frames, it enters the forwarding state. It still monitors for topology changes that could require it to transition back to the blocking state to prevent a loop.

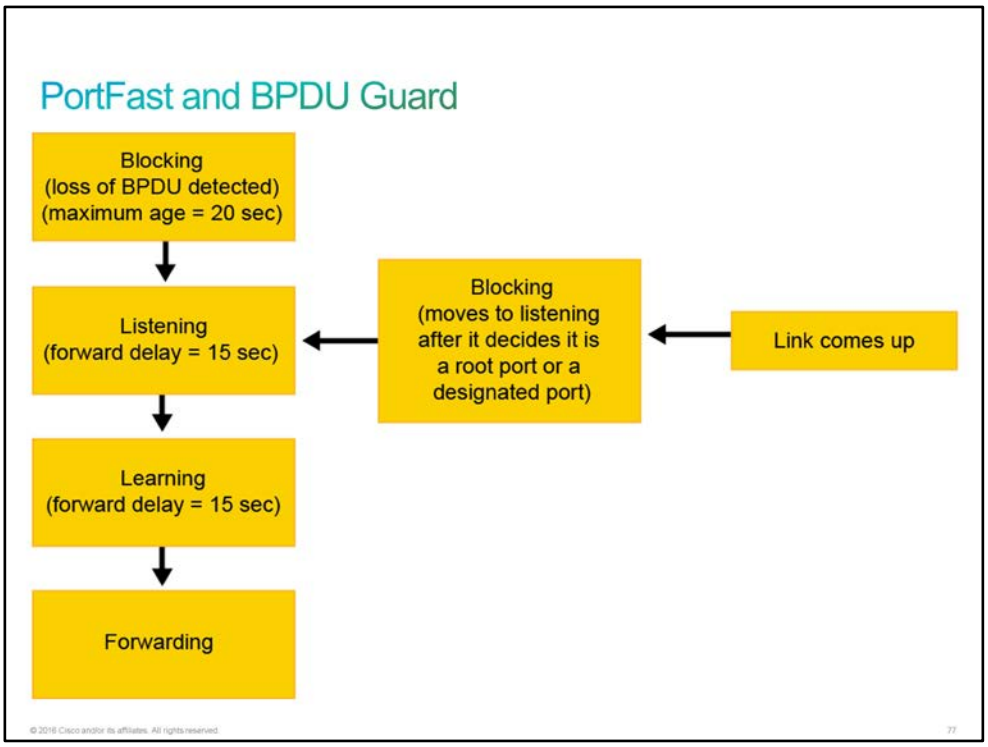

If a switch port connects to another switch, the STP initialization cycle must transition from state to state to ensure a loop-free topology.

However, for access devices such as PCs, laptops, servers, and printers, the delays that incurred with STP initialization can cause problems such as DHCP timeouts. Cisco designed the PortFast and BPDU features as enhancements to STP to reduce the time that is required for an access device to enter the forwarding state.

STP is designed to prevent loops. Because there can be no loop on a port that is connected directly to a host or server, the full function of STP is not needed for that port. PortFast is a Cisco enhancement to STP that allows a switchport to begin forwarding much faster than a switchport in normal STP mode.

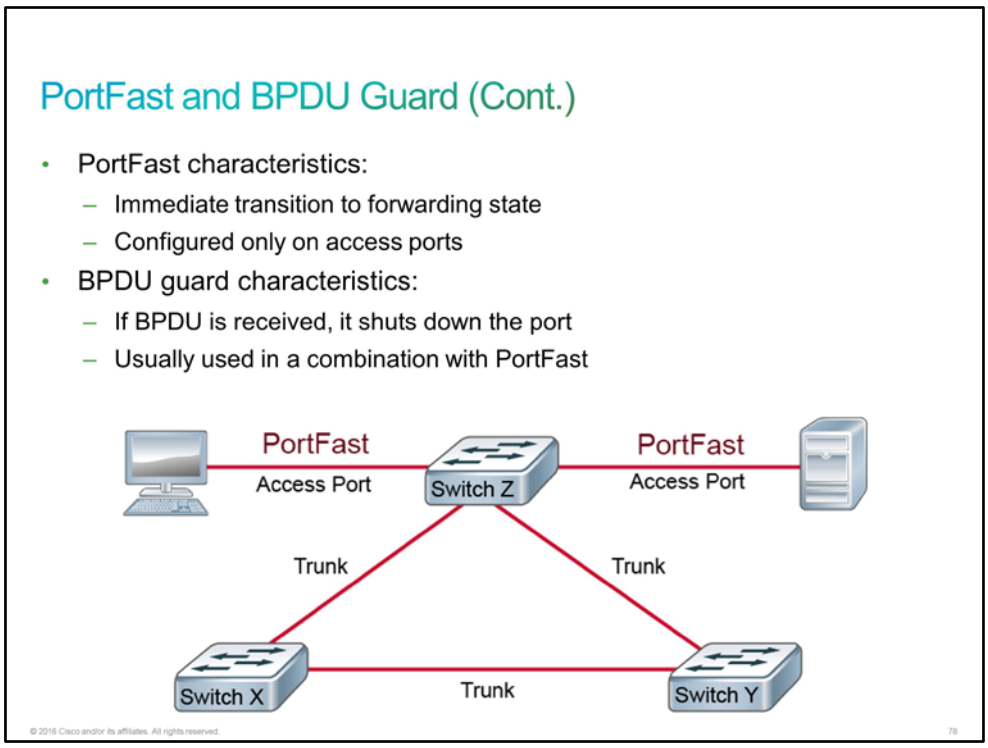

When the PortFast feature is enabled on a switch port that is configured as an access port, that port bypasses the typical STP listening and learning states. This feature allows the port to transitions from the blocking to the forwarding state immediately. You can use PortFast on access ports that are connected to a single workstation or to a server to allow those devices to connect to the network immediately rather than waiting for spanning tree to converge.

In a valid PortFast configuration, configuration BPDUs should never be received, because access devices do not generate BPDUs. A BPDU that a port receives would indicate that another bridge or switch is connected to the port. This event could happen if a user plugged a switch on their desk into the port where the user PC was already plugged into.

Assuming that users decide they want more bandwidth. Since there are two network access connections in their office, they decide to use both of them. To use them both, they unplug their individual PCs from the network switches and plug it into their own switch. They then plug the new switch into both of the network access ports. If portfast is enabled on both ports of the network switch, this action could cause a loop and bring the network to a halt.

If the users in the example have realized they are causing network issues. If the users only disconnect one of the links from their switch to network switch—have they eliminated *all* the issues they were causing? What would be the result if their switch had a lower BID than the root bridge in the network? Wouldn't their switch become the root bridge?

The STP PortFast BPDU guard enhancement allows network designers to enforce the STP domain borders and keep the active topology predictable. The devices behind the ports that have STP PortFast enabled are not able to influence the STP topology. At the reception of BPDUs, the BPDU guard operation disables the port that has PortFast configured. The BPDU guard transitions the port into errdisable state, and a message appears on the console. For example, the following message might appear:

```
2000 May 12 15:13:32 %SPANTREE-2-RX_PORTFAST:Received BPDU on PortFast enable port.
Disabling 2/1
2000 May 12 15:13:32 %PAGP-5-PORTFROMSTP:Port 2/1 left bridge port 2/1
```

Note	Because the purpose of PortFast is to minimize the time that access ports that are connecting to user equipment and servers must wait for spanning tree to converge, you should use it only on access ports. If you enable PortFast on a port that is connecting to another switch, you risk creating a spanning-tree loop.

Configuring PortFast and BPDU Guard

PortFast and BPDU guard can be configured on a port-by-port basis, or globally for all ports on a switch.

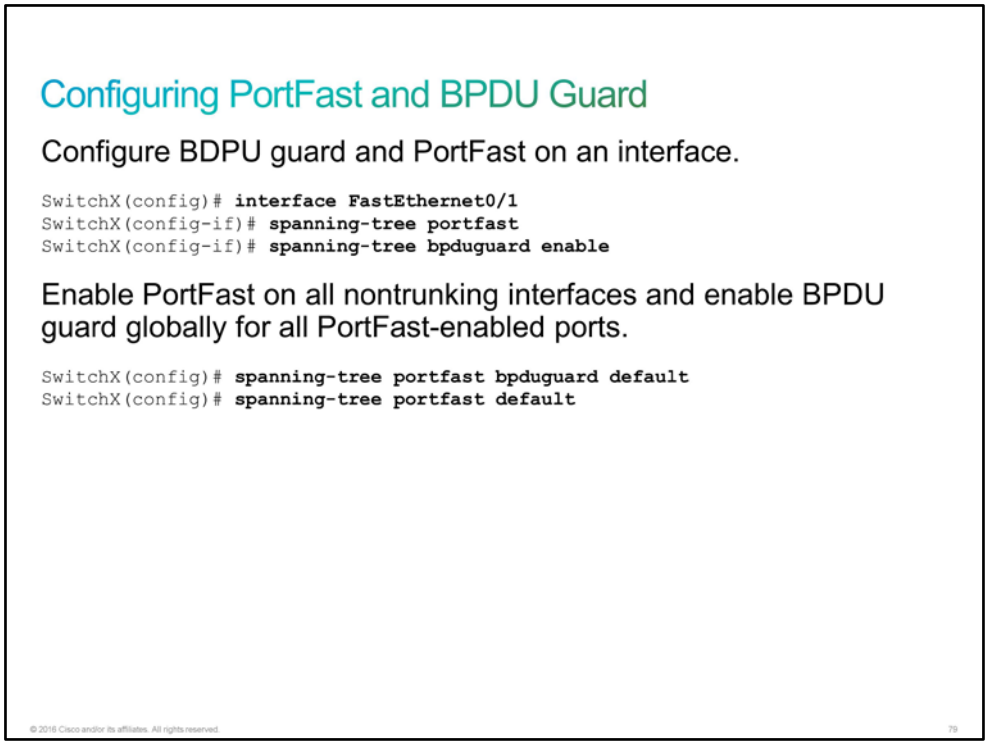

The **spanning-tree bpduguard enable** interface configuration command configures BPDU guard on an interface. The **spanning-tree portfast bpduguard default** global configuration command enables BPDU guard globally for all PortFast-enabled ports.

The **spanning-tree portfast** interface configuration command configures PortFast on an interface. The **spanning-tree portfast default** global configuration command enables PortFast on all nontrunking interfaces.

Verifying PortFast and BPDU Guard

Use the **show running-config interface** command to validate the PortFast and BPDU guard configuration for a given interface.

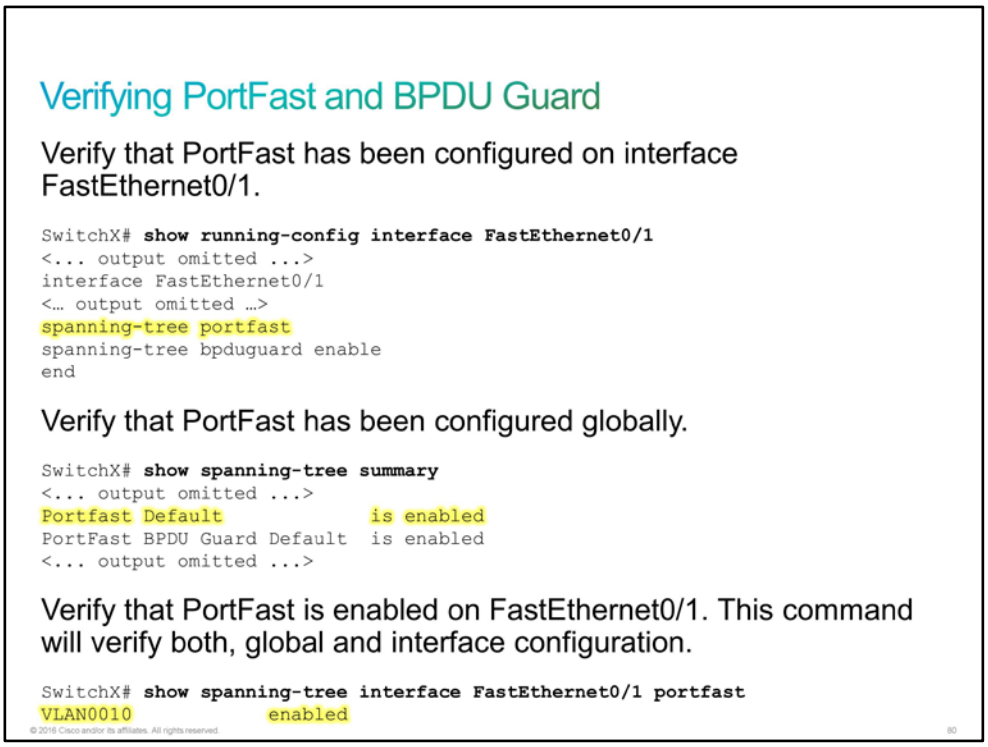

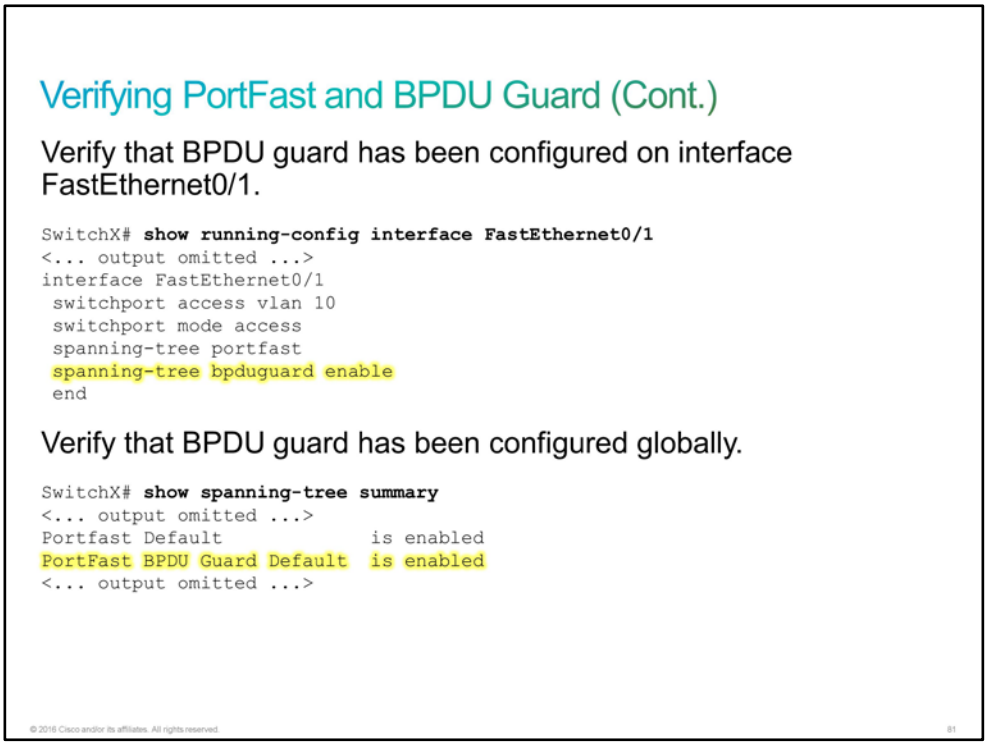

The table lists the commands that you use to implement and verify PortFast and BPDU guard.

PortFast and BPDU Guard Commands

Command	Description
spanning-tree portfast	Enables PortFast on a Layer 2 access port and forces it to enter the forwarding state immediately.
spanning-tree portfast default	Globally enables the PortFast feature on all nontrunking ports. When the PortFast feature is enabled, the port changes from a blocking state to a forwarding state without making the intermediate spanning-tree state changes.
spanning-tree bpduguard enable	Enables BPDU guard on a Layer 2 access port.
spanning-tree portfast bpduguard default	Globally enables the BPDU guard feature.
show running-config interface *typeslot/port*	Indicates whether PortFast and BPDU guard have been configured on a port.
show spanning-tree interface *typeslot/port* **portfast**	Indicates whether PortFast has been configured on a port. The command verifies both global and interface configuration.
show spanning-tree summary	Indicates whether PortFast and BPDU guard have been configured globally.

Note When you enable the PortFast feature globally, you will not see it under the interface configuration using **show running-config interface** *type slot/port* command. For this case, you should use **show spanning-tree interface** *type slot/port* **portfast** command or **show spanning-tree summary** command.

Discovery 19: Troubleshoot STP Issues

Introduction

The biggest problem with STP is not the fact that it can fail, because any protocol can. In fact, STP is one of the most reliable protocols available. The main concern is that when a problem that is related to STP exists, there are usually major negative consequences. Unlike with many protocols, where the only thing that happens when a malfunction occurs is that you lose some of the functionality that you gained through this protocol. For instance, if the routing protocol is malfunctioning on one of your routers, you might lose connectivity to networks that are reachable through that particular router. However, this loss generally does not affect the rest of your network. If you have some way to connect to that router, you can still perform your troubleshooting routines to diagnose and fix the problem.

In this discovery, you will demonstrate how to use different commands to troubleshoot STP.

Topology

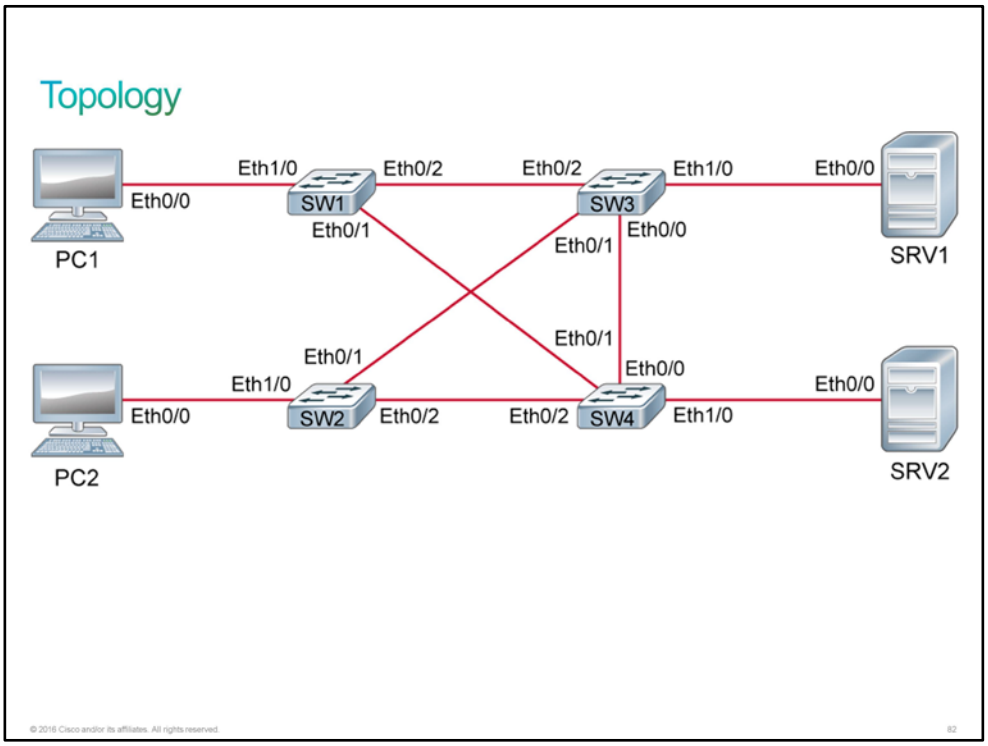

Job Aid

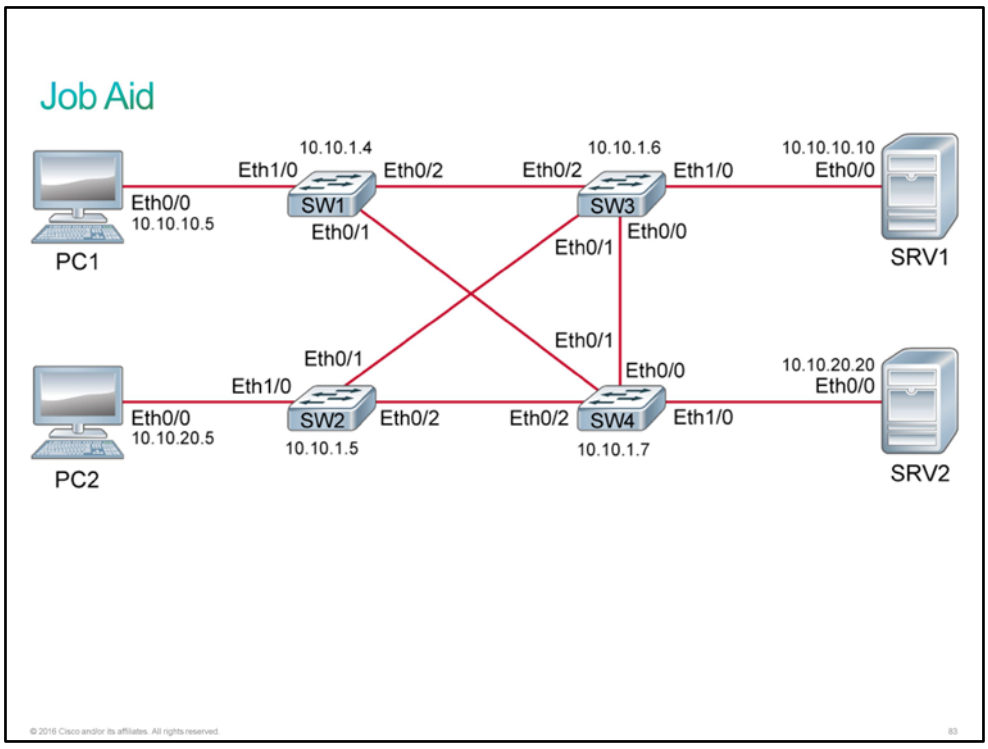

The configuration is as follows:

- All devices have their basic configurations in place, including hostnames and IP addresses.

Device Information

Device Details

Device	Interface	Neighbor	IP Address
PC1	Ethernet0/0	SW1	10.10.10.5/24
PC2	Ethernet0/0	SW2	10.10.20.5/24
SW1	VLAN 1	—	10.10.1.4/24
SW2	VLAN 1	—	10.10.1.5/24
SW3	VLAN 1	—	10.10.1.6/24
SW4	VLAN 1	—	10.10.1.7/24
SVR1	Ethernet0/0	SW3	10.10.10.10/24
SVR2	Ethernet0/0	SW4	10.10.20.20/24

Device Cabling Details

Switch	Port	Switch	Port
SW1	Ethernet0/1	SW4	Ethernet0/1
SW1	Ethernet0/2	SW3	Ethernet0/2
SW2	Ethernet0/1	SW3	Ethernet0/1
SW2	Ethernet0/2	SW4	Ethernet0/2
SW3	Ethernet0/0	SW4	Ethernet0/0

Note PCs and SRVs in the virtual lab environment are simulated as routers, so you should use Cisco IOS commands to configure them or make verifications.

Task 1: Spanning-Tree Failure Consequences

With STP, you can observe two different types of failures:

- The first one is similar to the routing problem that was just described. STP may erroneously block certain ports that should have gone to the forwarding state. This block will cause problems that are similar to the routing problem: You might lose connectivity to certain parts of your network, but the rest of the network is unaffected. If you are able to access the switch, you can troubleshoot and attempt to resolve the issue.
- The second type of failure is when STP erroneously moves one or more ports to the forwarding state. This type of failure can be very disruptive.

Spanning-Tree Failure Consequences

What will happen to this network if Switch D erroneously transitions both its ports to the forwarding state?

- Any frame that enters a bridging loop will continue to be forwarded by the switches indefinitely.

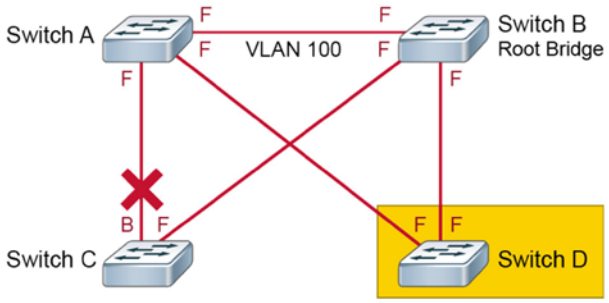

An Ethernet frame header does not include a TTL field. Therefore, any frame that enters a bridging loop will continue to be forwarded by the switches indefinitely. The only exceptions are the frames whose destination addresses are recorded in the MAC address table of the switches. These frames will be forwarded to the port that the MAC address is associated with and will not go into an endless loop. However, any frame that is flooded by a switch (such as broadcasts, multicasts, and unicasts) with an unknown destination MAC address, will go into an endless loop.

Spanning-Tree Failure Consequences (Cont.)

The consequences of STP failure are severe.

- The load on all links in the switched LAN quickly starts increasing.
- Due to the very high load for the CPU, the switch becomes unreachable.

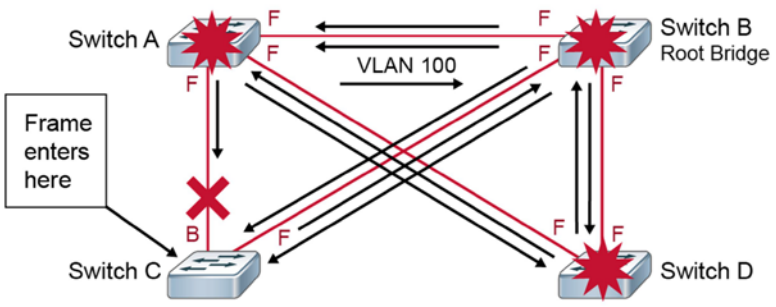

- The figure shows how the load on all links in the switched LAN quickly starts increasing as more frames enter the loop. This problem is not limited to the links that form the loop. The problem also affects any other links in the switched domain, because the frames are flooded on all links. When the spanning-tree failure is limited to a single VLAN, only links in that VLAN are affected. Switches and trunks that do not carry that VLAN operate normally.
- If the spanning-tree failure has caused more than one bridging loop, traffic increases exponentially, because frames not only start circling but also start getting duplicated. This problem happens because when you have multiple loops, you will also have switches that receive a frame on a port and then flood it out on multiple ports, essentially creating a copy of the frame every time they forward it.
- The switches will experience frequent MAC address table changes. This problem happens because frames usually start looping in both directions. This action causes a switch to see a frame with a certain source MAC address coming in on a port and then see a frame with the same source MAC address coming in on a different port just a fraction of a second later.
- Because of the combination of a very high load on all links and the switch CPUs running at maximum load, these devices typically become unreachable. As a result, diagnosing this problem while it is happening is nearly impossible.

A viable approach is to take over the role of the failing spanning tree by manually removing redundant links in the switched network, either physically or through configuration (if that is still possible), until all loops are eliminated from the topology. When you have broken the loops, the traffic and CPU loads should quickly drop to normal levels, and you should regain connectivity to your devices.

Although this intervention restores connectivity to the network, you cannot consider it the end of your troubleshooting process. You have removed all redundancy from your switched network, and you need to restore the redundant links.

Of course, if the underlying cause of the spanning-tree failure has not been fixed, chances are that restoring the redundant links will trigger a new broadcast storm. Before you restore the redundant links, you should spend sufficient time to investigate what happened at the moment when the broadcast storm started. When you eventually start restoring the redundant links, you should carefully monitor the network and have an emergency plan to fall back on if you see a new broadcast storm developing.

Note	Since it is difficult to simulate a STP failure, you will only verify the operation of STP for VLAN 10.

Activity

Complete the following steps:

Step 1 Map out the spanning tree for VLAN 10.

Start by accessing the console of SW1 and displaying the spanning-tree status for VLAN 10.


```
SW1# show spanning-tree vlan 10

VLAN0010
  Spanning tree enabled protocol ieee
  Root ID    Priority    24586
             Address     aabb.cc00.5500
             Cost        200
             Port        2 (Ethernet0/1)
             Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec

  Bridge ID  Priority    32778  (priority 32768 sys-id-ext 10)
             Address     aabb.cc00.5400
             Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec
             Aging Time  300 sec

Interface           Role Sts Cost      Prio.Nbr Type
------------------- ---- --- --------- -------- --------------------------------
-
Et0/1               Root FWD 100       128.2    Shr
Et0/2               Altn BLK 100       128.3    Shr
Et1/0               Desg FWD 100       128.5    Shr
```

Access the console of SW2 and display the spanning-tree status for VLAN 10.

```
SW2# show spanning-tree vlan 10

VLAN0010
  Spanning tree enabled protocol ieee
  Root ID    Priority    24586
             Address     aabb.cc00.5500
             This bridge is the root
             Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec

  Bridge ID  Priority    24586  (priority 24576 sys-id-ext 10)
             Address     aabb.cc00.5500
             Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec
             Aging Time  300 sec

Interface           Role Sts Cost      Prio.Nbr Type
------------------- ---- --- --------- -------- --------------------------------
-
Et0/1               Desg FWD 100       128.2    Shr
Et0/2               Desg FWD 100       128.3    Shr
```

Access the console of SW3 and display the spanning-tree status for VLAN 10.

```
SW3# show spanning-tree vlan 10

VLAN0010
  Spanning tree enabled protocol ieee
  Root ID    Priority    24586
             Address     aabb.cc00.5500
             Cost        100
             Port        2 (Ethernet0/1)
             Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec

  Bridge ID  Priority    32778  (priority 32768 sys-id-ext 10)
             Address     aabb.cc00.5600
             Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec
             Aging Time  300 sec

Interface           Role Sts Cost      Prio.Nbr Type
------------------- ---- --- --------- -------- --------------------------------
-
Et0/0               Altn BLK 100       128.1    Shr
Et0/1               Root FWD 100       128.2    Shr
Et0/2               Desg FWD 100       128.3    Shr
Et1/0               Desg FWD 100       128.5    Shr
```

Access the console of SW4 and display the spanning-tree status for VLAN 10.

```
SW4# show spanning-tree vlan 10

VLAN0010
  Spanning tree enabled protocol ieee
  Root ID    Priority    24586
             Address     aabb.cc00.5500
             Cost        100
             Port        3 (Ethernet0/2)
             Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec

  Bridge ID  Priority    28682  (priority 28672 sys-id-ext 10)
             Address     aabb.cc00.5a00
             Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec
             Aging Time  300 sec

Interface           Role Sts Cost      Prio.Nbr Type
------------------- ---- --- --------- -------- --------------------------------
-
Et0/0               Desg FWD 100       128.1    Shr
Et0/1               Desg FWD 100       128.2    Shr
Et0/2               Root FWD 100       128.3    Shr
```

Note: The MAC addresses might differ in your output.

Step 2 The topology, as you currently understand it, looks like the following example:

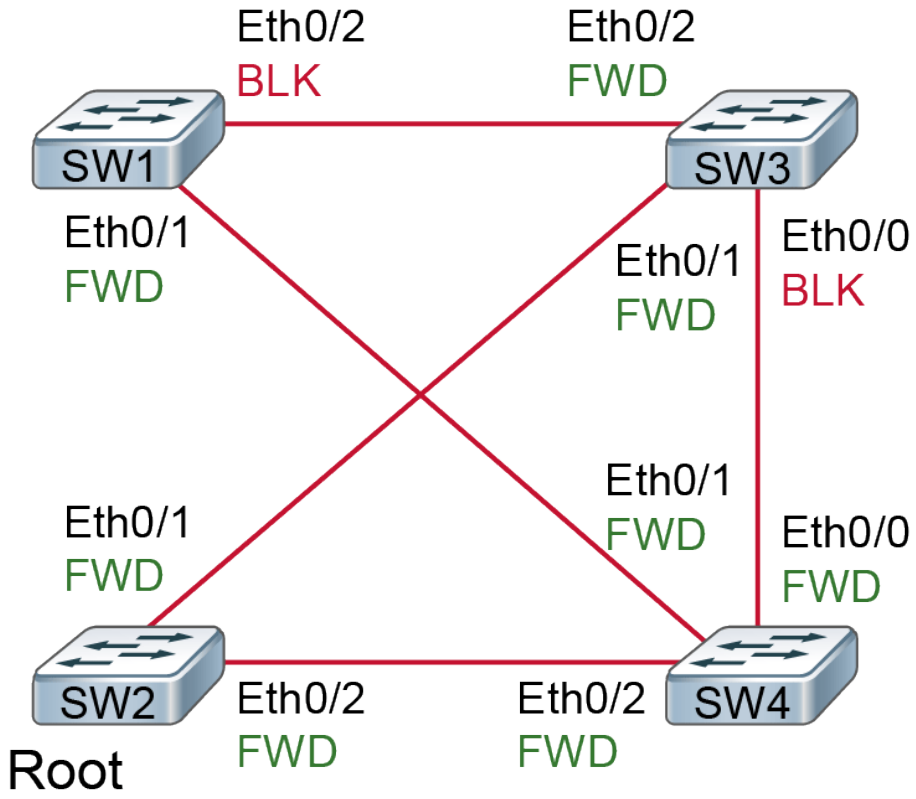

You can determine that the spanning tree is functioning as it should—each switch has only one path to the root bridge.

This is the end of the discovery lab.

Challenge

1. Which risk is posed by operating a switched network with redundant paths?

   A. loops
   B. CRC Errors
   C. late collisions
   D. interface input errors

2. Which feature of PVST+ is not available in RSTP?

   A. fast convergence on topology changes
   B. per-port STP
   C. per-VLAN STP instance
   D. edge ports

3. Place the phases of normal Spanning Tree initialization into the correct order.

Learning	1st
Listening	2nd
Forwarding	3rd
Blocking	4th

4. Which item represents a problem solved by PortFast?

   A. DHCP Timeout
   B. Bandwidth Throttling
   C. Duplex Mismatch
   D. Native VLAN Mismatch

5. Which two symptoms indicate that a loop might exist in the network? (Choose two.)

   A. The CPU load of the switches approaches 100 percent utilization.
   B. MAC addresses flap frequently between ports of the switches.
   C. Expired messages are received by the hosts.
   D. The load on the WAN links in the network approaches 100 percent utilization.

6. Which BID would win election as the root, assuming that the switches with these BIDs were in the same network?

   A. 32769:0200.1111.1111
   B. 32769:0200.2222.2222
   C. 4097:0200.1111.1111
   D. 4097:0200.2222.2222

7. Which of the following is true?

   A. None of the switches are root switch until the election.
   B. All switches operate as the root switch when they boot up.
   C. The root bridge is the switch with the highest BID.
   D. The root bridge is the switch which has bridge priority 65535.

8. When an access port is enabled with Portfast feature, which STP states are bypassed ? (Choose two)

   A. Learning
   B. Blocking
   C. Forwarding
   D. Listening

Answer Key

Challenge

1. A
2. C
3.

Blocking	1st
Listening	2nd
Learning	3rd
Forwarding	4th

4. A
5. A, B
6. C
7. B
8. A, D

Lesson 3: Improving Redundant Switched Topologies with EtherChannel

Introduction

The law firm calls CCS complaining of very slow data transfer and asks if there is something you can do to help. You and Bob tell them that the solution lies in a technology called EtherChannel. They agree to the implementation of EtherChannel, and Bob asks if you are ready to go onsite with him to perform the implementation.

Introduction to EtherChannel

With the proliferation of bandwidth-intensive applications such as video and interactive messaging, comes a need for greater network speeds and scalable bandwidth. You can increase network speed by using faster links, but faster links are more expensive. Furthermore, this solution cannot scale indefinitely and finds its limitation where the fastest possible port is no longer fast enough.

You can also increase network speeds by using more physical links between switches. One downside of this method is that you must be strictly consistent in the configuration of each physical link. The second one is that STP will block one of the links.

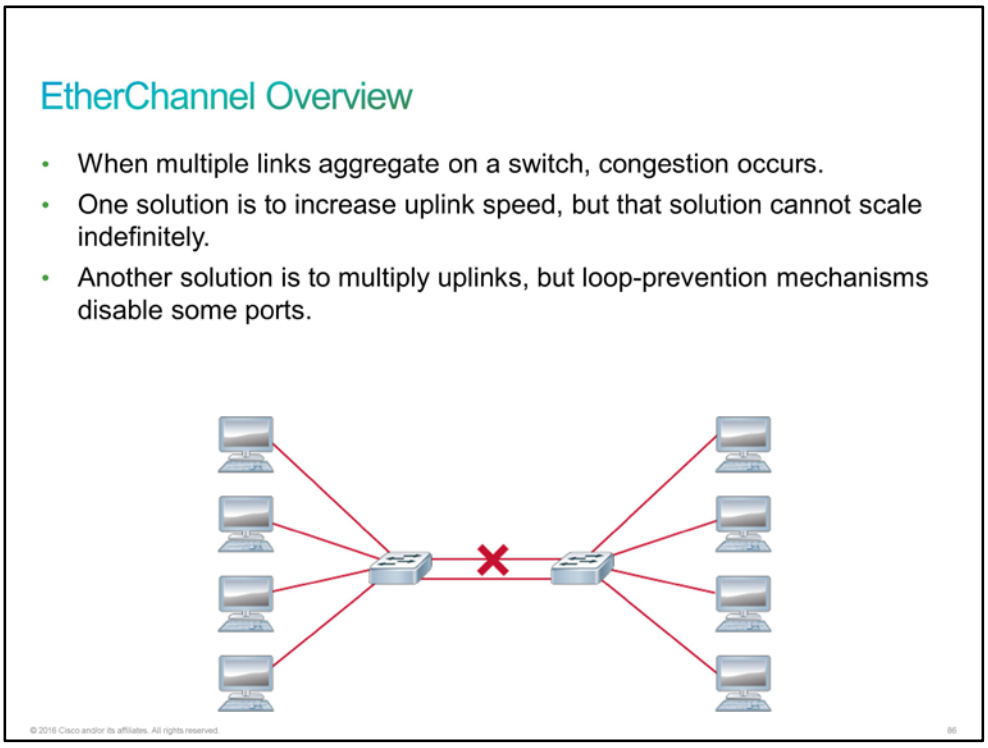

EtherChannel technology provides a solution. EtherChannel technology was originally developed by Cisco as a means of increasing speed between switches by grouping several FastEthernet or GigabitEthernet ports into one logical EtherChannel link, as shown in the following figure. Since the two physical links are bundled into a single EtherChannel, STP no longer sees the two physical links. Instead it sees a single EtherChannel. As a result, STP does not need to block one of the physical links to prevent a loop. Because all physical links in the EtherChannel are active, bandwidth is increased. EtherChannel provides the additional bandwidth without upgrading links to a faster and more expensive connection, because it relies on existing switch ports.

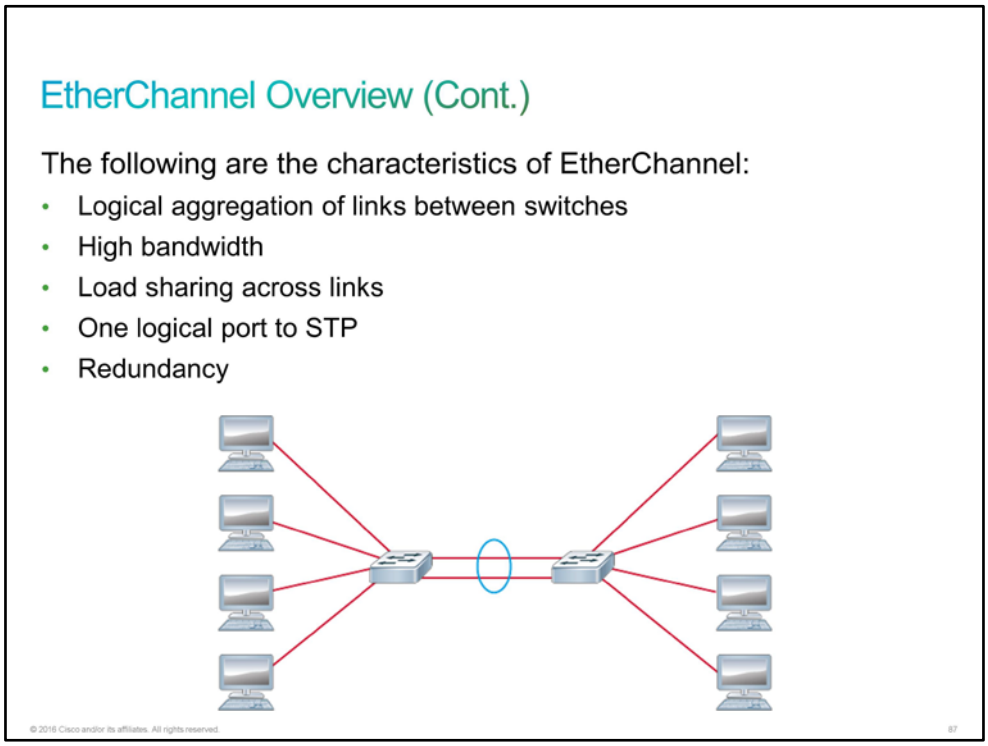

Some devices other than switches support link aggregation into an EtherChannel link. In any case, EtherChannel creates a one-to-one relationship. You can create an EtherChannel link between two switches or between an EtherChannel-enabled server and a switch. However, you cannot send traffic to two different switches through the same EtherChannel link. One EtherChannel link always connects only two devices.

You can group from two to eight physical ports into a logical EtherChannel link, but you cannot mix port types within a single EtherChannel. For example, you could group four Fast Ethernet ports into one logical Ethernet link, but you could not group two FastEthernet ports and two GigabitEthernet ports into one logical Ethernet link.

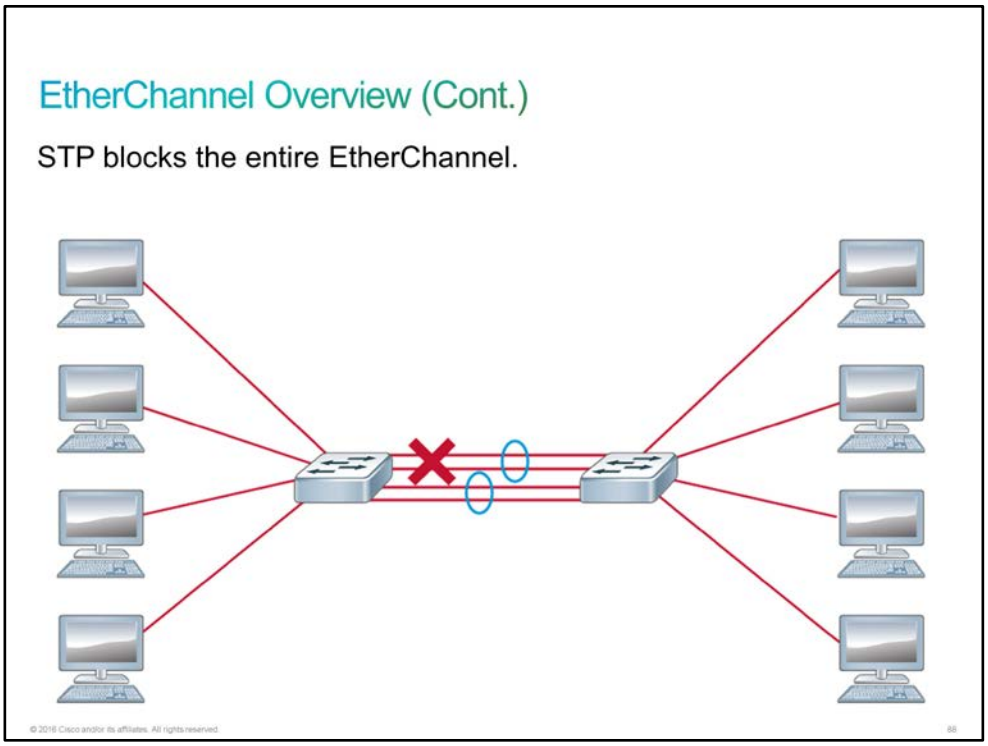

You can also configure multiple EtherChannel links between two devices. When several EtherChannels exist between two switches, STP may block one of the EtherChannels to prevent redundant links. When STP blocks one of the redundant links, it blocks one entire EtherChannel, thus blocking all the ports belonging to that EtherChannel link.

In addition to higher bandwidth, EtherChannel provides several other advantages:

- You can perform most configuration tasks on the EtherChannel interface instead of on each individual port, which ensures configuration consistency throughout the links.
- Because EtherChannel relies on the existing switch ports, you do not need to upgrade the link to a faster and more expensive connection to obtain more bandwidth.
- Load balancing is possible between links that are part of the same EtherChannel. Depending on your hardware platform, you can implement one or several load balancing methods, such as source MAC-to-destination MAC or source IP-to-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 redundant links. When STP blocks one of the redundant links, it blocks one EtherChannel, thus blocking 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. The loss of a physical link within an EtherChannel does not create a change in the topology, and you don't need a spanning-tree recalculation. As long as at least one physical link is active, the EtherChannel is functional, even if its overall throughput decreases.

EtherChannel Protocols

You can use two different protocols for link aggregation. These protocols allow ports with similar characteristics to form a channel through dynamic negotiation with adjoining switches.

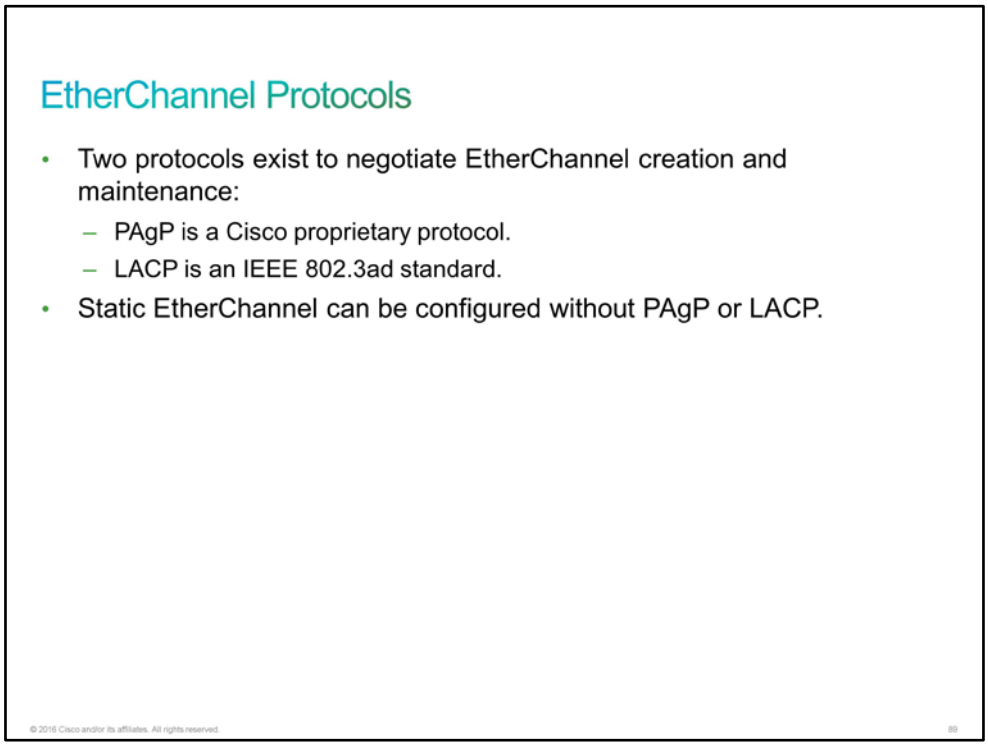

PAgP is a Cisco proprietary protocol that aids in the automatic creation of EtherChannel links. When you configure an EtherChannel link using PAgP, PAgP packets are sent between EtherChannel-capable ports to negotiate the forming of a channel. When PAgP identifies matched Ethernet links, it groups the links into an EtherChannel. The EtherChannel is then added to the spanning tree as a single bridge port.

When enabled, PAgP also manages 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 you create an EtherChannel, all ports have the same type of configuration. 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 will also change the configuration on the physical interfaces.

LACP is part of an IEEE specification (802.3ad) that allows several physical ports to be bundled to form a single logical channel. LACP allows a switch to negotiate an automatic bundle by sending LACP packets to the peer. It performs a function that is similar to PAgP with Cisco EtherChannel. Because LACP is an IEEE standard, you can use it to facilitate EtherChannels in multivendor environments. Cisco devices support both protocols.

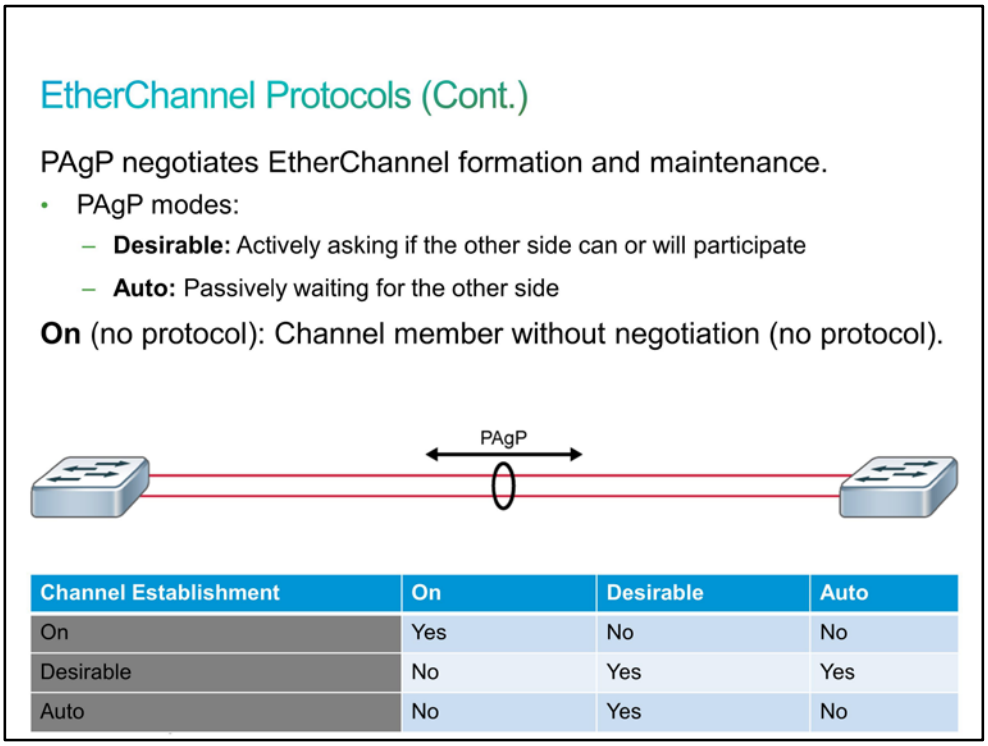

PAgP helps create the EtherChannel link by detecting the configuration of each side and making sure that they are compatible so that the EtherChannel link can be enabled when needed. The table shows the settings for PAgP.

Mode	Purpose
PAgP auto	This PAgP mode 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.
PAgP desirable	This PAgP mode places an interface in an active negotiating state in which the interface initiates negotiations with other interfaces by sending PAgP packets.
On	This mode forces the interface to channel without PAgP. Interfaces that you configure in the on mode do not exchange PAgP packets.

The modes must be compatible on each side. If you configure one side to be in *auto* mode, it will be placed in a 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 you disable all modes by using the **no** command or if no mode is configured, then the interface is placed in the off mode and EtherChannel is disabled.

Note that 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 will form, because the side that is set to *on* mode will not negotiate.

EtherChannel Protocols (Cont.)

LACP negotiates EtherChannel formation and maintenance.

- LACP modes:
  - **Active:** Actively asking if the other side can or will participate
  - **Passive:** Passively waiting for the other side

On (no protocol): Channel member without negotiation (no protocol).

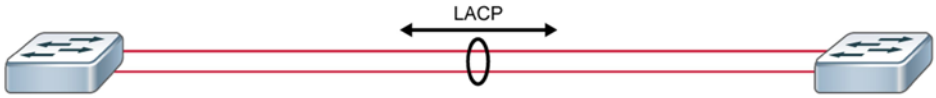

Channel Establishment	On	Active	Passive
On	YES	NO	NO
Active	NO	YES	YES
Passive	NO	YES	NO

LACP provides the same negotiation benefits as PAgP. LACP helps create the EtherChannel link by detecting the configuration of each side and making sure that they are compatible, so that the EtherChannel link can be enabled when needed. The table shows the settings for LACP.

Mode	Purpose
LACP passive	This LACP mode places a port in a passive negotiating state. In this state, the port responds to the LACP packets that it receives but does not initiate LACP packet negotiation.
LACP active	This LACP mode places a port in an active negotiating state. In this state, the port initiates negotiations with other ports by sending LACP packets.
On	This mode forces the interface to channel without LACP. Interfaces that you configure in the on mode do not exchange LACP packets.

Like PAgP, modes must be compatible on both sides for the EtherChannel link to form. The *on* mode is mentioned here again because it creates the EtherChannel configuration unconditionally, without PAgP or LACP dynamic negotiation.

Discovery 20: Configure and Verify EtherChannel

Introduction

The purpose of this discovery is to provide you with some experience working with EtherChannel. The live virtual lab is prepared with the switches represented in the topology diagram and the connectivity table. All devices have their basic configurations in place, including hostnames and IP addresses. Note that all the links between the switches use pairs of connections. You will see that this fact does not lead to doubling the bandwidth by default. You will configure EtherChannel on some of the links and verify the results.

Topology

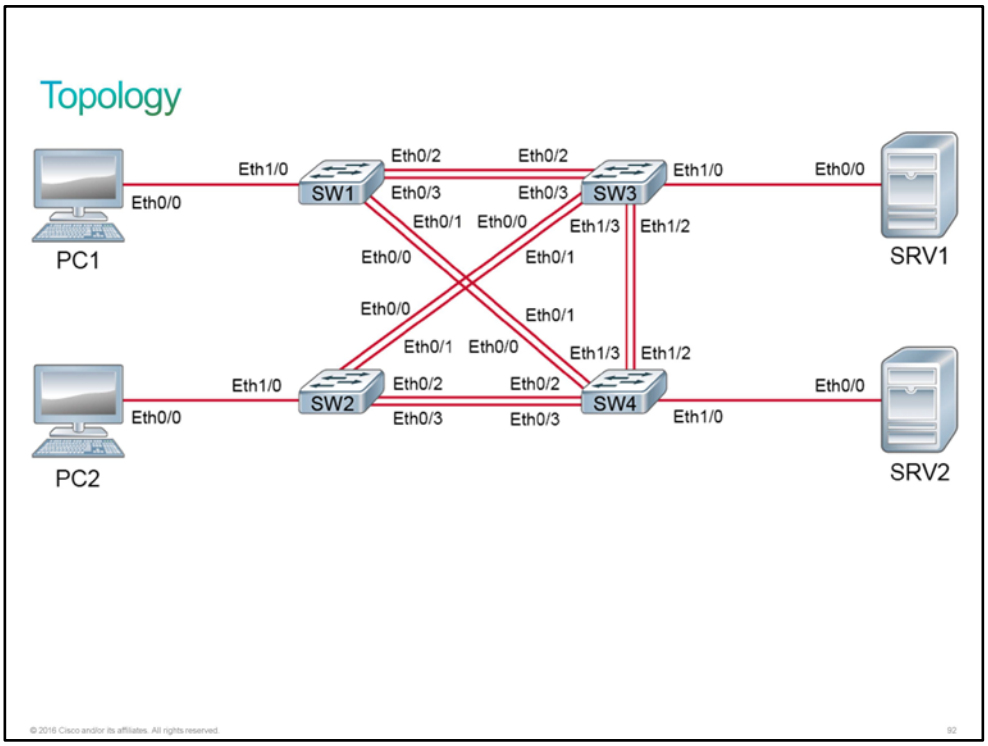

Job Aid

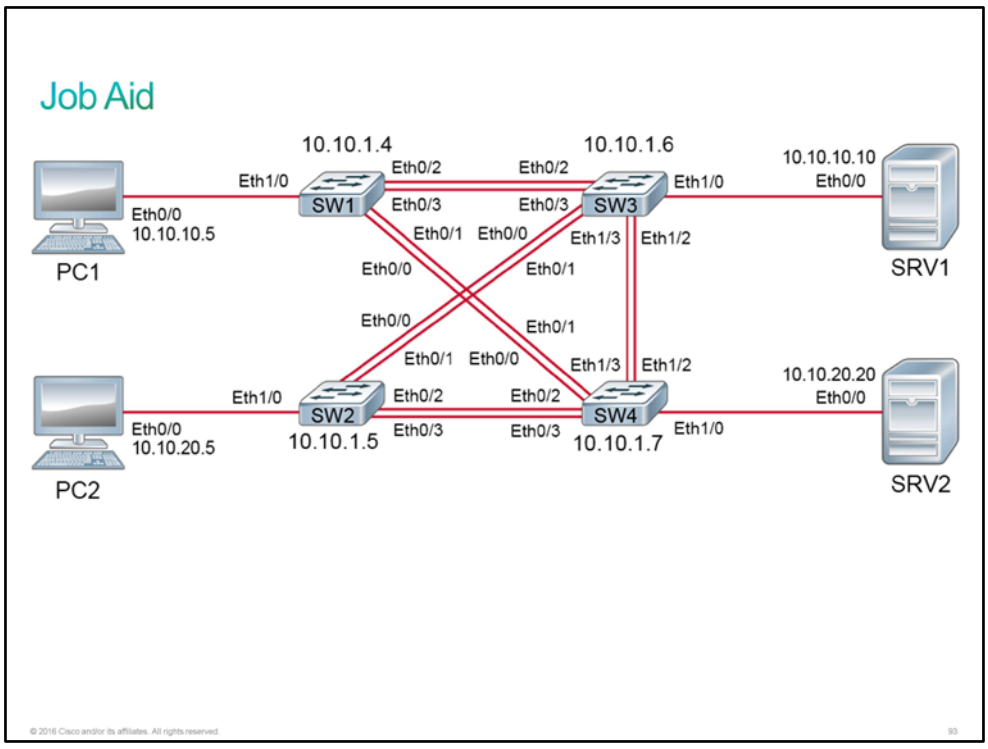

The configuration is as follows:

- All devices have their basic configurations in place, including hostnames and IP addresses.

Device Information

Device Details

Device	Interface	Neighbor	IP Address
PC1	Ethernet0/0	SW1	10.10.10.5/24
PC2	Ethernet0/0	SW2	10.10.20.5/24
SW1	VLAN 1	—	10.10.1.4/24
SW2	VLAN 1	—	10.10.1.5/24
SW3	VLAN 1	—	10.10.1.6/24
SW4	VLAN 1	—	10.10.1.7/24
SVR1	Ethernet0/0	SW3	10.10.10.10/24
SVR2	Ethernet0/0	SW4	10.10.20.20/24

Device Cabling Details

Switch	Port	Switch	Port
SW1	Ethernet0/0	SW4	Ethernet0/0
SW1	Ethernet0/1	SW4	Ethernet0/1
SW1	Ethernet0/2	SW3	Ethernet0/2
SW1	Ethernet0/3	SW3	Ethernet0/3
SW2	Ethernet0/0	SW3	Ethernet0/0
SW2	Ethernet0/1	SW3	Ethernet0/1
SW2	Ethernet0/2	SW4	Ethernet0/2
SW2	Ethernet0/3	SW4	Ethernet0/3
SW3	Ethernet1/2	SW4	Ethernet1/2
SW3	Ethernet1/3	SW4	Ethernet1/3

Note PCs and SRVs in the virtual lab environment are simulated as routers, so you should use Cisco IOS commands to configure them or make verifications.

