CCNAX

Interconnecting Cisco Networking Devices: Accelerated

Student Guide Volume 3

Version 3.0

Part Number: 97-3638-01

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 6: Troubleshooting Basic Connectivity

Introduction

Here you will learn how to troubleshoot end-to-end connectivity in an IPv4 network and connectivity in an IPv6 network.

Lesson 1: Troubleshooting IPv4 Network Connectivity

Introduction

Various customers have called CCS with complaints involving network connectivity problems, and several trouble tickets have been created. Bob has assigned all the network connectivity trouble tickets to you.

Troubleshooting Guidelines

It is impossible to write a set of troubleshooting procedures that will solve any IP connectivity problem. The troubleshooting process can be guided by structured methods, but the exact steps that are taken at each point along the way cannot be prescribed because they depend on many different factors. Each network is different, each problem is different, and the skill set and experience of each engineer that is involved in a troubleshooting process are different.

When end-to-end connectivity is not operational, the user will inform the network administrator. The administrator will start the troubleshooting process, as shown in the figure.

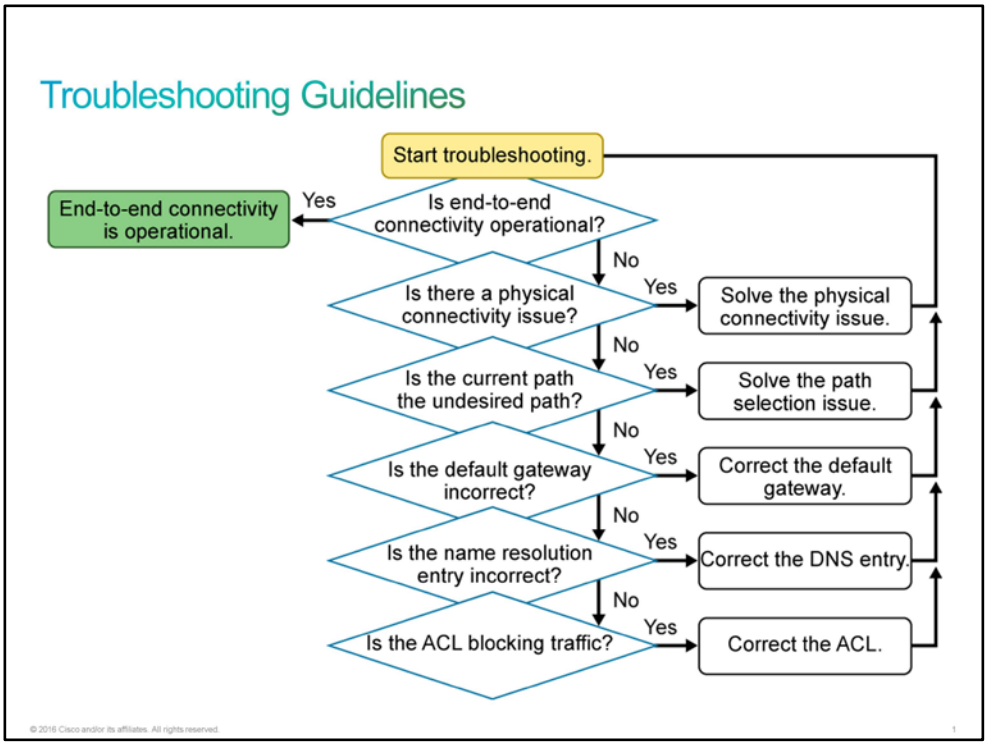

When there is no end-to-end connectivity, the following are some items that you should investigate:

- Check the cables, because there might be a faulty cable or interface. This is a link by link test. You may need to check each cable that lies in the packet path (the path between the source and destination devices that are experiencing connectivity problems).
- Make sure that the devices are determining the correct path from the source to the destination. Manipulate the routing information, if needed.
- Verify that the default gateway is correct.
- Verify that the name resolution settings are correct.
- Verify that there are no ACLs that are blocking traffic.

After every failed troubleshooting step, you should provide a solution to make the step successful. The outcome of this process is operational end-to-end connectivity.

Discovery 30: Use Troubleshooting Tools

Introduction

In this discovery, you will learn how to use some basic commands for verifying end-to-end connectivity in an IP network. 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 routers. There are no issues to troubleshoot with the network. The goal of this discovery is not to complete troubleshooting tasks but to become familiar with some basic troubleshooting tools.

Topology

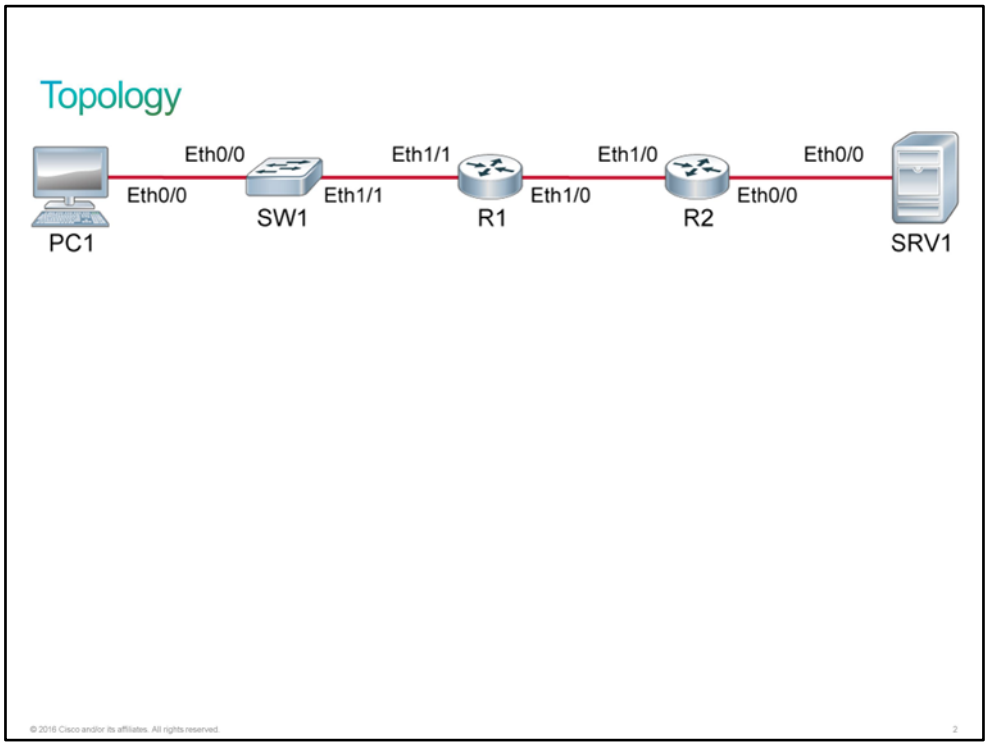

Job Aid

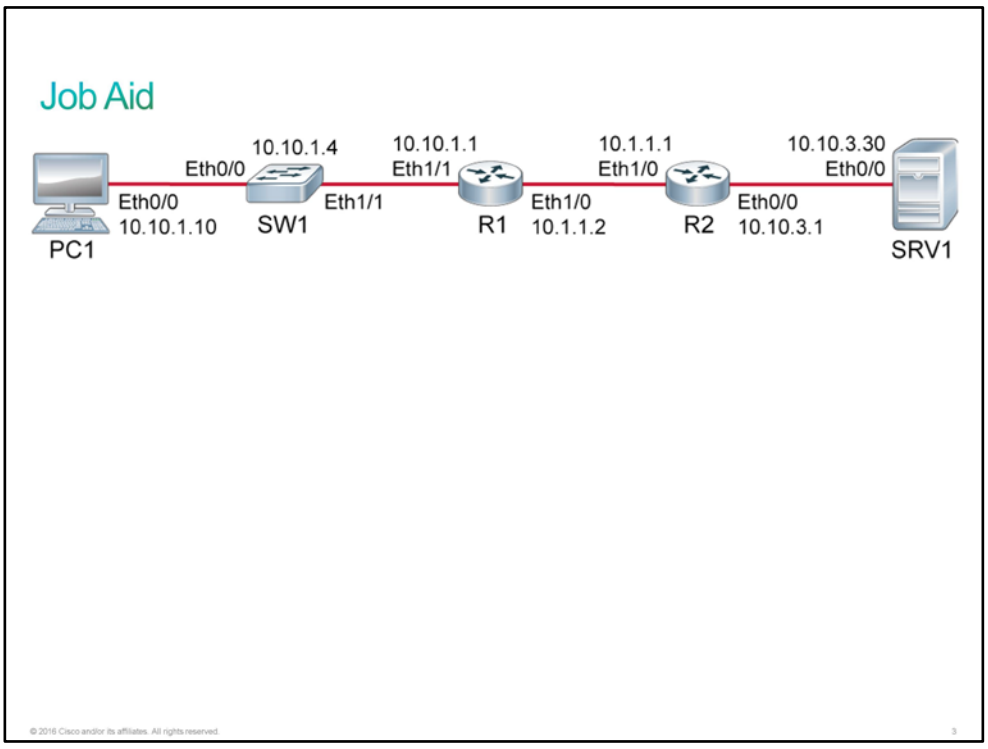

The configuration is as follows:

- All devices have their basic configurations in place, including hostnames and IP addresses.
- RIP is configured on R1 and R2.

Device Details

Device	Interface	Neighbor	IP Address
PC1	Ethernet0/0	SW1	10.10.1.10/24
SRV1	Ethernet0/0	R2	10.10.3.30/24
SW1	VLAN 1	—	10.10.1.4/24
SW1	Ethernet0/0	PC1	—
SW1	Ethernet1/1	R1	—
R1	Ethernet1/1	SW1	10.10.1.1/24
R1	Ethernet1/0	R2	10.1.1.2/30
R2	Ethernet1/0	R1	10.1.1.1/30
R2	Ethernet0/0	SRV1	10.10.3.1/24

Note	The 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: Use Troubleshooting Tools

Activity

Complete the following steps:

Step 1 Access the console of PC1. Ping SRV1 by its IP address.

The IP address of SRV1 is 10.10.3.30. You can verify this information in the Job Aid section.

```
PC1# 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 80 percent (4/5), round-trip min/avg/max = 1/1/1 ms
```

It is common for the first one or two probes of a ping attempt to time out if there are devices in the path that do not currently have ARP cache entries for their peers. When all ARP caches are properly populated, the ping attempts should be consistently successful.

Step 2 Attempt to ping the address 10.10.3.40. This address is on a valid subnet, but there is no host that is using the address.

Remember to take advantage of the IOS command history feature. It is easier to press the Page Up key and edit the previous command than it is to type this command.

```
PC1# ping 10.10.3.40
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.10.3.40, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
```

If there is no response to the ICMP echo request within the timeout interval, the IOS ping displays the period (.) character.

Step 3 Attempt to ping the address 10.10.4.40. This address is on a nonexistent subnet.

```
PC1# ping 10.10.4.40
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.10.4.40, timeout is 2 seconds:
U.U.U
Success rate is 0 percent (0/5)
```

In this case, because the network did not exist in the routing table of R1, R1 returned an ICMP unreachable error message to PC1. As a result, the **ping** command displays the "U" character. The difference between a timeout and an explicit unreachable message can be significant for troubleshooting.

Using IP SLA for Troubleshooting

Instead of using **ping** manually, you can use an IP SLA ICMP echo test to test the availability of far-end devices. The far-end device can be any device with IP capabilities—a router, switch, PC, server, and so on.

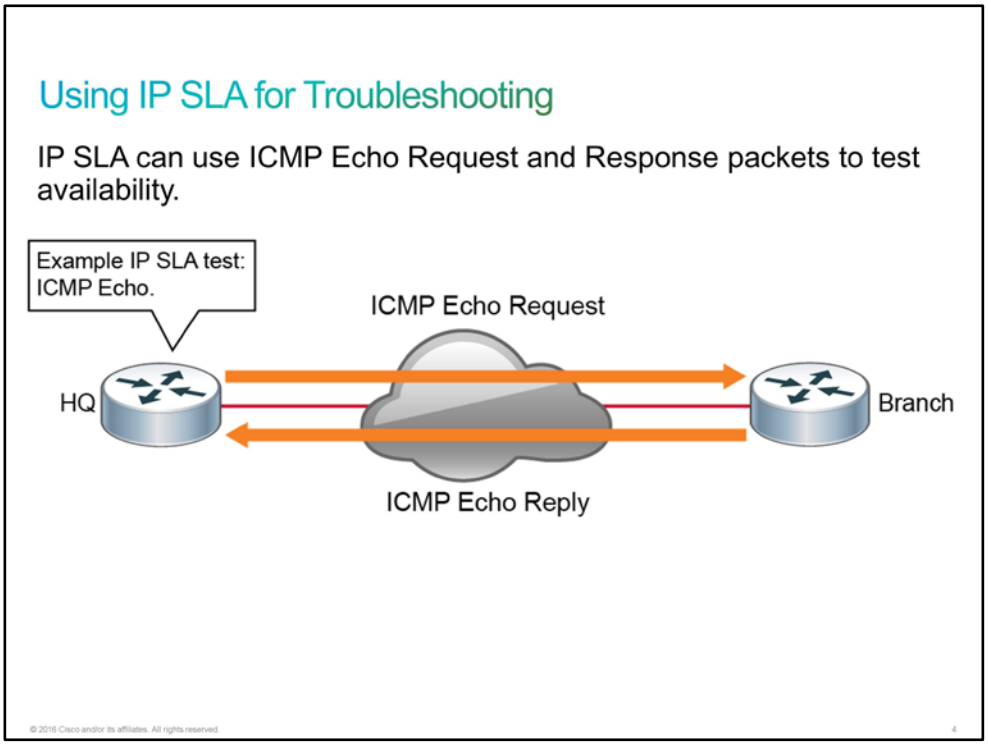

There are several common functions for the IP SLA measurements:

- Edge-to-edge network availability monitoring
  - For example, packet loss statistics
- Network performance monitoring and network performance visibility
  - For example, network latency and response time
- Troubleshooting of network operation
  - For example, end-to-end network connectivity

The ICMP Echo is only one of the available IP SLA tests. You can have multiple IP SLA operations (measurements) running in a network at any given time.

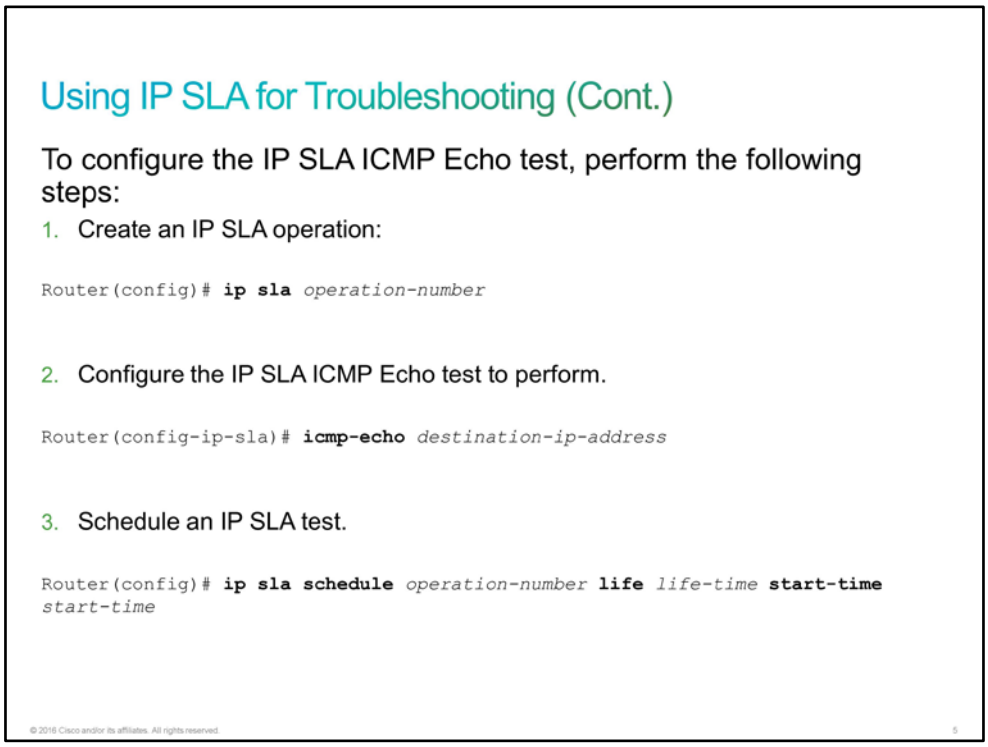

The following table describes the commands that you can use to configure an IP SLA ICMP Echo test.

Command	Description
ip sla *operation-number*	Creates an IP SLAs operation and enters the IP SLAs configuration mode.
icmp-echo *destination-ip-address*	Configures an ICMP Echo test for the specified destination.
frequency *seconds*	(Optional) Sets the rate at which a specified IP SLAs operation repeats. The range is from 1 to 604800 seconds; the default is 60 seconds.
ip sla schedule *operation-number* [**life** {**forever** | *seconds*}] [**start-time** {*hh*:*mm*[:*ss*] [*month day* | *day month*] | **pending** | **now** | **after** *hh*:*mm*:*ss*}] [**ageout** *seconds*] [**recurring**]	Configures the scheduling parameters for an individual IP SLAs operation. • With the **life** keyword, you set how long the IP SLA test will run. If you choose **forever**, the test will run until you manually remove it. By default, the IP SLA test will run for 1 hour. • With the **start-time** keyword, you will set when the IP SLA test should start. You can start the test right away by issuing the **now** keyword, or you can configure a delayed start. • With the **ageout** keyword, you can control how long the collected data is kept. • With the **recurring** keyword, you can schedule a test to run periodically—for example, at the same time each day.

Note	After an IP SLA test is scheduled to run, you will not be able to modify it.

Step 4 Access the console of R1 and configure an IP SLA ICMP Echo test to the SRV1 IP address (10.10.3.30).

Define the IP SLA with the number 1 and set the frequency to 10 seconds.

```
R1# conf t
R1(config)# ip sla 1
R1(config-ip-sla)# icmp-echo 10.10.3.30
R1(config-ip-sla-echo)# frequency 10
R1(config-ip-sla-echo)# exit
```

Step 5 Schedule IP SLA 1 on R1 to perform an ICMP Echo test forever and to start running now.

```
R1(config)# ip sla schedule 1 life forever start-time now
R1(config)# exit
```

Verifying IP SLA Operation

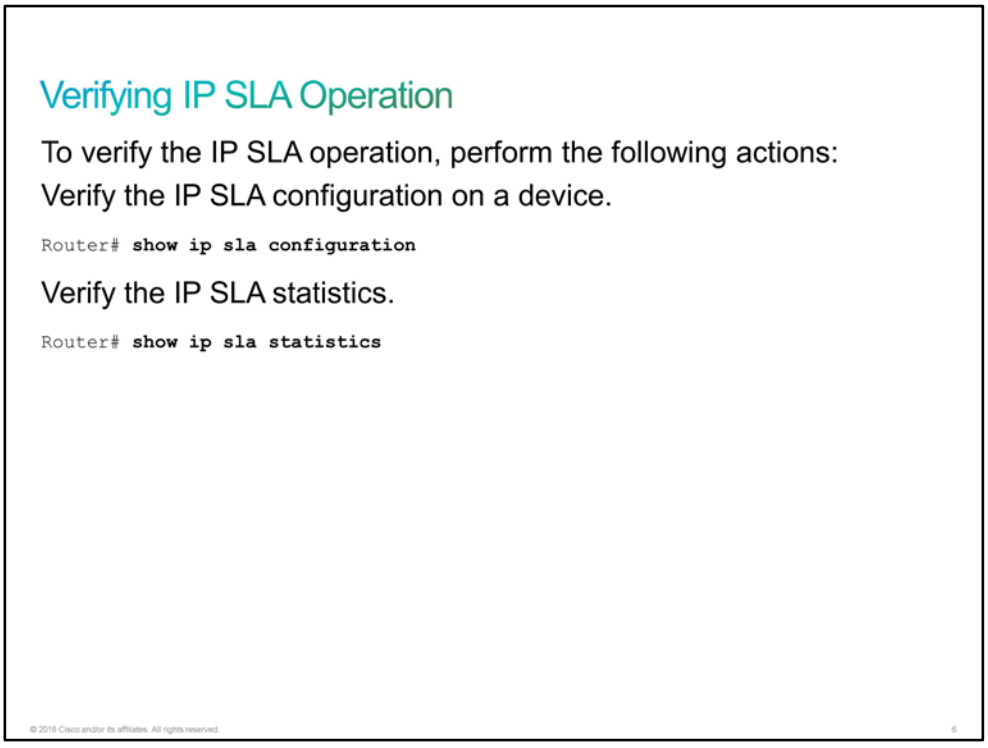

Use the **show ip sla configuration** command to verify the configured parameters. If you want to investigate the results of the test, you should use the **show ip sla statistics** command.

Step 6 On R1, verify the IP SLA configuration.

R1 should have an ICMP Echo test configured to the SRV1 IP address. The test should run every 10 seconds and should be scheduled to run forever.

```
R1# show ip sla configuration
IP SLAs Infrastructure Engine-III
Entry number: 1
Owner:
Tag:
Operation timeout (milliseconds): 5000
Type of operation to perform: icmp-echo
Target address/Source address: 10.10.3.30/0.0.0.0
Type Of Service parameter: 0x0
Request size (ARR data portion): 28
Verify data: No
Vrf Name:
Schedule:
   Operation frequency (seconds): 10  (not considered if randomly scheduled)
   Next Scheduled Start Time: Start Time already passed
   Group Scheduled : FALSE
   Randomly Scheduled : FALSE
   Life (seconds): Forever
   Entry Ageout (seconds): never
   Recurring (Starting Everyday): FALSE
   Status of entry (SNMP RowStatus): Active
<... output omitted ...>
```

Step 7 On R1, verify the IP SLA statistics to verify that SRV1 is reachable.

```
R1# show ip sla statistics
IPSLAs Latest Operation Statistics

IPSLA operation id: 1
        Latest RTT: 1 milliseconds
Latest operation start time: 15:08:23 PST Thu Nov 5 2015
Latest operation return code: OK
Number of successes: 91
Number of failures: 0
Operation time to live: Forever
```

The IP SLA 1 test on R1 has been successfully performed 91 times and the test never failed. Note that these numbers may differ in your output.

Step 8 Execute a **traceroute** command that targets the SRV1 IP address.

```
PC1# traceroute 10.10.3.30
Type escape sequence to abort.
Tracing the route to SRV1 (10.10.3.30)
VRF info: (vrf in name/id, vrf out name/id)
  1 10.10.1.1 1 msec 0 msec 0 msec
  2 10.1.1.1 1 msec 0 msec 1 msec
  3 SRV1 (10.10.3.30) 0 msec *  1 msec
```

The traceroute displays the "near-side" IP address of every router in the path to the destination IP address.

The traceroute attempts to display both the DNS hostname and the IP address of each hop in the path. This information is evident in the last line in the example output. There is no DNS service in the virtual lab environment, but a static IP host entry for SRV1 has been set in the PC1 configuration.

Note: In the emulated virtual lab environment, it is normal for the middle probe to the final destination to time out.

Step 9 Attempt a traceroute to the nonexistent address 10.10.3.40. Because the destination cannot be reached, the traceroute will continue to send probes with consistently higher TTL values.

The traceroute will terminate after 30 hops. However, you can interrupt it at any time by pressing the Ctrl-Shift-6 keys simultaneously.

```
PC1# traceroute 10.10.3.40
Type escape sequence to abort.
Tracing the route to 10.10.3.40
VRF info: (vrf in name/id, vrf out name/id)
  1 10.10.1.1 1 msec 0 msec 1 msec
  2 10.1.1.1 0 msec 0 msec 0 msec
  3  *  *  *
  4  *  * <Ctrl-Shift-6>
```

A traceroute sends a series of IP probe packets. It first sends three probes with a TTL = 1. The probes will reach the first hop, which will decrement the TTL to 0. Because the first hop is not allowed to forward the packet with an expired TTL, it returns ICMP unreachable messages, which the traceroute program processes. The traceroute will then send three probes with a TTL = 2, which will make it to the second hop. It continues to increase the TTL until the final destination responds.

Step 10 Verify Telnet reachability for SRV1. Verify that the prompt shows SRV1, then terminate the Telnet session with the **exit** command.

Log in with the password **Cisco123**.

```
PC1# telnet 10.10.3.30
Trying 10.10.3.30 ... Open


User Access Verification

Password:
SRV1> exit

[Connection to 10.10.3.30 closed by foreign host]
PC1#
```

Step 11 Verify that SRV1 is running an HTTP service on the TCP port 80 by using the **telnet** command.

Because you cannot mimic the behavior of a web browser from the Telnet CLI, enter a few random characters and press **Enter**. SRV1 returns an error message and terminates the connection.


```
PC1# telnet 10.10.3.30 80
Trying 10.10.3.30, 80 ... Open
aaa <Enter>
HTTP/1.1 400 Bad Request
Date: Thu, 05 Nov 2015 12:42:11 GMT
Server: cisco-IOS
Accept-Ranges: none

400 Bad Request
[Connection to 10.10.3.30 closed by foreign host]
PC1#
```

Remember that Telnet uses TCP to test connectivity. By default, it will connect to port 23, but you can also specify other ports.

Step 12 Demonstrate that SRV1 is not running an FTP service on TCP port 21 by using the **telnet** command.

```
PC1# telnet 10.10.3.30 21
Trying 10.10.3.30, 21 ...
% Connection refused by remote host

PC1#
```

Step 13 Display the ARP cache on PC1, verifying that it has an entry that associates the IP address and MAC address of its default gateway.

The default gateway is the IP address of the Ethernet1/1 interface (10.10.1.1) of R1.

```
PC1# show arp
Protocol  Address          Age (min)  Hardware Addr   Type   Interface
Internet  10.10.1.1              19   aabb.cc00.4511  ARPA   Ethernet0/0
Internet  10.10.1.10              -   aabb.cc00.4200  ARPA   Ethernet0/0
```

Note: The MAC addresses might differ in your output.

Step 14 Display the ARP cache on R1, verifying that it has an entry that associates the IP address and MAC address of PC1.

```
R1# show arp
Protocol  Address          Age (min)  Hardware Addr   Type   Interface
Internet  10.1.1.1               39   aabb.cc00.4601  ARPA   Ethernet1/0
Internet  10.1.1.2                -   aabb.cc00.4501  ARPA   Ethernet1/0
Internet  10.10.1.1               -   aabb.cc00.4511  ARPA   Ethernet1/1
Internet  10.10.1.10             22   aabb.cc00.4200  ARPA   Ethernet1/1
```

Note: The MAC addresses might differ in your output.

Step 15 Access the console of SW1 and display its MAC address table. Observe the switch ports that are associated with the MAC addresses of PC1 and R1.

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

Vlan    Mac Address       Type        Ports
----    -----------       --------    -----
   1    aabb.cc00.4200    DYNAMIC     Et0/0
   1    aabb.cc00.4511    DYNAMIC     Et1/1
Total Mac Addresses for this criterion: 2
```

Note: The MAC addresses might differ in your output.

This is the end of the discovery lab.

Troubleshooting Physical Connectivity Issue

Inevitably, troubleshooting processes involve a component of hardware troubleshooting. There are three main categories of issues that could be the cause of a failure on the network: hardware failures, software failures (bugs), and configuration errors. A fourth category might be performance problems, but performance problems are more a symptom than a cause of a problem.

After you have used the **ping** and **traceroute** utilities to determine that a network connectivity problem exists and where it exists, check to see if there are physical connectivity issues before you get involved in more complex troubleshooting. You could spend hours troubleshooting a situation only to find that a network cable is loose or malfunctioning.

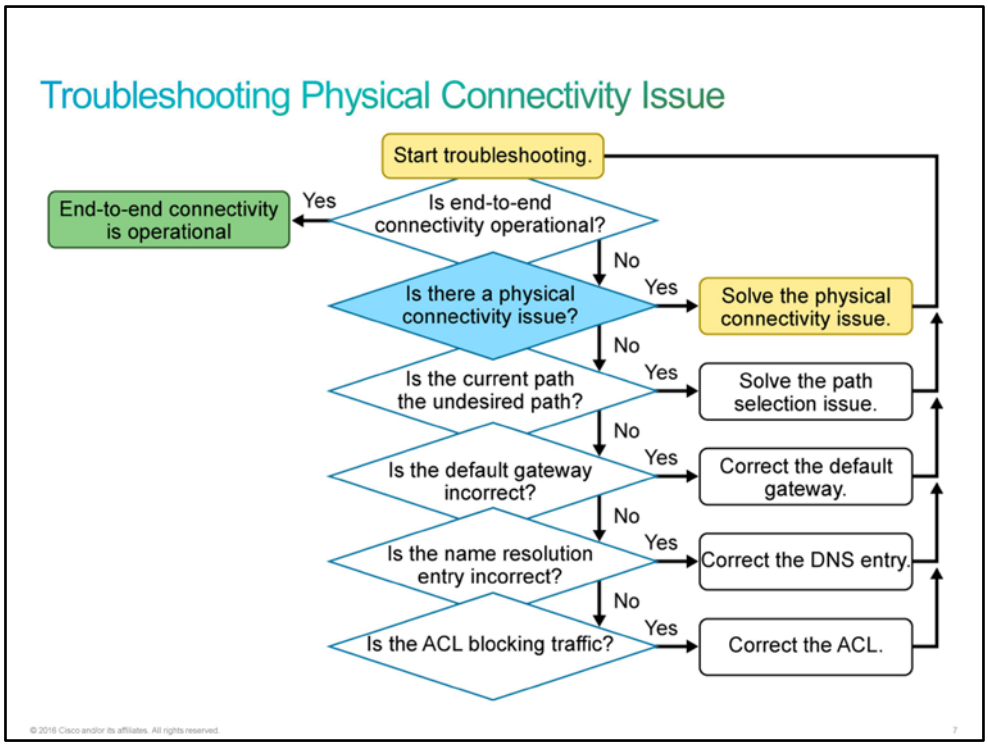

If you have physical access to devices that you suspect are causing network problems, you can save troubleshooting time by looking at the port LEDs. The port LEDs show the link status and can indicate an error condition. If a link light for a port is not on, make sure that both ends of the cable are plugged into the correct ports.

The interfaces that the traffic passes through are another component that is always worth verifying when you are troubleshooting performance-related issues and you suspect the hardware to be at fault. Usually, the interfaces are one of the first things that you would verify while tracing the path between devices.

The output of the **show interface** command lists these important statistics that should be checked. The first line of the output from this command tells you whether an interface is up or down.

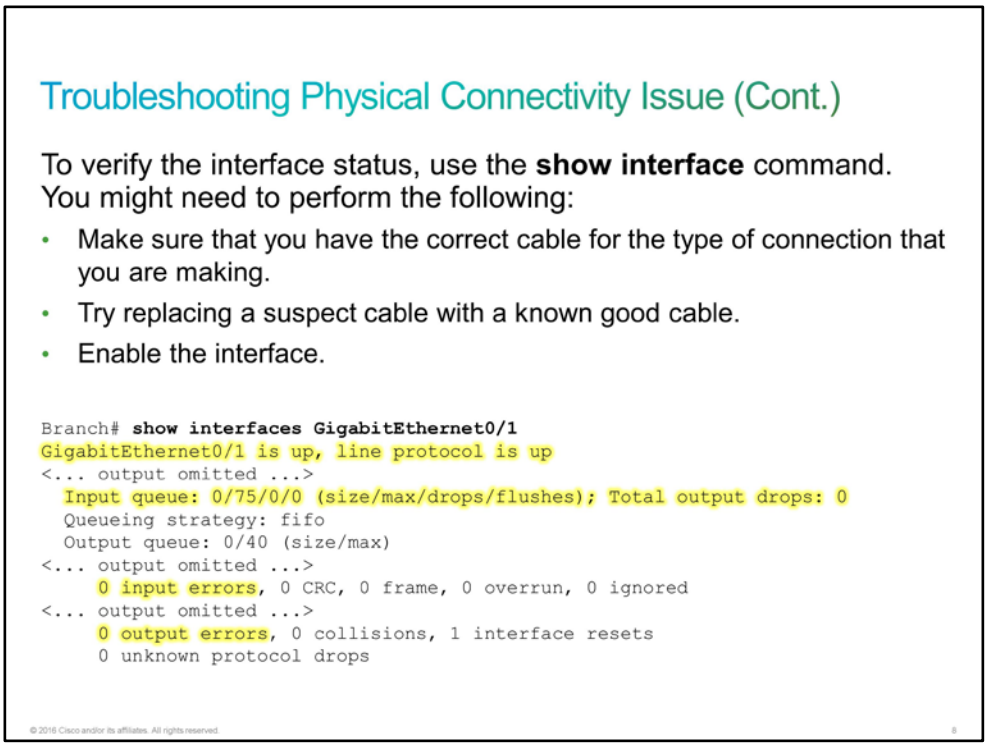

The output of the **show interface** command also displays the following important statistics:

- **Input queue drops:** Input queue drops (and the related ignored and throttle counters) signify the fact that at some point more traffic was delivered to the router than it could process. This situation does not necessarily indicate a problem because it could be normal during traffic peaks. However, it could be an indication that the CPU cannot process packets in time. So if this number is consistently high, you should try to determine at which moments these counters are increasing and how this increase relates to the CPU usage.
- **Output queue drops:** Output queue drops indicate that packets were dropped due to a congestion on the interface. Seeing output drops is normal at any point where the aggregate input traffic is higher than the output traffic. During traffic peaks, the packets are dropped if traffic is delivered to the interface faster than the interface can send it out. However, although this setting is considered normal behavior, it leads to packet drops and queuing delays, so applications that are sensitive to packet drops and queuing delays, such as VoIP, might suffer from performance issues. Consistent output drops might indicate that you need to implement an advanced queuing mechanism to provide good QoS to each application.
- **Input errors:** Input errors indicate errors that are experienced during the reception of the frame, such as CRC errors. High numbers of CRC errors could indicate cabling problems, interface hardware problems, or, in an Ethernet-based network, duplex mismatches.
- **Output errors:** Output errors indicate errors, such as collisions, during the transmission of a frame. In most Ethernet-based networks, full-duplex transmission is the norm and half-duplex transmission is the exception. In full-duplex transmission, operation collisions cannot occur. Therefore, collisions, especially late collisions, often indicate duplex mismatches.

A common cause of interface errors is mismatched duplex setting between two ends of an Ethernet link. Most Ethernet links today operate in the full-duplex mode. Also, point-to-point Ethernet links should always run in the full-duplex mode. While collisions were formerly seen as normal occurrences for an Ethernet link, collisions today often indicate that duplex negotiation has failed and that the link is not operating in the correct duplex mode. The half-duplex mode is relatively rare today and you can typically see it in environments that use hubs. However, half duplex on both ends of a connection still performs better than a duplex mismatch.

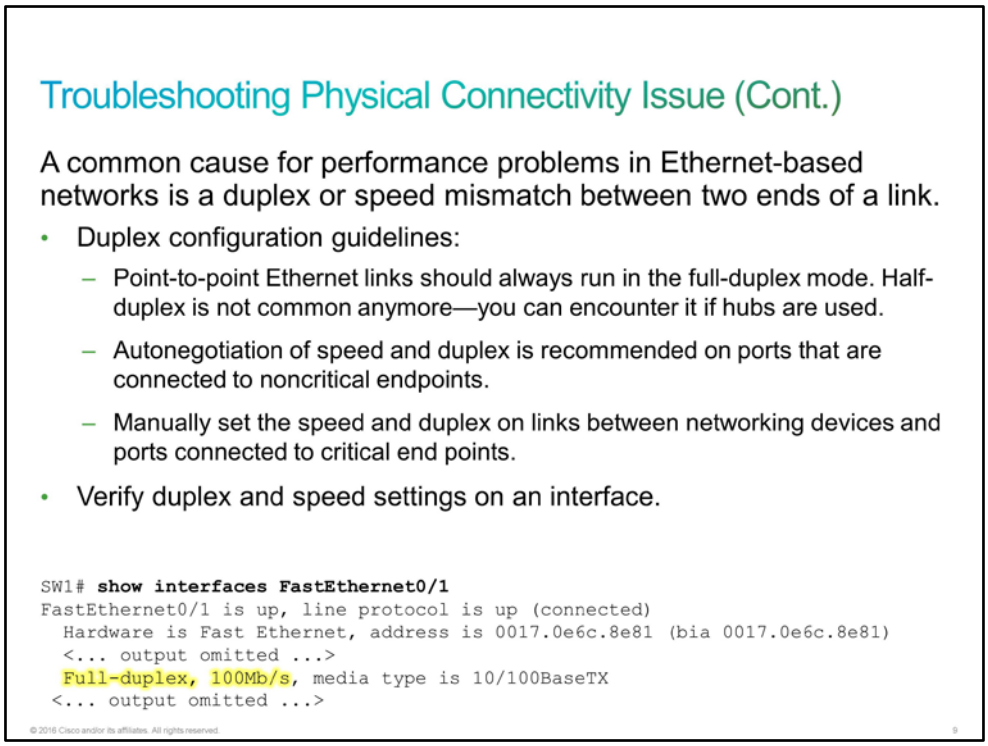

The IEEE 802.3ab Gigabit Ethernet standard mandates the use of autonegotiation for speed and duplex. Also, although autonegotiation is not mandatory, practically all Fast Ethernet NICs also use it by default. The use of autonegotiation for speed and duplex is the current recommended practice for ports that are connected to noncritical endpoints. You should manually set the speed and duplex on links between networking devices and ports that are connected to critical endpoints, such as servers.

However, if duplex negotiation fails for some reason, you might have to set the speed and duplex manually on both ends. Typically, it would mean setting the duplex mode to full duplex on both ends of the connection.

The table summarizes possible settings of speed and duplex for a connection between a switch port and an end-device NIC. The table gives just a general idea about speed and duplex misconfiguration combinations.

Speed and Duplex Settings for End-Device NIC and Switch Connections

Configuration NIC (Speed, Duplex)	Configuration Switch (Speed, Duplex)	Resulting NIC (Speed, Duplex)	Resulting Switch (Speed, Duplex)	Comments
AUTO	AUTO	1000 Mbps, full duplex	1000 Mbps, full duplex	Assuming that the maximum capability of a Cisco Catalyst switch and NIC is 1000 Mbps, full duplex.

Configuration NIC (Speed, Duplex)	Configuration Switch (Speed, Duplex)	Resulting NIC (Speed, Duplex)	Resulting Switch (Speed, Duplex)	Comments
1000 Mbps, full duplex	AUTO	1000 Mbps, full duplex	1000 Mbps, full duplex	A link is established, but the switch does not see any autonegotiation information from the NIC. Because Cisco Catalyst switches support only a full-duplex operation with 1000 Mbps, they default to full duplex. This change happens only when operating at 1000 Mbps.
AUTO	1000 Mbps, full duplex	1000 Mbps, full duplex	1000 Mbps, full duplex	Assuming that the maximum capability of a NIC is 1000 Mbps, full duplex.
1000 Mbps, full duplex	1000 Mbps, full duplex	1000 Mbps, full duplex	1000 Mbps, full duplex	Correct manual configuration.
100 Mbps, full duplex	1000 Mbps, full duplex	No link	No link	Neither side establishes a link due to a speed mismatch.
100 Mbps, full duplex	AUTO	100 Mbps, full duplex	100 Mbps, half duplex	A duplex mismatch can result in performance issues, intermittent connectivity, and loss of communication.
AUTO	100 Mbps, full duplex	100 Mbps, half-duplex	100 Mbps, full duplex	A duplex mismatch can result in performance issues, intermittent connectivity, and loss of communication.
100 Mbps, full duplex	100 Mbps, full duplex	100 Mbps, full duplex	100 Mbps, full duplex	Correct manual configuration.
100 Mbps, half duplex	AUTO	100 Mbps, half duplex	100 Mbps, half duplex	A link is established, but the switch does not see any autonegotiation information from the NIC and defaults to half duplex when operating at 10/100 Mbps.
10 Mbps, half duplex	AUTO	10 Mbps, half duplex	10 Mbps, half duplex	A link is established, but the switch does not see FLP. It defaults to 10 Mbps, half duplex.
10 Mbps, half duplex	100 Mbps, half duplex	No link	No link	Neither side establishes a link due to a speed mismatch.
AUTO	100 Mbps, half duplex	100 Mbps, half duplex	100 Mbps, half duplex	A link is established, but the NIC does not see any autonegotiation information. It defaults to 100 Mbps, half duplex.

Configuration NIC (Speed, Duplex)	Configuration Switch (Speed, Duplex)	Resulting NIC (Speed, Duplex)	Resulting Switch (Speed, Duplex)	Comments
AUTO	10 Mbps, half duplex	10 Mbps, half duplex	10 Mbps, half duplex	A link is established, but the NIC does not see FLP. It defaults to 10 Mbps, half duplex.

Identification of Current and Desired Path

When you are sure that you have eliminated any physical connectivity issues, you can move on to more in-depth troubleshooting, such as troubleshooting routing and switching issues.

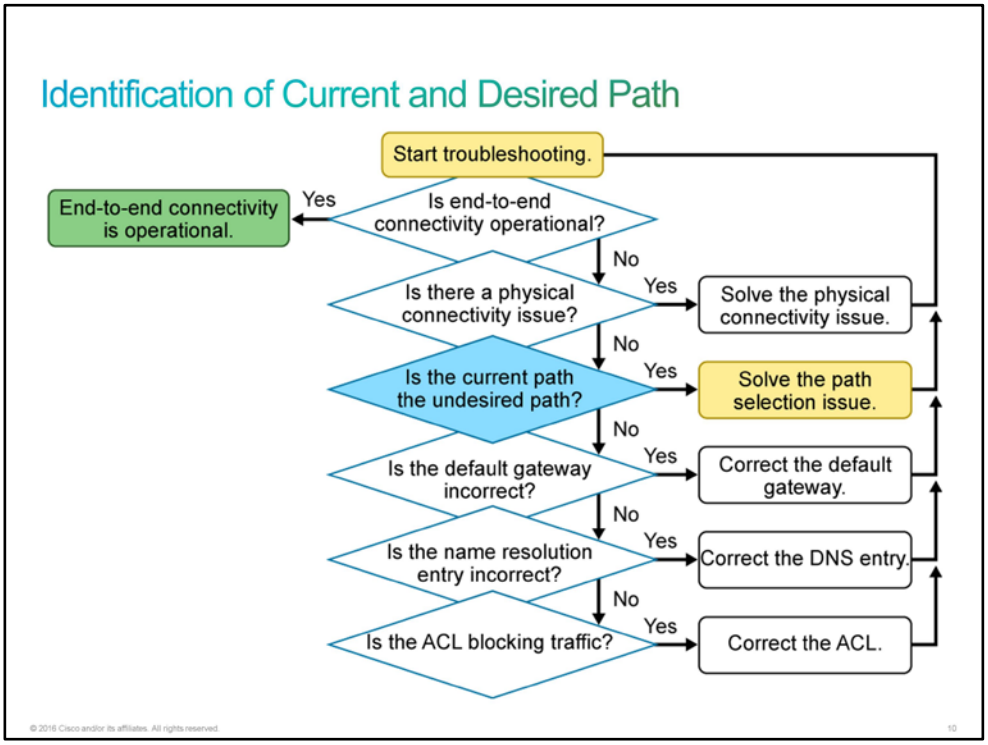

To troubleshoot Layer 3 connectivity, you need to have a good understanding of the processes that are involved in routing a packet from a host across multiple routers to the final destination.

Consider the scenario in which you are unable to send an email through the SMTP server at 172.16.1.100.

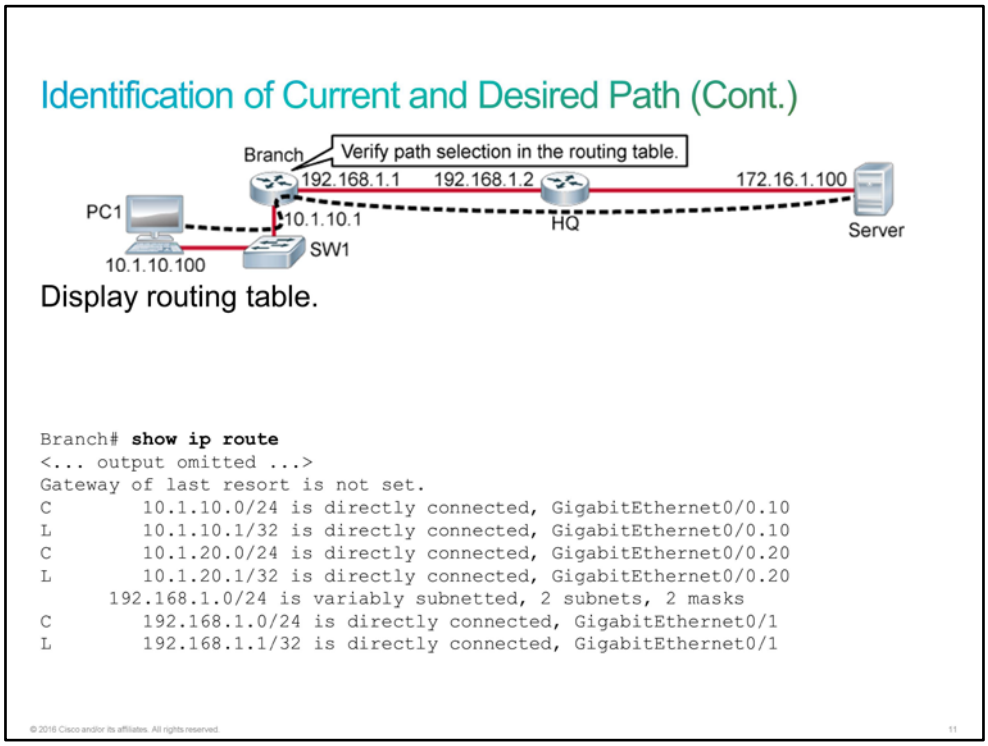

As you study the network that the figure shows, you should ask yourself these questions:

- Which decisions will PC1 make, which information does it need, and which actions will it perform to successfully send a packet that is destined for the Server to the first-hop router Branch?
- Which decisions will the router Branch make, which information does it need, and which actions will it perform to successfully send the packet from PC1 that is destined for the Server to the router Headquarters?

On the router, use the **show ip route** command to examine the routing table. In the example, the problem is that the routing table on the Branch router does not have the route to Server (172.16.1.100).

Routing Table

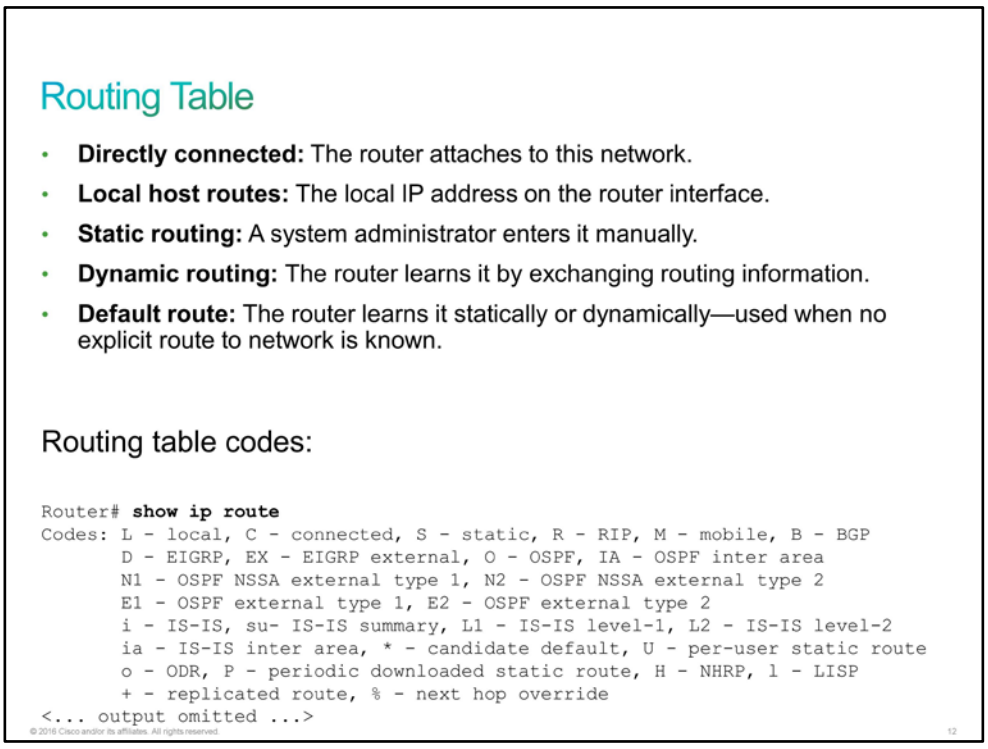

The routing tables can be populated by these methods:

- **Directly connected networks:** This entry comes from having router interfaces that are directly attached to network segments. This method is the most certain method of populating a routing table. If the interface fails or is administratively shut down, the device will remove the entry for this network from the routing table. The administrative distance is 0 and will therefore pre-empt all other entries for this destination network. Entries with the lowest administrative distance are the best, most-trusted sources.
- **Local host routes:** This entry comes from the local IP address on the router interface. The subnet mask represents the host route.
- **Static routes:** A system administrator manually enters static routes directly into the configuration of a router. The default administrative distance for a static route is 1. Therefore, the static routes will be included in the routing table, unless there is a direct connection to this network. Static routes can be an effective method for small, simple networks that do not change frequently. For larger and unstable networks, the solution with static routes does not scale.
- **Dynamic routes:** The router learns dynamic routes automatically when you configure the routing protocol and a neighbor relationship to other routers is established. The information is responsive to changes in the network and updates constantly. There is, however, always a lag between the time that a network changes and when all the routers become aware of the change. The time delay for a router to match a network change is called the convergence time. A shorter convergence time is better for users of the network. Different routing protocols perform differently in this regard. Larger networks require the dynamic routing method because there are usually many addresses and constant changes. These changes require updates to routing tables across all routers in the network, or connectivity is lost.
- **Default routes:** A default route is an optional entry that is used when no explicit path to a destination is found in the routing table. You can manually insert the default route, or it can be populated from a dynamic routing protocol.

The **show ip route** command displays the routing table in a router. The first part of the output explains the codes, presenting the letters and the associated sources of the entries in the routing table.

- **L:** Reserved for the local host route.
- **C:** Reserved for directly connected networks.
- **S:** Reserved for static routes.
- **R:** Reserved for RIP.
- **O:** Reserved for the OSPF routing protocol.
- **D:** Reserved for EIGRP. The letter "D" stands for DUAL, which is the update algorithm that EIGRP uses.

These scenarios show the different actions that a router takes if the destination address in a packet matches or does not match a routing table entry:

- If the destination address in a packet does not match an entry in the routing table, then the device uses the default route. If no default route is configured on the router, the device discards the packet.
- If the destination address in a packet matches a single entry in the routing table, the router forwards the packet through the interface that is defined in this route.
- If the destination address in a packet matches more than one entry in the routing table and the routing entries have the same prefix (network mask), the router can distribute the packets for this destination among the routes that are defined in the routing table.
- If the destination address in a packet matches more than one entry in the routing table and the routing entries have different prefixes (network masks), the router forwards the packets for this destination out of the interface that is associated with the route that has the longer prefix match.

Using SPAN for Troubleshooting

A traffic sniffer can be a valuable tool for monitoring and troubleshooting a network. Properly placing a traffic sniffer to capture a traffic flow but not interrupt it can prove challenging.

When local area networks were based on hubs, connecting a traffic sniffer was simple. When a hub receives a packet on one port, the hub sends out a copy of that packet on all ports except on the one where the hub received the packet. Therefore, the traffic sniffer that is connected a hub port could receive all traffic in the network.

Modern local networks are essentially switched networks. After a switch boots, it starts to build up a Layer 2 forwarding table based of the source MAC addresses of the different packets that the switch receives. After the switch builds this forwarding table, it then forwards traffic that is destined for a MAC address directly to the corresponding port. This way, it prevents a traffic sniffer that is connected to another port to receive the unicast traffic. The SPAN feature was therefore introduced on switches.

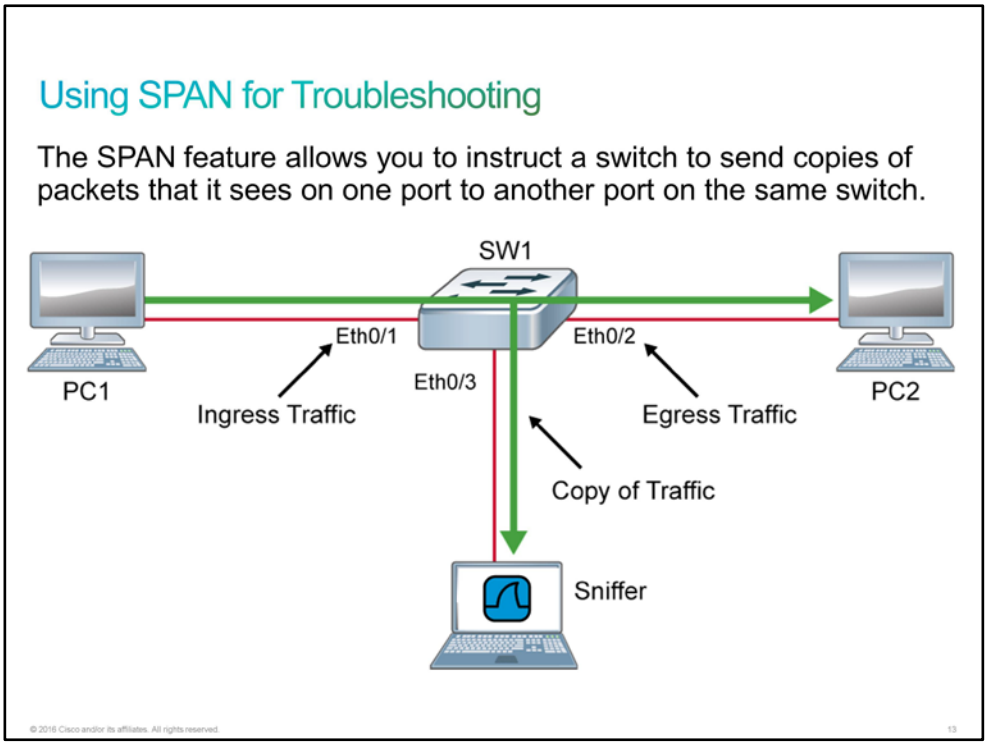

The SPAN feature allows you to analyze network traffic passing through the port and sending a copy of the traffic to another port on the switch that has been connected to a network analyzer or other monitoring device. SPAN copies the traffic that the device receives and/or sends on source ports to a destination port for analysis. SPAN does not affect the switching of network traffic on the source ports.

If you would like to analyze the traffic flowing from PC1 to PC2 on the figure, you need to specify a source port. You can either configure the interface Ethernet0/1 to capture the ingress traffic or the interface Ethernet0/2 to capture the egress traffic. Second, specify the interface Ethernet0/3 as a destination port. The traffic flowing from PC1 to PC2 will then be copied to that interface, and you will be able to analyze it with a traffic sniffer.

Configuring SPAN

With SPAN, the switch is instructed to copy all the traffic that it sends and receives on a source port to a destination port by configuring a SPAN session.

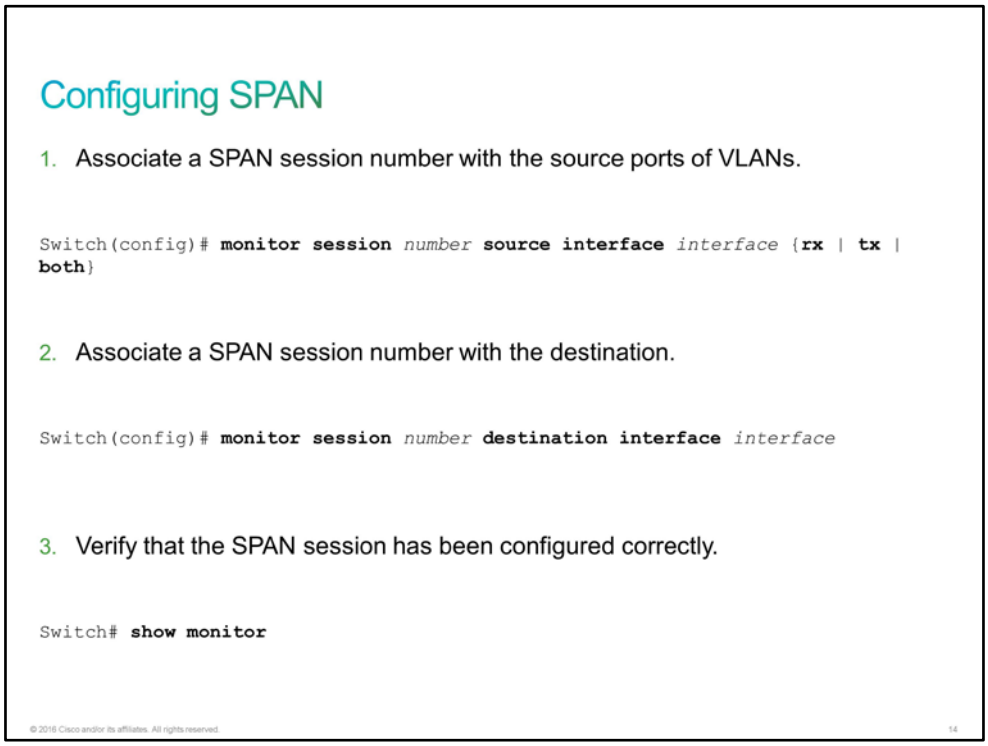

The SPAN session is identified by a session number. The first step is that you associate a SPAN session with source ports by using the **monitor session** *number* **source interface** *interface* command. You can optionally specify which traffic you want to monitor on the source interface—if you want to monitor only received traffic, use **rx** keyword, if you want to monitor only transmitted traffic then use the **tx** command. If you want to monitor both, received and transmitted traffic, use the **both** keyword. If you do not specify anything, received and transmitted traffic is captured on an interface.

Similarly, you associate a destination port with a SPAN session number by using the **monitor session** *number* **destination interface** *interface* command.

At the end, you can verify that you specified the correct source and destination ports by using the **show monitor** command.

When configuring a SPAN, you have to take notice of the following facts:

- A destination port cannot be a source port, or vice versa.
- The destination port is no longer a normal switch port—only monitored traffic passes through that port.

Configuring SPAN (Cont.)

1. Associate SPAN session number with source ports.
2. Associate SPAN session number with the destination.

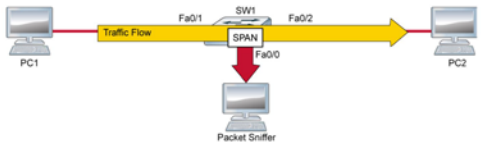


```
SW1(config)# monitor session 1 source interface FastEthernet0/2 both
SW1(config)# monitor session 1 destination interface FastEthernet0/0

SW1# show monitor
Session 1
---------
Type               : Local Session
Source Ports       :
    Both           : Fa0/2
Destination Ports  : Fa0/0
    Encapsulation  : Native
          Ingress: Disabled
```


In the example that is shown in the figure, the objective is to capture all the traffic that is sent between PC1 and PC2, both connected to the SW1. A packet sniffer is connected to port FastEthernet0/0. The switch is instructed to copy all the traffic that it sends and receives on port FastEthernet0/2 to port FastEthernet0/0 by configuring a SPAN session.

Troubleshooting Default Gateway Issues

In the absence of a detailed route on the router or an incorrect default gateway on the host, communication between two endpoints in different networks will not work.

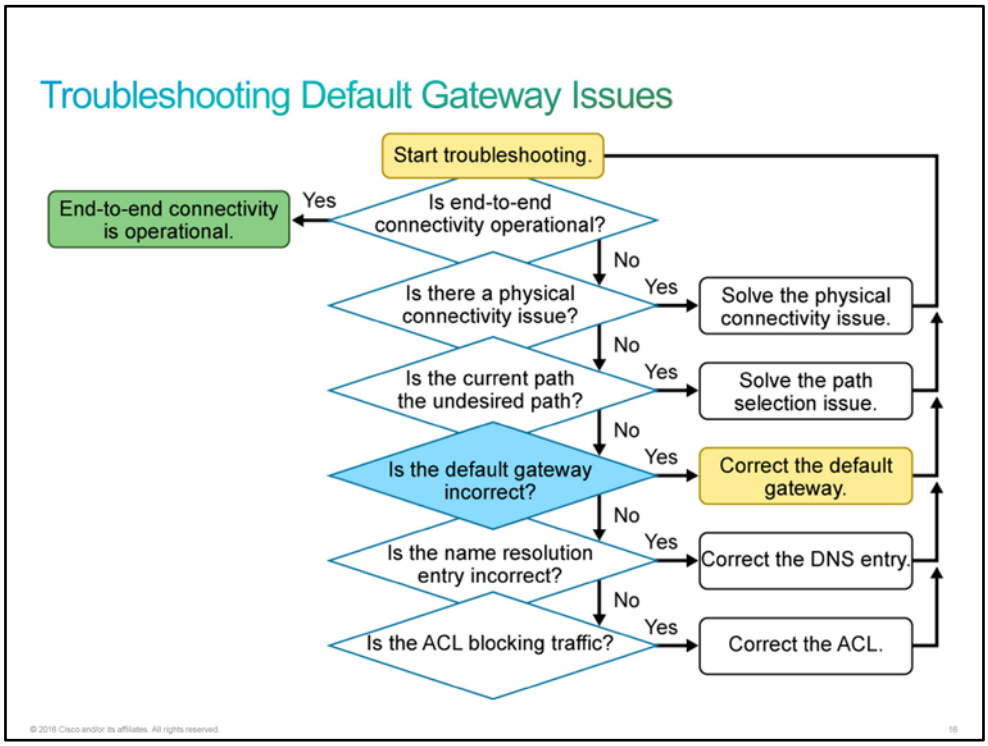

In the example, PC1 needs connectivity to the Server. The figure shows the configuration of default gateways on the PC and the Branch router. For communication between the PC and the Server to work, the PC and the Branch router must have one of the following:

- Specific routes to the network 172.16.1.0
- Correctly configured default gateways

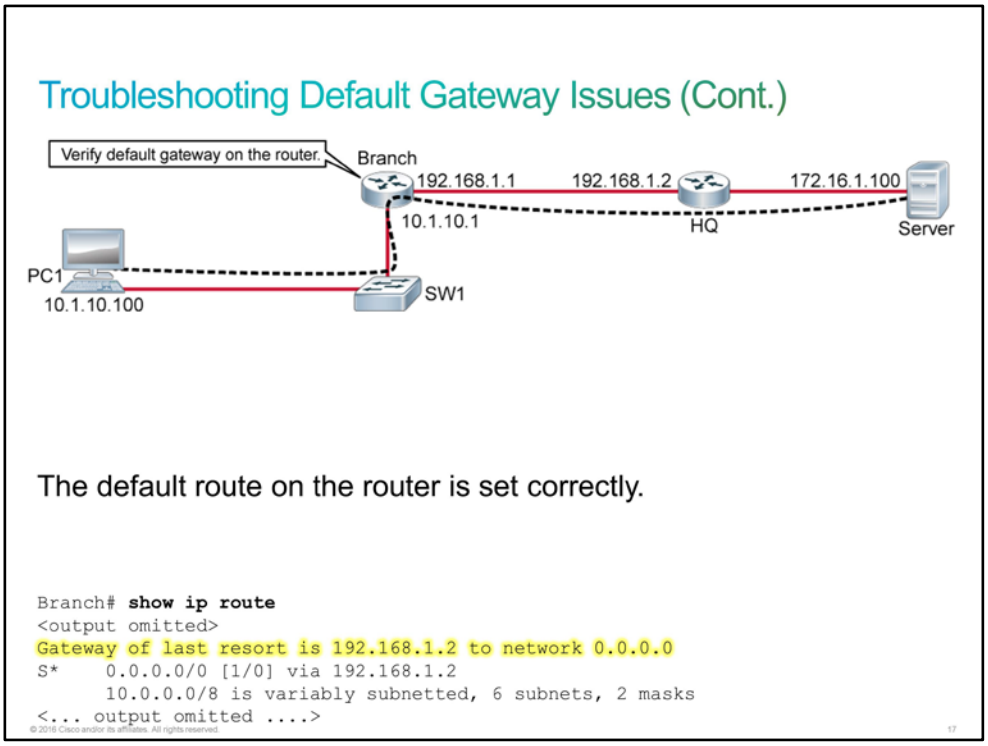

To verify the default gateway on a Cisco IOS device, use the **show ip route** command. To verify the default gateway on a Windows host, use the **route print** command.

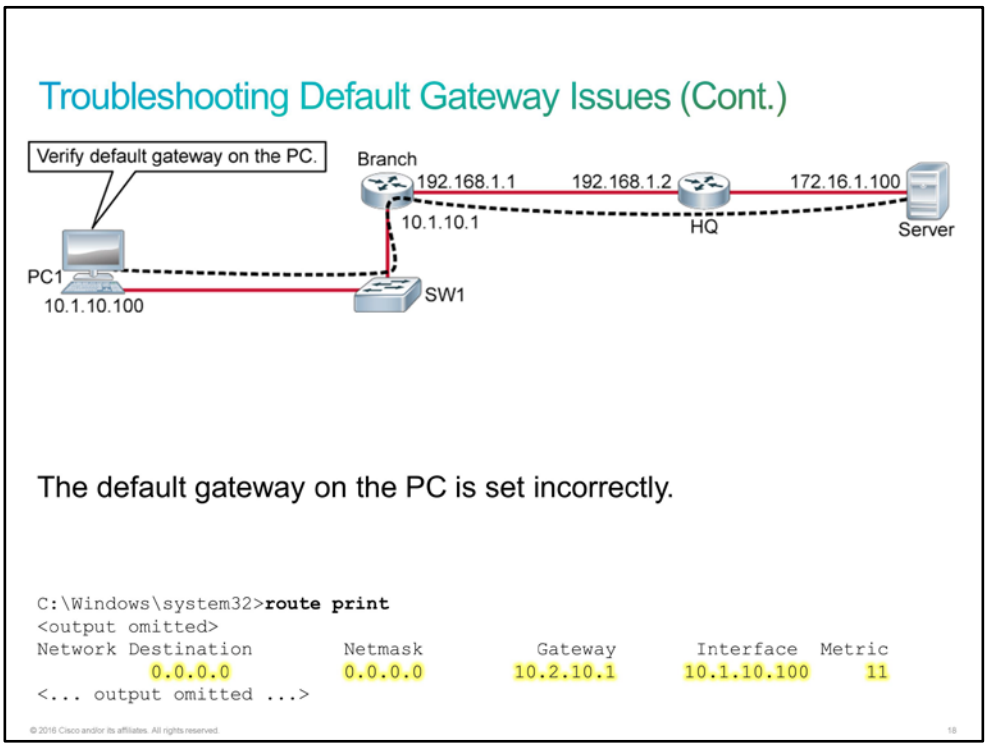

In the example, the Branch router has the correct default gateway, which is the IP address of the HQ router. PC1 has the wrong default gateway. PC1 should have the default gateway of the Branch router 10.1.10.1.

Troubleshooting Name Resolution Issue

The next troubleshooting step on the troubleshooting flow chart involves determining whether there is a name resolution issue on the network. Name resolution is the mapping of IP addresses to names, and vice versa. Name resolution is very important for networks because you often use names instead of IP addresses in order to access resources. For example, you typically access websites by using their names, such as www.somedomain.com, instead of their IP addresses because it is much easier to remember names.

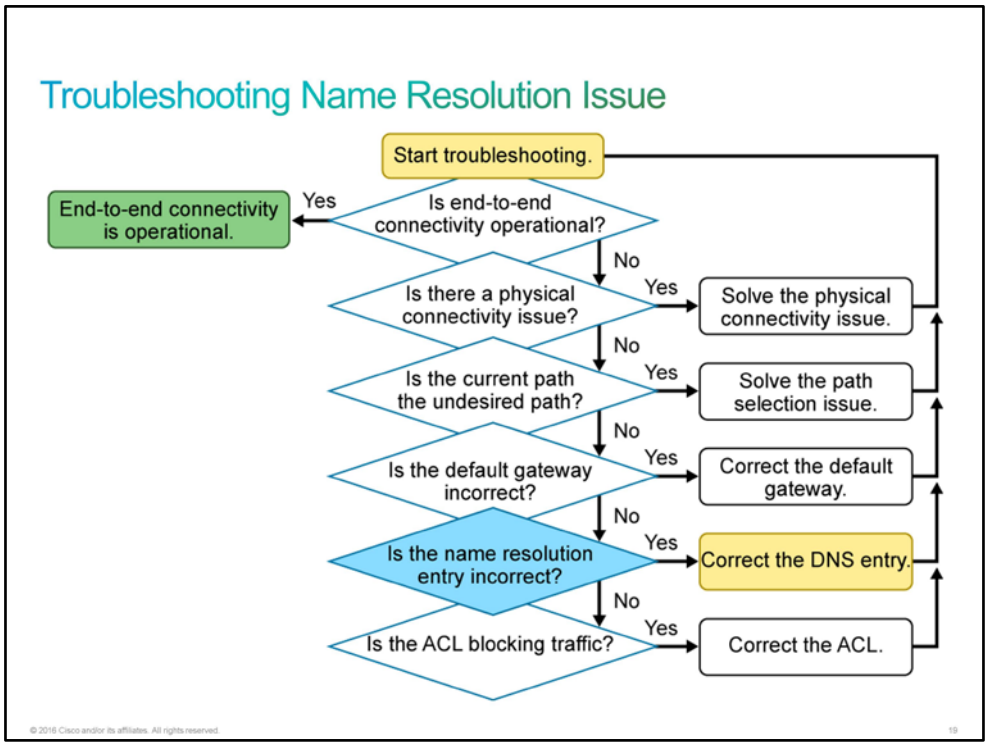

The IP defines a hierarchical naming scheme that allows a device to be identified by its location or domain. Domain names are formed with periods (.) as the delimiting characters. For example, Cisco is a commercial organization that IP identifies by the ".com" domain name, so its domain name is cisco.com. A specific device in this domain, for example, the FTP system, is identified as ftp.cisco.com.

The mapping of computer names to IP addresses can be done in two ways:

- **Static:** The system administrator creates a text file, which is called the hosts file, and enters each computer name and IP address. The file is then distributed on the network. When a user makes a request for a connection to another computer, the system uses the file to resolve the name to the correct IP address. This system works well for simple networks that change infrequently.
- **Dynamic:** The DNS protocol controls the DNS—a distributed database with which you can map host names to IP addresses.

When you configure name resolution on the device, you can substitute the host name for the IP address with all IP commands, such as **ping** or **telnet**.

To keep track of domain names, IP has defined the concept of a domain name server, which holds a cache (or database) of names that is mapped to IP addresses. To map domain names to IP addresses, you must first identify the host names, specify the name server that is present on your network, and enable the DNS.

It is possible for IP connectivity to work but for the name resolution to fail. If you are unable to access a website by its name, you might still be able to access it by its IP address. To determine if you are experiencing a name resolution problem, ping the destination by its IP address and then by its name. If you can ping the device by its IP address but not its name, there is a name resolution problem.

If you discover that you have a name resolution issue on your network, you can create or modify the mappings between names and IP addresses in three different places:

- **In the hosts file on a PC:** The hosts file is simply a text file that maps names to IP addresses. In a Windows operating system, the file is located at *C:\Windows\System32\drivers\etc*. Other operating systems might have the hosts file in a different location, they might use a different file, or may not have it at all. You can open and edit the hosts file with a text editor such as Notepad. This file works well for simple networks that change infrequently.
- **In your DNS:** The DNS protocol controls the DNS, a distributed database in which you can map hostnames to IP addresses.
  - You can configure DNS server information within DHCP pool, using the **dns-server** *ip_address* command. Make sure that you specify the correct IP address for the DNS server.
- **On a Cisco switch or router:** You can create static name resolution entries on a switch or a router by using the **ip host** *name ip_address* command. For example, if you want to add an entry that is named "Server" that will resolve to the IP address 172.16.1.100, the syntax would be **ip host Server 172.16.1.100**.

In the example, PC1 is configured with static a mapping of the name and IP address, then the name resolution is verified using the **ping** command.

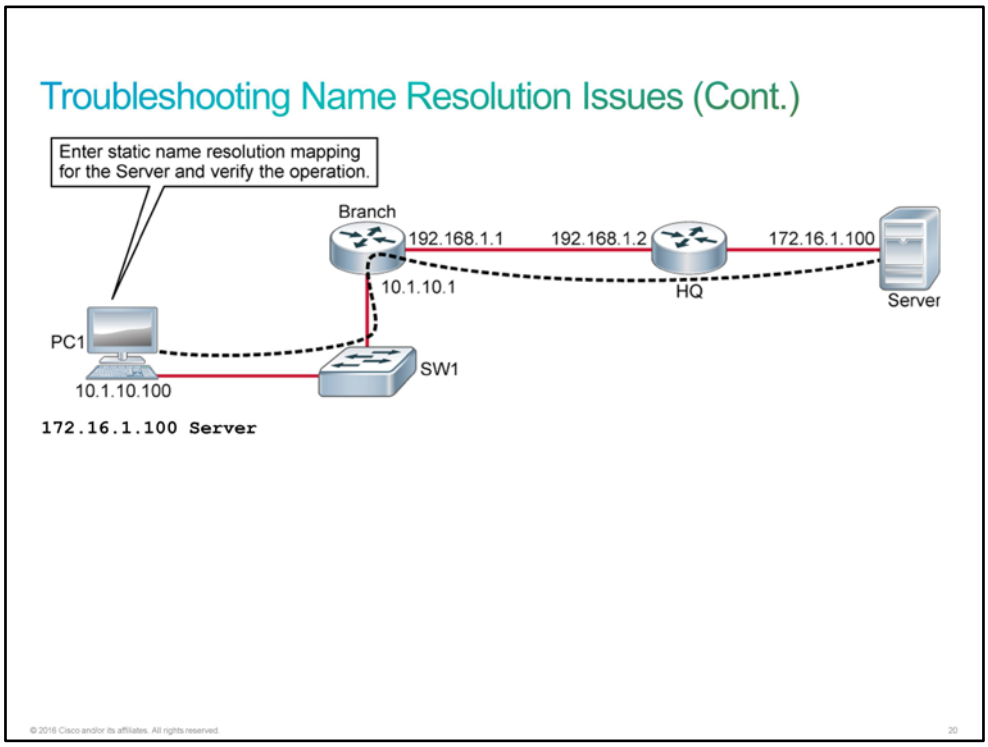

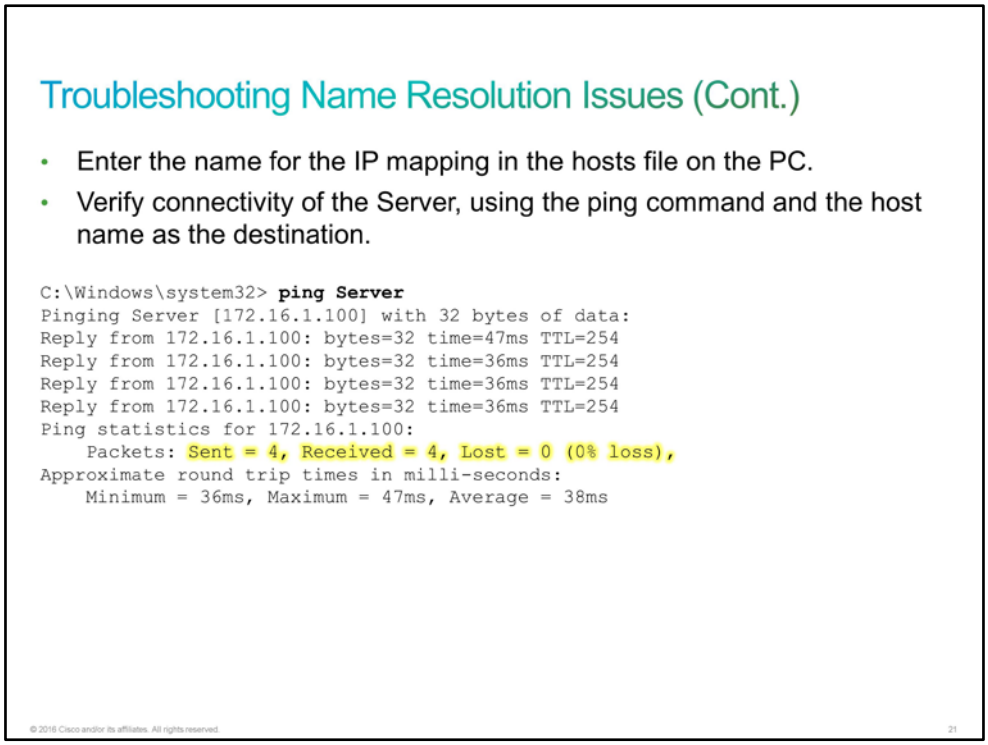

Discovery 31: Configure and Verify IPv4 Extended Access Lists

Introduction

A common mechanism that is used for traffic filtering is ACL. ACLs enable you to control access based on Layer 3 packet-header information. Standard ACLs cannot fulfill all traffic-filtering requirements, they provide only limited options for network traffic filtering.

A standard ACL can specify only source IP addresses and source networks, so it is not possible to filter to a specific destination. For more precise traffic filtering, you should use extended ACLs.

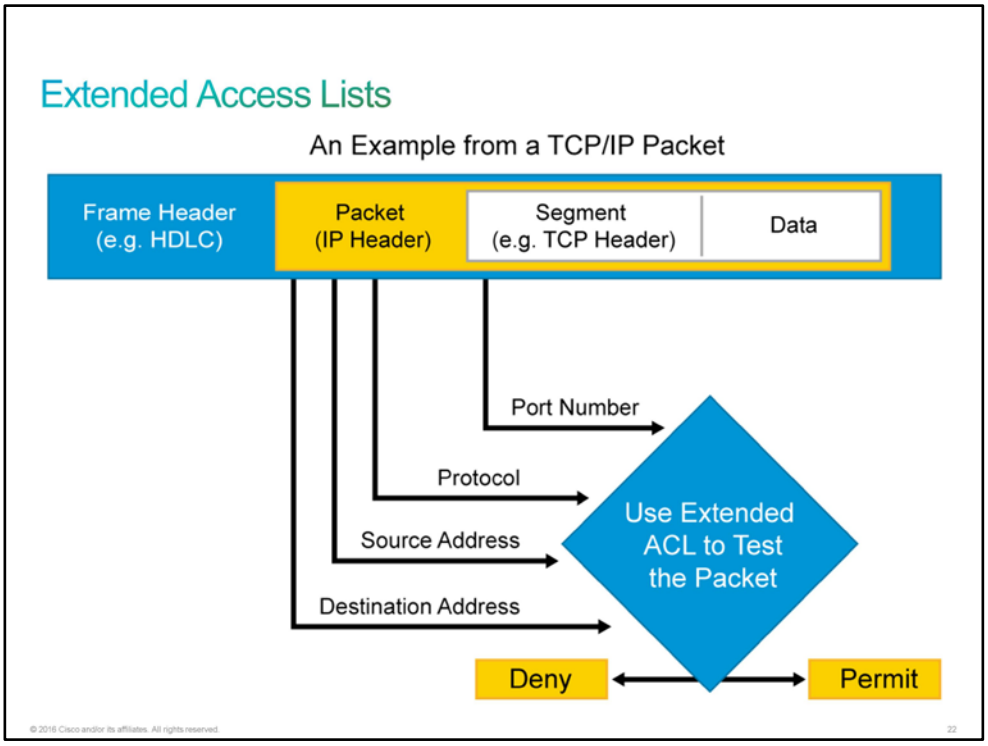

Extended ACLs provide a greater range of control. In addition to verifying packet source addresses, extended ACLs also check destination addresses, protocols, and port numbers, as shown in the figure. They provide more criteria on which to base the ACL. For example, an extended ACL can simultaneously allow email traffic from a network to a specific destination and deny file transfers and web browsing for a specific host.

The ability to filter on a protocol and port number allows you to build very specific extended ACLs. Using the appropriate port number, you can specify an application by configuring either the port number or the name of a well-known port.

You have two types of extended ACLs:

- **Named:** More common.
- **Numbered:** Ranges from 100 to 199, and from 2000 to 2699 (providing a total of 800 possible extended ACLs).

This discovery will guide you through the extended IPv4 ACL configuration. The virtual lab environment 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. The configuration of both ACL will be on R1 and it will be applied inbound on the interface Ethernet0/0 to influence the traffic from PC1.

Note	The policy that is defined in the ACL was chosen to demonstrate how ACLs work. The policy does not reflect any real world application.

Topology

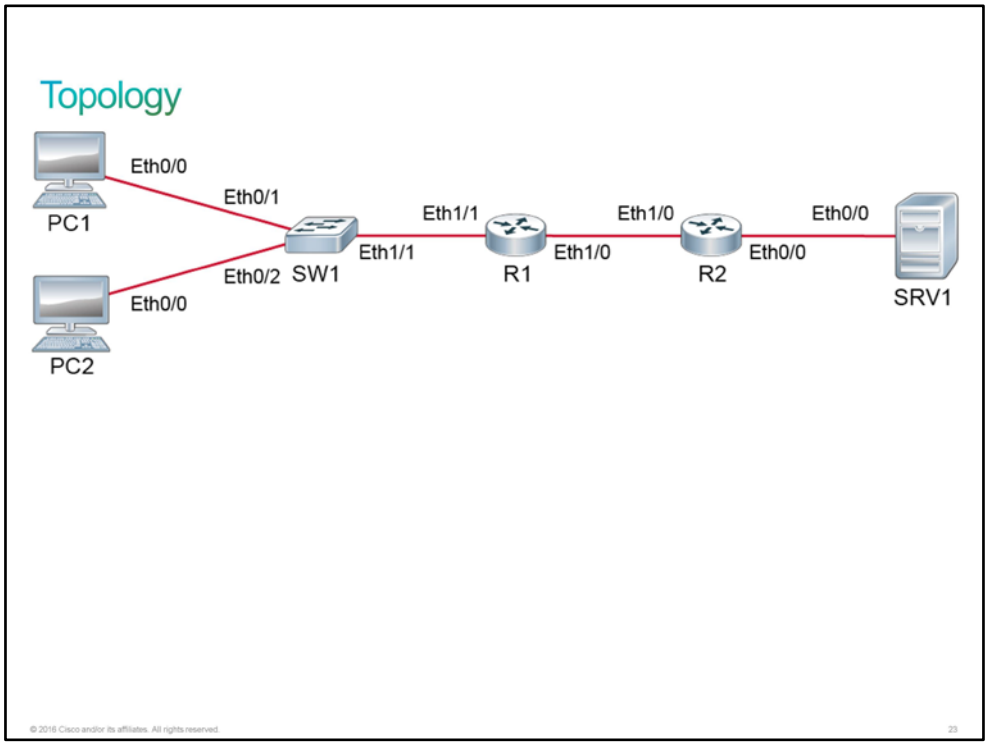

Job Aid

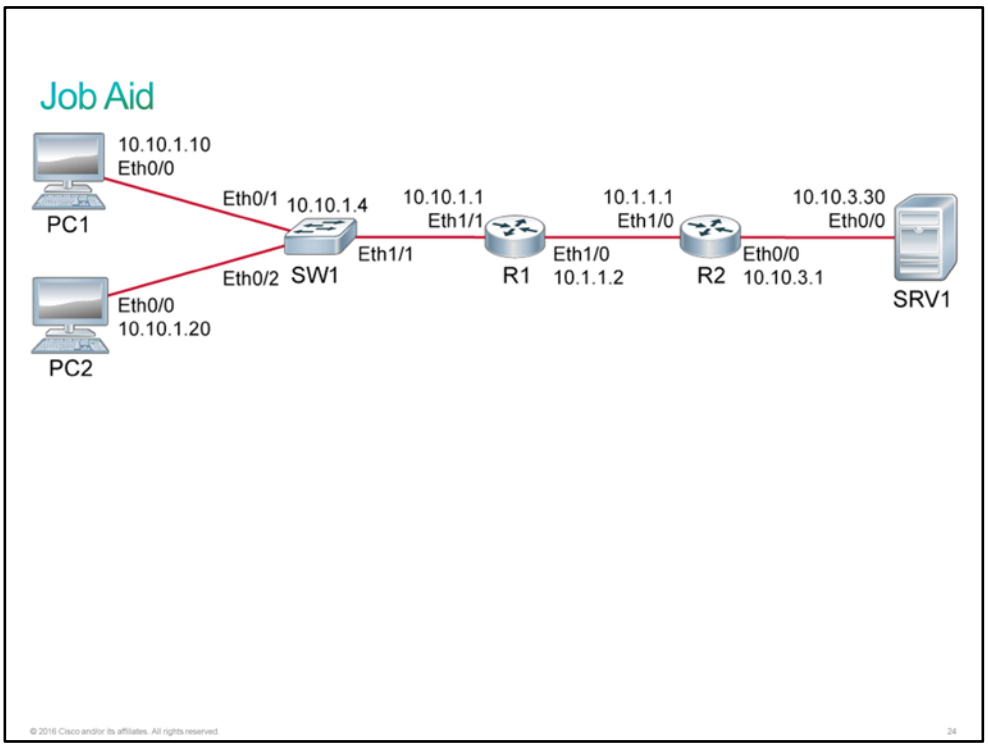

The configuration is as follows:

- All devices have their basic configurations in place including hostnames, IPv4, and IPv6 addresses.
- RIP is configured on R1 and R2 to provide IPv4 routing.
- Static routes are configured on R1 and R2 to provide IPv6 routing.

Device Details

Device	Interface	Neighbor	IPv4 Address	IPv6 Address
PC1	Ethernet0/0	SW1	10.10.1.10/24	2001:DB8:0:10::/64 Auto
PC2	Ethernet0/0	SW1	10.10.1.20/24	2001:DB8:0:10::/64 Auto
SRV1	Ethernet0/0	R2	10.10.3.30/24	2001:DB8:0:3::30/64
SW1	VLAN 1		10.10.1.4/24	2001:DB8:0:10::/64 Auto
SW1	Ethernet0/1	PC1	—	—
SW1	Ethernet0/2	PC2	—	—
SW1	Ethernet1/1	R1	—	—
R1	Ethernet1/1	SW1	10.10.1.1/24	2001:DB8:0:10::1/64

Device	Interface	Neighbor	IPv4 Address	IPv6 Address
R1	Ethernet1/0	R2	10.1.1.2/30	2001:DB8:0:2::1/64
R2	Ethernet1/0	R1	10.1.1.1/30	2001:DB8:0:2::2/64
R2	Ethernet0/0	SRV1	10.10.3.1/24	2001:DB8:0:3::1/64

Note	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 IPv4 Extended Access Lists

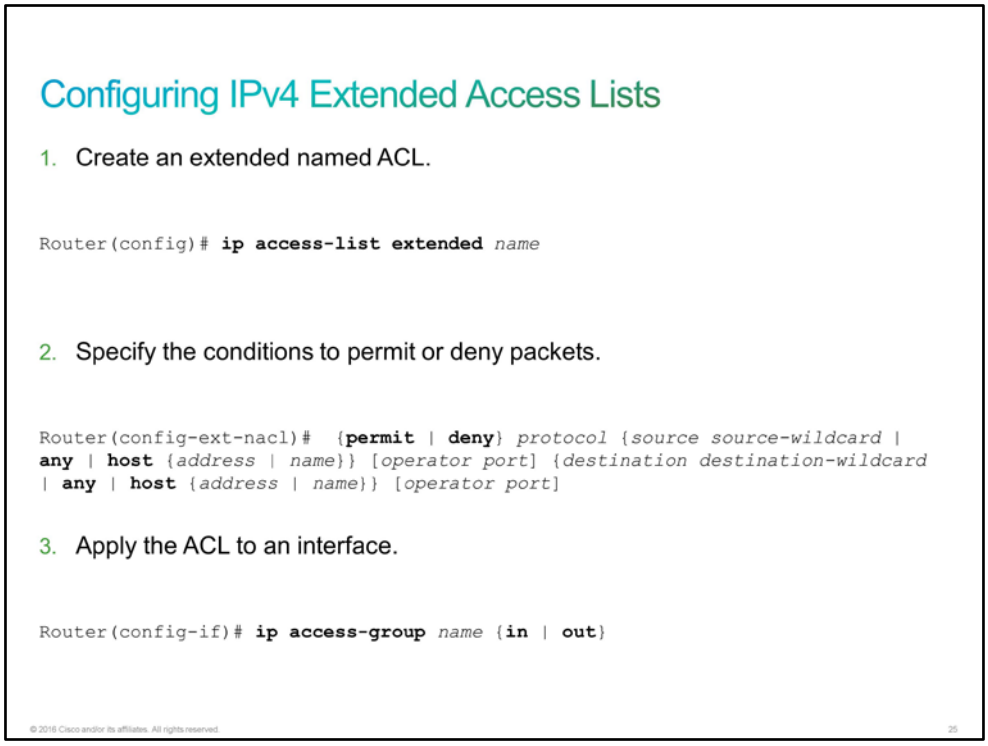

The previous examples show the steps to configure an extended named ACL. The following table explains the commands that you will use and their parameters.

Command	Description
ip access-list extended *name*	Defines an extended IP access list using a name and enters extended named access list configuration mode.
{**permit** | **deny**} *protocol* {*source source-wildcard* | **any** | **host** {*address* | *name*}} [*operator port*] {*destination destination-wildcard* | **any** | **host** {*address* | *name*}} [*operator port*]	• Permits or denies all packets that match all conditions that the remark specifies. • You can specify either the name or the number of the protocol. The most commonly used keywords are **ip**, **tcp**, **udp**, and **icmp**. • The operator is an optional parameter that compares source and destination ports when TCP or UDP is specified as the protocol. Possible operands include **lt** (less than), **gt** (greater than), **eq** (equal), **neq** (not equal), and **range** (inclusive range). • The port is an optional decimal number or name of a TCP or UDP port.
ip access-group *name* {**in** | **out**}	Applies the specified access list to the interface in the inbound or outbound direction.

Activity

Complete the following steps:

Step 1 Access the console on R1 and configure a named extended IPv4 ACL. The ALC should be named "Example4".

ACL should have these four statements:

- The first should deny all UDP traffic.
- The second should permit TCP from PC1 to any destination as long as the destination port is 23 (Telnet).
- The third should deny all other TCP traffic from PC1.
- The last should explicitly permit all IP traffic.

Note: At the end of every created ACL is an implicit deny statement.