Task 1: Configure and Verify EtherChannel

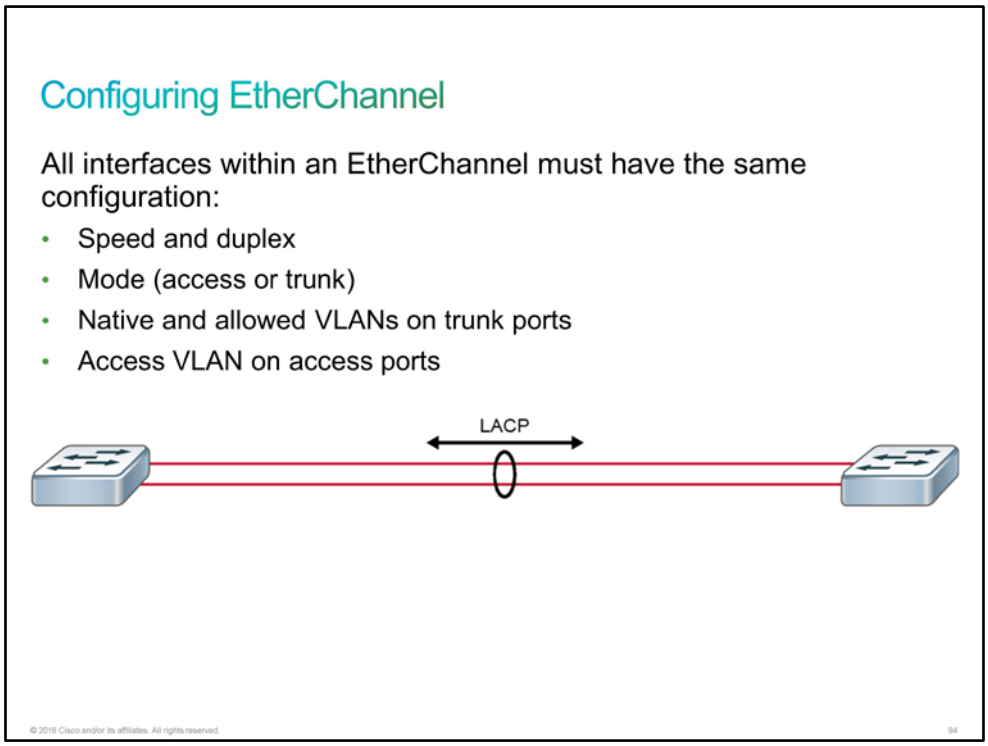

Follow these guidelines and restrictions when configuring the EtherChannel interfaces:

- **EtherChannel support:** All Ethernet interfaces on all modules support EtherChannel (maximum of eight interfaces), with no requirement that the interfaces should 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 VLAN:** An EtherChannel supports the same allowed range of VLANs on all the interfaces in a trunking Layer 2 EtherChannel.

If you have to change these settings, configure them in the port-channel interface configuration mode. After you configure the port-channel interface, any configuration that you apply to the port-channel interface affects individual interfaces as well. The opposite does not apply and will cause interface incompatibility in the EtherChannel.

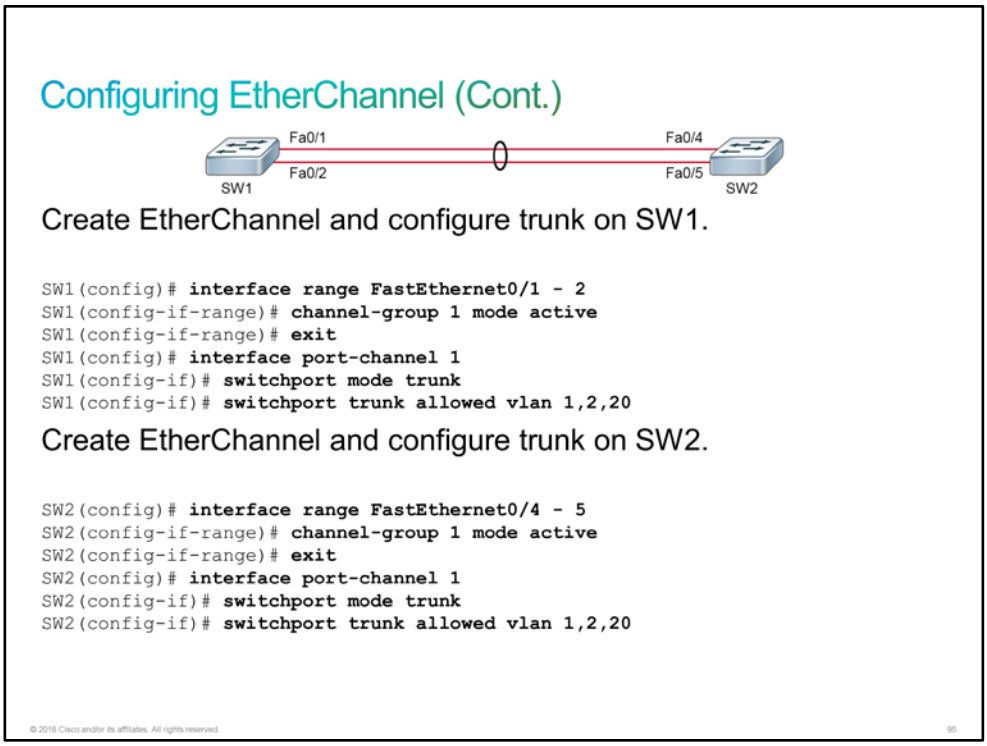

The configuration of an EtherChannel is based on two steps, as described in the table.

Command	Description
interface range *interface*	Specifies the interfaces that will compose the EtherChannel group. The **range** keyword allows you to select several interfaces and configure them all together. A good practice is to start by shutting down those interfaces, so that incomplete configuration will not start to create activity on the link.
channel-group *identifier* **mode active**	Creates the port-channel interface, if necessary, and assigns the specified interfaces to it. The identifier specifies a channel group number.

Note	The channel-group identifier does not need to match on both sides of the port channel. However, it is a good practice to do so because it makes it easier to manage the configuration.

In the example, FastEthernet0/1 and FastEthernet0/2 are bundled into EtherChannel interface port channel 1. To change Layer 2 settings on the EtherChannel interface, enter the EtherChannel interface configuration mode using the **interface port-channel** command, followed by the interface identifier. In the example, EtherChannel is configured as a trunk interface with allowed VLANs as specified.

Activity

Complete the following steps:

Step 1 Start by accessing the console of SW1 and displaying the interface status summary on SW1.

```
SW1# show interfaces status

Port      Name               Status       Vlan       Duplex  Speed Type
Et0/0     Link to SW4        connected    trunk        auto   auto unknown
Et0/1     Link to SW4        connected    trunk        auto   auto unknown
Et0/2     Link to SW3        connected    trunk        auto   auto unknown
Et0/3     Link to SW3        connected    trunk        auto   auto unknown
Et1/0     Link to PC1        connected    10           auto   auto unknown
Et1/1                        connected    1            auto   auto unknown
Et1/2                        connected    1            auto   auto unknown
Et1/3                        connected    1            auto   auto unknown
```

Both Ethernet0/2 and 0/3 are connected to SW3.

Ethernet1/0 is assigned to VLAN 10. The **show spanning-tree** examples that are used in this discovery arbitrarily specify VLAN 10, so Ethernet1/0 will be listed in a later example output.

Step 2 Display the spanning tree for VLAN 10 on SW1.

Both Ethernet0/2 and 0/3 connect to SW3, but only Ethernet0/2 is forwarding. The spanning tree is blocking on Ethernet0/3 to prevent a bridging loop. Only half of the potential bandwidth in this pair of links is in use.

```
SW1# show spanning-tree vlan 10

VLAN0010
  Spanning tree enabled protocol ieee
  Root ID    Priority    24586
             Address     aabb.cc00.0d00
             Cost        100
             Port        3 (Ethernet0/2)
             Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec

  Bridge ID  Priority    32778  (priority 32768 sys-id-ext 10)
             Address     aabb.cc00.0b00
             Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec
             Aging Time  300 sec

Interface           Role Sts Cost      Prio.Nbr Type
------------------- ---- --- --------- -------- --------------------------------
-
Et0/0               Altn BLK 100       128.1    Shr
Et0/1               Altn BLK 100       128.2    Shr
Et0/2               Root FWD 100       128.3    Shr
Et0/3               Altn BLK 100       128.4    Shr
Et1/0               Desg FWD 100       128.5    Shr
```

With a little more exploration, you could determine that the root switch for VLAN 10 is SW3.

Note: The MAC addresses might differ in your output.

Step 3 Shut down interfaces Ethernet0/2 and 0/3 on switch SW1.

On SW1, enter the following commands:

```
SW1#  conf t
Enter configuration commands, one per line. End with CNTL/Z.
SW1(config)# interface range Ethernet0/2 - 3
SW1(config-if-range)# shutdown
*Dec 28 09:09:31.692: %LINK-5-CHANGED: Interface Ethernet0/2, changed state to
administratively down
*Dec 28 09:09:31.693: %LINK-5-CHANGED: Interface Ethernet0/3, changed state to
administratively down
*Dec 28 09:09:32.693: %LINEPROTO-5-UPDOWN: Line protocol on Interface
Ethernet0/2, changed state to down
*Dec 28 09:09:32.694: %LINEPROTO-5-UPDOWN: Line protocol on Interface
Ethernet0/3, changed state to down
```

Step 4 Shut down interfaces Ethernet0/2 and 0/3 on switch SW3.

On SW3, enter the following commands:

```
SW3# conf t
Enter configuration commands, one per line. End with CNTL/Z.
SW3(config)# interface range Ethernet0/2 - 3
SW3(config-if-range)# shutdown
*Dec 28 09:10:17.356: %LINK-5-CHANGED: Interface Ethernet0/2, changed state to
administratively down
*Dec 28 09:10:17.356: %LINK-5-CHANGED: Interface Ethernet0/3, changed state to
administratively down
SW3(config-if-range)#
*Dec 28 09:10:18.360: %LINEPROTO-5-UPDOWN: Line protocol on Interface
Ethernet0/2, changed state to down
*Dec 28 09:10:18.360: %LINEPROTO-5-UPDOWN: Line protocol on Interface
Ethernet0/3, changed state to down
```

Step 5 Assign Ethernet0/2 and 0/3 to port channel 1 on the switch SW1. Use LACP protocol.

On SW1, enter the following commands:

```
SW1(config-if-range)# channel-group 1 mode active
Creating a port-channel interface Port-channel 1
SW1(config-if-range)#
```

Step 6 Assign Ethernet0/2 and 0/3 to port channel 1 on switch SW3. Use LACP protocol.

On SW3, enter the following commands:

```
SW3(config-if-range)# channel-group 1 mode active
Creating a port-channel interface Port-channel 1
```

Step 7 Enable interfaces Ethernet0/2 and 0/3 on switch SW1.

On SW1, enter the following commands:

```
SW1(config-if-range)# no shutdown
```

Step 8 Enable interfaces Ethernet0/2 and 0/3 on switch SW3.

On SW3, enter the following commands:

```
SW3(config-if-range)# no shutdown
*Dec 28 09:13:11.268: %LINK-3-UPDOWN: Interface Ethernet0/2, changed state to
up
*Dec 28 09:13:11.268: %LINK-3-UPDOWN: Interface Ethernet0/3, changed state to
up
*Dec 28 09:13:12.272: %LINEPROTO-5-UPDOWN: Line protocol on Interface
Ethernet0/2, changed state to up
*Dec 28 09:13:12.272: %LINEPROTO-5-UPDOWN: Line protocol on Interface
Ethernet0/3, changed state to up
SW3(config-if-range)#
*Dec 28 09:13:18.543: %LINEPROTO-5-UPDOWN: Line protocol on Interface Port-
channel1, changed state to up
```

Line protocol for physical interfaces Ethernet 0/2 and Ethernet 0/3 goes up. Logical interface Port-channel 1 also transitions to up state.

Step 9 Assign the description "EChannel to SW3" to port channel 1 on SW1.

On SW1, enter the following commands:

```
SW1(config-if-range)# exit
SW1(config)# interface port-channel 1
SW1(config-if)# description EChannel to SW3
SW1(config-if)# end
SW1#
```

Step 10 Assign the description "EChannel to SW1" to port channel 1 on SW3.

On SW3, enter the following commands:

```
SW3(config-if-range)# exit
SW3(config-if)# interface port-channel 1
SW3(config-if)# description EChannel to SW1
SW3(config-if)# end
SW3#
```

Step 11 Display the interface status summary on SW1.

The port channel is up on SW1.

```
SW1# show interfaces status

Port      Name               Status       Vlan       Duplex  Speed Type
Et0/0     Link to SW4        connected    trunk        auto   auto unknown
Et0/1     Link to SW4        connected    trunk        auto   auto unknown
Et0/2     Link to SW3        connected    trunk        auto   auto unknown
Et0/3     Link to SW3        connected    trunk        auto   auto unknown
Et1/0     Link to PC1        connected    10           auto   auto unknown
Et1/1                        connected    1            auto   auto unknown
Et1/2                        connected    1            auto   auto unknown
Et1/3                        connected    1            auto   auto unknown
Po1       EChannel to SW3    connected    trunk        auto   auto
```

Ethernet0/2 and 0/3 are still recognized as physical interfaces in Cisco IOS commands.

Step 12 Display the interface status summary on SW3.

The port channel is up on SW3.

```
SW3# show interfaces status

Port      Name               Status       Vlan       Duplex  Speed Type
Et0/0     Link to SW2        connected    trunk        auto   auto unknown
Et0/1     Link to SW2        connected    trunk        auto   auto unknown
Et0/2     Link to SW1        connected    trunk        auto   auto unknown
Et0/3     Link to SW1        connected    trunk        auto   auto unknown
Et1/0     Link to SRV1       connected    10           auto   auto unknown
Et1/1                        connected    1            auto   auto unknown
Et1/2     Link to SW4        connected    trunk        auto   auto unknown
Et1/3     Link to SW4        connected    trunk        auto   auto unknown
Po1       EChannel to SW1    connected    trunk        auto   auto
```

Ethernet0/2 and 0/3 are still recognized as physical interfaces in Cisco IOS commands.

Step 13 Display the spanning tree for VLAN 10 on SW1. This will be revealing.

Ethernet0/2 and 0/3 are no longer visible to the spanning tree. Instead, they have been replaced with the virtual port channel 1 interface.

```
SW1# show spanning-tree vlan 10

VLAN0010
  Spanning tree enabled protocol ieee
  Root ID    Priority    24586
             Address     aabb.cc00.0800
             Cost        56
             Port        65 (Port-channel1)
             Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec

  Bridge ID  Priority    32778  (priority 32768 sys-id-ext 10)
             Address     aabb.cc00.0500
             Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec
             Aging Time  15  sec

Interface           Role Sts Cost      Prio.Nbr Type
------------------- ---- --- --------- -------- --------------------------------
-
Et0/0               Desg FWD 100       128.1    Shr
Et0/1               Desg FWD 100       128.2    Shr
Et1/0               Desg FWD 100       128.5    Shr
Po1                 Root FWD 56        128.65   Shr
```

If you are quick, Ethernet0/0 and 0/1 may be listening or learning. If so, repeat the command until all interfaces are forwarding. The cost of the port channel is 56, which is much lower than the cost of 100 that is assigned to individual interfaces.

The port channel is forwarding. The forwarding state implies that the port channel is forwarding on all member interfaces. Remember that at the beginning of this discovery both Ethernet0/0 and 0/1 were blocking. They are now forwarding. These interfaces connect to SW4. Since the SW1 path back to the root now costs less due to the port channel, its ports have been selected as the designated ports for these two links. SW4 is now alternate and blocking on these two links.

Note: The MAC addresses might differ in your output.

Verifying EtherChannel

You can use several commands to verify an EtherChannel configuration. You can first use the **show interface port-channel** command to display the general status of the EtherChannel interface. In the example, the port channel 1 interface is up.

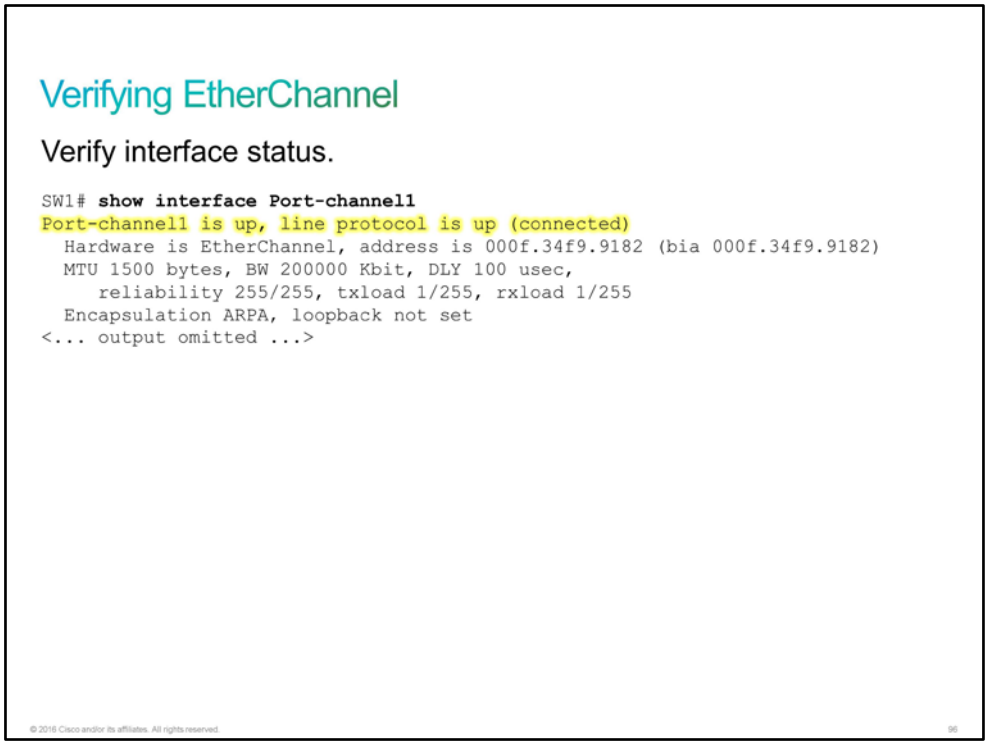

When several port channel interfaces are configured on the same device, you can use the **show etherchannel summary** command to simply display one line of information per port channel. In this example, the switch has one EtherChannel configured; group 1 uses LACP. The interface bundle consists of the FastEthernet0/1 and FastEthernet0/2 interfaces. You can see that the group is Layer 2 EtherChannel and that it is in use (shown by the letters SU next to the port channel number).

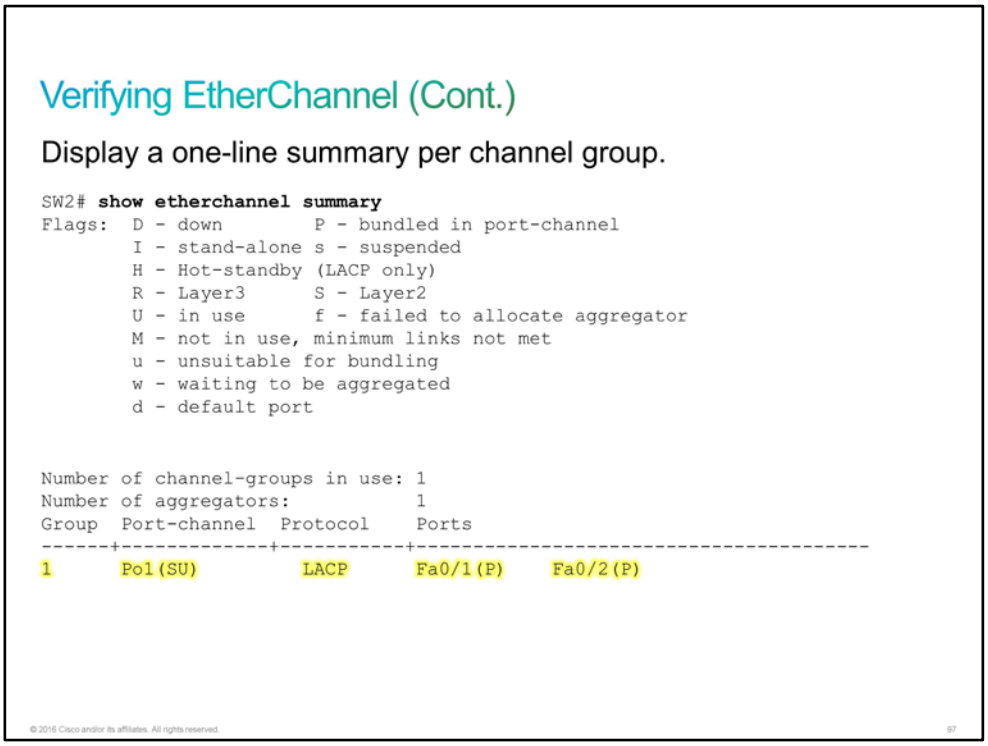

Use the **show etherchannel port-channel** command to display information about the specific port channel interface. In the example, the port channel 1 interface consists of two physical interfaces: FastEthernet0/1 and FastEthernet0/2. It uses LACP in active mode. It is properly connected to another switch with a compatible configuration, which is why the port channel is said to be in use.

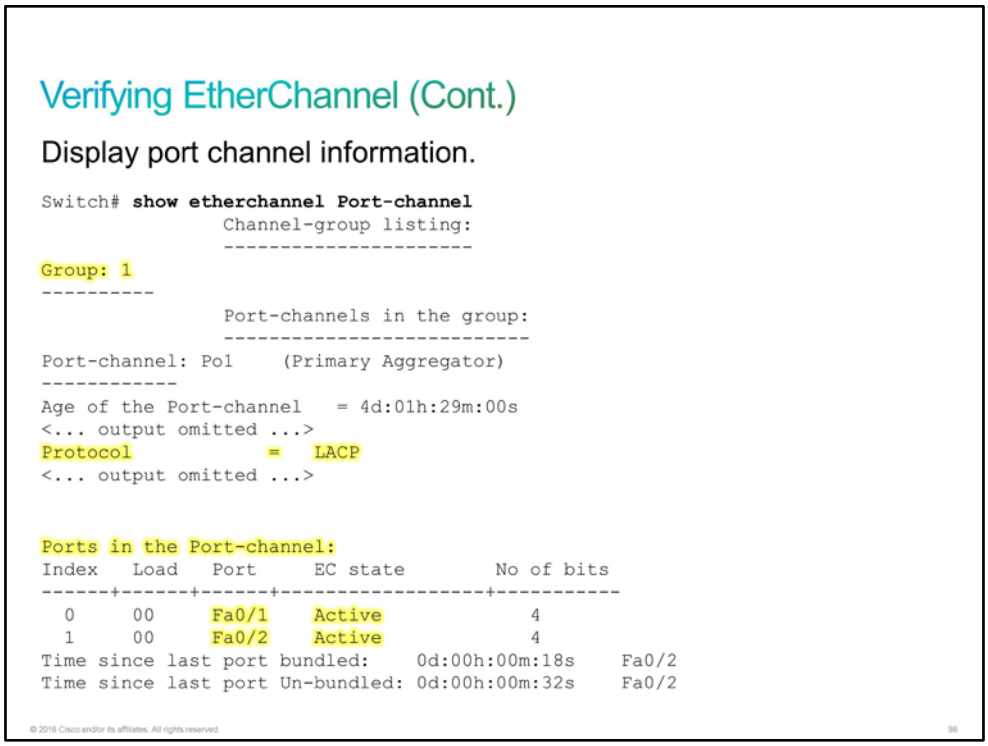

Note	*Load* does not actually indicate the load over an interface. It is meant to be a hexadecimal value that decodes which interface will be chosen for a specific flow of traffic.

Step 14 Display the full status of the port channel 1 interface on SW1.

From this output, you can determine that the port channel is made up of Ethernet0/2 and 0/3 and that the logical bandwidth on the channel is 20 Mbps.

```
SW1# show interfaces Port-channel 1
Port-channel1 is up, line protocol is up (connected)
  Hardware is Ethernet, address is aabb.cc00.0530 (bia aabb.cc00.0530)
  Description: EChannel to SW3
  MTU 1500 bytes, BW 20000 Kbit/sec, DLY 1000 usec,
     reliability 255/255, txload 1/255, rxload 1/255
  Encapsulation ARPA, loopback not set
  Keepalive set (10 sec)
  Auto-duplex, Auto-speed, media type is unknown
  input flow-control is off, output flow-control is unsupported
  Members in this channel: Et0/2 Et0/3
  ARP type: ARPA, ARP Timeout 04:00:00
  Last input 00:00:00, output never, output hang never
  Last clearing of "show interface" counters never
  Input queue: 0/2000/0/0 (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
     2041 packets input, 162930 bytes, 0 no buffer
     Received 1858 broadcasts (0 multicasts)
     0 runts, 0 giants, 0 throttles
     0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored
     0 input packets with dribble condition detected
     2069 packets output, 158394 bytes, 0 underruns
     0 output errors, 0 collisions, 0 interface resets
     0 unknown protocol drops
     0 babbles, 0 late collision, 0 deferred
     0 lost carrier, 0 no carrier
     0 output buffer failures, 0 output buffers swapped out
```

Step 15 Display the full EtherChannel status on SW1.

From this output, you can determine that the number of ports in this port channel is two. The members are the Ethernet0/2 and Ethernet0/3 interfaces. The protocol that was used to build the bundle is LACP.


```
SW1# 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:11m:56s
Logical slot/port   = 16/0          Number of ports = 2
HotStandBy port = null
Port state          = Port-channel Ag-Inuse
Protocol            =   LACP
Port security       = Disabled

Ports in the Port-channel:

Index   Load   Port     EC state        No of bits
------+------+------+------------------+-----------
  0     00     Et0/2    Active             0
  0     00     Et0/3    Active             0

Time since last port bundled:    0d:01h:11m:39s    Et0/2
```

Step 16 Display the summary of EtherChannel status on SW1.

From this output, you can determine that the Layer 2 port channel 1 is made up of Ethernet0/2 and 0/3.

```
SW1# 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      f - failed to allocate aggregator

        M - not in use, minimum links not met
        u - unsuitable for bundling
        w - waiting to be aggregated
        d - default port


Number of channel-groups in use: 1
Number of aggregators:           1

Group  Port-channel  Protocol    Ports
------+-------------+-----------+-----------------------------------------------
-
1      Po1(SU)         LACP      Et0/2(P)    Et0/3(P)
```

This is the end of the discovery lab.

Challenge

1. You have just configured two EtherChannels between two switches. Each EtherChannel contains four physical links. Which option describes how STP will react?

   A. It will block one physical link within one of the EtherChannels.
   B. It will block one entire EtherChannel.
   C. It will block one physical link within each EtherChannel.
   D. It will not block any links.

2. Which of the following is not true about "*on*" mode in Etherchannel ?

   A. *On* mode manually places the interface in an EtherChannel, without any negotiation.
   B. It works only if the other side is also set to *on*.
   C. If the other side is set to negotiate parameters through PAgP, EtherChannel will form, because the side that is set to *on* mode will negotiate by default.
   D. None of the above

3. Which option describes the primary purpose of LACP?

   A. to enable switch ports with similar characteristics to form an EtherChannel through dynamic negotiation with adjoining switches
   B. to maintain EtherChannels that are configured by PAgP
   C. to tear down EtherChannels that are created by PAgP after they are no longer needed
   D. to work with PAgP to combine switch ports with similar characteristics into EtherChannels
   E. to dynamically configure ports that have the same trunking status and trunk type with identical speed, duplex, and VLAN settings so that they can be combined into an EtherChannel link

4. Refer to the figure. Which combination will result in an EtherChannel being established between the two switches, which are running PAgP?

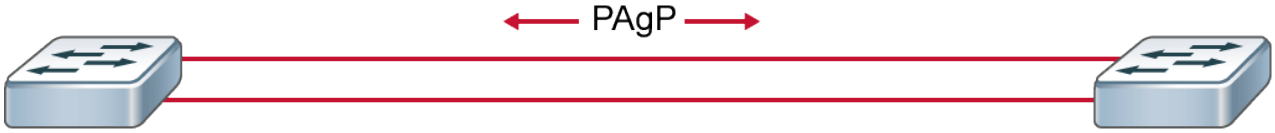

   A. switch 1 DESIRABLE, switch 2 AUTO
   B. switch 1 AUTO, switch 2 AUTO
   C. switch 1 ACTIVE, switch 2 PASSIVE
   D. switch 1 ACTIVE, switch 2 ACTIVE

5. You are about to configure one EtherChannel link between two switches. The EtherChannel will consist of four physical links. Which option describes how STP will interoperate with this configuration?

   A. STP will block all but one of the physical links.
   B. STP will enable links to meet the traffic bandwidth requirements.
   C. STP will prevent load balancing among the four links.
   D. STP will not block any links.

6. Which option describes the primary purpose of EtherChannel?

   A. enabling you to mix AC and DC power supply units in the same switch chassis
   B. enabling network devices to send and receive data across shared networks with all the functionality and security of a private network
   C. providing additional bandwidth without the expense of link upgrades
   D. enabling you to provide preferential treatment to high-priority traffic

7. Refer to the figure. Which two combinations will result in an EtherChannel being established between the two switches, which is running LACP? (Choose two.)

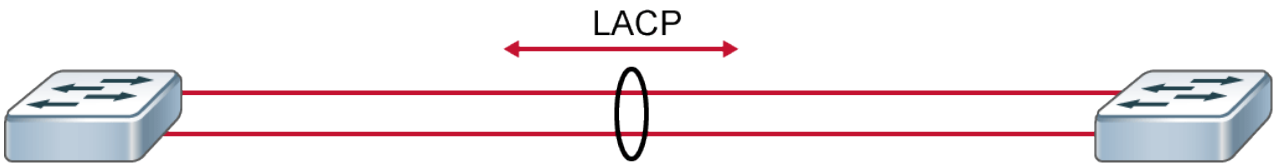

   A. switch 1 DESIRABLE, switch 2 AUTO
   B. switch 1 ACTIVE, switch 2 ACTIVE
   C. switch 1 DESIRABLE, switch 2 DESIRABLE
   D. switch 1 ACTIVE, switch 2 PASSIVE
   E. switch 1 PASSIVE, switch 2 PASSIVE

Answer Key

Challenge

1. B
2. C
3. A
4. A
5. D
6. C
7. B, D

Lesson 4: Routing Between VLANs

Introduction

Your boss sends you to your customer to enable routing between VLANs. You will need to understand why Layer 3 routing is needed. Also, you will have to describe different Layer 3 routing solutions. You will propose to the customer a router on a stick and a Layer 3 switch as viable solutions. You will also need to demonstrate basic configuration examples for both solutions.

Purpose of Inter-VLAN Routing

Each VLAN is a unique broadcast domain. Computers on separate VLANs are, by default, not able to communicate. You can permit these end stations to communicate by using a solution that is called *inter-VLAN routing*. Inter-VLAN communication occurs between broadcast domains via a Layer 3 device.

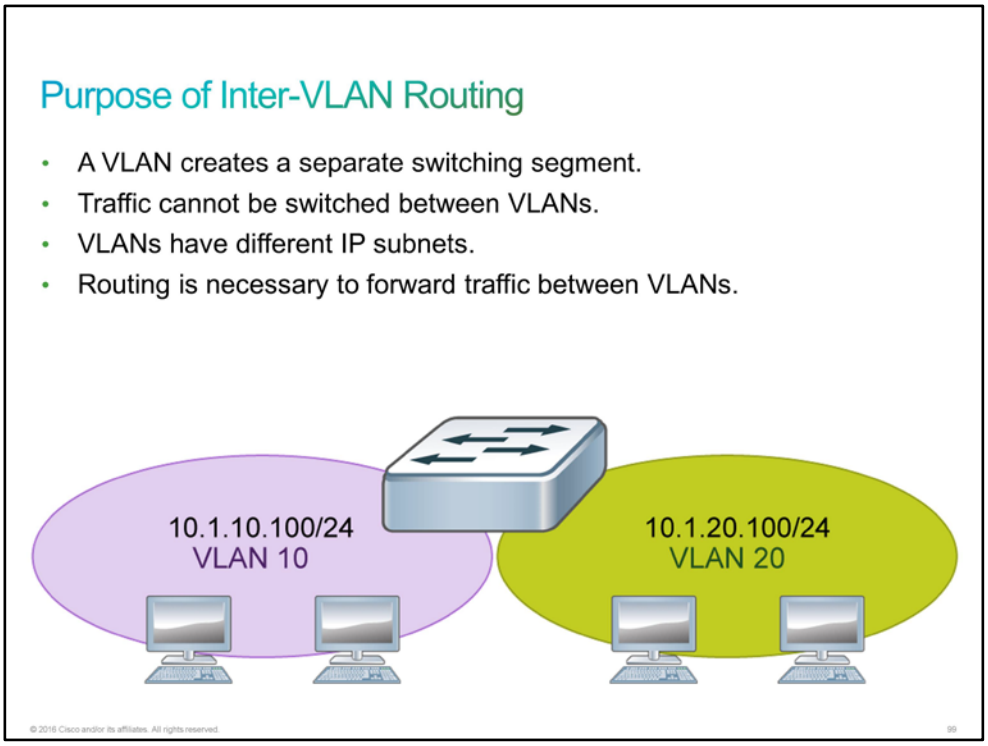

VLANs perform network partitioning and traffic separation at Layer 2 and are usually associated with unique IP subnets on the network. This subnet configuration facilitates the routing process in a multi-VLAN environment. Inter-VLAN communication cannot occur without a Layer 3 device. When you use a router to facilitate inter-VLAN routing, the router interfaces can be connected to separate VLANs.

Options for Inter-VLAN Routing

Inter-VLAN routing is a process of forwarding network traffic from one VLAN to another VLAN using a Layer 3 device.

Option 1: Router with a Separate Interface in Each VLAN

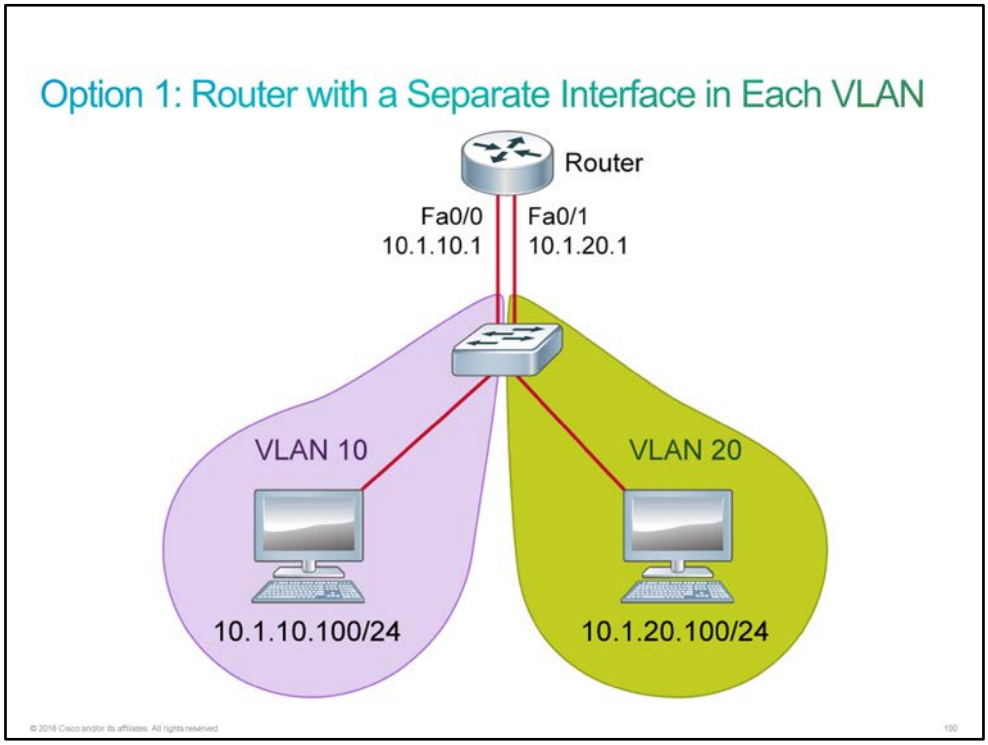

Traditional inter-VLAN routing requires multiple physical interfaces on both the router and the switch. VLANs are associated with unique IP subnets on the network. This subnet configuration facilitates the routing process in a multi-VLAN environment. When you use a router to facilitate inter-VLAN routing, the router interfaces can be connected to separate VLANs. Devices on those VLANs send traffic through the router to reach other VLANs. However, when you use a separate interface for each VLAN on a router, you can quickly run out of interfaces. This solution is not very scalable.

Option 2: Layer 3 Switch

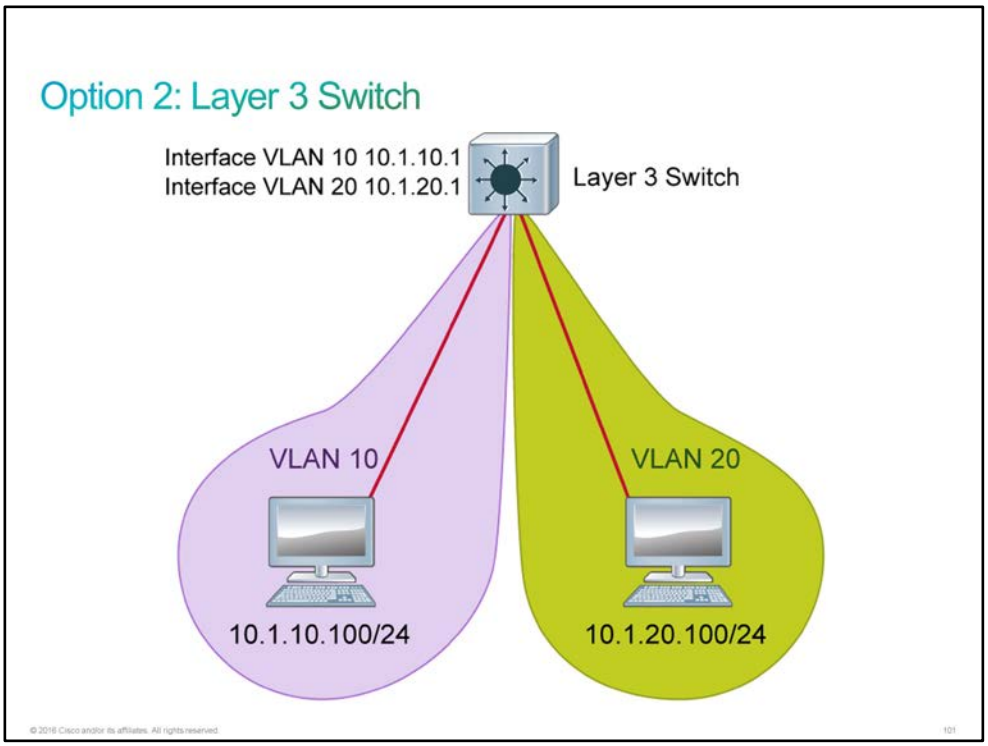

Some switches can perform Layer 3 functions, replacing the need for dedicated routers to perform basic routing on a network. Layer 3 switches are capable of performing inter-VLAN routing. Traditionally, a switch makes forwarding decisions by looking at the Layer 2 header, whereas a router makes forwarding decisions by looking at the Layer 3 header. A Layer 3 switch combines the functionality of a switch and a router in one device. It switches traffic when the source and destination are in the same VLAN and routes traffic when the source and destination are in different VLANs (that is, on different IP subnets). To enable a Layer 3 switch to perform routing functions, you must properly configure VLAN interfaces on the switch. You must use the IP addresses that match the subnet that the VLAN is associated with on the network. The Layer 3 switch must also have IP routing enabled.

Layer 3 switching is more scalable than router on a stick, because the latter can pass only so much traffic through the trunk link. In general, a Layer 3 switch is primarily a Layer 2 device that has been upgraded to have some routing capabilities. A router is a Layer 3 device that can perform some switching functions.

However, the line between switches and routers becomes hazier every day. Some Layer 2 switches, such as the switches in the Cisco Catalyst 2960 Series, support limited Layer 3 functionality. The Cisco Catalyst 2960 Switch supports static routing on SVIs. So, you can configure static routes, but routing protocols are not supported.

Following is an example configuration on the Layer 3 switch with PCs connected to the VLAN 10 and VLAN 20. PCs in the VLAN 10 will have default gateway 10.1.10.1 and PCs in the VLAN 20 will have default gateway 10.1.20.1. Layer 3 switch will perform routing between VLAN 10 and VLAN 20.


```
ip routing
!
interface Vlan10
 ip address 10.1.10.1 255.255.255.0
 no shutdown
!
interface Vlan20
 ip address 10.1.20.1 255.255.255.0
 no shutdown
```

Option 3: Router on a Stick

Not all inter-VLAN routing configurations require multiple physical interfaces. Some router software permits configuring router interfaces as trunk links. Trunk links open up new possibilities for inter-VLAN routing. A *router on a stick* is a type of router configuration in which a single physical interface routes traffic among multiple VLANs on a network.

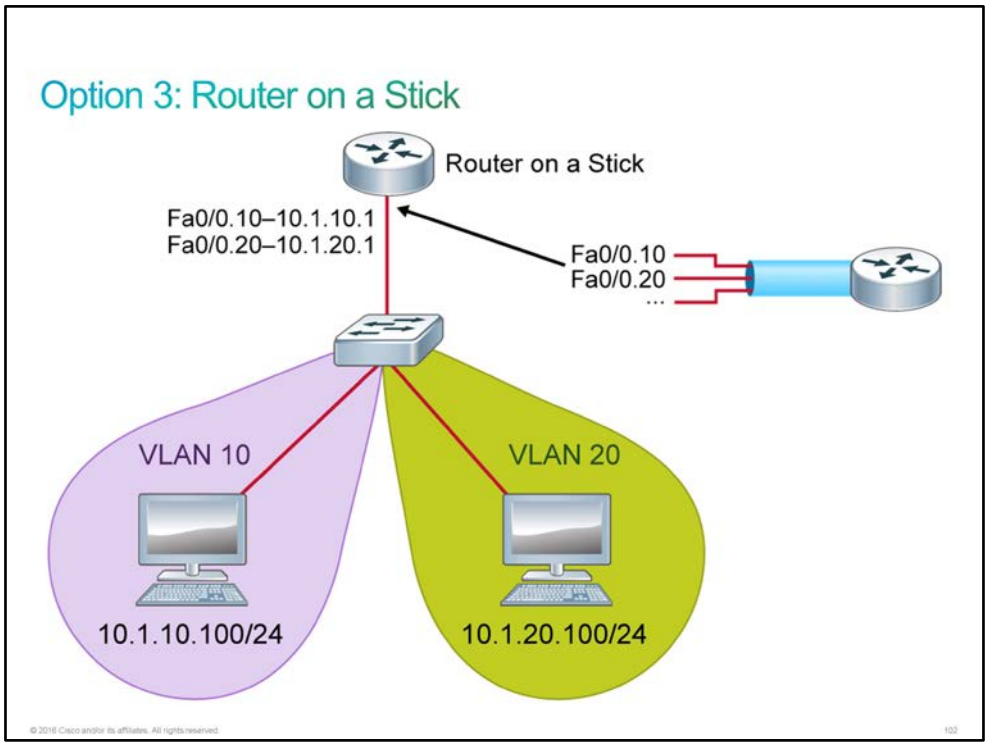

The figure shows a router that is attached to a core switch. The configuration between a router and a core switch is sometimes referred to as a router on a stick. The router interface is configured to operate as a trunk link and is connected to a switch port that is configured in the trunk mode. The router performs inter-VLAN routing by accepting VLAN-tagged traffic on the trunk interface coming from the adjacent switch and internally routing between the VLANs using subinterfaces. (Subinterfaces are multiple virtual interfaces that are associated with one physical interface.) To perform inter-VLAN routing functions, the router must know how to reach all VLANs that are being interconnected. There must be a separate logical connection on the router for each VLAN. Also, you must enable VLAN trunking (such as IEEE 802.1Q) on those connections. The router already knows about the directly connected networks. The router then forwards the routed VLAN traffic that is tagged for the destination VLAN out the same physical interface.

These subinterfaces are configured in software. Each is independently configured with its own IP addresses and a VLAN assignment to operate on a specific VLAN. Subinterfaces are configured for different subnets corresponding to their VLAN assignment to facilitate logical routing before the data frames are VLAN-tagged and sent back out the physical interface.

Discovery 21: Configure a Router on a Stick

Introduction

This discovery lab will guide you through routing between VLANs. The devices are configured as pictured in the topology diagram. Currently, devices have IP addresses in the 10.10.1.0/24 or 10.10.2.0/24 subnets. You will start by migrating this configuration to a one that uses two VLANs and two physical interfaces on R1 to route between them. You will then continue the migration to implement three VLANs and the use of trunking on R1 to allow one physical interface to have a logical presence on multiple VLANs. In the end, the switches will maintain their IP presence on VLAN 1, PC2 and PC4 on VLAN 2, and PC1 and PC3 will move to VLAN 3. R1 will be the default gateway for all the hosts and it will route between the VLANs.

Topology

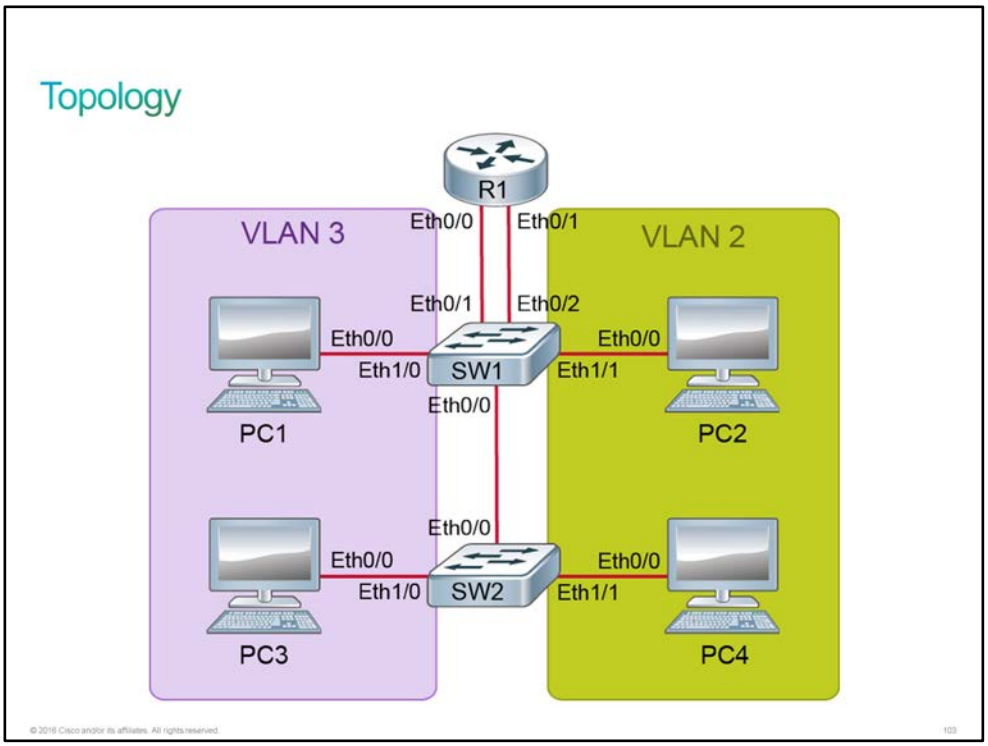

Job Aid

Device Information

Device Information Table

Device	Characteristic	Value
R1	Hostname	R1
R1	Ethernet0/0 description	Link to SW1 VLAN 1
R1	Ethernet0/0 IP address	10.10.1.1/24

Device	Characteristic	Value
R1	Ethernet0/1 description	Link to SW1 VLAN 2
R1	Ethernet0/1 IP address	NA
R1	Loopback0 IP address	10.10.99.1/24
PC1	Hostname	PC1
PC1	IP address	10.10.1.10/24
PC1	IP default gateway	10.10.1.1
PC2	Hostname	PC2
PC2	VLAN	2
PC2	IP address	10.10.2.20/24
PC3	Hostname	PC3
PC3	IP address	10.10.1.30/24
PC3	IP default gateway	10.10.1.1
PC4	Hostname	PC4
PC4	VLAN	2
PC4	IP address	10.10.2.40/24
SW1	Hostname	SW1
SW1	VLAN 1 IP address	10.10.1.4/24
SW1	IP default gateway	10.10.1.1
SW1	Ethernet0/0 description	Link to SW2
SW1	Ethernet0/1 description	Link to R1 VLAN 1
SW1	Ethernet0/2 description	Link to R1 VLAN 2
SW1	Ethernet1/0 description	Link to PC1
SW1	Ethernet1/1 description	Link to PC2
SW2	Hostname	SW2
SW2	VLAN 1 IP address	10.10.1.5/24

Device	Characteristic	Value
SW2	IP default gateway	10.10.1.1
SW2	Ethernet0/0 description	Link to SW1
SW2	Ethernet1/0 description	Link to PC3
SW2	Ethernet1/1 description	Link to PC4

PCs in the virtual lab environment are simulated as routers, so you should use Cisco IOS commands to configure them or make verifications.

Device Information Table (Changes)

Device	Characteristic	Value
R1	Ethernet0/1.2 IP address	10.10.2.1/24 VLAN 2
R1	Ethernet0/1.3 IP address	10.10.3.1/24 VLAN 3
PC1	VLAN	3
PC1	IP address	10.10.3.10/24
PC1	IP default gateway	10.10.3.1
PC2	IP default gateway	10.10.2.1
PC3	VLAN	3
PC3	IP address	10.10.3.30/24
PC3	IP default gateway	10.10.3.1
PC4	IP default gateway	10.10.2.1

Task 1: Configure a Router with a Trunk Link

Activity

Step 1 One way to implement routing between VLANs is to connect physical interfaces on routers to access ports that are assigned to the appropriate VLANs on the switches. R1 already has its Ethernet0/0 interface connected to a VLAN 1 access port (Ethernet0/1) on SW1. In the following series of steps, you will configure a second physical connection from R1 to an access port on VLAN 2. Access the console of SW1 and verify the current interface status.

Enter the following command to the SW1 switch:


```
SW1# sh int status

Port      Name               Status       Vlan       Duplex  Speed Type
Et0/0     Link to SW2        connected    trunk        auto   auto unknown
Et0/1     Link to R1 VLAN 1  connected    1            auto   auto unknown
Et0/2     Link to R1 VLAN 2  connected    1            auto   auto unknown
Et0/3                        connected    1            auto   auto unknown
Et1/0     Link to PC1        connected    1            auto   auto unknown
Et1/1     Link to PC2        connected    2            auto   auto unknown
Et1/2                        connected    1            auto   auto unknown
Et1/3                        connected    1            auto   auto unknown
```

R1 Ethernet0/1 is connected to SW1 Ethernet0/2.

Step 2 Configure Ethernet0/2 to be an access port that is assigned to VLAN 2.

Enter the following commands to the SW1 switch:

```
SW1# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
SW1(config)# int e 0/2
SW1(config-if)# switchport mode access
SW1(config-if)# switchport access vlan 2
SW1(config-if)# end
SW1#
```

Step 3 Access the console of R1 and configure the IP address of its Ethernet0/1 interface as 10.10.2.1/24. Enable the interface.

Enter the following commands to the R1 router:

```
R1# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
R1(config)# inte e 0/1
R1(config-if)# ip address 10.10.2.1 255.255.255.0
R1(config-if)# no shutdown
R1(config-if)# end
R1#
*Oct 30 07:57:23.805: %LINK-3-UPDOWN: Interface Ethernet0/1, changed state to
up
*Oct 30 07:57:24.810: %LINEPROTO-5-UPDOWN: Line protocol on Interface
Ethernet0/1, changed state to up
```

Step 4 R1 can now act as a gateway for VLAN 2. On PC4, configure 10.10.2.1 as its default gateway.

Enter the following commands to PC4:

```
PC4# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
PC4(config)# ip default-gateway 10.10.2.1
PC4(config)# end
PC4#
```

PC4 is already configured in VLAN 2.

Step 5 Verify that PC4 can now reach hosts on VLAN 1 by pinging PC1 (10.10.1.10).

Enter the following command to PC4:

```
PC4# ping 10.10.1.10
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.10.1.10, timeout is 2 seconds:
.!!!!
Success rate is 80 percent (4/5), round-trip min/avg/max = 1/1/1 ms
```

Consult the topology diagram to understand the physical path from PC4 to PC1. The VLAN 2 path flows from PC4 to SW2. Then it crosses the trunk link to SW1 and flows up to the R1 Ethernet0/1 interface. R1 then performs the route forwarding from VLAN 2 to VLAN 1. It sends the packet out its Ethernet0/0 interface back to SW1 on VLAN 1. SW1 then delivers the packet to PC1.

Step 6 Access the console of PC2 and repeat the gateway configuration.

Enter the following commands to PC2:

```
PC2# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
PC2(config)# ip default-gateway 10.10.2.1
PC2(config)# end
PC2#
```

PC2 is already configured in VLAN 2.

Step 7 Verify that PC2, which is connected to VLAN 2 on SW1, can reach hosts that are connected to VLAN 1 on SW2. Ping PC3 (10.10.1.30). The attempt should succeed.

Enter the following command to the PC2:

```
PC2# ping 10.10.1.30
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.10.1.30, timeout is 2 seconds:
.!!!!
Success rate is 80 percent (4/5), round-trip min/avg/max = 1/1/1 ms
```

Consult the topology diagram. Use the same logic that was described in a previous step to trace the path between PC2 and PC3.

Step 8 At this point, there is successful routing between two VLANs using physical interfaces on R1. In the next series of steps, you will add a third VLAN to the mix. You will leave the switch management IP addresses in VLAN 1, but all the PCs will be distributed between VLAN 2 and VLAN 3. There are not enough physical interfaces available on R1 to connect to the three VLANs individually. You will have to enable trunking on a physical interface on R1 to allow it to have a logical connection to multiple VLANs. Start the process by accessing the console of SW1 and creating VLAN 3, assigning it the name "Marketing."

Enter the following commands to the SW1 switch:

```
SW1# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
SW1(config)# vlan 3
SW1(config-vlan)# name Marketing
SW1(config-vlan)# exit
```


Step 9 Assign the port that is serving PC1 to VLAN 3.

Enter the following commands to the SW1 switch:

```
SW1(config)# int e 1/0
SW1(config-if)# switchport access vlan 3
SW1(config-if)# end
SW1#
```

Step 10 Verify the status of the interface Ethernet1/0.

Enter the following command to the SW1 switch:

```
SW1# show vlan id 3

VLAN Name                             Status    Ports
---- -------------------------------- --------- -------------------------------
3    Marketing                        active    Et0/0, Et1/0

VLAN Type  SAID       MTU   Parent RingNo BridgeNo Stp  BrdgMode Trans1 Trans2
---- ----- ---------- ----- ------ ------ -------- ---- -------- ------ ------
3    enet  100003     1500  -      -      -        -    -        0      0

Primary Secondary Type              Ports
------- --------- ----------------- ---------------------------------------
```

Step 11 Access the console of PC1 and configure its IP address and default gateway to be consistent with VLAN 3 (10.10.3.10/24 and 10.10.3.1)

Enter the following commands to PC1:

```
PC1# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
PC1(config)# inte e 0/0
PC1(config-if)# ip address 10.10.3.10 255.255.255.0
PC1(config-if)# exit
PC1(config)# ip default-gateway 10.10.3.1
PC1(config)# end
PC1#
```

Step 12 Access the console of SW2 and create VLAN 3 as you did on SW1. Assign the PC3 switch port to VLAN 3.

Enter the following commands to the SW2 switch:

```
SW2# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
SW2(config)# vlan 3
SW2(config-vlan)# name Marketing
SW2(config-vlan)# exit
SW2(config)# int e 1/0
SW2(config-if)# switchport access vlan 3
SW2(config-if)# end
SW2#
```

Step 13 Verify the port status of Ethernet1/0.

Enter the following command to the SW2 switch:

```
SW2# sh int e1/0 status

Port      Name               Status       Vlan       Duplex  Speed Type
Et1/0     Link to PC3        connected    3            auto   auto unknown
```

Step 14 Verify the status of VLAN 3.

Enter the following command to the SW2 switch:

```
SW2# show vlan id 3

VLAN Name                             Status    Ports
---- -------------------------------- --------- -------------------------------
3    Marketing                        active    Et0/0, Et1/0

VLAN Type  SAID       MTU   Parent RingNo BridgeNo Stp  BrdgMode Trans1 Trans2
---- ----- ---------- ----- ------ ------ -------- ---- -------- ------ ------
3    enet  100003     1500  -      -      -        -    -        0      0

Primary Secondary Type              Ports
------- --------- ----------------- ------------------------------------------
```

Step 15 Access the console of PC3 and configure its IP address and default gateway to be consistent with VLAN 3 (10.10.3.30/24 and 10.10.3.1).

Enter the following commands to PC3:

```
PC3# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
PC3(config)# int e 0/0
PC3(config-if)# ip address 10.10.3.30 255.255.255.0
PC3(config-if)# exit
PC3(config)# ip default-gateway 10.10.3.1
PC3(config)# end
PC3#
```

Step 16 PC3 and PC1 are now both correctly configured for VLAN 3 and assigned to it. Verify that there is connectivity between them. Attempt to ping PC1 from PC3. The attempt should succeed.

Enter the following command to the PC3:

```
PC3# ping 10.10.3.10
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.10.3.10, timeout is 2 seconds:
.!!!!
Success rate is 80 percent (4/5), round-trip min/avg/max = 1/1/1 ms
```


Trunk Link Configuration Example

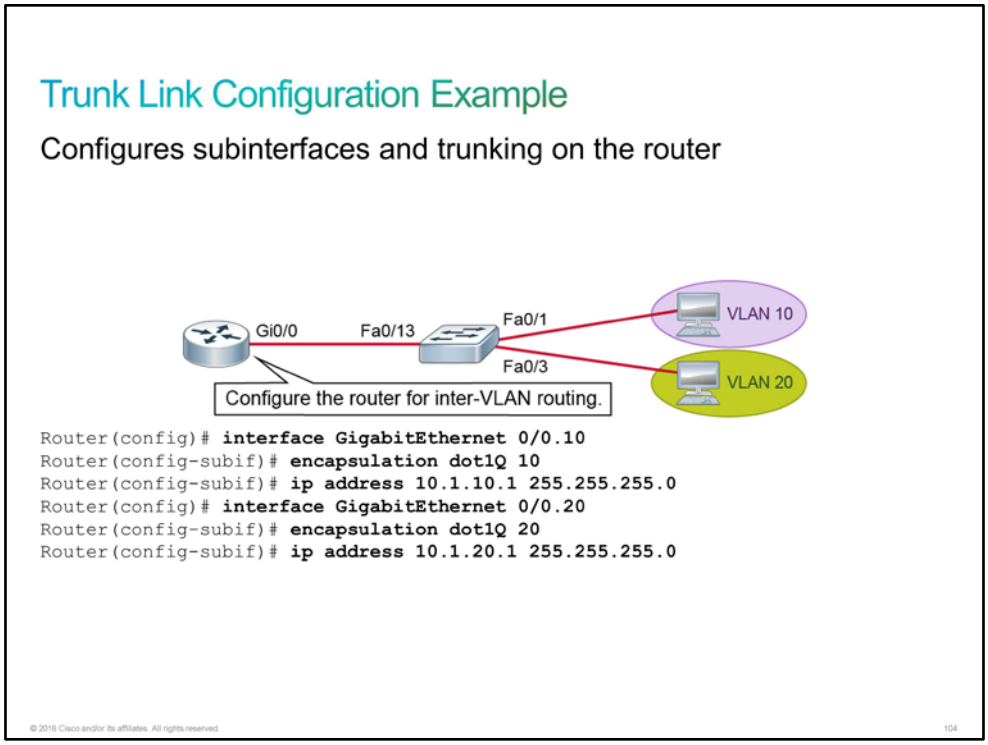

Command and Variable	Description
interface *interface*	Enters interface configuration mode
encapsulation dot1Q *vlan_number*	Defines the encapsulation format as IEEE 802.1Q and specifies the VLAN identifier
ip address *ip_address network_mask*	Assigns an IP address and network mask to an interface

In the figure, the GigabitEthernet0/0 interface is divided into subinterfaces—GigabitEthernet0/0.10 and GigabitEthernet0/0.20. Each subinterface represents the router in each of the VLANs for which it routes.

In the example, the **encapsulation dot1q 20** command enables 802.1Q encapsulation trunking on the GigabitEthernet0/0.20 subinterface. The value 20 represents the VLAN number (or VLAN identifier), therefore associating 802.1Q-tagged traffic from this VLAN with the subinterface.

Each 802.1Q-tagged VLAN on the trunk link requires a subinterface with 802.1Q encapsulation trunking that is enabled in this manner. The subinterface number does not have to be the same as the dot1q VLAN number. However, management and troubleshooting are easier when the two numbers are the same.

In this example, devices in different VLANs use the subinterfaces of the router as gateways to access the devices that are connected to the other VLANs.

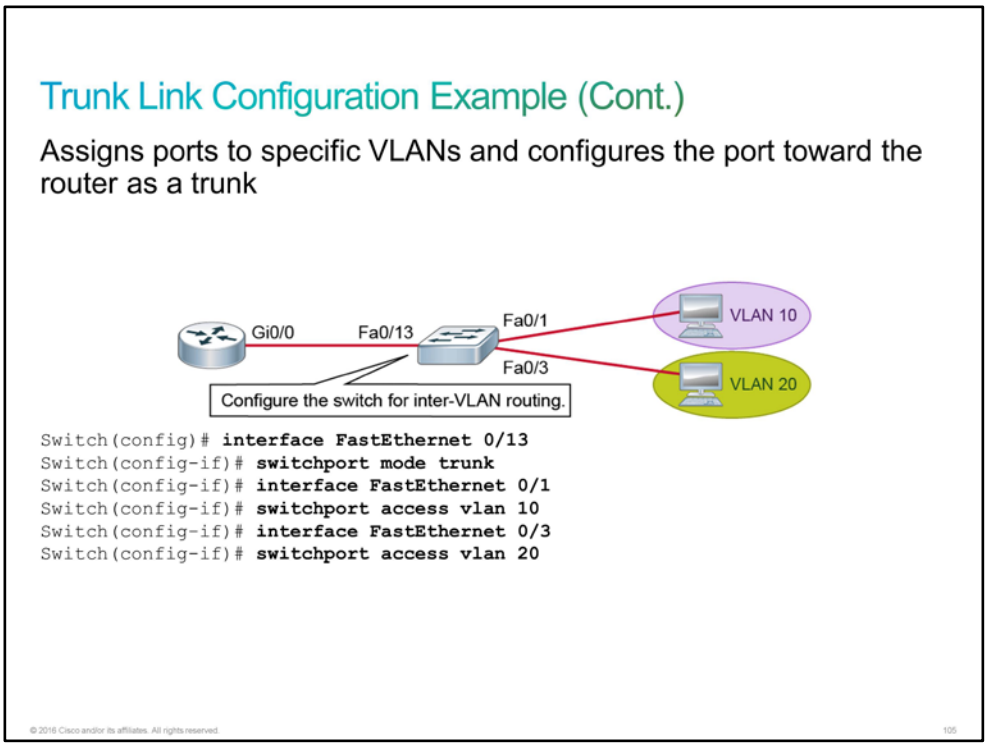

Command and Variable	Description
interface *interface*	Enters interface configuration mode.
switchport mode trunk	Sets the interface type. The **trunk** keyword specifies a trunking VLAN Layer 2 interface.
switchport access *vlan_number*	Sets the access VLAN when the interface is in the access mode. To reset the access-mode VLAN to the appropriate default VLAN for the switch, use the **no** form of this command.

On the switch, assign interfaces to the appropriate VLANs and configure the interface toward the router as a trunk. The trunk link will carry traffic from different VLANs, and the router will route between those VLANs.

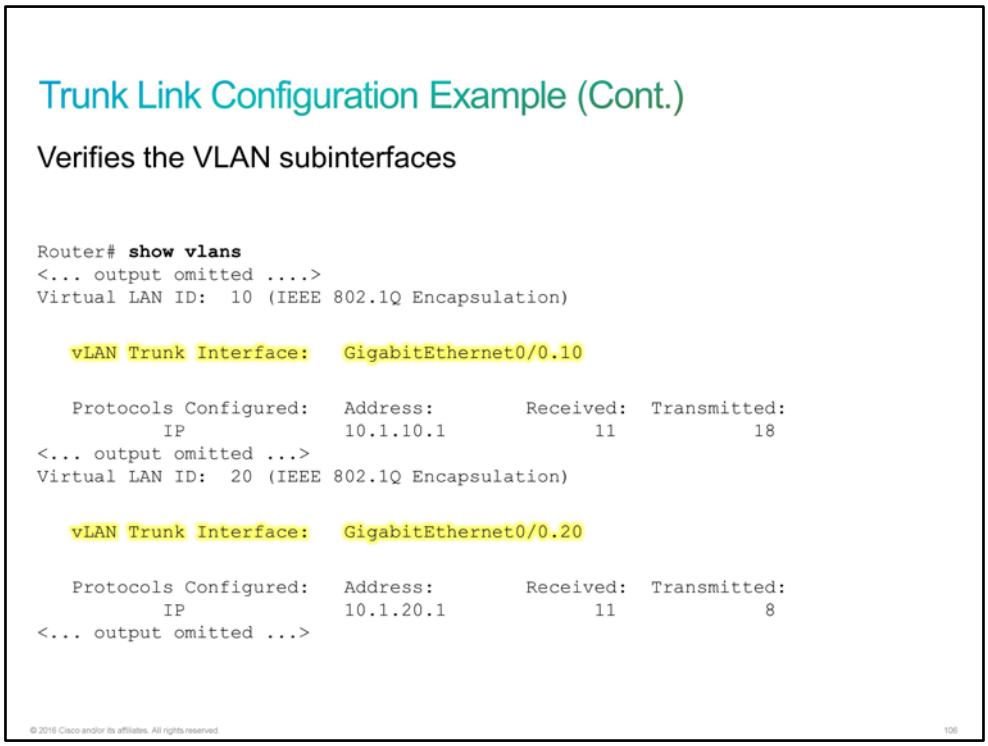

To verify the router configuration, use the **show** commands to display the running configuration, IP routing information, and IP protocol information for each VLAN to verify that the routing table represents the subnets of all VLANs.

The **show vlans** command displays the information about the Cisco IOS VLAN subinterfaces. The sample output shows two VLAN subinterfaces—FastEthernet0/0.10 and FastEthernet0/0.20.

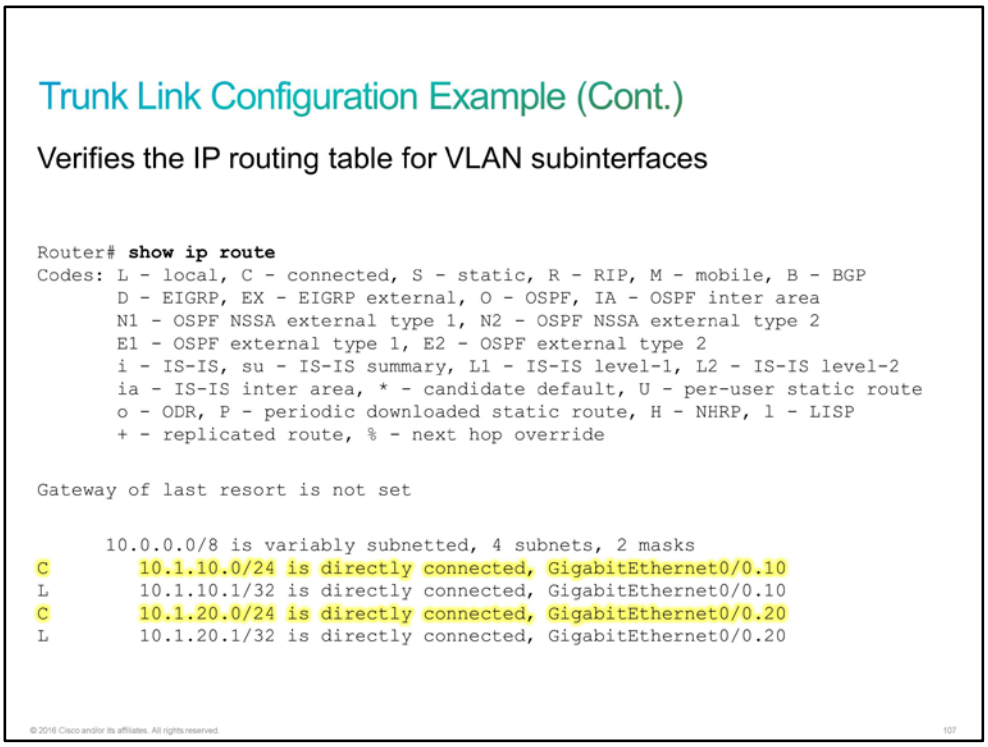

The **show ip route** command displays the state of the routing table. The sample output shows two subinterfaces. The GigabitEthernet0/0.10 and GigabitEthernet0/0.20 VLAN subinterfaces are directly connected to the router.

Step 17 To prepare for the trunk port configuration on R1, you must configure SW1 Ethernet0/2 as a trunk port. Connect to the SW1 console and configure Ethernet0/2.

Only two VLANs will be configured on a single interface.

Enter the following commands to the SW1 switch:

```
SW1# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
SW1(config)# int e 0/2
SW1(config-if)# switchport trunk encapsulation dot1q
SW1(config-if)# switchport trunk native vlan 256
SW1(config-if)# switchport mode trunk
SW1(config-if)# end
SW1#
*Oct 30 09:55:41.399: %LINEPROTO-5-UPDOWN: Line protocol on Interface
Ethernet0/2, changed state to down
*Oct 30 09:55:44.049: %LINEPROTO-5-UPDOWN: Line protocol on Interface
Ethernet0/2, changed state to up
```

Step 18 Now configure R1 so that it can connect to both VLAN 2 and VLAN 3 via its physical Ethernet0/1 interface. Connect to the console of R1 and remove the IP address that is currently configured on the physical interface.

Enter the following commands to the R1 router:


```
R1# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
R1(config)# int eth 0/1
R1(config-if)# no ip address
```

Step 19 Create the logical subinterface Ethernet0/1.2, assign it to VLAN 2, and give it the appropriate IP address.

Enter the following commands to the R1 router:

```
R1(config-if)# int eth0/1.2
R1(config-subif)# encapsulation dot1Q 2
R1(config-subif)# ip address 10.10.2.1 255.255.255.0
```

In this example, the subinterface number, the VLAN ID, and the third octet of the IP address are all consistent with each other. This practice is common, but it is not a technical requirement.

Step 20 Define the logical subinterface Ethernet0/1.3, assign it to VLAN 3, and give it the appropriate IP address. Leave the configuration mode when you are done.

Enter the following commands to the R1 router:

```
R1(config-subif)# int eth0/1.3
R1(config-subif)# encapsulation dot1Q 3
R1(config-subif)# ip address 10.10.3.1 255.255.255.0
R1(config-subif)# end
R1#
```

Step 21 If everything was configured successfully, R1 should have connectivity to devices in both VLAN 2 and VLAN 3. Verify this connectivity by pinging all PC systems. All pings should succeed.

Enter the following commands to the R1 router:

```
R1# ping 10.10.3.10
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.10.3.10, timeout is 2 seconds:
..!!!
Success rate is 60 percent (3/5), round-trip min/avg/max = 1/1/1 ms
R1# ping 10.10.2.20
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.10.2.20, timeout is 2 seconds:
.!!!!
Success rate is 80 percent (4/5), round-trip min/avg/max = 1/1/1 ms
R1# ping 10.10.3.30
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.10.3.30, timeout is 2 seconds:
..!!!
Success rate is 60 percent (3/5), round-trip min/avg/max = 1/1/1 ms
R1# ping 10.10.2.40
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.10.2.40, timeout is 2 seconds:
.!!!!
Success rate is 80 percent (4/5), round-trip min/avg/max = 1/1/1 ms
```

Step 22 Verify that the PCs can reach each other also. Access the console of PC1 and ping PC2.

Enter the following command to the PC1:

```
PC1# ping 10.10.2.20
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.10.2.20, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
```

Step 23 Verify that R1 is in the path between PC1 and PC2 using the **traceroute** command.

Enter the following command to PC1:

```
PC1# traceroute 10.10.2.20
Type escape sequence to abort.
Tracing the route to 10.10.2.20
VRF info: (vrf in name/id, vrf out name/id)
  1 10.10.3.1 0 msec 0 msec 0 msec
  2 10.10.2.20 0 msec *  1 msec
```

Step 24 Also verify that PC1 can still reach the switch management IP addresses that remain in VLAN 1. Ping 10.10.1.4 and 10.10.1.5. Both attempts should succeed.

Enter the following commands to PC1:

```
PC1# ping 10.10.1.4
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.10.1.4, timeout is 2 seconds:
.!!!!
Success rate is 80 percent (4/5), round-trip min/avg/max = 1/1/1 ms
PC1# ping 10.10.1.5
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.10.1.5, timeout is 2 seconds:
.!!!!
Success rate is 80 percent (4/5), round-trip min/avg/max = 1/1/1 ms
```

This is the end of the discovery lab.

Challenge

1. What is the purpose of Inter-VLAN Routing?

   A. To transport packets within VLANs
   B. To convert packets from Layer 2 to Layer 3
   C. To transport packets between VLANs
   D. To increase the security of VLANs
   E. To increase the scalability of VLANs

2. Each VLAN is usually associated with multiple subnets. True or False?

   A. True
   B. False

3. The Router on a stick method typically requires multiple physical interfaces each one physical interface for corresponding VLANs. True or False.?

   A. True
   B. False

4. Which of the following is the most scalable solution for Inter-VLAN Routing?

   A. Router with separate interfaces in each VLAN
   B. L3 Switch
   C. Router on a Stick

5. Which of the following involves only a single connection between a router and a switch for the purpose of Inter-VLAN Routing?

   A. Router with a separate Interface in each VLAN
   B. Layer 3 Switch
   C. Router on a stick

6. To correctly implement Router on a stick, which feature is critical?

   A. Router Sub-interfaces that are each configured with a dot1Q VLAN tag.
   B. Router Interface that is tagged with multiple dot1Q VLAN tags.
   C. Router Interface that is configured to be a trunk

7. Which command enables 802.1Q encapsulation trunking on Ethernet subinterface e0/0.20?

   A. **int e0/0.20**
      encapsulation 20 dot1q
   B. **int e0/0.20**
      encapsulation dot1q 20
   C. **int e0/0.20**
      encapsulation vlan 20 8021q
   D. **int e0/0.20**
      encapsulation vlan 20 dot1q

Answer Key

Challenge

1. C
2. B
3. B
4. B
5. C
6. A
7. B

Lesson 5: Using a Cisco IOS Network Device as a DHCP Server

Introduction

Your boss sends you to your customer to explain why DHCP is a practical alternative to manual allocation of IP addresses. You will need to show a DHCP server configuration example and explain DHCP terms such as the default router, DNS, lease, domain name, and excluded addresses. You will describe to your customer why having a central DHCP server is a good idea. You will also prepare to answer any DNS-related questions that your customer may have.

Need for a DHCP Server

While you can manually assign IP addresses to network hosts in small networks, it can present an administrative burden in medium-sized LANs.

It can be a time-consuming task for administrators to manually set all network connectivity parameters on the end host whenever it is moved or replaced. For mobile employees, who usually work from home and occasionally come into the office, manually setting the correct network parameters can be challenging. Also, manual settings may be incorrect, or the equipment that is brought to the office may already have settings in place. The result could be poor network connectivity on the host or even a problem for other users if a host with a duplicate IP address is connected to the local network.

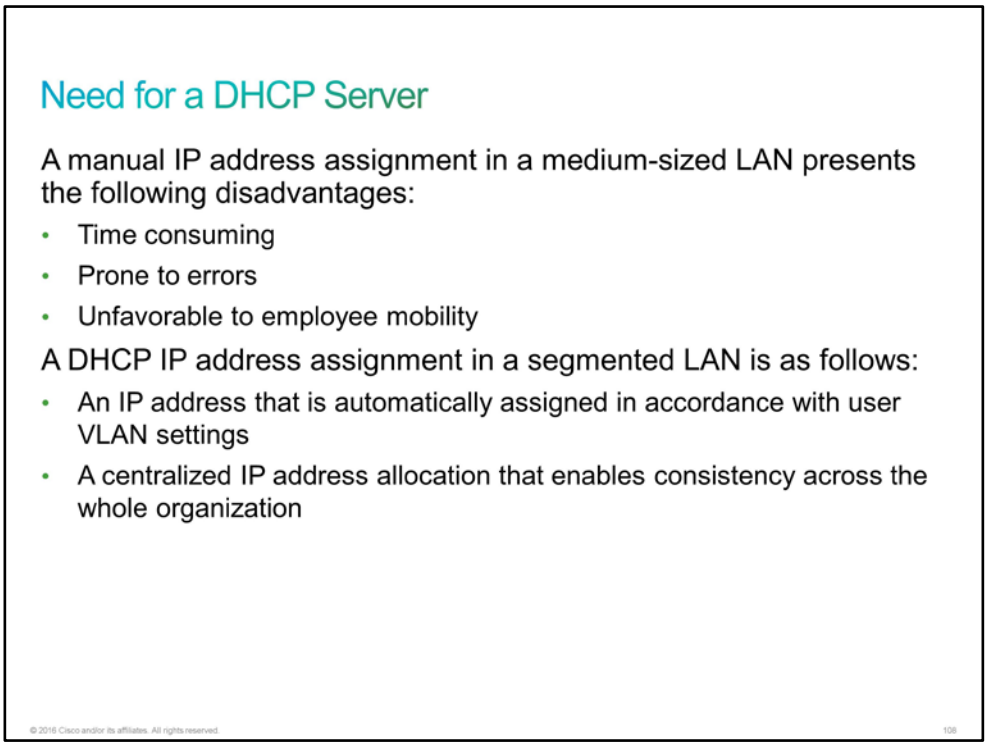

Introducing a DHCP server to the local network simplifies the IP address assignment. You usually use a DHCP server in small networks to support frequent changes and to assign correct IP addresses to guest hosts that are connecting to a LAN. A DHCP server greatly contributes to simplified administration when LANs are segmented using VLANs. A DHCP server automatically assigns IP addresses to end hosts according to the VLAN assignment of the host.

If you use a centralized DHCP server in your organization, it will enable the organization to administer all dynamic IP address assignments in one place. This practice ensures consistency across the organization. You can also easily manage branch offices, for example.

Understanding DHCP

The DHCP dynamic allocation of IP addresses is based on a client-server model. The figure displays an example of how an end host obtains an IP address from a DHCP server.

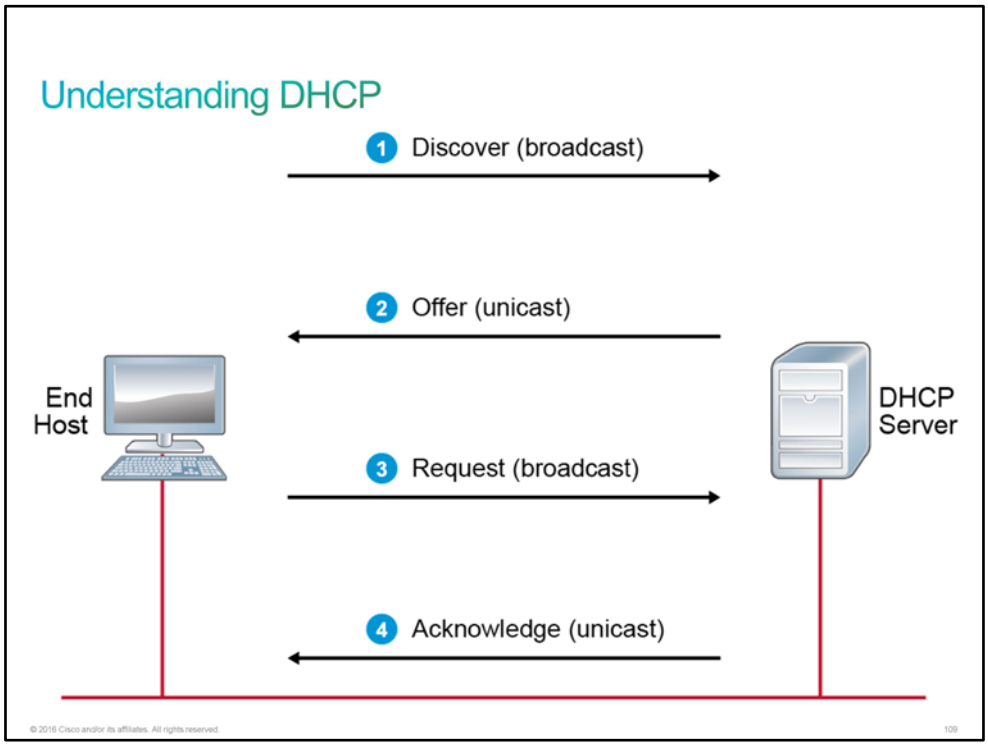

When DHCP dynamically allocates an IP address, the operation can be divided into the following phases:

1. **DHCP discover:** A client broadcasts a DHCP discover message with its own hardware MAC address to discover available DHCP servers.
2. **DHCP offer:** When a DHCP server receives a DHCP discover from a client, it reserves an IP address for the client and sends a DHCP offer to the client. This message contains the client MAC address, the IP address that the server is offering, the subnet mask, the lease duration, and the IP address of the DHCP server that is making the offer.
3. **DHCP request:** When a client receives a DHCP offer message, it responds with a DHCP request message, indicating its intent to accept the parameters in the DHCP offer. A DHCP request message is broadcast because the DHCP client has still not received an acknowledged IP address.
4. **DHCP acknowledgment (ACK):** After the DHCP server receives the DHCP request message, it acknowledges the request with a unicast DHCP acknowledgment message. The packet includes confirmation for all requested parameters. At this point, the IP configuration process is completed.

Note	The following examples show a simplified packet capture of a DHCP request.

```
Source              Destination         Protocol   Info
0.0.0.0             255.255.255.255     DHCP       DHCP Discover
78:ac:c0:52:e8:bd   ff:ff:ff:ff:ff:ff

Source              Destination          Protocol   Info
10.10.1.1           255.255.255.255      DHCP       DHCP Offer 10.10.1.241
00:1b:d5:9c:34:27   ff:ff:ff:ff:ff:ff

Source              Destination         Protocol   Info
0.0.0.0             255.255.255.255     DHCP       DHCP Request
78:ac:c0:52:e8:bd   ff:ff:ff:ff:ff:ff

Source              Destination         Protocol    Info
10.10.1.1           255.255.255.255     DHCP        DHCP ACK
00:1b:d5:9c:34:27   ff:ff:ff:ff:ff
```


Configuring a DHCP Server

In this example configuration scenario, you will configure a DHCP server on a Cisco IOS router. Guest LAN users need to receive an IP address from the specified address pool along with a default gateway, domain name, and IP address of a DNS server. The IP address assignment should be valid for 12 hours.

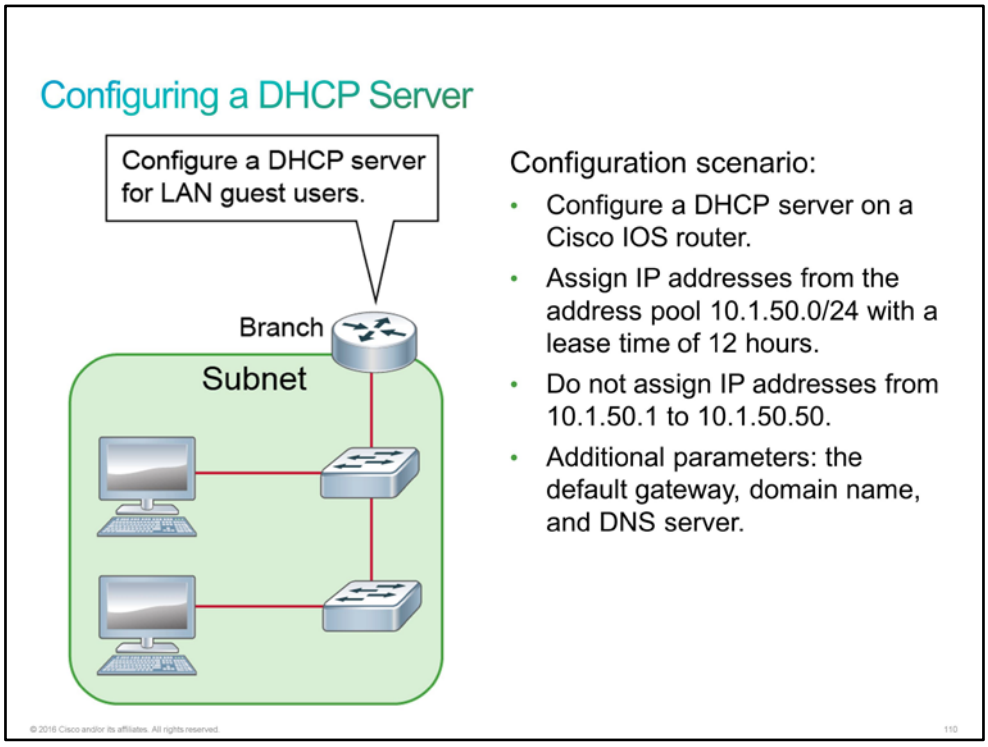

You can also implement a DHCP server on Cisco Catalyst switches.

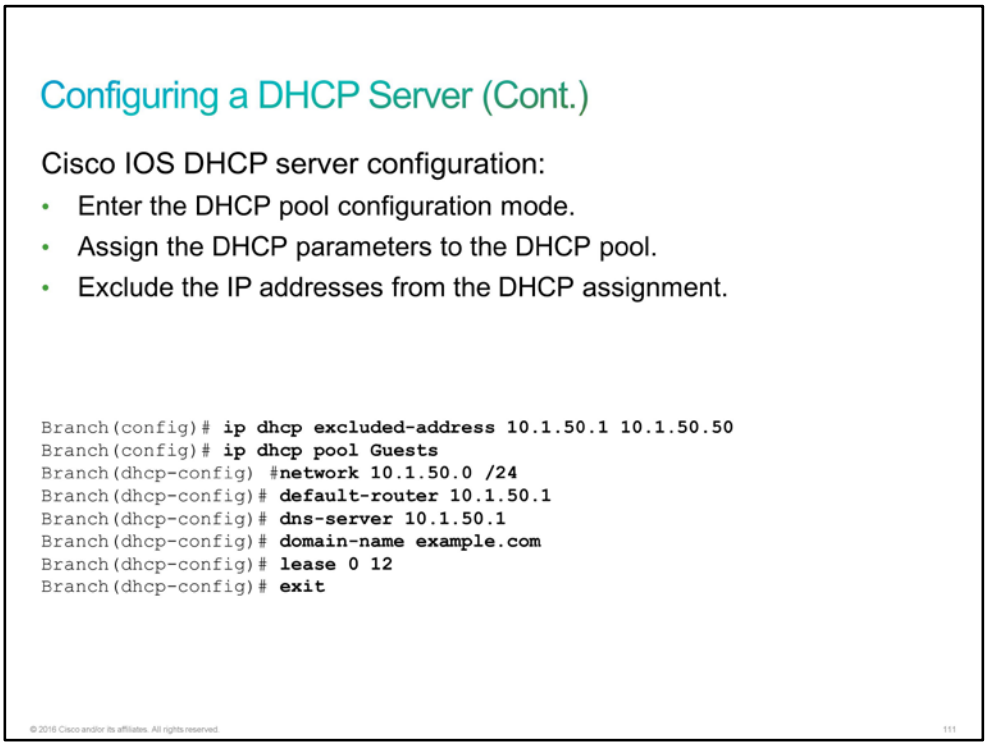

To enable the Cisco IOS DHCP server, enter the DHCP configuration mode by defining a DHCP pool. Use the commands that the table shows to define the pool parameters.

Command	Description
ip dhcp pool *name*	Defines the pool name and enters the DHCP configuration mode.
network *network-number* [*mask* | *prefix-length*]	Defines addresses in the DHCP pool. Optionally, define a subnet mask or prefix length to define the network part.
default-router *address*	Specifies the IP address of the default router for a DHCP client.
domain-name *domain*	Specifies the domain name for the DHCP client.
lease {*days* [*hours*] [*minutes*] | *infinite*}	Specifies the duration of the lease. The default is a one-day lease.
ip dhcp excluded-address *ip-address* [*last-ip-address*]	Specifies a single excluded IP address or range of addresses that a DHCP server should not assign to DHCP clients. Use it in the global configuration mode.

Discovery 22: Configure a Cisco Router as a DHCP Server

Introduction

This discovery lab will guide you through DHCP services using Cisco IOS devices. Review the topology diagram. R1 is configured to route between VLANs 1, 2, and 3. SRV1 and PC2 are on VLAN 2, while PC1 and PC3 are on VLAN 3. SRV1 is configured with a static IP address. PC1, PC2, and PC3 initially have no IP configuration. During this discovery lab, you will configure SRV1 as a DHCP server for its local VLAN—VLAN 2. You will then configure PC2 as a DHCP client and observe the DHCP process. Next, you will configure a second DHCP pool on SRV1. The pool will be applicable to VLAN 3. You will configure DHCP relay on R1 so that it will forward DHCP requests from VLAN 3 to SRV1. You will then configure PC1 and PC3 as DHCP clients and observe the DHCP process with R1 as a DHCP relay.

Topology

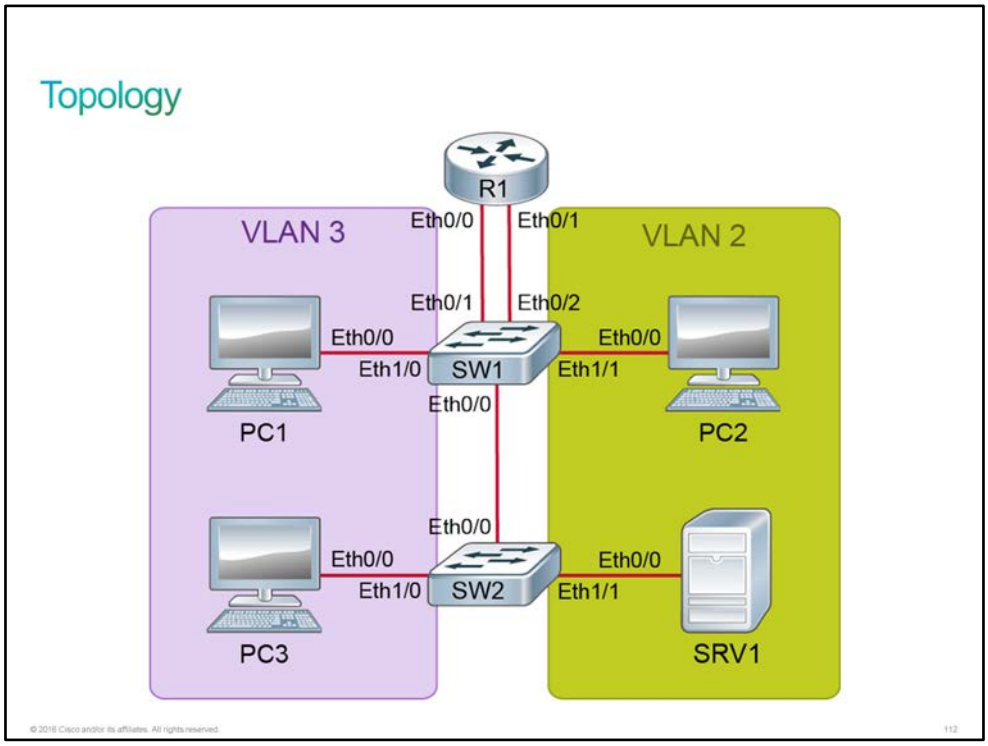

Job Aid

Device Information

Device Information Table

Device	Characteristic	Value
R1	Hostname	R1

Device	Characteristic	Value
R1	Ethernet0/0 description	Link to SW1 VLAN 1
R1	Ethernet0/0 IP address	10.10.1.1/24
R1	Ethernet0/1.2 description	Link to SW1 VLAN 2
R1	Ethernet0/1.2 IP address	10.10.2.1/24
R1	Ethernet0/1.3 description	Link to SW1 VLAN 3
R1	Ethernet0/1.3 IP address	10.10.3.1/24
R1	Loopback0 IP address	10.10.99.1/24
PC1	Hostname	PC1
PC2	Hostname	PC2
PC3	Hostname	PC3
SRV1	Hostname	SRV1
SRV1	IP address	10.10.2.40/24
SRV1	IP Default gateway	10.10.2.1
SW1	Hostname	SW1
SW1	VLAN 1 IP address	10.10.1.4/24
SW1	IP default gateway	10.10.1.1
SW1	Ethernet0/0 description	Link to SW2
SW1	Ethernet0/1 description	Link to R1 VLAN 1
SW1	Ethernet0/2 description	Link to R1 VLAN 2 and VLAN 3
SW1	Ethernet1/0 description	Link to PC1
SW1	Ethernet1/1 description	Link to PC2
SW2	Hostname	SW2
SW2	VLAN 1 IP address	10.10.1.5/24
SW2	IP default gateway	10.10.1.1
SW2	Ethernet0/0 description	Link to SW1

Device	Characteristic	Value
SW2	Ethernet1/0 description	Link to PC3
SW2	Ethernet1/1 description	Link to PC4

PCs in the virtual lab environment are simulated as routers, so you should use Cisco IOS commands to configure them or make verifications.

Task 1: Configure a Cisco Router as a DHCP Server

Activity

Step 1 Define a DHCP pool on SRV1. Name the pool Subnet_10.10.2 and configure it to serve the 10.10.2.0/24 subnet, specifying 10.10.2.1 as the default gateway and a lease of 7 days. Also, exclude the first 40 addresses from the pool (exclude 10.10.2.1–40). When you are done, leave the global configuration mode.

Enter the following commands to SRV1:

```
SRV1(config)# ip dhcp pool Subnet_10.10.2
SRV1(dhcp-config)# network 10.10.2.0 /24
SRV1(dhcp-config)# default-router 10.10.2.1
SRV1(dhcp-config)# lease 7
SRV1(dhcp-config)# exit
SRV1(config)# ip dhcp excluded-address 10.10.2.1 10.10.2.40
SRV1(config)# end
SRV1#
```

The most commonly used commands are abbreviated in this guided discovery. For example, **conf t** is the abbreviation for **configure terminal**. You can also attempt **tab** completion of commands to see the full commands during the discovery execution. For example, **conf<tab> t<tab>** would expand to **configure terminal**.

Step 2 Enable DHCP server packet debugging on SRV1 to provide visibility into the DHCP process.

Enter the following command to SRV1:

```
SRV1# debug ip dhcp server packet
DHCP server packet debugging is on.
```

Step 3 Before starting the configuration process on PC2, display the interface configuration and routing table to verify that there is no IP configuration for Ethernet0/0.

Enter the following commands to PC2:

```
PC2# sh ip int brief
Interface                  IP-Address      OK? Method Status
Protocol
Ethernet0/0                unassigned      YES NVRAM  administratively down
down
Ethernet0/1                unassigned      YES NVRAM  administratively down
down
Ethernet0/2                unassigned      YES NVRAM  administratively down
down
Ethernet0/3                unassigned      YES NVRAM  administratively down
down

PC2# sh ip route
Default gateway is not set

Host               Gateway           Last Use    Total Uses  Interface
ICMP redirect cache is empty
```

Step 4 On PC2, configure Ethernet0/0 to obtain its IP address via DHCP. Also, define the string "PC2" as its DHCP client ID. In a later step, you will compare this part to the Cisco IOS default client ID. Leave the interface in the shutdown state until the next step.

Enter the following commands to PC2:

```
PC2# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
PC2(config)# int e0/0
PC2(config-if)# ip address dhcp
PC2(config-if)# ip dhcp client client-id ascii PC2
```

Step 5 Enable Ethernet0/0 and wait approximately 15 seconds. You should see a syslog message indicating that Ethernet0/0 has received an IP address via DHCP.

Enter the following command to PC2:

```
PC2(config-if)# no shut
PC2(config-if)#
*Nov  3 13:56:59.238: %LINK-3-UPDOWN: Interface Ethernet0/0, changed state to
up
*Nov  3 13:57:00.242: %LINEPROTO-5-UPDOWN: Line protocol on Interface
Ethernet0/0, changed state to up
PC2(config-if)#
*Nov  3 13:57:14.158: %DHCP-6-ADDRESS_ASSIGN: Interface Ethernet0/0 assigned
DHCP address 10.10.2.41, mask 255.255.255.0, hostname PC2
```

Step 6 Leave the configuration mode and display the status of Ethernet0/0 and the routing table on PC2 to verify the results of the DHCP process.

Enter the following commands to PC2:


```
PC2(config-if)# end
PC2# show ip int brief
Interface                  IP-Address      OK? Method Status
Protocol
Ethernet0/0                10.10.2.41      YES DHCP   up                    up
Ethernet0/1                unassigned      YES NVRAM  administratively down
down
Ethernet0/2                unassigned      YES NVRAM  administratively down
down
Ethernet0/3                unassigned      YES NVRAM  administratively down
down
PC2# show ip route
Default gateway is 10.10.2.1

Host               Gateway           Last Use    Total Uses  Interface
ICMP redirect cache is empty
```

Step 7 Display the DHCP lease information on PC2.

Enter the following command to PC2:

```
PC2# show dhcp lease
Temp IP addr: 10.10.2.41  for peer on Interface: Ethernet0/0
Temp  sub net mask: 255.255.255.0
   DHCP Lease server: 10.10.2.40, state: 5 Bound
   DHCP transaction id: 9FB
   Lease: 604800 secs,  Renewal: 302400 secs,  Rebind: 529200 secs
Temp default-gateway addr: 10.10.2.1
   Next timer fires after: 3d11h
   Retry count: 0   Client-ID: PC2
   Client-ID hex dump: 504332
   Hostname: PC2
```

The lease is 604,800 seconds, which is equivalent to 7 days.

The client ID is represented in hexadecimal notation. In the ASCII table, the hexadecimal value 50 (decimal value 80) is capital "P." Similarly, hexadecimal 43 is capital "C" and hexadecimal 32 is the numeral "2."

Step 8 Ping the loopback interface of R1 (10.10.99.1) to verify IP and routing functionality from PC2. The ping attempt should succeed.

Enter the following command to PC2:

```
PC2# ping 10.10.99.1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.10.99.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/201/1005 ms
```

Step 9 Before moving back to SRV1, display the MAC address of Ethernet0/0 on PC2.

Enter the following command to PC2:

```
PC2# show int e0/0 | inc address
  Hardware is AmdP2, address is aabb.cc00.0300 (bia aabb.cc00.0300)
  Internet address is 10.10.2.41/24
```

MAC address may be different in your lab setup.

Step 10 Return to the console of SRV1. Review the sequence of debug messages waiting for you.

Note that even if you do not understand all the debug output, you can often find answers by studying the details. Notice the messages that indicate the DHCP DORA process that is associated with the client ID 504332 (ASCII PC2 in hexadecimal notation) with the MAC address aabb.cc00.0300.

```
SRV1#
*Nov  3 13:57:03.030: DHCPD: client's VPN is .
*Nov  3 13:57:03.030: DHCPD: No option 125
*Nov  3 13:57:03.030: DHCPD: DHCPDISCOVER received from client 0050.4332 on
interface Ethernet0/0.
*Nov  3 13:57:03.030: DHCPD: Allocate an address without class information
(10.10.2.0)
*Nov  3 13:57:06.038: DHCPD: Saving workspace (ID=0xAE000001)
*Nov  3 13:57:06.109: DHCPD: New packet workspace 0xEFAA5460 (ID=0x15000002)
*Nov  3 13:57:06.109: DHCPD: client's VPN is .
*Nov  3 13:57:06.109: DHCPD: No option 125
*Nov  3 13:57:06.109: DHCPD: DHCPDISCOVER received from client 0050.4332 on
interface Ethernet0/0.
*Nov  3 13:57:10.056: DHCPD: Reprocessing saved workspace (ID=0xAE000001)
*Nov  3 13:57:10.056: DHCPD: DHCPDISCOVER received from client 0050.4332 on
interface Ethernet0/0.
*Nov  3 13:57:10.057: DHCPD: Sending DHCPOFFER to client 0050.4332
(10.10.2.41).
*Nov  3 13:57:10.057: DHCPD: no option 125
*Nov  3 13:57:10.057: DHCPD: broadcasting BOOTREPLY to client aabb.cc00.0300.
*Nov  3 13:57:10.057: DHCPD: client's VPN is .
*Nov  3 13:57:10.057: DHCPD: No option 125
*Nov  3 13:57:10.057: DHCPD: DHCPREQUEST received from client 0050.4332.
*Nov  3 13:57:10.057: DHCPD: No default domain to append - abort update
*Nov  3 13:57:10.057: DHCPD: Sending DHCPACK to client 0050.4332 (10.10.2.41).
*Nov  3 13:57:10.057: DHCPD: no option 125
*Nov  3 13:57:10.057: DHCPD: broadcasting BOOTREPLY to client aabb.cc00.0300.
```

DHCP operations fall into four phases: server discovery, IP lease offer, IP request, and IP lease acknowledgement. These stages are often abbreviated as DORA for discovery, offer, request, and acknowledgement.

There are three DHCP discovery messages before DHCP offer is sent.

Monitoring DHCP Server Functions

You can verify the configured DHCP parameters using the **show ip dhcp pool** command in the privileged EXEC mode. The output displays the total number of available addresses, the configured address range, and number of leased addresses. Keep in mind that the total address number does not take the excluded IP addresses into account.

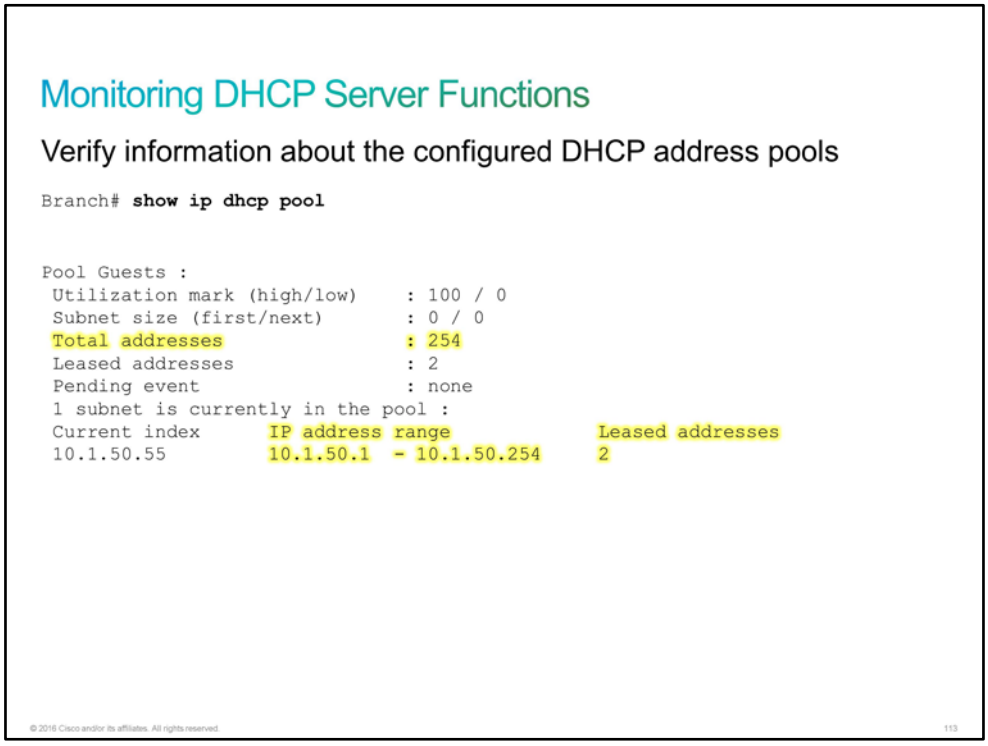

For more details about the **show ip dhcp pool** command, refer to *Cisco IOS IP Addressing Services Command Reference* at http://www.cisco.com/en/US/docs/ios-xml/ios/ipaddr/command/ipaddr-r1.html#GUID-23A47402-6EB5-4945-8DEB-ABCB7BCF3D68.

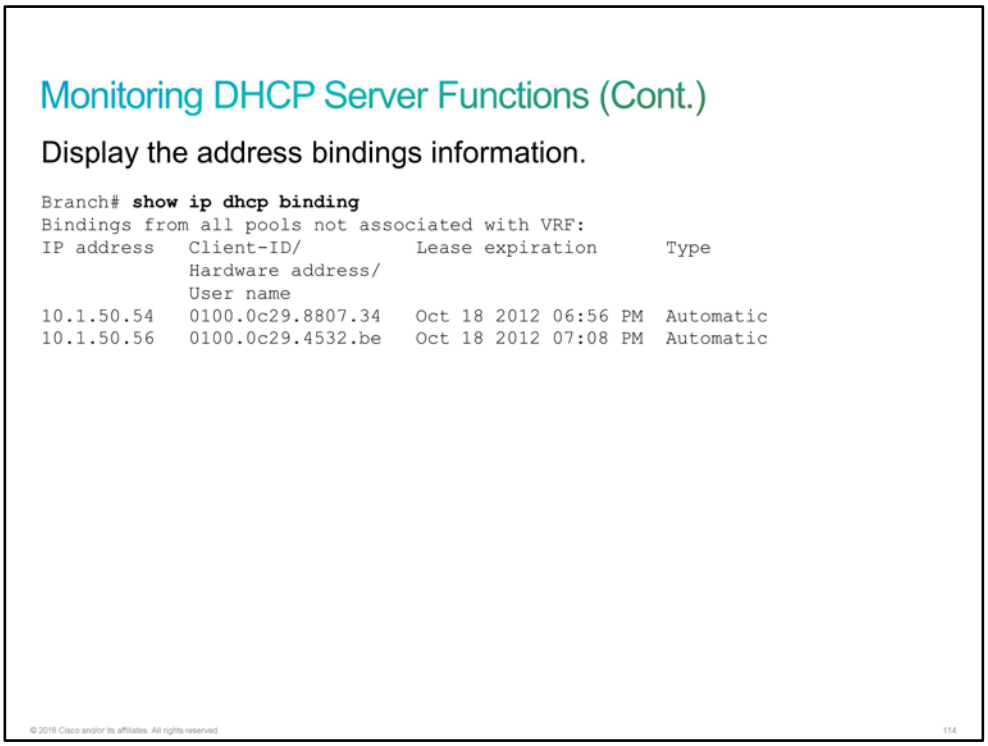

To verify the operation of a DHCP server, use the **show ip dhcp binding** command, which displays a list of all IP address-to-MAC address bindings that the DHCP service has provided. Also, the lease expiration time and type of DHCP allocation are listed.

Cisco Device as a DHCP Relay Agent

Because DHCP clients do not have information about the network to which they are attached, they use UDP broadcasts to send their initial DHCP discover message. Routers do not normally forward UDP broadcasts because they are configured to not forward broadcast traffic. A primary goal of router configuration is to control unnecessary proliferation of broadcast packets. It means that a DHCP client may be unable to obtain its IP address or other configuration parameters from a DHCP server if the DHCP server is on a different subnet from the DHCP client.

It is common, and often desirable, for DHCP servers and clients to reside on different subnets. Instead of using a different DHCP server for each subnet, you might want to use a centralized DHCP solution. A centralized DHCP server enables you to manage IP address assignment in one place for an entire organization. This solution is less time consuming, and it reduces the chance of human error.

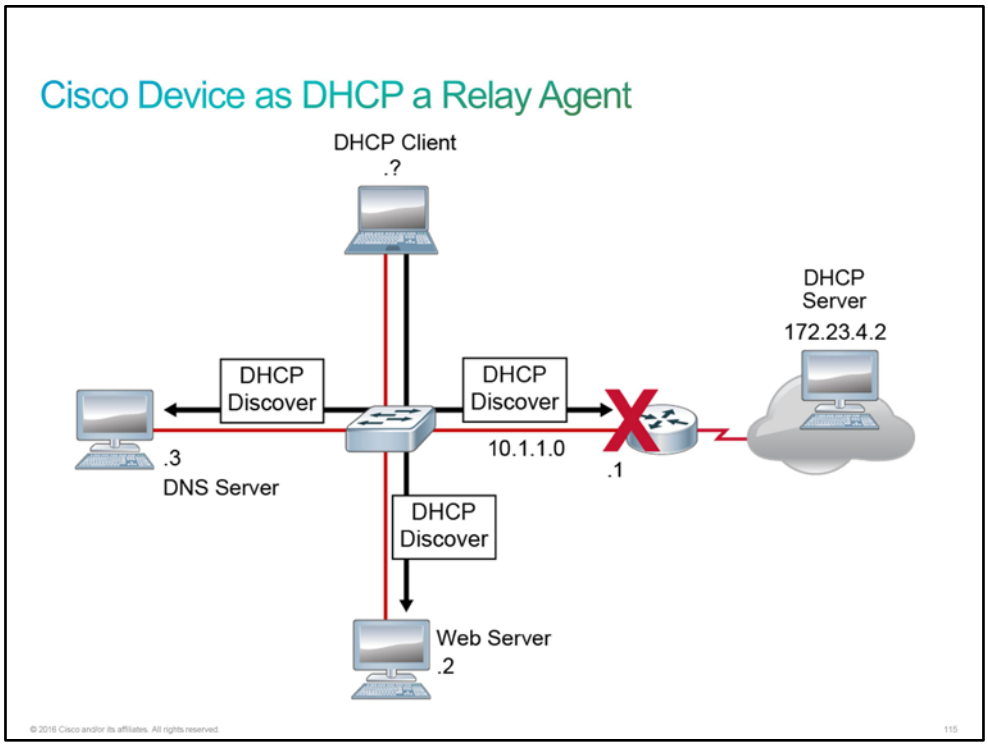

In this scenario, the DHCP process is as follows:

1. The DHCP client transmits a broadcast DHCP discover message.
2. When the router receives the broadcast DHCP discover message, it transmits it to the DHCP server, after it stores one of its own IP addresses in the gateway IP address (the giaddr field) of the DHCP packet. The address that the router stores in the giaddr field is the address of the interface on which the router received the DHCP discover message.
3. The DHCP server uses the giaddr to determine the subnet of the IP address that it will offer to the client. The DHCP server then reserves an appropriate IP address for the client and sends a DHCP Offer message to the giaddr address.
4. The DHCP relay agent on the router receives the DHCP offer message and retransmits it to the DHCP client.
5. If the IP address is acceptable, the DHCP client sends a DHCP request message indicating its intent to accept the parameters in the DHCP offer.
6. The router forwards the DHCP request message to the DHCP server.
7. When the DHCP server receives the DHCP request message, it marks the IP address as "in use" in its database and sends a DHCP ACK message. The message includes a confirmation for all requested parameters.
8. The router forwards the DHCP ACK message to the DHCP client.

At this point, the IP configuration process is complete.

To enable the DHCP relay agent on a Cisco device, use the **ip helper-address** command to configure a helper address on the router interface that is connected to the client. The helper address should specify the IP address of the DHCP server, as the figure shows:

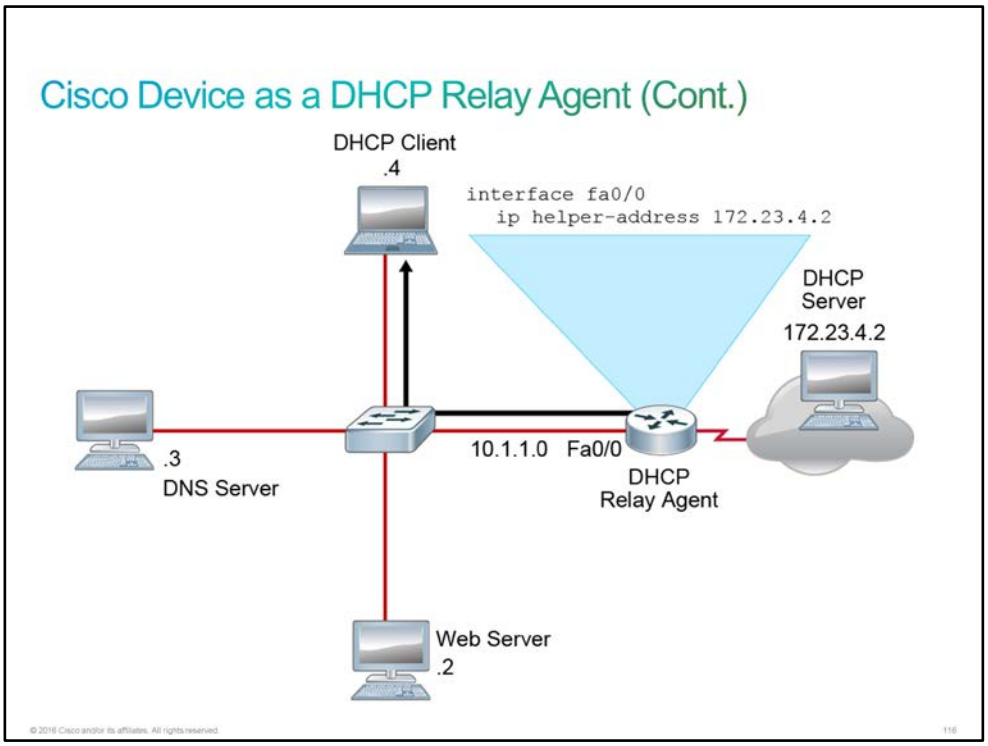

Step 11 Display the DHCP server binding tables on SRV1. Note that the IP address is assigned to the client ID that you configured on PC2 (504332 is the hexadecimal representation of the ASCII string "PC2").

Enter the following commands to SRV1:

```
SRV1# sh ip dhcp binding
Bindings from all pools not associated with VRF:
IP address          Client-ID/              Lease expiration        Type
                    Hardware address/
                    User name
10.10.2.41          0050.4332               Nov 10 2015 05:57 AM    Automatic
```

Step 12 Configure a second DHCP pool on SRV1. This pool will be applied to VLAN 3, using the network 10.10.3.0/24 with 10.10.3.1 as the default gateway. As you did for VLAN 2, specify a 7-day lease and exclude the first 40 addresses from the pool.

Enter the following commands to SRV1:


```
SRV1# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
SRV1(config)# ip dhcp pool Subnet_10.10.3
SRV1(dhcp-config)# network 10.10.3.0 /24
SRV1(dhcp-config)# default-router 10.10.3.1
SRV1(dhcp-config)# lease 7
SRV1(dhcp-config)# exit
SRV1(config)# ip dhcp excluded-address 10.10.3.1 10.10.3.40
SRV1(config)# end
SRV1#
```

Step 13 On the R1 router, configure the Ethernet0/1.3 interface on VLAN 3 to forward DHCP requests to SRV1.

Enter the following commands to the R1 router:

```
R1# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
R1(config)# int e0/1.3
R1(config-subif)# ip helper-address 10.10.2.40
R1(config-subif)# end
R1#
```

Step 14 On PC1, configure its Ethernet0/0 interface to obtain its IP address via DHCP, enable the interface, and wait for the DHCP assignment to complete.

Enter the following commands to PC1:

```
PC1# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
PC1(config)# int e0/0
PC1(config-if)# ip address dhcp
PC1(config-if)# no shut
*Nov  3 14:25:59.950: %DHCP-6-ADDRESS_ASSIGN: Interface Ethernet0/0 assigned
DHCP address 10.10.3.41, mask 255.255.255.0, hostname PC1
```

Step 15 On PC1, leave the configuration mode and display the interface status, routing table, and DHCP lease information.

Enter the following commands to PC1:

```
PC1(config-if)# end
PC1# show ip int brief
Interface                  IP-Address      OK? Method Status
Protocol
Ethernet0/0                10.10.3.41      YES DHCP   up                    up
Ethernet0/1                unassigned      YES NVRAM  administratively down
down
Ethernet0/2                unassigned      YES NVRAM  administratively down
down
Ethernet0/3                unassigned      YES NVRAM  administratively down
down
PC1# show ip route
Default gateway is 10.10.3.1

Host              Gateway           Last Use    Total Uses  Interface
ICMP redirect cache is empty
PC1# show dhcp lease
Temp IP addr: 10.10.3.41  for peer on Interface: Ethernet0/0
Temp  sub net mask: 255.255.255.0
   DHCP Lease server: 10.10.2.40, state: 5 Bound
   DHCP transaction id: 73E
   Lease: 604800 secs,  Renewal: 302400 secs,  Rebind: 529200 secs
Temp default-gateway addr: 10.10.3.1
   Next timer fires after: 3d11h
   Retry count: 0   Client-ID: cisco-aabb.cc00.0200-Et0/0
   Client-ID hex dump: 636973636F2D616162622E636330302E
                       303230302D4574302F30
   Hostname: PC1
```

You did not configure a client ID on PC1. The default client ID on Cisco IOS device is the string "cisco", followed by the MAC address of the interface, followed by a string that specifies the interface. The hexadecimal notation for this ID is complex.

Step 16 Verify connectivity from PC1 by attempting to ping the Loopback interface of R1 (10.10.99.1). Also ping the dynamic address that was assigned to PC2 (10.10.2.41). The ping attempts should succeed.

Enter the following commands to PC1:

```
PC1# ping 10.10.99.1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.10.99.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/201/1003 ms
PC1# ping 10.10.2.41
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.10.2.41, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
```

Step 17 Return to SRV1 to view the debug messages that were produced during the DHCP process.

The debug messages reveal the normal DORA process, but this time using the relay agent 10.10.3.1. The relay agent is the R1 interface on VLAN 3. The DHCP server can identify the appropriate pool to use based on the IP address of the relay agent.


```
SRV1#
*Nov  3 14:25:53.320: DHCPD: client's VPN is .
*Nov  3 14:25:53.320: DHCPD: No option 125
*Nov  3 14:25:53.320: DHCPD: DHCPDISCOVER received from client
0063.6973.636f.2d61.6162.622e.6363.3030.2e30.3230.302d.4574.302f.30 through
relay 10.10.3.1.
*Nov  3 14:25:53.320: DHCPD: Allocate an address without class information
(10.10.3.0)
*Nov  3 14:25:54.322: DHCPD: Saving workspace (ID=0x15000002)
*Nov  3 14:25:55.835: DHCPD: Reprocessing saved workspace (ID=0x15000002)
*Nov  3 14:25:55.835: DHCPD: DHCPDISCOVER received from client
0063.6973.636f.2d61.6162.622e.6363.3030.2e30.3230.302d.4574.302f.30 through
relay 10.10.3.1.
*Nov  3 14:25:55.835: DHCPD: Sending DHCPOFFER to client
0063.6973.636f.2d61.6162.622e.6363.3030.2e30.3230.302d.4574.302f.30
(10.10.3.41).
*Nov  3 14:25:55.835: DHCPD: no option 125
*Nov  3 14:25:55.835: DHCPD: unicasting BOOTREPLY for client aabb.cc00.0200 to
relay 10.10.3.1.
*Nov  3 14:25:55.836: DHCPD: New packet workspace 0xEFAAC900 (ID=0xF5000003)
*Nov  3 14:25:55.836: DHCPD: client's VPN is .
*Nov  3 14:25:55.836: DHCPD: No option 125
*Nov  3 14:25:55.836: DHCPD: DHCPREQUEST received from client
0063.6973.636f.2d61.6162.622e.6363.3030.2e30.3230.302d.4574.302f.30.
*Nov  3 14:25:55.836: DHCPD: No default domain to append - abort update
*Nov  3 14:25:55.836: DHCPD: Sending DHCPACK to client
0063.6973.636f.2d61.6162.622e.6363.3030.2e30.3230.302d.4574.302f.30
(10.10.3.41).
*Nov  3 14:25:55.836: DHCPD: no option 125
*Nov  3 14:25:55.836: DHCPD: unicasting BOOTREPLY for client aabb.cc00.0200 to
relay 10.10.3.1.
```

There are two DHCP discovery messages before DHCP offer is sent.

Step 18 Display the DHCP server binding tables on SRV1. You should see that it has leased two addresses at this point.

Enter the following command to SRV1:

```
SRV1# sh ip dhcp binding
Bindings from all pools not associated with VRF:
IP address          Client-ID/              Lease expiration        Type
                    Hardware address/
                    User name
10.10.2.41          0050.4332               Nov 10 2015 05:57 AM    Automatic
10.10.3.41          0063.6973.636f.2d61.    Nov 10 2015 06:25 AM    Automatic
                    6162.622e.6363.3030.
                    2e30.3230.302d.4574.
                    302f.30
```

Step 19 Disable the debug output on SRV1.