```
R1# conf t
R1(config)# ip access-list extended Example4
R1(config-ext-nacl)# deny udp any any
R1(config-ext-nacl)# permit tcp host 10.10.1.10 any eq 23
R1(config-ext-nacl)# deny tcp host 10.10.1.10 any
R1(config-ext-nacl)# permit ip any any
R1(config-ext-nacl)# exit
```

Step 2 Apply the ACL to the interface Ethernet1/1 in the inbound direction.

At the end, leave the configuration mode.


```
R1(config)# interface Ethernet1/1
R1(config-if)# ip access-group Example4 in
R1(config-if)# end
```

Step 3 Display the ACL.

Look over the ACL definition—any slight variation in its definition can lead to large differences in its behavior.

```
R1# show ip access-lists Example4
Extended IP access list Example4
    10 deny udp any any
    20 permit tcp host 10.10.1.10 any eq telnet
    30 deny tcp host 10.10.1.10 any
    40 permit ip any any
```

The access list has all four statements in the correct order as you configured. Note that the output does not display the implicit deny statement that is at the end of every ACL.

Step 4 Now test the ACL performance by executing all the types of traffic that you specified in the ACL statements.

The first line of the ACL will block all UDP traffic. SRV1 is configured as the NTP server, but because NTP uses the UDP protocol, the first line in the ACL should block access for PC1.

To verify this case, access the console of PC1 and configure it to use the SRV1 IPv4 address as an NTP server and then display the status of NTP on PC1.

```
PC1# conf t
PC1(config)# ntp server 10.10.3.30
PC1(config)# end
```

Because NTP traffic from PC1 is blocked, you should find that it has not synchronized to SRV1.

```
PC1# show ntp status
Clock is unsynchronized, stratum 16, no reference clock
nominal freq is 250.0000 Hz, actual freq is 250.0000 Hz, precision is 2**10
ntp uptime is 3500 (1/100 of seconds), resolution is 4000
reference time is 00000000.00000000 (00:00:00.000 PST Mon Jan 1 1900)
clock offset is 0.0000 msec, root delay is 0.00 msec
root dispersion is 0.52 msec, peer dispersion is 0.00 msec
loopfilter state is 'FSET' (Drift set from file), drift is 0.000000000 s/s
system poll interval is 8, never updated.
```

Step 5 The second line in the ACL explicitly permits Telnet traffic from PC1.

Verify that PC1 can successfully **telnet** to SRV1. Use the username "admin" and password "Cisco123".

At the SRV1 system prompt, use the **exit** command to terminate the connection.

```
PC1# telnet 10.10.3.30
Trying 10.10.3.30 ... Open

User Access Verification

Username: admin
Password:
SRV1> exit

[Connection to 10.10.3.30 closed by foreign host]
PC1#
```

Step 6 The third line in the ACL is denying all other TCP traffic from PC1. Verify that PC1 cannot use SSH to reach SRV1.

The **–l** is a dash and a capital "L", not a numeral 1. Think of "L" to specify the login ID for the SSH session.

```
PC1# ssh -l admin 10.10.3.30
% Destination unreachable; gateway or host down
```

Step 7 The fourth line in the ACL, which explicitly permits all IP traffic, should permit any non-UDP traffic from PC1. Verify that PC1 can ping the server.

```
PC1# 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 100 percent (5/5), round-trip min/avg/max = 1/1/2 ms
```

The first three lines do not explicitly specify the ICMP protocol. Hence, any ICMP traffic should be permitted by the fourth line in the ACL which explicitly permits all IP traffic that did not match any previous line.

Step 8 The ACL only applies to traffic coming from PC1. Access the console of PC2 and attempt the same test sequence that you did from PC1.

The test uses ICMP and TCP, not UDP. All the tests should succeed.


```
PC2# telnet 10.10.3.30
Trying 10.10.3.30 ... Open

User Access Verification

Username: admin
Password:
SRV1> exit

[Connection to 10.10.3.30 closed by foreign host]
PC2# ssh -l admin 10.10.3.30
Password:
SRV1> exit

[Connection to 10.10.3.30 closed by foreign host]
PC2# 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 100 percent (5/5), round-trip min/avg/max = 1/1/2 ms
```

Notice that PC2 can both telnet and SSH to the server, whereas PC1 could not. This is because no ACL is applied toward the PC2.

Step 9 Display the ACL again and observe the updated hit counters that are associated with the activity that you just initiated.

```
R1# show ip access-lists Example4
Extended IP access list Example4
    10 deny udp any any (6 matches)
    20 permit tcp host 10.10.1.10 any eq telnet (58 matches)
    30 deny tcp host 10.10.1.10 any (1 match)
    40 permit ip any any (89 matches)
```

Due to the dynamic nature of the lab environment, the hit counters that you observe are likely to differ from what the example shows.

This is the end of the discovery lab.

Troubleshooting ACL Issues

If you have eliminated physical connectivity, routing, and name resolution issues and you are still experiencing network connectivity problems, your next step is to troubleshoot ACLs.

The routers may have ACLs configured that prohibit a protocol to pass the interface in the inbound or outbound direction.

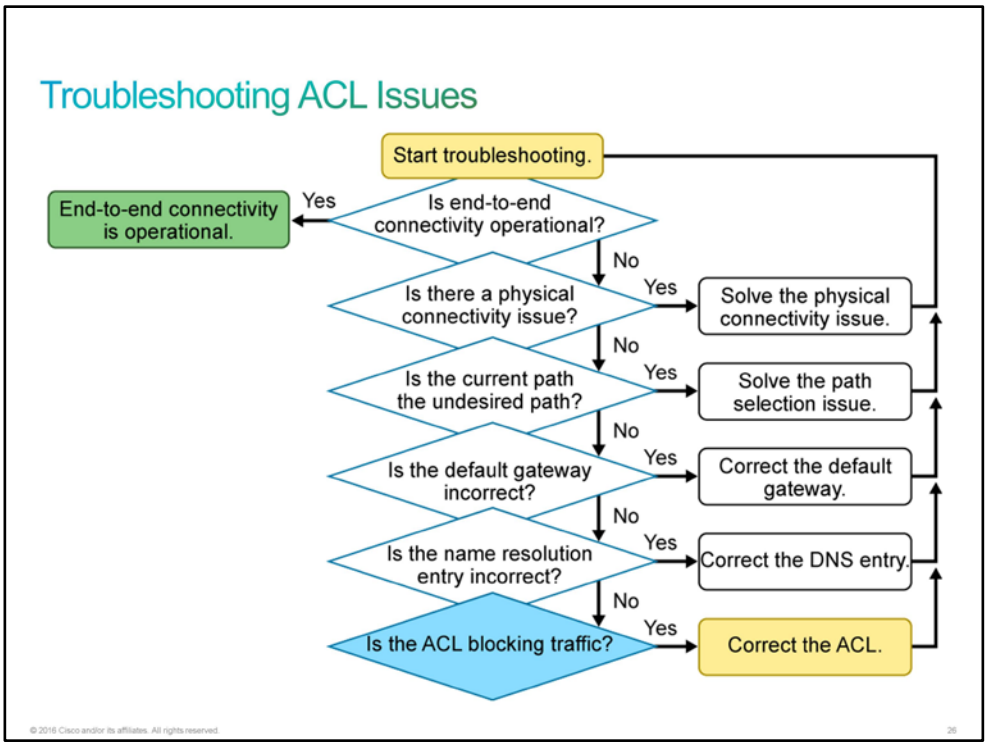

In the example, PC1 in unable to use Telnet to connect to the Server.

To begin, you might want to use the **show ip access-lists** command to display the contents of all ACLs that are configured on the router. By entering the ACL name or number as an option for this command, you can display a specific ACL.

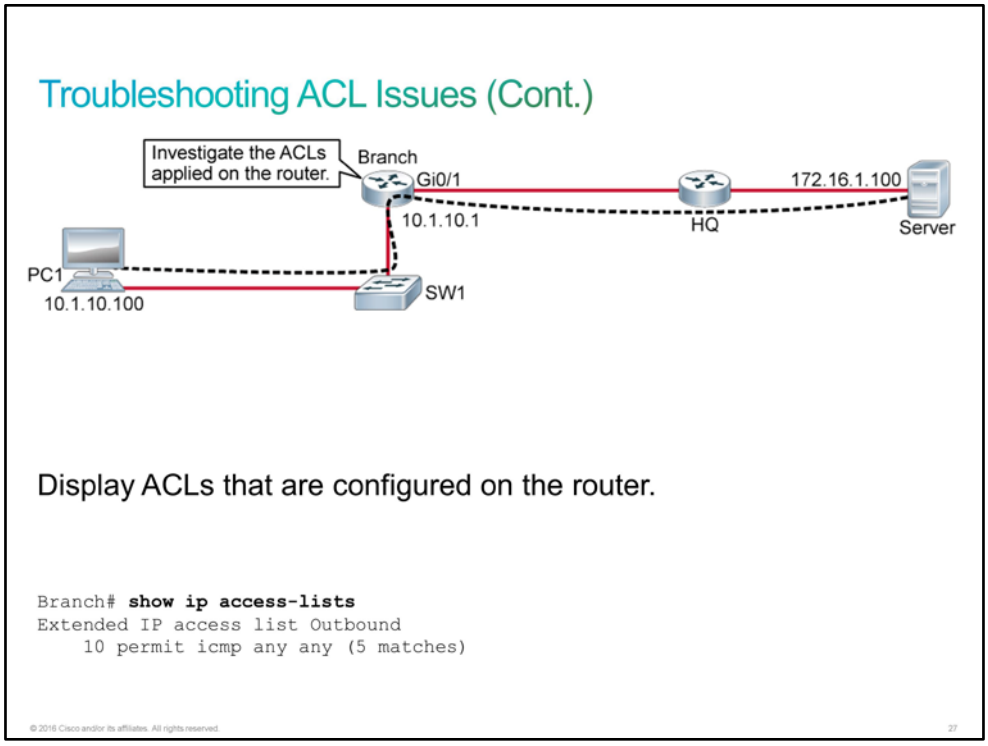

In this example, there is an ACL that is named "Outbound." It is implicitly denying Telnet and all other traffic except ICMP.

When you discover that an ACL on a router is blocking traffic that you want to permit, you can use the **show ip interface** command to determine where the ACL is applied.

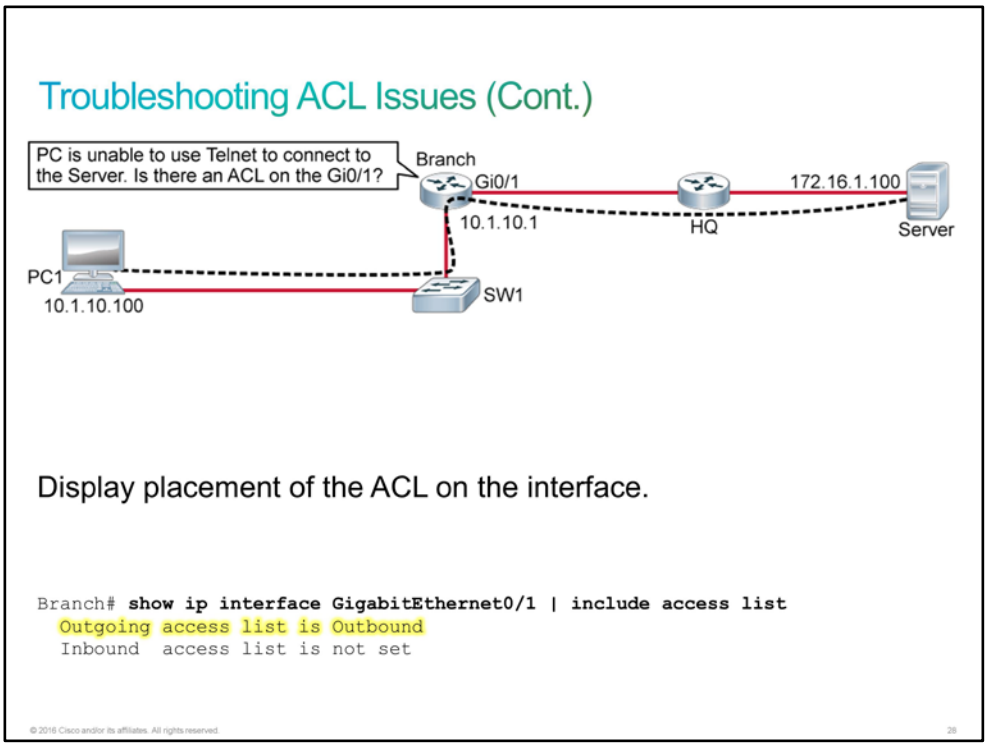

In the following example, the IP ACL that is named "Outbound" has been configured on the interface GigabitEthernet0/1 as an outbound ACL.

To have the Branch router permit Telnet, you would need to add an ACL entry that allows Telnet. Currently, the Outbound ACL permits only the ICMP protocol. In order to allow a Telnet connection from PC1 to the server, add an entry in the Outbound ACL to allow the TCP protocol and port 23 for Telnet as follows:

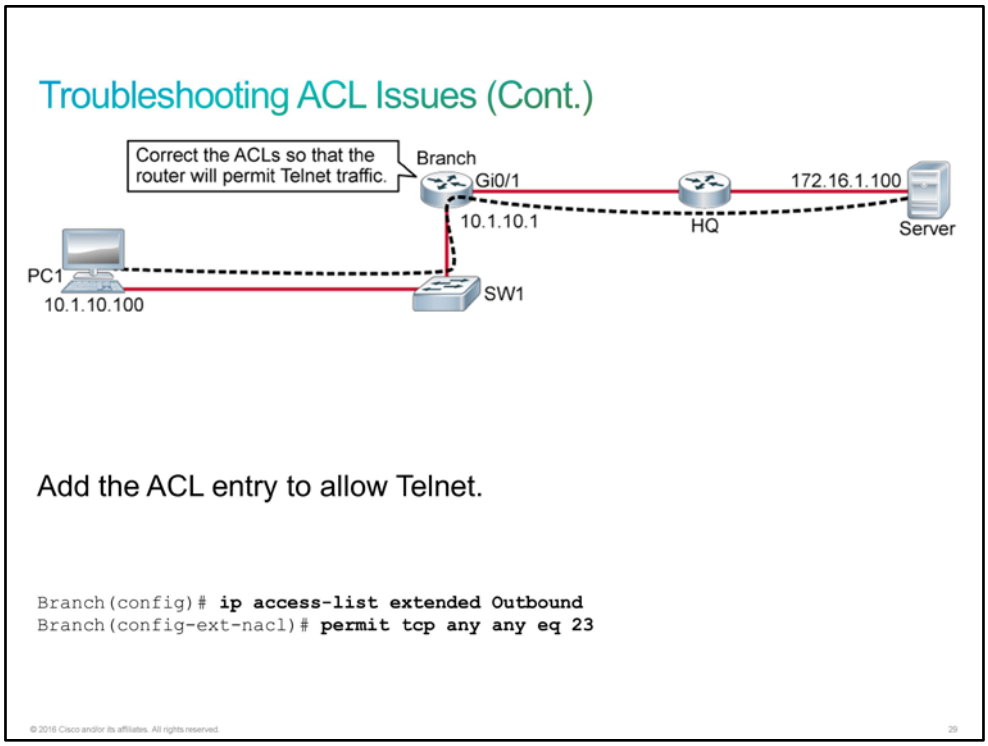

After correcting the Outbound ACL, a Telnet connection from PC1 to the server should be successful.

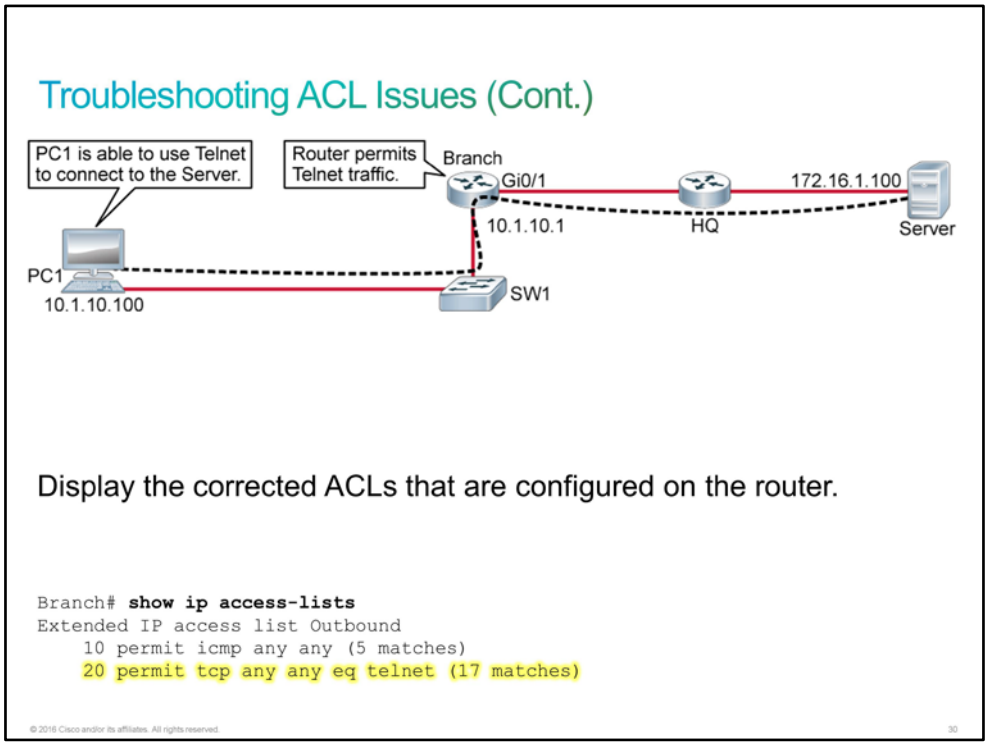

Discovery 32: Troubleshoot IPv4 Network Connectivity

Introduction

This discovery will guide you through troubleshooting connectivity in an IPv4 network. The 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 has been configured as the dynamic routing protocol.

Four issues have been introduced on different devices in the live virtual lab environment. Your job is to find and fix these issues. There are only four steps in this discovery. The step describes the complaint that you must address. To get the feeling of troubleshooting activities, try to uncover and resolve the problems before you use the Answer Key for each step.

Resolve each issue before moving to the next issue. Sometimes, you will need to resolve the issue to be able to go to the following issue.

Topology

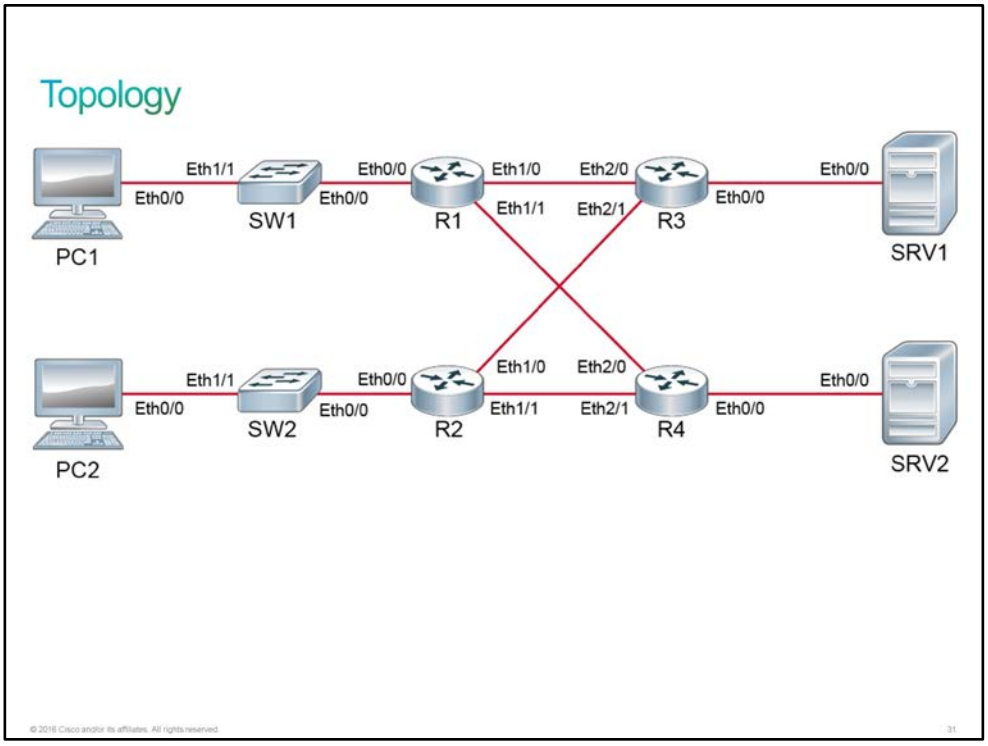

Job Aid

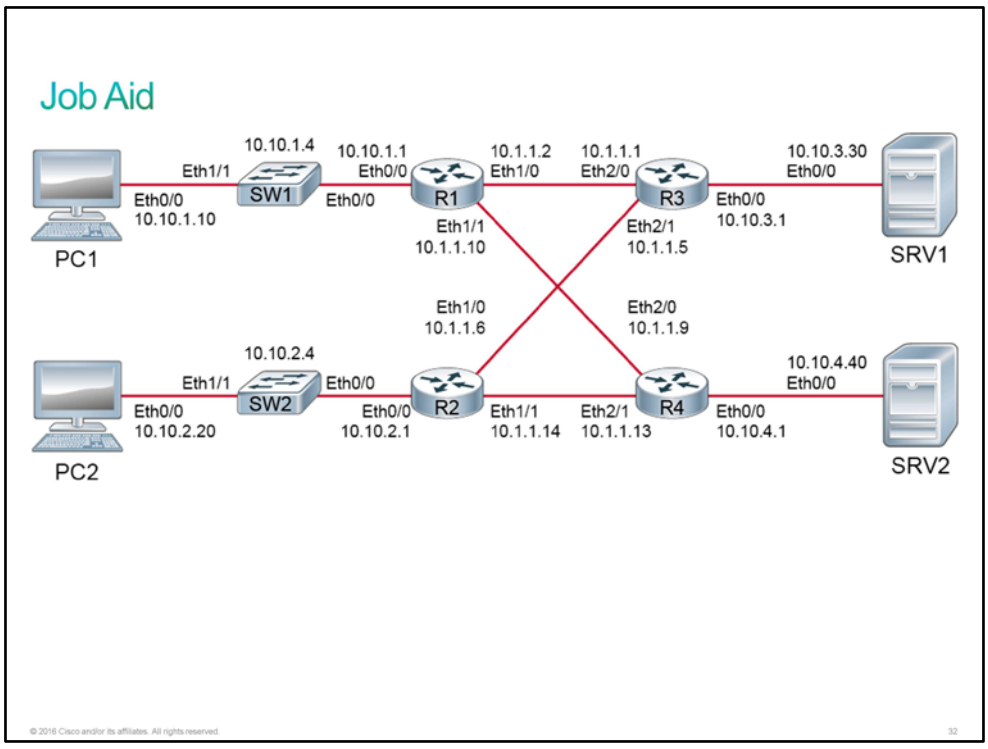

The configuration is as follows:

- All devices have their basic configurations in place, including hostnames and IP addresses.
- RIP is configured on all four routers.
- Four issues, related to the PCs connectivity to the SRVs, exist in the network.

Device Information

Device	Interface	Neighbor	IPv4 Address
PC1	Ethernet0/0	SW1	10.10.1.10/24
PC2	Ethernet0/0	SW2	10.10.2.20/24
SRV1	Ethernet0/0	R3	10.10.3.30/24
SRV2	Ethernet0/0	R4	10.10.4.40/24
SW1	VLAN 1	—	10.10.1.4/24
SW2	VLAN 1	—	10.10.2.4/24
R1	Ethernet0/0	SW1	10.10.1.1/24
R1	Ethernet1/0	R3	10.1.1.2/30

Device	Interface	Neighbor	IPv4 Address
R1	Ethernet1/1	R4	10.1.1.10/30
R2	Ethernet 0/0	SW2	10.10.2.1/24
R2	Ethernet1/0	R3	10.1.1.6/30
R2	Ethernet1/1	R4	10.1.1.14/30
R3	Ethernet2/0	R1	10.1.1.1/30
R3	Ethernet2/1	R2	10.1.1.5/30
R3	Ethernet0/0	SRV1	10.10.3.1/24
R4	Ethernet2/0	R1	10.1.1.9/30
R4	Ethernet2/1	R2	10.1.1.13/30
R4	Ethernet0/0	SRV2	10.10.4.1/24

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: Troubleshoot IPv4 Network Connectivity

Activity

Complete the following steps:

Step 1 The user at PC1 is complaining of not being able to connect to SRV1. The user is using Telnet for connectivity.

```
PC1# telnet 10.10.3.30
Trying 10.10.3.30 ...
% Destination unreachable; gateway or host down
```

PC1 will take the path through R1 and R3 to reach the SRV1, so you should investigate those two routers on the path.

Here are some steps that you might take in the troubleshooting process. Sometimes, the steps show that an item is in a normal working order. This data is valuable to have when troubleshooting. In other cases, the step may point out something that is out of order and gets you closer to determining the root cause:

- On R1, verify the following:
  - You have a valid route to the SRV1—use the **show ip route** command.

  - Ethernet0/0 and Ethernet1/0 interfaces are in the admin "up/up" state—use the **show ip interface brief** command.
  - No access list is applied to the Ethernet0/0 and Ethernet1/0 interfaces—use the **show run | section interface** or **show ip interface** command.
- On R3, verify the following:
  - You have a valid route to the SRV1—use the **show ip route** command.
  - Ethernet0/0 and Ethernet2/0 interfaces are in the admin "up/up" state—use the **show ip interface brief** command.

Notice that R3 has the interface toward SRV1 (Ethernet0/0) administratively disabled.