Enter the following command to SRV1:

```
SRV1# undebug all
All possible debugging has been turned off
```

Step 20 In the next several steps, you will configure PC3 as a DHCP client and verify the results. Access the console of PC3. Configure its Ethernet0/0 interface for DHCP and leave the configuration mode.

Enter the following commands to PC3:

```
PC3# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
PC3(config)# int e0/0
PC3(config-if)# ip address dhcp
PC3(config-if)# no shut
*Nov  3 14:39:19.095: %LINK-3-UPDOWN: Interface Ethernet0/0, changed state to
up
*Nov  3 14:39:20.099: %LINEPROTO-5-UPDOWN: Line protocol on Interface
Ethernet0/0, changed state to up
PC3(config-if)# end
*Nov  3 14:39:26.056: %SYS-5-CONFIG_I: Configured from console by console
*Nov  3 14:39:31.458: %DHCP-6-ADDRESS_ASSIGN: Interface Ethernet0/0 assigned
DHCP address 10.10.3.42, mask 255.255.255.0, hostname PC3
```

Step 21 On PC3, display the interface configuration and routing table.

Enter the following commands to PC3:

```
PC3# sh ip int brie
Interface                  IP-Address      OK? Method Status
Protocol
Ethernet0/0                10.10.3.42      YES DHCP   up                    up
Ethernet0/1                unassigned      YES NVRAM  administratively down
down
Ethernet0/2                unassigned      YES NVRAM  administratively down
down
Ethernet0/3                unassigned      YES NVRAM  administratively down
down
PC3# sh ip route
Default gateway is 10.10.3.1

Host               Gateway           Last Use    Total Uses  Interface
ICMP redirect cache is empty
```

Step 22 On PC3, display the DHCP lease information.

Enter the following command to PC3:

```
PC3# sh dhcp lease
Temp IP addr: 10.10.3.42  for peer on Interface: Ethernet0/0
Temp  sub net mask: 255.255.255.0
   DHCP Lease server: 10.10.2.40, state: 5 Bound
   DHCP transaction id: ADD
   Lease: 604800 secs,  Renewal: 302400 secs,  Rebind: 529200 secs
Temp default-gateway addr: 10.10.3.1
   Next timer fires after: 3d11h
   Retry count: 0   Client-ID: cisco-aabb.cc00.0c00-Et0/0
   Client-ID hex dump: 636973636F2D616162622E636330302E
                       306330302D4574302F30
   Hostname: PC3
```

Step 23 Verify that PC3 can ping the dynamic IP address of PC2 (10.10.2.41).

Enter the following command to PC3:

```
PC3# ping 10.10.2.41
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.10.2.41, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/202/1004 ms
```

Step 24 On SRV1, display the DHCP server binding table. You should now see that it has leased three IP addresses, including one for PC3.

Enter the following command to SRV1:

```
SRV1# sh ip dhcp binding
Bindings from all pools not associated with VRF:
IP address          Client-ID/              Lease expiration        Type
                    Hardware address/
                    User name
10.10.2.41          0050.4332               Nov 10 2015 05:57 AM    Automatic
10.10.3.41          0063.6973.636f.2d61.    Nov 10 2015 06:25 AM    Automatic
                    6162.622e.6363.3030.
                    2e30.3230.302d.4574.
                    302f.30
10.10.3.42          0063.6973.636f.2d61.    Nov 10 2015 06:39 AM    Automatic
                    6162.622e.6363.3030.
                    2e30.6330.302d.4574.
                    302f.30
```

During this discovery lab, you have explored the DHCP client, server, and relay features on Cisco IOS devices. SRV1 was configured as a DHCP server with two pools, one for its local subnet (VLAN 2) and one for a remote subnet (VLAN 3). R1 was configured as a DHCP relay agent to forward the DHCP requests broadcast on VLAN 3 to SRV1. You configured PC1, PC2, and PC3 as DHCP clients. Feel free to continue with an independent exploration of the use of DHCP within the lab environment.

This is the end of the discovery lab.

Understanding DNS

DNS provides an efficient way to convert human-readable names of IP end systems into machine-readable IP addresses that are necessary for routing.

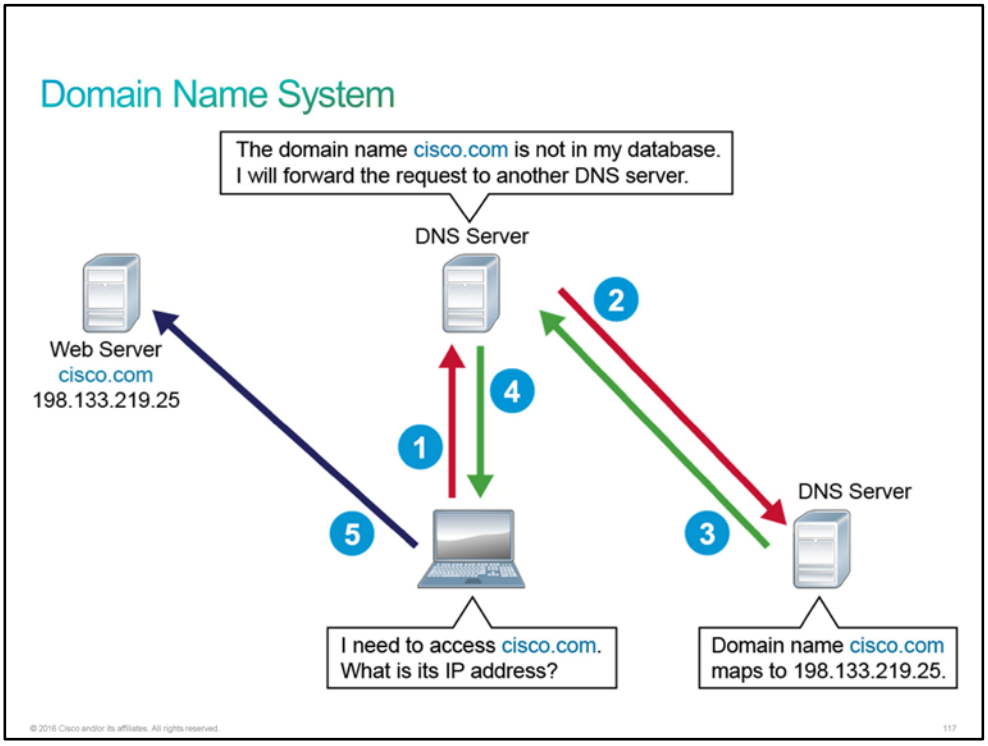

In data networks, devices are labeled with numeric IP addresses so that they can send and receive messages over the network. However, most people find it difficult to remember this numeric address. Therefore, domain names were created to convert the numeric address into a simple, recognizable name.

DNS was created for domain name-to-address resolution for networks. DNS uses a set of servers to resolve the names that are associated with numbered addresses. The DNS protocol defines an automated service that matches resource names with the required numeric network address.

Discovery 23: Troubleshoot DHCP Issues

Introduction

This discovery lab will guide you through DHCP services troubleshooting using Cisco IOS tools. R1 is configured to route between VLANs 1, 2, and 3. SRV1 and PC2 are on VLAN 2, while PC1 and PC3 are on VLAN 3. SRV1 is configured with a static IP address. PC1, PC2, and PC3 are configured as DHCP clients. SRV1 is configured as a DHCP server with two pools, one for its local subnet (VLAN 2) and one for a remote subnet (VLAN 3). R1 is configured as a DHCP relay agent to forward DHCP request broadcasts on VLAN 3 to SRV1.

There are a couple of mistakes in the initial configuration and it is up to you to troubleshoot them.

Topology

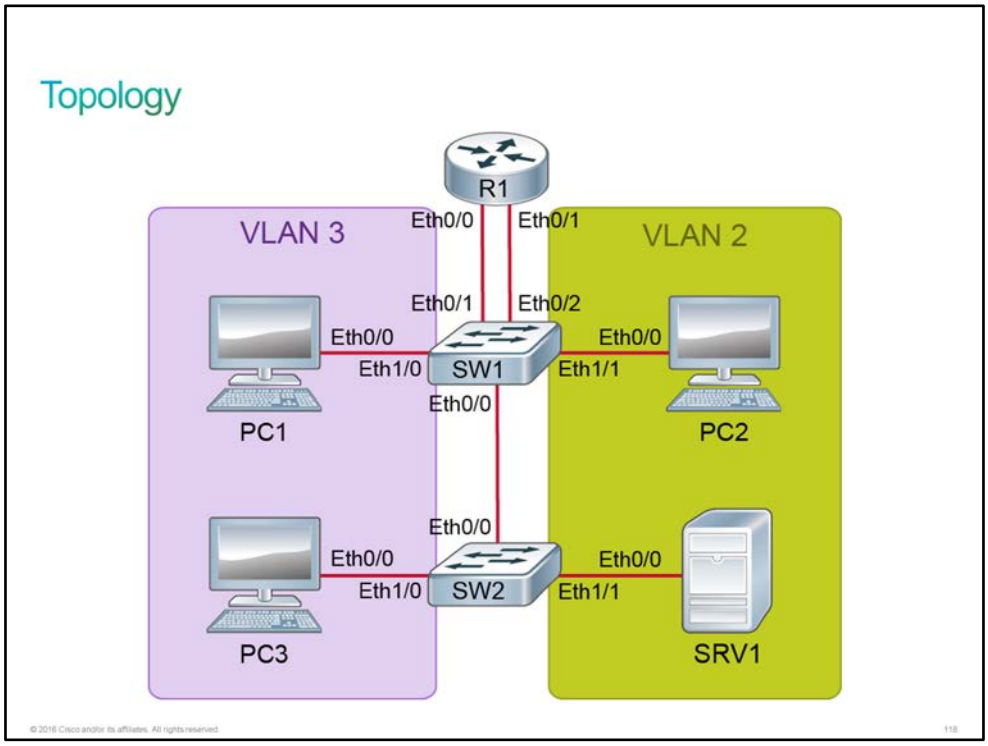

Job Aid

Device Information

Device Information Table

Device	Characteristic	Value
R1	Hostname	R1
R1	Ethernet0/0 description	Link to SW1 VLAN 1

Device	Characteristic	Value
R1	Ethernet0/0 IP address	10.10.1.1/24
R1	Ethernet0/1.2 description	Link to SW1 VLAN 2
R1	Ethernet0/1.2 IP address	10.10.2.1/24
R1	Ethernet0/1.3 description	Link to SW1 VLAN 3
R1	Ethernet0/1.3 IP address	10.10.3.1/24
R1	Loopback0 IP address	10.10.99.1/24
PC1	Hostname	PC1
PC1	IP address	DHCP
PC1	IP default gateway	DHCP
PC2	Hostname	PC2
PC2	IP address	DHCP
PC2	IP default gateway	DHCP
PC3	Hostname	PC3
PC3	IP address	DHCP
PC3	IP default gateway	DHCP
PC4	Hostname	PC4
PC4	IP address	10.10.2.40/24
PC4	IP default gateway	10.10.2.1
SW1	Hostname	SW1
SW1	VLAN 1 IP address	10.10.1.4/24
SW1	IP default gateway	10.10.1.1
SW1	Ethernet0/0 description	Link to SW2
SW1	Ethernet0/1 description	Link to R1 VLAN 1
SW1	Ethernet0/2 description	Link to R1 VLAN 2 and VLAN 3
SW1	Ethernet1/0 description	Link to PC1

Device	Characteristic	Value
SW1	Ethernet1/1 description	Link to PC2
SW2	Hostname	SW2
SW2	VLAN 1 IP Address	10.10.1.5/24
SW2	IP default gateway	10.10.1.1
SW2	Ethernet0/0 description	Link to SW1
SW2	Ethernet1/0 description	Link to PC3
SW2	Ethernet1/1 description	Link to PC4

PCs in the virtual lab environment are simulated as routers, so you should use Cisco IOS commands to configure them or make verifications.

Task 1: Troubleshooting DHCP Issues

Activity

Step 1 From PC2, verify connectivity to SRV1 (10.10.2.40). PC2 and SRV1 are in the same VLAN (VLAN 20).

Enter the following command to PC2:

```
PC2# ping 10.10.2.40
% Unrecognized host or address, or protocol not running.
```

Step 2 On PC2, verify the Ethernet0/0 interface IP address.

With the **include** option, you can use the | symbol for a logical OR between two strings. Enter the following command to PC2:

```
PC2# sh ip int brie | inc Interface|0/0
Interface                  IP-Address      OK? Method Status
Protocol
Ethernet0/0                unassigned      YES DHCP   up                    up
```

The Ethernet 0/0 method to learn an IP address is DHCP, but there is no IP address assigned.

Step 3 Display the DHCP lease information on PC2.

Enter the following command to PC2:

```
PC2# show dhcp lease
Temp IP addr: 0.0.0.0  for peer on Interface: Ethernet0/0
Temp  sub net mask: 0.0.0.0
   DHCP Lease server: 0.0.0.0, state: 3 Selecting
   DHCP transaction id: 10AF
   Lease: 0 secs,  Renewal: 0 secs,  Rebind: 0 secs
   Next timer fires after: 00:00:04
   Retry count: 2   Client-ID: cisco-aabb.cc00.1700-Et0/0
   Client-ID hex dump: 636973636F2D616162622E636330302E
                       313730302D4574302F30
   Hostname: PC2
```

There is no DHCP lease information on PC2. You may also get empty output.

```
PC2# show dhcp lease
PC2#
```

Step 4 Display the DHCP server binding table on SRV1. Note that no IP address is assigned.

Enter the following commands to SRV1:

```
SRV1# sh ip dhcp binding
Bindings from all pools not associated with VRF:
IP address          Client-ID/              Lease expiration        Type
                    Hardware address/
                    User name
```

Step 5 On SRV1, verify the information about the configured DHCP address pool and compare it to the DHCP configuration.

Enter the following commands to SRV1:

```
SRV1# sh ip dhcp pool

Pool Subnet_10.10.2 :
 Utilization mark (high/low)    : 100 / 0
 Subnet size (first/next)       : 0 / 0
 Total addresses                : 254
 Leased addresses               : 0
 Pending event                  : none
 1 subnet is currently in the pool :
 Current index        IP address range                    Leased addresses
 0.0.0.0              10.10.2.1        - 10.10.2.254      0
<... output omitted ...>
SRV1# sh run | sec dhcp
ip dhcp excluded-address 10.10.2.1 10.10.2.255
ip dhcp excluded-address 10.10.3.1 10.10.3.40
ip dhcp excluded-address 10.10.3.42 10.10.3.255
ip dhcp pool Subnet_10.10.2
 network 10.10.2.0 255.255.255.0
 default-router 10.10.2.40
 lease 7
ip dhcp pool Subnet_10.10.3
 network 10.10.30.0 255.255.255.0
 default-router 10.10.3.1
 lease 7
```


In the configuration, you can see that all IP addresses are excluded from the DHCP pool Subnet_10.10.2.

Also note that the wrong default router IP address is used in the DHCP pool Subnet_10.10.2

Step 6 On SRV1, correct the DHCP pool of excluded addresses for the DHCP pool Subnet_10.10.2. Exclude the IP addresses 10.10.2.1 to 10.10.2.40. Change the default router IP address in the DHCP pool Subnet_10.10.2 into 10.10.2.1.

Enter the following commands to SRV1:

```
SRV1# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
SRV1(config)# no ip dhcp excluded-address 10.10.2.1 10.10.2.255
SRV1(config)# ip dhcp excluded-address 10.10.2.1 10.10.2.40
SRV1(config)# ip dhcp pool Subnet_10.10.2
SRV1(dhcp-config)# no default-router 10.10.2.40
SRV1(dhcp-config)# default-router 10.10.2.1
```

Step 7 On PC2, disable and re-enable the Ethernet0/0 interface. Observe the DHCP message on the console.

Enter the following commands to PC2:

```
PC2# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
PC2(config)# int e 0/0
PC2(config-if)# sh
*Nov  4 08:52:47.261: %LINK-5-CHANGED: Interface Ethernet0/0, changed state to
administratively down
*Nov  4 08:52:48.261: %LINEPROTO-5-UPDOWN: Line protocol on Interface
Ethernet0/0, changed state to down
PC2(config-if)# no sh
*Nov  4 08:52:53.925: %LINK-3-UPDOWN: Interface Ethernet0/0, changed state to
up
*Nov  4 08:52:54.929: %LINEPROTO-5-UPDOWN: Line protocol on Interface
Ethernet0/0, changed state to up
*Nov  4 08:53:03.075: %DHCP-6-ADDRESS_ASSIGN: Interface Ethernet0/0 assigned
DHCP address 10.10.2.42, mask 255.255.255.0, hostname PC2
```

Step 8 Display the DHCP server binding table on SRV1. Note that one IP address is assigned.

Enter the following command to SRV1:

```
SRV1# sh ip dhc binding
Bindings from all pools not associated with VRF:
IP address          Client-ID/              Lease expiration        Type
                    Hardware address/
                    User name
10.10.2.42          0063.6973.636f.2d61.    Nov 11 2015 12:52 AM    Automatic
                    6162.622e.6363.3030.
                    2e30.3330.302d.4574.
                    302f.30
```

Step 9 Display the DHCP lease information on PC2.

Enter the following command to PC2:

```
PC2# show dhcp lease
Temp IP addr: 10.10.2.42  for peer on Interface: Ethernet0/0
Temp  sub net mask: 255.255.255.0
   DHCP Lease server: 10.10.2.40, state: 5 Bound
   DHCP transaction id: 618
   Lease: 604800 secs,  Renewal: 302400 secs,  Rebind: 529200 secs
Temp default-gateway addr: 10.10.2.1
   Next timer fires after: 3d11h
   Retry count: 0   Client-ID: cisco-aabb.cc00.0300-Et0/0
   Client-ID hex dump: 636973636F2D616162622E636330302E
                       303330302D4574302F30
   Hostname: PC2
```

PC2 has an IP address assigned from the correct pool. Also, the IP default gateway is correct (10.10.2.1.)

Step 10 Continue troubleshooting why PC1 and PC3 do not get DHCP responses from SRV1. PC1 and PC3 are both in VLAN 3, but they have no IP addresses is assigned to them.

Enter the following command to PC1:

```
PC1# sh ip int brie | in Interface|0/0
Interface                  IP-Address      OK? Method Status
Protocol
Ethernet0/0                unassigned      YES DHCP   up                    up
```

Enter the following command to PC3:

```
PC3# sh ip int brie | in Interface|0/0
Interface                  IP-Address      OK? Method Status
Protocol
Ethernet0/0                unassigned      YES DHCP   up                    up
```

Step 11 On the R1 router, verify the helper IP address configuration on Ethernet0/1.3.

Enter the following command to the R1 router:

```
R1# sh run int e0/1.3
Building configuration...

Current configuration : 152 bytes
!
interface Ethernet0/1.3
 description Link to SW1 VLAN 3
 encapsulation dot1Q 3
 ip address 10.10.3.1 255.255.255.0
 ip helper-address 10.10.3.40
end
```

Note that the IP helper address is incorrect.

Step 12 On the R1 router, correct the helper IP address on the Ethernet0/1.3 interface.

Enter the following commands to the R1 router:


```
R1# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
R1(config)# interface Ethernet0/1.3
R1(config-subif)# no ip helper-address 10.10.3.40
R1(config-subif)# ip helper-address 10.10.2.40
```

Verify that PC1 and PC3 are still not getting DHCP responses from SRV1.

```
PC1# sh ip int brie | in Interface|0/0
Interface                  IP-Address      OK? Method Status
Protocol
Ethernet0/0                unassigned      YES DHCP   up                    up


PC3# sh ip int brie | in Interface|0/0
Interface                  IP-Address      OK? Method Status
Protocol
Ethernet0/0                unassigned      YES DHCP   up                    up
```

Step 13 On SRV1, verify the information about the configured DHCP address pool and compare it to the DHCP configuration.

Enter the following commands to SRV1:

```
SRV1# sh ip dhcp pool
<... output omitted ...>
Pool Subnet_10.10.3 :
 Utilization mark (high/low)    : 100 / 0
 Subnet size (first/next)       : 0 / 0
 Total addresses                : 254
 Leased addresses               : 0
 Pending event                  : none
 1 subnet is currently in the pool :
 Current index        IP address range                    Leased addresses
 10.10.30.1           10.10.30.1       - 10.10.30.254      0
SRV1# sh run  | sec dhcp
ip dhcp excluded-address 10.10.3.1 10.10.3.40
ip dhcp excluded-address 10.10.3.42 10.10.3.255
ip dhcp excluded-address 10.10.2.1 10.10.2.40
ip dhcp pool Subnet_10.10.2
 network 10.10.2.0 255.255.255.0
 default-router 10.10.2.1
 lease 7
ip dhcp pool Subnet_10.10.3
 network 10.10.30.0 255.255.255.0
 default-router 10.10.3.1
 lease 7
```

An incorrect subnet is configured in the DHCP pool Subnet_10.10.3. Only one IP address is allowed in the subnet 10.10.3.0/24.

Step 14 On SRV1, correct the DHCP pool of excluded addresses for the DHCP pool Subnet_10.10.3. Exclude the IP addresses from 10.10.3.1 to 10.10.3.40. Change the network IP address in the DHCP pool from Subnet_10.10.3 into 10.10.3.0 /24.

Enter the following commands to SRV1:

```
SRV1# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
SRV1(config)# no ip dhcp excluded-address 10.10.3.42 10.10.3.255
SRV1(config)# ip dhcp pool Subnet_10.10.3
SRV1(dhcp-config)# no  network 10.10.30.0 255.255.255.0
SRV1(dhcp-config)# network 10.10.3.0 255.255.255.0
```

Step 15 It may take up to 60 seconds for PCs to get a DHCP response. Note that both PC1 and PC3 will get a DHCP response.

```
PC1#
*Nov  4 09:26:38.880: %DHCP-6-ADDRESS_ASSIGN: Interface Ethernet0/0 assigned
DHCP address 10.10.3.42, mask 255.255.255.0, hostname PC1
```

```
PC3#
*Nov  4 09:26:38.329: %DHCP-6-ADDRESS_ASSIGN: Interface Ethernet0/0 assigned
DHCP address 10.10.3.41, mask 255.255.255.0, hostname PC3
```

This is the end of the discovery lab.

Challenge

1. Why would you need a DHCP Server in an Enterprise environment?

   A. An Automatic IP Address assignment system can cause inconsistency across the whole organization
   B. An Automatic IP Address assignment system can automatically assign an IP Address in accordance with user VLAN Settings
   C. An Automatic IP Address assignment system is time consuming
   D. An Automatic IP Address assignment system is prone to errors

2. In which of the following phases does the actual decision happen of which IP Address will be assigned to a device?

   A. DHCP discover
   B. DHCP offer
   C. DHCP request
   D. DHCP acknowledgement

3. Which of the following commands is used to determine how much time the IP Address will be assigned to a device, when configuring a DHCP Server on a Cisco Router?

   A. **ip dhcp pool**
   B. **network**
   C. **lease**
   D. **domain-name**
   E. **clock**

4. Which of the following commands shows you the IP Addresses in use on a Cisco Router configured to be a DHCP Server?

   A. **show ip dhcp pool**
   B. **show ip dhcp bindings**
   C. **show ip dhcp assignments**
   D. **show ip dhcp conflicts**
   E. **show ip dhcp address**

5. To enable the DHCP Relay feature, which command should you enable?

   A. **ip dhcp relay**
   B. **ip dhcp transmit**
   C. **ip dhcp helper**
   D. **ip helper-address**
   E. There is no command required. Cisco Routers perform DHCP Relay by default.

6. What is the purpose of DNS?

   A. To map Assigned Addresses to MAC Addresses
   B. To map Assigned Addresses to Hostnames
   C. To map private IP Addresses to public IP Addresses
   D. To map public and private IP Addresses to domain names

7. You have been asked to plan for the implementation of DHCP Server on a Cisco Router.
   The IP Addresses have to be between 10.1.50.0 to 10.1.50.254.
   The Default Gateway is 10.1.50.1
   The DNS Server is 10.1.50.2
   The IP Addresses should be assigned for 6 hours at a time.
   The domain name is cisco.com.

   Which of the following commands are correct for this implementation? (Choose four.)

   A. **ip dhcp pool Users**
   B. **network 10.1.50.0/24**
   C. **network 10.1.50.1 255.255.255.0**
   D. **ip dhcp excluded-address 10.1.50.1 10.1.50.2**
   E. **ip dhcp excluded-address 10.1.50.1 10.1.50.10**
   F. **lease 6**
   G. **lease 0 6**
   H. **dns-server 10.1.50.1**
   I. **default-router 10.1.50.2**

Answer Key

Challenge

1. B
2. C
3. C
4. B
5. D
6. B
7. A, B, D, G

Lesson 6: Understanding Layer 3 Redundancy

Introduction

The law firm calls CCS to ask about default gateway redundancy. Although the law firm has dual redundant routers that connect to the Internet, none of the PCs at the firm can access the Internet because the primary router went down due to a UPS failure. Bob asks if you are ready to go onsite with him to explain the various FHRP options and then implement one.

Need for Default Gateway Redundancy

When the host determines that a destination IP network is not on its local subnet, it forwards the packet to the default gateway. Although an IP host can run a dynamic routing protocol to build a list of reachable networks, most IP hosts rely on a locally configured or dynamically learned through DHCP default gateway.

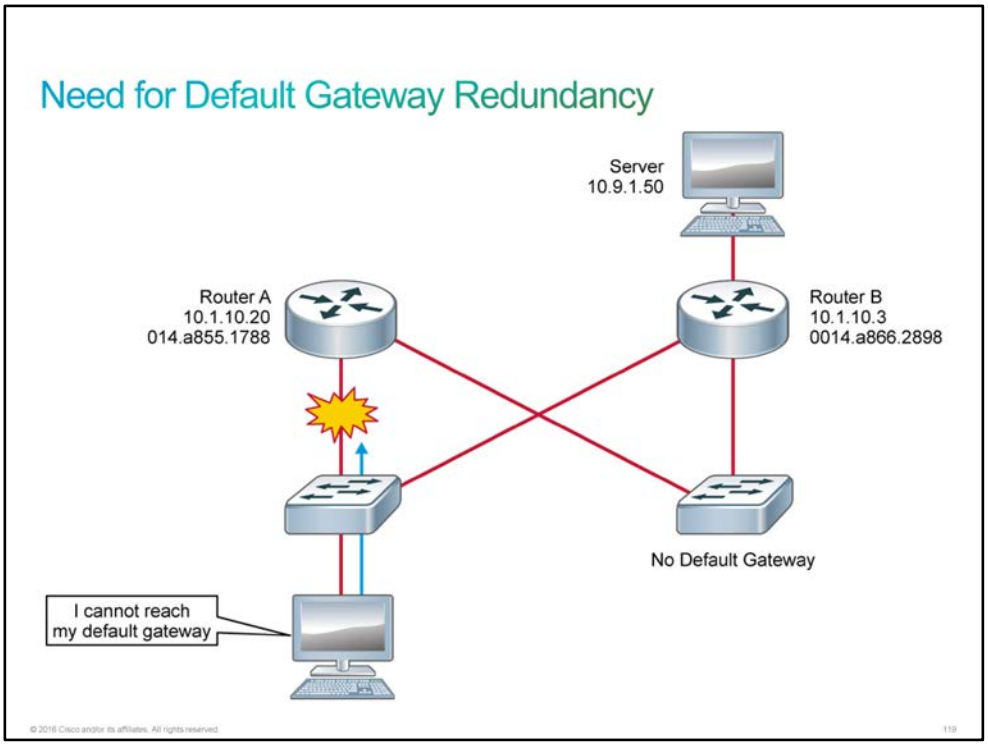

Having redundant equipment alone does not guarantee uptime. In this example, both router A and router B are responsible for routing packets for the 10.1.10.0/24 subnet. Because the routers are deployed as a redundant pair, if router A becomes unavailable, the IGP can quickly and dynamically converge and determine that router B will now transfer packets that would otherwise have gone through router A. Most workstations, servers, and printers, however, do not receive this dynamic routing information.

Each end device is configured with a single default gateway IP address that does not dynamically update when the network topology changes. If the default gateway fails, the local device is unable to send packets off the local network segment. As a result, the host is isolated from the rest of the network. Even if a redundant router exists 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	Even though the example is illustrated on routers, it is equally valid on Layer 3 switches.

Understanding FHRP

The figure represents a generic router FHRP with a set of routers working together to present the illusion of a single router to the hosts on the LAN. By sharing an IP (Layer 3) address and a MAC (Layer 2) address, two or more routers can act as a single "virtual" router.

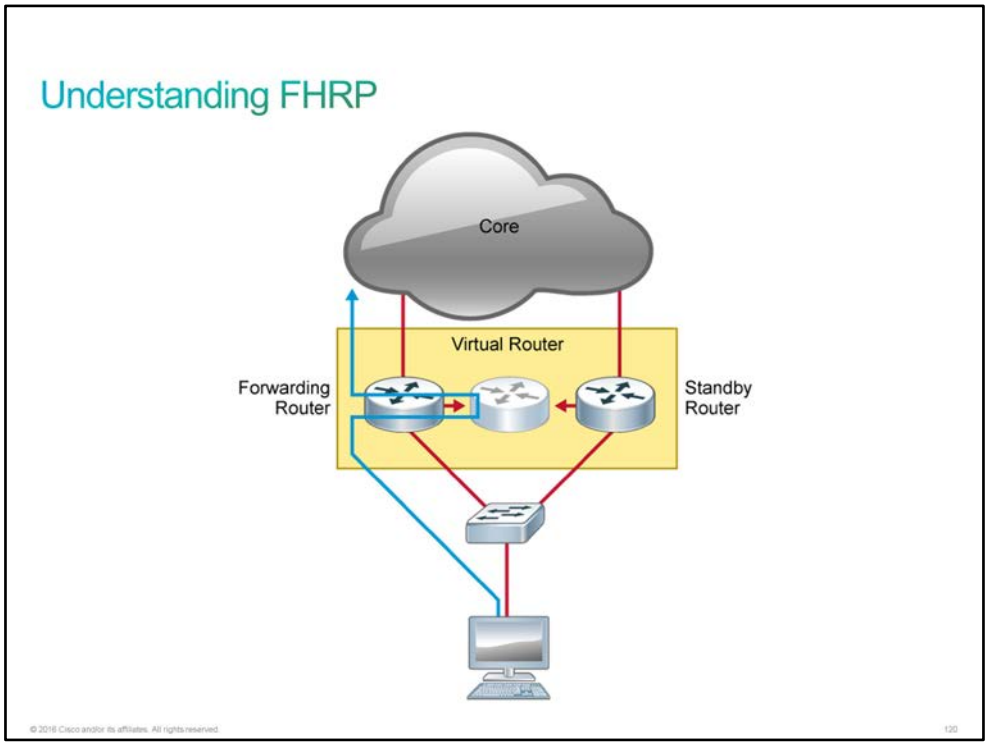

Hosts that are on the local subnet configure the IP address of the virtual router as their default gateway. When a host needs to communicate to another IP host on a different subnet, it will use ARP to resolve the MAC address of the default gateway. The ARP resolution returns the MAC address of the virtual router. The packets that devices send to the MAC address of the virtual router can then be routed to their destination by any active or standby router that is part of that virtual router group.

You use an FHRP to coordinate two or more routers as the devices that are responsible for processing the packets that are sent to the virtual router. The host devices send traffic to the address of the virtual router. The actual (physical) router that forwards this traffic is transparent to the end stations.

The redundancy protocol provides the mechanism for determining which router should take the active role in forwarding traffic and determining when a standby router should take over that role. The transition from one forwarding router to another is also transparent to the end devices.

Cisco routers and switches commonly use three FHRPs. A common feature of FHRPs is to provide a default gateway failover that is transparent to hosts.

1. **HSRP:** HSRP is an FHRP that Cisco designed to create a redundancy framework between network routers or switches in order to achieve default gateway failover capabilities. Only one router forwards traffic. HSRP is defined in RFC 2281.
2. **VRRP:** VRRP is an open FHRP standard that offers the ability to add more than two routers for additional redundancy. Only one router forwards traffic. VRRP is defined in RFC 5798.
3. **GLBP:** GLBP is an FHRP that Cisco designed to allow multiple active forwarders to load-balance outgoing traffic.

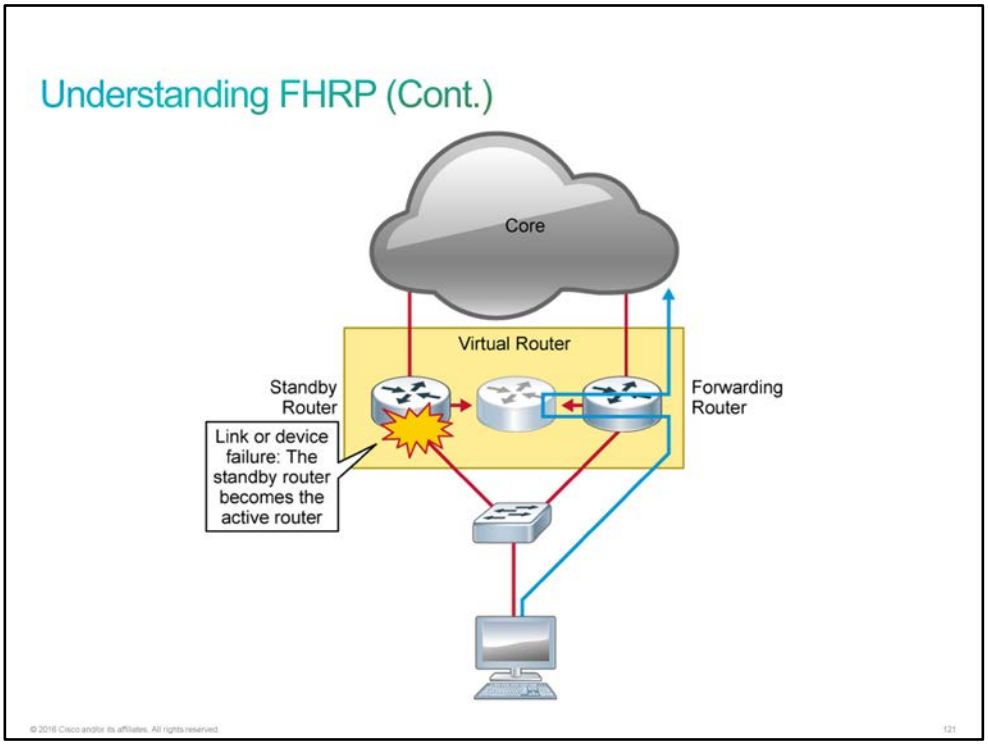

When a route fails, the following steps take place:

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 IP and MAC addresses of the virtual router, the end stations see no disruption in service.

Understanding HSRP

HSRP is an FHRP that allows for transparent failover of the first-hop IP device (default gateway). Most IP hosts have an IP address of a single router configured as the default gateway. When you use HSRP, the HSRP virtual IP address is configured as the default gateway for the host instead of the IP address of the router.

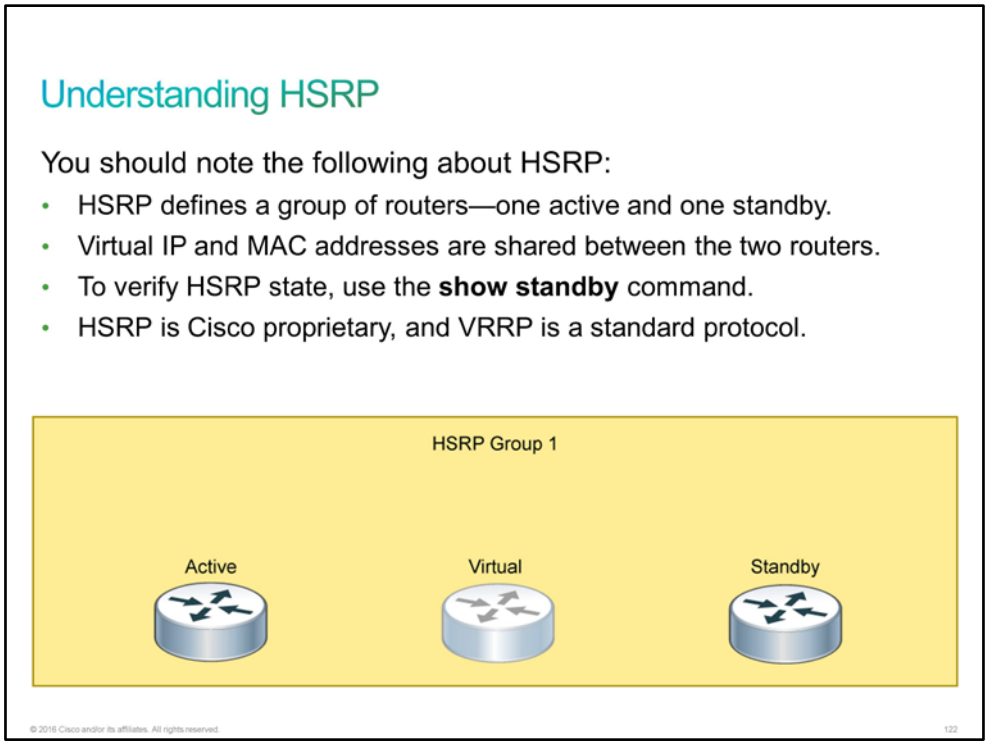

HSRP defines a standby group of routers, with one router that is designated as the active router. HSRP provides gateway redundancy by sharing IP and MAC addresses between redundant gateways. The protocol consists of virtual MAC and IP addresses that two routers that belong to the same HSRP group share between each other.

Hosts on the IP subnet that are protected by HSRP configure their default gateway with the HSRP group virtual IP address. The packets that are received on the virtual IP address are forwarded to the active router.

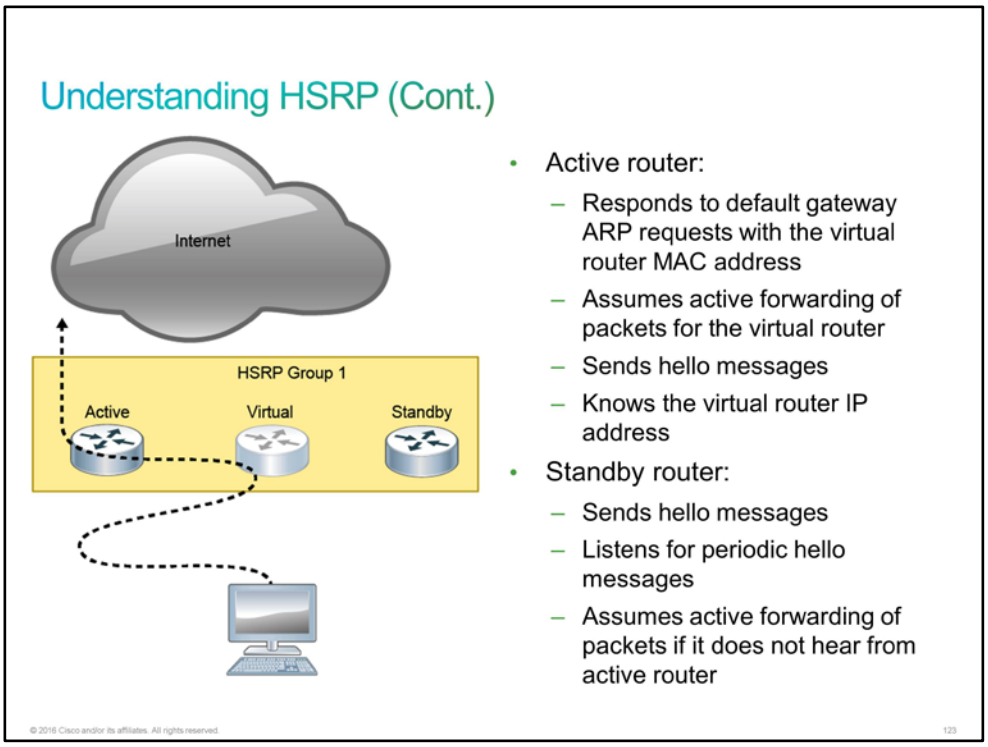

HSRP Terminology

Term	Definition
Active router	The router that is currently forwarding packets for the virtual router
Standby router	The primary backup router
Standby group	The set of routers participating in HSRP that jointly emulate a virtual router

The function of the HSRP standby router is to monitor the operational status of the HSRP group and to quickly assume the packet-forwarding responsibility if the active router becomes inoperable.

HSRP is a Cisco proprietary protocol, and VRRP is a standard protocol. Beyond this fact, the differences between HSRP and VRRP are very slight.

Besides the default behavior, you can configure some other HSRP features to increase you network availability and performance:

- Interface tracking: When a tracked interface becomes unavailable, the HSRP tracking feature ensures that a router with an unavailable key interface will relinquish the active router role.
- Load Balancing: Routers can simultaneously provide redundant backup and perform load sharing across various subnets and VLANs.

Discovery 24: Configure and Verify HSRP

Introduction

In this guided discovery, you will work with HSRP. Hosts on IP networks usually only have a single IP address that is configured as their default gateway. HSRP allows two physical routers to work together in an HSRP group to provide a virtual IP address and an associated virtual MAC address.

The end hosts use the virtual IP address as their default gateway and learn the virtual MAC address via ARP. One of the routers in the group is active and responsible for the virtual addresses. The other router is in a standby state and monitors the active router.

If there is a failure on the active router, the standby router assumes the active state. The virtual addresses are always functional, regardless of which physical router is responsible for them. The end hosts are not aware of any changes in the physical routers.

Consult the topology diagram. The live virtual lab is prepared with the devices that are represented in the topology diagram and the connectivity table. All devices have their basic configurations in place, including hostnames and IP addresses. RIP is configured on the three routers, making both R1 and R2 aware of the 10.10.99.0 subnet that is connected to R3.

The two PCs are configured with 10.10.1.1 as their default gateway. Note that this address does not yet exist in the topology. R1 uses 10.10.1.2 and R2 uses 10.10.1.3. In this discovery, you will configure and verify HSRP on R1 and R2, using 10.10.1.1 as the virtual IP address.

You will start by verifying the initial state on PC1 and R1. You will then configure and verify HSRP on R1. It only takes one functional router in an HSRP group to provide forwarding services for the end hosts on the network. You will then configure and verify HSRP on R2. Finally, you will cause a fault in R1 and then verify that R2 takes over the HSRP active role.

Topology

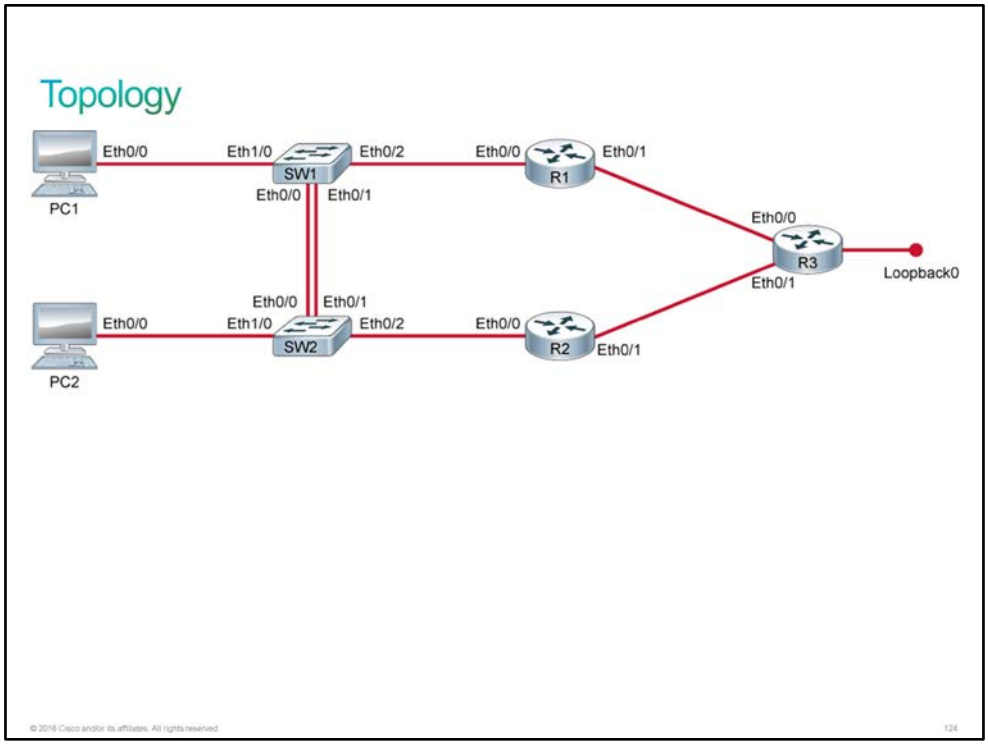

Job Aid

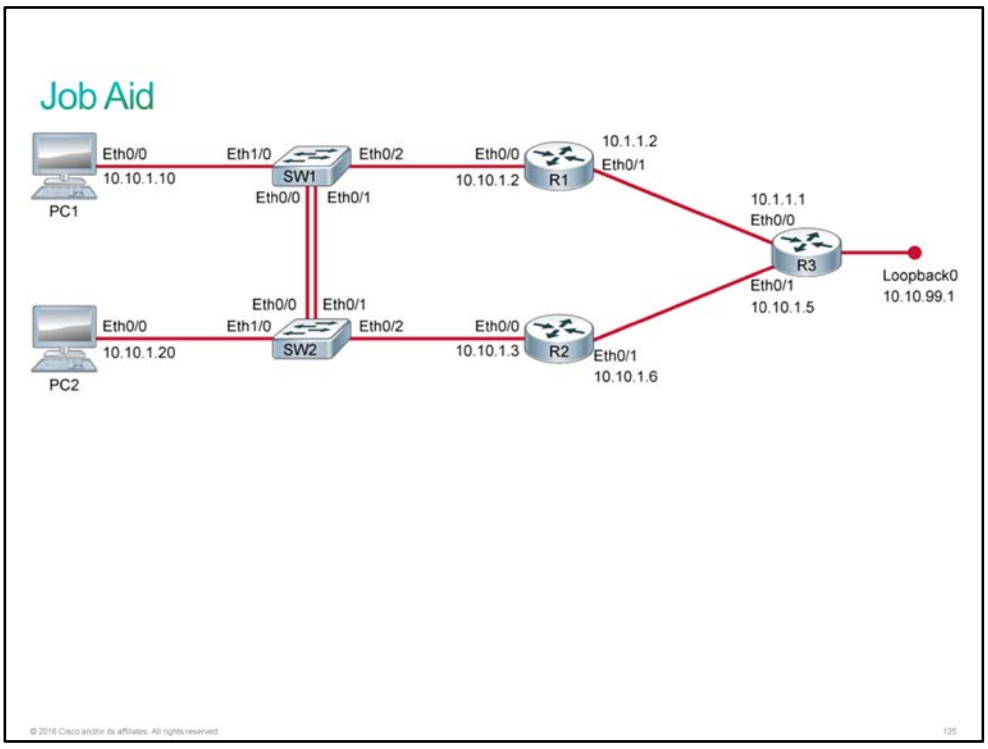

The configuration is as follows:

- All devices have their basic configurations in place, including hostnames and IP addresses.
- RIP is configured on the three routers, making both R1 and R2 aware of the 10.10.99.0 subnet that is connected to R3.
- The two PCs are configured with 10.10.1.1 as their default gateway. Note that this address does not yet exist in the topology

Device Details

Device	Interface	Neighbor	IP Address
PC1	Ethernet0/0	SW1	10.10.1.10/24
PC2	Ethernet0/0	SW2	10.10.1.20/24
R1	Ethernet0/0	SW1	10.10.1.2/24
R1	Ethernet0/1	R3	10.1.1.2/30
R2	Ethernet0/0	SW2	10.10.1.3/24
R2	Ethernet0/1	R3	10.1.1.6/30
R3	Ethernet0/0	R1	10.1.1.1/30

Device	Interface	Neighbor	IP Address
R3	Ethernet0/1	R2	10.1.1.5/30
R3	Loopback0	—	10.10.99.1/24

Note	PCs in the virtual lab environment are simulated as routers, so you should use Cisco IOS commands to configure them or make verifications.

Task 1: Configure and Verify HSRP

Activity

Complete the following steps:

Step 1 Access the console of PC1. View its routing table to verify that 10.10.1.1 is indeed its default gateway. Also, attempt to ping 10.10.1.1 to verify that it does not yet exist on the network.

Enter these commands on PC1:

```
PC1# show ip route
Default gateway is 10.10.1.1

Host               Gateway           Last Use    Total Uses  Interface
ICMP redirect cache is empty

PC1# ping 10.10.1.1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.10.1.1, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
```

PC1 is not running any dynamic routing protocols. Default gateway is configured with IP address 10.10.1.1, which does not yet exist. This is why ping operation fails.

Step 2 Access the console of R1 and verify that the IP address of its Ethernet0/0 interface is 10.10.1.2.

Enter this command on the R1 router:

```
R1# show ip interface brief
Interface                  IP-Address      OK? Method Status
Prot                ocol
Ethernet0/0                10.10.1.2       YES NVRAM  up                    up
Ethernet0/1                10.1.1.2        YES NVRAM  up                    up
Ethernet0/2                unassigned      YES NVRAM  administratively down
down
Ethernet0/3                unassigned      YES NVRAM  administratively down
down
```

Step 3 Access the console of R2 and verify that the IP address of its Ethernet0/0 interface is 10.10.1.3.

Enter this command on the R1 router:

```
R2# show ip interface brief
Interface                  IP-Address      OK? Method Status
Protocol
Ethernet0/0                10.10.1.3       YES NVRAM  up                    up
Ethernet0/1                10.1.1.6        YES NVRAM  up                    up
Ethernet0/2                unassigned      YES NVRAM  administratively down
down
Ethernet0/3                unassigned      YES NVRAM  administratively down
down
```

HSRP Configuration

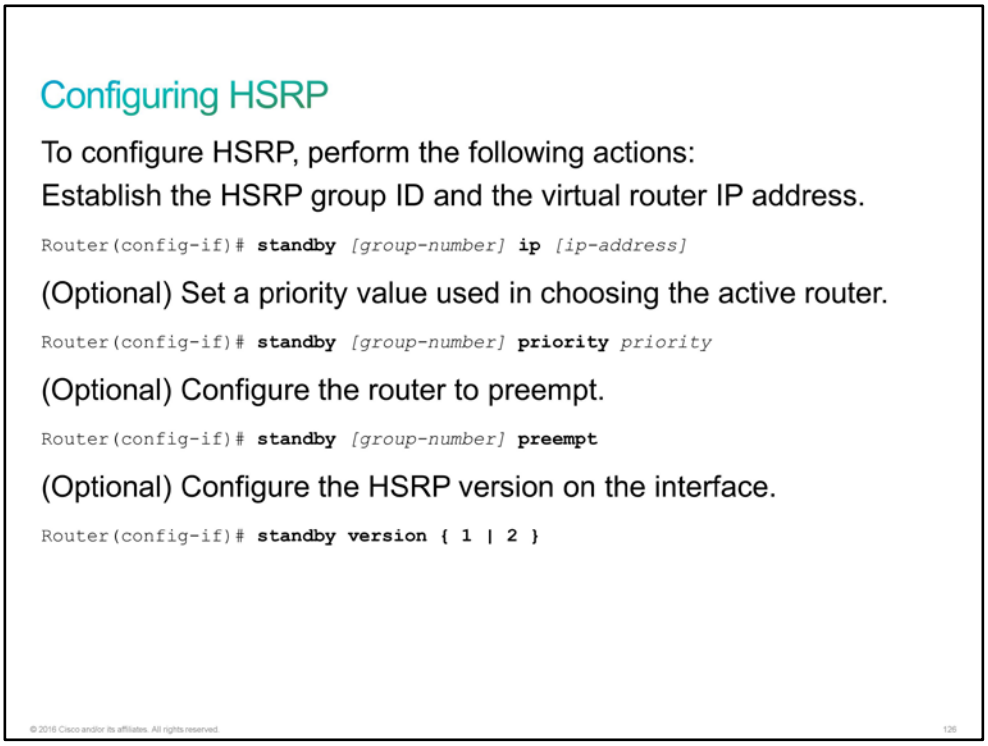

All basic HSRP configuration is performed in the interface configuration mode by using the **standby** command. The **standby ip** interface configuration command activates HSRP on the configured interface. If an IP address is specified, that address is used as the designated address for the Hot Standby group.

Note	A router in an HSRP group can be any routed interface that supports HSRP, including routed ports on Layer 3 switches and SVIs.

Command	Description
standby *[group-number]* **ip** *[ip-address]*	Establishes the HSRP group ID and the virtual router IP address. The group number on the interface specifies the group for which HSRP is being enabled. The range is 0 to 255; the default is 0. If there is only one HSRP group, you do not need to enter a group number. The IP address specifies the virtual IP address of the hot standby router interface.

Command	Description
standby *[group-number]* **priority** *priority*	Set a priority value used in choosing the active router. The range is 1 to 255; the default priority is 100. The highest number represents the highest priority.
standby *[group-number]* **preempt**	Configure the router to preempt , which means that when the local router has a higher priority than the active router, the local router becomes the active router.
standby version { *1* | *2* }	Configure the HSRP version on the interface. If you do not enter this command or do not specify a keyword, the interface runs the default HSRP version, HSRPv1.

Assigning a priority allows you to select the active and standby routers. If preemption is enabled, the router with the highest priority becomes the active router. If priorities are equal, the current active router does not change. The highest number (1 to 255) represents the highest priority (most likely to become the active router).

Step 4 Configure HSRP for R1 on Ethernet0/0. Assign the virtual IP address 10.10.1.1 and the HSRP group number 1.

Enter these commands on the R1 router:

```
R1# conf t
R1(config)# interface Ethernet0/0
R1(config-if)# standby 1 ip 10.10.1.1
R1(config-if)# end
R1#
*Nov 20 07:57:10.405: %HSRP-5-STATECHANGE: Ethernet0/0 Grp 1 state Standby ->
Active
```

Wait for the syslog message that indicates that R1 has transitioned to the active HSRP state before continuing. It will take 20 seconds from the time that you enabled HSRP.

Once HSRP is enabled on R1, it listens for HSRP Hello messages to determine if there is another HSRP device on the subnet. Since R1 is currently the only HSRP device, it becomes the active HSRP node and then begins sending Hello messages.

HSRP Verification

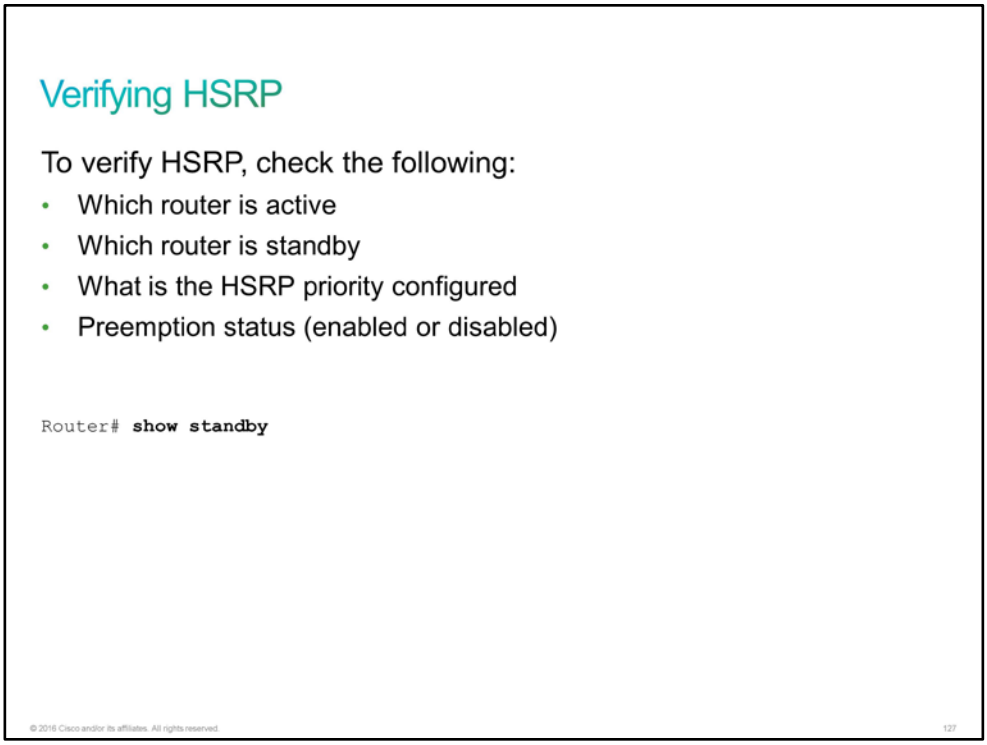

The **show standby** command is used to monitor the HSRP state on each router in the standby group.

Step 5 Display the HSRP status on R1.

Enter this command on the R1 router:

```
R1# show standby
Ethernet0/0 - Group 1
  State is Active
    2 state changes, last state change 00:17:49
  Virtual IP address is 10.10.1.1
  Active virtual MAC address is 0000.0c07.ac01
    Local virtual MAC address is 0000.0c07.ac01 (v1 default)
  Hello time 3 sec, hold time 10 sec
    Next hello sent in 0.240 secs
  Preemption disabled
  Active router is local
  Standby router is unknown
  Priority 100 (default 100)
  Group name is "hsrp-Et0/0-1" (default)
```

R1 is the active HSRP router and the standby router is unknown. R2 has not yet been configured, so there is no standby router.

A virtual MAC address 0000.0c07.ac01 is associated with the virtual IP address 10.10.1.1. Both the MAC address and the IP address will be shared by the HSRP routers.

The Active virtual MAC address is formed from the well-known 5-byte prefix 0000.0c07.ac and the HSRP group number is encoded into the 1-byte suffix.

Step 6 Verify the real MAC address that is assigned to Ethernet0/0 on R1.

Enter this command on the R1 router:

```
R1# show interfaces Ethernet0/0
Ethernet0/0 is up, line protocol is up
  Hardware is AmdP2, address is aabb.cc00.0300 (bia aabb.cc00.0300)
  Description: Link to SW1
  Internet address is 10.10.1.2/24
<... output omitted ...>
```

This MAC address is associated with the real IP address 10.10.1.2 on R1. The MAC address might differ in your output.

Step 7 Access the console of PC1. Verify that you can now ping 10.10.1.1, which is the default gateway configured on PC1. Also ping 10.10.99.1, which should only be reachable from PC1 if its default gateway is available.

Enter these commands on the PC1 router:

```
PC1# ping 10.10.1.1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.10.1.1, timeout is 2 seconds:
.!!!!
Success rate is 80 percent (4/5), round-trip min/avg/max = 1/1/1 ms

PC1# ping 10.10.99.1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.10.99.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
```

Step 8 View the ARP cache on PC1.

Enter this command on the PC1 router:

```
PC1# show arp
Protocol  Address          Age (min)  Hardware Addr   Type   Interface
Internet  10.10.1.1               0   0000.0c07.ac01  ARPA   Ethernet0/0
<...output omitted...>
```

Note that the MAC address that is associated with 10.10.1.1 is the HSRP virtual MAC address and not the Ethernet0/0 physical MAC address for R1.

Step 9 Access the console of R2. Configure HSRP on Ethernet0/0. Assign the virtual IP address 10.10.1.1 and the HSRP group number to 1.

Enter this command on the R2 router:

```
R2# conf t
R2(config)# interface Ethernet0/0
R2(config-if)# standby 1 ip 10.10.1.1
R2(config-if)# end
R2#
*Nov 20 09:05:13.312: %HSRP-5-STATECHANGE: Ethernet0/0 Grp 1 state Speak ->
Standby
```

Because R1 is already in the active state, R2 will transition to the standby state. Wait for the syslog message that indicates this transition.

R2 receives a Hello message from R1 and assumes the standby state. If R2 detects three missed Hello messages from R1, R2 will promote itself to the active state.

Step 10 Display the HSRP status on R2.