```
R3# show ip interface brief
Interface                  IP-Address      OK? Method Status
Protocol
Ethernet0/0                10.10.3.1       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
Ethernet1/0                unassigned      YES NVRAM  administratively down
down
Ethernet1/1                unassigned      YES NVRAM  administratively down
down
Ethernet1/2                unassigned      YES NVRAM  administratively down
down
Ethernet1/3                unassigned      YES NVRAM  administratively down
down
Ethernet2/0                10.1.1.1        YES NVRAM  up                    up
Ethernet2/1                10.1.1.5        YES NVRAM  up                    up
Ethernet2/2                unassigned      YES NVRAM  administratively down
down
Ethernet2/3                unassigned      YES NVRAM  administratively down
down
```

You can resolve the problem by enabling the Ethernet0/0 interface on R3.

```
R3# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
R3(config)# interface Ethernet0/0
R3(config-if)# no shut
R3(config-if)# end
```

When you fix the configuration on R3, you should now be able to telnet to SRV1 from PC1. Use "admin" for the username and "Cisco123" for the password.

```
PC1# telnet 10.10.3.30
Trying 10.10.3.30 ... Open


User Access Verification

Username: admin
Password:
SRV1>exit

[Connection to 10.10.3.30 closed by foreign host]
```

Step 2 The user at PC1 is complaining of not being able to connect to SRV2. In fact, if the user attempts to ping SRV2, the ping shows that the server is unreachable.

```
PC1# ping SRV2
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.10.1.40, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
```

The IP address of SRV2 that PC1 is pinging is 10.10.1.40. However, this address is not the IP address of the server. You can verify it by showing the interface status on SRV2. You can also do verify it by comparing the address to the information in the topology diagram and the connectivity table.

```
SRV2# show interfaces Ethernet0/0
Ethernet0/0 is up, line protocol is up
  Hardware is AmdP2, address is aabb.cc00.2100 (bia aabb.cc00.2100)
  Description: Link to R4
  Internet address is 10.10.4.40/24
  MTU 1500 bytes, BW 10000 Kbit/sec, DLY 1000 usec,
     reliability 255/255, txload 1/255, rxload 1/255
<... output omitted ...>
```

If you ping the SRV2 using the IP address instead of its host name, you will see that the server is reachable.

```
PC1# ping 10.10.4.40
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.10.4.40, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
```

The problem is an incorrect entry for SRV2 in the local host configuration on PC1.

```
PC1# show running-config | include host
hostname PC1
ip host SRV2 10.10.1.40
ip host SRV1 10.10.3.30
```

You can resolve the problem by configuring the host entry properly.


```
PC1# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
PC1(config)# ip host SRV2 10.10.4.40
PC1(config)# end
```

When you configure the entry properly, you should be able to ping SRV2 by hostname from PC1.

```
PC1# ping SRV2
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.10.4.40, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
```

Step 3 The user at PC2 is complaining of not being able to connect to SRV1. If the user attempts to ping SRV1 IP address, the ping shows that the server is unreachable.

```
PC2# 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 0 percent (0/5)
```

PC2 will take the path through R2 and R3 to reach SRV1, so you should investigate those two routers on the path.

Here are some steps that you might take in the troubleshooting process:

- On R3, verify the following:
  - You have a valid route to the SRV1—use the **show ip route** command.
  - Ethernet0/0 and Ethernet2/1 interfaces are in the admin "up/up" state—use the **show ip interface brief** command.
  - No access list is applied to the Ethernet0/0 and Ethernet2/1 interfaces—use the **show run | section interface** or **show ip interface** command.
  - You can ping the SRV1 IP address.
- On R2, verify the following:
  - You have a valid route to SRV1—use the **show ip route** command.
  - Ethernet0/0 and Ethernet1/0 interfaces are in the admin "up/up" state—use the **show ip interface brief** command.
  - No access list is applied to the Ethernet0/0 and Ethernet1/0 interfaces—use the **show run | section interface** command.
  - You can ping the SRV1 IP address.

Because SRV1 is up, and R2 and R3 are properly configured, you can determine that the problem lies on PC2. You can verify that PC2 has the interface toward the SRV1 enabled.

Using the **show ip route** command, determine if PC2 has a valid route to the SRV1.

```
PC2# show ip route
Default gateway is 10.0.2.1

Host               Gateway           Last Use    Total Uses  Interface
ICMP redirect cache is empty
```

PC2 has been configured with an incorrect default gateway—the default gateway is set to some nonexisting IP address in the network.

You can resolve the problem by configuring the default gateway on PC2. The R2 Ethernet0/0 IP address should be set as the default gateway on 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
```

When you configure PC2 with the correct default gateway, you should be able to ping SRV1 from PC2

```
PC2# 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 80 percent (4/5), round-trip min/avg/max = 1/1/1 ms
```

Step 4 The user on PC2 is complaining of not being able to connect to SRV2. The user is using Telnet for connectivity.

```
PC2# telnet 10.10.4.40
Trying 10.10.4.40 ...
% Destination unreachable; gateway or host down
```

```
PC2# ping 10.10.4.40
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.10.4.40, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/2 ms
```

Telnet uses TCP to test connectivity. By default it will connect to port 23. SRV2 on port 23 is not reachable; however, you can see that the ping which uses ICMP to test connectivity to SRV2 is successful.

Using the **traceroute** command, you can determine that PC2 takes the path via R1 and R4 to reach SRV2.

```
PC2# traceroute 10.10.4.40
Type escape sequence to abort.
Tracing the route to 10.10.4.40
VRF info: (vrf in name/id, vrf out name/id)
  1 10.10.2.1 2 msec 0 msec 1 msec
  2 10.1.1.13 1 msec 1 msec 1 msec
  3 10.1.1.13 !A  *  !A
```

Note that "!A" indicates that there is an ACL applied, that is blocking access to the SRV2.

You can determine, that packet comes through R1 to the R4, where it gets blocked:

- On R1, you can still verify that there is no ACL configured—use the **show ip access-lists** command,
- On R4, verify if there is an ACL configured.

```
R4# show ip access-lists
Extended IP access list Server
    10 deny udp any any
    20 deny tcp any any eq telnet (1 match)
    30 deny tcp any any eq www
    40 permit ip any any (5 matches)
```

R4 has an IP ACL configured that is blocking the telnet access to SRV2. Verify if this ACL is applied on the Ethernet0/0 or Ethernet2/1 interface on R4.

```
R4# show ip interface Ethernet0/0
Ethernet0/0 is up, line protocol is up
  Internet address is 10.10.4.1/24
  Broadcast address is 255.255.255.255
  Address determined by non-volatile memory
  MTU is 1500 bytes
  Helper address is not set
  Directed broadcast forwarding is disabled
  Multicast reserved groups joined: 224.0.0.9
  Outgoing access list is Server
  Inbound  access list is not set
 <... output omitted ...>

R4# show ip interface Ethernet2/1
Ethernet2/1 is up, line protocol is up
  Internet address is 10.1.1.13/30
  Broadcast address is 255.255.255.255
  Address determined by non-volatile memory
  MTU is 1500 bytes
  Helper address is not set
  Directed broadcast forwarding is disabled
  Multicast reserved groups joined: 224.0.0.9
  Outgoing access list is not set
  Inbound  access list is not set
<... output omitted ...>
```

The IP access list is applied in the outbound direction to the Ethernet0/0 interface, the one connecting to the SRV2.

Solve the problem, by either removing the statement or changing it from "deny" to "permit." The example shows the second option.

```
R4# conf t
R4(config)# ip access-list extended Server
R4(config-ext-nacl)# no 20
R4(config-ext-nacl)# 20 permit tcp any any eq 23
R4(config-ext-nacl)# end
```

PC1 should now be able to connect to SRV2 using telnet, with the "admin" username and "Cisco123" password.

```
PC2# telnet 10.10.4.40
Trying 10.10.4.40 ... Open


User Access Verification

Username: admin
Password:
SRV2> exit

[Connection to 10.10.4.40 closed by foreign host]
```

Note: The **traceroute** will still not work, because the ACL is denying UDP, which is what the traceroute uses.

This is the end of the discovery lab.

Challenge

1. Which command would you use to determine whether there are any input or output errors on a GigabitEthernet0/0 interface?

   A. **show ip route GigabitEthernet0/0**
   B. **show ip interfaces GigabitEthernet0/0**
   C. **show interfaces GigabitEthernet0/0**
   D. **show mac-address-table**

2. Which command would you use to identify the current path to a given destination on a router?

   A. **show ip route**
   B. **route print**
   C. **show ip interfaces brief**
   D. **show arp**

3. Which Cisco IOS command will enable you to see the path that packets are taking on a hop-by-hop basis?

   A. **path**
   B. **traceroute**
   C. **ping**
   D. **show route**

4. Which of the following statements that are related to configuring SPAN is true ? (Choose two)

   A. The destination port cannot be a source port, or vice versa.
   B. The destination port can be same as source port.
   C. Destination port is no longer a normal switch port—only monitored traffic passes through that port.
   D. Source port is no longer a normal switch port—only monitored traffic passes through that port.

5. Which command would show you whether an ACL is applied to an interface, GigabitEthernet 0/1 ?

   A. **show access lists GigabitEthernet 0/1**
   B. **show access lists**
   C. **show ip interface GigabitEthernet 0/1**
   D. **show interface brief**

6. Which of the following commands will you use to deny telnet access from IP address 10.1.1.1 into 10.1.1.2 ?

   A. **access-list 90 deny tcp 10.1.1.1 0.0.0.0 10.1.1.2 0.0.0.0 eq 21**
   B. **access-list 99 deny telnet 10.1.1.1 0.0.0.0 10.1.1.2 0.0.0.0**
   C. **access-list 101 deny ip 10.1.1.1 0.0.0.0 10.1.1.2 0.0.0.0 telnet**
   D. **access-list 101 deny tcp 10.1.1.1 0.0.0.0 10.1.1.2 0.0.0.0 eq 23**

7. Where should extended ACLs be placed in a network ?

A. As close to the packet's destination as possible
B. As close to the default gateway as possible
C. As close to the source of the packet as possible
D. As close to a border gateway router as possible

Answer Key

Challenge

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

Lesson 2: Troubleshooting IPv6 Network Connectivity

Introduction

A customer has called CCS with a complaint involving IPv6 network connectivity problems. A trouble ticket has been issued.

After reviewing the trouble ticket, decide whether you are ready to go onsite to solve the problem or whether you first need to do research on troubleshooting IPv6 network connectivity.

IPv6 Unicast Addresses

IPv6 unicast addresses are assigned to each node (interface). Their uses are discussed in RFC 4291. The five types of unicast addresses are listed below.

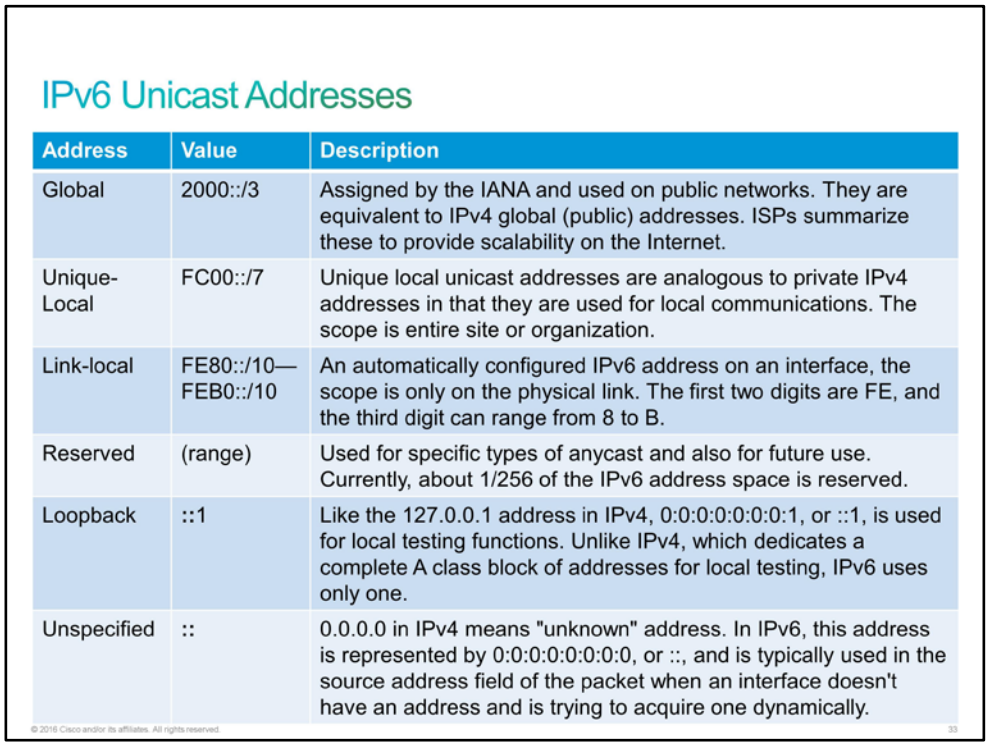

Global Addresses

RFC 4291 specifies the 2000::/3 prefix to be the global unicast address space that the IANA may allocate to the RIRs. A global unicast address is an IPv6 address that is created from the global unicast prefix. The structure of global unicast addresses enables the aggregation of routing prefixes, which limits the number of routing table entries in the global routing table. Global unicast addresses that are used on links are aggregated upward through organizations and eventually to the ISPs.

The IANA assigns a global address. The global address starts with 2000::/3. The /3 prefix length implies that only the first 3 bits are significant in matching the prefix 2000. The first 3 bits of the first hexadecimal value 2 are 001x. The fourth bit, x, is insignificant and can be either a 0 or a 1. It results in the first hextet being a 2 (0010) or a 3 (0011). The remaining 24 bits in the hextet (16-bit segment) can be a 0 or a 1.

The figure shows how address space can be allocated to the RIR and ISP. These values are minimum allocations, which means that an RIR will get a /23 or shorter, an ISP will get a /32 or shorter, and a site will get a /48 or shorter. A shorter prefix length allows more available address space. For example, a site could get a /40 instead of a /48, giving it more addresses if it can justify it to its ISP. The figure shows a provider aggregatable model where the end customer obtains its IPv6 address from the ISP. The end customer can also choose a provider-independent address space by going straight to the RIR. In this case, it is not uncommon for an end customer to be able to justify a /32 prefix.

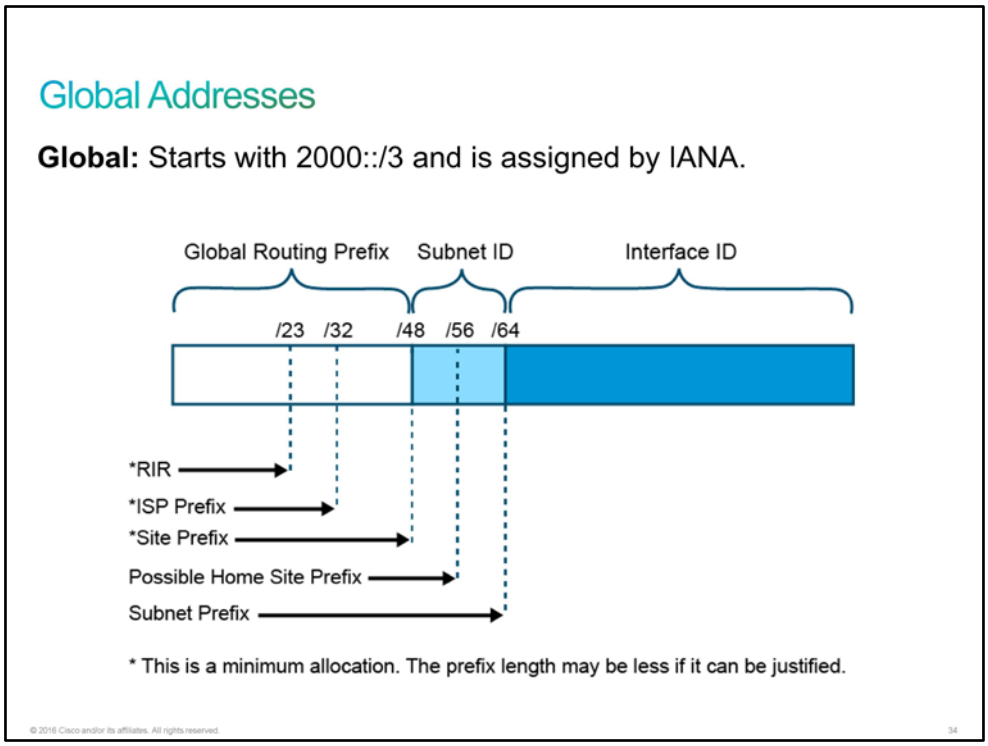

Note	The ICANN, the operator for IANA, allocates IPv6 address blocks to the five RIRs. The current global unicast address assignment from IANA begins with the binary value 001 or the prefix 2000::/3. This value allocation results in a range of global unicast addresses of 2000::/3 through 3FFF::/3

Local Addresses

A block of IPv6 addresses is set aside for local addresses, just as is done with private addresses in IPv4. These local addresses are local only to a particular link or site; therefore, they are never routed outside of a particular company network. There are two kinds of local addresses:

- **Unique local addresses:** These addresses are similar to RFC 1918, *Address Allocation for Private Internets*, in IPv4 today. The scope of these addresses is an entire site or organization. They allow addressing within an organization without needing to use a public prefix. Routers forward datagrams using site-local addresses within the site, but not outside the site, to the public Internet.

  In hexadecimal, site-local addresses begin with FE and then "C" to "F" for the third hexadecimal digit. So, these addresses begin with FEC, FED, FEE, or FEF.

- **Link-local addresses:** The concept of the link-local scope is new to IPv6. These addresses have a smaller scope than site-local addresses—they refer only to a particular physical link (physical network). Routers do not forward datagrams using link-local addresses, not even within the organization; they are only for local communication on a particular physical network segment.

  These addresses are used for link communications such as automatic address configuration, neighbor discovery, and router discovery. Many IPv6 routing protocols also use link-local addresses. A link-local address typically begins with FE80::/10.

Note	Technically speaking, an address within the prefix FE80::/10 is considered a link-local address. This scope includes addresses beginning with FE80:: through FEBF::.—this last address prefix bumps up next to the fec0::/10 range that is assigned to the deprecated site-local address scope. In common practice though, link-local addresses will typically begin with 0xFE80.

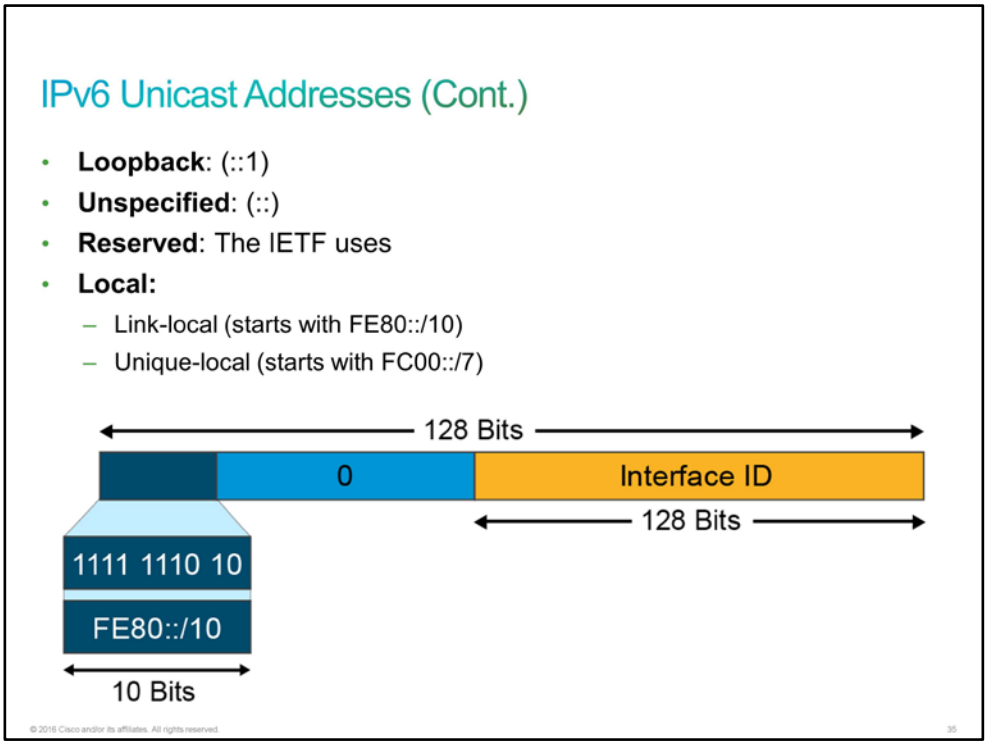

Loopback Addresses

Just as with IPv4, a provision has been made for a special loopback IPv6 address for testing. Datagrams that are sent to this address "loop back" to the sending device. However, in IPv6, there is just one address, not a whole block, for this function. The loopback address is 0:0:0:0:0:0:0:1, which is normally expressed as "::1".

Unspecified Addresses

In IPv4, an IP address containing all zeroes has a special meaning—it refers to the host itself and is used when a device does not know its own address. In IPv6, this concept has been formalized, and the all-zeros address is named the *unspecified* address. It is typically used in the source field of a datagram that a device that seeks to have its IP address configured sends. You can apply address compression to this address. Because the address is all zeroes, the address becomes just "::".

Reserved Addresses

The IETF reserved a portion of the IPv6 address space for various uses, both present and in the future. Reserved addresses represent 1/256th of the total IPv6 address space. The lowest address within each subnet prefix (the interface identifier set to all zeroes) is reserved as the *subnet-router* anycast address. The 128 highest addresses within each /64 subnet prefix are reserved to be used as anycast addresses.

Assigning IPv6 Addresses

Interface identifiers in IPv6 addresses are used to identify interfaces on a link. They can also be thought of as the "host portion" of an IPv6 address. Interface identifiers need to be unique on a specific link. Interface identifiers are always 64 bits and can be dynamically derived from a Layer 2 media and encapsulation. There are several ways to assign an IPv6 address to a device:

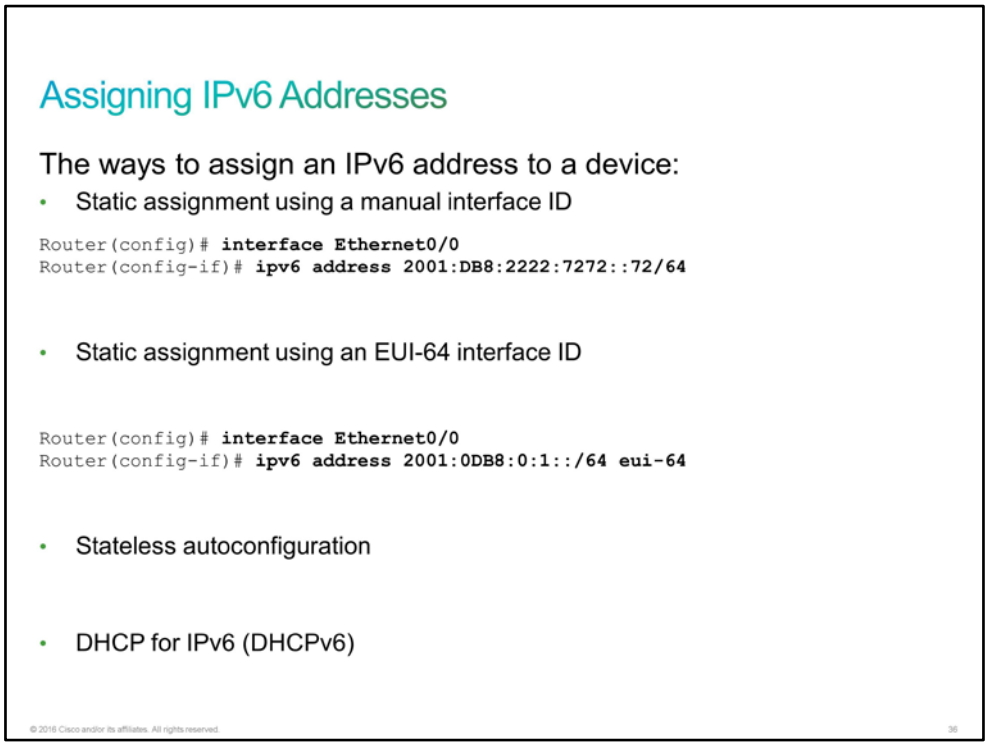

- **Static assignment using a manual interface ID:** One way to statically assign an IPv6 address to a device is to manually assign both the prefix (network) and interface ID (host) portions of the IPv6 address. To configure an IPv6 address on a Cisco router interface and enable IPv6 processing on that interface, use the **ipv6 address** *ipv6-address/prefix-length* command in the interface configuration mode.
- **Static assignment using an EUI-64 interface ID:** Another way to statically assign an IPv6 address is to configure the prefix (network) portion of the IPv6 address and derive the interface ID (host) portion from the Layer 2 MAC address of the device, which is known as the EUI-64 interface ID.

  To configure an IPv6 address for an interface and enable IPv6 processing on the interface using an EUI-64 interface ID in the low order 64 bits of the address (host), use the **ipv6 address** *ipv6-prefix/prefix-length* **eui-64** command in the interface configuration mode.
- **Stateless autoconfiguration:** As the name implies, *autoconfiguration* is a mechanism that automatically configures the IPv6 address of a node. In IPv6, it is assumed that non-PC devices, and also computer terminals, will be connected to the network. The autoconfiguration mechanism was introduced to enable plug-and-play networking of these devices to help reduce administration overhead.
- DHCPv6: DHCP for IPv6 enables DHCP servers to pass configuration parameters such as IPv6 network addresses to IPv6 nodes. It offers the capability of automatic allocation of reusable network addresses and additional configuration flexibility. This protocol is a stateful counterpart to IPv6 stateless address autoconfiguration (RFC 2462). Devices can use it separately or concurrently with IPv6 stateless address autoconfiguration to obtain configuration parameters.

Use of EUI-64 Format in IPv6 Addresses

The 64-bit interface identifier in an IPv6 address identifies a unique interface on a link. A link is a network medium over which network nodes communicate using the link layer. The interface identifier can also be unique over a broader scope. Often, an interface identifier is the same as or is based on the link layer (MAC) address of an interface. As in IPv4, a subnet prefix in IPv6 is associated with one link.

The EUI-64 standard explains how to stretch IEEE 802 MAC addresses from 48 to 64 bits. The following figure illustrates this process.

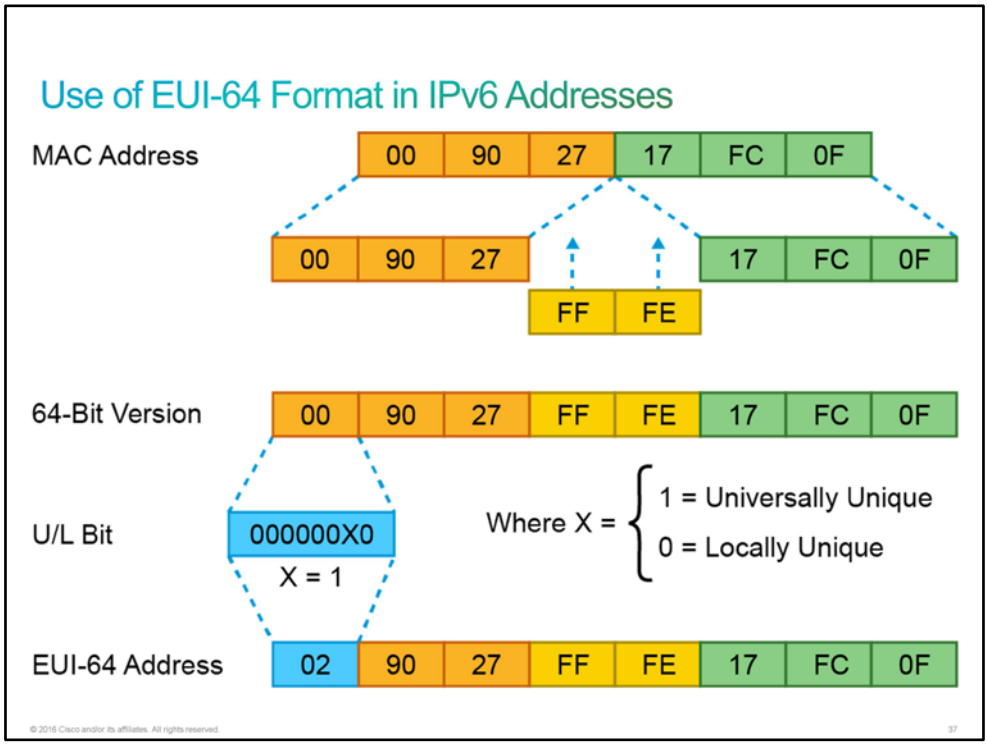

Interface identifiers in the global unicast and other IPv6 address types must be 64 bits long and can be constructed in the 64-bit EUI-64 format. The EUI-64 format interface ID is derived from the 48-bit link layer (MAC) address by inserting the hexadecimal number FFFE between the upper 3 bytes (OUI field) and the lower 3 bytes (serial number) of the link layer address.

Troubleshooting End-to-End IPv6 Connectivity

As with troubleshooting IPv4 connectivity, the troubleshooting process for IPv6 can be guided by structured methods. The overall troubleshooting procedure is the same as troubleshooting IPv4, with differences that are related to IPv6 specifics.

When end-to-end connectivity is not operational, the user will inform the network administrator. The administrator will start the troubleshooting process, as the figure shows.

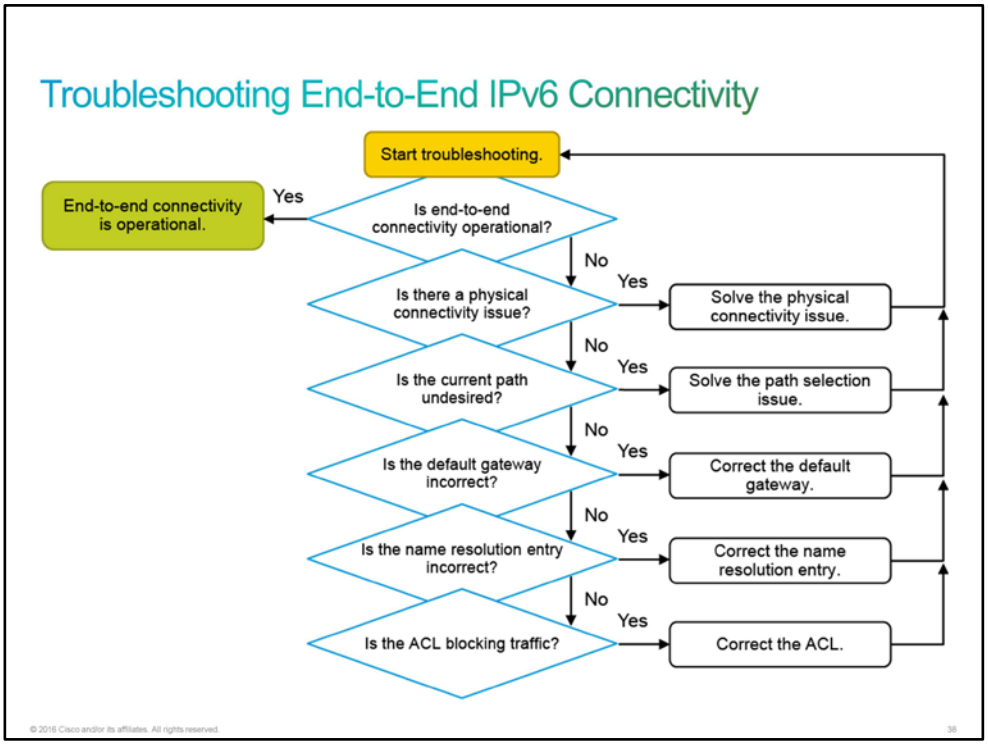

When there is no end-to-end connectivity, you would want to investigate some of the following items:

- If there is an issue with the physical connectivity, solve it by adjusting the configuration or changing the hardware.
- Make sure that devices are determining the correct path from the source to the destination. Manipulate the routing information if needed.
- Verify that the default gateway is correct.
- Check if everything is correct about the name resolution settings. There should be a name resolution server that is accessible over IPv4 or IPv6.
- Verify that there are no ACLs blocking traffic.

After every failed troubleshooting step, a solution should be provided to make the step successful. The outcome of this process is operational, end-to-end connectivity.

Verification of End-to-End IPv6 Connectivity

You can use several verification tools to verify end-to-end IPv6 connectivity:

- **Ping**: A successful ping means that the device endpoints are able to communicate. This result does not mean that there are no problems, but it simply proves that the basic IP connectivity is working.
- **Traceroute**: The results of traceroute can help you determine how far along the path data can successfully reach. Knowing at what point the data fails can help you determine where the issue is.
- **Telnet**: Used to test the transport layer connectivity for any TCP port over IPv6.
- **Neighbor discovery**: Does the same as ARP in IPv4.

In the following scenario, a PC1 wants to access applications on the server. The figure shows the desirable path.

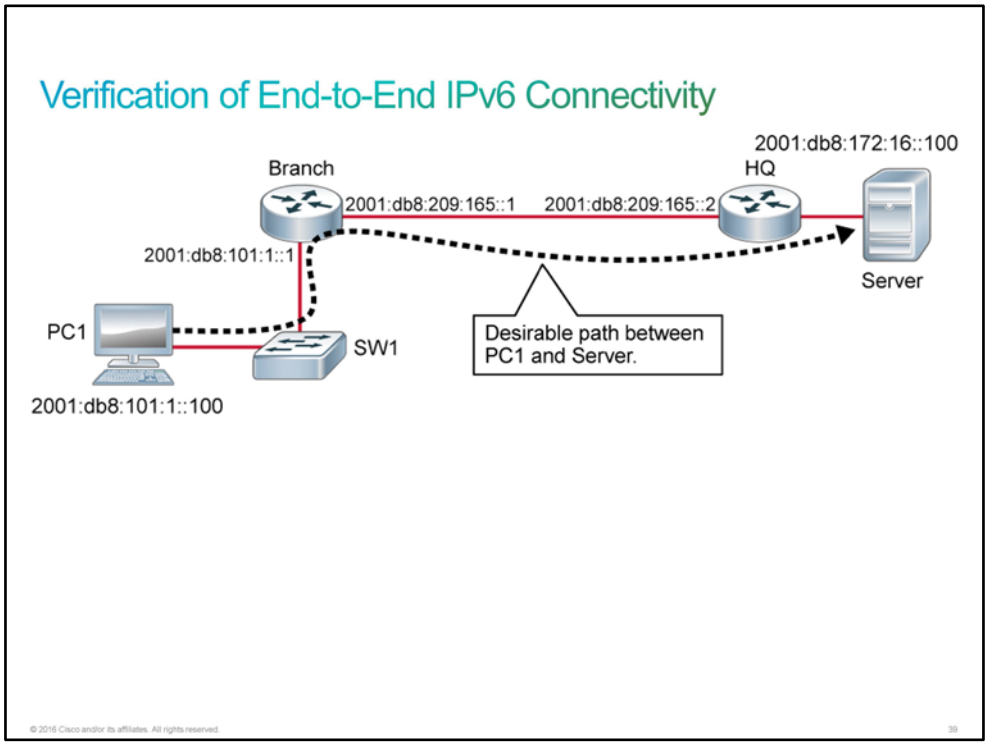

You can use the ping utility to test end-to-end IPv6 connectivity by providing the IPv6 address as the destination address. The utility recognizes the IPv6 address when one is provided and uses IPv6 as a protocol to test connectivity.

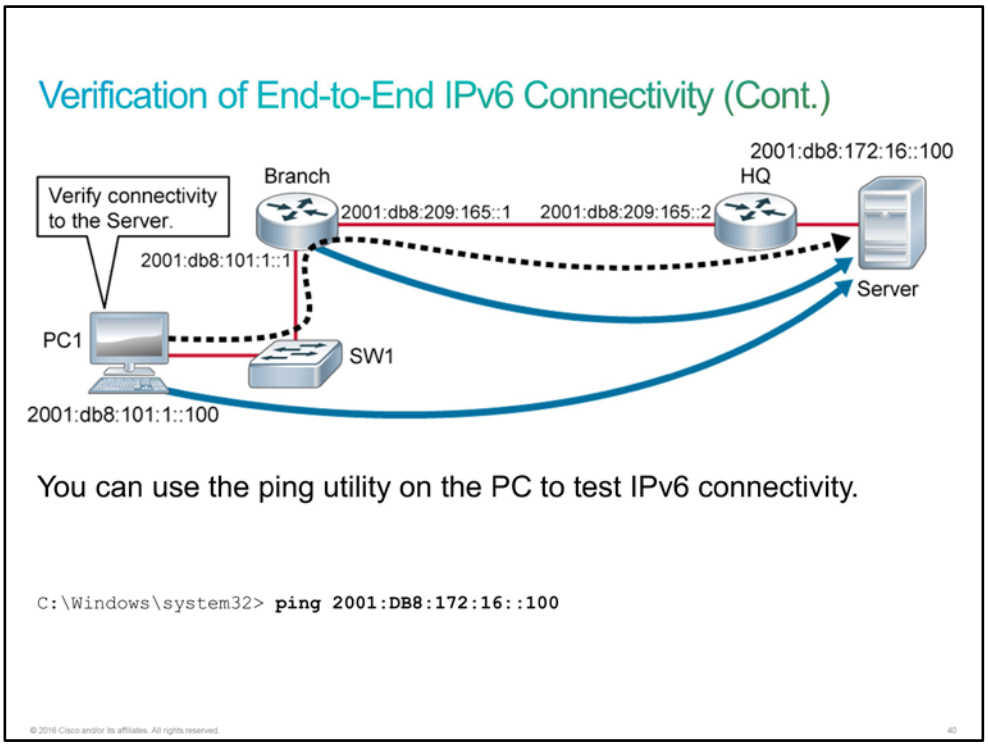

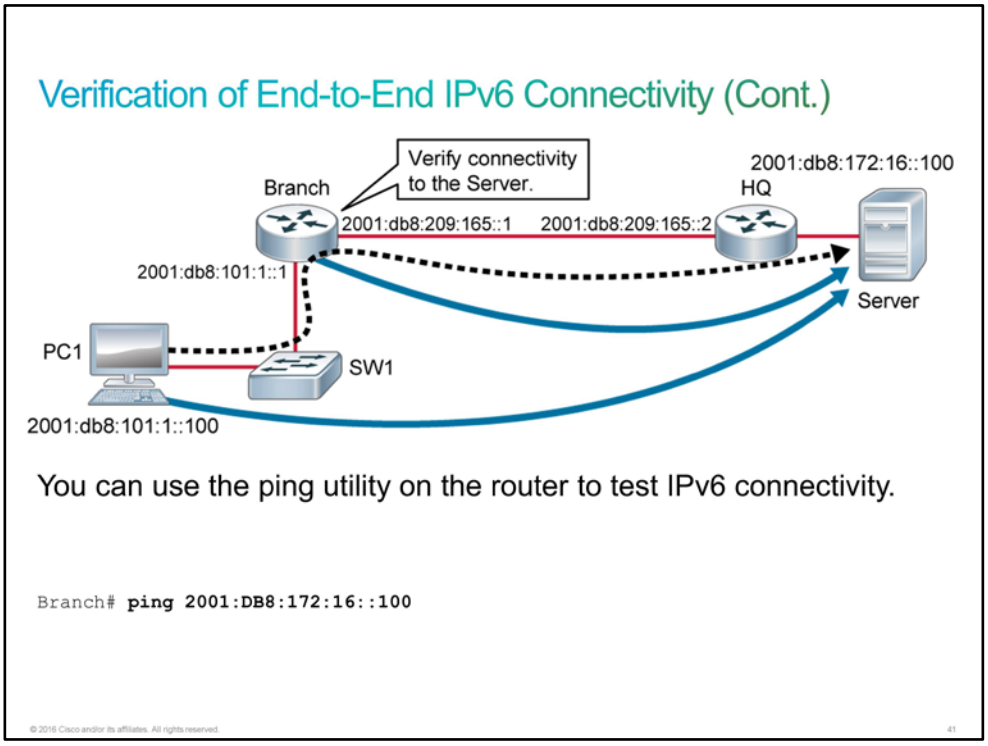

Use the ping utility on the PC to test IPv6 connectivity:

```
C:\Windows\system32> ping 2001:DB8:172:16::100

Pinging 2001:db8:172:16::100 with 32 bytes of data:
Reply from 2001:db8:172:16::100: time=19ms
Reply from 2001:db8:172:16::100: time=1ms
Reply from 2001:db8:172:16::100: time=1ms
Reply from 2001:db8:172:16::100: time=1ms

Ping statistics for 2001:db8:172:16::100:
    Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
Approximate round trip times in milli-seconds:
    Minimum = 1ms, Maximum = 19ms, Average = 5ms
```

You can also use the ping utility on the router to test IPv6 connectivity:

```
Branch# ping 2001:DB8:172:16::100
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2001:DB8:172:16::100, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 0/0/4 ms
```

Traceroute is a utility that allows observation of the path between two hosts and supports IPv6. Use the **traceroute** Cisco IOS command or **tracert** Windows command, followed by the IPv6 destination address, to observe the path between two hosts. The trace generates a list of IPv6 hops that are successfully reached along the path. This list provides important verification and troubleshooting information.