Enter this command on the R2 router:

```
R2# show standby
Ethernet0/0 - Group 1
  State is Standby
    1 state change, last state change 00:06:34
  Virtual IP address is 10.10.1.1
  Active virtual MAC address is 0000.0c07.ac01
    Local virtual MAC address is 0000.0c07.ac01 (v1 default)
  Hello time 3 sec, hold time 10 sec
    Next hello sent in 0.928 secs
  Preemption disabled
  Active router is 10.10.1.2, priority 100 (expires in 11.152 sec)
  Standby router is local
  Priority 100 (default 100)
  Group name is "hsrp-Et0/0-1" (default)
```

R2 is in the standby state and uses the same virtual IP address and virtual MAC address that R1 uses. R2 is aware that R1 (10.10.1.2) is the active router.

The virtual MAC address is shared by active and standby HSRP devices. When an IP host sends an ARP request for its default gateway, the active HSRP router responds with the virtual MAC address. If the active HSRP router fails, the MAC address of the default IP address does not change.

Step 11 Access the console of R1 and display the status of HSRP again.

Enter this command on the R1 router:

```
R1# show standby
Ethernet0/0 - Group 1
  State is Active
    2 state changes, last state change 01:27:26
  Virtual IP address is 10.10.1.1
  Active virtual MAC address is 0000.0c07.ac01
    Local virtual MAC address is 0000.0c07.ac01 (v1 default)
  Hello time 3 sec, hold time 10 sec
    Next hello sent in 0.288 secs
  Preemption disabled
  Active router is local
  Standby router is 10.10.1.3, priority 100 (expires in 9.456 sec)
  Priority 100 (default 100)
  Group name is "hsrp-Et0/0-1" (default)
```

As it was before, R1 is in the active state and is using the virtual IP address and virtual MAC address. R1 is now aware of R2 (10.10.1.3) as the standby router in the HSRP group.

Step 12 Cause a fault on R1 by disabling its Ethernet0/0 interface. Observe the syslog messages indicating the state changes.

Enter these commands on the R1 router:


```
R1# conf t
R1(config)# interface Ethernet0/0
R1(config-if)# shutdown
*Nov 20 10:06:52.369: %HSRP-5-STATECHANGE: Ethernet0/0 Grp 1 state Active ->
Init
R1(config-if)#
*Nov 20 10:06:54.375: %LINK-5-CHANGED: Interface Ethernet0/0, changed state to
administratively down
*Nov 20 10:06:55.379: %LINEPROTO-5-UPDOWN: Line protocol on Interface
Ethernet0/0, changed state to down
```

HSRP reacts even before the interface changes state.

Step 13 View the status of HSRP on R1.

Enter this command on the R1 router:

```
R1(config-if)# end
R1# show standby
Ethernet0/0 - Group 1
  State is Init (interface down)
    3 state changes, last state change 00:03:19
  Virtual IP address is 10.10.1.1
  Active virtual MAC address is unknown
    Local virtual MAC address is 0000.0c07.ac01 (v1 default)
  Hello time 3 sec, hold time 10 sec
  Preemption disabled
  Active router is unknown
  Standby router is unknown
  Priority 100 (default 100)
  Group name is "hsrp-Et0/0-1" (default)
```

R1 goes to HSRP state "Init."

Step 14 Observe the syslog messages on R2 that were produced in association with the HSRP activity. Also display the status of HSRP, verifying that R2 is now in the active state.

Enter these commands on the R2 router:

```
R2#
*Nov 20 10:06:52.365: %HSRP-5-STATECHANGE: Ethernet0/0 Grp 1 state Standby ->
Active

R2# show standby
Ethernet0/0 - Group 1
  State is Active
    2 state changes, last state change 00:14:00
  Virtual IP address is 10.10.1.1
  Active virtual MAC address is 0000.0c07.ac01
    Local virtual MAC address is 0000.0c07.ac01 (v1 default)
  Hello time 3 sec, hold time 10 sec
    Next hello sent in 2.496 secs
  Preemption disabled
  Active router is local
  Standby router is unknown
  Priority 100 (default 100)
  Group name is "hsrp-Et0/0-1" (default)
```

R2 is now the active HSRP router. The standby router is unknown because R1 is offline. The virtual IP address and the virtual MAC address remain unchanged.

Step 15 Access the console of PC1. Verify that, even though there has been a change in the physical routers, PC1 still has access to the 10.10.99.1 IP address.

Enter this command on the PC1 router:

```
PC1# ping 10.10.99.1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.10.99.1, timeout is 2 seconds:
!!!!!
```

The network behind the default gateway is still accessible.

Step 16 View the ARP cache on PC1.

Enter this command on the PC1 router:

```
PC1# show arp
Protocol  Address          Age (min)  Hardware Addr   Type   Interface
Internet  10.10.1.1               18   0000.0c07.ac01  ARPA   Ethernet0/0
<...output omitted...>
```

Note that the MAC address that is associated with 10.10.1.1 is still the HSRP virtual MAC address.

Step 17 Return to R1. Enable the Ethernet0/0 interface. Wait for the link and HSRP status messages.

Enter this command on the R1 router:

```
R1# conf t
R1(config)# interface Ethernet0/0
R1(config-if)#  no shutdown
*Nov 20 10:41:54.804: %LINK-3-UPDOWN: Interface Ethernet0/0, changed state to
up
*Nov 20 10:41:55.804: %LINEPROTO-5-UPDOWN: Line protocol on Interface
Ethernet0/0, changed state to up
*Nov 20 10:42:17.140: %HSRP-5-STATECHANGE: Ethernet0/0 Grp 1 state Speak ->
Standby
```

R1 goes to HSRP standby state.

Step 18 View the status of HSRP on R1.

Enter this command on the R1 router:


```
R1(config-if)# end
R1# show standby
Ethernet0/0 - Group 1
  State is Standby
    4 state changes, last state change 00:02:43
  Virtual IP address is 10.10.1.1
  Active virtual MAC address is 0000.0c07.ac01
    Local virtual MAC address is 0000.0c07.ac01 (v1 default)
  Hello time 3 sec, hold time 10 sec
    Next hello sent in 1.568 secs
  Preemption disabled
  Active router is 10.10.1.3, priority 100 (expires in 9.456 sec)
  Standby router is local
  Priority 100 (default 100)
  Group name is "hsrp-Et0/0-1" (default)
```

You may have expected R1 to return to the active state. This change does not happen by default. HSRP does support the concept of priority-based preemption but it is disabled by default. Even if preemption was enabled, both R1 and R2 have the same (default) priority of 100.

This is the end of the discovery lab.

Discovery 25: Troubleshoot HSRP

Introduction

In this guided discovery, you will work with typical HSRP configuration issues. You will see a duplicated IP address issue on both R1 and R2 routers. The reason for this issue is HSRP misconfiguration.

The desired HSRP configuration uses 10.10.1.1 as the virtual IP address in the HSRP group 1, and R1 is the active HSRP router. This is not the case, so you will perform troubleshooting steps to isolate the configuration issues.

Topology

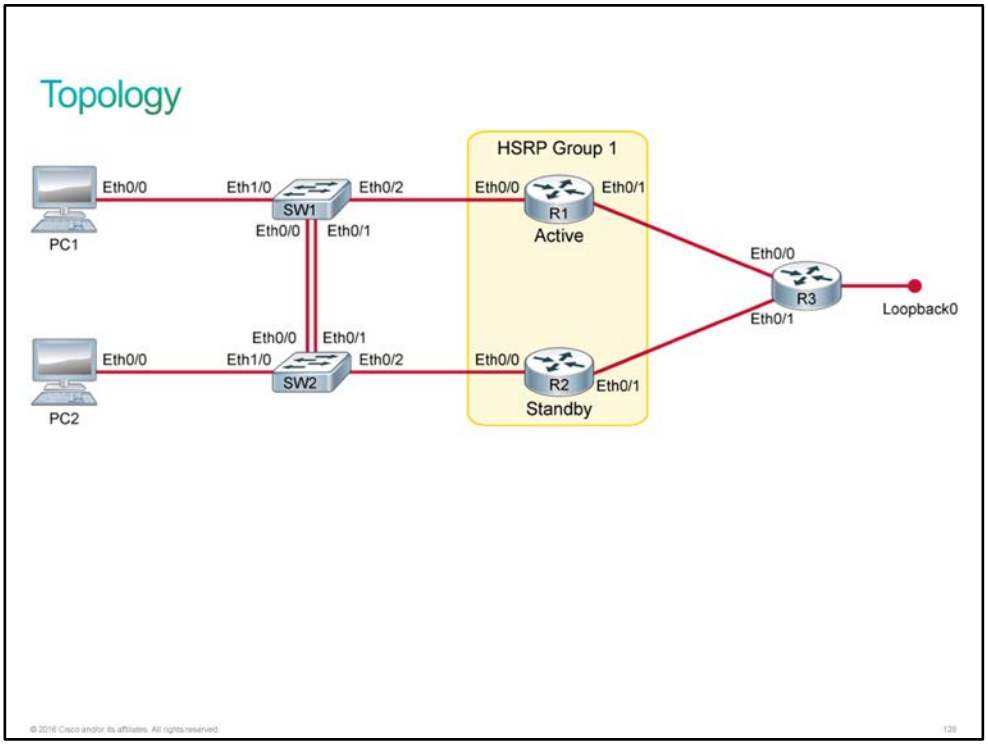

Job Aid

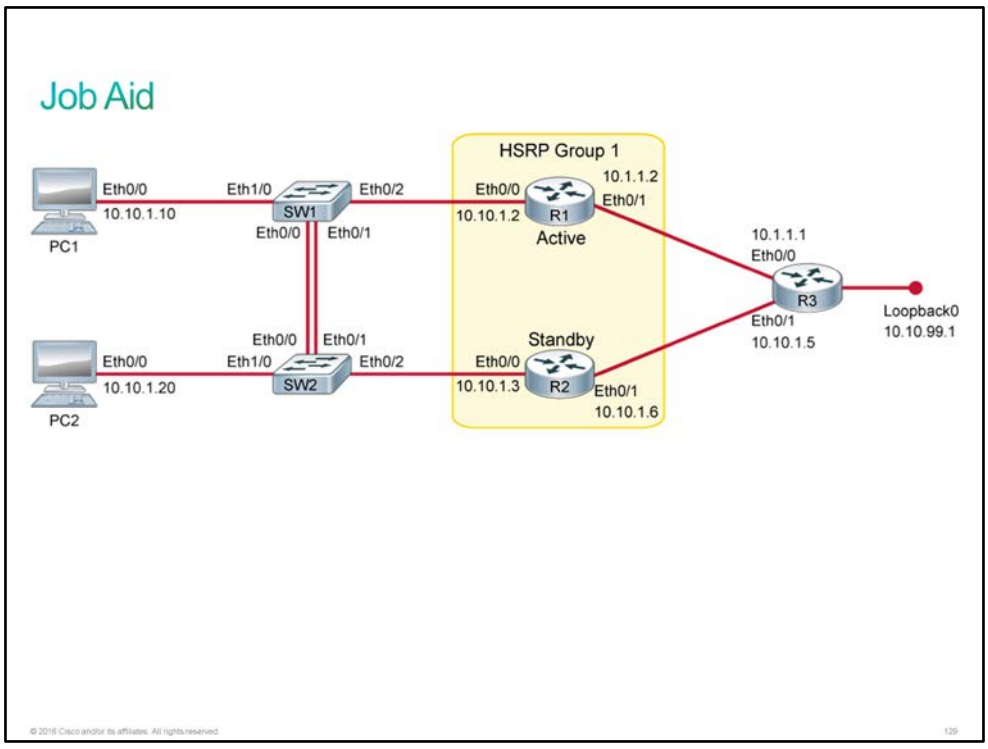

The configuration is as follows:

- All devices have their basic configurations in place, including hostnames and IP addresses.
- RIP is configured on the three routers, making both R1 and R2 aware of the 10.10.99.0 subnet that is connected to R3.
- The two PCs are configured with 10.10.1.1 as their default gateway.
- HSRP is preconfigured but it does not behave as expected. The desired behavior is as follows:
  - Virtual IP address should be 10.10.1.1.
  - The used HSRP group should be 1.
  - The used HSRP version should be 1.
  - R1 should be the active HSRP router, while the R2 should be in the standby state.

Device Details

Device	Interface	Neighbor	IP Address
PC1	Ethernet0/0	SW1	10.10.1.10/24
PC2	Ethernet0/0	SW2	10.10.1.20/24
R1	Ethernet0/0	SW1	10.10.1.2/24
R1	Ethernet0/1	R3	10.1.1.2/30

Device	Interface	Neighbor	IP Address
R2	Ethernet0/0	SW2	10.10.1.3/24
R2	Ethernet0/1	R3	10.1.1.6/30
R3	Ethernet0/0	R1	10.1.1.1/30
R3	Ethernet0/1	R2	10.1.1.5/30
R3	Loopback0		10.10.99.1/24

Note	PCs in the virtual lab environment are simulated as routers, so you should use Cisco IOS commands to configure them or make verifications.

Task 1: Troubleshoot HSRP

Activity

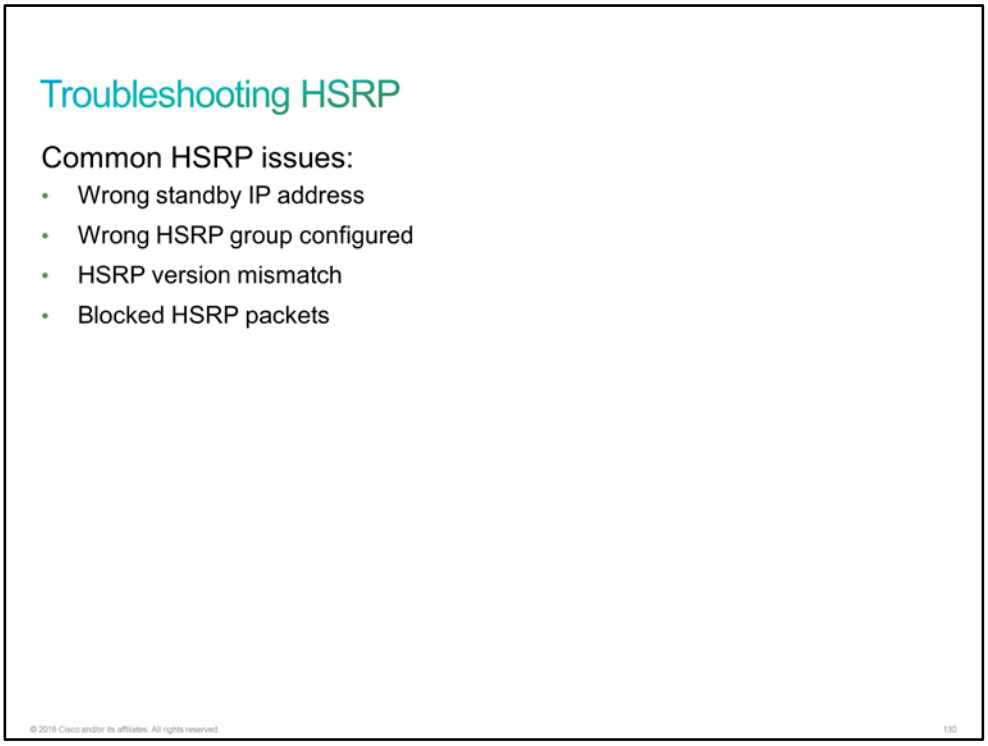

There are several possible misconfigurations of HSRP:

- Different HSRP virtual IP addresses could be configured. Console messages will notify you about this situation. With such a configuration, when the active router fails, the standby router takes over with a virtual IP address, which is different to the one used previously, and different to the one configured as the default-gateway address for end devices.
- If a wrong HSRP group is configured on the peers, this misconfiguration leads to both peers becoming active. This issue will manifest as a duplicate IP address problem.
- HSRP comes in 2 versions, 1 and 2. If there is a version mismatch, both routers will become active. This mismatch results in duplicate IP addresses.

Most of the HSRP misconfiguration problems can be solved by checking the output of the **show standby** command. In the output, you can notice the active IP and the MAC address, the timers, the active router, and several others parameters.

HSRP messages are sent to the multicast IP address 224.0.0.2 and UDP port 1985 in version 1 and the multicast IP address 224.0.0.102 and UDP port 1985 in version 2. These IP addresses and ports need to be permitted in the inbound access lists. If the packets are blocked, the peers will not see each other and there will be no HSRP redundancy. To check the interface access list, use the **show ip interface** command.

Assigning a priority allows you to select the active and standby routers. If preemption is enabled, the router with the highest priority becomes the active router. If priorities are equal, the current active router does not change. The highest number (1 to 255) represents the highest priority (most likely to become the active router).

Complete the following steps:

Step 1 Access the console of R1 and observe the output.

Observe the output on R1 router:

```
R1#
*Nov 30 17:04:40.901: %IP-4-DUPADDR: Duplicate address 10.10.1.1 on
Ethernet0/0, sourced by 0000.0c9f.f002
```

The console messages indicate that there is a duplicate IP address problem in the network. Source of the duplicate address is device R2 with the MAC address 0000.0C9F.F002, which is an HSRP version 2 format MAC address. HSRP version 2 uses a new MAC address range 0000.0C9F.F0XX, where XX is the group number.

Step 2 Access the console of R2 and observe the output.

Observe the output on R2 router:

```
R2#
*Nov 30 17:05:06.904: %IP-4-DUPADDR: Duplicate address 10.10.1.1 on
Ethernet0/0, sourced by 0000.0c07.ac01
```

The console messages indicate that there is a duplicate IP address problem in the network. Source of the duplicate address is device R1 with MAC address 0000.0C07.AC01, which is an HSRP version 1 format MAC address. HSRP version 1 uses MAC addresses 0000.0c07.acXX, where XX is the HSRP group number.

Step 3 Display the HSRP status on R1.

Enter this command on the R1 router:

```
R1# show standby
Ethernet0/0 - Group 1
  State is Active
    2 state changes, last state change 04:06:41
  Virtual IP address is 10.10.1.1
  Active virtual MAC address is 0000.0c07.ac01
    Local virtual MAC address is 0000.0c07.ac01 (v1 default)
  Hello time 3 sec, hold time 10 sec
    Next hello sent in 1.696 secs
  Preemption enabled
  Active router is local
  Standby router is unknown
  Priority 99 (configured 99)
  Group name is "hsrp-Et0/0-1" (default)
```

HSRP was configured on the Ethernet0/0 interface for HSRP group 1. This router is the active router for this group, but the standby router is unknown. The HSRP version in use is 1.

Also note that the virtual MAC address equals the source MAC address of the duplicated IP address that is received on the R2 router.

Step 4 Display the HSRP status on R2.

Enter this command on the R2 router:

```
R2# show standby
Ethernet0/0 - Group 2 (version 2)
  State is Active
    2 state changes, last state change 04:28:34
  Virtual IP address is 10.10.1.1
  Active virtual MAC address is 0000.0c9f.f002
    Local virtual MAC address is 0000.0c9f.f002 (v2 default)
  Hello time 3 sec, hold time 10 sec
    Next hello sent in 0.544 secs
  Preemption disabled
  Active router is local
  Standby router is unknown
  Priority 100 (default 100)
  Group name is "hsrp-Et0/0-2" (default)
```

HSRP was configured on the Ethernet0/0 interface for HSRP group 2. This router is the active router for this group, but the standby router is unknown. R1 and R2 are not HSRP aware of each other, as they are configured for different HSRP groups.

The HSRP version that R2 uses is 2. The R1 router uses version 1.

Also note that the virtual MAC address equals the source MAC address of the duplicated IP address received on the R2 router.

Step 5 Verify the configuration of the Ethernet0/0 interface on R1.

Enter this command on the R1 router:


```
R1# show running-config interface Ethernet0/0
Building configuration...

Current configuration : 156 bytes
!
interface Ethernet0/0
 description Link to SW1
 ip address 10.10.1.2 255.255.255.0
 standby 1 ip 10.10.1.1
 standby 1 priority 99
 standby 1 preempt
end
```

The virtual IP address 10.10.1.1 is configured for HSRP group 1. The default HSRP version is 1.

Step 6 Verify the configuration of the Ethernet0/0 interface on R2.

Enter this command on the R2 router:

```
R2# show running-config interface Ethernet0/0
Building configuration...

Current configuration : 152 bytes
!
interface Ethernet0/0
 description Link to SW2
 ip address 10.10.1.3 255.255.255.0
 standby version 2
 standby 1 preempt
 standby 2 ip 10.10.1.1
```

The virtual IP address 10.10.1.1 is configured for HSRP group 2. The HSRP version in use is 2. There is clearly a mismatch in HSRP group and version configuration between routers R1 and R2.

Step 7 Access the console of the R2 router and fix the configuration.

Enter these commands on the R2 router:

```
R2# conf t
R2(config)# interface Ethernet0/0
R2(config-if)# no standby version 2
R2(config-if)#
*Nov 30 18:40:41.389: %HSRP-5-STATECHANGE: Ethernet0/0 Grp 2 state Active ->
Init
R2(config-if)# no standby 2 ip 10.10.1.1
R2(config-if)# standby 1 ip 10.10.1.1
R2(config-if)#
*Nov 30 18:41:11.629: %HSRP-5-STATECHANGE: Ethernet0/0 Grp 1 state Listen ->
Active
R2(config-if)# end
R2#
```

You need to change the HSRP version from 2 to 1 on the R2 router. The default HSRP version is 1.

You also need to change the HSRP group from 2 to 1 for virtual IP address 10.10.1.1.

After you perform these two changes, routers R1 and R2 become HSRP-aware of each other. R2 becomes the active router.

Step 8 Display the HSRP status on R1.

Enter this command on the R1 router:

```
R1# show standby
Ethernet0/0 - Group 1
  State is Standby
    4 state changes, last state change 00:10:42
  Virtual IP address is 10.10.1.1
  Active virtual MAC address is 0000.0c07.ac01
    Local virtual MAC address is 0000.0c07.ac01 (v1 default)
  Hello time 3 sec, hold time 10 sec
    Next hello sent in 2.224 secs
  Preemption enabled
  Active router is 10.10.1.3, priority 100 (expires in 10.976 sec)
  Standby router is local
  Priority 99 (configured 99)
  Group name is "hsrp-Et0/0-1" (default)
```

Routers are now aware of each other as they are both configured for the same HSRP group with matching HSRP versions. R1 is the standby router because of the lower priority.

Step 9 Display the HSRP status on R2.

Enter this command on the R2 router:

```
R2# show standby
Ethernet0/0 - Group 1
  State is Active
    1 state change, last state change 00:16:59
  Virtual IP address is 10.10.1.1
  Active virtual MAC address is 0000.0c07.ac01
    Local virtual MAC address is 0000.0c07.ac01 (v1 default)
  Hello time 3 sec, hold time 10 sec
    Next hello sent in 1.968 secs
  Preemption enabled
  Active router is local
  Standby router is 10.10.1.2, priority 99 (expires in 7.872 sec)
  Priority 100 (default 100)
  Group name is "hsrp-Et0/0-1" (default)
```

R2 is the active router because the configured priority is higher.

Step 10 Access the console of R1 and increase the HSRP priority, so that R1 will become HSRP-active.

Enter these commands on the R1 router:

```
R1# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
R1(config)# interface Ethernet0/0
R1(config-if)# standby 1 priority 101
R1(config-if)#
*Nov 30 19:03:07.863: %HSRP-5-STATECHANGE: Ethernet0/0 Grp 1 state Standby ->
Active
R1(config-if)# end
R1#
```

As preemption is enabled on both routers R1 and R2, HSRP state changes. R1 becomes the active HSRP router.

Step 11 Display the HSRP status on R1.

Enter this command on the R1 router:

```
R1# show standby
Ethernet0/0 - Group 1
  State is Active
    5 state changes, last state change 00:04:49
  Virtual IP address is 10.10.1.1
  Active virtual MAC address is 0000.0c07.ac01
    Local virtual MAC address is 0000.0c07.ac01 (v1 default)
  Hello time 3 sec, hold time 10 sec
    Next hello sent in 2.256 secs
  Preemption enabled
  Active router is local
  Standby router is 10.10.1.3, priority 100 (expires in 9.312 sec)
  Priority 101 (configured 101)
  Group name is "hsrp-Et0/0-1" (default)
```

R1 is the active router with priority 101, R2 is the standby router with priority 100.

Also note that preemption is enabled.

Step 12 Display the HSRP status on R2.

Enter this command on the R2 router:

```
R2# show standby
Ethernet0/0 - Group 1
  State is Standby
    3 state changes, last state change 00:07:29
  Virtual IP address is 10.10.1.1
  Active virtual MAC address is 0000.0c07.ac01
    Local virtual MAC address is 0000.0c07.ac01 (v1 default)
  Hello time 3 sec, hold time 10 sec
    Next hello sent in 1.968 secs
  Preemption enabled
  Active router is 10.10.1.2, priority 101 (expires in 8.720 sec)
  Standby router is local
  Priority 100 (default 100)
  Group name is "hsrp-Et0/0-1" (default)
```

R2 is the standby router with priority 100, R1 is the active router with priority 101.

Also note that preemption is enabled.

This is the end of the discovery lab.

Challenge

1. What is the function of an FHRP?

   A. The FHRP supplies hosts with routing information.
   B. The FHRP is a routing protocol.
   C. The FHRP provides default gateway redundancy.
   D. The FHRP is standards-based.

2. Which of the following is not an HSRP state ? (Choose two)

   A. INIT
   B. ACTIVE
   C. ESTABLISHED
   D. IDLE

3. Which command configures an interface to enable HSRP with the virtual router IP address 10.10.1.1?

   A. **standby 1 ip 10.10.1.1**
   B. **ip hsrp 1 standby 10.10.1.1**
   C. **hsrp 1 ip 10.10.1.1**
   D. **standby 1 hsrp ip 10.10.1.1**

4. Which command displays the status of all HSRP groups on a Cisco router or Layer 3 switch?

   A. **show ip hsrp**
   B. **show hsrp**
   C. **show standby hsrp**
   D. **show standby**
   E. **show hsrp groups**

5. Two routers are part of HSRP standby group. There was no priority configured on the routers for the HSRP group. Which of the statements below is correct ?

   A. Both routers will be in ACTIVE state.
   B. Both routers will be in STANDBY state.
   C. Both routers will be in LISTEN state.
   D. One router will be ACTIVE and other STANDBY state
   E. None of the above.

6. Which of the following statement is true about the HSRP version 1 hello packet ?

   A. HSRP hello packets are sent to multicast address 224.0.0.2 with UDP port 1985.
   B. HSRP hello packets are sent to multicast address 224.0.0.5.
   C. HSRP hello packets are sent to multicast address 224.0.0.2 with TCP port 1985
   D. HSRP hello packets are sent to multicast address 224.0.0.10 with UDP port 1986.

7. R1 and R2 are in HSRP group 1. R1 is the active router with a priority of 120 and R2 has the default priority. Now, R1 reboots and so R2 becomes the active router. Once R1 is back up , which of the following statement will be true ?

   A. R1 will become the active router.
   B. R1 will become the active router again if preempt is enabled.
   C. Both routers will be in active state.
   D. Both routers will be in standby state.

Answer Key

Challenge

1. C
2. C, D
3. A
4. D
5. D
6. A
7. B

Lesson 7: Implementing RIPv2

Introduction

Conceptually, routing information takes the form of entries in a routing table, with one entry for each identified route. The network administrator can statically configure the entries so that they appear in the routing table, or the router can use a routing protocol to create and maintain the routing table dynamically to accommodate network changes whenever they occur.

Distance vector routing protocols include RIPv2, which is one of the most enduring of all routing protocols. RIPv2 is a relatively old, but still commonly used, IGP for use in small, homogeneous networks. RIPv2 is a classic distance vector routing protocol.

You need to show your boss that you understand the characteristics of RIPv2 and how to configure and verify it. You will also need to troubleshoot any issues that your customer might experience.

Introduction to Routing Protocols

A routing protocol is a set of processes, algorithms, and messages that is used to exchange routing information and populate the routing table with the choice of best paths for the routing protocol. Routing protocols are a set of rules by which routers dynamically share their routing information. As routers become aware of changes to the networks for which they act as the gateway, or changes to links between routers, they pass on this information to other routers. When a router receives information about new or changed routes, it updates its own routing table and, in turn, passes the information to other routers. In this way, all routers have accurate routing tables that are updated dynamically and can learn about routes to remote networks that are many hops away.

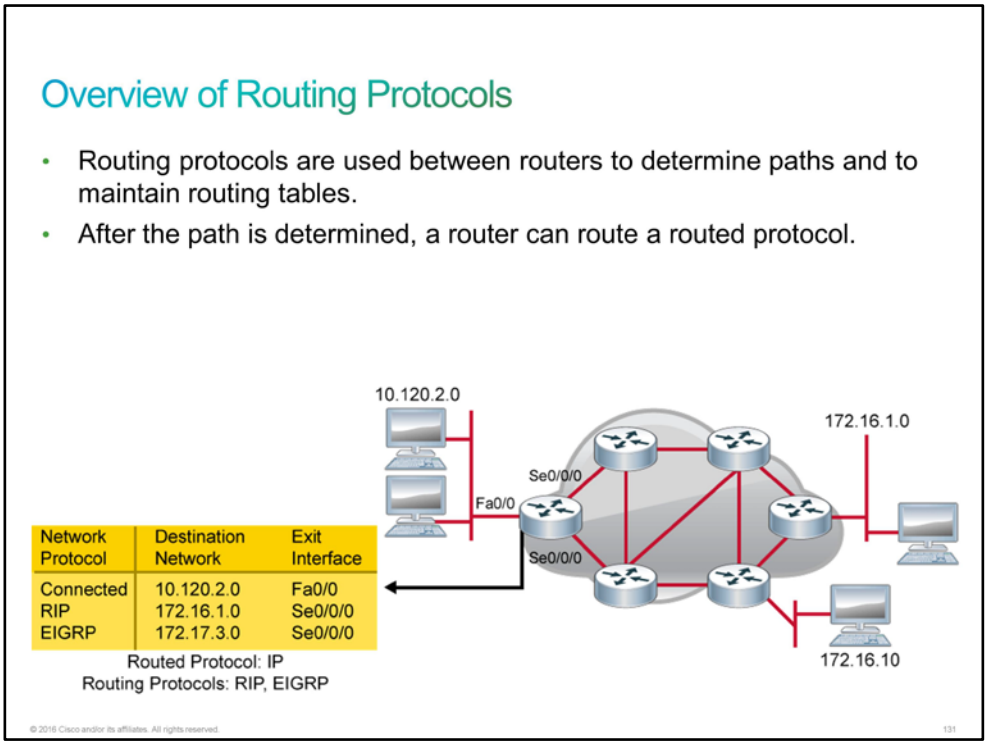

Further examples of the information that routing protocols determine are as follows:

- How updates are conveyed
- Which knowledge is conveyed
- When to convey knowledge
- How to locate recipients of the updates

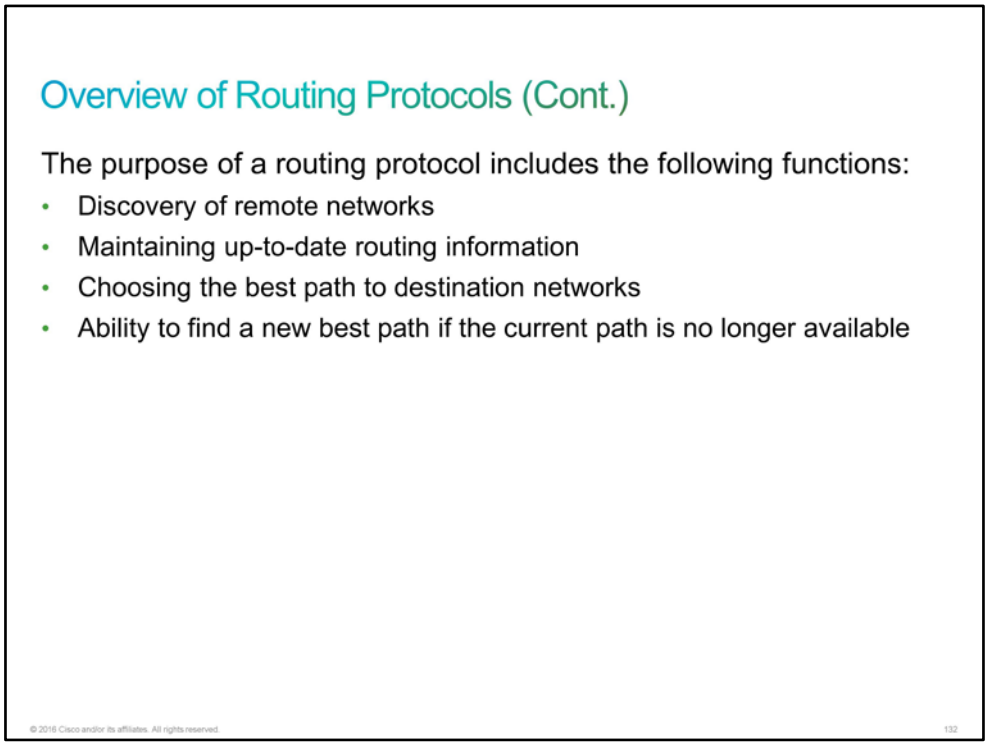

All routing protocols have the same purpose—to learn about remote networks and to quickly adapt whenever there is a change in the topology. The method that a routing protocol uses to accomplish this purpose depends upon the algorithm that it uses and the operational characteristics of that protocol.

The operations of a dynamic routing protocol vary depending on the type of routing protocol and on the routing protocol itself. In general, the operations of a dynamic routing protocol can be described as follows:

- The router sends and receives routing messages on its interfaces.
- The router shares routing messages and routing information with other routers that are using the same routing protocol.
- Routers exchange routing information to learn about remote networks.
- When a router detects a topology change, the routing protocol can advertise this change to other routers.

Although routing protocols provide routers with up-to-date routing tables, there are costs that put additional demands on the memory and processing power of the router. First, the exchange of route information adds overhead that consumes network bandwidth. This overhead can be a problem, particularly for low-bandwidth links between routers. Second, after the router receives the route information, protocols such as EIGRP and OSPF process it extensively to make routing table entries. It means that routers that are employing these protocols must have sufficient processing capacity to implement the algorithms of the protocol and to perform timely packet routing and forwarding.

Distance Vector and Link-State Routing Protocols

In the following example, you can see a hierarchical view of dynamic routing protocol classification.

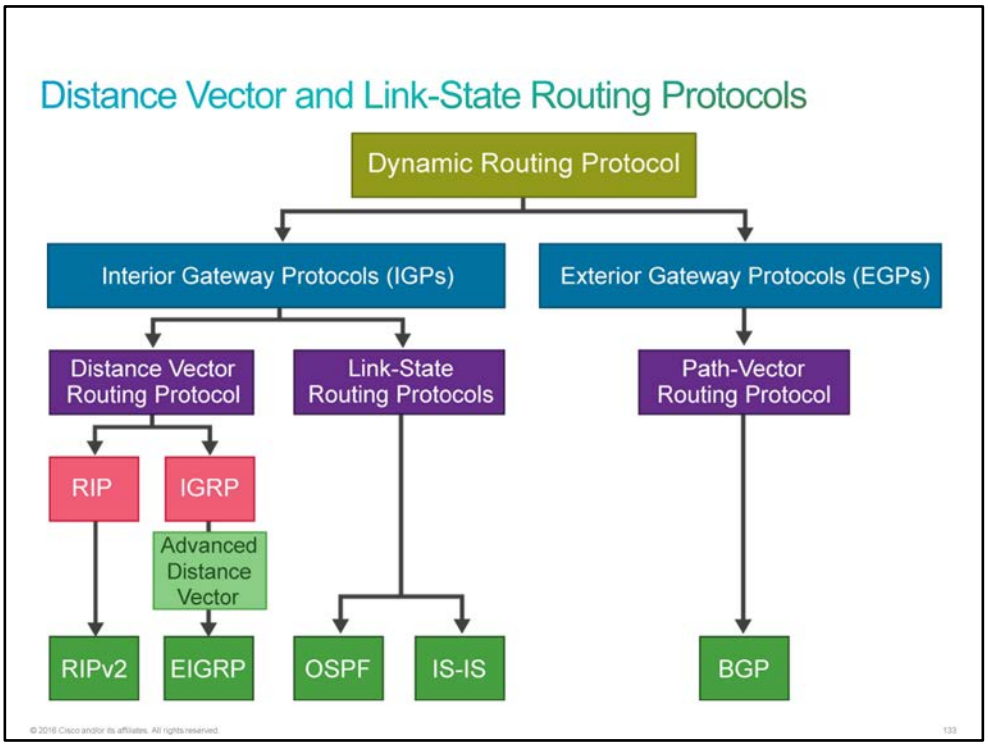

Within an AS, you can classify most IGP routing as conforming to one of these algorithms:

- **Distance vector:** The distance vector routing approach determines the direction (vector) and distance (a metric, such as hop count in the case of RIP) to any link in the internetwork. Pure distance vector protocols periodically send complete routing tables to all connected neighbors. This mode of operation is key in defining what a distance vector routing protocol is. In large networks, these routing updates can become enormous, causing significant traffic on the links. The only information that a router knows about a remote network is the distance or metric to reach this network and which path or interface to use to get there. Different distance vector routing protocols may use different kinds of metrics. Distance vector routing protocols do not have an actual map of the network topology. For a router, the view of the network is based on the information that its neighbors provide.
- **Advanced distance vector:** The advanced distance vector approach combines aspects of the link-state and distance vector algorithms. EIGRP is a Cisco proprietary routing protocol that combines the advantages of link-state and distance vector routing protocols. EIGRP may act like a link-state routing protocol because it uses a Hello protocol to discover neighbors and form neighbor relationships and because only partial updates are sent when a change occurs. However, EIGRP is still based on the key distance vector routing protocol principle that information about the rest of the network is learned from directly connected neighbors.
- **Link-state:** The link-state approach, which uses the SPF algorithm, creates an abstraction of the exact topology of the entire internetwork, or at least of the partition in which the router is situated. A link-state router uses the link-state information to create a topology map and to select the best path to all destination networks in the topology. All link-state routers use an identical "map" of the network and calculate the shortest paths to reach the destination networks in relation to where they are on this map. Unlike their distance vector counterparts, complete routing tables are *not* exchanged periodically. Instead, event-based, "triggered" updates containing only specific link-state information are sent. Periodic keepalives that are small and efficient, in the form of hello messages, are exchanged between directly connected neighbors to establish and maintain reachability to this neighbor.

Note	Autonomous System, also known as AS, is a collection of networks within a common administrative domain.

Understanding RIPv2

RIPv2 is a standardized routing protocol that works in a mixed vendor router environment.

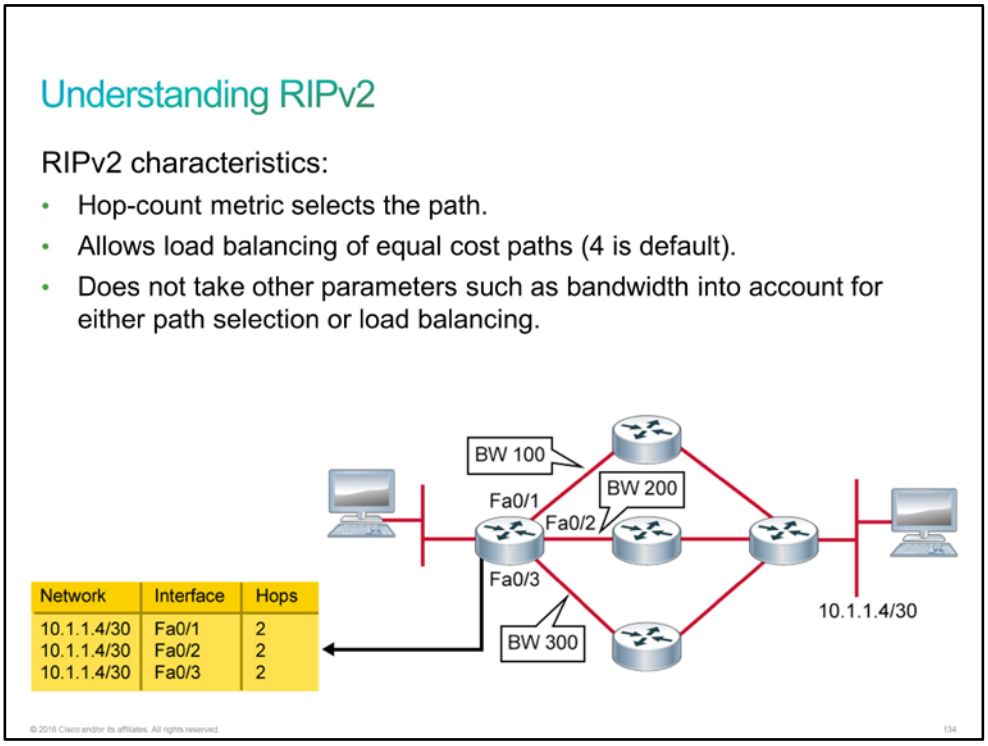

RIP is one of the easiest routing protocols to configure, making it a good choice for small networks. However, RIPv2 still has limitations. RIPv2 has a route metric that is based only on hop count and is limited to 15 hops.

Configure RIPv2

The following are the steps that you need to take to configure RIP version 2 (also known as RIPv2) on a Cisco IOS router:

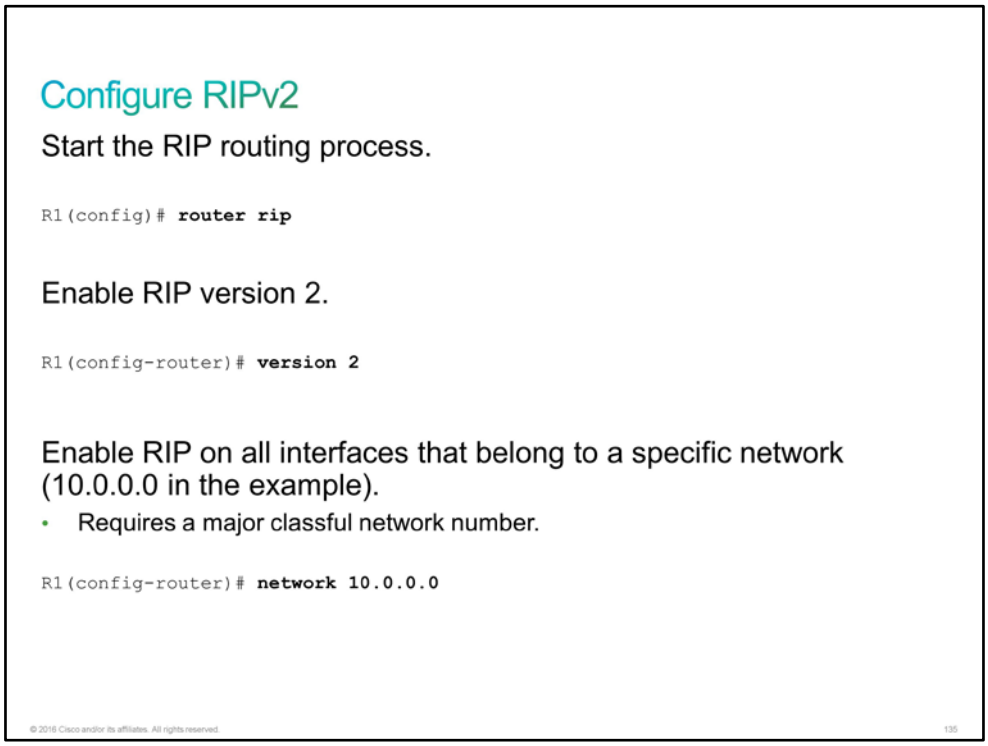

To enter the router configuration mode for RIP, enter the **router rip** command at the global configuration prompt. Notice that the prompt changes from a global configuration prompt to the router configuration prompt.

This command does not directly start the RIP process. Instead, it provides access to configure routing protocol settings. The device does not send any routing updates.

To enable RIP routing for a network, use the **network** command in the router configuration mode. The **network** command assigns a major network number to which the router is directly connected. The RIP routing process associates interface addresses with the advertised network number and will begin RIP packet processing on the specified interfaces.

The **network** command performs the following functions:

- Enables RIP on all interfaces that belong to a specific network. Associated interfaces will now both send and receive RIP updates.
- Advertises the specified network in RIP routing updates that a device sends to other routers every 30 seconds.

The **network 0.0.0.0** command specifies all networks.

By default, when you configure a RIP process on a Cisco router, the router is running RIPv1. To use RIPv2, enter the **version 2** command. You should configure this command on all routers in the routing domain. The RIP process will now include the subnet mask in all updates, making RIPv2 a classless routing protocol.

Verify RIPv2

The **show ip protocols** command displays values about the routing protocols and the routing protocol timer information that is associated with the router.

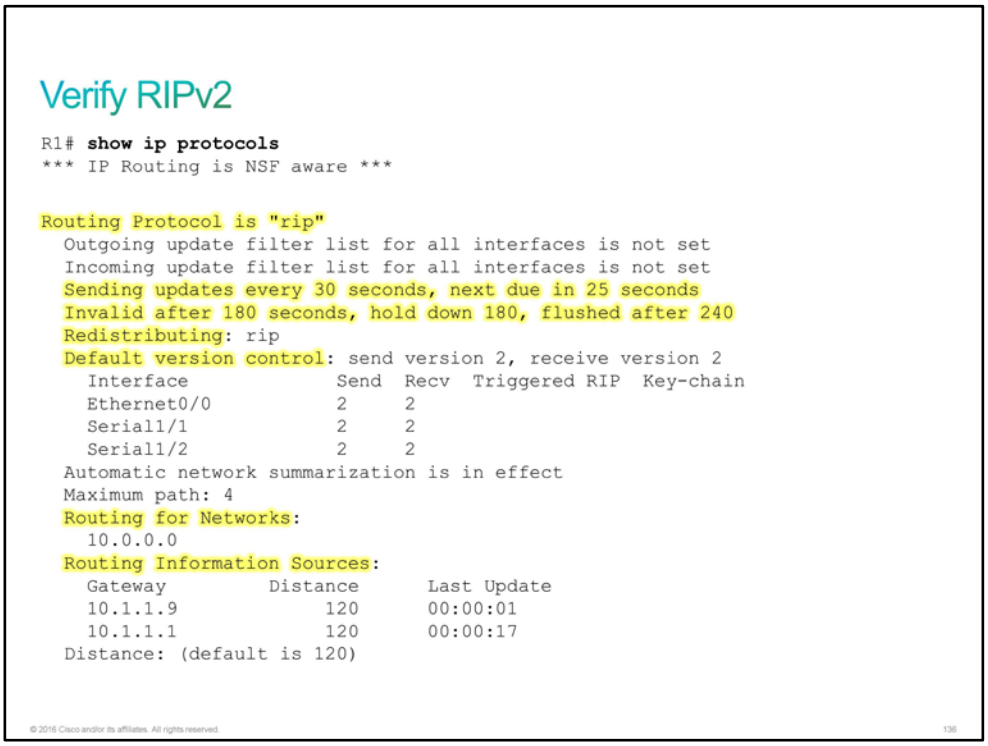

This table describes the significant fields that the display shows.

Field	Description
Routing Protocol is "rip"	Specifies the routing protocol that is used.
Sending updates every 30 seconds	Specifies the time between sending updates.
next due in 25 seconds	Specifies when the device will send the next update.
Invalid after 180 seconds	Specifies the value of the invalid parameter. Invalid timer is an interval of time after which a route is declared invalid; it should be at least three times the value of the update argument. A route becomes invalid when no updates refresh the route. The route then enters into a holddown state where it is marked as inaccessible and advertised as unreachable. However, the route is still used to forward packets. The range is from 1 to 4,294,967,295. The default is 180 seconds.

Field	Description
hold down 180	Specifies the current value (in seconds) of the hold-down parameter. Holddown timer is an interval during which routing information regarding better paths is suppressed; it should be at least three times the value of the update argument. A route enters into a holddown state when an update packet is received that indicates that the route is unreachable. The route is marked as inaccessible and advertised as unreachable. However, the route is still used to forward packets. When holddown expires, routes advertised by other sources are accepted and the route is no longer inaccessible. The range is from 0 to 4,294,967,295. The default is 180 seconds.
flushed after 240	Specifies the time (in seconds) after which the individual routing information will be thrown (flushed) out. Flushed timer is an amount of time that must pass before the route is removed from the routing table; the interval that is specified should be greater than the sum of the invalid argument plus the holddown argument. If it is less than this sum, the proper holddown interval cannot elapse, which results in a new route being accepted before the holddown interval expires. The range is from 1 to 4,294,967,295. The default is 240 seconds.
Redistributing	Lists the protocol that is being redistributed.
Default version control	Specifies the version of RIP packets that the device is sending and receiving.
Routing for Networks	Specifies the networks for which the routing process is currently injecting routes.
Routing Information Sources	Lists all the routing sources that the Cisco IOS Software is using to build its routing table. For each source, you will see the following information displayed: IP address, administrative distance, and time when the device received the last update from this source.

The **show ip route** command displays the contents of the IP routing table.

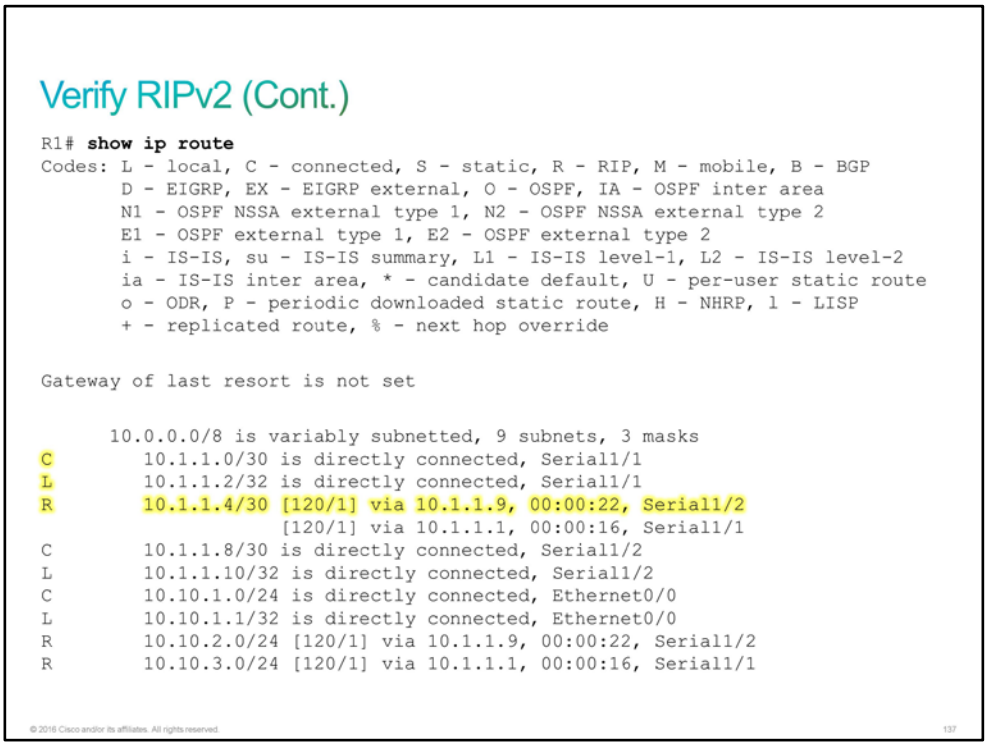

The routing table contains entries for all known networks and subnetworks, and a code that indicates how that information was learned. This table explains the output and function of key fields from the **show ip route** command.

Output	Description
C, L, or R	Identifies the source of the route. For example, a "C" indicates that the route came from a direct connection of the route to a router interface. The "L" indicates a local route. The "R" indicates that RIP is the protocol that determined the route.
10.1.1.4/30	Indicates the address of the remote network.
120/1	The first number is the administrative distance of the information source. The second number is the metric for the route (1 hop in this example).
via 10.1.1.9	Specifies the address of the next-hop router to the remote network.
00:00:22	Specifies the amount of time since the route was last updated (22 sec in this example).
Serial1/2	Specifies the interface through which the specified network can be reached.

If the device is not exchanging the routing information, use the **show running-config** or **show ip protocols** commands on the router to check for a possible misconfigured routing protocol.

Discovery 26: Configure and Verify RIPv2

Introduction

In this discovery lab, you will configure and verify RIPv2 for IPv4. You will adjust RIP timers, disable automatic summarization, configure a passive interface, and generate a default route into RIP.

The lab is prepared with the devices as represented in the topology diagram and connectivity table. All devices have their basic configurations in place including the hostnames and IP addresses. Default gateways are defined on PC1, PC2, and SRV1, but no other routing has been configured.

Topology

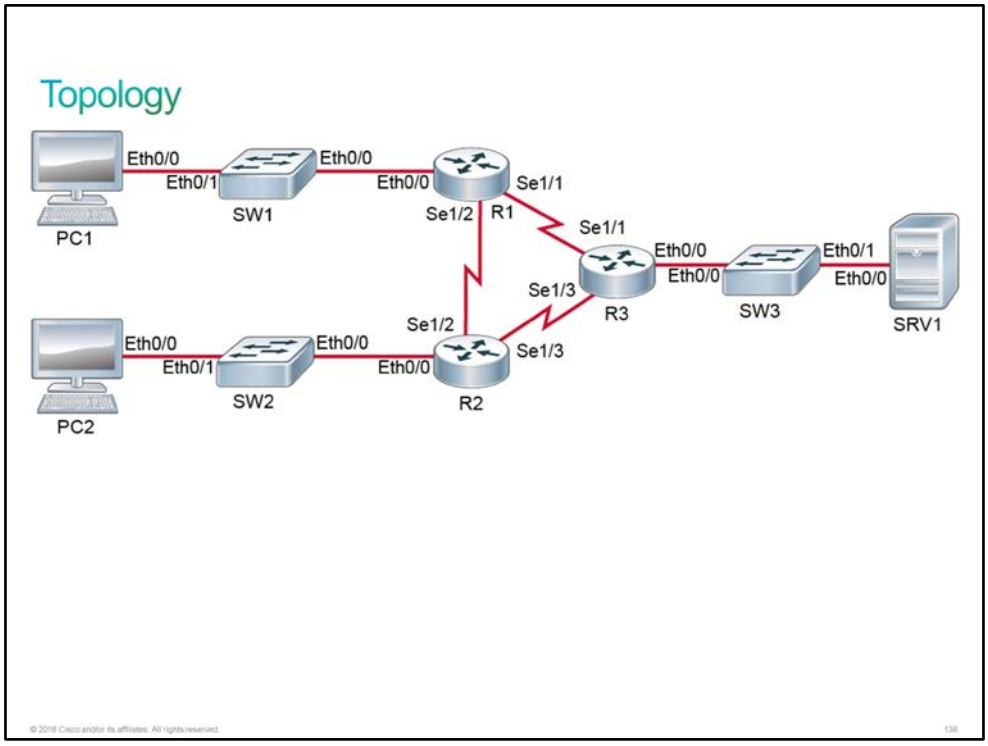

Job Aid

Device Information

Device Information Table

Device	Characteristic	Value
PC1	Hostname	PC1
PC1	IP address	10.10.1.10/24
PC1	Default gateway	10.10.1.1

Device	Characteristic	Value
PC2	Hostname	PC2
PC2	IP address	10.10.2.20/24
PC2	Default gateway	10.10.2.1
SRV1	Hostname	SRV1
SRV1	IP address	10.10.3.30/24
SRV1	Default gateway	10.10.3.1
SW1	Hostname	SW1
SW1	VLAN 1 IP address	10.10.1.4/24
SW1	Default gateway	10.10.1.1
SW1	Ethernet0/0 description	Link to R1
SW1	Ethernet0/1 description	Link to PC1
SW2	Hostname	SW2
SW2	VLAN 1 IP address	10.10.2.4/24
SW2	Default gateway	10.10.2.1
SW2	Ethernet0/0 description	Link to R2
SW2	Ethernet0/1 description	Link to PC2
SW3	Hostname	SW3
SW3	VLAN 1 IP address	10.10.3.4/24
SW3	Default gateway	10.10.3.1
SW3	Ethernet0/0 description	Link to R3
SW3	Ethernet0/1 description	Link to SRV1
R1	Hostname	R1
R1	Ethernet0/0 description	Link to SW1
R1	Ethernet0/0 IP address	10.10.1.1/24
R1	Serial1/1 description	Link to R3

Device	Characteristic	Value
R1	Serial1/1 IP address	10.1.1.2/30
R1	Serial1/2 description	Link to R2
R1	Serial1/2 IP address	10.1.1.10/30
R2	Hostname	R2
R2	E0/0 description	Link to SW2
R2	E0/0 IP address	10.10.2.1/24
R2	Serial1/2 description	Link to R1
R2	Serial1/2 IP address	10.1.1.9/30
R2	Serial1/3 description	Link to R3
R2	Serial1/3 IP address	10.1.1.6/30
R3	Hostname	R3
R3	Ethernet0/0 description	Link to SW3
R3	Ethernet0/0 IP address	10.10.3.1/24
R3	Serial1/1 description	Link to R1
R3	Serial1/1 IP address	10.1.1.1/30
R3	Serial1/3 description	Link to R2
R3	Serial1/3 IP address	10.1.1.5/30

PCs and SRV in the virtual lab environment are simulated as routers, so you should use Cisco IOS commands to configure them or make verifications.

Task 1: Configure and Verify RIPv2

Activity

Step 1 On the R1, R2, and R3 routers, enable RIPv2 and include the network 10.0.0.0/8 into RIP.

On R1, enter the following commands:

```
R1# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
R1(config)# router rip
R1(config-router)# version 2
R1(config-router)# network 10.0.0.0
R1(config-router)# end
R1#
```

On R2, enter the following commands:

```
R2# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
R2(config)# router rip
R2(config-router)# version 2
R2(config-router)# network 10.0.0.0
R2(config-router)# end
R2#
```

On R3, enter the following commands:

```
R3# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
R3(config)# router rip
R3(config-router)# version 2
R3(config-router)# network 10.0.0.0
R3(config-router)# end
R3#
```

Step 2 On R1, verify the IP routing table.

On R1, enter the following command:

```
R1# sh 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
       + - replicated route, % - next hop override

Gateway of last resort is not set

      10.0.0.0/8 is variably subnetted, 9 subnets, 3 masks
C        10.1.1.0/30 is directly connected, Serial1/1
L        10.1.1.2/32 is directly connected, Serial1/1
R        10.1.1.4/30 [120/1] via 10.1.1.9, 00:00:22, Serial1/2
                     [120/1] via 10.1.1.1, 00:00:16, Serial1/1
C        10.1.1.8/30 is directly connected, Serial1/2
L        10.1.1.10/32 is directly connected, Serial1/2
C        10.10.1.0/24 is directly connected, Ethernet0/0
L        10.10.1.1/32 is directly connected, Ethernet0/0
R        10.10.2.0/24 [120/1] via 10.1.1.9, 00:00:22, Serial1/2
R        10.10.3.0/24 [120/1] via 10.1.1.1, 00:00:16, Serial1/1
```

You should see many RIP routes in the routing table.

Task 2: Adjust RIP Timers

Activity

Step 1 On R1, verify RIP timers.

On R1, enter the following command:

```
R1# show ip protocols
*** IP Routing is NSF aware ***

Routing Protocol is "rip"
  Outgoing update filter list for all interfaces is not set
  Incoming update filter list for all interfaces is not set
  Sending updates every 30 seconds, next due in 25 seconds
  Invalid after 180 seconds, hold down 180, flushed after 240
  Redistributing: rip
  Default version control: send version 2, receive version 2
    Interface             Send  Recv  Triggered RIP  Key-chain
    Ethernet0/0           2     2
    Serial1/1             2     2
    Serial1/2             2     2
  Automatic network summarization is in effect
  Maximum path: 4
  Routing for Networks:
    10.0.0.0
  Routing Information Sources:
    Gateway         Distance      Last Update
    10.1.1.9             120      00:00:01
    10.1.1.1             120      00:00:17
  Distance: (default is 120)
```

Updates are sent every 30 second. Invalid and hold-down timers are 180 seconds. The flush timer is 240 seconds.

Step 2 On R1, change the update timer to 60 seconds and leave other timers default.

On R1, enter the following commands:

```
R1# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
R1(config)# router rip
R1(config-router)# timers basic 60 180 180 240
R1(config-router)# end
R1#
```

On R1, examine how interface configuration changed. You should see **ip rip advertise** command line under RIP enabled interfaces.

```
R1# show running-config | section rip
 description Link to SW1
 ip rip advertise 60
 description Link to R3
 ip rip advertise 60
 description Link to R2
 ip rip advertise 60
router rip
 version 2
 timers basic 60 180 180 240
 network 10.0.0.0
```

Router sets **ip rip advertise** command under interface configuration mode only first time when timer is changed from default value. If you want to further adjust timer, you will need to change interface configuration as well as router rip configuration.

Step 3 On R1, verify the RIP timers again.

On R1, enter the following command:

```
R1# show ip protocols
*** IP Routing is NSF aware ***

Routing Protocol is "rip"
  Outgoing update filter list for all interfaces is not set
  Incoming update filter list for all interfaces is not set
  Sending updates every 60 seconds, next due in 47 seconds
  Invalid after 180 seconds, hold down 180, flushed after 240
  Redistributing: rip
  Default version control: send version 2, receive version 2
    Interface             Send  Recv  Triggered RIP  Key-chain
    Ethernet0/0           2     2
    Serial1/1             2     2
    Serial1/2             2     2
  Automatic network summarization is in effect
  Maximum path: 4
  Routing for Networks:
    10.0.0.0
  Routing Information Sources:
    Gateway         Distance      Last Update
    10.1.1.9             120      00:00:06
    10.1.1.1             120      00:00:01
  Distance: (default is 120)
```

You can see that the update timer has changed to 60 seconds. All other timers remain default.

You may continue this task, by changing update timer to 60 seconds, on R2 and R3.

Task 3: Disable RIP Auto Summary

Activity

Step 1 On R2, create the Loopback0 interface, assign the IP address 192.168.1.5/30, and include this network into RIP.

On R2, enter the following commands:


```
R2# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
R2(config)# int lo 0
R2(config-if)# ip address 192.168.1.5 255.255.255.252
R2(config-if)# router rip
R2(config-router)# network 192.168.1.0
R2(config-router)# end
R2#
```

Step 2 On R1, examine the IP routing table. You should see a new RIP route.

On R1, enter the following command:

```
R1# sh 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
       + - replicated route, % - next hop override

Gateway of last resort is not set

      10.0.0.0/8 is variably subnetted, 9 subnets, 3 masks
C        10.1.1.0/30 is directly connected, Serial1/1
L        10.1.1.2/32 is directly connected, Serial1/1
R        10.1.1.4/30 [120/1] via 10.1.1.9, 00:00:05, Serial1/2
                     [120/1] via 10.1.1.1, 00:00:18, Serial1/1
C        10.1.1.8/30 is directly connected, Serial1/2
L        10.1.1.10/32 is directly connected, Serial1/2
C        10.10.1.0/24 is directly connected, Ethernet0/0
L        10.10.1.1/32 is directly connected, Ethernet0/0
R        10.10.2.0/24 [120/1] via 10.1.1.9, 00:00:05, Serial1/2
R        10.10.3.0/24 [120/1] via 10.1.1.1, 00:00:18, Serial1/1
R     192.168.1.0/24 [120/1] via 10.1.1.9, 00:00:05, Serial1/2
```

You will not see the original R2 Loopback0 subnet mask (/30) but rather a whole network mask (/24), because of the automatic summarization on R1.

Step 3 On R2, disable automatic RIP summarization.

On R2, enter the following commands:

```
R2# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
R2(config)# router rip
R2(config-router)# no auto-summary
R2(config-router)# end
R2#
```

Step 4 On R1, examine the IP routing table again. You should see the R2 Loopback subnet (/30) advertised with its original subnet mask.

On R1, enter the following command:

```
R1# sh 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
       + - replicated route, % - next hop override

Gateway of last resort is not set

      10.0.0.0/8 is variably subnetted, 9 subnets, 3 masks
C        10.1.1.0/30 is directly connected, Serial1/1
L        10.1.1.2/32 is directly connected, Serial1/1
R        10.1.1.4/30 [120/1] via 10.1.1.9, 00:00:03, Serial1/2
                     [120/1] via 10.1.1.1, 00:00:22, Serial1/1
C        10.1.1.8/30 is directly connected, Serial1/2
L        10.1.1.10/32 is directly connected, Serial1/2
C        10.10.1.0/24 is directly connected, Ethernet0/0
L        10.10.1.1/32 is directly connected, Ethernet0/0
R        10.10.2.0/24 [120/1] via 10.1.1.9, 00:00:03, Serial1/2
R        10.10.3.0/24 [120/1] via 10.1.1.1, 00:00:22, Serial1/1
      192.168.1.0/24 is variably subnetted, 2 subnets, 2 masks
R        192.168.1.0/24 [120/2] via 10.1.1.1, 00:00:22, Serial1/1
R        192.168.1.4/30 [120/1] via 10.1.1.9, 00:00:03, Serial1/2
```

You will still see the route for the network 192.168.1.0/24 because R3 is automatically summarizing the subnet 192.168.1.4/30 from R2.

To completely eliminate the network 192.168.1.0/24 from the R1 routing table, you will need to disable automatic RIP summarization on R3. Feel free to continue with independent exploration of the use of RIP automatic summarization within the lab environment.

If you do not want to wait for RIP timers to expire, clear IP routing table on R1.

```
R1# clear ip route *
```

Task 4: Configure a RIP Passive Interface

Activity

Step 1 On R3, examine to which interfaces RIP is sending updates.

On R3, enter the following command:


```
R3# sh ip prot
*** IP Routing is NSF aware ***

Routing Protocol is "rip"
  Outgoing update filter list for all interfaces is not set
  Incoming update filter list for all interfaces is not set
  Sending updates every 30 seconds, next due in 17 seconds
  Invalid after 180 seconds, hold down 180, flushed after 240
  Redistributing: rip
  Default version control: send version 2, receive version 2
    Interface             Send  Recv  Triggered RIP  Key-chain
    Ethernet0/0           2     2
    Serial1/1             2     2
    Serial1/3             2     2
  Automatic network summarization is not in effect
  Maximum path: 4
  Routing for Networks:
    10.0.0.0
  Routing Information Sources:
    Gateway         Distance      Last Update
    10.1.1.2             120      00:00:10
    10.1.1.6             120      00:00:18
  Distance: (default is 120)
```

RIP is sending updates to all interfaces that are configured in the 10.0.0.0/8 network.

Step 2 Because there are no RIP speaking routers that are connected to the Ethernet0/0 interface on R3, disable RIP updates on this interface.

On R3, enter the following commands:

```
R3# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
R3(config)# router rip
R3(config-router)# passive-interface Ethernet 0/0
R3(config-router)# end
R3#
```

Step 3 On R3, examine to which interfaces RIP is sending updates again.

On R3, enter the following command:

```
R3# sh ip prot
*** IP Routing is NSF aware ***

Routing Protocol is "rip"
  Outgoing update filter list for all interfaces is not set
  Incoming update filter list for all interfaces is not set
  Sending updates every 30 seconds, next due in 8 seconds
  Invalid after 180 seconds, hold down 180, flushed after 240
  Redistributing: rip
  Default version control: send version 2, receive version 2
    Interface             Send  Recv  Triggered RIP  Key-chain
    Serial1/1             2     2
    Serial1/3             2     2
  Automatic network summarization is not in effect
  Maximum path: 4
  Routing for Networks:
    10.0.0.0
  Passive Interface(s):
    Ethernet0/0
  Routing Information Sources:
    Gateway         Distance      Last Update
    10.1.1.2             120      00:00:44
    10.1.1.6             120      00:00:09
  Distance: (default is 120)
```

The R3 router does not send RIP updates to the Ethernet0/0 interface anymore.

Step 4 On R2, examine to which interfaces RIP is sending updates.

On R2, enter the following command:

```
R2# sh ip protocols
*** IP Routing is NSF aware ***

Routing Protocol is "rip"
  Outgoing update filter list for all interfaces is not set
  Incoming update filter list for all interfaces is not set
  Sending updates every 30 seconds, next due in 23 seconds
  Invalid after 180 seconds, hold down 180, flushed after 240
  Redistributing: rip
  Default version control: send version 2, receive version 2
    Interface             Send  Recv  Triggered RIP  Key-chain
    Ethernet0/0           2     2
    Serial1/2             2     2
    Serial1/3             2     2
    Loopback0             2     2
  Automatic network summarization is not in effect
  Maximum path: 4
  Routing for Networks:
    10.0.0.0
    192.168.1.0
  Routing Information Sources:
    Gateway         Distance      Last Update
    10.1.1.10            120      00:00:39
    10.1.1.5             120      00:00:15
  Distance: (default is 120)
```

There are many unnecessary interfaces (Ethernet0/0 and Loopback0) to which RIP is sending updates.

Step 5 On R2, configure RIP with the all passive interfaces functionality and then disable the passive interface functionality for Serial1/2 and Serial1/3 interfaces.

On R2, enter the following commands:

```
R2# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
R2(config)# router rip
R2(config-router)# passive-interface default
R2(config-router)# no passive-interface Serial 1/2
R2(config-router)# no passive-interface Serial 1/3
R2(config-router)# end
R2#
```

Step 6 On R2, examine to which interfaces RIP is sending updates again.

On R2, enter the following command:

```
R2# sh ip protocols
*** IP Routing is NSF aware ***

Routing Protocol is "rip"
  Outgoing update filter list for all interfaces is not set
  Incoming update filter list for all interfaces is not set
  Sending updates every 30 seconds, next due in 26 seconds
  Invalid after 180 seconds, hold down 180, flushed after 240
  Redistributing: rip
  Default version control: send version 2, receive version 2
    Interface             Send  Recv  Triggered RIP  Key-chain
    Serial1/2             2     2
    Serial1/3             2     2
  Automatic network summarization is not in effect
  Maximum path: 4
  Routing for Networks:
    10.0.0.0
    192.168.1.0
  Passive Interface(s):
    Ethernet0/0
    Ethernet0/1
    Ethernet0/2
    Ethernet0/3
    Serial1/0
  Passive Interface(s):
    Serial1/1
    Loopback0
    RG-AR-IF-INPUT1
    VoIP-Null0
  Routing Information Sources:
    Gateway         Distance      Last Update
    10.1.1.10            120      00:00:05
    10.1.1.5             120      00:00:09
  Distance: (default is 120)
```

The R2 router sends RIP updates only to Serial1/2 and Serial1/3 interfaces.

Task 5: Generate a Default RIP Route

Activity

Step 1 On R3, configure RIP to originate the default route.

On R3, enter the following commands:

```
R3# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
R3(config)# router rip
R3(config-router)# default-information originate
R3(config-router)# end
R3#
```

Step 2 On R1, verify the routing table. You should see the default RIP route in the routing table.

On R1, enter the following command:

```
R1# sh 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
       + - replicated route, % - next hop override

Gateway of last resort is 10.1.1.1 to network 0.0.0.0

R*    0.0.0.0/0 [120/1] via 10.1.1.1, 00:00:04, Serial1/1
      10.0.0.0/8 is variably subnetted, 9 subnets, 3 masks
C        10.1.1.0/30 is directly connected, Serial1/1
L        10.1.1.2/32 is directly connected, Serial1/1
R        10.1.1.4/30 [120/1] via 10.1.1.9, 00:00:13, Serial1/2
                     [120/1] via 10.1.1.1, 00:00:24, Serial1/1
C        10.1.1.8/30 is directly connected, Serial1/2
L        10.1.1.10/32 is directly connected, Serial1/2
C        10.10.1.0/24 is directly connected, Ethernet0/0
L        10.10.1.1/32 is directly connected, Ethernet0/0
R        10.10.2.0/24 [120/1] via 10.1.1.9, 00:00:13, Serial1/2
R        10.10.3.0/24 [120/1] via 10.1.1.1, 00:00:24, Serial1/1
      192.168.1.0/24 is variably subnetted, 2 subnets, 2 masks
R        192.168.1.0/24 [120/4] via 10.1.1.9, 00:00:13, Serial1/2
R        192.168.1.4/30 [120/1] via 10.1.1.9, 00:00:13, Serial1/2
```

During this discovery lab, you have explored RIPv2 configuration, RIP timers, RIP auto summary, RIP passive interface, and how to introduce a default route into RIP. Feel free to continue with independent exploration of the use of RIP within the lab environment.

This is the end of the discovery lab.

Discovery 27: Troubleshoot RIPv2

Introduction

In this discovery lab, you will troubleshoot RIP. The lab is prepared with the devices as represented in the topology diagram and connectivity table. All devices have their basic configurations in place including hostnames and IP addresses. Default gateways are defined on PC1, PC2, and SRV1. The RIP routing protocol is configured between R1, R2, and R3 routers.

In the RIP that is configured on the R1, R2, and R3 routers there are some issues.

Topology

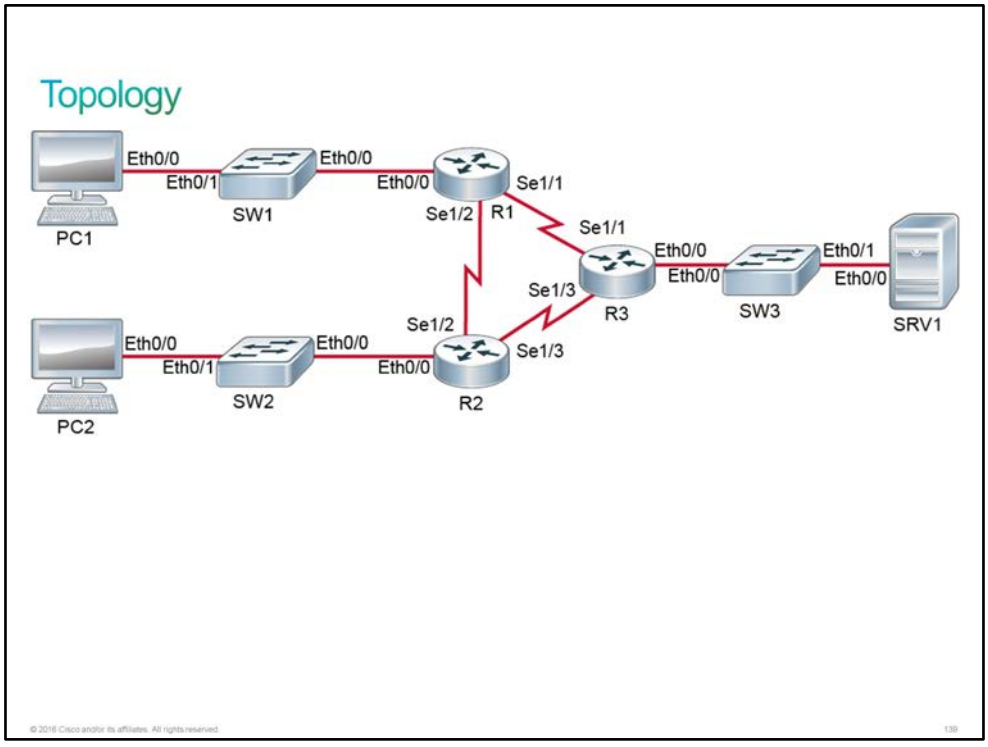

Job Aid

Device Information

Device Information Table

Device	Characteristic	Value
PC1	Hostname	PC1
PC1	IP address	10.10.1.10/24
PC1	Default gateway	10.10.1.1

Device	Characteristic	Value
PC2	Hostname	PC2
PC2	IP address	10.10.2.20/24
PC2	Default gateway	10.10.2.1
SRV1	Hostname	SRV1
SRV1	IP address	10.10.3.30/24
SRV1	Default gateway	10.10.3.1
SW1	Hostname	SW1
SW1	VLAN 1 IP address	10.10.1.4/24
SW1	Default gateway	10.10.1.1
SW1	Ethernet0/0 description	Link to R1
SW1	Ethernet0/1 description	Link to PC1
SW2	Hostname	SW2
SW2	VLAN 1 IP address	10.10.2.4/24
SW2	Default gateway	10.10.2.1
SW2	Ethernet0/0 description	Link to R2
SW2	Ethernet0/1 description	Link to PC2
SW3	Hostname	SW3
SW3	VLAN 1 IP address	10.10.3.4/24
SW3	Default gateway	10.10.3.1
SW3	Ethernet0/0 description	Link to R3
SW3	Ethernet0/1 description	Link to SRV1
R1	Hostname	R1
R1	Ethernet0/0 description	Link to SW1
R1	Ethernet0/0 IP address	10.10.1.1/24
R1	Serial1/1 description	Link to R3

Device	Characteristic	Value
R1	Serial1/1 IP address	10.1.1.2/30
R1	Serial1/2 description	Link to R2
R1	Serial1/2 IP address	10.1.1.10/30
R2	Hostname	R2
R2	Ethernet0/0 description	Link to SW2
R2	Ethernet0/0 IP address	10.10.2.1/24
R2	Serial1/2 description	Link to R1
R2	Serial1/2 IP address	10.1.1.9/30
R2	Serial1/3 description	Link to R3
R2	Serial1/3 IP address	10.1.1.6/30
R3	Hostname	R3
R3	Ethernet0/0 description	Link to SW3
R3	Ethernet0/0 IP address	10.10.3.1/24
R3	Serial1/1 description	Link to R1
R3	Serial1/1 IP address	10.1.1.1/30
R3	Serial1/3 description	Link to R2
R3	Serial1/3 IP address	10.1.1.5/30

PCs and SRV in the virtual lab environment are simulated as routers, so you should use Cisco IOS commands to configure them or make verifications.

Task 1: Troubleshoot RIPv2

Activity

Step 1 From the R1 router, verify connectivity to the R3 router Ethernet0/0 interface IP address (10.10.3.1).

On R1, enter the following command:

```
R1# ping 10.10.3.1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.10.3.1, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
```

Ping is not successful.

Step 2 On the R1 router, verify IP routing table

On R1, enter the following command:

```
R1# sh 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
       + - replicated route, % - next hop override

Gateway of last resort is not set

      10.0.0.0/8 is variably subnetted, 6 subnets, 3 masks
C        10.1.1.0/30 is directly connected, Serial1/1
L        10.1.1.2/32 is directly connected, Serial1/1
C        10.1.1.8/30 is directly connected, Serial1/2
L        10.1.1.10/32 is directly connected, Serial1/2
C        10.10.1.0/24 is directly connected, Ethernet0/0
L        10.10.1.1/32 is directly connected, Ethernet0/0
```

Note that the R1 router has no RIP routes in the routing table.

Step 3 Because the R1 router is not receiving any RIP routes from the R3 router, verify the RIP database on the R3 router.

On R3, enter the following command:

```
R3# show ip rip database
R3#
```

You will see that the R3 router has an empty RIP database.

Step 4 On R3, examine for which networks RIP is enabled.

On R3, enter the following command:


```
R3# sh ip protocols
*** IP Routing is NSF aware ***

Routing Protocol is "rip"
  Outgoing update filter list for all interfaces is not set
  Incoming update filter list for all interfaces is not set
  Sending updates every 300 seconds, next due in 0 seconds
  Invalid after 180 seconds, hold down 180, flushed after 240
  Redistributing: rip
  Default version control: send version 2, receive version 2
  Automatic network summarization is in effect
  Maximum path: 4
  Routing for Networks:
    1.0.0.0
  Passive Interface(s):
    Ethernet0/0
  Routing Information Sources:
    Gateway         Distance      Last Update
  Distance: (default is 120)
```

Note that RIP should be enabled for the network 10.0.0.0/8 and not 1.0.0.0/8.

Step 5 On R3, correct the configuration by adding the correct network into RIP.

On R3, enter the following commands:

```
R3# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
R3(config)# router rip
R3(config-router)# no network 1.0.0.0
R3(config-router)# network 10.0.0.0
R3(config-router)# end
R3#
```

Step 6 On R3, verify the RIP database again.

On R3, enter the following command:

```
R3# show ip rip database
10.0.0.0/8    auto-summary
10.1.1.0/30    directly connected, Serial1/1
10.1.1.4/30    directly connected, Serial1/3
10.1.1.8/30
    [1] via 10.1.1.2, 00:00:00, Serial1/1
10.10.1.0/24
    [1] via 10.1.1.2, 00:00:00, Serial1/1
10.10.3.0/24    directly connected, Ethernet0/0
```

You can see many subnets from the 10.0.0.0/8 network, which indicates that RIP on R3 is enabled for the correct network.

Note that there is no subnet coming from the R2 router (10.10.2.0/24).

Step 7 On R2, verify the RIP protocol information to examine if RIP is enabled for the correct network.

On R2, enter the following command:

```
R2# show ip protocols
*** IP Routing is NSF aware ***

Routing Protocol is "rip"
  Outgoing update filter list for all interfaces is not set
  Incoming update filter list for all interfaces is not set
  Sending updates every 300 seconds, next due in 285 seconds
  Invalid after 180 seconds, hold down 180, flushed after 240
  Redistributing: rip
  Default version control: send version 2, receive version 2
  Automatic network summarization is in effect
  Maximum path: 4
  Routing for Networks:
    10.0.0.0
  Passive Interface(s):
    Ethernet0/0
    Ethernet0/1
    Ethernet0/2
    Ethernet0/3
    Serial1/0
    Serial1/1
    Serial1/2
    Serial1/3
    RG-AR-IF-INPUT1
    VoIP-Null0
  Routing Information Sources:
    Gateway         Distance      Last Update
    10.1.1.10            120      00:00:01
    10.1.1.5             120      00:02:32
  Distance: (default is 120)
```

RIP is enabled for the correct network, but there are many interfaces set as passive, including the interfaces toward the R1 (Serial1/2) and R3 (Serial1/3) routers.

Step 8 On R2, disable the passive interface functionality for interfaces toward R1 (Serial1/2) and R3 (Serial1/3) routers.

On R2, enter the following commands:

```
R2# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
R2(config)# router rip
R2(config-router)# no passive-interface Serial 1/2
R2(config-router)# no passive-interface Serial 1/3
R2(config-router)# end
R2#
```

Step 9 On R3, verify the RIP database again.

On R3, enter the following command:


```
R3# show ip rip database
10.0.0.0/8    auto-summary
10.1.1.0/30    directly connected, Serial1/1
10.1.1.4/30    directly connected, Serial1/3
10.1.1.8/30
    [1] via 10.1.1.6, 00:00:53, Serial1/3
    [1] via 10.1.1.2, 00:00:00, Serial1/1
10.10.1.0/24
    [1] via 10.1.1.2, 00:00:00, Serial1/1
10.10.2.0/24
    [1] via 10.1.1.6, 00:00:53, Serial1/3
10.10.3.0/24    directly connected, Ethernet0/0
```

You can see that there is a subnet coming from the R2 router (10.10.2.0/24). It may take some time for the R2 router to send a RIP update.

To speed up convergence, clear IP routing table on R3.

```
R3# clear ip route *
R3#
```

Step 10 On R1, verify the IP routing table again.

On R1, enter the following command:

```
R1# sh 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
       + - replicated route, % - next hop override

Gateway of last resort is not set

      10.0.0.0/8 is variably subnetted, 6 subnets, 3 masks
C        10.1.1.0/30 is directly connected, Serial1/1
L        10.1.1.2/32 is directly connected, Serial1/1
C        10.1.1.8/30 is directly connected, Serial1/2
L        10.1.1.10/32 is directly connected, Serial1/2
C        10.10.1.0/24 is directly connected, Ethernet0/0
L        10.10.1.1/32 is directly connected, Ethernet0/0
```

There are still no RIP routes on the R1 router.

Step 11 On R1, verify RIP timers

On R1, enter the following command:

```
R1# show ip protocols
*** IP Routing is NSF aware ***

Routing Protocol is "rip"
  Outgoing update filter list for all interfaces is not set
  Incoming update filter list for all interfaces is not set
  Sending updates every 3 seconds, next due in 2 seconds
  Invalid after 5 seconds, hold down 5, flushed after 10
  Redistributing: rip
  Default version control: send version 2, receive version 2
    Interface             Send  Recv  Triggered RIP  Key-chain
    Serial1/1             2     2
    Serial1/2             2     2
  Automatic network summarization is in effect
  Maximum path: 4
  Routing for Networks:
    10.0.0.0
  Passive Interface(s):
    Ethernet0/0
  Routing Information Sources:
    Gateway         Distance      Last Update
    10.1.1.9             120      00:00:26
    10.1.1.1             120      00:00:42
  Distance: (default is 120)
```

The RIP timers on R1 are set to very low values. The update timer is 3 seconds, the invalid and hold timers are 5 seconds, and the flush timer is 10 seconds.

Step 12 Compare the R1 RIP timers with the RIP timers on the R2 and R3 routers.

On R2, enter the following command:


```
R2# sh ip protocols
*** IP Routing is NSF aware ***

Routing Protocol is "rip"
  Outgoing update filter list for all interfaces is not set
  Incoming update filter list for all interfaces is not set
  Sending updates every 300 seconds, next due in 26 seconds
  Invalid after 180 seconds, hold down 180, flushed after 240
  Redistributing: rip
  Default version control: send version 2, receive version 2
    Interface             Send  Recv  Triggered RIP  Key-chain
    Serial1/2             2     2
    Serial1/3             2     2
  Automatic network summarization is in effect
  Maximum path: 4
  Routing for Networks:
    10.0.0.0
  Passive Interface(s):
    Ethernet0/0
    Ethernet0/1
    Ethernet0/2
    Ethernet0/3
    Serial1/0
    Serial1/1
  Passive Interface(s):
    RG-AR-IF-INPUT1
    VoIP-Null0
  Routing Information Sources:
    Gateway         Distance      Last Update
    10.1.1.10            120      00:00:00
    10.1.1.5             120      00:01:20
  Distance: (default is 120)
```

On R3, enter the following command:

```
R3# sh ip protocols
*** IP Routing is NSF aware ***

Routing Protocol is "rip"
  Outgoing update filter list for all interfaces is not set
  Incoming update filter list for all interfaces is not set
  Sending updates every 300 seconds, next due in 112 seconds
  Invalid after 180 seconds, hold down 180, flushed after 240
  Redistributing: rip
  Default version control: send version 2, receive version 2
    Interface             Send  Recv  Triggered RIP  Key-chain
    Serial1/1             2     2
    Serial1/3             2     2
  Automatic network summarization is in effect
  Maximum path: 4
  Routing for Networks:
    10.0.0.0
  Passive Interface(s):
    Ethernet0/0
  Routing Information Sources:
    Gateway         Distance      Last Update
    10.1.1.2             120      00:00:00
    10.1.1.6             120      00:00:20
  Distance: (default is 120)
```

RIP timers on the R2 and R3 routers are higher than RIP timers on the R1 router. Also, because the RIP update timer is higher than other RIP timers, it will make the RIP routes flap.

Step 13 On the R1, R2 and R3 routers, synchronize the RIP timers. Also, change the RIP update timer to be less than the other RIP timers.

On R1, enter the following commands:

```
R1# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
R1(config)# router rip
R1(config-router)# timers basic 100 180 180 240
R1(config-router)# interface Serial1/1
R1(config-if)# no ip rip advertise 3
R1(config-if)# interface Serial1/2
R1(config-if)# no ip rip advertise 3
R1(config-router)# end
R1#
```

To change RIP update timer, you will need to adjust **ip rip advertise** command under interface configuration mode.

On R2, enter the following commands:

```
R2# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
R2(config)# router rip
R2(config-router)# timers basic 100 180 180 240
R2(config-router)# interface Serial1/2
R2(config-if)# no ip rip advertise 300
R2(config-if)# interface Serial1/3
R2(config-if)# no ip rip advertise 300
R2(config-router)# end
R2#
```

On R3, enter the following commands:

```
R3# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
R3(config)# router rip
R3(config-router)# timers basic 100 180 180 240
R3(config-router)# interface Serial1/1
R3(config-if)# no ip rip advertise 300
R3(config-if)# interface Serial1/3
R3(config-if)# no ip rip advertise 300
R3(config-router)# end
R3#
```

Step 14 On R1, verify the IP routing table again.

On R1, enter the following command:


```
R1# sh 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
       + - replicated route, % - next hop override

Gateway of last resort is not set

      10.0.0.0/8 is variably subnetted, 9 subnets, 3 masks
C        10.1.1.0/30 is directly connected, Serial1/1
L        10.1.1.2/32 is directly connected, Serial1/1
R        10.1.1.4/30 [120/1] via 10.1.1.9, 00:00:22, Serial1/2
                     [120/1] via 10.1.1.1, 00:02:30, Serial1/1
C        10.1.1.8/30 is directly connected, Serial1/2
L        10.1.1.10/32 is directly connected, Serial1/2
C        10.10.1.0/24 is directly connected, Ethernet0/0
L        10.10.1.1/32 is directly connected, Ethernet0/0
R        10.10.2.0/24 [120/1] via 10.1.1.9, 00:00:22, Serial1/2
R        10.10.3.0/24 [120/1] via 10.1.1.1, 00:02:30, Serial1/1
```

It may take up to 100 seconds for R1 to receive RIP routes from the R2 and R3 routers. It may be much more suitable to configure all routers with lower RIP timers. Feel free to continue with independent adjustments of the RIP timers within the lab environment.

Step 15 From the R1 router, verify connectivity to the R3 router Ethernet0/0 interface IP address (10.10.3.1) again.

On R1, enter the following command:

```
R1# ping 10.10.3.1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.10.3.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 8/8/9 ms
```

The ping is successful.

This is the end of the discovery lab.

Challenge

1. In which of the following networks would it be a bad idea to implement a dynamic routing protocol? (Choose two.)

   A. A High-speed Data Center Core Network
   B. A Remote Network with low-bandwidth connections
   C. A Network with a mix of Cisco and Non-Cisco Equipment
   D. A Large Network with several hundred subnets

2. Routers that are configured for distance vector routing protocol such as RIPv2 have an actual map of the network topology. True or False?

   A. True
   B. False

3. The default administrative distance for RIPv2 is which of the following?

   A. 80
   B. 100
   C. 120
   D. 140

4. Analyze the following configuration:

```
router rip
version 2
network 192.168.1.0
no auto
```

   Which of the following statements is true?

   A. The **router rip statement** is incorrect as you need to specify the instance number e.g. **router rip 1**
   B. The **version 2** statement is not required, as RIP defaults to version 2 when an instance is created.
   C. The **network 192.168.1.1** statement is incorrect, as you need to specify the subnet mask. e.g. **network 192.168.1.0 255.255.255.0**
   D. The **network 192.168.1.0** statement is incorrect, as you need to specify the subnet mask using slash format. e.g. network 192.168.1.0/24
   E. There is nothing wrong with this configuration. it would work perfectly.

5. While troubleshooting RIP, which command would you use to show the RIP table database?

   A. **show ip rip topology**
   B. **show ip rip traffic**
   C. **show ip rip route**
   D. **show ip rip database**

6. By default, how many paths are supported for load balancing traffic by RIP routing protocol?

   A. 2
   B. 4
   C. 6
   D. 8

7. Analyze the following configuration:

```
R1(config)# router rip
R1(config-router)# timers basic 60 180 180 240
```

   What timer is assigned the value of 180?

   A. update
   B. invalid
   C. holddown
   D. flush

Answer Key

Challenge

1. B
2. B
3. C
4. E
5. D
6. B
7. C

Module 5: Introducing IPv6

Introduction

This module starts with an explanation of why IPv6 is needed and how it is different from IPv4. To understand IPv6 operation, its header, neighbor discovery, and stateless autoconfiguration are discussed. An overview of different routing options for IPv6 is done, and static IPv6 routing is explained.

Lesson 1: Introducing Basic IPv6

Introduction

Your boss sends you to your customer to explain the limitations of IPv4 and propose workaround solutions for these problems. You will need to introduce IPv6 as a more permanent solution to the customer.

IPv6 satisfies the increasingly complex requirements of hierarchical addressing that IPv4 does not satisfy. With a 128-bit address length, the IPv6 address space is significantly larger and more diverse, and is therefore more complicated to manage.

IPv4 Addressing Exhaustion Workarounds

IPv4 provides approximately four billion unique IP addresses. While that is many addresses, it is not enough to keep up with the growth of the Internet.

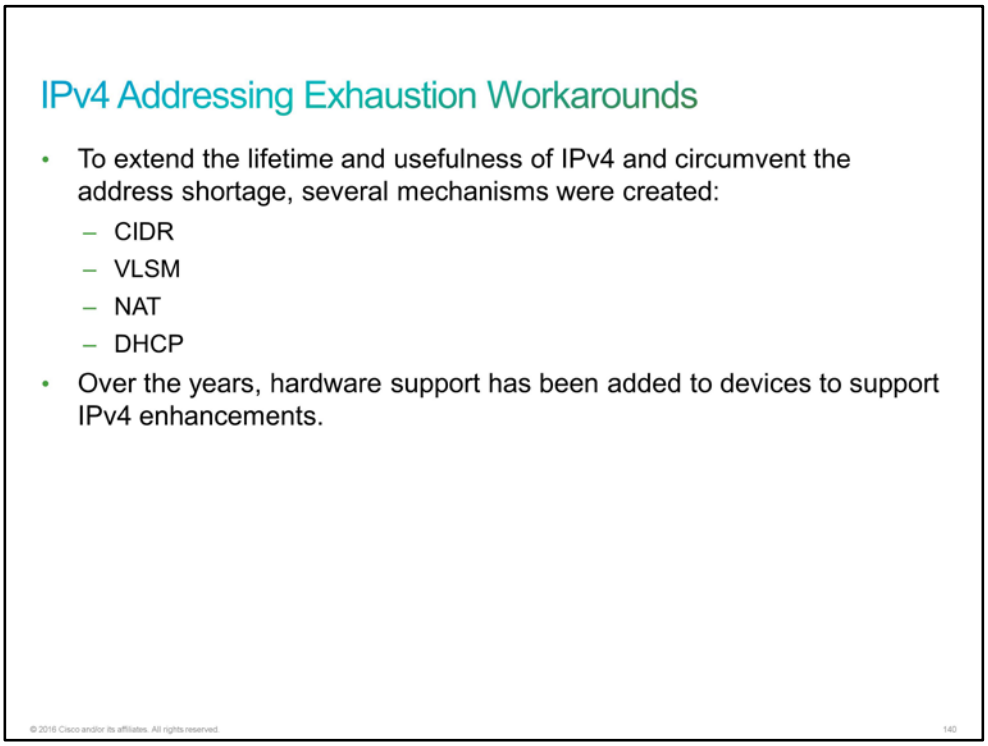

To allocate IPv4 addresses efficiently, CIDR was developed. CIDR allows the address space to be divided into smaller blocks.

VLSMs allow more efficient use of IP addresses, specifically on small segments such as point-to-point serial links. VLSM usage was recommended in RFC 1817. CIDR and VLSM support is a prerequisite for ISPs to receive more allocations.

NAT introduced a model in which a device facing outward to the Internet has a globally routable IPv4 address, while the internal network is configured with private addresses. These private addresses can never leave the site, so they can be identical in many different enterprise networks. In this way, even large enterprises with thousands of systems can hide behind a few routable public networks.

The DHCP is used by a client to acquire configuration information, such as an IP address, a default route, and DNS setup from a server.

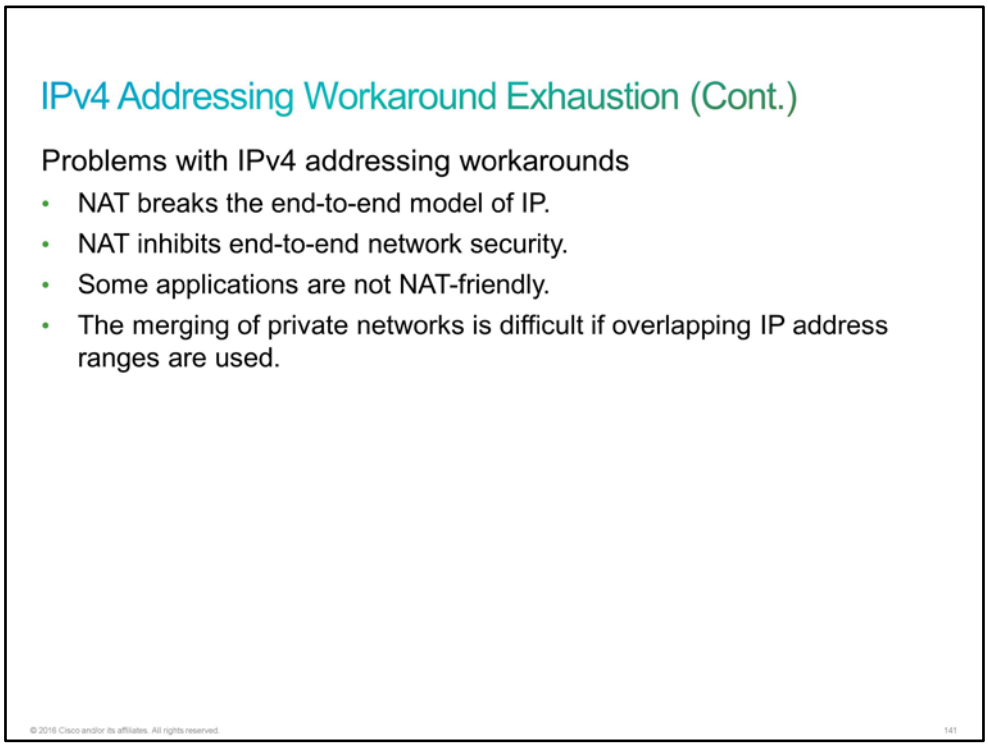

One of the arguments against deploying IPv6 is that NAT will solve the problems of limited address space in IPv4. The use of NAT merely delays the exhaustion of the IPv4 address space by using global addresses for large internal networks.

There are several negative implications of using NAT, some of which are identified in RFC 2775 and RFC 2993, as follows:

- NAT breaks the end-to-end model of IP. IP was defined so that the underlying layers do not process the connection; only the endpoints process the connection.
- NAT inhibits end-to-end network security. To protect the integrity of the IP header by some cryptographic functions, the IP header cannot be changed between the origin of the packet (to protect the integrity of the header) and the final destination (to check the integrity of the received packet). Any translation of parts of a header on the path will break the integrity check.
- When applications are not NAT-friendly—which means that, for a specific application, more than just the port and address mapping are necessary to forward the packet through the NAT device—NAT has to embed complete knowledge of all the applications to perform correctly. This fact is especially true for dynamically allocated ports with rendezvous ports, embedded IP addresses in application protocols, security associations, and so on. Therefore, the NAT device needs to be upgraded each time that a new non-NAT-friendly application is deployed (for example, peer-to-peer).
- When different networks use the same private address space and they have to merge or connect, there is an address-space collision. Hosts that are different but have the same address cannot communicate with each other. This problem can be resolved by techniques such as renumbering or Twice NAT. Twice NAT is the practice of changing both the source and destination address of a packet. However, these techniques are costly and, later on, increase NAT complications.

IPv6 Features

Although VLSM, NAT, and other workarounds (for avoiding the transition to IPv6) are available, networks with Internet connectivity must begin the transition to IPv6 as soon as possible. For IPv4 networks that provide goods and services to Internet users, it is especially important because the transition by the Internet community is already under way. New networks may be unable to acquire IPv4 addresses, and networks running IPv6 exclusively will not be able to communicate with IPv4-only networks unless you configure an intermediary gateway or another transition mechanism. IPv6 and IPv4 are completely separate protocols, and IPv6 is not backward-compatible with IPv4. As the Internet evolves, organizations must adopt IPv6 to support future business continuity, growth, and global expansion.

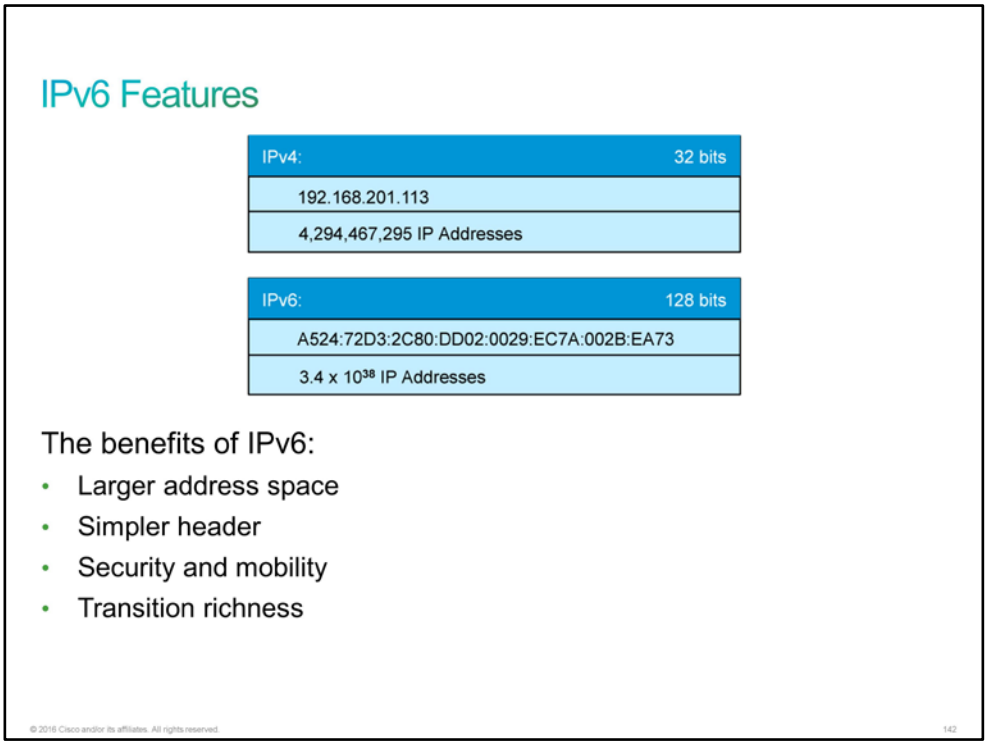

IPv6 includes several features that make it attractive for building global-scale, highly effective networks.

- **Larger address space:** The expanded address space includes several IP addressing enhancements:
  - It provides improved global reachability and flexibility.
  - A better aggregation of IP prefixes is announced in the routing tables. The aggregation of routing prefixes limits the number of routing table entries, which creates efficient and scalable routing tables.
  - Multihoming increases the reliability of the Internet connection of an IP network. With IPv6, a host can have multiple IP addresses over one physical upstream link. For example, a host can connect to several ISPs.
  - Autoconfiguration is enabled.
  - There are more "plug-and-play" options (which will work as soon as they are active in the network) for more devices.
  - Simplified mechanisms are available for address renumbering and modification.

- **Simpler header:** Streamlined header structures make the processing of IPv6 packets faster and more efficient for intermediate routers within the network. This fact is especially true when large numbers of packets are routed in the core of the IPv6 Internet.
- **Security and mobility:** Features that were not part of the original IPv4 specification, such as security and mobility, are now built into IPv6. IPsec is mandatory in IPv6, making the IPv6 Internet more secure. Mobility enables mobile network devices to move around in networks without breaks in established network connections.
- **Transition richness:** IPv6 also includes a rich set of transition tools, such as "dual stacking," to allow an easy, nondisruptive transition over time to IPv6-dominant networks.

IPv6 Addresses

IPv6 addresses are represented as a series of eight 16-bit hexadecimal fields that are separated by colons.

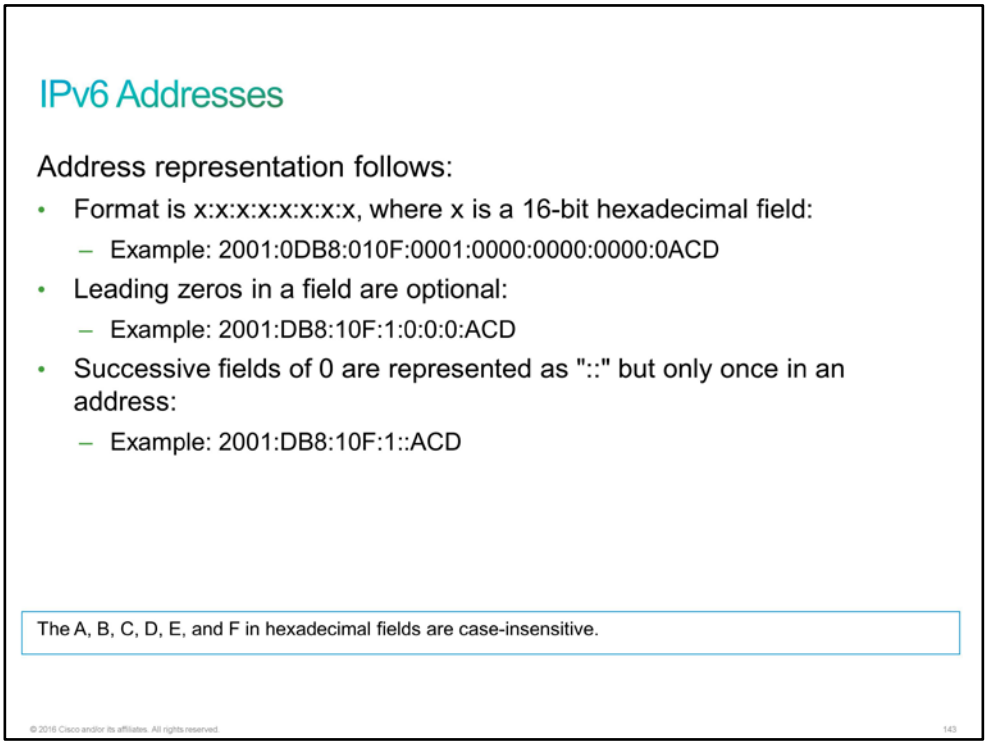

The following are two ways to shorten the writing of IPv6 addresses:

- The leading zeros in a field are optional, so 010F can be written as 10F. A field that contains all zeros (0000) can be written as 0.
- Successive fields of zeros can be represented as a double colon (::) but only once in an address. An address parser can identify the number of missing zeros by separating the two parts and filling in zeros until the 128 bits are completed. However, if two double colons are placed in the address, there is no way to identify the size of each block of zeros. Therefore, only one double colon is possible in a valid IPv6 address.

The use of the double-colon technique makes many addresses very small; for example, FF01:0:0:0:0:0:0:1 becomes FF01::1. The unspecified address is written as a double colon because it contains only zeros.

IPv6 Address Types

IPv6 supports three basic types of addresses. Each address type has specific rules regarding its construction and use.

IPv6 supports three types of addresses:

- **Unicast:** Unicast addresses are used in a one-to-one context.
- **Multicast:** A multicast address identifies a group of interfaces. Traffic that is sent to a multicast address is sent to multiple destinations at the same time. An interface may belong to any number of multicast groups.
- **Anycast:** An IPv6 anycast address is assigned to an interface on more than one node. When a packet is sent to an anycast address, it is routed to the nearest interface that has this address. The nearest interface is found according to the measure of distance of the particular routing protocol. All nodes that share the same address should behave the same way so that the service is offered similarly regardless of the node that services the request.

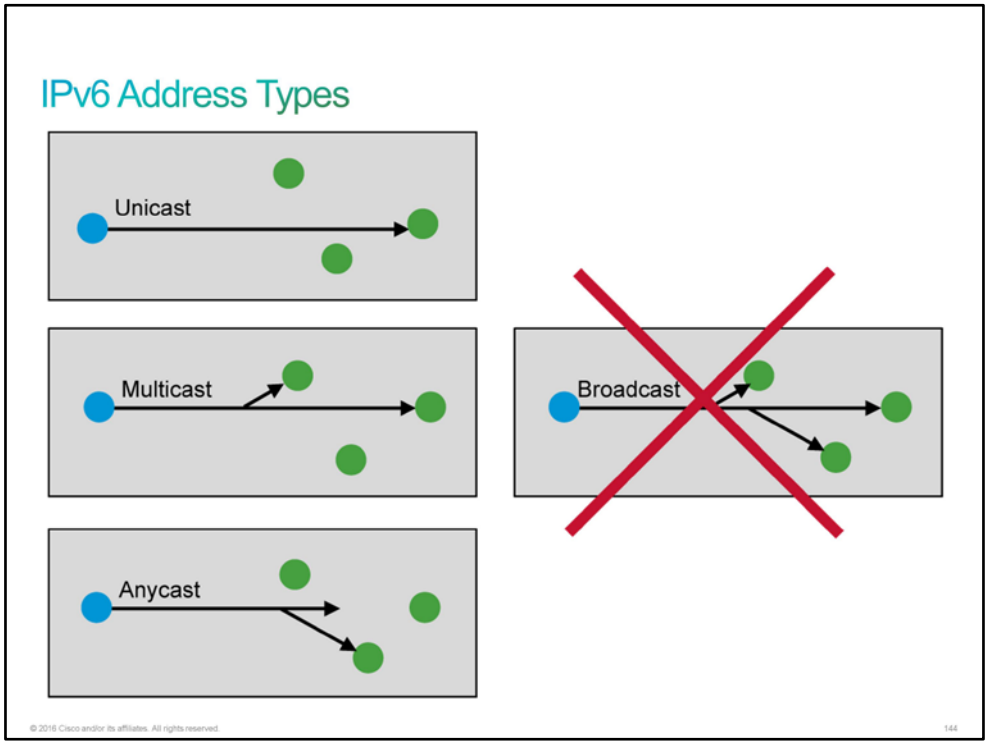

IPv6 has no support for broadcast addresses in the way that they are used in IPv4. Instead, specific multicast addresses (such as the all-nodes multicast address) are used.

IPv6 Address Scopes and Prefixes

To fully understand IPv6 addressing, it is important to have a solid understanding of IPv6 scopes and prefixes. An IPv6 address scope specifies the region of the network in which the address is valid. For example, the loopback address has a scope that is called "link-local," which means that it should be used in a directly attached network (link). There are three different scopes or regions: the link scope, the site network scope, and the global network scope.

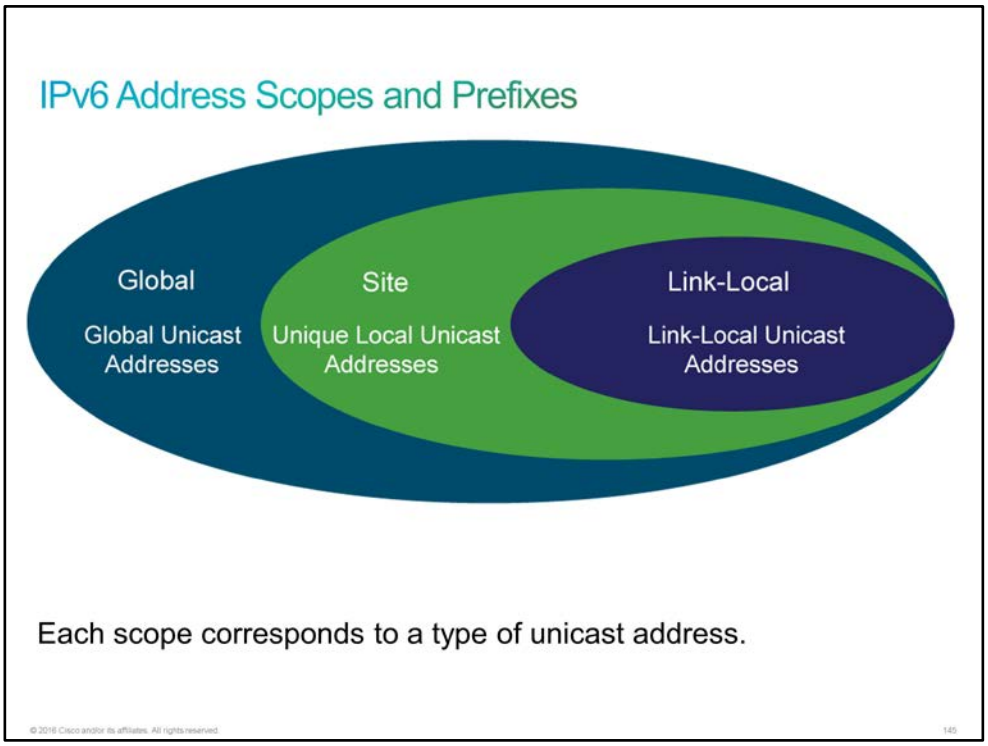

Addresses in the link scope are called link-local unicast addresses, addresses in the site network scope are called unique local unicast addresses, and addresses in the global network scope are called global unicast addresses.

IPv6 Unicast Addresses

An IPv6 unicast address generally uses 64 bits for the network ID and 64 bits for the host ID. The network ID is administratively assigned, and the host ID can be configured manually or autoconfigured.

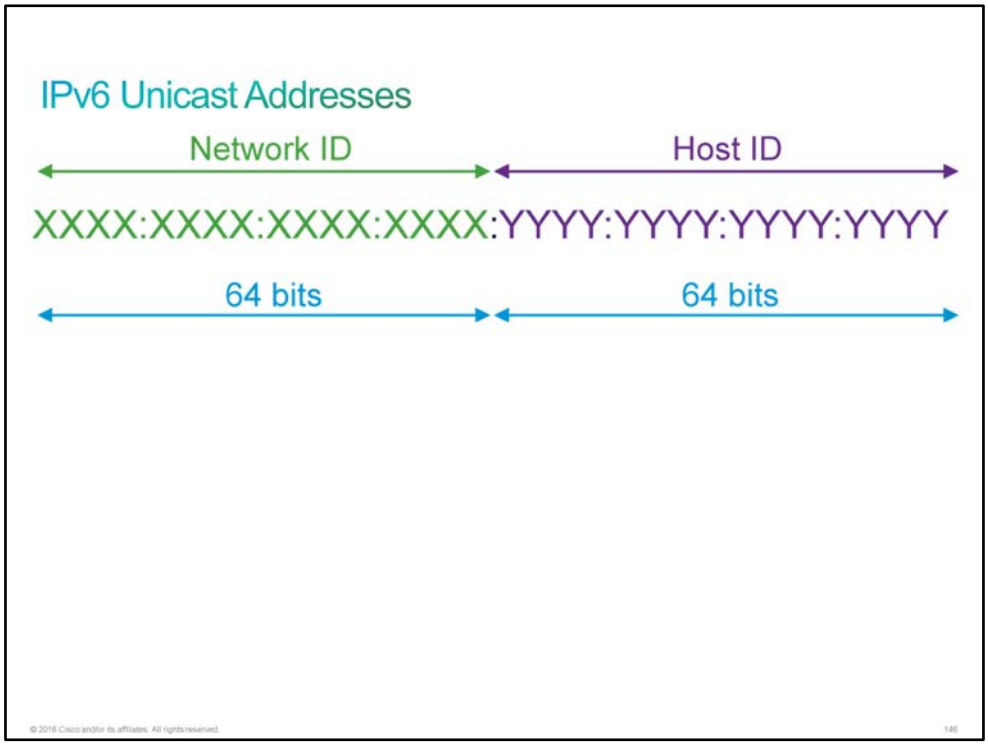

IPv6 Global Unicast Address

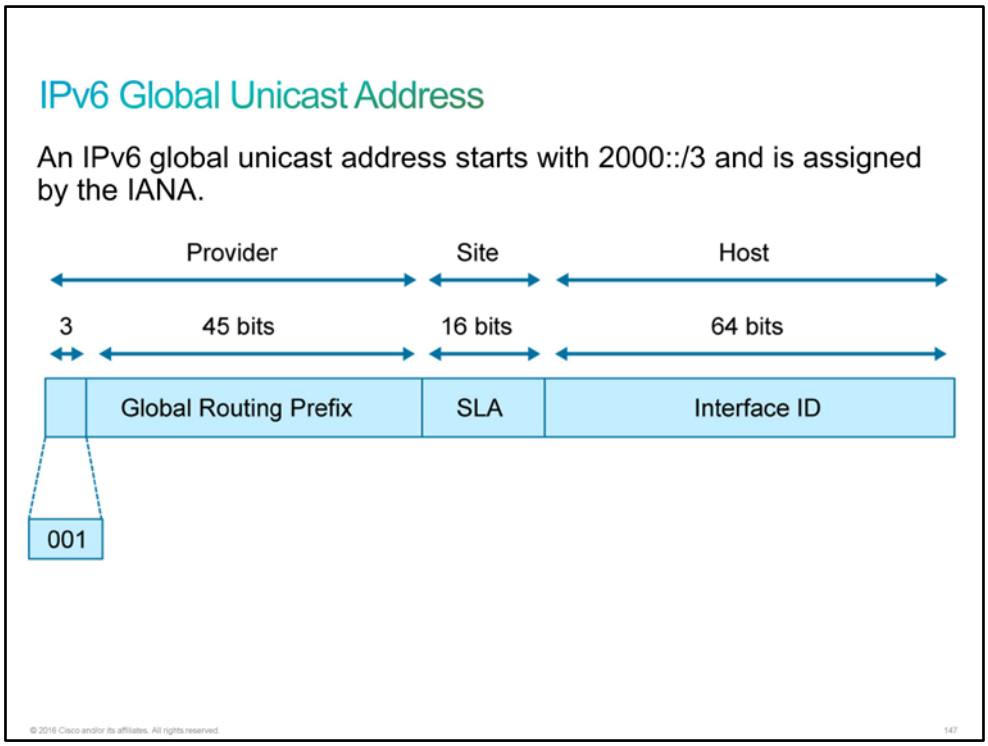

Global unicast addresses are routable and reachable across the Internet. They are intended for widespread generic use. A global unicast address is structured hierarchically to allow address aggregation, and it can be identified by the fact that its three high-level bits are set to 001 (2000::/3).

The global routing prefix is assigned to a service provider by the IANA. An SLA, or subnet ID, is assigned to a customer by its service provider. A subnet ID can be used by an individual organization to create its own local addressing hierarchy and to identify subnets. A subnet ID is similar to a subnet in IPv4, except that an organization with an IPv6 subnet ID can support up to 65,535 individual subnets. The interface ID (or host ID) has the same meaning for all unicast addresses. It is used to identify the interfaces that are on a link and that must be unique to the link. The interface ID is 64 bits long and is typically created by using the EUI-64 format. An example of a global unicast address is 2001:0DB8:BBBB:CCCC:0987:65FF:FE01:2345.

Note	An IPv6 address prefix, in the format *ipv6-prefix/prefix-length*, can be used to represent bitwise contiguous blocks of the entire address space. The prefix length is a decimal value that indicates how many of the high-order contiguous bits of the address compose the prefix. An IPv6 address network prefix is represented in the same way as the network prefix (as in 10.1.1.0/24) in IPv4. For example, 2001:DB8:8086:6502::/32 is a valid IPv6 prefix.

IPv6 Unique Local Unicast Address

Unique local unicast addresses are analogous to private IPv4 addresses in that they are used for local communications, intersite VPNs, and so on. They are not routable on the Internet without IPv6 NAT, but they are routable inside a limited area, such as a site. They may also be routed between a limited set of sites. A unique local unicast address has the following characteristics:

- It has a globally unique prefix (that is, it has a high probability of uniqueness).
- It has a well-known prefix to enable easy filtering at site boundaries.
- It allows combining or privately interconnecting sites without creating any address conflicts or requiring a renumbering of interfaces that use these prefixes.
- It is ISP-independent and can be used for communications inside a site without having any permanent or intermittent Internet connectivity.
- If it is accidentally leaked outside of a site via routing or the DNS, there is no conflict with any other addresses.
- Applications may treat unique local addresses like global scoped addresses.