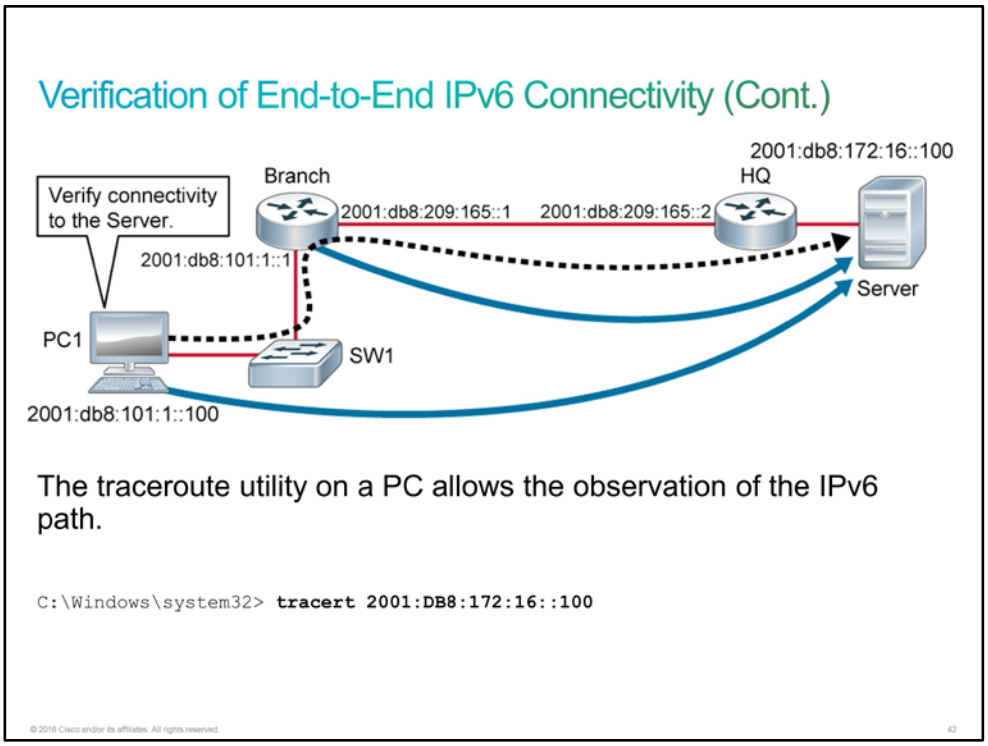

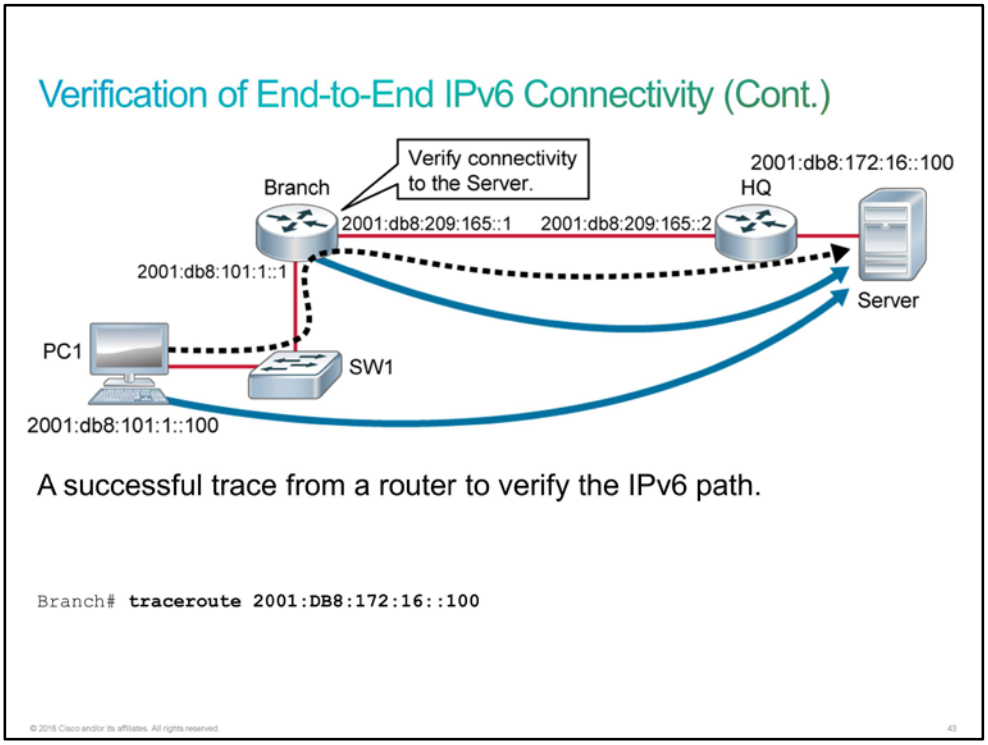

The traceroute utility on the PC allows you to observe the IPv6 path:


```
C:\Windows\system32> tracert 2001:DB8:172:16::100

Tracing route to 2001:db8:172:16::100 over a maximum of 30 hops

  1     1 ms     1 ms    <1 ms  2001:db8:101:1::1
  2    10 ms    1 ms      1 ms  2001:db8:209:165::2
  3    10 ms     1 ms     1 ms  2001:db8:172:16::100

Trace complete.
```

You can also use the traceroute utility on the router to observe the IPv6 path:

```
Branch# traceroute 2001:DB8:172:16::100
Type escape sequence to abort.
Tracing the route to 2001:DB8:172:16::100

  1 2001:DB8:209:165::2 0 msec 0 msec 0 msec
  2 2001:DB8:172:16::100 0 msec 0 msec 0 msec
```

Similar to IPv4, you can use Telnet to test end-to-end transport layer connectivity over IPv6 using the **telnet** command from a PC, router, or a switch. When you provide the IPv6 destination address, the protocol stack determines that the IPv6 protocol has to be used. If you omit the port number, the client will connect to port 23. You can also specify a specific port number on the client and connect to any TCP port that you want to test.

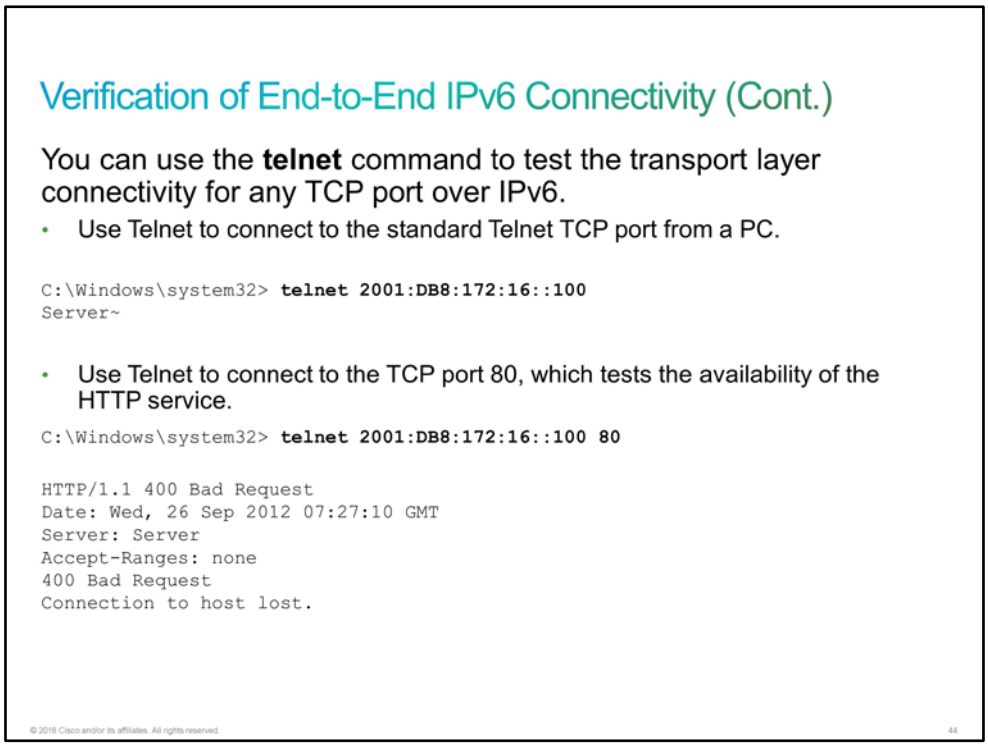

In the example, you can see two connections from a PC to the Server. The first one connects to port 23 and tests Telnet over IPv6. The second connects to port 80 and tests HTTP over IPv6.

When troubleshooting end-to-end connectivity, it is useful to verify mappings between destination IP addresses and Layer 2 Ethernet addresses on individual segments. In IPv4, ARP provides this functionality. In IPv6, the neighbor discovery process and ICMPv6 replace the ARP functionality. The neighbor discovery table caches IP addresses and their resolved Ethernet physical (MAC) addresses. As shown in the figure, the **netsh interface ipv6 show neighbor** Windows command lists all devices that are currently in the neighbor discovery table cache. The information that the CLI displays for each device includes the IP address, physical (MAC) address, and the type of addressing. By examining the neighbor discovery table, you can verify that the destination IPv6 addresses map to the correct Ethernet addresses.

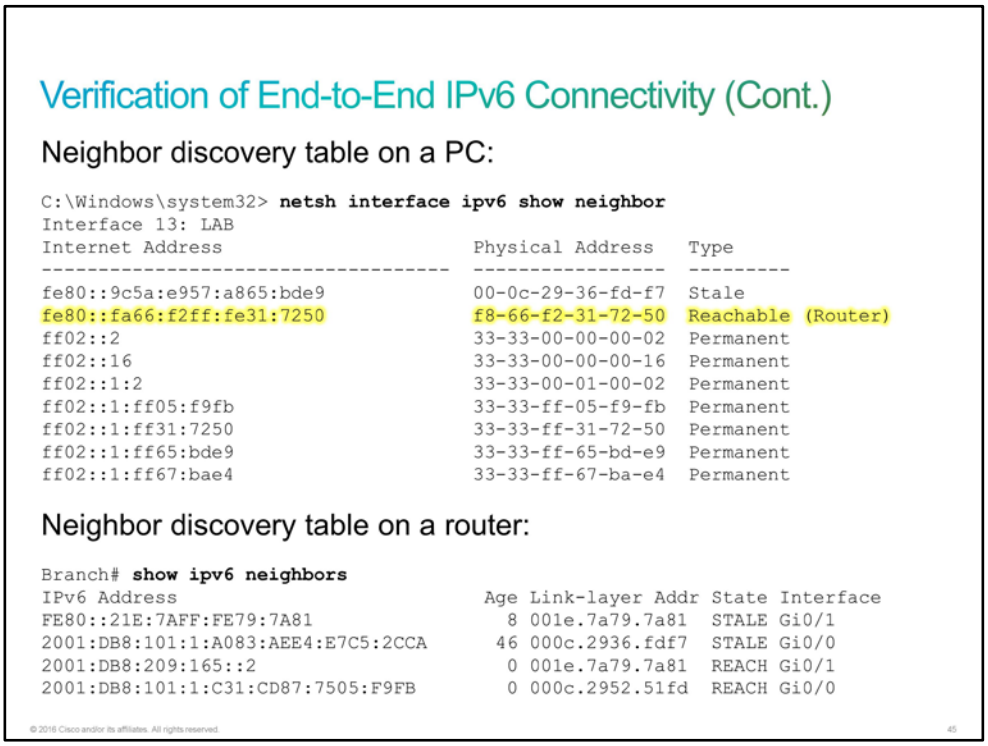

The figure also shows an example of the neighbor discovery table on the Cisco IOS router. The table includes the IPv6 address of the neighbor, age in minutes because the address was confirmed as reachable, and the state. The states are explained in the table:

State	Description
INCMP (Incomplete)	Address resolution is being performed on the entry. The source has sent a neighbor solicitation message to the solicited-node multicast address of the target, but it has not received the corresponding neighbor advertisement message.
REACH (Reachable)	The source has received positive confirmation within the last ReachableTime milliseconds that the forward path to the neighbor was functioning correctly. While in the REACH state, the device takes no special action as it is sending packets.
STALE	More than ReachableTime milliseconds have elapsed since the device received the last positive confirmation that the forward path was functioning properly. While in the STALE state, the device takes no action until a packet is sent.
DELAY	More than ReachableTime milliseconds have elapsed since the device received the last positive confirmation that the forward path was functioning properly. A packet was sent within the last DELAY_FIRST_PROBE_TIME seconds. If the device receives no reachability confirmation within DELAY_FIRST_PROBE_TIME seconds of entering the DELAY state, send a neighbor solicitation message and change the state to PROBE.

State	Description
PROBE	The device actively seeks a reachability confirmation by resending neighbor solicitation messages in RetransTimer milliseconds until a reachability confirmation is received.

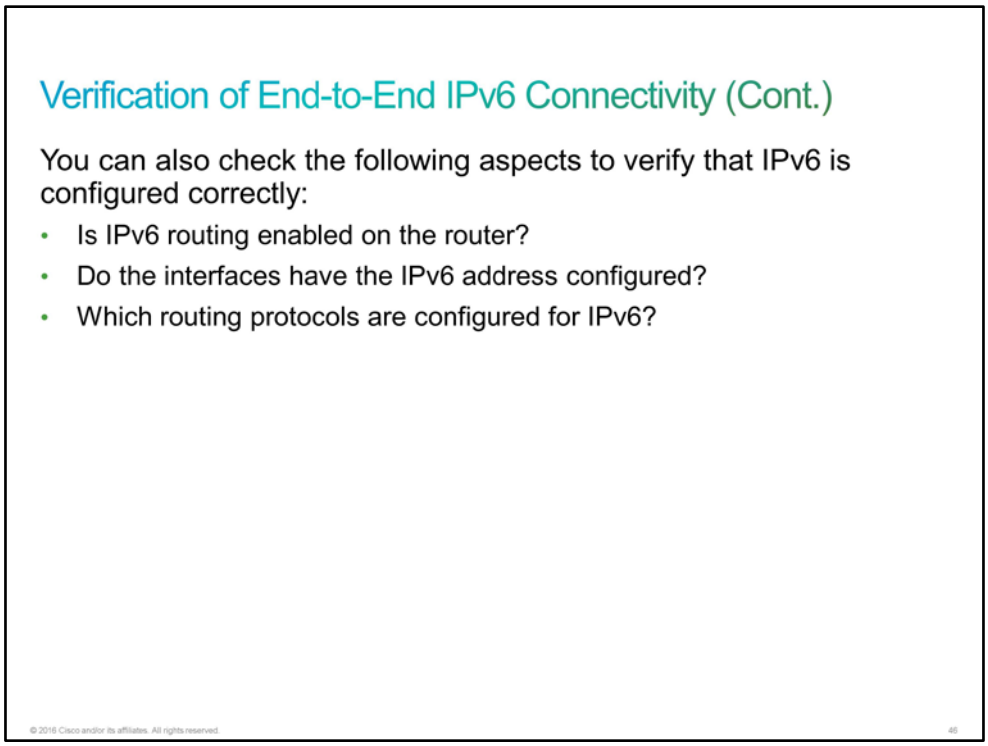

You can use several other commands to verify that IPv6 is configured correctly on routers:

- Verify that IPv6 routing has been enabled on the router. In the **show running-config** command look for the **ipv6 unicast-routing** command.
- Verify that the interfaces have been configured with the correct IPv6 addresses. You can use the **show ipv6 interface** command to display the statuses and configurations for all IPv6 interfaces.
- Verify the IPv6 routing protocols that are running on the router using the **show ipv6 protocols** command.

Identification of Current and Desired IPv6 Path

To verify that the current IPv6 path matches the desired path to reach destinations, use the **show ipv6 route** command on a router to examine the routing table.

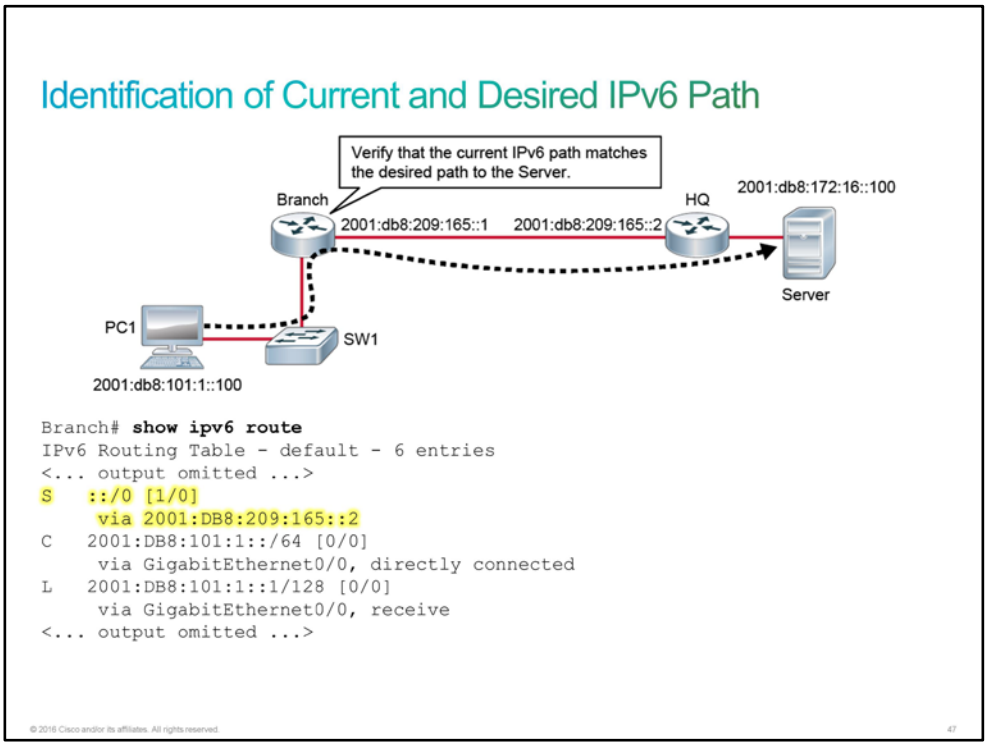

The routing table on the Branch router in the example has a default route that is configured. The router will use it to route packets to the server (2001:db8:172:16::100).

Troubleshooting Default Gateway Issues in IPv6

In the absence of the default gateway on a host, communication between two endpoints in a different network will not work.

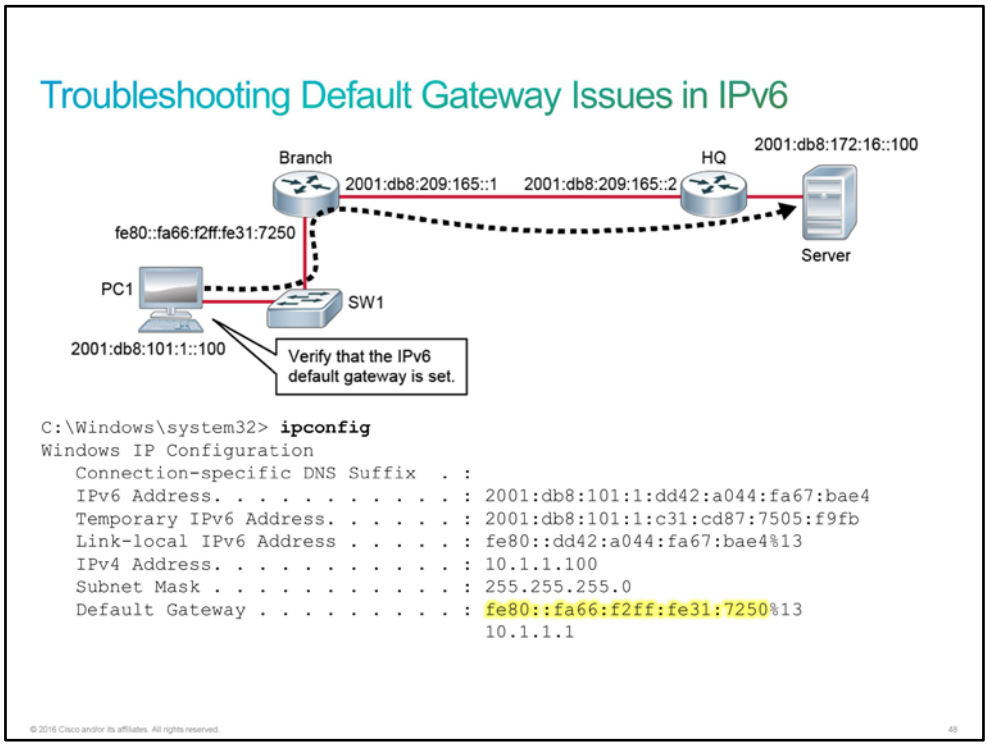

If a PC needs access to other networks in addition to the directly connected network, a correct configuration of the default gateway is very important. If a PC has to send a packet to a network that is *not* directly connected, it has to send the packet to the default gateway, which is the first router on the path to the destinations. The default gateway then forwards the packet toward the destination.

Note	You will see a percent sign (%), followed by a number, at the end of the IPv6 link-local address and at the end of the default gateway. The number that follows the percent sign identifies an interface on the PC and is not part of the IPv6 address. It should be ignored when determining the IPv6 address of the default gateway.

In IPv6, you can manually configure the default gateway or use stateless autoconfiguration.

- In the case of stateless autoconfiguration, the default gateway is advertised to PCs that are using route advertisements. In IPv6, the IPv6 address that the device advertises inside route advertisements as a default gateway is the link-local IPv6 address of a router interface.
- If you decide to configure the default gateway, which is unlikely, you can set the default gateway either to the global IPv6 address or to the link-local IPv6 address.

Note	A link-local address is intended only for communications within the segment of a local network or a point-to-point connection that a host is connected to. The link-local IPv6 addresses are assigned with the fe80::/64 prefix.

To verify that a PC has the default gateway set, you can use the **ipconfig** command on a Microsoft Windows PC or the **ifconfig** command on Linux and Mac OS X. In the example, the PC has the IPv6 default gateway set to the link-local address of the Branch router.

Troubleshooting Name Resolution Issues in IPv6

Because IPv6 networks are long and difficult to remember, DNS is even more important for IPv6 than for IPv4.

The hosts file serves the function of translating human-friendly host names into IPv6 addresses that identify and locate a host in an IPv6 network. In some operating systems, the hosts file content is preferred over other methods, such as the DNS. Unlike the DNS, the hosts file is under the direct control of the local computer administrator.

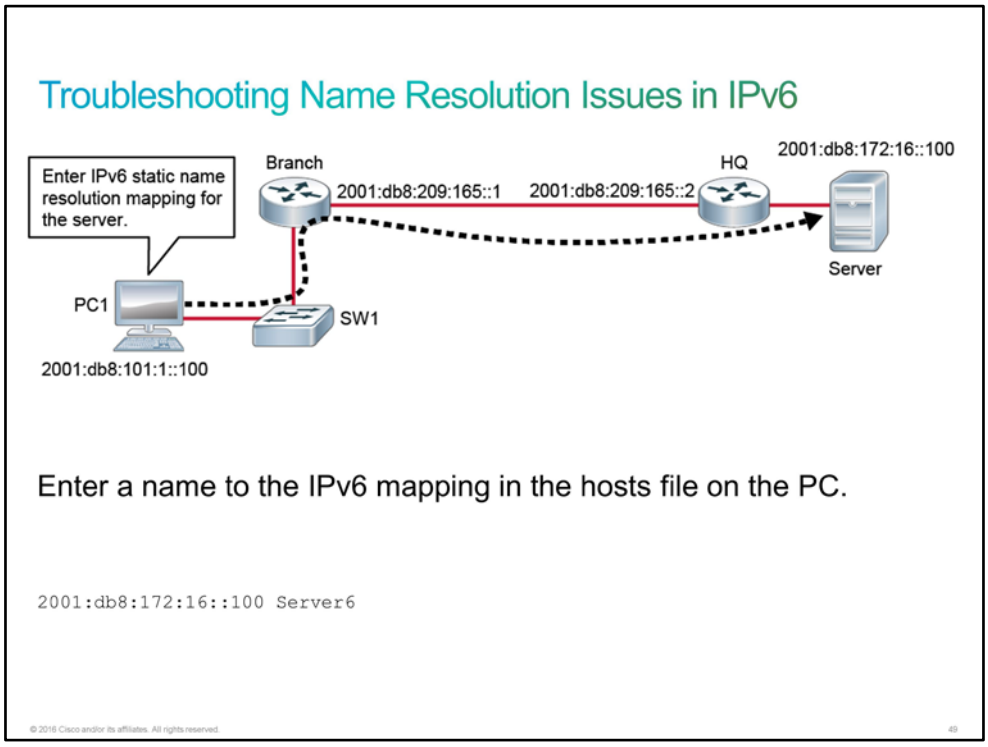

For a Windows operating system, the file is located at *C:\Windows\System32\drivers\etc\hosts*. Other operating systems may have the hosts file in a different location, or they may use a different file, or may not have it at all. You can open the hosts file in a text editor such as Notepad.

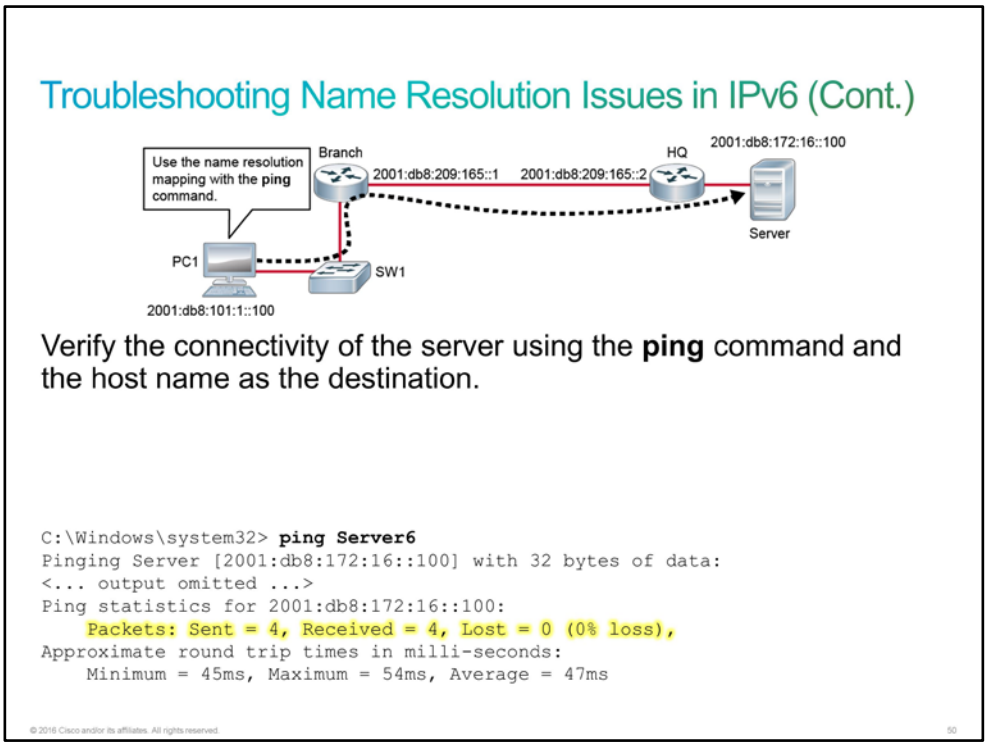

To verify the static name resolution, verify the connectivity to the server using the host name **Server6** instead of its IPv6 address. The ping should be successful.

Discovery 33: Configure and Verify IPv6 Extended Access Lists

Introduction

This discovery will guide you through the extended IPv6 ACLs configuration. The virtual lab environment 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 IPv6 addresses. The configuration of ACL will be on R1 and it will be applied inbound on the interface Ethernet0/0, to influence traffic from PC2.

Note	The policy that is defined in the ACL was chosen to demonstrate how ACLs work. The policy does not reflect any real world application.

Topology

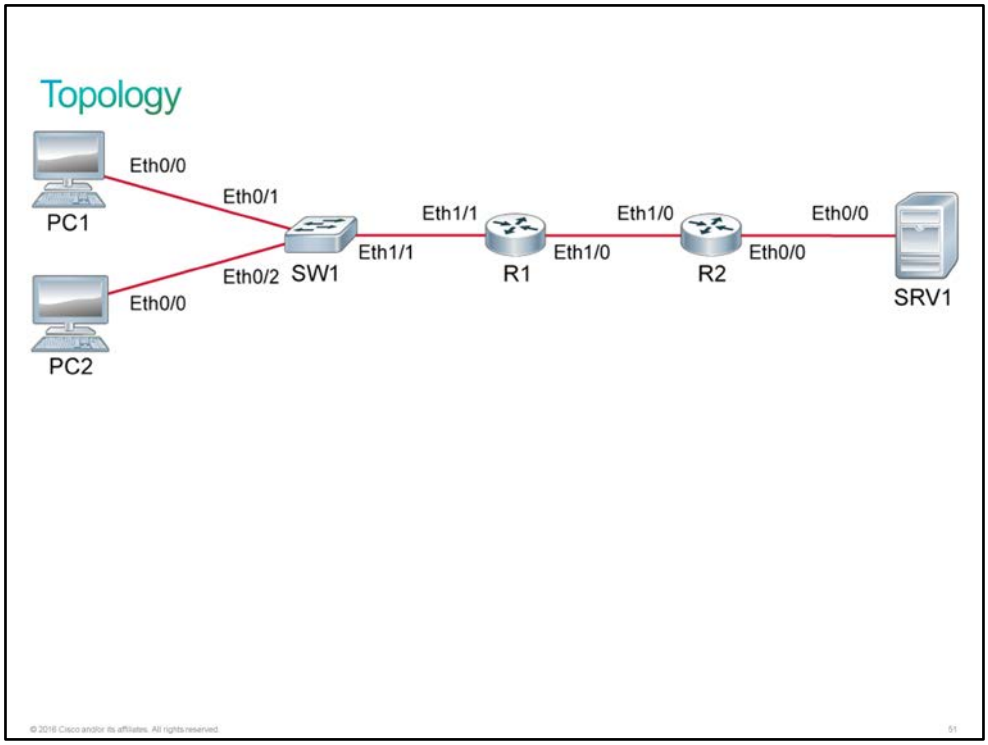

Job Aid

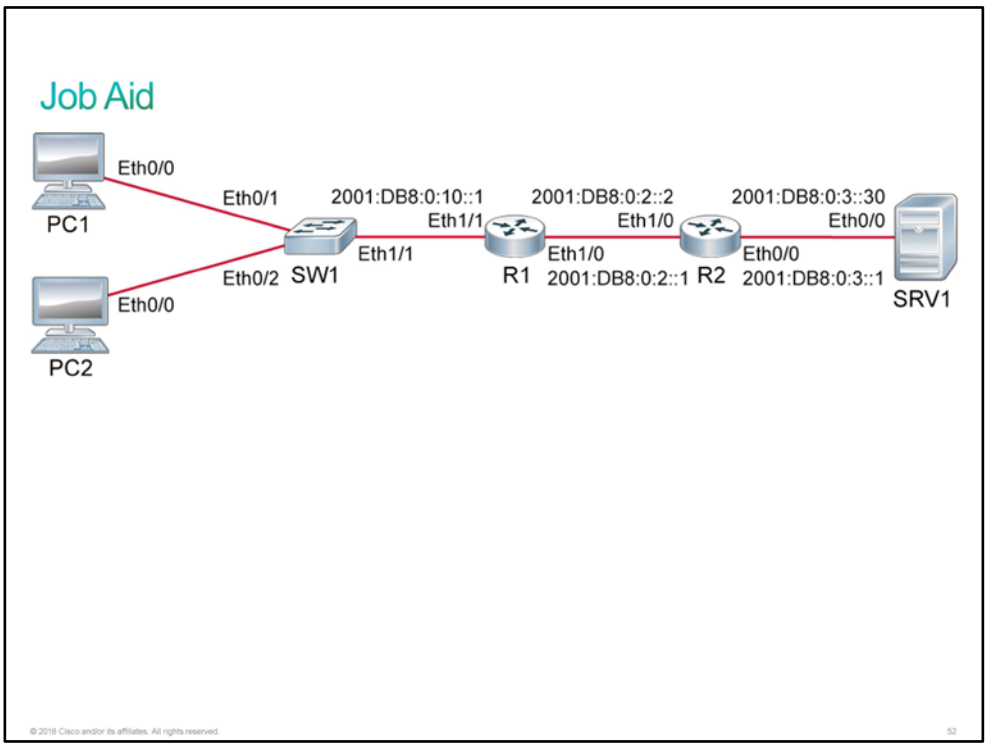

The configuration is as follows:

- All devices have their basic configurations in place including hostnames, IPv4, and IPv6 addresses.
- RIP is configured on R1 and R2 to provide IPv4 routing.
- Static routes are configured on R1 and R2 to provide IPv6 routing.

Device Details

Device	Interface	Neighbor	IPv4 Address	IPv6 Address
PC1	Ethernet0/0	SW1	10.10.1.10/24	2001:DB8:0:10::/64 Auto
PC2	Ethernet0/0	SW1	10.10.1.20/24	2001:DB8:0:10::/64 Auto
SRV1	Ethernet0/0	R2	10.10.3.30/24	2001:DB8:0:3::30/64
SW1	VLAN 1		10.10.1.4/24	2001:DB8:0:10::/64 Auto
SW1	Ethernet0/1	PC1	—	—
SW1	Ethernet0/2	PC2	—	—
SW1	Ethernet1/1	R1	—	—
R1	Ethernet1/1	SW1	10.10.1.1/24	2001:DB8:0:10::1/64

Device	Interface	Neighbor	IPv4 Address	IPv6 Address
R1	Ethernet1/0	R2	10.1.1.2/30	2001:DB8:0:2::1/64
R2	Ethernet1/0	R1	10.1.1.1/30	2001:DB8:0:2::2/64
R2	Ethernet0/0	SRV1	10.10.3.1/24	2001:DB8:0:3::1/64

Note	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 IPv6 Extended Access Lists

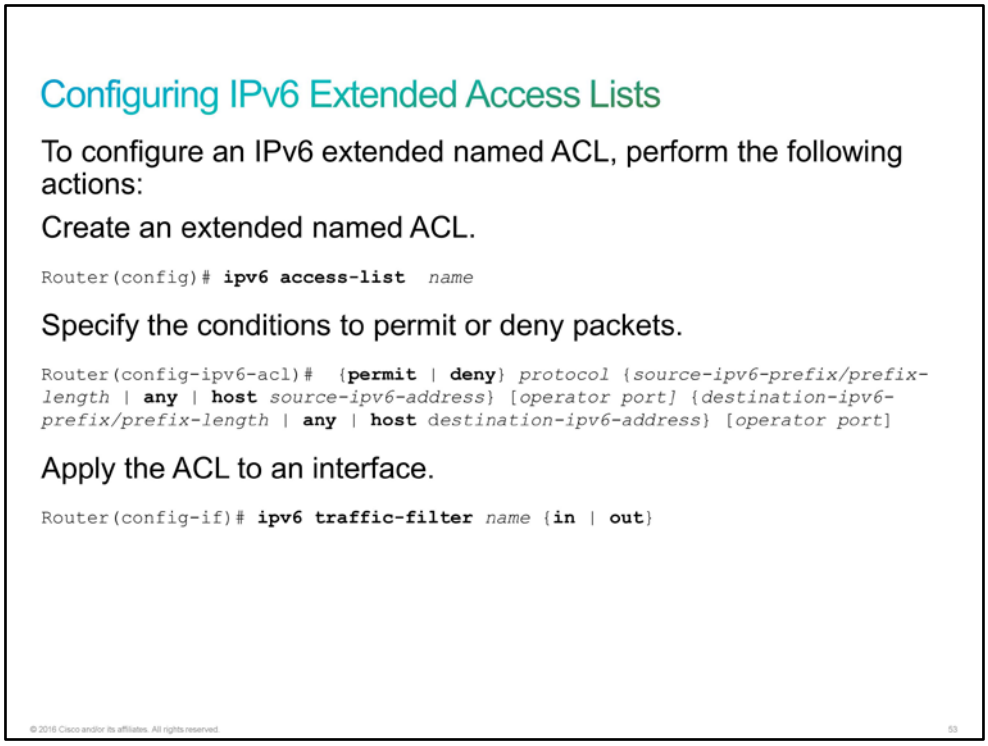

The examples show the steps to configure IPv6 ACL. The following table explains the commands that you will use in the configuration.

Command	Description
ipv6 access-list *name*	Defines an IPv6 access list using a name and enters the IPv6 access list configuration mode.
{**permit** | **deny**} *protocol* {*source-ipv6-prefix/prefix-length* | **any** | **host** *source-ipv6-address*} [*operator port*] {*destination-ipv6-prefix/prefix-length* | **any** | **host** *destination-ipv6-address*} [*operator port*]	Specifies permit or deny conditions for an IPv6 ACL.
ipv6 traffic-filter *name* {**in** | **out**}	Applies the specified access list to the interface in the inbound or outbound direction.

Note	Each IPv6 ACL has implicit permit rules to enable IPv6 neighbor discovery (**permit icmp any any nd-na** and **permit icmp any any nd-ns**). IPv6 ACLs implicitly allow IPv6 neighbor discovery packets to be sent and received on an interface. At the end of this implicit permit there is an implicit deny any rule (**deny ipv6 any any**).

Activity

Complete the following steps:

Step 1 Access the console on R1 and configure a named extended IPv6 ACL. The ACL should be named "Example6".

The ACL should have the following four statements:

- The first should deny all UDP traffic.
- The second should permit TCP from PC2 to any destination as long as the destination port is 23 (Telnet).
- The third should deny all other TCP traffic from PC2.
- The last should explicitly permit all IPv6 traffic.

First, you need to access the console of PC2 and obtain its IPv6 address.

```
PC2# show ipv6 interface brief
Ethernet0/0            [up/up]
    FE80::A8BB:CCFF:FE00:1900
    2001:DB8:0:10:A8BB:CCFF:FE00:1900
Ethernet0/1            [administratively down/down]
    unassigned
Ethernet0/2            [administratively down/down]
    unassigned
Ethernet0/3            [administratively down/down]
    unassigned
```

Note: The IPV6 address in your output may differ, so make sure you will use your IPv6 address, not the one provided in this output!

Now configure the specified ACL on R1.

```
R1# conf t
R1(config)# ipv6 access-list Example6
R1(config-ipv6-acl)# deny udp any any
R1(config-ipv6-acl)# permit tcp host 2001:DB8:0:10:A8BB:CCFF:FE00:1900 any eq
23
R1(config-ipv6-acl)# deny tcp host 2001:DB8:0:10:A8BB:CCFF:FE00:1900 any
R1(config-ipv6-acl)# permit ipv6 any any
R1(config-ipv6-acl)# exit
```

Step 2 Apply the ACL to the interface Ethernet1/1 in the inbound direction.

At the end, leave the configuration mode.

```
R1(config)# interface Ethernet1/1
R1(config-if)# ipv6 traffic-filter Example6 in
R1(config-if)# end
```

Step 3 Display the configured IPv6 ACL.

```
R1# show ipv6 access-list Example6
IPv6 access list Example6
    deny udp any any sequence 10
    permit tcp host 2001:DB8:0:10:A8BB:CCFF:FE00:1900 any eq telnet sequence 20
    deny tcp host 2001:DB8:0:10:A8BB:CCFF:FE00:1900 any sequence 30
    permit ipv6 any any sequence 40
```

The access list has all four statements in the correct order as you have configured them. Note that the output does not display the implicit permit statements for neighbor discovery, and deny any statement that is at the end of every ACL.

Step 4 The first line of the ACL will block all UDP traffic. SRV1 is configured as the NTP server, but because NTP uses the UDP protocol, the first line in the ACL should block IPv6 access for PC2. To verify it, access the console of PC2 and configure it to use the SRV1 IPv6 address as an NTP server. Then display the status of NTP on PC2.

```
PC2# conf t
PC2(config)# ntp server 2001:DB8:0:3::30
PC2(config)# end
```

Because NTP traffic from PC2 is blocked, you should find that it has not synchronized to SRV1.

```
PC2# show ntp status
Clock is unsynchronized, stratum 16, no reference clock
nominal freq is 250.0000 Hz, actual freq is 250.0000 Hz, precision is 2**10
ntp uptime is 6500 (1/100 of seconds), resolution is 4000
reference time is 00000000.00000000 (00:00:00.000 PST Mon Jan 1 1900)
clock offset is 0.0000 msec, root delay is 0.00 msec
root dispersion is 0.99 msec, peer dispersion is 0.00 msec
loopfilter state is 'FSET' (Drift set from file), drift is 0.000000000 s/s
system poll interval is 8, never updated.
```

Step 5 The second line in the ACL explicitly permits Telnet traffic from PC2. Verify that PC2 can successfully telnet to the SRV1 IPv6 address. Use the username "admin" and password "Cisco123".

At the SRV1 system prompt, use the **exit** command to terminate the connection.

```
PC2# telnet 2001:DB8:0:3::30
Trying 2001:DB8:0:3::30 ... Open

User Access Verification

Username: admin
Password:
SRV1> exit

[Connection to 2001:DB8:0:3::30 closed by foreign host]
PC2#
```

Step 6 The third line in the ACL denies all other TCP traffic from PC2. Verify that PC2 cannot use SSH to reach the SRV1 IPv6 address.

```
PC2# ssh -l admin 2001:DB8:0:3::30
% Destination unreachable; gateway or host down
```

Step 7 The fourth line in the ACL, which explicitly permits all IPv6 traffic, should permit any non-UDP traffic from PC2. Verify that PC2 can ping the server.

```
PC2# ping 2001:DB8:0:3::30
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2001:DB8:0:3::30, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/3/15 ms
```

The first three lines do not explicitly specify the ICMP protocol. So, any ICMP traffic should be permitted by the fourth line in the ACL which explicitly permits all IPv6 traffic that did not match any of the previous lines.

Step 8 The ACL only applies to traffic coming from PC2. Access the console of PC1 and attempt the same test sequence that you did from PC1.

The test uses ICMP and TCP, not UDP. All the tests should succeed.


```
PC1# telnet 2001:DB8:0:3::30
Trying 2001:DB8:0:3::30 ... Open

User Access Verification

Username: admin
Password:
SRV1> exit

[Connection to 2001:DB8:0:3::30 closed by foreign host]
PC1# ssh -l admin 2001:DB8:0:3::30
Password:
SRV1> exit

[Connection to 2001:DB8:0:3::30 closed by foreign host]
PC1# ping 2001:DB8:0:3::30
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2001:DB8:0:3::30, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/3/15 ms
```

Notice that PC1 can both telnet and SSH to the server, whereas PC2 could not because no ACL is applied toward the PC1.

Step 9 Display the ACL again and observe the updated hit counters that are associated with the activity that you just initiated.

```
R1# show ipv6 access-list Example6
IPv6 access list Example6
    deny udp any any (9 matches) sequence 10
    permit tcp host 2001:DB8:0:10:A8BB:CCFF:FE00:1900 any eq telnet (39
matches) sequence 20
    deny tcp host 2001:DB8:0:10:A8BB:CCFF:FE00:1900 any (1 match) sequence 30
    permit ipv6 any any (173 matches) sequence 40
```

Due to the dynamic nature of the lab environment, the hit counters that you observe are likely to differ from what the example shows.

This is the end of the discovery lab.

Troubleshooting ACL Issues in IPv6

Another cause of a malfunction of an IPv6 network can be an ACL misconfiguration.

In the given scenario, the Telnet connection to the server is not working and you need to investigate the ACLs that are configured on the router.

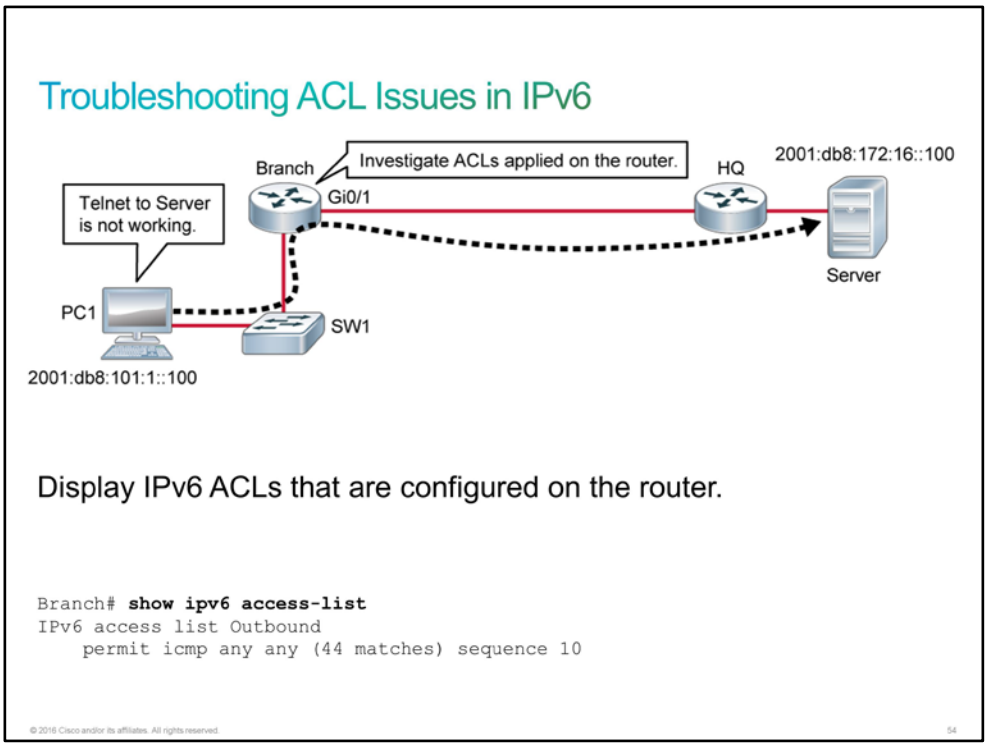

First, you can verify whether there are any IPv6 ACLs configured on a router. You can use the **show ipv6 access-list** command. In the example, an ACL that is named Outbound is configured on the router.

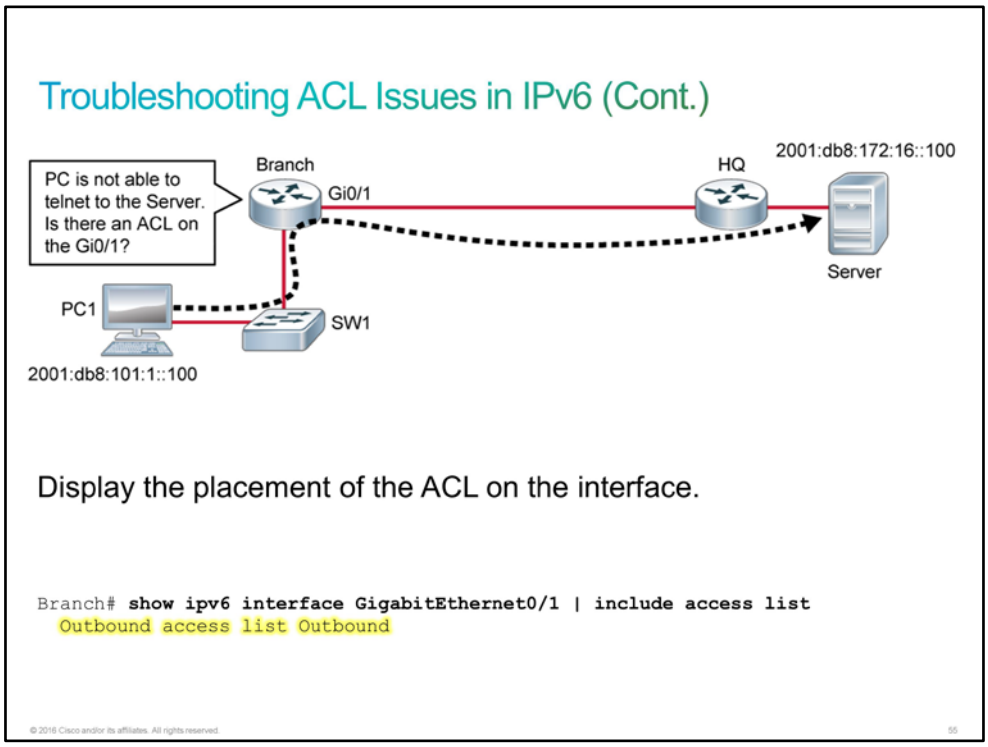

Next, verify if an ACL is attached to an interface. Use the **show ipv6 interface** command.

In the example, an ACL that is named Outbound *is* configured on the router. The ACL is applied to the GigabitEthernet0/1 interface in the outbound direction.

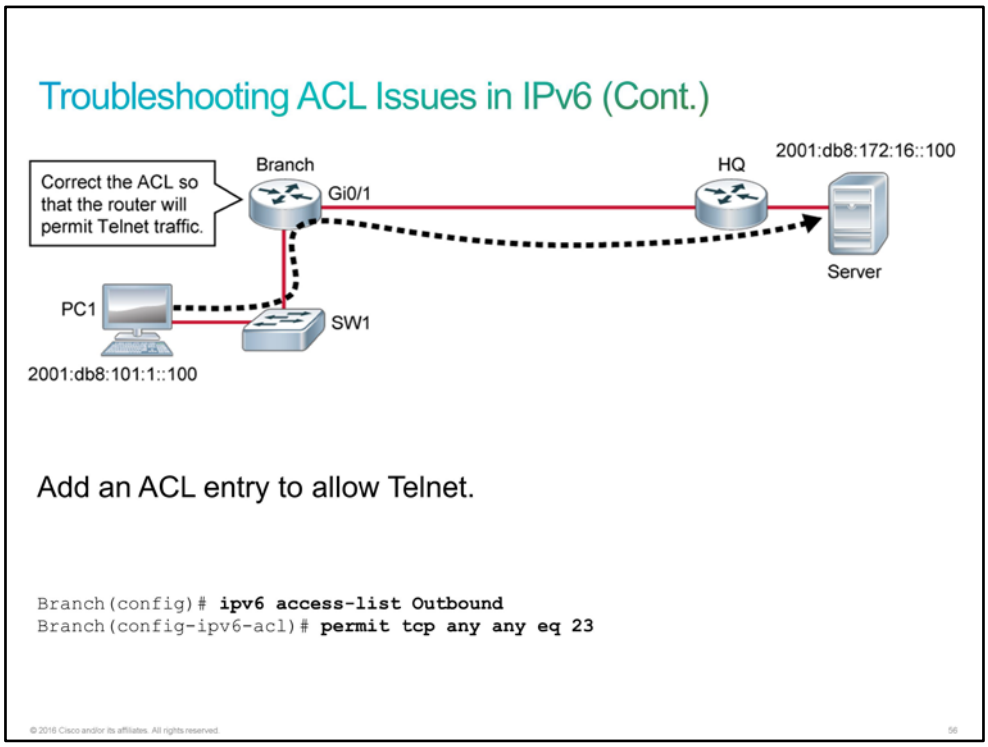

In the example, you verified that an ACL that is named Outbound is configured on the router. The ACL is applied to the GigabitEthernet0/1 interface in the outbound direction. The ACL permits only ICMP protocol, which is why ping will work. In order to allow Telnet from PC1 to the server, you need to add an entry in the Outbound ACL to allow the protocol TCP and port 23 for Telnet.

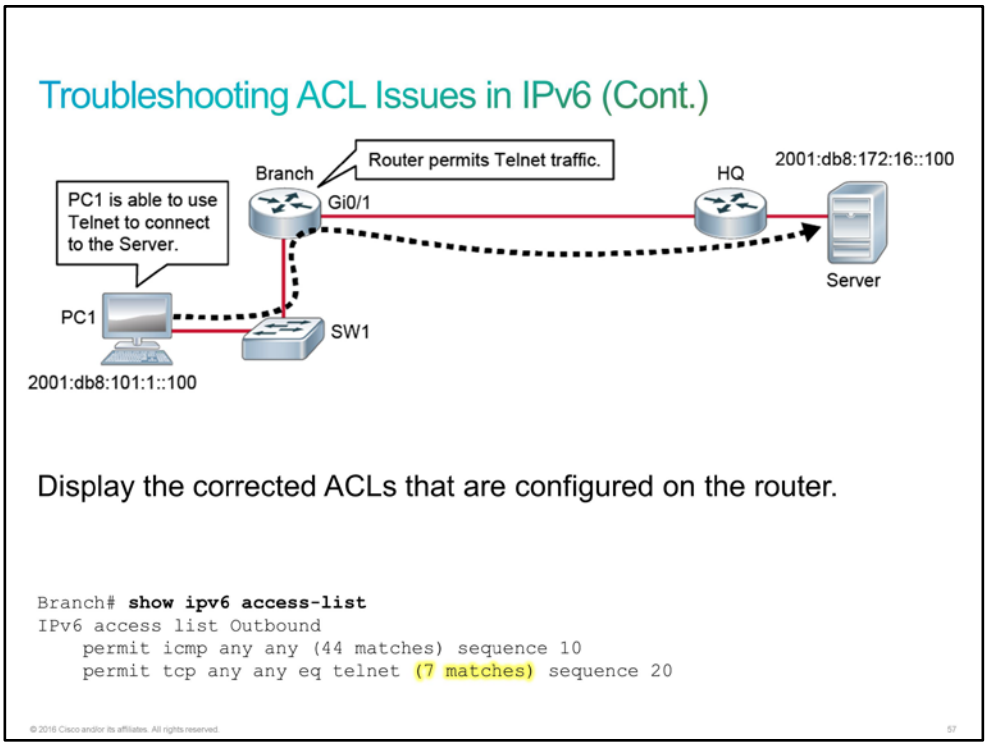

After correcting the ACL, a Telnet connection from PC1 to the server will be successful.

Discovery 34: Troubleshoot IPv6 Network Connectivity

Introduction

This discovery will give you a chance to do some troubleshooting in an IPv6 environment. 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. IPv4 and IPv6 coexist in this network in a dual stack environment. RIP is configured on the routers to provide IPv4 routing. For IPv6, static routes are configured.

Four issues have been introduced on different devices. Your job is to find and fix these issues. There are only four steps in this discovery. A step describes the complaint that you must address. To get the feeling of troubleshooting activities, try to uncover and resolve the problems before you use the Answer Key for each step.

Resolve each issue before moving to the next issue. Sometimes, you may have to resolve a previous issue so that the following issues are demonstrated.

Topology

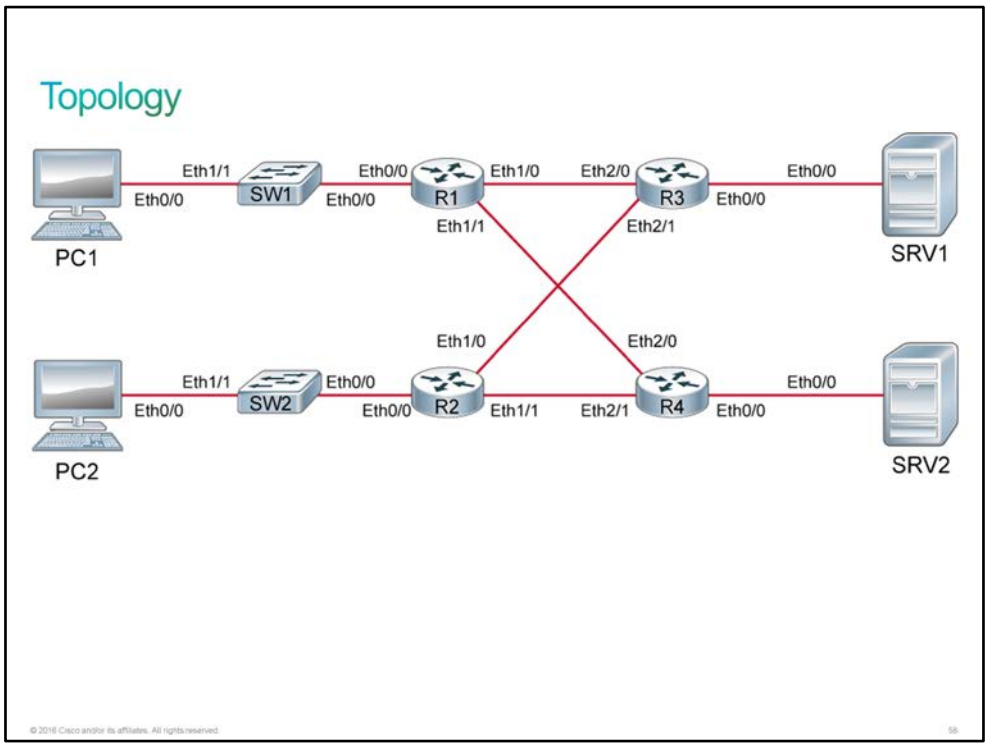

Job Aid

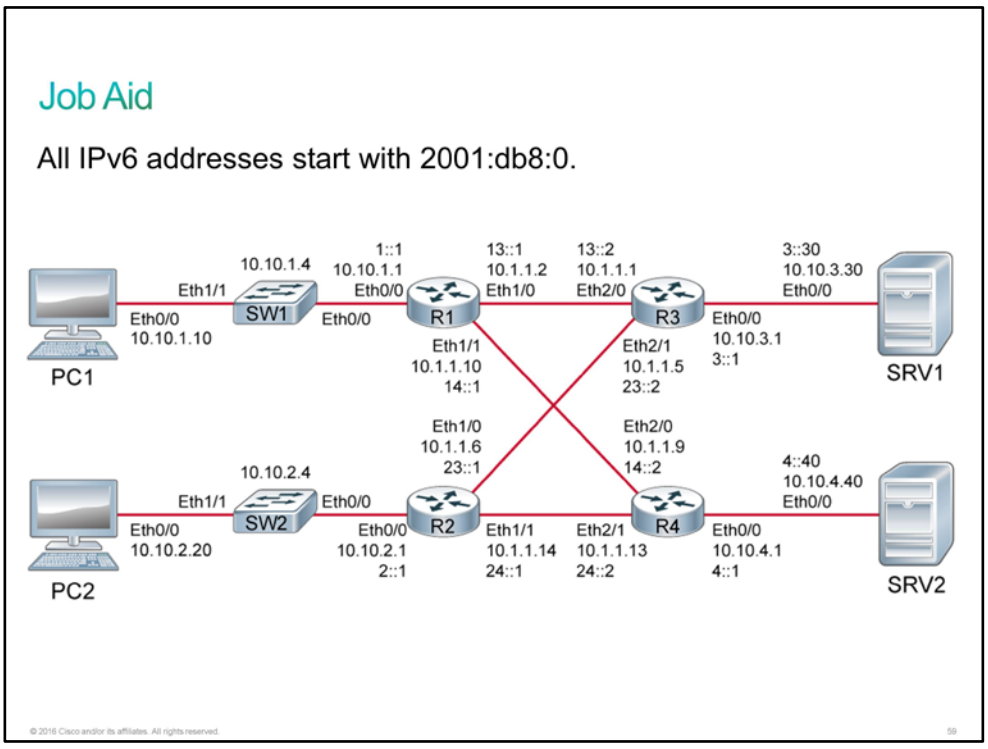

The configuration is as follows:

- All devices have their basic configurations in place, including hostnames, IPv4, and IPv6 addresses.
- RIP is configured on all four routers to provide IPv4 routing.
- Static routes are configured on all four routers to provide IPv6 routing.
- Four issues, related to the PCs connectivity to the SRVs, exist in the network.

Device Information

Device	Interface	Neighbor	IPv4 Address	IPv6 Address
PC1	Ethernet0/0	SW1	10.10.1.10/24	2001:DB8:0:1::/64 Auto
PC2	Ethernet0/0	SW2	10.10.2.20/24	2001:DB8:0:2::/64 Auto
SRV1	Ethernet0/0	R3	10.10.3.30/24	2001:DB8:0:3::30/64
SRV2	Ethernet0/0	R4	10.10.4.40/24	2001:DB8:0:4::40/64
SW1	VLAN 1	—	10.10.1.4/24	2001:DB8:0:1::/64 Auto
SW2	VLAN 1	—	10.10.2.4/24	2001:DB8:0:2::/64 Auto
R1	Ethernet0/0	SW1	10.10.1.1/24	2001:DB8:0:1::1/64

Device	Interface	Neighbor	IPv4 Address	IPv6 Address
R1	Ethernet1/0	R3	10.1.1.2/30	2001:DB8:0:13::1/64
R1	Ethernet1/1	R4	10.1.1.10/30	2001:DB8:0:14::1/64
R2	Ethernet 0/0	SW2	10.10.2.1/24	2001:DB8:0:2::1/64
R2	Ethernet1/0	R3	10.1.1.6/30	2001:DB8:0:23::1/64
R2	Ethernet1/1	R4	10.1.1.14/30	2001:DB8:0:24::1/64
R3	Ethernet2/0	R1	10.1.1.1/30	2001:DB8:0:13::2/64
R3	Ethernet2/1	R2	10.1.1.5/30	2001:DB8:0:23::2/64
R3	Ethernet0/0	SRV1	10.10.3.1/24	2001:DB8:0:3::1/64
R4	Ethernet2/0	R1	10.1.1.9/30	2001:DB8:0:14::2/64
R4	Ethernet2/1	R2	10.1.1.13/30	2001:DB8:0:24::2/64
R4	Ethernet0/0	SRV2	10.10.4.1/24	2001:DB8:0:4::1/64

Note	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 IPv6 Network Connectivity

Activity

Complete the following steps:

Step 1 The user at PC1 is complaining of not being able to connect to SRV1. In fact, if the user attempts to ping SRV1, the ping shows that the server is unreachable.

```
PC1# ping SRV1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2001:DB8:0:4::30, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
```

The IPv6 address of SRV1 that PC1 is pinging is 2001:DB8:0:4::30. However, this address is not the IPv6 address of the server. You can verify it by showing the interface status on SRV1 or by comparing the address to the information in the topology diagram and the connectivity table.

```
SRV1# show ipv6 interface Ethernet0/0
Ethernet0/0 is up, line protocol is up
  IPv6 is enabled, link-local address is FE80::A8BB:CCFF:FE00:4500
  No Virtual link-local address(es):
  Description: Link to R3
  Global unicast address(es):
    2001:DB8:0:3::30, subnet is 2001:DB8:0:3::/64
  Joined group address(es):
    FF02::1
    FF02::1:FF00:30
    FF02::1:FF00:4500
  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 NS retransmit interval is 1000 milliseconds
  Default router is FE80::A8BB:CCFF:FE00:2100 on Ethernet0/0
```

If you ping the SRV1 using the IPv6 address, you will see that the server is reachable.

```
PC1# ping 2001:DB8:0:3::30
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2001:DB8:0:3::30, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
```

So, the problem is an incorrect entry for SRV1 in the local host configuration on PC1.

```
PC1# show running-config | include host
hostname PC1
ipv6 host SRV2 2001:DB8:0:4::40
ipv6 host PC1 2001:DB8:0:1:A8BB:CCFF:FE00:100
ipv6 host PC2 2001:DB8:0:2:A8BB:CCFF:FE00:200
ipv6 host SRV1 2001:DB8:0:4::30
```

You can resolve the problem by configuring the host entry properly.

```
PC1# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
PC1(config)# ipv6 host SRV1 2001:DB8:0:3::30
PC1(config)# end
```

When you configure the entry properly, you should be able to ping SRV1 by hostname from PC1.

```
PC1# ping SRV1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2001:DB8:0:3::30, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
```

Step 2 The user on PC2 is having trouble getting to most network resources. In particular, the user needs to access SRV1.


```
PC2# ping SRV1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2001:DB8:0:3::30, timeout is 2 seconds:

% No valid route for destination
Success rate is 0 percent (0/1)
```

The **ping** command indicates that there is no route to the destination. That is, PC2 does not have a route, not that its gateway indicates that the gateway lacks a route.

```
PC2# show ipv6 route
IPv6 Routing Table - default - 1 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
L   FF00::/8 [0/0]
     via Null0, receive
```

The only entry in the IPv6 routing table for PC2 is the multicast for Null0.

```
PC2# show ipv6 interface Ethernet0/0
Ethernet0/0 is up, line protocol is up
  IPv6 is enabled, link-local address is FE80::A8BB:CCFF:FE00:2600
  No Virtual link-local address(es):
  Description: Link to SW2
  Stateless address autoconfig enabled
  No global unicast address is configured
  Joined group address(es):
    FF02::1
    FF02::2
    FF02::1:FF00:2600
  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.
```

PC2 is configured for stateless autoconfiguration. It needs to see a router advertisement to properly configure its own IPv6 address and gateway assignment. Why does R2 not send the advertisements?

```
R2# show ipv6 interface Ethernet0/0
Ethernet0/0 is administratively down, line protocol is down
  IPv6 is tentative, link-local address is FE80::A8BB:CCFF:FE00:2000 [TEN]
  No Virtual link-local address(es):
  Description: Link to SW2
  Global unicast address(es):
    2001:DB8:0:2::1, subnet is 2001:DB8:0:2::/64 [TEN]
  Joined group address(es):
    FF02::1
    FF02::2
  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 cause is that the Ethernet0/0 interface on R2 (the one facing PC2) is administratively down.

```
R2# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
R2(config)# interface Ethernet0/0
R2(config-if)# no shut
R2(config-if)# end
*Oct 29 09:51:24.134: %LINK-3-UPDOWN: Interface Ethernet0/0, changed state to
up
*Oct 29 09:51:25.139: %LINEPROTO-5-UPDOWN: Line protocol on Interface
Ethernet0/0, changed state to up
```

When you enable the Ethernet0/0 interface, stateless autoconfiguration works properly on PC2. Note that IPv6 address in your output may be different.


```
PC2# show ipv6 interface Ethernet0/0
Ethernet0/0 is up, line protocol is up
  IPv6 is enabled, link-local address is FE80::A8BB:CCFF:FE00:2600
  No Virtual link-local address(es):
  Description: Link to SW2
  Stateless address autoconfig enabled
  Global unicast address(es):
    2001:DB8:0:2:A8BB:CCFF:FE00:2600, subnet is 2001:DB8:0:2::/64 [EUI/CAL/PRE]
      valid lifetime 2591988 preferred lifetime 604788
  Joined group address(es):
    FF02::1
    FF02::2
    FF02::1:FF00:2600
  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.
```

PC2 should now have access to network resources (SRV1 in this case).

```
PC2# ping SRV1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2001:DB8:0:3::30, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
```

Step 3 The user at PC2 is much happier now because of being able to access SRV1. However, the user is still having difficulty reaching SRV2. Connectivity is terrible. When the user attempts to ping SRV2, half of the packets time out.

```
PC2# ping SRV2
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2001:DB8:0:4::40, timeout is 2 seconds:
.!.!.
Success rate is 40 percent (2/5), round-trip min/avg/max = 1/1/1 ms
```

When the packets consistently alternate between success and timeout, it indicates that there is load balancing going on at a point where one of the paths is valid and the other path is not. Where might this be? You know that the path from PC2 to SRV2 should traverse R2 and R4.

Observe R2 to determine if R2 is the point where load balancing occurs.

```
R2# show ipv6 route
IPv6 Routing Table - default - 12 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
<... output omitted ...>
C   2001:DB8:0:2::/64 [0/0]
     via Ethernet0/0, directly connected
<... output omitted ...>
S   2001:DB8:0:4::/64 [1/0]
     via 2001:DB8:0:24::2, Ethernet1/1
<... output omitted ...>
```

There is no load balancing going on for the SRV2 subnet on R2. R2 is directly connected to 2001:DB8:0:2::/64 (the PC2 network) and it forwards all traffic for that network directly from Ethernet0/0. R2 also forwards all traffic that is destined to 2001:DB8:0:4::/64 (the SRV2 network) to R4 via Ethernet/1.

Move to R4 and observe if R4 is the point where load balancing occurs.

```
R4# show ipv6 route
IPv6 Routing Table - default - 12 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
<... output omitted ...>
S   2001:DB8:0:2::/64 [1/0]
     via 2001:DB8:0:24::1, Ethernet2/1
<... output omitted ...>
C   2001:DB8:0:4::/64 [0/0]
     via Ethernet0/0, directly connected
<... output omitted ...>
```

The situation is the same on R4. It is directly connected to the SRV2 network and it forwards all traffic to the PC2 network to R2 via Ethernet2/1.

If the problem is not on the routers, it may be on the endpoints.


```
SRV2# 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:DB8:0:4::1
     via 2001:DB8:0:4::2
C   2001:DB8:0:4::/64 [0/0]
     via Ethernet0/0, directly connected
L   2001:DB8:0:4::40/128 [0/0]
     via Ethernet0/0, receive
L   FF00::/8 [0/0]
     via Null0, receive
```

Two static default routes are configured on SRV2. One points to R4 and the other points to a nonexistent address. You have to remove the invalid route.

```
SRV2# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
SRV2(config)# no ipv6 route ::/0 2001:DB8:0:4::2
SRV2(config)# end
```

This should ensure consistent communication between PC2 and SRV2.

```
PC2# ping SRV2
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2001:DB8:0:4::40, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
```

Step 4 Even though you have successfully solved the problems that the user at PC2 had with access to SRV2, now the user at PC1 is complaining about access to the SRV2. SRV2 is running HTTP services on port 80; however, the user is complaining that the web access to the server is not working.

```
PC1# telnet SRV2 80
Translating "SRV2"...domain server (255.255.255.255)
Trying 2001:DB8:0:4::40, 80 ...
% Destination unreachable; gateway or host down
```

```
PC1# ping SRV2
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2001:DB8:0:4::40, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
```

Remember that Telnet uses TCP to test connectivity. By default, it will connect to port 23, but you can also specify other ports. SRV2 on port 80 is not reachable; however, you can see that **ping**, which uses ICMP to test connectivity, to SRV2 is successful.

SRV2 is reachable, so you have to determine who in the network is blocking the connectivity.

Using the **traceroute** command, you can determine which path does the PC1 take to reach SRV2. This action will give you the list of routers to investigate.

```
PC1# traceroute SRV2
Type escape sequence to abort.
Tracing the route to SRV2 (2001:DB8:0:4::40)

  1 2001:DB8:0:1::1 1 msec 1 msec 0 msec
  2 2001:DB8:0:14::2 1 msec 1 msec 1 msec
  3 SRV2 (2001:DB8:0:4::40) 0 msec 1 msec 0 msec
```

PC1 takes the path via R1 and R4 to reach SRV2. Investigate if any of these two routers are blocking the web access to SRV2.

```
R1# show ipv6 access-list
R1#
```

There is no IPv6 access list configured on R1. What about R4?

```
R4# show ipv6 access-list
IPv6 access list Outbound
    deny tcp any host 2001:DB8:0:4::40 eq www (2 matches) sequence 10
    deny tcp any host 2001:DB8:0:4::40 eq 443 sequence 20
    permit tcp any host 2001:DB8:0:4::40 sequence 30
    permit icmp any any (32 matches) sequence 40
    permit tcp any any eq telnet sequence 50
    permit ipv6 any any (3 matches) sequence 60
```

R4 has an IPv6 access list configured that is blocking the www access to SRV2. Verify where this access list is applied.

```
R4# show running-config interface Ethernet2/0
Building configuration...

Current configuration : 124 bytes
!
interface Ethernet2/0
 description Link to R1
 ip address 10.1.1.9 255.255.255.252
 ipv6 address 2001:DB8:0:14::2/64
end

R4# show running-config interface Ethernet0/0
Building configuration...

Current configuration : 158 bytes
!
interface Ethernet0/0
 description Link to SRV2
 ip address 10.10.4.1 255.255.255.0
 ipv6 address 2001:DB8:0:4::1/64
 ipv6 traffic-filter Outbound out
end
```

The IPv6 access list is applied in the outbound direction to the Ethernet0/0 interface, the one connecting to the SRV2.

Note: To see if an access-list is applied to an interface, you could also use the **show ipv6 interface** command. However, the output for IPv6 is not the same as for IPv4—the access-list part appears only if there is an access-list applies to this interface.

To solve the problem, you have to options. You can either remove the first statement in the access list completely, or you can change it to "permit." The first option is shown here.

```
R4# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
R4(config)# ipv6 access-list Outbound
R4(config-ipv6-acl)# no sequence 10
```

Note: The command to add or remove the specific rule from the access-list is not the same as for IPv4 access-lists. You also have to specify the **sequence** keyword before the sequence-number.

PC1 should now be able to connect to SRV2 on port 80—the user should have web access to the server.

```
PC1# telnet SRV2 80
Translating "SRV2"...domain server (255.255.255.255)
Trying 2001:DB8:0:4::40, 80 ... Open
```

This is the end of the discovery lab.

Challenge

1. Which command verifies end-to-end transport layer connectivity for SMTP from a PC over an IPv6 path?

   A. **ping** *IPv6_address 25*
   B. **telnet** *IPv4_address 23*
   C. **telnet** *IPv6_address 25*
   D. **tracert** *IPv6_address*

2. Based on this output, the router is able to send a packet to the server at 2001:db8:172:16::100.