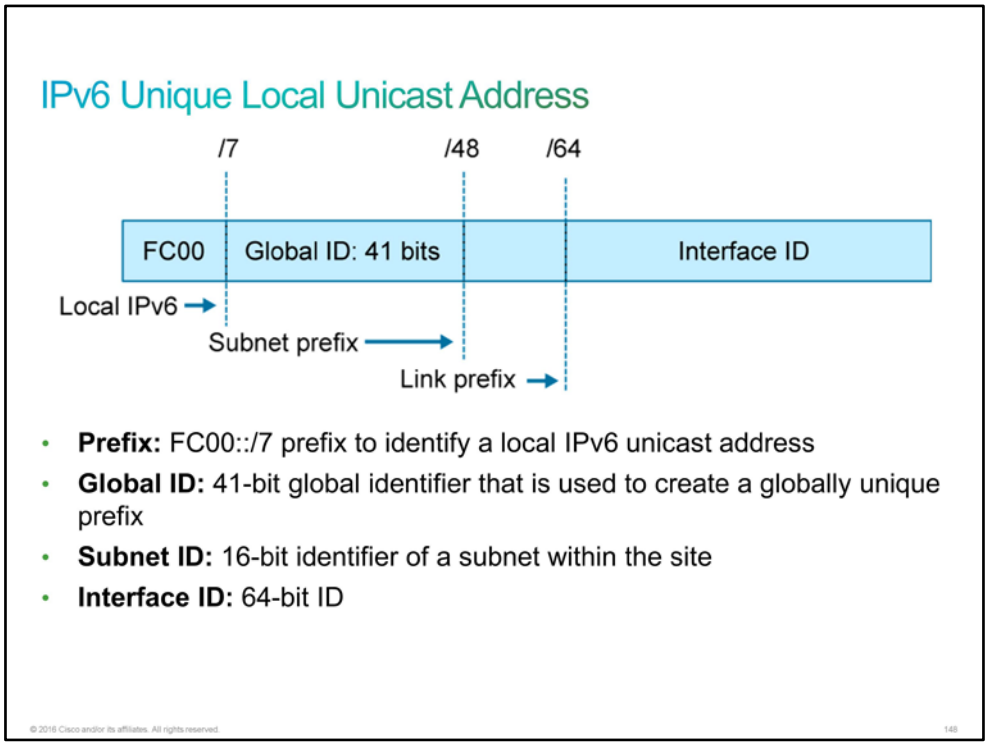

In unique local unicast addresses, global IDs are defined by the administrator of the local domain. Subnet IDs are also defined by the administrator of the local domain. Subnet IDs are typically defined by using a hierarchical addressing plan to enable route summarization. An example of a unique local unicast address is FD00:aaaa:bbbb:CCCC:0987:65FF:FE01:2345.

IPv6 Link Local Unicast Address

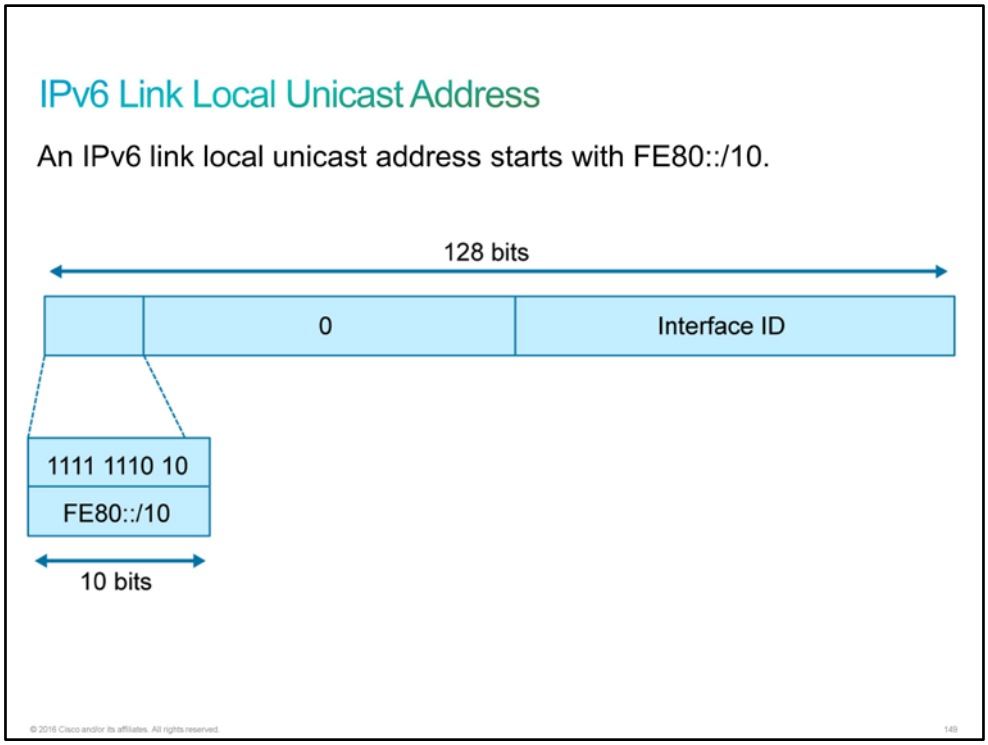

A link-local address is an IPv6 unicast address that is automatically configured on any interface that is using the link-local prefix FE80::/10 (1111 1110 10) and the interface identifier in the modified EUI-64 format. Each interface on an IPv6 node must have a link-local address. Link-local addresses are used in the neighbor discovery protocol and the stateless autoconfiguration process. Nodes on a local link can use link-local addresses to communicate; the nodes do not need globally unique addresses to communicate.

An example of a link-local unicast address is FE80:0000:0000:0000:0987:65FF:FE01:2345, which would generally be represented in shorthand notation as FE80::987:65FF:FE01:2345.

Note	In addition to global, link-local, and unique local unicast addresses, Cisco IOS Software supports an IPv4-compatible IPv6 address. This is an IPv6 unicast address that has zeros in the high-order 96 bits of the address and an IPv4 address in the low-order 32 bits of the address. The format of an IPv4-compatible IPv6 address is 0:0:0:0:0:0:A.B.C.D or ::A.B.C.D. The entire 128-bit IPv4-compatible IPv6 address is used as the IPv6 address of a node, and the IPv4 address that is embedded in the low-order 32 bits is used as the IPv4 address of the node. IPv4-compatible IPv6 addresses are assigned to nodes that support both the IPv4 and IPv6 protocol stacks and are used in automatic tunnels.

IPv6 Multicast Addresses

The next figure illustrates the format of an IPv6 multicast address.

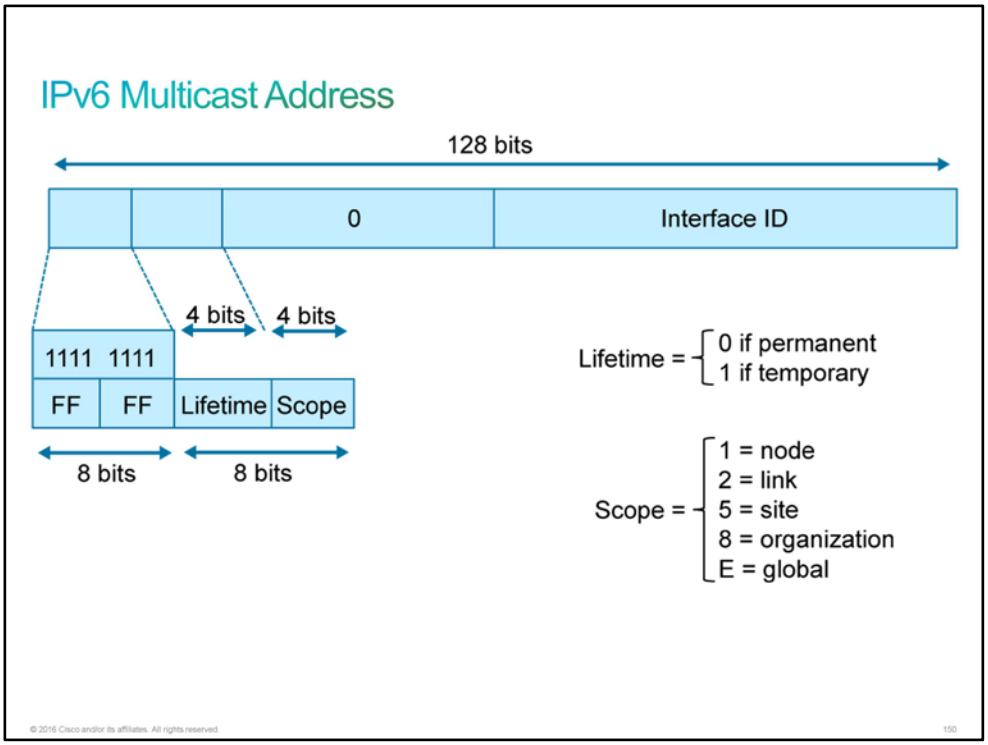

An IPv6 multicast address has an 8-bit prefix, FF00::/8 (1111 1111). The second octet following the prefix defines the lifetime and scope of the multicast address. IPv6 nodes are required to join (receive packets that are destined for) certain multicast groups.

IPv6 Anycast Addresses

Anycast addresses are syntactically indistinguishable from unicast addresses, because anycast addresses are allocated from the unicast address space. Assigning a unicast address to more than one interface makes a unicast address an anycast address. The nodes to which the anycast address is assigned must be explicitly configured to recognize that the address is an anycast address.

Anycast addresses can only be used by devices, not hosts, and an anycast address must not be used as the source address of an IPv6 packet.

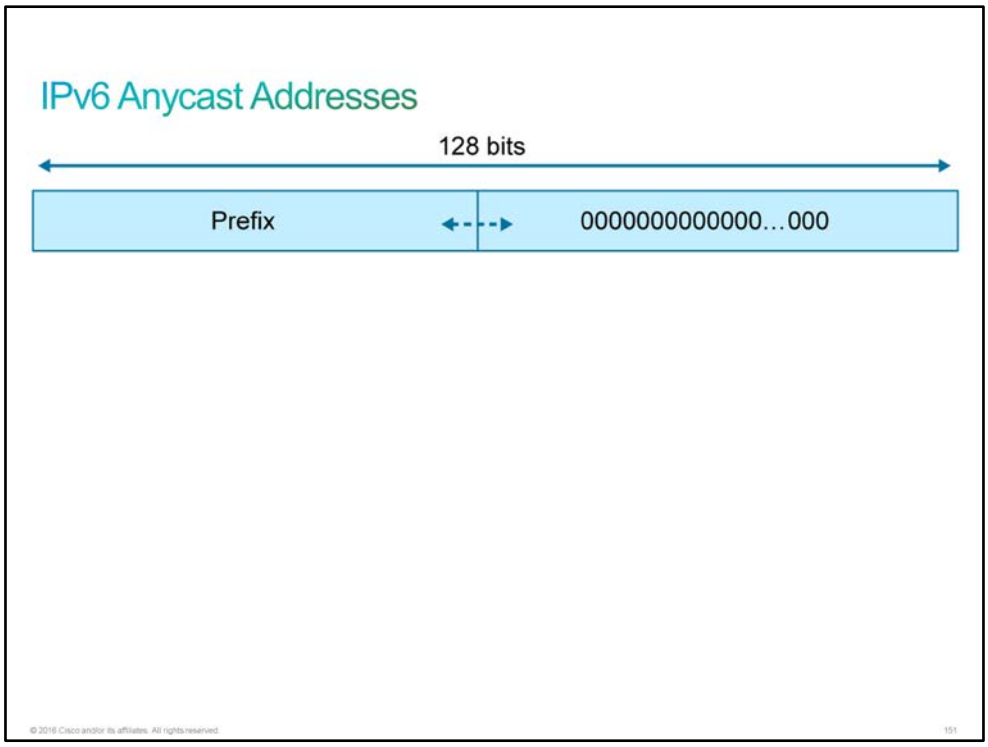

The figure shows the format of the anycast address of the subnet device; the address has a prefix connected by a series of zeros (the interface ID). The anycast address of the subnet device can be used to reach a device on the link that is identified by the prefix in the anycast address.

As with IPv4, IPv6 addresses are assigned to interfaces; however, unlike IPv4, an IPv6 interface is expected to have multiple addresses. The IPv6 addresses assigned to an interface can be any of the basic types: unicast, multicast, or anycast.

IPv6 Address Allocation

IPv6 addresses can be configured manually or autoconfigured.

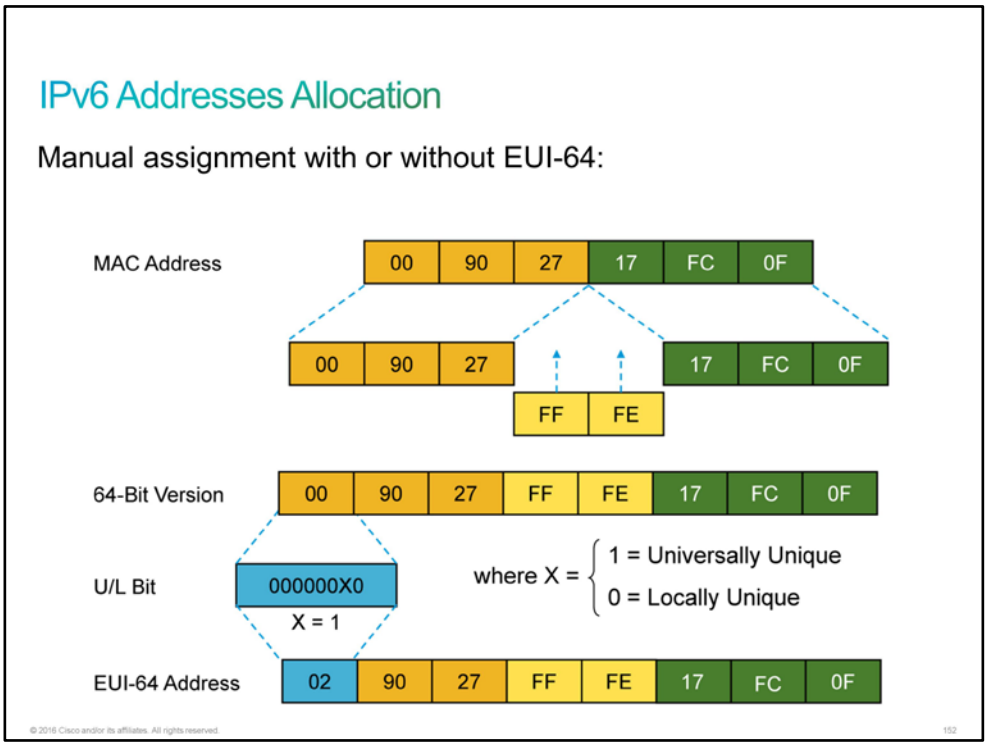

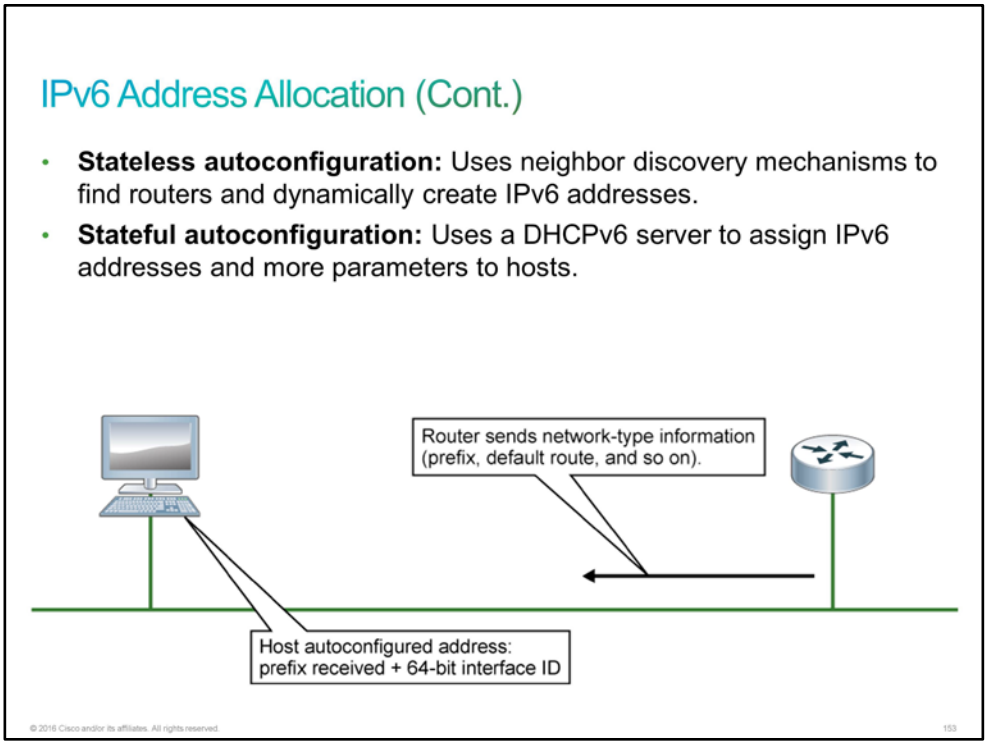

There are three methods for assigning IPv6 addresses:

- **Manual:** The entire IPv6 address can be manually configured, or the host identifier (the rightmost 64 bits) can be computed from the EUI-64 of the interface. The EUI-64 format expands the 48-bit MAC address of the device interface to 64 bits by inserting FFFE into the middle 16 bits, as shown in the figure. The EUI-64 host ID format is commonly used for Cisco IP Phones, gateways, and routers.
- **Autoconfiguration:** Autoconfiguration enables "plug and play," which connects devices to the network without any configuration and without any stateful servers (such as DHCP servers). Autoconfiguration is an important feature for enabling deployment of new devices on the Internet, such as cell phones, wireless devices, home appliances, networks, and so on. Autoconfiguration can be accomplished in two ways:
  - **Stateless autoconfiguration**: Stateless autoconfiguration uses neighbor discovery mechanisms to find routers and dynamically create IPv6 addresses. To use this method for an IPv6 node, it is important to connect the IPv6 node to a network that uses at least one IPv6 router. The router transmits router advertisements to the link. These announcements can allow the on-link connected IPv6 nodes to configure themselves with an IPv6 address and routing parameters, as specified in RFC 2462, without further human intervention. A node on the link can automatically configure its global IPv6 address by appending its interface identifier (64 bits) to the prefix (64 bits) that is included in the router advertisement messages. Stateless autoconfiguration enables "plug and play," which connects devices to the network without any configuration and without any stateful servers (such as DHCP servers). It is an important feature for enabling the deployment of new devices on the Internet, such as cell phones, wireless devices, home appliances, and networks.
  - **Stateful autoconfiguration**: Stateful autoconfiguration uses a DHCPv6 server to assign IPv6 addresses and additional parameters to hosts. Stateful autoconfiguration keeps a record of which addresses are assigned to which hosts, while the stateless method maintains no such records.

A router announcement can indicate to hosts whether more configuration parameters are available via stateful configuration (DHCPv6), such as DNS, IP options, and so on.

DNS is a distributed Internet directory service that is used to translate between domain names and IP addresses and also between IP addresses and domain names. The DNS protocol had to be updated to support IPv6 in addition to IPv4. Using DDNS, DHCPv6 clients can dynamically update their records in DNS.

An IPv6 address must be configured on an interface before the interface can forward IPv6 traffic. Configuring a global unicast IPv6 address on an interface automatically configures a link-local address and activates IPv6 for that interface. Also, the configured interface automatically joins the following required multicast groups for that link:

- Solicited-node multicast group FF02:0:0:0:0:1:FF00::/104 (for each unicast and anycast address that is assigned to the interface)
- All-nodes link-local multicast group FF02::1
- All-routers link-local multicast group FF02::2

Challenge

1. Which four mechanisms are used to extend the lifetime and usefulness of IPv4? (Choose four.)

   A. CIDR
   B. VLSM
   C. NAT
   D. DHCP
   E. DNS
   F. IPv6

2. What are the four problems with IPv4 addressing workarounds? (Choose four.)

   A. NAT breaks the end-to-end model of IP.
   B. DHCP breaks the end-to-end model of IP.
   C. NAT inhibits end-to-end network security.
   D. VLSM inhibits end-to-end network security.
   E. Some applications are not NAT friendly.
   F. Some applications are NAT friendly.
   G. The merging of private networks is difficult if overlapping IP address ranges are used.

3. What are the four benefits of IPv6? (Choose four.)

   A. larger address space
   B. smaller address space
   C. complex header
   D. simpler header
   E. security and mobility
   F. transition richness
   G. TCP and UDP included in the basic header

4. Which address type is not supported in the IPv6?

   A. unicast
   B. multicast
   C. anycast
   D. broadcast

5. How does an IPv6 link local unicast address start?

   A. FE80::/10
   B. FC00::/7
   C. FC00::/8
   D. FD00::/8

6. How does an IPv6 multicast address start?

   A. FF00::/8
   B. FC00::/7
   C. FE80/10
   D. FD00::/8

7. How does the MAC address 00A1-6789-ABCD translate into the right-most 64 bits of the IPv6 address when EUI-64 is used?

   A. 01A1:67FF:FE89:ABCD
   B. 01A1:67EE:EE89:ABCD
   C. 02A1:67FF:FF89:ABCD
   D. 02A1:67FF:FE89:ABCD

Answer Key

Challenge

1. A, B, C, D
2. A, C, E, G
3. A, D, E, F
4. D
5. A
6. A
7. D

Lesson 2: Understanding IPv6 Operation

Introduction

Any device that attaches to a network goes through numerous processes to identify itself and to obtain services from the network. This premise is true in either an IPv4 or IPv6 network. However, people who design and manage IPv6 networks will discover that although the processes that are used in IPv6 have some similarities to those processes that are used in IPv4, they are different. Understanding these processes is fundamental to properly supporting an IPv6-enabled environment.

Your customer has decided to introduce IPv6 into their network. So, your boss sends you to the customer to configure basic IPv6 connectivity. But before you do so, they would like to know more about IPv6. You will need to explain to your customer the IPv6 header and compare it to the IPv4 header. You will need to understand ICMPv6 messages and the neighbor discovery process. You will also explain to your customer how stateless autoconfiguration helps to automatically assign IPv6 addresses to devices in the network.

Comparison of IPv4 and IPv6 Headers

The IPv6 header design significantly differs from the IPv4 header in several ways.

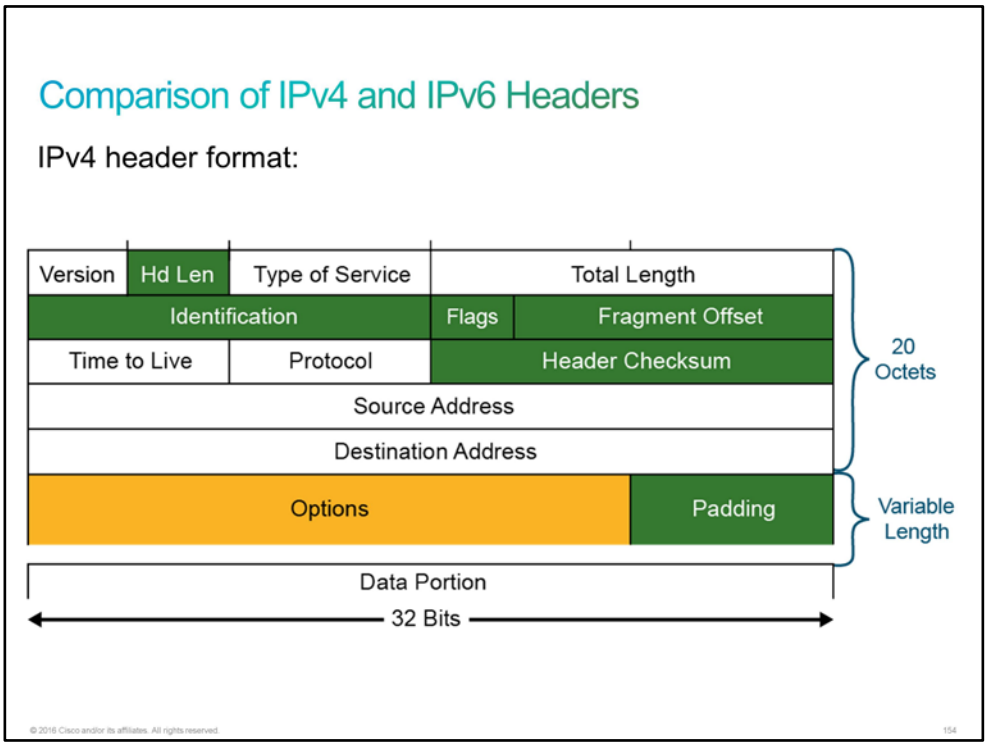

The IPv4 header contains 12 fields. Following these fields is an Options field of variable length that the figure shows in yellow and a data portion that is usually the transport layer segment. The basic IPv4 header has a size of 20 octets. The Options field increases the size of the IP header.

Of the 12 IPv4 header fields, six are removed in IPv6; these fields are shown in green and blue in the figure. The main reasons for removing these fields in IPv6 are as follows:

- The Internet Header Length field was removed because it is no longer required. Unlike the variable-length IPv4 header, the IPv6 header is fixed at 40 bytes.
- Fragmentation is processed differently in IPv6 and does not need the Flags field in the basic IPv4 header. In IPv6, routers no longer process fragmentation. IPv6 hosts are responsible for path MTU discovery. If the host needs to send data that exceed the path MTU, the host is responsible for fragmentation. The related Flags field option appears in the Fragmentation Extension Header in IPv6. This header is attached only to a packet that is fragmented.
- The Header Checksum field at the IP layer was removed because most data link layer technologies already perform checksum and error control. This change forces formerly optional upper-layer checksums (such as UDP) to become mandatory.

The Options field is not present in IPv6. Also, in IPv6, a chain of extension headers processes any additional services. Examples of extension headers include Fragmentation, AH, and ESP.

Most other fields were either unchanged or changed only slightly.

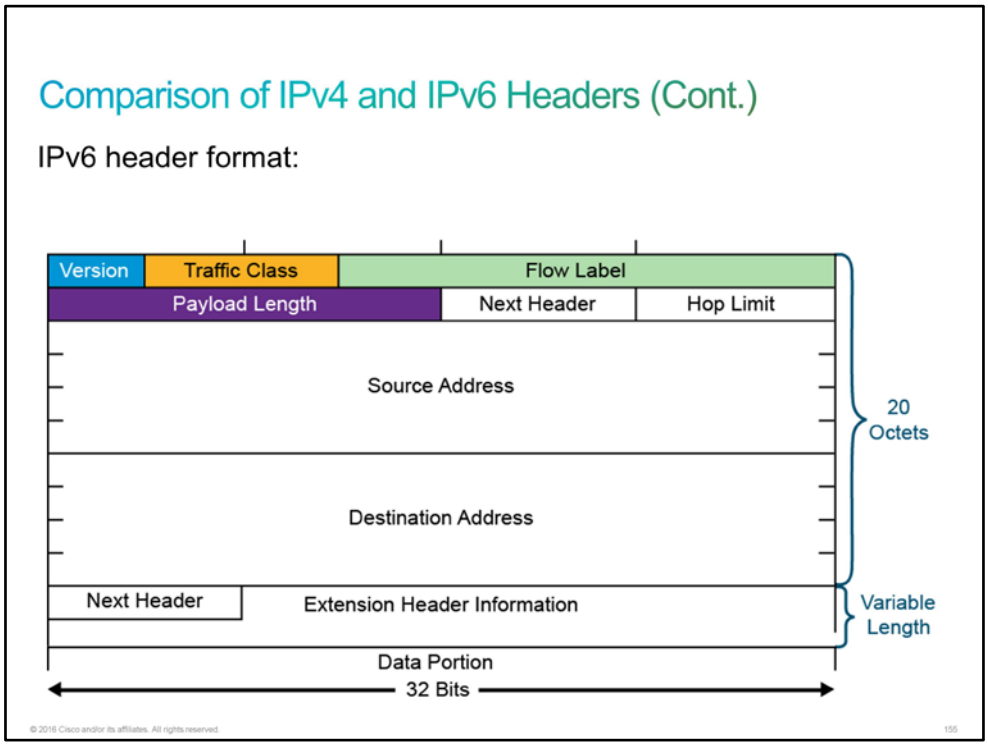

The IPv6 header has 40 octets, instead of 20 octets as in IPv4. The IPv6 header has fewer fields, and the header is aligned on 64-bit boundaries to enable fast processing by current and next-generation processors. The Source and Destination IP fields are four times larger than in IPv4.

Note	The Source and Destination IP fields are the most important headers to understand.

The IPv6 header contains eight fields:

1. **Version:** This 4-bit field contains the number 6, instead of the number 4 as in IPv4.
2. **Traffic Class:** This 8-bit field is similar to the ToS field in IPv4. The source node uses this field to mark the priority of outbound packets.
3. **Flow Label:** This new field has a length of 20 bits and is used to mark individual traffic flows with unique values. Routers are expected to apply an identical QoS treatment to each packet in a flow.
4. **Payload Length:** This field is like the Total Length field for IPv4, but because the IPv6 base header is a fixed size, this field describes the length of the payload only, not of the entire packet.
5. **Next Header:** The value of this field determines the type of information that follows the basic IPv6 header.
6. **Hop Limit:** This field specifies the maximum number of hops that an IP packet can take. This field begins at 255 and is decremented by each IPv6 router along the path to the destination. An IPv6 packet can pass through a maximum of 254 hops before it is deleted. The hop limit is designed to prevent packets from circulating forever if there is a routing error. In normal routing, this limit should never be reached.
7. **Source Address:** This field of 16 octets, or 128 bits, identifies the source of the packet.
8. **Destination Address:** This field of 16 octets, or 128 bits, identifies the destination of the packet.

The extension headers, if there are any, follow these eight fields. The number of extension headers is not fixed, so the total length of the extension header chain is variable.

For further exploration of IPv6 header fields and their functions, use your favorite search engine to locate RFC 2460, *Internet Protocol, Version 6 (IPv6) Specification*.

Internet Control Message Protocol Version 6

The ICMPv6 provides the same diagnostic services as ICMPv4, and it extends the functionality for some specific IPv6 functions that did not exist in IPv4.

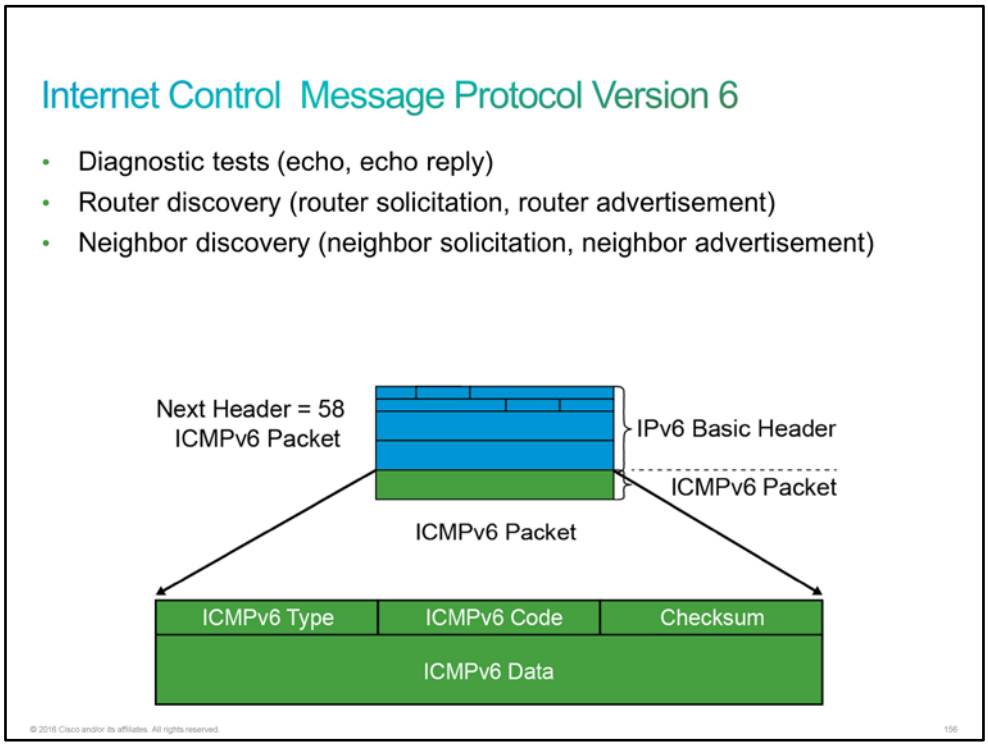

ICMPv6 is like ICMPv4. ICMPv6 enables nodes to make diagnostic tests and report problems. Like ICMPv4, ICMPv6 implements two kinds of messages—error messages, such as Destination Unreachable, Packet Too Big, or Time Exceeded, and informational messages, such as Echo Request and Echo Reply.

ICMPv6 Descriptions

ICMPv6 Type	Description
1	Destination Unreachable
128	Echo Request
129	Echo Reply
133	Router Solicitation
134	Router Advertisement
135	Neighbor Solicitation
136	Neighbor Advertisement

The ICMPv6 packet is identified as 58 in the Next Header field. Inside the ICMPv6 packet, the Type field identifies the type of ICMP message. The Code field further details the specifics of this type of message. The Data field contains information that is sent to the receiver for diagnostics or information purposes.

The Next Header field of the basic IPv6 header is set to a value of 58 to signal to the receiver that an ICMPv6 header and payload is being transported in this packet.

ICMPv6 is used on-link for router solicitation and advertisement, for neighbor solicitation and advertisement (acquisition of data link layer addresses for IPv6 neighbors), and for the redirection of nodes to the best gateway.

Neighbor Discovery

Neighbor discovery is used on-link for router solicitation and advertisement, for neighbor solicitation and advertisement, and for the redirection of nodes to the best gateway.

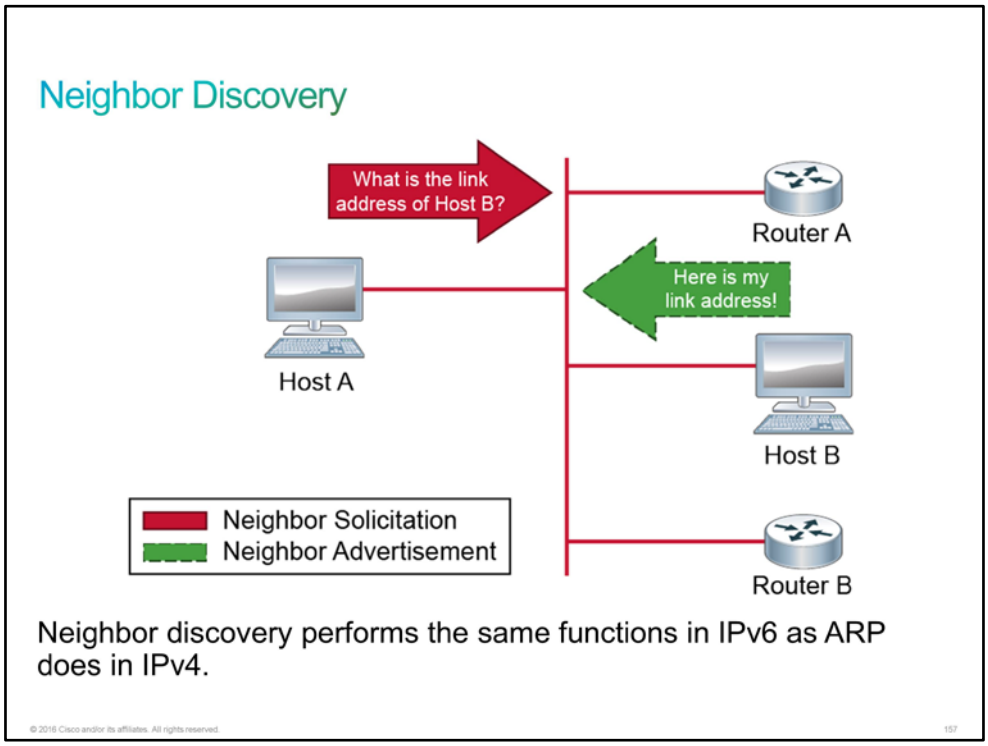

Neighbor discovery is a process that enables these functions:

- Determining the data link layer address of a neighbor on the same link, like ARP does in IPv4
- Finding neighbor routers on a link
- Keeping track of neighbors
- Querying for duplicate addresses

Neighbor discovery achieves these results by using ICMPv6 with IPv6 multicast addresses.

Solicited-Node Multicast Address

The solicited-node address is a multicast address. Any node must join the multicast group that corresponds to each of its unicast and anycast addresses. The solicited-node address is composed of the FF02:0:0:0:0:1:FF/104 prefix, which is concatenated with the right-most 24 bits of the corresponding unicast or anycast address.

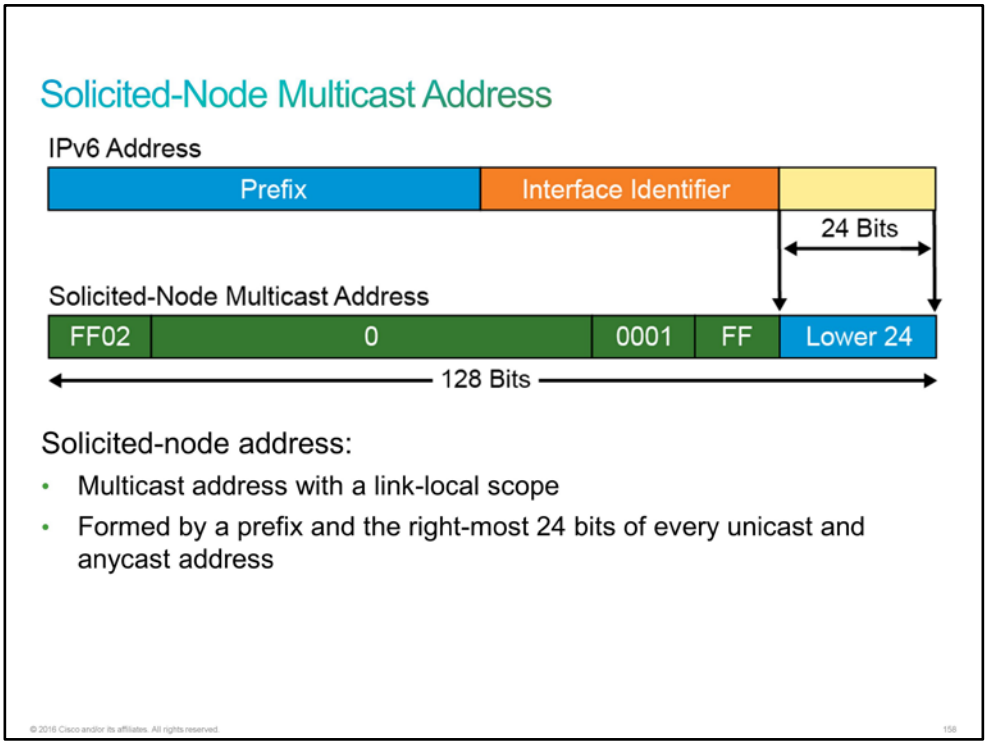

As an FF02::/16 address, a solicited-node multicast address has a link-local scope.

Solicited-node addresses are used for Neighbor Solicitation messages, when another node needs the data link layer address of an IPv6 address to send the right frame on the data link layer. The source node takes the right-most 24 bits of the IPv6 address of the destination node and sends a Neighbor Solicitation message to the multicast group on the link-local address. The corresponding node responds with its data link layer address.

This function avoids the broadcasts that IPv4 ARP uses. Here all nodes receive the requests.

Multicast Mapping over Ethernet

If an IPv6 address is known, then the associated IPv6 solicited-node multicast address is known. The example in the figure gives the IPv6 address 2001:DB8:1001:F:2C0:10FF:FE17:FC0F. The associated solicited-node address is FF02::1:FF17:FC0F.

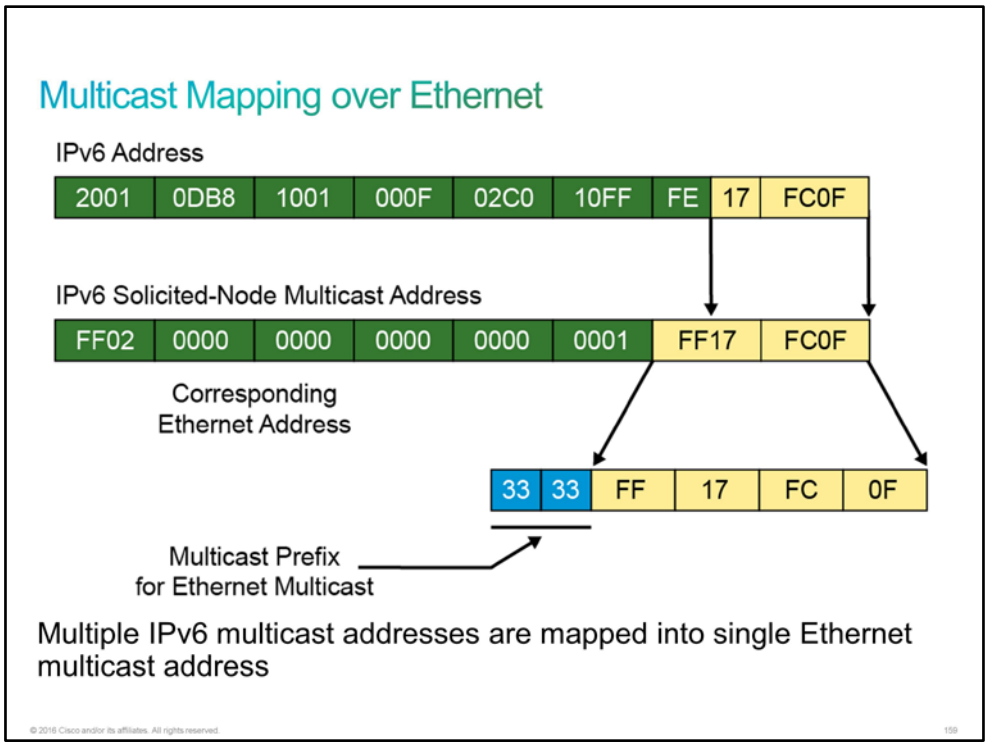

If an IPv6 multicast address is known, then the associated Ethernet MAC address is known. Multicast Ethernet addresses are formed by adding the last 32 bits of the IPv6 multicast address to 33:33.

As the figure shows, the IPv6 solicited-node multicast address is FF02::1:FF17:FC0F. The associated Ethernet MAC address is 33:33:FF:17:FC:0F.

You must understand that the resulting MAC address is a virtual MAC address: It is not burned into any Ethernet card. Depending on the IPv6 unicast address, which determines the IPv6 solicited-node multicast address, any Ethernet card may be instructed to listen to any of the 2^{24} possible virtual MAC addresses that begin with 33-33-FF. In IPv6, Ethernet cards often listen to multiple virtual multicast MAC addresses and their own burned-in unicast MAC addresses.

Stateless Address Autoconfiguration

Stateless autoconfiguration, also known as SLAAC, uses neighbor discovery mechanisms to find routers and dynamically assign IPv6 addresses.

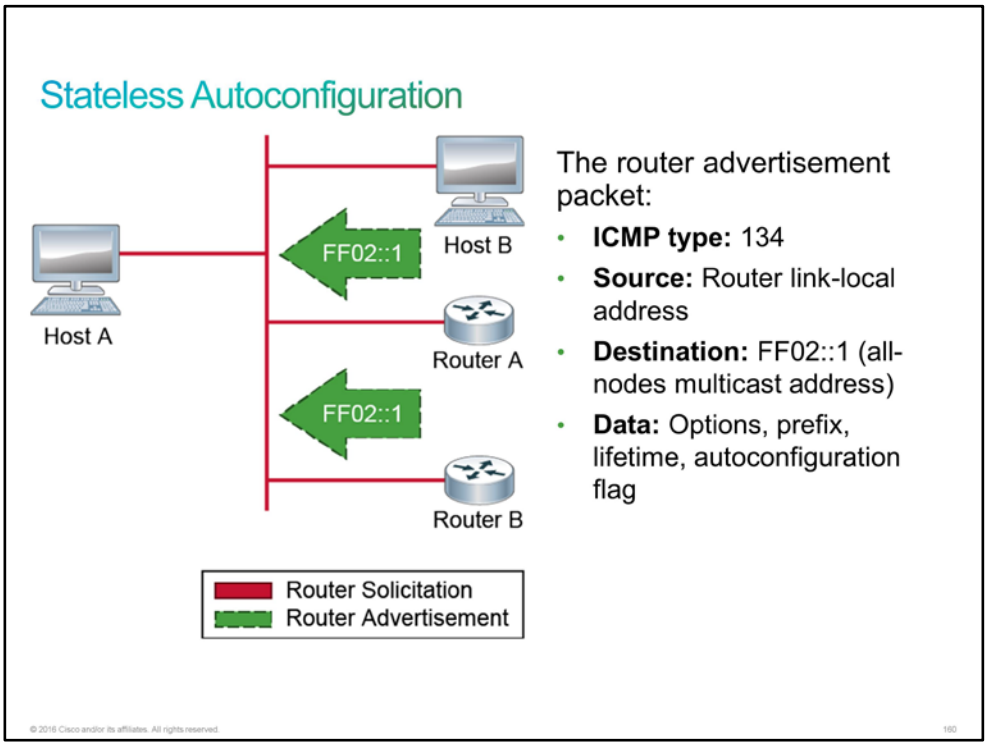

Routers periodically send router advertisements on all their configured interfaces. The router sends a router advertisement to the all-nodes multicast address.

Examples of information that the message might contain:

- **Prefixes that can be used on the link:** This information enables stateless autoconfiguration of the hosts. These prefixes must be /64 for stateless autoconfiguration.
- **Lifetime of the prefixes:** The default valid lifetime is 30 days, and the default preferred lifetime is 7 days.
- **Flags:** Flags indicate the kind of autoconfiguration that the hosts can perform.
- **Default router:** The default router advertises its existence and lifetime.
- **Other types of information for hosts:** This information can include the default MTU and hop count.

By sending prefixes, router advertisements allow host autoconfiguration. By assigning lifetimes to prefixes, a router advertisement enables administrative renumbering of hosts. When the lifetime of a deprecated prefix has decreased to zero, a new prefix will be assigned.

You can configure outer advertisement timing and other parameters on routers.

A router sends router advertisements immediately after a router solicitation. Router solicitations ask routers that are connected to the local link to send an immediate router advertisement so that the host can receive the autoconfiguration information without waiting for the next scheduled router advertisement.

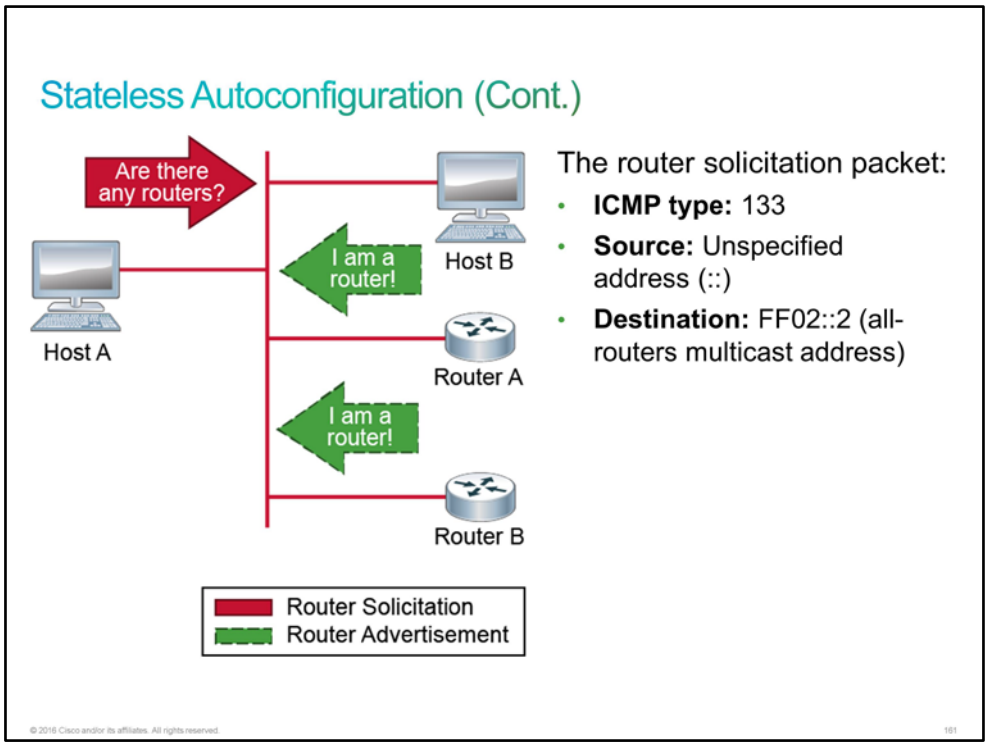

The router solicitation message is defined as follows:

- The ICMP type is 133.
- The source address is the unspecified address (or the IP address that is assigned to the sending interface when the IP address is known, which is usually not the case).
- The destination address is the all-routers multicast address with the link-local scope.

When a router sends an answer to a router solicitation, the destination address of the router advertisement is the unicast address of the requester.

A router should send a router solicitation only at the host boot time and only three times. This practice avoids flooding of router solicitation packets if there is no router on the local network.

Configuring Stateless Autoconfiguration

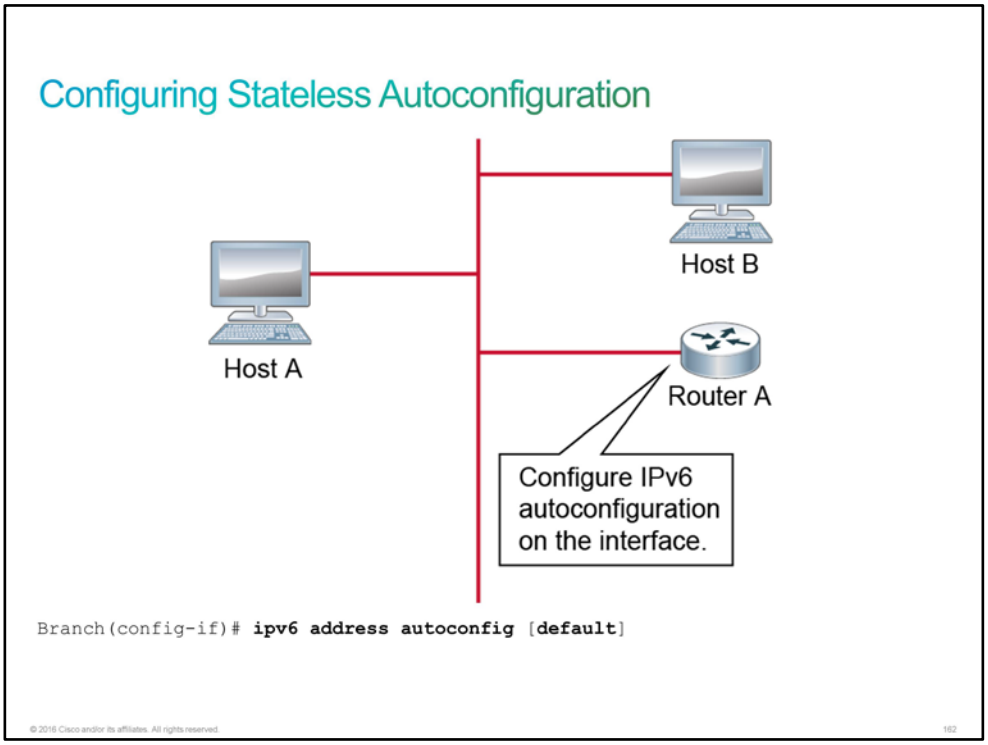

The **ipv6 address autoconfig** command enables stateless autoconfiguration on IPv6 routers on an interface-by-interface basis.

Command	Description
ipv6 address autoconfig [**default**]	Configures stateless autoconfiguration on the interface. If you add the **default** keyword, the router will advertise itself as the default route for the local link.

Discovery 28: Configure Basic IPv6 Connectivity

Introduction

In this discovery lab, you will explore the configuration of IPv6 in a small network containing three routers and three end hosts. Consult the topology diagram and address table to understand the network connectivity and addressing. All systems are currently configured with IPv4 addresses and RIP routing. During migration, IPv4 and IPv6 are commonly implemented in parallel with dual stacks on IPv6-capable systems. You will leave the IPv4 configuration in place during this exercise. Initially, IPv6 is also fully configured on R2 and PC2. This discovery lab will guide you through configuring IPv6 on the rest of the network.

First you will configure static IPv6 addresses on R1 and R3. Note, for simplicity, that all static IPv6 addresses in the topology differ in only 2 bytes. They all start with 2001:0DB8:0000:00. The eighth byte completes the 64-bit prefix and represents the network (01, 02, 03, 04, 05, or 06) within the topology. The next 7 bytes are all 00. The final byte specifies the host on the network; in this example, the byte is either 01 or 02.

After configuring the IPv6 addresses on R1 and R3, you will configure PC1 and SRV1 for IPv6 stateless autoconfiguration. You will then verify connectivity between PC1 and R1 and between SRV1 and R3.

Topology

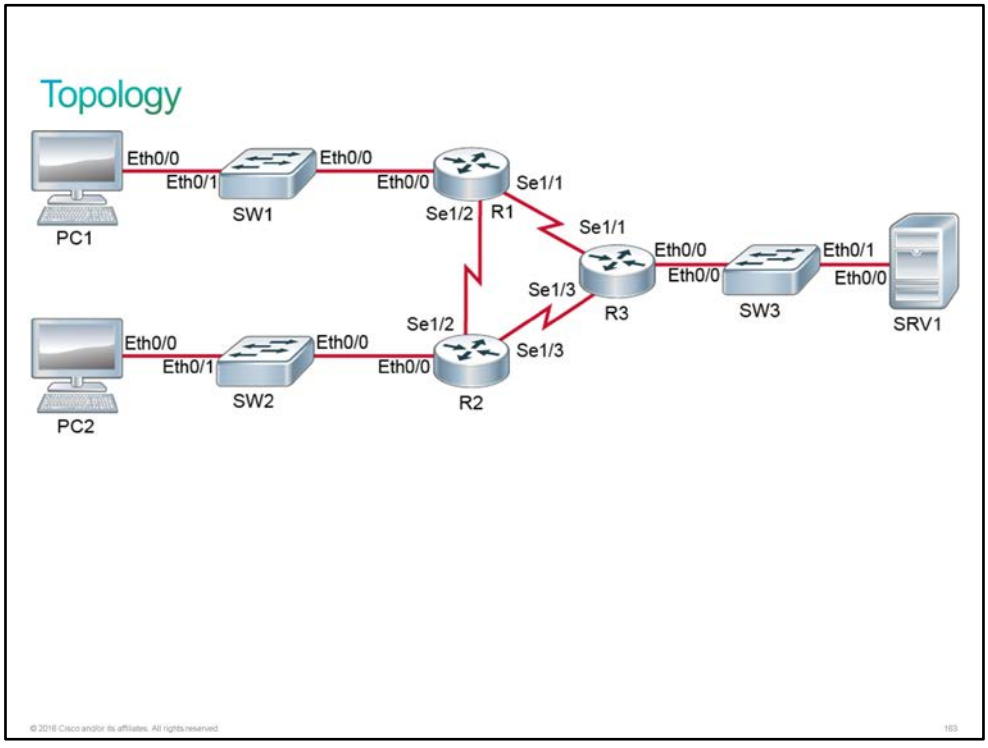

Job Aid

Device Information

Device Information Table

Device	Characteristic	Value
PC1	Hostname	PC1
PC1	IPv4 address	10.10.1.10/24
PC1	Default gateway	10.10.1.1
PC1	IPv6 address	2001:DB8:0:1::/64 Auto
PC2	Hostname	PC2
PC2	IPv4 address	10.10.2.20/24
PC2	Default gateway	10.10.2.1
PC2	IPv6 address	2001:DB8:0:2::/64 Auto
SRV1	Hostname	SRV1
SRV1	IPv4 address	10.10.3.30/24
SRV1	Default gateway	10.10.3.1
SRV1	IPv6 address	2001:DB8:0:3::/64 Auto
SW1	Hostname	SW1
SW1	VLAN 1 IPv4 address	10.10.1.4/24
SW1	Default gateway	10.10.1.1
SW1	Ethernet0/0 description	Link to R1
SW1	Ethernet0/1 description	Link to PC1
SW2	Hostname	SW2
SW2	VLAN 1 IPv4 address	10.10.2.4/24
SW2	Default gateway	10.10.2.1
SW2	Ethernet0/0 description	Link to R2
SW2	Ethernet0/1 description	Link to PC2

Device	Characteristic	Value
SW3	Hostname	SW3
SW3	VLAN 1 IPv4 address	10.10.3.4/24
SW3	Default gateway	10.10.3.1
SW3	Ethernet0/0 description	Link to R3
SW3	Ethernet0/1 description	Link to SRV1
R1	Hostname	R1
R1	Ethernet0/0 description	Link to SW1
R1	Ethernet0/0 IPv4 address	10.10.1.1/24
R1	Ethernet0/0 IPv6 address	2001:DB8:0:1::1/64
R1	Serial1/1 description	Link to R3
R1	Serial1/1 IPv4 address	10.1.1.2/30
R1	Serial1/1 IPv6 address	2001:DB8:0:4::1/64
R1	Serial1/2 description	Link to R2
R1	Serial1/2 IPv4 address	10.1.1.10/30
R1	Serial1/2 IPv6 address	2001:DB8:0:5::1/64
R2	Hostname	R2
R2	Ethernet0/0 description	Link to SW2
R2	Ethernet0/0 IPv4 address	10.10.2.1/24
R2	Ethernet0/0 IPv6 address	2001:DB8:0:2::1/64
R2	Serial1/2 description	Link to R1
R2	Serial1/2 IPv4 address	10.1.1.9/30
R2	Serial1/2 IPv6 address	2001:DB8:0:5::2/64
R2	Serial1/3 description	Link to R3
R2	Serial1/3 IPv4 address	10.1.1.6/30
R2	Serial1/3 IPv6 address	2001:DB8:0:6::1/64

Device	Characteristic	Value
R3	Hostname	R3
R3	Ethernet0/0 description	Link to SW3
R3	Ethernet0/0 IPv4 address	10.10.3.1/24
R3	Ethernet0/0 IPv6 address	2001:DB8:0:3::1/64
R3	Serial1/1 description	Link to R1
R3	Serial1/1 IPv4 address	10.1.1.1/30
R3	Serial1/1 IPv6 address	2001:DB8:0:4::2/64
R3	Serial1/3 description	Link to R2
R3	Serial1/3 IPv4 address	10.1.1.5/30
R3	Serial1/3 IPv6 address	2001:DB8:0:6::2/64

PCs and SRV in the virtual lab environment are simulated as routers, so you should use Cisco IOS commands to configure them or make verifications.

Task 1: Configure IPv6 Addresses

Activity

Step 1 On the R1 router, enable IPv6 routing.

By default routing for IPv6 is not enabled on Cisco router. You will need to enable IPv6 routing on routers. If IPv6 routing is not enabled, router may still play a role of IPv6 host.

On R1, enter the following commands:

```
R1# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
R1(config)# ipv6 unicast-routing
```

The most commonly used commands are abbreviated in this guided discovery—for example, **conf t** for **configure terminal**. If there is any confusion, you can attempt tab completion on commands to see the full commands during the discovery execution. For example, **conf<tab> t<tab>** would expand to **configure terminal**.

Step 2 On R1, configure the IPv6 address 2001:db8:0:5::1/64 on the Serial1/2 interface.

On R1, enter the following commands:


```
R1(config)# int Serial 1/2
R1(config-if)# ipv6 address 2001:db8:0:5::1/64
```

Step 3 R2 is fully IPv6-configured, and Serial1/1 is the R1 link to R2. If you correctly configured the address of R1, you should be able to ping the R2 IPv6 address (2001:db8:0:5::2). Enter the **do** command to execute an EXEC mode **ping** to confirm that there is connectivity from R1 to R2.

On R1, enter the following command:

```
R1(config-if)# do ping 2001:db8:0:5::2
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2001:DB8:0:5::2, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 6/8/9 ms
```

Step 4 Configure the R1 IPv6 addresses on Ethernet0/0 (2001:db8:0:1::1/64) and Serial1/1 (2001:db8:0:4::1/64). Remember to take advantage of the Cisco IOS command recall. The IPv6 addresses in the topology are very similar. Currently, there are no configured IPv6 peers on Ethernet0/0 or Serial1/1, so you cannot use the **ping** command for verification. Leave the configuration mode when the addressing is complete.