```
Branch#show ipv6 route
IPv6 Routing Table - default - 6 entries
Codes: C - Connected, L - Local, S - Static, U - Per-user Static route
       B - BGP, HA - Home Agent, MR - Mobile Router, R - RIP
       I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea, IS - ISIS summary
       D - EIGRP, EX - EIGRP external, NM - NEMO, ND - Neighbor Discovery
       l - LISP, O - OSPF Intra, OI - OSPF Inter, OE1 - OSPF ext 1
       OE2 - OSPF ext 2,  ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2
C   2001:DB8:101:1::/64 [0/0]
     via GigabitEthernet0/0, directly connected
L   2001:DB8:101:1::1/128 [0/0]
     via GigabitEthernet0/0, receive
C   2001:DB8:209:165::/64 [0/0]
     via GigabitEthernet0/1, directly connected
L   2001:DB8:209:165::1/128 [0/0]
     via GigabitEthernet0/1, receive
L   FF00::/8 [0/0]
     via Null0, receive
```

   A. True
   B. False

3. Which three options are valid representations of the IPv6 address 2035:0001:2BC5:0000:0000:087C:0000:000A? (Choose three.)

   A. 2035:0001:2BC5::087C::000A
   B. 2035:1:2BC5::87C:0:A
   C. 2035:0001:2BC5::087C:0000:000A
   D. 2035:1:2BC5:0:0:87C::A
   E. 2035:1:2BC5::087C:A

4. Which statement is true about the EUI-64 address format of the system ID for stateless autoconfiguration that is used by Cisco?

   A. It is the MAC address plus the Site-Level Aggregator
   B. It is the MAC address plus the ISO OUI
   C. It expands the 48-bit MAC address to 64 bits by inserting FFFE into the middle 16 bits.
   D. It does not follow IEEE standards for uniqueness of the address.
   E. It is only used by Cisco

5. Which command will show that the current IPv6 path matches the desired path to reach destinations?

   A. **show ipv6 address**
   B. **show ipv6 route**
   C. **show ipv6 interface**
   D. **show ipv6 inspect**

6. Which type of IPv6 address is advertised inside route advertisements as a default gateway?

   A. global unicast
   B. loopback
   C. reserved
   D. link-local

7. Which command verifies whether any IPv6 ACLs are configured on a router?

   A. **show ipv6 configuration**
   B. **show ipv6 interface**
   C. **show ipv6 access-list**
   D. **show ipv6 route**

Answer Key

Challenge

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

Module 7: Implementing Network Device Security

Introduction

This module describes the steps that are required to secure local and remote access to network devices. It discusses general recommendations on how to improve device hardening. It describes how to configure syslog and how to safely run debugs on Cisco IOS. This module also describes the different stages of the router bootup process, Cisco IOS File System, and how to manage Cisco IOS images or configuration files. The universality of Cisco IOS images and the idea behind licensing are explained, and students are also shown how to verify the current license and install a new one.

Lesson 1: Securing Administrative Access

Introduction

Your boss sends you to your customer to verify potential security threats. On the customer network devices, you will secure access to a privileged level. The customer may want to know the difference between enabling a password and enabling a secret. The customer may ask you how to secure access to the console line and how to secure remote access by enabling and limiting access to SSH. You will also explain how to protect vty with a standard numbered access control list.

Introduction to Network Device Security

Many forms of security threats have emerged because of the rapid growth of the Internet. Viruses, Trojan horse attacks, malicious hackers, and even the employees of an organization are potential security hazards to corporate networks. These threats have the potential to steal and destroy sensitive corporate data, tie up valuable resources, and inflict major damage due to network downtime. This situation may lead to a cost crisis and cripple the company financially. Security breaches are also encountered more frequently in home or private networks. Everyone has a reason to be concerned.

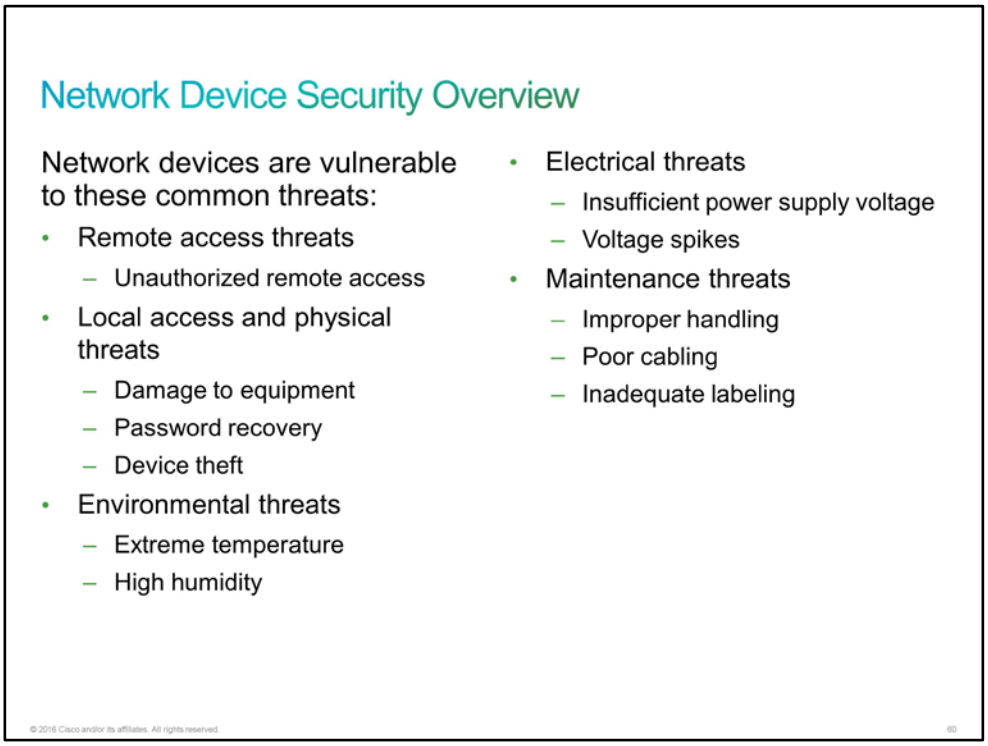

Common threats to network device security and mitigation strategies can be summarized as follows:

- **Remote access threats:** Unauthorized remote access is a threat when security is weak in remote access configuration. Mitigation techniques for this type of threat include configuring strong authentication and encryption for remote access policy and rules, configuration of login banners, use of ACLs, and VPN access.
- **Local access and physical threats:** These threats include physical damage to network device hardware, password recovery that is allowed by weak physical security policies, and device theft. Mitigation techniques for this type of threat include locking the wiring closet and allowing access only to authorized personnel. It also includes blocking physical access through a dropped ceiling, raised floor, window, duct work, or other possible point of entry. Use electronic access control, and log all entry attempts. Monitor facilities with security cameras.
- **Environmental threats:** Temperature extremes (heat or cold) or humidity extremes (too wet or too dry) can present a threat. Mitigation techniques for this type of threat include creating the proper operating environment through temperature control, humidity control, positive air flow, remote environmental alarms, and recording, and monitoring.

- **Electrical threats:** Voltage spikes, insufficient supply voltage (brownouts), unconditioned power (noise), and total power loss are potential electrical threats. Mitigation techniques for this type of threat include limiting potential electrical supply problems by installing UPS systems and generator sets, following a preventative maintenance plan, installing redundant power supplies, and using remote alarms and monitoring.
- **Maintenance threats:** These threats include improper handling of important electronic components, lack of critical spare parts, poor cabling, and inadequate labeling. Mitigation techniques for this type of threat include using neat cable runs, labeling critical cables and components, stocking critical spares, and controlling access to console ports.

Securing Access to Privileged EXEC Mode

You can secure a router or a switch by using passwords to restrict access. Using passwords and assigning privilege levels is a way to provide terminal access control in a network. It is a form of management plane hardening. You can establish passwords on individual lines, such as the console, and to the privileged EXEC mode. Passwords are case-sensitive.

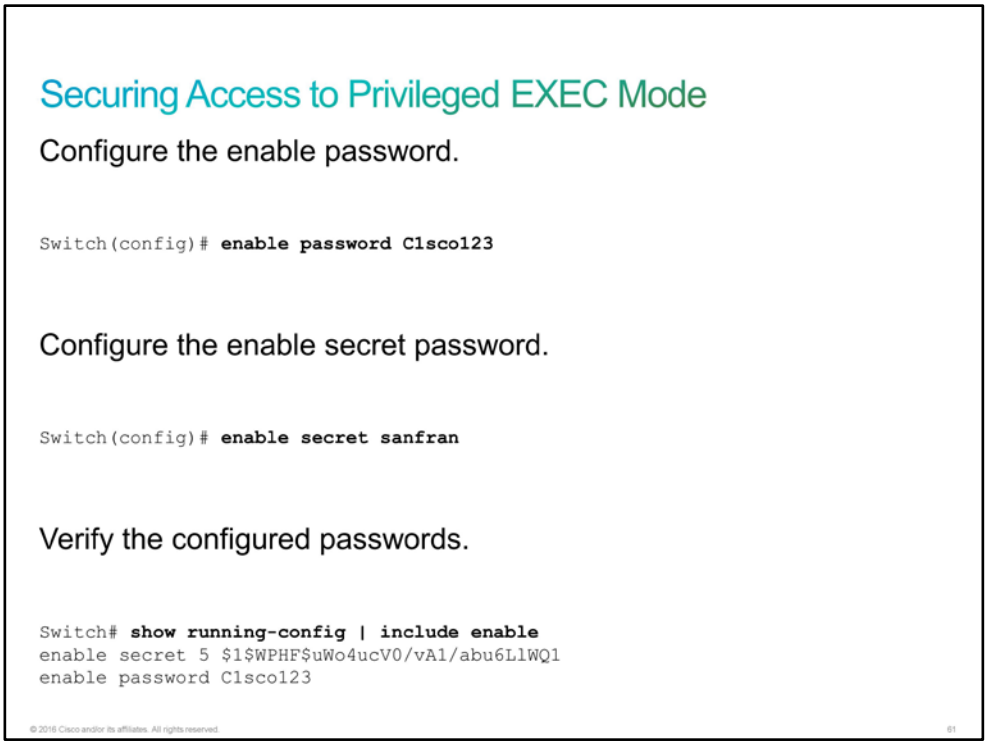

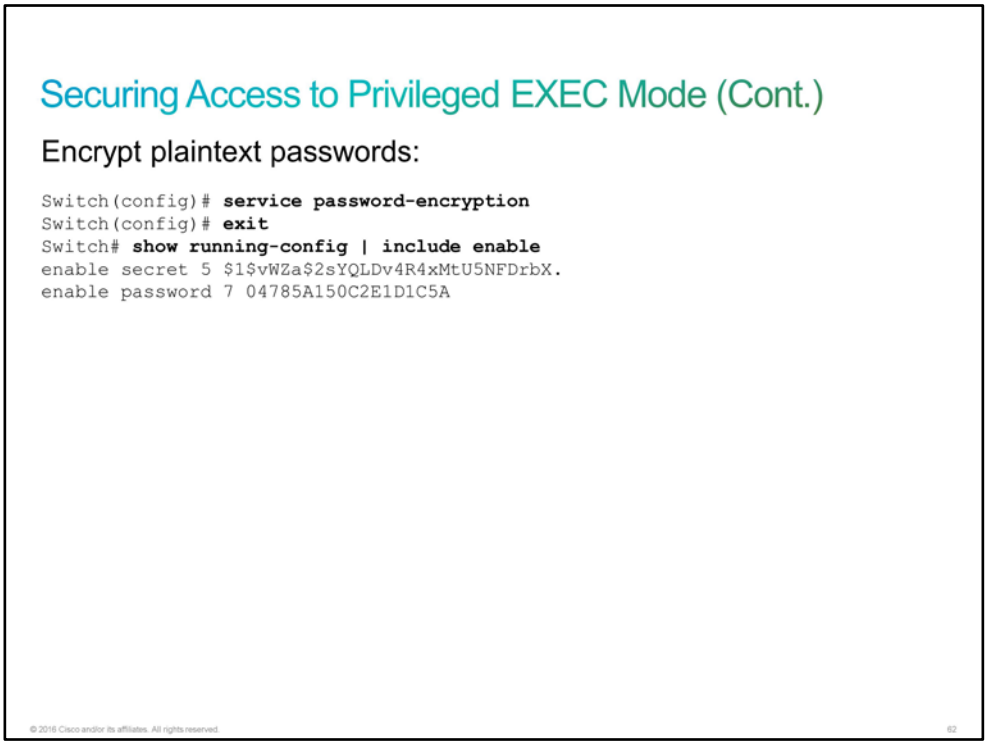

Note	The passwords that the figure shows are for instructional purposes only. Passwords that are used in an actual implementation should meet the requirements of strong passwords.

The **enable password** global command restricts access to the privileged EXEC mode. You can assign an encrypted form of the enable password, which is called the enable secret password, by entering the **enable secret** *password* command at the global configuration mode prompt with the desired password. When you configure the enable secret password, it is used instead of the enable password rather than in addition to it.

You can also add a further layer of security, which is particularly useful for passwords that cross the network or are stored on a TFTP server. Cisco provides a feature that allows the use of encrypted passwords. To set password encryption, enter the **service password-encryption** command in the global configuration mode.

Passwords that are displayed or set after you configure the **service password-encryption** command will be encrypted. Service password encryption uses type-7 encryption, which is not very secure. There are several tools and web pages available that convert an encrypted password into a plaintext string.

On the other hand, the **enable secret** command uses the MD5-type encryption that, to this point, has not been broken. It is recommended that you always use the **enable secret** *password* instead of the **enable password** command.

Securing Console Access

Use the **line console 0** command followed by the **password** and **login** subcommands to require login and establish a login password on a console terminal. By default, logging in is not enabled on the console.

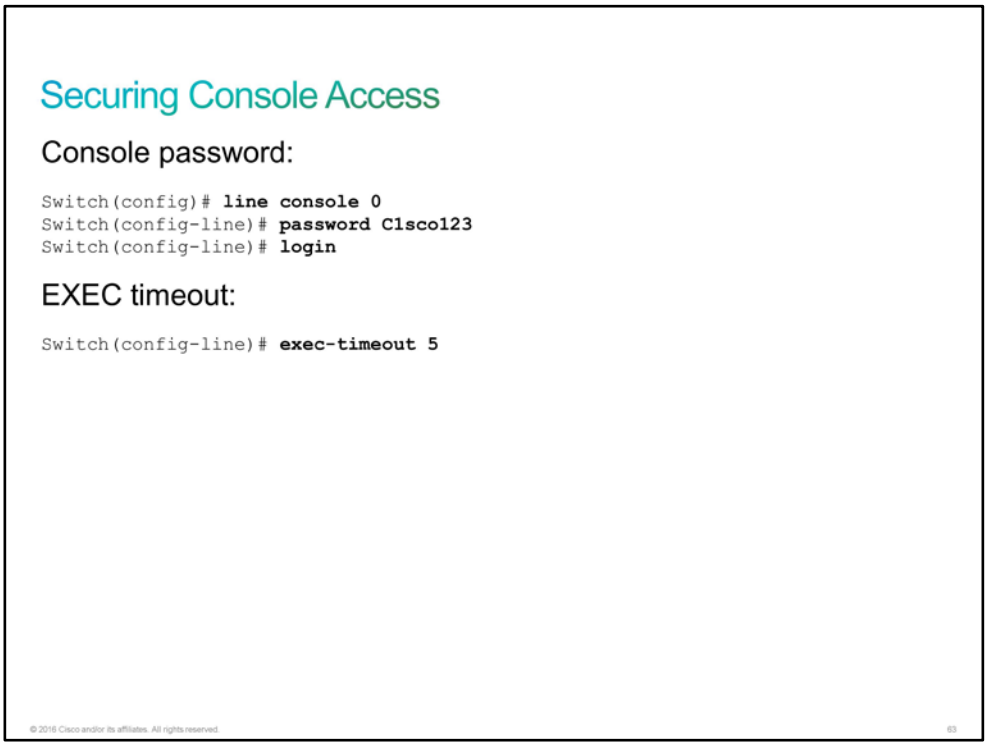

Note	Enter the **service password-encryption** command in the global configuration mode to encrypt the console password. Although this encryption is weak and can be easily decrypted, it is still better than a cleartext password. At least you are protected against exposing the password to casual observers.

The **exec-timeout** command prevents users from remaining connected to a console port when they leave a station. In the example, when no user input is detected on the console for 5 minutes, the user that is connected to the console port is automatically disconnected.

Securing Remote Access

You can establish an SSH connection to the SSH-enabled device using an SSH client on your PC, such as PuTTY. When you establish a connection for the first time from a specific computer, you are presented with a security alert window that indicates that the server host key is not cached in the PuTTY cache. By adding a key to the cache, you will avoid seeing this security alert window every time that you establish an SSH connection from this computer.

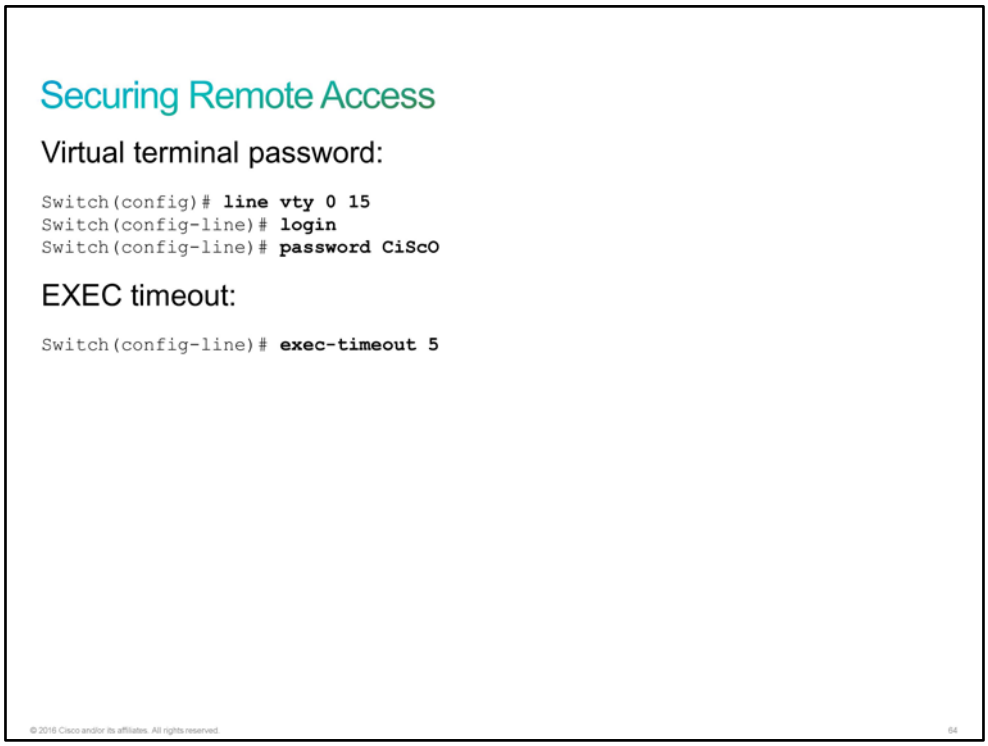

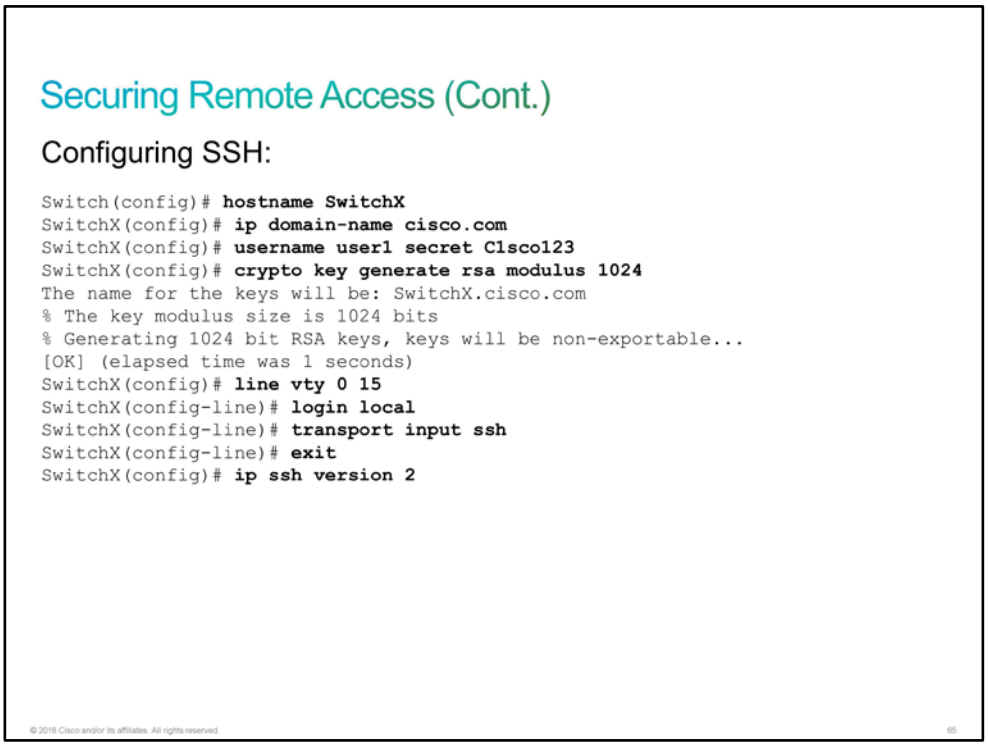

The **line vty 0 15** command, followed by the **login** and **password** subcommands, requires login and establishes a login password on incoming Telnet sessions.

You can use the **login local** command to enable password checking on each user by using the username and secret password that are specified with the **username** global configuration command. The **username** command establishes username authentication with encrypted passwords.

The **exec-timeout** command prevents users from remaining connected to a vty port when they leave a station. In the example, when no user input is detected on a vty line for 5 minutes, the vty session is automatically disconnected.

To configure SSH on a Cisco switch or router, you need to complete the following steps:

1. Use the **hostname** command to configure the hostname of the device so that it is not *Switch* (on a Cisco switch) or *Router* (on a Cisco router).
2. Configure the DNS domain with the **ip domain name** command. The domain name is required to be able to generate certificate keys.
3. Generate RSA keys that the user will use in authentication—use the **crypto key generate rsa** command.
4. Configure the user credentials that the user will use for authentication. By specifying the **login local** command for vty lines, you are essentially telling the network device to use locally defined credentials for authentication. Configure locally defined credentials using the **username** *username* **secret** *password* command.
5. (Optional) You can also limit access to a device to users that use SSH and block Telnet with the **transport input ssh** vty mode command. If you want to support login banners and enhanced security encryption algorithms, force SSH version 2 on your device with the **ssh version 2** command in the global configuration mode.

Securing Remote Access (Cont.)

Verify that SSH is enabled:

```
Switch# show ip ssh
SSH Enabled - version 2.0
Authentication timeout: 120 secs; Authentication retries: 3
```

Check the SSH connection to the device:

```
Switch# show ssh
Connection  Version  Encryption     State            Username
0           2.0      3DES           Session started  cisco
```


Securing Remote Access (Cont.)

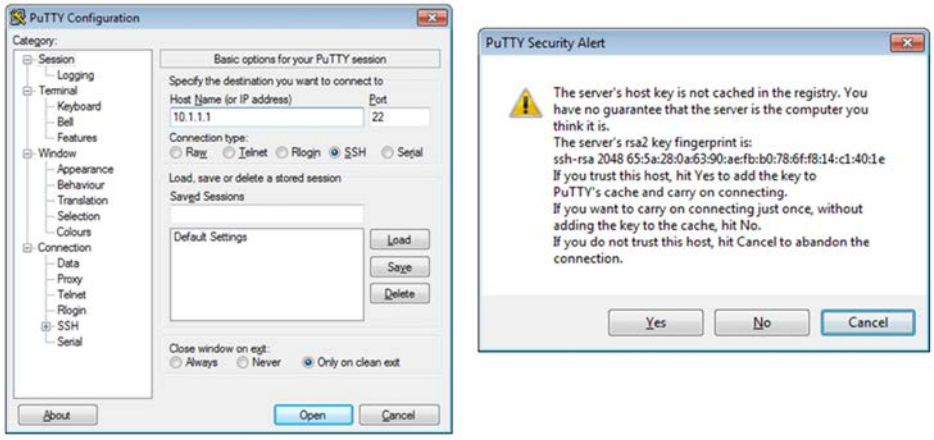

To display the version and configuration data for SSH on the device that you configured as an SSH server, use the **show ip ssh** command. In the example, SSHv2 is enabled.

To check the SSH connection to the device, use the **show ssh** command.

Discovery 35: Enhance Security of Initial Configuration

Introduction

This discovery lab will guide you through the various aspects of securing administrative access to Cisco IOS devices. You will secure access to the privileged EXEC, and see the difference between enable password and enable secret. You will also secure access to the console port. You will enable remote access to the vty lines via Telnet and SSH. You will set SSH as the only acceptable remote access protocol.

The devices are configured as represented in the topology diagram, including their IP addresses. This discovery lab will focus on R1. You will use other devices as sources of remote access connections.

Topology

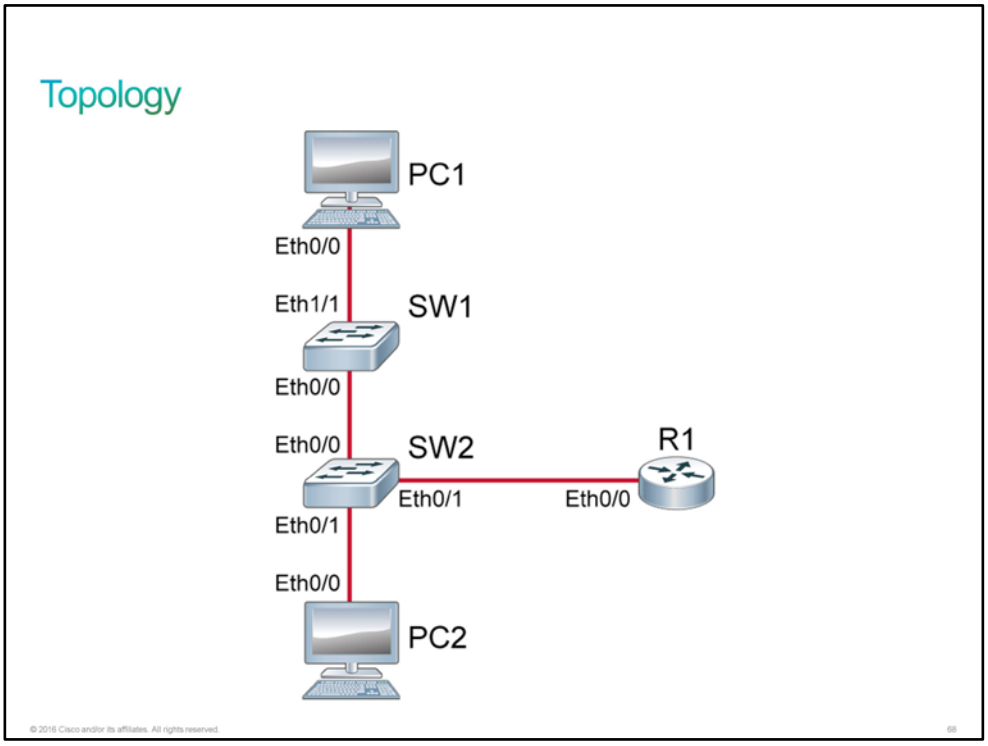

Job Aid

Device Information

Device Information Table

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

Device	Characteristic	Value
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
SW1	Hostname	SW1
SW1	VLAN 1 IP address	10.10.1.2/24
SW1	Default gateway	10.10.1.1
SW1	Ethernet0/0 description	Link to SW2
SW1	Ethernet0/1 description	Link to PC1
SW2	Hostname	SW2
SW2	VLAN 1 IP address	10.10.1.3/24
SW2	Default gateway	10.10.1.1
SW2	Ethernet0/0 description	Link to SW1
SW2	Ethernet0/1 description	Link to R1
SW2	Ethernet0/2 description	Link to PC2
R1	Hostname	R1
R1	Ethernet0/0 description	Link to SW2
R1	Ethernet0/0 IP address	10.10.1.1/24
R1	Loopback 0 IP	10.10.3.1/24

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: Secure Access to Privileged EXEC Mode

Activity

Step 1 On R1, access the privileged EXEC with the **enable** command and the global configuration with the **configure terminal** command.

On R1, enter the following commands:

```
R1> en
R1# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
R1(config)#
```

The most commonly used commands are abbreviated in this guided discovery. For example, **en** for **enable** and **conf t** for **configure terminal**. If there is any confusion, you can attempt tab completion of 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 Set the enable password to "Password123" and leave the configuration mode.

On R1, enter the following commands:

```
R1(config)# enable password Password123
R1(config)# end
R1#
```

Step 3 The enable password will now protect access to the privileged EXEC. Verify this fact by leaving the privileged EXEC with the **disable** command, and then use **enable** again, and authenticate with the password "Password123."

On R1, enter the following commands:

```
R1# disable
R1> enable
Password: Password123
R1#
```

Step 4 View the enable password in the running configuration.

On R1, enter the following command:

```
R1# sh run | inc enable
enable password Password123
```

By default, the enable password is stored in the configuration as clear text.

Step 5 Configure an enable secret, setting it to "Secret123."

On R1, enter the following commands:

```
R1# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
R1(config)# enable secret Secret123
R1(config)# end
R1#
```

Step 6 When both are present, the enable secret takes precedence over the enable password. Verify that this fact is correct.

On R1, enter the following commands:

```
R1# disable
R1> enable
Password: Password123
Password: Secret123
R1#
```

The enable password was not accepted to access privileged EXEC. The enable secret was required.

Step 7 View the enable password and the enable secret in the configuration.

On R1, enter the following command:

```
R1# sh run | inc enable
enable secret 4 9h/bNbZRK8Hm9J2ONmwUdf0KoztPJewuR2NseOBKzM6
enable password Password123
```

The enable secret is always stored in a protected fashion in the configuration file. Cisco IOS on routers supports several encryption types. On production routers, you will most likely find type 8 or type 9. Using type 4 is not recommended due to security risks.

Step 8 Enable the **service password-encryption** in the configuration mode. Then revisit how the enable credentials appear in the running configuration.

On R1, enter the following commands:

```
R1# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
R1(config)# service password-encryption
R1(config)# end
R1# sh run | inc enable
enable secret 4 9h/bNbZRK8Hm9J2ONmwUdf0KoztPJewuR2NseOBKzM6
enable password 7 03345A1815182E5E4A584B56
```

The **enable password** is now protected using the Cisco IOS type 7 encryption algorithm. Understand that type 7 encryption is better than nothing, but it is nowhere near as strong as type 5 MD5. The service password encryption will also protect other cleartext passwords that may appear in the configuration file.

Task 2: Secure Console and Remote Access

Activity

Step 1 Enable a password on the console of R1 (line console 0) by using the **login** command with the **password** command. Set the password to "Console123." Also, set the exec-timeout value to 5 minutes.

On R1, enter the following commands:

```
R1# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
R1(config)# line con 0
R1(config-line)# login
% Login disabled on line 0, until 'password' is set
R1(config-line)# password Console123
R1(config-line)# exec-timeout 5
R1(config-line)# end
R1#
```

The default value for exec-timeout is 10 minutes. Setting the time longer may be a convenience for the administrator. Setting the time shorter improves security by limiting the time that a session stays up if the administrator physically walks away from the terminal.

Step 2 View the configuration that is now in place on "line con 0."

On R1, enter the following command:

```
R1# sh run | section line
line con 0
 exec-timeout 5 0
 password 7 080243401A16091243595F
 logging synchronous
 login
line aux 0
line vty 0 4
 login
 transport input all
```

Service password-encryption continues to encrypt new passwords as they are defined.

Step 3 Verify the console password by logging out completely from the Cisco IOS CLI session on R1 and then logging back in. Continue by using the **enable** command to access the privileged EXEC.

On R1, enter the following commands:

```
R1# logout

R1 con0 is now available

Press RETURN to get started.
<Enter>

User Access Verification

Password: Console123
R1> en
Password: Secret123
R1#
```

Step 4 In a similar fashion, add a password to the five vty lines (line vty 0 4), setting the credential that is required for remote access to the CLI of R1. Also, for demonstration purposes, set the exec-timeout to the very small value of 0 minutes and 30 seconds.

On R1, enter the following commands:


```
R1# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
R1(config)# line vty 0 4
R1(config-line)# login
% Login disabled on line 2, until 'password' is set
% Login disabled on line 3, until 'password' is set
% Login disabled on line 4, until 'password' is set
% Login disabled on line 5, until 'password' is set
% Login disabled on line 6, until 'password' is set
R1(config-line)# password VTYPass
R1(config-line)# exec-timeout 0 30
R1(config-line)# end
R1#
```

Step 5 Verify that you can now access the CLI of R1 via Telnet from other systems. Access the console of PC1 and telnet to 10.10.1.1.

On PC1, enter the following commands:

```
PC1> telnet 10.10.1.1
Trying 10.10.1.1 ... Open

User Access Verification

Password: VTYPass
R1>
```

The prompt changed from PC1 to R1. You are currently accessing the console of PC1, but using PC1 to remotely access the CLI of R1.

Step 6 Verify that you can use this remote connection to access the R1 privileged EXEC with the **enable** command and the "Secret123" enable secret. Then remain idle for 30 seconds and verify the functioning of the exec-timeout.

On PC1, enter the following commands:

```
R1> en
Password: Secret123
R1# <wait 30 seconds for exec-timeout>
[Connection to 10.10.1.1 closed by foreign host]
PC1>
```

After the exec-timeout expired, the telnet session was closed. You have returned to the PC1 CLI.

Step 7 Return to the console of R1.

If it has been longer than the 5 minute exec-timeout set on line con 0, you will have to reauthenticate using "Console123" as the console password and "Secret123" as the enable password.

Step 8 You will now increase the sophistication of the login process. Instead of using simply a password for remote access, you will require a username and a password. The first step is to define a username in the configuration. Enter the configuration mode, and then use the **?** to display the options that are available as you configure a username.

On R1, enter the following commands:

```
R1> en
R1# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
R1(config)# username ?
  WORD  User name
```

The next element of the command line is to specify the username as a freeform WORD.

Step 9 Continue by specifying "admin" as the username followed by the **?** to display the next set of options.

On R1, enter the following command:

```
R1(config)# username admin ?
  aaa                 AAA directive
  access-class        Restrict access by access-class
  autocommand         Automatically issue a command after the user logs in
  callback-dialstring Callback dialstring
  callback-line       Associate a specific line with this callback
  callback-rotary     Associate a rotary group with this callback
  dnis                Do not require password when obtained via DNIS
  nocallback-verify   Do not require authentication after callback
  noescape            Prevent the user from using an escape character
  nohangup            Do not disconnect after an automatic command
  nopassword          No password is required for the user to log in
  one-time            Specify that the username/password is valid for only one
                      time
  password            Specify the password for the user
  privilege           Set user privilege level
  secret              Specify the secret for the user
  user-maxlinks       Limit the user's number of inbound links
  view                Set view name
  <cr>
```

There are several options available, but for this purpose, focus only on "password" and "secret." Understand that the differences regarding the **username** command are the same as they are with the enable password and the enable secret.

Step 10 Continue by specifying secret as the credential storage option and use the **?** to display the next set of options.

On R1, enter the following command:

```
R1(config)# username admin secret ?
  0     Specifies an UNENCRYPTED secret will follow
  4     Specifies a SHA256 ENCRYPTED secret will follow
  5     Specifies a MD5 ENCRYPTED secret will follow
  LINE  The UNENCRYPTED (cleartext) user secret
```

You can specify a 4 followed by an SHA-256-protected secret or a 5 followed by an MD5-protected secret. These options allow you to copy the protected secret from one configuration to another. There is also the option to specify a 0 followed by the clear text secret. Specifying the 0 is optional and generally not used. In this case, you do not have a protected secret to work with. You will simply enter the cleartext secret next.

Step 11 Complete the definition of the username "admin" with the "Cisco123" secret.

On R1, enter the following command:


```
R1(config)# username admin secret Cisco123
```

Step 12 Remain in configuration mode. Use the **do** command to execute the privileged EXEC **show running-config** command from within configuration mode. Send the output through the include filter specifying the "user" string.

On R1, enter the following command:

```
R1(config)# do show run | inc user
username admin secret 4 vwcGVdcUZcRMCyxaH2U9Y/PTujsnQWPSbt.LFG8lhTw
```

The username "admin" is now defined and its password is stored in the configuration as a Cisco IOS type 4, SHA-256-protected secret.

Step 13 Currently, the vty lines have the **login** command set without an argument. In this state, authentication is done using the password that is defined on the line itself. If the **login** command is enhanced with the **local** argument, then authentication will be accomplished using usernames stored in the local running configuration. Enter the vty line configuration mode and configure the **login local** command. Then leave the configuration mode.

On R1, enter the following commands:

```
R1(config)# line vty 0 4
R1(config-line)# login local
R1(config-line)# end
R1(config-line)# no password
R1#
```

Step 14 Remote access authentication should now require a valid username and password. Access the console of PC1 and attempt to telnet to R1 (10.10.1.1) to verify this fact.

On PC1, enter the following commands:

```
PC1> telnet 10.10.1.1
Trying 10.10.1.1 ... Open

User Access Verification

Username: admin
Password: Cisco123
R1>
```

Step 15 Recall that the exec-timeout on the vty lines is set to 0 minutes and 30 seconds. Quickly continue to use the remote connection to modify the configuration. As long as less than 30 seconds pass between keystrokes, the connection should remain open. If you do get logged out, you can return to R1 to enter the configuration changes.

On PC1 console while telneted to R1, enter the following commands:

```
R1> en
Password: Secret123
R1# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
R1(config)# line vty 0 4
R1(config-line)# exec-timeout 10
R1(config-line)# end
R1#
```

You may be accustomed to seeing a "SYS-5-CONFIG_I" syslog message displayed when leaving the configuration mode. That message did not appear here. By default, syslog messages are displayed on the console (line con 0) only. In this case, you have a session to a vty line. You can see the syslog message on the R1 console screen.

Changing the configuration via a remote access session is itself a demonstration of the importance of securing administrative access. Without proper security, network attackers could access your network devices and take control of your network.

Step 16 From this remote connection, review the configuration of the vty lines with a filtered **show running-config** command.

On PC1 console while telneted to R1, enter the following commands:

```
R1# show run | section line
line con 0
 exec-timeout 5 0
 password 7 080243401A16091243595F
 logging synchronous
 login
line aux 0
line vty 0 4
 login local
 transport input all
```

There is no longer an exec-timeout command on line vty 0 4. Setting the timeout value back to its default value of 10 minutes causes the command to be hidden in the running configuration.

Step 17 Close the remote access connection using either the **logout** or **exit** command.

On PC1 console while telneted to R1, enter the following commands:

```
R1# logout
[Connection to 10.10.1.1 closed by foreign host]
PC1>
```

From the user or privileged EXEC, you can use the **logout** and **exit** commands interchangeably to terminate remote access connections.

Task 3: Enable SSH

Activity

Step 1 On R1, enable SSH. The prerequisite of SSH on Cisco IOS vty lines is having an RSA public/private key pair. The prerequisite to defining the key pair is to have a hostname and a

domain name defined. R1 already has a hostname configured. Configure "icnd.lab" as the domain name, then generate a 1024-bit RSA public/private key pair.

On R1, enter the following commands:

```
R1# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
R1(config)# ip domain-name icnd.lab
R1(config)# crypto key generate rsa
The name for the keys will be: R1.icnd.lab
Choose the size of the key modulus in the range of 360 to 4096 for your
  General Purpose Keys. Choosing a key modulus greater than 512 may take
  a few minutes.