On R1, enter the following commands:

```
R1(config-if)# int e0/0
R1(config-if)# ipv6 address 2001:db8:0:1::1/64
R1(config-if)# int s 1/1
R1(config-if)# ipv6 address 2001:db8:0:4::1/64
R1(config-if)# end
R1#
```

Step 5 On R3, enable IPv6 routing.

On R3, enter the following commands:

```
R3# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
R3(config)# ipv6 unicast-routing
```

Step 6 Configure the R3 IPv6 address (2001:db8:0:4::2/64) on Serial1/1, then verify that you can ping R1 (2001:db8:0:4::1) from R3.

On R3, enter the following commands:

```
R3(config)# int Serial 1/1
R3(config-if)# ipv6 address 2001:db8:0:4::2/64
R3(config-if)# do ping 2001:db8:0:4::1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2001:DB8:0:4::1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 8/8/9 ms
```

Step 7 Configure the R3 IPv6 address (2001:db8:0:6::2/64) on Serial1/3, then verify that you can ping R2 (2001:db8:0:6::1) from R3. Remember to take advantage of the Cisco IOS command recall feature.

On R3, enter the following commands:

```
R3(config-if)# int Serial 1/3
R3(config-if)# ipv6 address 2001:db8:0:6::2/64
R3(config-if)# do ping 2001:db8:0:6::1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2001:DB8:0:6::1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 8/8/9 ms
```

Step 8 Configure the R3 IPv6 address (2001:db8:0:3::1/64) on Ethernet0/0. There are currently no IPv6 peers on Ethernet0/0, so you cannot use the **ping** command for verification. Leave the configuration mode when you are done configuring Ethernet0/0.

On R3, enter the following commands:

```
R3(config-if)# int Ethernet 0/0
R3(config-if)# ipv6 address 2001:db8:0:3::1/64
R3(config-if)# end
R3#
```

Step 9 On R3, display the Ethernet0/0 MAC address using the **show interface** command. The output can be run through the **include** filter using **address** as the filter string to reduce the amount of command output.

On R3, enter the following command:

```
R3# show int e0/0 | include address
  Hardware is AmdP2, address is aabb.cc00.0300 (bia aabb.cc00.0300)
  Internet address is 10.10.3.1/24
```

The MAC address in your output may differ.

Step 10 On R3, use the **show ipv6 interface brief** command to display the IPv6 addresses that are assigned to the R3 interfaces.

On R3, enter following command:


```
R3# sh ipv6 int brief
Ethernet0/0            [up/up]
    FE80::A8BB:CCFF:FE00:300
    2001:DB8:0:3::1
Ethernet0/1            [administratively down/down]
    unassigned
Ethernet0/2            [administratively down/down]
    unassigned
Ethernet0/3            [administratively down/down]
    unassigned
Serial1/0              [administratively down/down]
    unassigned
Serial1/1              [up/up]
    FE80::A8BB:CCFF:FE00:300
    2001:DB8:0:4::2
Serial1/2              [administratively down/down]
    unassigned
Serial1/3              [up/up]
    FE80::A8BB:CCFF:FE00:300
    2001:DB8:0:6::2
```

There are two IPv6 addresses on each of the three configured interfaces. There is a link-local address that was statelessly autoconfigured. There is also the global unicast address that you configured.

The statelessly autoconfigured link-local address is configured using the EUI-64 standard on the FE80::/10 prefix. The algorithm that the EUI-64 standard uses to stretch the 48-bit MAC address to 64 bits is to invert the seventh bit of the MAC address and to insert FFFE into the middle of the MAC address. So, AA:BB:CC:00:03:00 becomes A8BB:CCFF:FE00:300.

The serial interfaces, being point-to-point links, do not use MAC addresses. IPv6 "borrows" the MAC address from an Ethernet interface to compute the link-local address for serial interfaces. The result is that R3 is using the same link-local address on multiple interfaces. This situation is acceptable because each interface is in its own broadcast domain; therefore, the link-local address is unique within the broadcast domain.

Step 11 On R1, display the full IPv6 information that is associated with Ethernet0/0.

On R1, enter the following command:

```
R1# sh ipv6 int e0/0
Ethernet0/0 is up, line protocol is up
  IPv6 is enabled, link-local address is FE80::A8BB:CCFF:FE00:100
  No Virtual link-local address(es):
  Description: Link to SW1
  Global unicast address(es):
    2001:DB8:0:1::1, subnet is 2001:DB8:0:1::/64
  Joined group address(es):
    FF02::1
    FF02::2
    FF02::1:FF00:1
    FF02::1:FF00:100
  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.
```

The output displays both the global unicast address and the link-local address.

IPv6 automatically joins several required multicast groups. All addresses that start with FF02 are IPv6 multicast addresses.

IPv6 neighbor discovery is automatically enabled on IPv6 interfaces. R1 will send neighbor discovery router advertisements containing the global unicast prefix on Ethernet0/0. The hosts on this network can use these advertisements for stateless autoconfiguration.

Task 2: Configure IPv6 Stateless Autoconfiguration

Activity

Step 1 With R3 sending neighbor discovery router advertisements on its Ethernet0/0 interface, SRV1 can use stateless autoconfiguration for IPv6. On SRV1, display its MAC address.

On SRV1, enter the following command:

```
SRV1# show int e0/0 | inc address
  Hardware is AmdP2, address is aabb.cc00.0e00 (bia aabb.cc00.0e00)
  Internet address is 10.10.3.30/24
```

You are observing the MAC address, where you can see how it is used with the EUI-64 standard to generate the SRV1 IPv6 address with stateless autoconfiguration.

MAC address in your output may be different.

Step 2 On SRV1, configure Ethernet0/0 to use stateless autoconfiguration for the IPv6 address assignment and for the IPv6 default route assignment.

PCs (for example, Windows) are typically enabled for SLAAC by default.

On SRV1, enter the following commands:

```
SRV1# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
SRV1(config)# int e 0/0
SRV1(config-if)# ipv6 address autoconfig default
SRV1(config-if)# end
SRV1#
```

Step 3 On SRV1, display the IPv6 addresses that are assigned to Ethernet0/0.

On SRV1, enter the following command:

```
SRV1# show ipv6 interface brief e0/0
Ethernet0/0            [up/up]
    FE80::A8BB:CCFF:FE00:E00
    2001:DB8:0:3:A8BB:CCFF:FE00:E00
```

There are two addresses: the link-local address using the standard FE80::/10 prefix, and the global unicast address using the 2001:DB8:0:3::/64 prefix that SRV1 received from the R3 router advertisement. Both use the EUI-64 standard to incorporate the Ethernet0/0 MAC address into the IPv6 address.

Step 4 Display the IPv6 routing table on SRV1.

On SRV1, enter the following command:

```
SRV1# sh ipv6 route
IPv6 Routing Table - default - 4 entries
Codes: C - Connected, L - Local, S - Static, U - Per-user Static route
       B - BGP, HA - Home Agent, MR - Mobile Router, R - RIP
       H - NHRP, I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea
       IS - ISIS summary, D - EIGRP, EX - EIGRP external, NM - NEMO
       ND - ND Default, NDp - ND Prefix, DCE - Destination, NDr - Redirect
       O - OSPF Intra, OI - OSPF Inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2
       ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2, l - LISP
ND  ::/0 [2/0]
     via FE80::A8BB:CCFF:FE00:300, Ethernet0/0
NDp 2001:DB8:0:3::/64 [2/0]
     via Ethernet0/0, directly connected
L   2001:DB8:0:3:A8BB:CCFF:FE00:E00/128 [0/0]
     via Ethernet0/0, receive
L   FF00::/8 [0/0]
     via Null0, receive
```

The default route (to prefix ::/0) is to the R3 link-local address.

Step 5 At this point, SRV1 should be able to ping the R3 global unicast addresses on Ethernet0/0 (2001:db8:0:3::1), on Serial1/1 (2001:db8:0:4::2), and Serial1/3 (2001:db8:0:6::2). Confirm this connectivity using the **ping** command. Again, be sure to take advantage of the Cisco IOS command recall feature.

On SRV1, enter the following commands:

```
SRV1# ping 2001:db8:0:3::1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2001:DB8:0:3::1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/4/20 ms
SRV1# ping 2001:db8:0:4::2
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2001:DB8:0:4::2, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
SRV1# ping 2001:db8:0:6::2
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2001:DB8:0:6::2, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
```

Step 6 However, you cannot ping addresses on R1 or R2. Attempt to ping the R1 Serial1/1 interface (2001:db8:0:4::1).

On SRV1, enter the following command:

```
SRV1# ping 2001:db8:0:4::1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2001:DB8:0:4::1, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
```

Because R3 is directly connected to the 2001:db8:0:4::/64 subnet, it can successfully send the probe to R1. The **ping** fails because R1 does not (yet) have a route back to the 2001:db8:0:3/64 network where SRV1 is connected.

Step 7 On PC1, configure Ethernet0/0 to use stateless autoconfiguration and default route assignment.

On PC1, enter the following commands :

```
PC1# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
PC1(config)# int Ethernet 0/0
PC1(config-if)# ipv6 address autoconfig default
PC1(config-if)# end
PC1#
```

Step 8 On PC1, display the IPv6 addresses that are assigned to Ethernet0/0.

On PC1, enter following command:

```
PC1# sh ipv6 int brief e0/0
Ethernet0/0            [up/up]
    FE80::A8BB:CCFF:FE00:C00
    2001:DB8:0:1:A8BB:CCFF:FE00:C00
```

Step 9 Display the IPv6 routing table on PC1 to verify that it has an IPv6 default route.

On PC1, enter following command:


```
PC1# show ipv6 route
IPv6 Routing Table - default - 4 entries
Codes: C - Connected, L - Local, S - Static, U - Per-user Static route
       B - BGP, HA - Home Agent, MR - Mobile Router, R - RIP
       H - NHRP, I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea
       IS - ISIS summary, D - EIGRP, EX - EIGRP external, NM - NEMO
       ND - ND Default, NDp - ND Prefix, DCE - Destination, NDr - Redirect
       O - OSPF Intra, OI - OSPF Inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2
       ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2, l - LISP
ND  ::/0 [2/0]
     via FE80::A8BB:CCFF:FE00:100, Ethernet0/0
NDp 2001:DB8:0:1::/64 [2/0]
     via Ethernet0/0, directly connected
L   2001:DB8:0:1:A8BB:CCFF:FE00:C00/128 [0/0]
     via Ethernet0/0, receive
L   FF00::/8 [0/0]
     via Null0, receive
```

Step 10 From PC1, verify that you can ping the R1 Ethernet0/0 (2001:db8:0:1::1), Serial1/1 (2001:db8:0:4::1), and Serial1/2 (2001:db8:0:5::1) interfaces.

On PC1, enter the following commands:

```
PC1# ping 2001:db8:0:1::1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2001:DB8:0:1::1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/3/15 ms
PC1# ping 2001:db8:0:4::1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2001:DB8:0:4::1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
PC1# ping 2001:db8:0:5::1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2001:DB8:0:5::1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
```

This is the end of the discovery lab.

Challenge

1. Which two IPv4 header fields are similar to fields in the IPv6 header? (Choose two.)

   A. header length
   B. flags
   C. fragment offset
   D. header checksum
   E. total length
   F. type of service

2. Which IPv6 header field specifies the maximum number of hops that an IPv6 packet can take?

   A. hop limit
   B. next header
   C. time to live
   D. flow label

3. Which ICMPv6 type is used for neighbor solicitation?

   A. ICMPv6 type 1
   B. ICMPv6 type 128
   C. ICMPv6 type 135
   D. ICMPv6 type 136

4. Which two ICMPv6 types are used for testing IPv6 reachability? (Choose two.)

   A. ICMPv6 type 1
   B. ICMPv6 type 128
   C. ICMPv6 type 129
   D. ICMPv6 type 135
   E. ICMPv6 type 136

5. What is the solicited-node multicast IP address for 2001:DB8:1001:F:2C0:10FF:FE17:FC0F?

   A. FF02:0000:0000:0000:0000:0001:FF17:FC0F
   B. 2001:0000:0000:0000:0000:0001:FF17:FC0F
   C. FF02:0000:0000:0000:0000:0001:FE17:FC0F
   D. FF02:0000:0000:0000:0000:10FF:FE17:FC0F

6. What is the corresponding Ethernet address for FF02:0000:0000:0000:0000:0001:FF17:FC0F?

   A. 33-33-33-17-FC-0F
   B. 00-00-33-17-FC-0F
   C. 33-33-FF-17-FC-0F
   D. FF-02-FF-17-FC-0F

7. Which Cisco IOS command do you use to configure stateless autoconfiguration?

   A. **ipv6 autoconfig**
   B. **ipv6 autoconfig-address**
   C. **ipv6 address stateless-autoconfig**
   D. **ipv6 address autoconfig**

Answer Key

Challenge

1. E, F
2. A
3. C
4. B, C
5. A
6. C
7. D

Lesson 3: Configuring IPv6 Static Routes

Introduction

Your boss sends you to your customer to describe the routing options for IPv6 and to show how to configure and verify IPv6 static routes.

Routing for IPv6

To support IPv6, all the IPv4 routing protocols had to go through varying degrees of changes, with the most obvious being that each had to be changed to support longer addresses and prefixes.

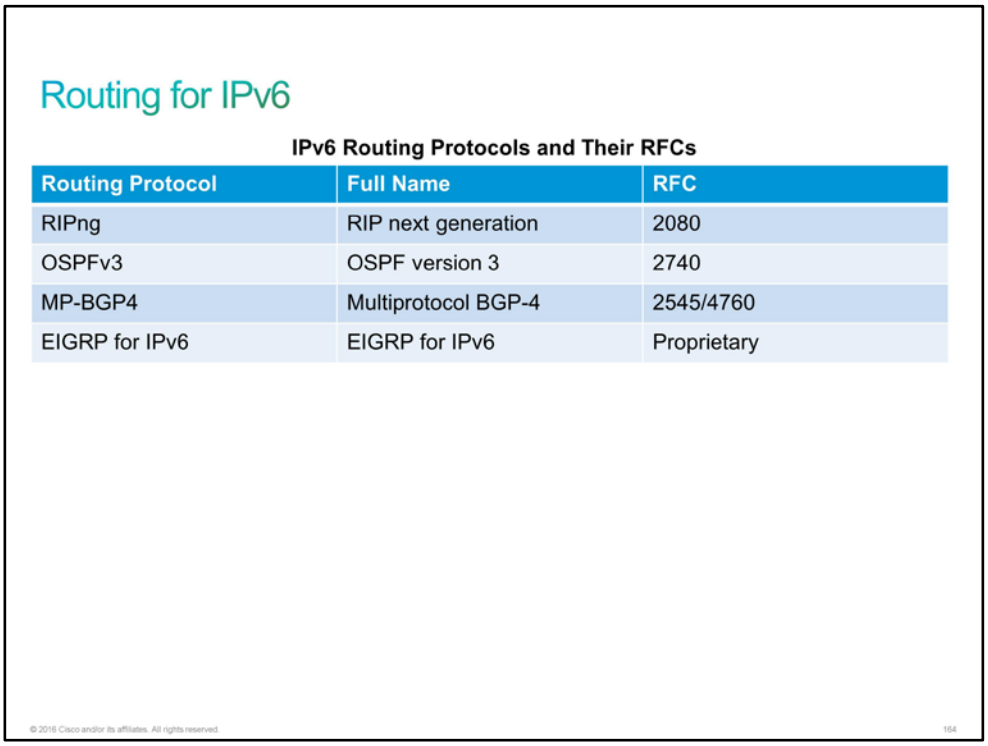

As with IPv4, most IPv6 routing protocols are IGPs, with BGP still being the only EGP of note. All these IGPs and BGP were updated to support IPv6. The table lists the routing protocols and their new RFCs.

Each of these routing protocols has to make several changes to support IPv6. The actual messages that are used to send and receive routing information have changed, using IPv6 headers instead of IPv4 headers, and using IPv6 addresses in those headers. For example, RIPng sends routing updates to the IPv6 destination address FF02::9, instead of to the former RIPv2 IPv4 224.0.0.9 address. Also, the routing protocols typically advertise their link-local IP address as the next hop in a route.

The routing protocols still retain many same internal features. For example, RIPng based on RIPv2 is still a distance vector protocol, with the hop count as the metric and 15 hops as the longest valid route (16 is infinity). OSPFv3, which was created specifically to support IPv6, is still a link-state protocol, with the cost as the metric but with many internals, including LSA types, changed. As a result, OSPFv2 is not compatible with OSPFv3. However, the core operational concepts remain the same.

You can also use and configure IPv6 static routing in the same way that you would with IPv4. There is an IPv6-specific requirement per RFC 2461 that a router must be able to determine the link-local address of each of its neighboring routers to ensure that the target address of a redirect message identifies the neighbor router by its link-local address. This requirement means that using a global unicast address as a next-hop address with IPv6 routing is not recommended.

Configuring IPv6 Static Routes

Configuring a static route for IPv6 is almost the same as it is in IPv4. In IPv4, the next-hop IP address and the exit interface can be specified. In IPv6, there is an extra feature. The next-hop IP address in IPv6 can either be the link-local address or the global address. The following example shows how to configure an IPv6 static route using these three different methods:

```
Router# conf t
Router(config)# ipv6 route 2001:0db8:BEEF::/32 FA1/0
Router(config)# ipv6 route 2001:0db8:BEEF::/32 FA1/0 fe80::2
Router(config)# ipv6 route 2001:0db8:BEEF::/32 2001:0db8:FEED::1
```

The first static route shows that the route to the network 2001:0db8:BEEF::/32 is configured via the FastEthernet1/0 interface. The second static route gives an option of the link-local next hop address, which is specified with the fe80 prefix. The third static route shows a route to the network that points to the global IPv6 address of 2001:0db8:FEED::1.

IPv6 Static Route Configuration Example

Consider the following example to understand IPv6 static route configuration:

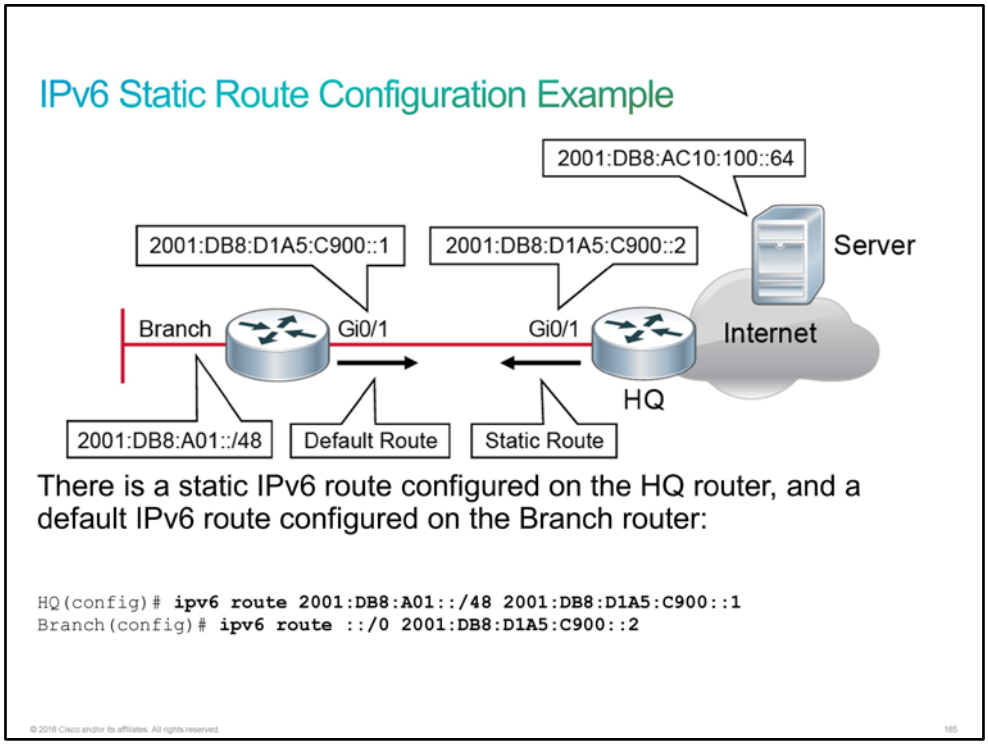

The following table shows IPv6 static and default route commands:

Configuring an IPv6 Static Route

Command and Variable	Description
ipv6 route *ipv6-network/ipv6-mask [outgoing_interface] ipv6-next-hop*	Configures an IPv6 static route. When next hop is link local address, outgoing interface must be specified.
ipv6 route ::/0 *[outgoing-interface] ipv6-next-hop*	Configures a default IPv6 route. When next hop is link local address, outgoing interface must be specified.

Verifying IPv6 Static Route Configuration

Use the **show ipv6 route static** or **show ipv6 route** command to verify the IPv6 static route configuration.

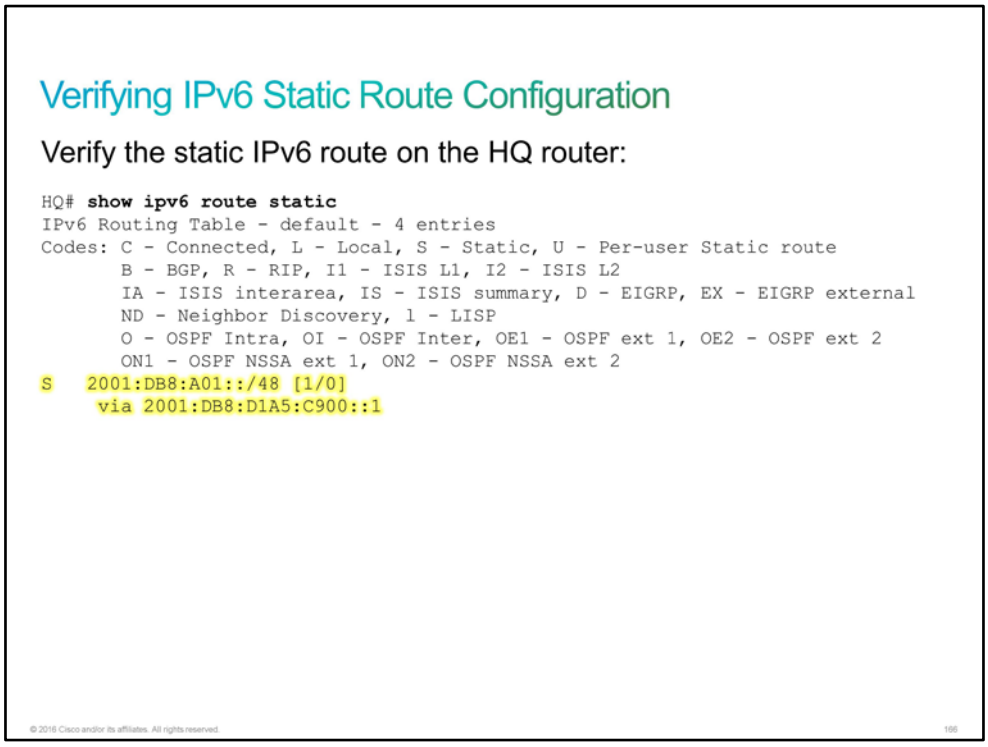

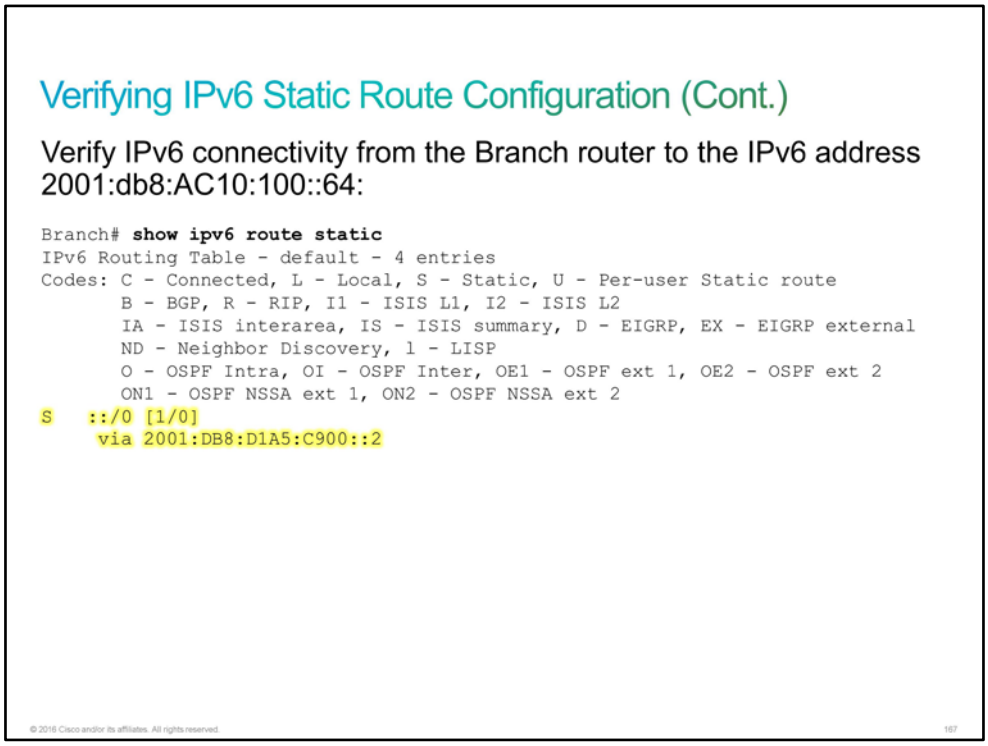

The following table shows IPv6 static route verification commands:

IPv6 Static Route Verification

Command and Variable	Purpose
show ipv6 static *[ipv6-address | ipv6-prefix/prefix-length][interface interface-type interface-number]* **[recursive]** *[***detail***]* or **show ipv6 route** *[ipv6-address | ipv6-prefix/prefix-length | protocol | interface-type interface-number]*	Displays the current contents of the IPv6 routing table. These examples show two different ways of displaying IPv6 static routes.

Example:

```
Router# show ipv6 static
```

or

```
Router# show ipv6 route static
```

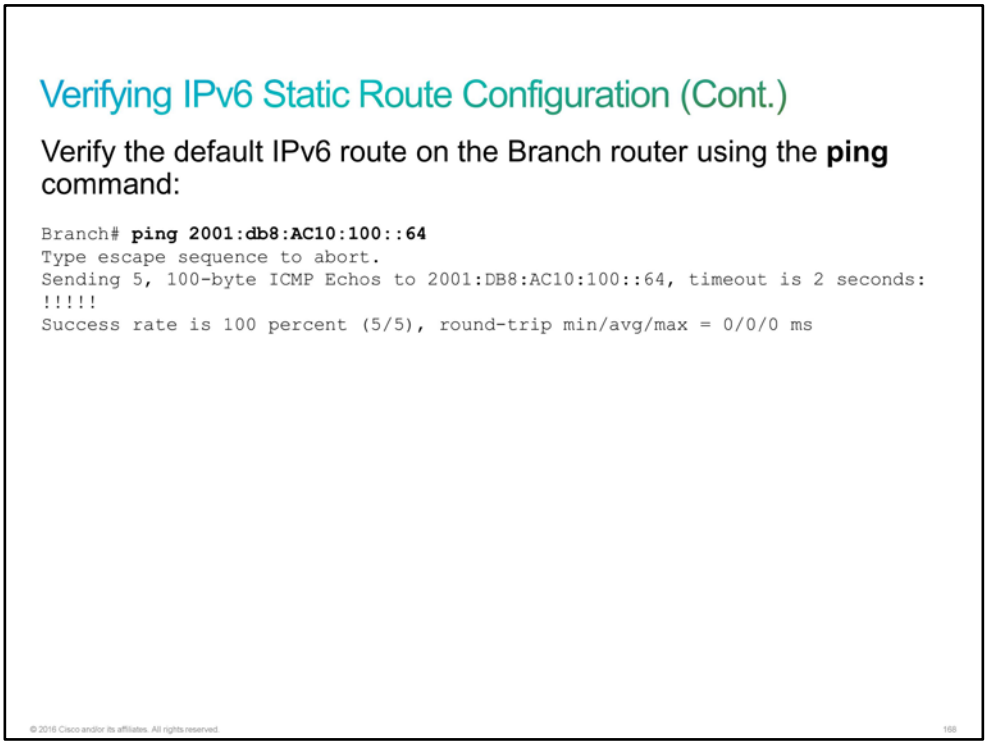

Discovery 29: Configure IPv6 Static Routes

Introduction

In this discovery lab, you will configure IPv6 static routing between R1 and R3 and verify connectivity between PC1 and SRV1. Consult the topology diagram and address table to understand the network connectivity and addressing. All systems are currently configured with IPv4 addresses, IPv6 addresses, and RIPv2 routing. You will leave the IPv4 configuration in place during this exercise. Also, you will configure a default static route on PC1.

Topology

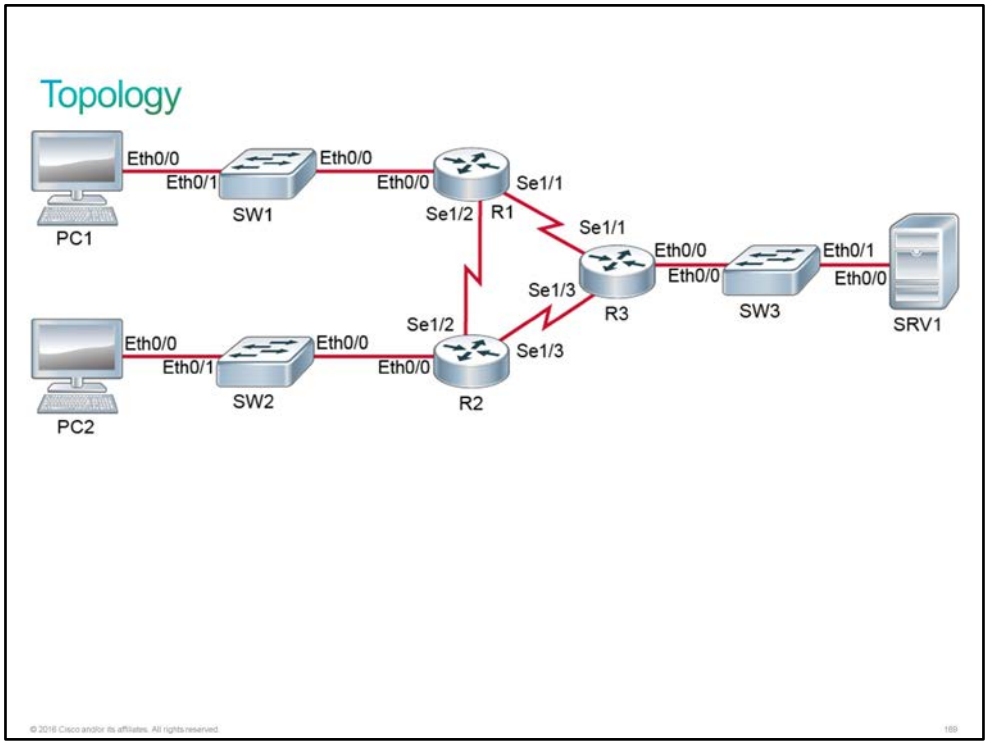

Job Aid

Device Information

Device Information Table

Device	Characteristic	Value
PC1	Hostname	PC1
PC1	IPv4 address	10.10.1.10/24
PC1	Default gateway	10.10.1.1
PC1	IPv6 address	2001:DB8:0:1::/64 Auto

Device	Characteristic	Value
PC2	Hostname	PC2
PC2	IPv4 address	10.10.2.20/24
PC2	Default gateway	10.10.2.1
PC2	IPv6 address	2001:DB8:0:2::/64 Auto
SRV1	Hostname	SRV1
SRV1	IPv4 address	10.10.3.30/24
SRV1	Default gateway	10.10.3.1
SRV1	IPv6 address	2001:DB8:0:3::100/64
SW1	Hostname	SW1
SW1	VLAN 1 IPv4 address	10.10.1.4/24
SW1	Default gateway	10.10.1.1
SW1	Ethernet0/0 description	Link to R1
SW1	Ethernet0/1 description	Link to PC1
SW2	Hostname	SW2
SW2	VLAN 1 IPv4 address	10.10.2.4/24
SW2	Default gateway	10.10.2.1
SW2	Ethernet0/0 description	Link to R2
SW2	Ethernet0/1 description	Link to PC2
SW3	Hostname	SW3
SW3	VLAN 1 IPv4 address	10.10.3.4/24
SW3	Default gateway	10.10.3.1
SW3	Ethernet0/0 description	Link to R3
SW3	Ethernet0/1 description	Link to SRV1
R1	Hostname	R1
R1	Ethernet0/0 description	Link to SW1

Device	Characteristic	Value
R1	Ethernet0/0 IPv4 address	10.10.1.1/24
R1	Ethernet0/0 IPv6 address	2001:DB8:0:1::1/64
R1	Serial1/1 description	Link to R3
R1	Serial1/1 IPv4 address	10.1.1.2/30
R1	Serial1/1 IPv6 address	2001:DB8:0:4::1/64
R1	Serial1/2 description	Link to R2
R1	Serial1/2 IPv4 address	10.1.1.10/30
R1	Serial1/2 IPv6 address	2001:DB8:0:5::1/64
R2	Hostname	R2
R2	E0/0 description	Link to SW2
R2	E0/0 IPv4 address	10.10.2.1/24
R2	E0/0 IPv6 address	2001:DB8:0:2::1/64
R2	S1/2 description	Link to R1
R2	S1/2 IPv4 address	10.1.1.9/30
R2	S1/2 IPv6 address	2001:DB8:0:5::2/64
R2	S1/3 description	Link to R3
R2	S1/3 IPv4 address	10.1.1.6/30
R2	S1/3 IPv6 address	2001:DB8:0:6::1/64
R3	Hostname	R3
R3	Ethernet0/0 description	Link to SW3
R3	Ethernet0/0 IPv4 address	10.10.3.1/24
R3	Ethernet0/0 IPv6 address	2001:DB8:0:3::1/64
R3	Serial1/1 description	Link to R1
R3	Serial1/1 IPv4 address	10.1.1.1/30
R3	Serial1/1 IPv6 address	2001:DB8:0:4::2/64

Device	Characteristic	Value
R3	Serial1/3 description	Link to R2
R3	Serial1/3 IPv4 address	10.1.1.5/30
R3	Serial1/3 IPv6 address	2001:DB8:0:6::2/64

PCs and SRV in the virtual lab environment are simulated as routers, so you should use Cisco IOS commands to configure them or make verifications.

Task 1: Configure IPv6 Default Routes

Activity

Step 1 On PC1, remove the IPv6 stateless autoconfiguration and set a fixed IPv6 address (2001:DB8:0:1::100/64).

On PC1, enter the following commands:

```
PC1# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
PC1(config)# inter e0/0
PC1(config-if)# no ipv6 address autoconfig default
PC1(config-if)# ipv6 address 2001:DB8:0:1::100/64
PC1(config-if)# exit
```

Step 2 On PC1, configure the IPv6 default route using R1's Ethernet0/0 address (2001:DB8:0:1::1).

On PC1, enter the following commands:

```
PC1(config)# ipv6 route ::/0 2001:DB8:0:1::1
PC1(config)# end
PC1#
```

Step 3 On PC1, verify the IPv6 routing table. You should see a static default route in the IPv6 routing table.

On PC1, enter the following command:


```
PC1# show ipv6 route
IPv6 Routing Table - default - 4 entries
Codes: C - Connected, L - Local, S - Static, U - Per-user Static route
       B - BGP, HA - Home Agent, MR - Mobile Router, R - RIP
       H - NHRP, I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea
       IS - ISIS summary, D - EIGRP, EX - EIGRP external, NM - NEMO
       ND - ND Default, NDp - ND Prefix, DCE - Destination, NDr - Redirect
       O - OSPF Intra, OI - OSPF Inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2
       ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2, l - LISP
S   ::/0 [1/0]
     via 2001:DB8:0:1::1
C   2001:DB8:0:1::/64 [0/0]
     via Ethernet0/0, directly connected
L   2001:DB8:0:1::100/128 [0/0]
     via Ethernet0/0, receive
L   FF00::/8 [0/0]
     via Null0, receive
```

Task 2: Configure IPv6 Static Routes

Activity

Step 1 Now it is time to configure IPv6 static routing between networks. On R3, configure a static route to the 2001:db8:0:1::/64 subnet using R1's Serial1/1 address (2001:db8:0:4::1).

On R3, enter the following commands:

```
R3# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
R3(config)# ipv6 route 2001:db8:0:1::/64 2001:db8:0:4::1
R3(config)# end
R3#
```

Step 2 On R3, view the IPv6 routing table.

On R3, enter the following command:

```
R3# show ipv6 route
IPv6 Routing Table - default - 8 entries
Codes: C - Connected, L - Local, S - Static, U - Per-user Static route
       B - BGP, HA - Home Agent, MR - Mobile Router, R - RIP
       H - NHRP, I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea
       IS - ISIS summary, D - EIGRP, EX - EIGRP external, NM - NEMO
       ND - ND Default, NDp - ND Prefix, DCE - Destination, NDr - Redirect
       O - OSPF Intra, OI - OSPF Inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2
       ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2, l - LISP
S   2001:DB8:0:1::/64 [1/0]
     via 2001:DB8:0:4::1
C   2001:DB8:0:3::/64 [0/0]
     via Ethernet0/0, directly connected
L   2001:DB8:0:3::1/128 [0/0]
     via Ethernet0/0, receive
C   2001:DB8:0:4::/64 [0/0]
     via Serial1/1, directly connected
L   2001:DB8:0:4::2/128 [0/0]
     via Serial1/1, receive
C   2001:DB8:0:6::/64 [0/0]
     via Serial1/3, directly connected
L   2001:DB8:0:6::2/128 [0/0]
     via Serial1/3, receive
L   FF00::/8 [0/0]
     via Null0, receive
```

Among the connected and local routes, you should see the static route that you just configured.

Step 3 Presently, R3 should be able to ping PC1 (2001:DB8:0:1::100). Confirm that this assumption is true.

On R3, enter the following command:

```
R3# ping 2001:DB8:0:1::100
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2001:DB8:0:1::100, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 8/16/45 ms
```

Step 4 There is now connectivity between R3 and PC1, but to establish connectivity between SRV1 and PC1, R1 needs to have a configured route. On R1, configure a static route to the 2001:db8:0:3::/64 subnet using R3's Serial1/1 address (2001:db8:0:4::2).

On R3, enter the following commands:

```
R1# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
R1(config)# ipv6 route 2001:db8:0:3::/64 2001:db8:0:4::2
R1(config)# end
R1#
```

Step 5 Validate the static routing between 2001:db8:0:1::/64 and 2001:db8:0:3::/64. From PC1, ping SRV1 (2001:DB8:0:3::100).

On PC1, enter the following command:


```
PC1# ping 2001:DB8:0:3::100
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2001:DB8:0:3::100, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 9/9/10 ms
```

Step 6 From PC1, verify that R1 and R3 are the intermediate hops to the SRV1.

On PC1, enter the following command:

```
PC1# traceroute 2001:DB8:0:3::100
Type escape sequence to abort.
Tracing the route to 2001:DB8:0:3::100

  1 2001:DB8:0:1::1 1 msec 1 msec 1 msec
  2 2001:DB8:0:4::2 9 msec 9 msec 10 msec
  3 2001:DB8:0:3::100 10 msec 10 msec 10 msec
```

Step 7 From PC1, use Telnet to connect to SRV1. Authenticate as **admin** with the password **Cisco123**. Once the telnet connection establishes, use the **exit** command to terminate the connection.

On PC1, enter the following commands:

```
PC1# telnet 2001:DB8:0:3::100
Trying 2001:DB8:0:3::100 ... Open

User Access Verification

Username: admin
Password: Cisco123
SRV1> exit

[Connection to 2001:DB8:0:3::100 closed by foreign host]
PC1#
```

This is the end of the discovery lab.

Challenge

1. RIPng sends routing updates to the destination address 224.0.0.9. True or False?

   A. True
   B. False

2. OSPFv3 was created to support IPv6 and also OSPFv3 is compatible with OSPFv2. True or False.?

   A. True
   B. False

3. In IPv6, creating static routes is just like creating IPv4 static routes, but there is an extra feature. What is this feature?

   A. The next hop can be a link local or a global address.
   B. The static route can be part of a routing protocol.
   C. The destination for the static route can be dynamically updated.
   D. The distance for the static route can be updated using a policy.

4. Which of the following is a valid, working IPv6 route that uses only an interface as the next hop?

   A. **Router(config)# ipv6 route FA1/0**
   B. **Router(config)# ipv6 route 2001:0db8:BEEF::/32 FA1/0**
   C. **Router(config)# ipv6 route 2001:0db8:BEEF::/32 FA1/0 fe80::2**
   D. **Router(config)# ipv6 route 2001:0db8:BEEF::/32 2001:0db8:FEED::1**

5. The command: **show ipv6 static** shows the IPv6 static routes from the routing table. True or False?

   A. True
   B. False

6. Which of the following is the way to make a IPv6 default route?

   A. **ipv6 route 0.0.0.0 0.0.0.0**
   B. **ipv6 route ::/0**
   C. **ipv6 route FF::0/32**
   D. **ipv6 route FF::0/128**

7. Which of the following is the IPv6 version of BGP?

   A. BGPng
   B. BGPv3
   C. BGP for IPv6
   D. MP-BGP4

Answer Key

Challenge

1. B
2. B
3. A
4. B
5. A
6. B
7. D

Glossary

3G mobile network

third-generation mobile network. Refers generically to a category of next-generation mobile networks, such as UMTS and IMT-2000.

4G

fourth generation of mobile phone mobile communications standards.

5G

fifth-generation mobile network.

ACK

acknowledgment. Notification sent from one network device to another to acknowledge that some event occurred (for example, the receipt of a message).

ACL

access control list. A list kept by routers to control access to or from the router for a number of services (for example, to prevent packets with a certain IP address from leaving a particular interface on the router).

AfriNIC

African Network Information Center. AfriNIC is the Regional Internet Registry (RIR) for Africa. They are responsible for the distribution and management of Internet number resources (such as IP addresses and ASN Numbers) for the African region.

AH

Authentication Header. A security protocol that provides data authentication and optional antireplay services. AH is embedded in the data to be protected (a full IP datagram).

anycast

In anycast routing, the IP packet is sent to the topologically nearest host.

APNIC

Asia Pacific Network Information Center. Nonprofit Internet registry organization for the Asia Pacific region. The other Internet registries are currently IANA, RIPE NCC, and InterNIC.

ARIN

American Registry for Internet Numbers. A nonprofit organization that administers and registers IP numbers for the geographical areas that are currently managed by Network Solutions (InterNIC). Those areas include, but are not limited to, North America, South America, South Africa, and the Caribbean.

ARP

Address Resolution Protocol. Internet protocol that is used to map an IP address to a MAC address. Defined in RFC 826.

AS

autonomous system. A collection of networks under a common administration sharing a common routing strategy. Autonomous systems are subdivided by areas. An autonomous system must be assigned a unique 16-bit number by the IANA.

ASCII

American Standard Code for Information Interchange. An 8-bit code for character representation (7 bits plus parity).

BGP

Border Gateway Protocol. Interdomain routing protocol that replaces EGP. BGP exchanges reachability information with other BGP systems. It is defined by RFC 1163.

BID

bridge ID.

BPDU

bridge protocol data unit. Spanning Tree Protocol hello packet that is sent out at configurable intervals to exchange information among bridges in the network.

broadcast address

A special address that is reserved for sending a message to all stations. Generally, a broadcast address is a MAC destination address of all 1s.

CATV

cable television. A communication system where multiple channels of programming material are transmitted to homes using broadband coaxial cable. Formerly called Community Antenna Television.

CIDR

classless interdomain routing. Technique supported by BGP4 and based on route aggregation. CIDR allows routers to group routes to reduce the quantity of routing information carried by the core routers. With CIDR, several IP networks appear to networks outside the group as a single, larger entity. With CIDR, IP addresses and their subnet masks are written as four octets, separated by periods, followed by a forward slash and a two-digit number that represents the subnet mask.

CPE

customer premises equipment. Terminating equipment such as terminals, telephones, and modems supplied by the telephone company, installed at customer sites, and connected to the telephone company network. Can also refer to any telephone equipment residing on the customer site.

CST

Common Spanning Tree.

DDNS

Dynamic Domain Name System.

DHCP

Dynamic Host Configuration Protocol (*common term*). Provides a mechanism for allocating IP addresses dynamically so that addresses can be reused when hosts no longer need them.

DHCPv6

Dynamic Host Configuration Protocol version 6.

DNS

Domain Name System. System used on the Internet for translating names of network nodes into addresses.

DOCSIS

Data-over-Cable Service Interface Specifications. Defines technical specifications for equipment at both subscriber locations and cable operator headends.

DORA

The DORA process in DHCP works as follows:
D ----------> Discover
O ----------> Offer
R ----------> Request
A ----------> Acknowledgment

DSL

digital subscriber line (*common term*). Public network technology that delivers high bandwidth over conventional copper wiring at limited distances. There are four types of DSL: ADSL, HDSL, SDSL, and VDSL. All are provisioned via modem pairs, with one modem located at a central office and the other at the customer site. Because most DSL technologies do not use the whole bandwidth of the twisted pair, there is room remaining for a voice channel.

DTP

Dynamic Trunking Protocol.

EGP

Exterior Gateway Protocol. It's the Internet protocol for exchanging routing information between autonomous systems. Documented in RFC 904. This is not to be confused with the general term *exterior gateway protocol*. EGP is an obsolete protocol that was replaced by BGP.

EIGRP

Enhanced Interior Gateway Routing Protocol. It's the advanced version of IGRP developed by Cisco. It provides superior convergence properties and operating efficiency, and it combines the advantages of link-state protocols with those of distance vector protocols.

EMI

electromagnetic interference. It's interference by electromagnetic signals that can cause reduced data integrity and increased error rates on transmission channels.

ESP

Encapsulating Security Payload. It's a security protocol that provides data privacy services, optional data authentication, and antiplay services. ESP encapsulates the data to be protected.

EtherChannel

Developed and copyrighted by Cisco Systems. It's the logical aggregation of multiple Ethernet interfaces used to form a single higher bandwidth routing or bridging endpoint.

Ethernet

Baseband LAN specification invented by Xerox Corporation and developed jointly by Xerox, Intel, and Digital Equipment Corporation. Ethernet networks use CSMA/CD and run over a variety of cable types at 10 Mbps. Ethernet is similar to the IEEE 802.3 series of standards. It is the most commonly used LAN technology because its protocol is easy to understand, implement, manage, and maintain. It allows low-cost network implementations, provides extensive topological flexibility for network installation, and guarantees successful interconnection and operation of standards-compliant products, regardless of manufacturer.

EUI-64

EUI 64-bit format

Fast Ethernet

Any of a number of 100-Mbps Ethernet specifications. Fast Ethernet offers a speed increase 10 times that of the 10BaseT Ethernet specification while preserving such qualities as frame format, MAC mechanisms, and MTU. Such similarities allow the use of existing 10BaseT applications and network management tools on Fast Ethernet networks. Based on an extension to the IEEE 802.3 specification.

FCS

frame check sequence. Extra characters added to a frame for error control purposes. Used in HDLC, Frame Relay, and other data link layer protocols.

FDDI

Fiber Distributed Data Interface. LAN standard, defined by ANSI X3T9.5, specifying a 100-Mbps token-passing network using fiber-optic cable, with transmission distances of up to 2 km. FDDI uses a dual-ring architecture to provide redundancy.

FHRP

First Hop Redundancy Protocol. A general category of protocols that includes GLBP, HSRP, and VRRP.

FTP

File Transfer Protocol. Protocol for exchanging files over the Internet.

giaddr

gateway IP address

Gigabit Ethernet

Standard for a high-speed Ethernet, approved by the IEEE (Institute of Electrical and Electronics Engineers) 802.3z standards committee in 1996.

GLBP

Gateway Load Balancing Protocol

Hello protocol

Protocol used by OSPF systems for establishing and maintaining neighbor relationships.

hexadecimal

In computer science, hexadecimal (also called base 16) is a numbering system with a base of 16. It uses 16 distinct symbols: 0-9 to represent values zero to nine, and A, B, C, D, E, and F to represent values 10 to 15.

HSRP

Hot Standby Router Protocol. Provides high network availability and transparent network topology changes. HSRP creates a hot standby router group with a lead router that services all packets sent to the hot standby address. The lead router is monitored by other routers in the group. If it fails, one of the standby routers inherits both the lead position and the hot standby address.

HSRPv1

Hot Standby Router Protocol version 1.

IANA

Internet Assigned Numbers Authority. Organization operated under the auspices of the ISOC as a part of the IAB. IANA delegates authority for IP address-space allocation and domain-name assignment to the InterNIC and other organizations. IANA also maintains a database of assigned protocol identifiers that is used in the TCP/IP stack, including autonomous system numbers.

ICMP

Internet Control Message Protocol. Network layer Internet protocol that reports errors and provides other information that is relevant to IP packet processing. Documented in RFC 792.

ICMPv4

Internet Control Message Protocol version 4.

ICMPv6

Internet Control Message Protocol version 6 is the implementation of ICMP for IPv6. It is defined in RFC 4443.

ID

identifier *(common term)*.

IEEE

Institute of Electrical and Electronics Engineers. Professional organization whose activities include the development of communications and network standards. IEEE LAN standards are the predominant LAN standards today.

IEEE 802.1D

Electrical and Electronics Engineers (IEEE) 802.1D is the MAC Bridges standard which includes Bridging, Spanning Tree and others. It is standardized by the 802.1 working group.

IEEE 802.1Q

The IEEE 802.1Q specification establishes a standard method for tagging Ethernet frames with VLAN membership information and defines the operation of VLAN bridges that permit the definition, operation, and administration of VLAN topologies within a bridged LAN infrastructure.

IEEE 802.1s

Electrical and Electronics Engineers (IEEE) 802.1s is an IEEE standard that is inspired by the earlier Cisco proprietary MISTP implementation.

IEEE 802.1w

Electrical and Electronics Engineers (IEEE) 802.1w is an IEEE standard that is inspired by the earlier 802.1D standard. It is also known as Rapid Spanning Tree Protocol (RSTP). It provides faster convergence of STP.

IEEE 802.3ad

The IEEE standard for link aggregation for parallel links, since moved to IEEE 802.1AX.

IGP

interior gateway protocol. Internet protocol used to exchange routing information within an autonomous system. Examples of common Internet IGPs include IGRP, OSPF, and RIP.

InterNIC

Internet Network Information Center. Organization that serves the Internet community by supplying user assistance, documentation, training, registration service for Internet domain names, and other services. Formerly called NIC.

IP

Internet Protocol. Network layer protocol in the TCP/IP stack offering a connectionless internetwork service. IP provides features for addressing, type-of-service specification, fragmentation and reassembly, and security. Defined in RFC 791.

IP address

A 32-bit address assigned to hosts using TCP/IP. An IP address belongs to one of five classes (A, B, C, D, or E) and is written as 4 octets separated by periods (dotted decimal format). Each address consists of a network number, an optional subnetwork number, and a host number. The network and subnetwork numbers together are used for routing, and the host number is used to address an individual host within the network or subnetwork. A subnet mask is used to extract network and subnetwork information from the IP address. CIDR provides a new way of representing IP addresses and subnet masks. Also called an Internet address.

IPsec

IP Security. A framework of open standards that provides data confidentiality, data integrity, and data authentication between participating peers. IPsec provides these security services at the IP layer. IPsec uses IKE to handle the negotiation of protocols and algorithms based on local policy and to generate the encryption and authentication keys to be used by IPsec. IPsec can protect one or more data flows between a pair of hosts, between a pair of security gateways, or between a security gateway and a host.

IPv4

IP version 4 *(common term)*. Internet Protocol version 4 is the fourth version in the development of IP and the first version of the protocol to be widely deployed. Along with IPv6, IPv4 is at the core of standards-based internetworking methods of the Internet. IPv4 is still used to route most traffic across the Internet. IPv4 is a connectionless protocol for use on packet-switched link layer networks (for example, Ethernet). It operates on a best-effort delivery model in that it does not guarantee delivery and does not assure proper sequencing or avoidance of duplicate delivery.

IPv6

IP version 6 *(common term)*. Replacement for the current version of IP (version 4). IPv6 includes support for flow ID in the packet header, which can be used to identify flows. Formerly called IPng (next generation).

ISL

Inter-Switch Link. Cisco proprietary protocol that maintains VLAN information as traffic flows between switches and routers.

ISP

Internet service provider. Company that provides Internet access to other companies and individuals.

LACNIC

Latin American and Caribbean Network Information Center.

LACP

Link Aggregation Control Protocol.

LAN

local-area network. High-speed, low-error data network covering a relatively small geographic area (up to a few thousand meters). LANs connect workstations, peripherals, terminals, and other devices in a single building or other geographically limited area. LAN standards specify cabling and signaling at the physical and data link layers of the OSI model. Ethernet, FDDI, and Token Ring are widely used LAN technologies.

LIR

local Internet registry.

LSA

link-state advertisement. A broadcast packet used by link-state protocols that contains information about neighbors and path costs. LSAs are used by the receiving routers to maintain their routing tables. Sometimes called an LSP.

MAC

Media Access Control. The lower of the two sublayers of the data link layer that is defined by the IEEE. The MAC sublayer handles access to shared media, such as whether token passing or contention will be used.

MAC address

a standardized data link layer address that is required for every port or device that connects to a LAN. Other devices in the network use these addresses to locate specific ports in the network and to create and update routing tables and data structures. A MAC address is 6 bytes long and is controlled by the IEEE. It is also known as a hardware address, MAC layer address, and physical address.

MISTP

Multi-Instance STP.

MST

Multiple Spanning Tree.

MSTP

Multiple Spanning Tree Protocol.

MTU

maximum transmission unit. The maximum packet size, in bytes, that a particular interface can handle.

multicast

single packets that are copied by the network and sent to a specific subset of network addresses. These addresses are specified in the Destination Address field.

NAT

Network Address Translation. A mechanism for reducing the need for globally unique IP addresses. NAT allows an organization with addresses that are not globally unique to connect to the Internet by translating these addresses into globally routable address space. Also known as Network Address Translator.

NVI

NAT Virtual Interface.

OSI

Open Systems Interconnection. International standardization program created by ISO and ITU-T to develop standards for data networking that facilitate multivendor equipment interoperability.

OSPF

Open Shortest Path First.
Link-state, hierarchical IGP routing algorithm proposed as a successor to RIP in the Internet community. OSPF features include least-cost routing, multipath routing, and load balancing. OSPF was derived from an early version of the IS-IS protocol.

OSPFv2

Open Shortest Path First version 2.

OSPFv3

Open Shortest Path First version 3.

PAgP

Port Aggregation Protocol.

PAT

port address translation. Translation method that allows the user to conserve addresses in the global address pool by allowing source ports in TCP connections or UDP conversations to be translated. Different local addresses then map to the same global address, with port translation providing the necessary uniqueness. When translation is required, the new port number is picked out of the same range as the original following the convention of Berkeley Standard Distribution (SD).

PVST+

Per VLAN Spanning Tree Plus. Support for dot1q trunks to map multiple spanning trees to a single spanning tree.

QoS

quality of service. Measure of performance for a transmission system that reflects its transmission quality and service availability.

Rapid PVST+

Rapid Per VLAN Spanning Tree Plus.

RFC

Request for Comments. Document series that is used as the primary means for communicating information about the Internet. Some RFCs are designated by the IAB as Internet standards. Most RFCs document protocol specifications, such as Telnet and FTP, but some RFCs are humorous or historical. RFCs are available online from numerous sources.

RFI

radio frequency interference. Radio frequencies that create noise that interferes with information being transmitted across unshielded copper cable.

RIP

Routing Information Protocol. A distance-vector routing protocol that uses hop count as a routing metric.

RIPE NCC

Réseaux IP Européens Network Coordination Centre. The RIPE NCC is one of five Regional Internet Registries (RIRs), providing Internet resource allocations, registration services, and coordination activities that support the operation of the Internet globally.

RIPng

Routing Information Protocol next generation.

RIPv1

Routing Information Protocol version 1.

RIPv2

Routing Information Protocol version 2.

RIR

regional Internet registry.

root bridge

Exchanges topology information with designated bridges in a spanning-tree implementation to notify all other bridges in the network when topology changes are required. This prevents loops and provides a measure of defense against link failure.

RSTP

Rapid Spanning Tree Protocol.

RSTP+

Rapid Spanning Tree Protocol Plus.

SLA

Site-Level Aggregator.

SLAAC

stateless address autoconfiguration. Allows IPv6 hosts to automatically acquire a valid IPv6 address when connected to an IPv6 network using the Neighbor Discovery Protocol.

SPF

Shortest Path First. Routing algorithm that iterates on the length of path to determine a shortest-path spanning tree. Commonly used in link-state routing algorithms. Sometimes called Dijkstra's algorithm.

STP

Spanning Tree Protocol. Bridge protocol that uses the spanning-tree algorithm, enabling a learning bridge to dynamically work around loops in a network topology by creating a spanning tree. Bridges exchange BPDU messages with other bridges to detect loops, and then remove the loops by shutting down selected bridge interfaces. Refers to both the IEEE 802.1 Spanning Tree Protocol standard and the earlier Digital Equipment Corporation Spanning Tree Protocol upon which it is based. The IEEE version supports bridge domains and allows the bridge to construct a loop-free topology across an extended LAN. The IEEE version generally is preferred over the Digital version.

STP

shielded twisted-pair. Two-pair wiring medium used in a variety of network implementations. STP cabling has a layer of shielded insulation to reduce EMI.

SVI

switch virtual interface.

syslog

system logging.

TCP

Transmission Control Protocol. Connection-oriented transport layer protocol that provides reliable full-duplex data transmission. TCP is part of the TCP/IP protocol stack.

TCP/IP

Transmission Control Protocol/Internet Protocol. Common name for the suite of protocols developed by the U.S. DoD in the 1970s to support the construction of worldwide internetworks. TCP and IP are the two best-known protocols in the suite.

Telnet

standard terminal emulation protocol in the TCP/IP protocol stack. Telnet is used for remote terminal connection, enabling users to log into remote systems and use resources as if they were connected to a local system. Telnet is defined in RFC 854.

Token Ring

Token-passing LAN that was developed and supported by IBM. Token Ring runs at 4 or 16 Mbps over a ring topology. Similar to IEEE 802.5.

ToS

type of service.

TTL

Time to Live. A mechanism that limits the lifespan or lifetime of data in a computer or network.

UDP

User Datagram Protocol. Connectionless transport layer protocol in the TCP/IP protocol stack. UDP is a simple protocol that exchanges datagrams without acknowledgments or guaranteed delivery, requiring that error processing and retransmission be handled by other protocols. UDP is defined in RFC 768.

unicast

message sent to a single network destination.

UTP

unshielded twisted-pair. Four-pair wire medium used in a variety of networks. UTP does not require the fixed spacing between connections that is necessary with coaxial-type connections. Five types of UTP cabling are commonly used: Category 1, Category 2, Category 3, Category 4, and Category 5.

VID

VLAN ID. The identification of the VLAN, which is used by the standard IEEE 802.1Q. Being 12 bits, it allows for the identification of 4096 VLANs.

VLAN

virtual LAN. Group of devices on one or more LANs that are configured (using management software) so that they can communicate as if they were attached to the same wire, when in fact they are located on a number of different LAN segments. Because VLANs are based on logical instead of physical connections, they are extremely flexible.

VLSM

variable-length subnet mask. Capability to specify a different subnet mask for the same network number on different subnets. VLSM can help optimize available address space.

VPN

virtual private network. Enables IP traffic to travel securely over a public TCP/IP network by encrypting all traffic from one network to another. A VPN uses tunneling to encrypt all information at the IP level.

VRRP

Virtual Router Redundancy Protocol.

VTP

VLAN Trunking Protocol.

WAN

wide-area network. Data communications network that serves users across a broad geographic area and often uses transmission devices provided by common carriers. Frame Relay, SMDS, and X.25 are examples of WANs.

WiMAX

Worldwide Interoperability for Microwave Access (WiMAX) is a family of wireless communications standards designed to provide up to 1 Gbit/s data rates.