How many bits in the modulus [512]: 1024
% Generating 1024 bit RSA keys, keys will be non-exportable...
[OK] (elapsed time was 0 seconds)

R1(config)#
*Nov 12 12:43:54.804: %SSH-5-ENABLED: SSH 1.99 has been enabled
```

Shortly after you generate the RSA keypair, SSH is automatically enabled on the router.

Step 2 There are some security issues that are associated with SSHv1. Limit the options to SSHv2 only, then leave configuration mode.

On R1, enter the following commands:

```
R1(config)# ip ssh version 2
R1(config)# end
R1#
```

Step 3 Both Telnet and SSH should now be options for remote access to R1. Access the console of PC1 to verify this fact. First telnet from PC1 to R1. Disconnect when you are connected.

On PC1, enter the following command:

```
PC1> telnet 10.10.1.1
Trying 10.10.1.1 ... Open

User Access Verification
Username: admin
Password: Cisco123

R1> exit
[Connection to 10.10.1.1 closed by foreign host]
PC1>
```

Step 4 Next, test an SSH connection from PC1 to R1.

On PC1, enter the following command:

```
PC1> ssh -l admin 10.10.1.1
Password: Cisco123
R1>
```

The "–l" is a dash and a lower-case "L" letter, not a dash with the number 1. Think lower-case "L" for "login."

Step 5 From this remote access connection, examine the vty configuration to understand why both Telnet and SSH are allowed. It is because "transport input all" is specified on the "line vty 0 4" configuration.

On R1, enter the following command:

```
R1> en
Password: Secret123
R1# show run | section line
line con 0
 exec-timeout 5 0
 password 7 080243401A16091243595F
 logging synchronous
 login
line aux 0
line vty 0 4
 password 7 0125323D6B0A151C
 login local
 transport input all
```

Step 6 Because SSH is superior to Telnet from a security perspective, change the transport input option from **all** to **ssh** under "line vty 0 4."

On R1, enter the following commands:

```
R1# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
R1(config)# line vty 0 4
R1(config-line)# transport input ssh
R1(config-line)# end
```

Step 7 Terminate this SSH session, to return to the local console of PC1.

On R1, enter the following command:

```
R1# logout
[Connection to 10.10.1.1 closed by foreign host]
PC1>
```

Step 8 Attempt a Telnet session from PC1 to R1. Because the transport input is now set to SSH only, the Telnet attempt should be rejected.

On PC1, enter the following command:

```
PC1> telnet 10.10.1.1
Trying 10.10.1.1 ...
% Connection refused by remote host
```

Step 9 Verify that SSH is still valid for new remote access sessions. Try to connect with SSH one more time from PC1 to R1. Terminate the session after it successfully initiates.

On PC1, enter the following command:

```
PC1> ssh -l admin 10.10.1.1
Password: Cisco123
R1> logout
[Connection to 10.10.1.1 closed by foreign host]
PC1>
```

You have explored many options for securing administrative access on Cisco IOS devices during this discovery lab. You used the enable password and the enable secret to protect the privileged EXEC. You used service password encryption to provide simple protection to cleartext passwords. You implemented a simple password protection on the console and the vty lines. You then went further with the vty lines, requiring a username and password for access, and configuring SSH. Feel free to continue exploring these concepts independently within the lab environment.

This is the end of the discovery lab.

Limiting Remote Access with ACLs

You can limit access to vty lines to specific IP addresses or subnets in order to control remote administration of network devices. Remote administration is commonly run over a Telnet or an SSH connection, where the SSH connection is an encrypted communication channel between the administrator workstation and the device.

Usually, there are two steps that you must complete to limit remote access with ACLs:

1. **Configure an ACL:** The following example shows an ACL being configured with two lines. The first line permits Telnet access from the network addresses in the 10.1.1.0/24 subnet. The second line is not mandatory because there is an implicit deny statement at the end of every ACL. However, creating an explicit deny statement and appending the **log** keyword allows you to monitor attempts of unauthorized sources trying to access the device.

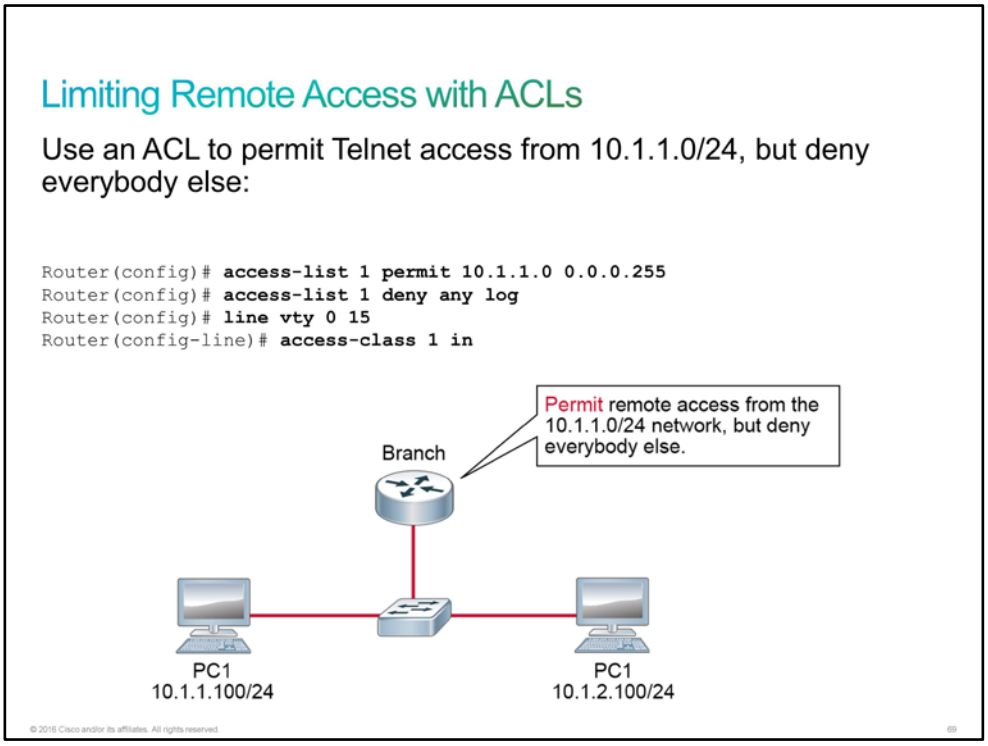

2. **Apply the ACL to the lines:** The **access-class** command applies the ACL on vty lines. Using the **in** keyword after the name of the ACL tells the router to limit vty connections that are coming into the network device.

Apply the ACL on vty lines:

```
Router(config)# line vty 0 15
Router(config-line)# access-class 1 in
```

The example uses 16 vty lines (range 0 to 15). In the configuration output, it will appear in two vty line ranges, first from 0 to 4 and second from 5 to 15. If there is no need for more than 5 vty lines, you may configure only first range.

Configuring the Login Banner

You can define a customized login banner to be displayed before the username and password login prompts.

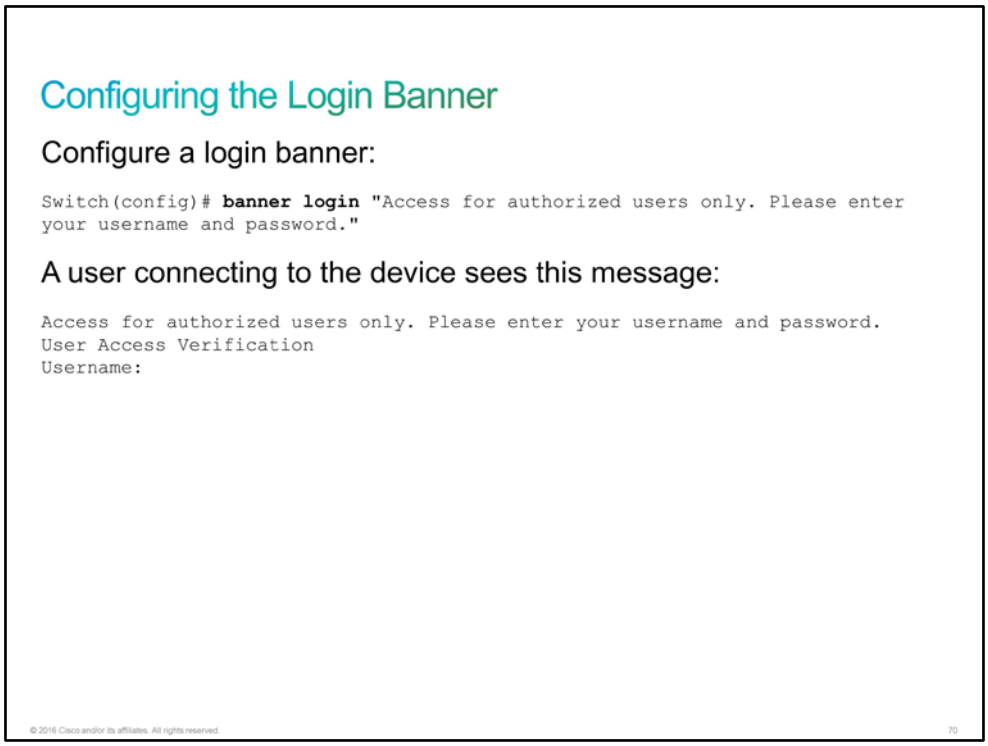

To configure a login banner, use the **banner login** command in global configuration mode. Enclose the banner text in quotation marks or use a delimiter that is different from any character appearing in the banner string.

Note	Use caution when you create the text that is used in the login banner. Words like "welcome" may imply that access is not restricted and may allow hackers some legal defense of their actions.

To define and enable an MOTD banner, use the **banner motd** command in global configuration mode.

This MOTD banner is displayed to all terminals that are connected and is useful for sending messages that affect all users (such as impending system shutdowns).

Discovery 36: Limit Remote Access Connectivity

Introduction

This discovery lab will guide you through the remote access limitation by using an ACL. You will implement login and exec banners.

The devices are configured as represented in the topology diagram, including IP addresses. This discovery lab will focus on R1. Other devices will be used as sources of remote access connections.

Topology

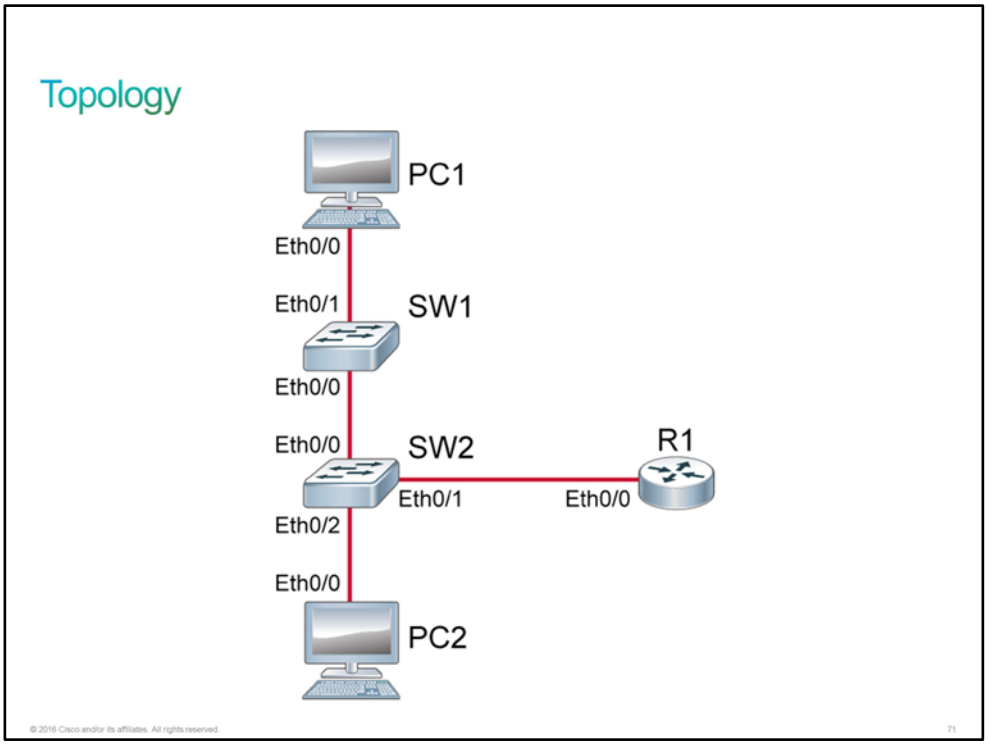

Job Aid

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

Device	Characteristic	Value
PC2	IP address	10.10.1.20/24
PC2	Default gateway	10.10.1.1
SW1	Hostname	SW1
SW1	VLAN 1 IP address	10.10.1.2/24
SW1	Default gateway	10.10.1.1
SW1	Ethernet0/0 description	Link to SW2
SW1	Ethernet0/1 description	Link to PC1
SW2	Hostname	SW2
SW2	VLAN 1 IP address	10.10.1.3/24
SW2	Default gateway	10.10.1.1
SW2	Ethernet0/0 description	Link to SW1
SW2	Ethernet0/1 description	Link to R1
SW2	Ethernet0/2 description	Link to PC2
R1	Hostname	R1
R1	Ethernet0/0 description	Link to SW2
R1	Ethernet0/0 IP address	10.10.1.1/24
R1	Loopback 0 IP	10.10.3.1/24
R1	Enable password	Password123
R1	Enable secret	Secret123
R1	Console password	Console123
R1	Console Login Mode	line
R1	VTY 0 4 Line Password	VTYPass
R1	VTY 0 4 Login Mode	local
R1	Username / Secret	admin / Cisco123

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: Limit Remote Access with ACLs

Activity

Step 1 Explore the use of ACLs to control the source IP addresses that are allowed to establish remote access sessions to a Cisco IOS device. Before defining new access lists, you should know which access lists exist to prevent accidental editing of an access list that is already defined. On R1, view the access lists that are in place. To access R1, use the console password "Console123" and the "Secret123" enable secret.

On R1, enter the following command:

```
R1 con0 is now available

Press RETURN to get started.
<Enter>

User Access Verification

Password: Console123
R1> en
Password: Secret123
R1# show access-list
R1#
```

On R1, no access lists are configured.

Step 2 Enter the global configuration mode and define a new access list number 1 that permits PC1 (10.10.1.10) and PC2 (10.10.1.20) and explicitly denies all other addresses with the log option enabled.

On R1, enter the following commands:

```
R1# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
R1(config)# access-list 1 permit 10.10.1.10
R1(config)# access-list 1 permit 10.10.1.20
R1(config)# access-list 1 deny any log
```

Step 3 Assign the access list 1 to the vty lines in the inbound direction using the **access-class** command, then leave the configuration mode.

On R1, enter the following commands:

```
R1(config)# line vty 0 4
R1(config-line)# access-class 1 in
R1(config-line)# end
R1#
```

Step 4 One at a time, access the consoles of PC1, PC2, and SW1 and attempt SSH sessions to the R1 (10.10.1.1). The sessions should be successful from PC1 and PC2, but not from SW1.

On PC1, enter the following command:

```
PC1> ssh -l admin 10.10.1.1
Password: Cisco123
R1> logout
[Connection to 10.10.1.1 closed by foreign host]
PC1>
```

On PC2, enter the following command:

```
PC2> ssh -l admin 10.10.1.1
Password: Cisco123
R1> logout
[Connection to 10.10.1.1 closed by foreign host]
PC2>
```

On SW1, enter the following command:

```
SW1> ssh -l admin 10.10.1.1
% Connection refused by remote host
SW1>
```

Step 5 Return to the console of R1. You should find a syslog message that is associated with the access attempt from SW1 that the explicit deny statement at the end of the access list denied.

Syslog message displayed on the R1 console:

```
R1#
*Nov 12 13:03:18.445: %SEC-6-IPACCESSLOGNP: list 1 denied 0 10.10.1.2 ->
0.0.0.0, 1 packet
```

Step 6 On the R1, show access list 1 and verify the match counters on each line.

On R1, enter the following command:

```
R1# show access-list 1
Standard IP access list 1
    10 permit 10.10.1.10 (2 matches)
    20 permit 10.10.1.20 (2 matches)
    30 deny   any log (1 match)
```

When an ACL is applied to the vty lines with the **access-class** command, each successful connection will increment the match counter on the associated permit statement by two while each rejected connection will only increment the match counter on the associated deny statement by one.

Task 2: Configure the Login and EXEC Banners

Activity

Step 1 Another access control option that you will explore during this discovery is the use of banner messages. The login banner is displayed before the user logs in, and the EXEC banner is displayed after a successful login. Start by configuring a login banner on R1.

On R1, enter the following commands:

```
R1# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
R1(config)# banner login "
Enter TEXT message.  End with the character '"'.
Access for authorized users only.
Enter your valid credentials for access:
"
```

Step 2 Also, configure an EXEC banner on R1, and then leave the configuration mode.

On R1, enter the following commands:

```
R1(config)# banner exec "
Enter TEXT message.  End with the character '"'.
AUTHORIZED ACCESS ONLY!
If you are not authorized, LOGOUT IMMEDIATELY
"
R1(config)# end
R1#
```

Step 3 Access the console of PC1 and execute an SSH connection to R1. You should see the login banner before entering the password and you should see the EXEC banner after authentication and before the first user EXEC prompt is displayed.

On PC1, enter the following commands:

```
PC1> ssh -l admin 10.10.1.1
Access for authorized users only.
Enter your valid credentials for access:

Password: Cisco123

AUTHORIZED ACCESS ONLY!
If you are not authorized, LOGOUT IMMEDIATELY
R1> logout
[Connection to 10.10.1.1 closed by foreign host]
PC1>
```

You have explored how to limit remote access connections during this discovery. You have limited authorized remote access systems with ACLs. You finished with a simple demonstration of the login and EXEC banners. Feel free to continue exploring these concepts independently within the lab environment.

This is the end of the discovery lab.

Challenge

1. High levels of humidity posing a danger to devices operating properly can be considered to be which of the following types of threats?

   A. Remote Access Threats
   B. Local Access and Physical Threats
   C. Environmental Threats
   D. Electrical Threats
   E. Maintenance Threats

2. Which of the following commands will encrypt plain text passwords on Routers?

   A. **password encryption**
   B. **service password-encryption**
   C. **service encryption**
   D. **enable secret**

3. Which of the following commands enables you to configure the parameters for console access?

   A. **line console 0**
   B. **line console**
   C. **login console 0**
   D. **login console**

4. Which of the following is the correct command to generate RSA keys that the user will use during authentication, when connecting over SSH to a router?

   A. **crypto key generate rsa**
   B. **crypto generate key rsa**
   C. **crypto rsa generate key**
   D. **crypto generate rsa key**

5. Choose the valid configuration to restrict remote users by applying ACL to vty lines.

   A. **router(config)# line vty 0 15**
      router(config-line)# access-group 1 in
   B. **router(config)# line vty 0 15**
      router(config-line)# access-list 1 in
   C. **router(config)# line vty 0 15**
      router(config-line)# access-class 1 in
   D. **router(config)# line vty 0 15**
      router(config-line)# ip access-group 1 in

6. Which of the following banners should be used to show information that should be hidden from unauthorized users?

   A. MOTD
   B. Login
   C. EXEC

7. Making sure that the cable runs are neat, is mitigation for which kind of threat?

   A. Remote Access Threats
   B. Environmental Threats
   C. Electrical Threats
   D. Maintenance Threats

Answer Key

Challenge

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

Lesson 2: Implementing Device Hardening

Introduction

Your boss sends you to your customer to secure unused ports. When discussing how to secure ports, you will introduce the **interface range** command. You will discuss the lack of control over utilized ports, and present port security as a possible solution. You will also explain the need to disable unused services. Your customer wants to implement the correct system time, so you will introduce NTP and demonstrate an NTP configuration example.

Securing Unused Ports

Unused ports on a switch can be a security risk. A hacker can plug a switch into an unused port and become part of the network. Therefore, unsecured ports can create a security hole.

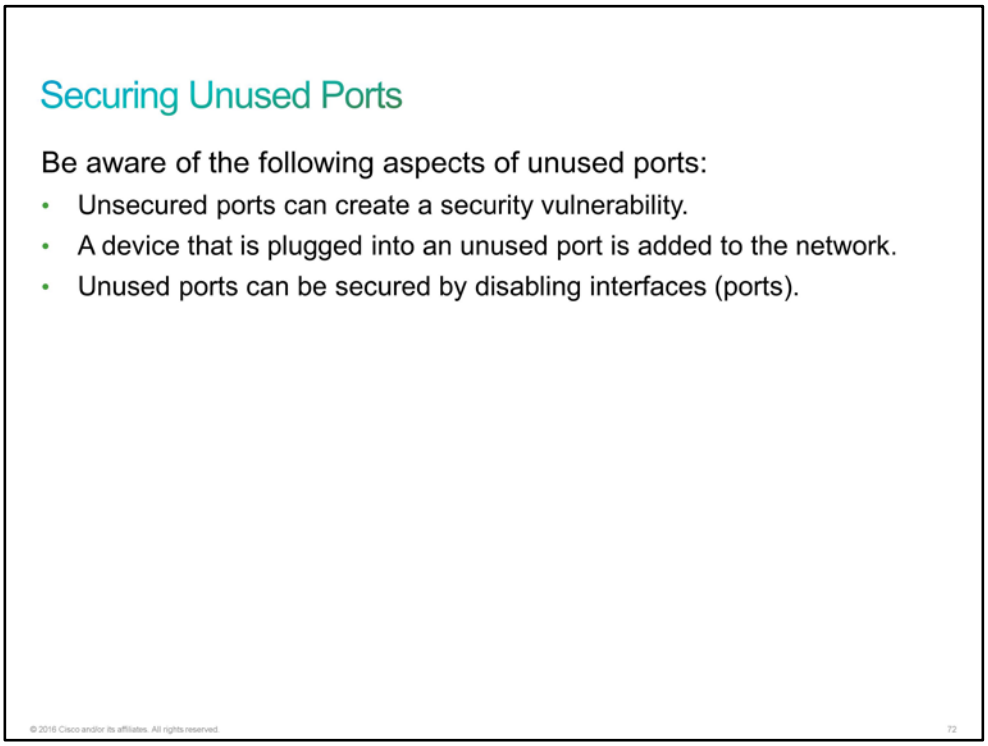

Disabling an Interface (Port)

A simple method that many administrators use to help secure their network from unauthorized access is to disable all unused ports on a network switch.

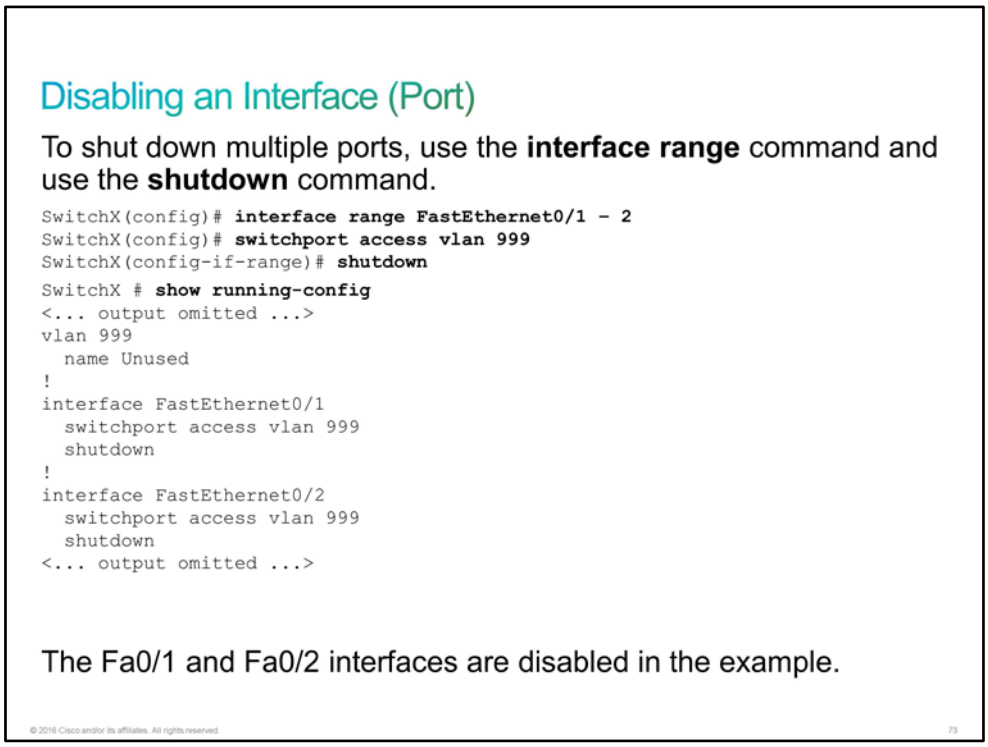

Imagine, for example, that the Cisco switch has 24 ports. If there are 3 Fast Ethernet connections in use, the practicing good security demand is that you disable the 21 unused ports.

It is simple to disable multiple ports on a switch. Navigate to each unused port and issue the Cisco IOS **shutdown** command. An alternate way to shut down multiple ports is to use the **interface range** command. If a port needs to be activated, you can manually enter the **no shutdown** command on this interface.

The process of enabling and disabling ports can become a tedious task, but the enhanced security on your network is well worth the effort.

To make configuration more secure add unused ports into an unused vlan. Use **vlan** command, to configure new VLAN and use **switchport access vlan** interface command, to add port into VLAN.

Port Security

Now that you know about protecting unused ports, you need to learn how to protect the ports that are in use. You can use the port security feature of Cisco IOS Software to restrict access to a switch port based on MAC addresses. A port that is configured with port security accepts frames only from secure MAC addresses. You can configure a device to learn these addresses dynamically, or you can configure them statically.

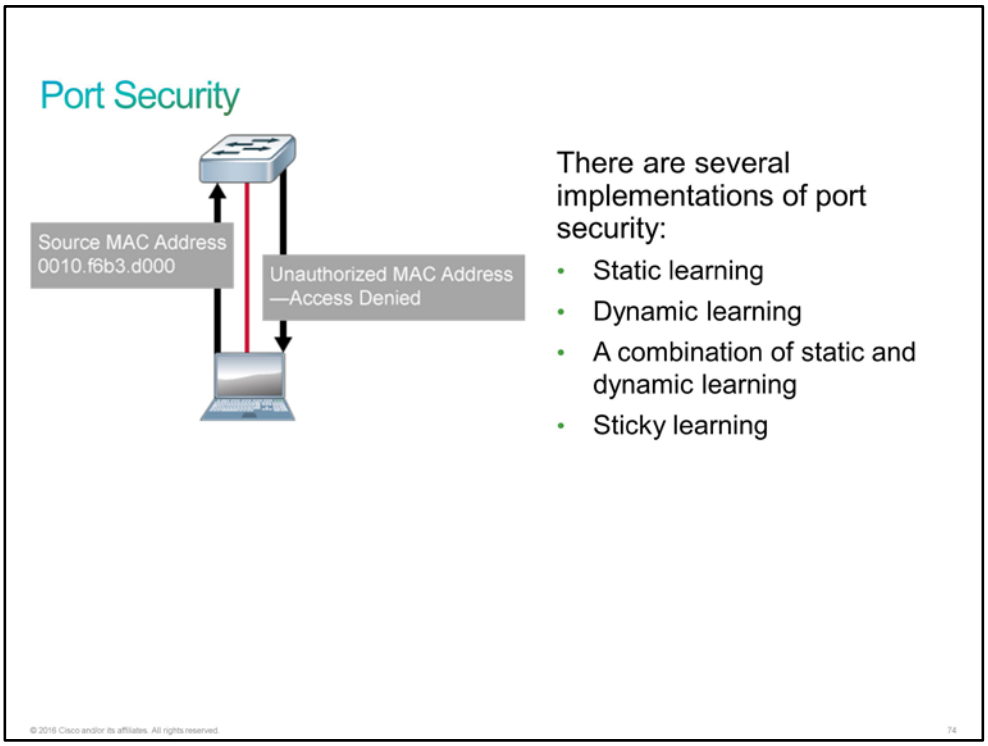

There are several implementations of port security:

- **Static learning:** You can statically configure specific MAC addresses that are permitted to use a port. The source MAC addresses that you do not specifically permit are not allowed to source frames to the port.
- **Dynamic learning:** You can specify how many MAC addresses are permitted to use a port at one time. Use the dynamic approach when you care only about how many MAC addresses are permitted to use the port, rather than which MAC addresses are permitted. If a port on which dynamic learning is configured has a link-down condition, all dynamically learned addresses are removed. Following bootup, a reload, or a link-down condition, port security does not populate the address table with dynamically learned MAC addresses until the port receives ingress traffic. Depending on how you configure the device, dynamically learned addresses age out after a certain period and new addresses are learned, up to the maximum that you have defined.

- **A combination of static and dynamic learning:** You can specify some of the permitted MAC addresses and let the switch learn the rest of the permitted MAC addresses. For example, you could limit the number of MAC addresses to four and statically configure two of the MAC addresses. The switch would then dynamically learn the next two MAC addresses that it received on that port. The two statically configured addresses would not age out, but the two dynamically learned addresses could, depending on your configuration.
- **Sticky learning:** When you configure sticky learning on an interface, the interface converts dynamically learned addresses to "sticky secure" addresses. This feature adds the dynamically learned addresses to the running configuration as if they were statically configured. If you save the running configuration to NVRAM, port security with sticky MAC addresses saves dynamically learned MAC addresses in the startup configuration file, and the port does not have to learn addresses from ingress traffic after a bootup or a restart. Sticky secure addresses do not age out.

When a frame arrives on a port for which port security is configured, its source MAC address is checked against the secure MAC address table. If the source MAC address matches an entry in the table for that port, the device forwards the frame to be processed. Otherwise, the device does not forward the frame.

When port security is configured on a port, the following situations are considered security violations:

- The maximum number of secure MAC addresses has been added to the address table, and a host whose MAC address is not in the address table attempts to access the interface.
- A host with a secure MAC address that is configured or learned on one secure port attempts to access another secure port in the same VLAN.

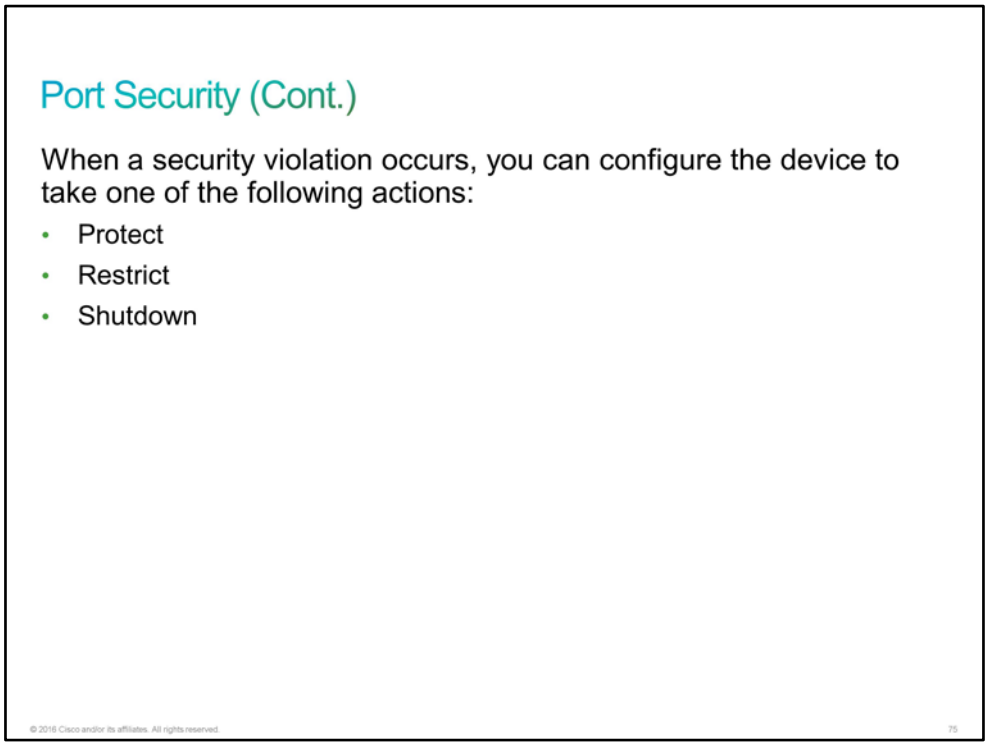

You can configure the device to take one of the following actions when a security violation occurs:

- **Protect:** The protect violation mode drops packets with unknown source addresses until you remove enough secure MAC addresses to drop below the maximum value.
- **Restrict:** The restrict violation mode also drops packets with unknown source addresses until you remove enough secure MAC addresses to drop below the maximum value. However, it also generates a log message and causes the security violation counter to increment.
- **Shutdown:** The shutdown violation mode puts the interface into an error-disabled state immediately. The entire port is shut down. Also, in this mode, the system generates a log message, sends an SNMP trap, and increments the violation counter. To make the interface usable, you must use a manual intervention or the error-disabled recovery. Shutdown is the default violation mode.

  When the port security violation mode is set to shutdown, the port with the security violation goes to the error-disabled state. You receive this notification on the device:

```
Sep 20 12:44:54.966: %PM-4-ERR_DISABLE: psecure-violation error detected on Fa0/5,
putting Fa0/5 in err-disable state
Sep 20 12:44:54.966: %PORT_SECURITY-2-PSECURE_VIOLATION: Security violation occurred,
caused by MAC
address 000c.292b.4c75 on port FastEthernet0/5.
Sep 20 12:44:55.973: %LINEPROTO-5-PPDOWN: Line protocol on Interface FastEthernet0/5,
changed state
 to down
Sep 20 12:44:56.971: %LINK-3-UPDOWN: Interface FastEthernet0/5, changed state to down
```

  To make the interface operational again, you need to disable the interface administratively and then enable it again, as shown here:

```
SwitchX(config)# interface FastEthernet 0/5
SwitchX(config-if)# shutdown
Sep 20 12:57:28.532: %LINK-5-CHANGED: Interface FastEthernet0/5, changed state to
administratively down
SwitchX(config-if)# no shutdown
Sep 20 12:57:48.186: %LINK-3-UPDOWN: Interface FastEthernet0/5, changed state to up
Sep 20 12:57:49.193: %LINEPROTO-5-UPDOWN: Line protocol on Interface FastEthernet0/5,
changed state
 to up
```

You can specify how secure MAC address aging occurs on a port by configuring either absolute or inactivity aging. When you configure absolute aging, all the dynamically learned secure addresses age out when the aging time expires. When you configure inactivity aging, the aging time defines the period of inactivity after which all the dynamically learned secure addresses age out. You can also specify the aging time.

You can configure port security only on static access ports or trunk ports. You cannot configure port security on an interface in the default mode (dynamic auto).

Configuring Port Security

To configure port security to limit and identify the MAC addresses of stations that are allowed to access the port, do as follows:

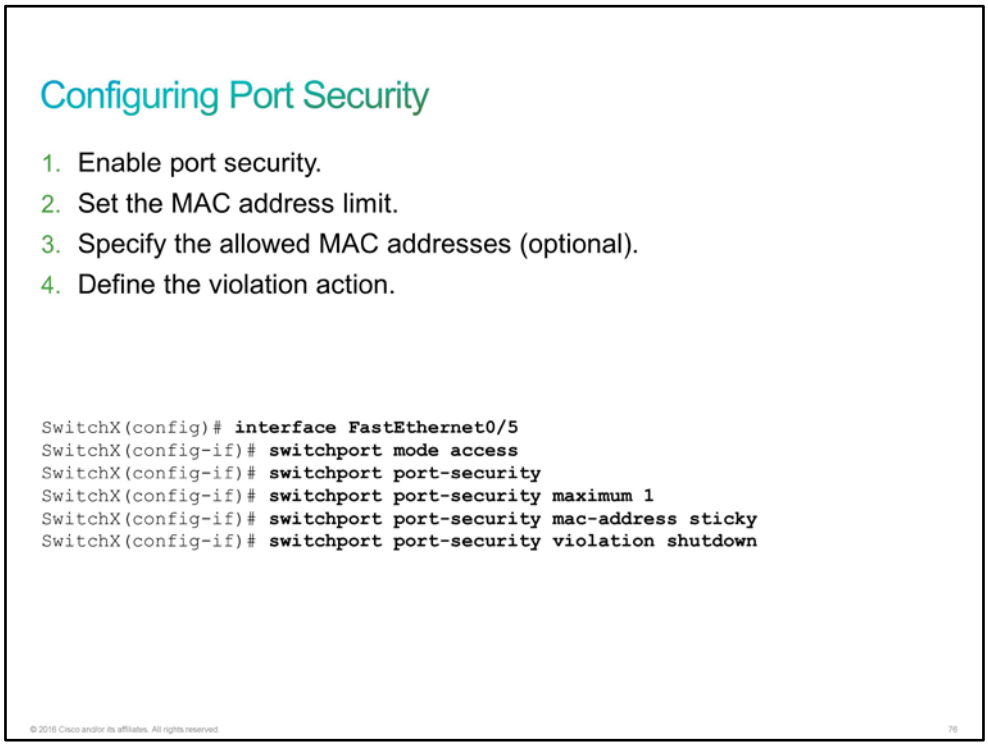

The figure shows how to enable sticky port security on the FastEthernet0/5 port of SwitchX.

Port security limits the number of valid MAC addresses that are allowed on a port. When you assign secure MAC addresses to a secure port, the port does not forward packets with source addresses outside the group of defined addresses.

If you limit the number of secure MAC addresses to one and assign a single MAC address to this port, only the workstation with this particular secure MAC address can successfully connect to this switch port.

If you configure a port as a secure port and the maximum number of secure MAC addresses is reached, a security violation occurs when the MAC address of a workstation that is attempting to access the port is different from any of the identified secure MAC addresses.

Note	Before port security can be activated, you must set the port mode to "access" or "trunk" using the **switchport mode access** | **trunk** command.

Use the **switchport port-security** interface command *without* keywords to enable port security on an interface. Use the **switchport port-security** interface command *with* keywords to configure a secure MAC address, a maximum number of secure MAC addresses, or the violation mode. Use the **no** form of this command to disable port security or to set the parameters to their default states.

You can also configure the maximum number of secure MAC addresses. In this figure, you can see the Cisco IOS command syntax that you use to set the maximum number of MAC addresses to one (**switchport port-security maximum 1**).

You can add secure addresses to the address table after setting the maximum number of secure MAC addresses that are allowed on a port in these ways:

- Manually configure all the addresses (**switchport port-security mac-address 0008.eeee.eeee**).
- Allow the port to dynamically configure all the addresses (**switchport port-security mac-address sticky**).
- Configure several MAC addresses and allow the rest of the addresses to be dynamically configured.

You can configure an interface to convert the dynamic MAC addresses to sticky secure MAC addresses and add them to the running configuration by enabling sticky learning. To enable sticky learning, enter the **switchport port-security mac-address sticky** interface configuration command. When you enter this command, the interface converts all the dynamic secure MAC addresses, including the MAC addresses that the device dynamically learned before you enabled sticky learning, to sticky secure MAC addresses.

The violation mode is set to shutdown (**switchport port-security violation shutdown**). This violation mode, which is the default mode, puts the interface into the error-disabled state immediately and shuts down the entire port.

Note	As mentioned, the other two violation modes are protect and restrict. These modes drop frames from the address that is not allowed, but unlike the shutdown mode, they do not put the interface into the error-disable state.
Note	Port security is disabled by default.

Verifying Port Security

After you have configured port security for your switch, verify that it has been configured correctly.

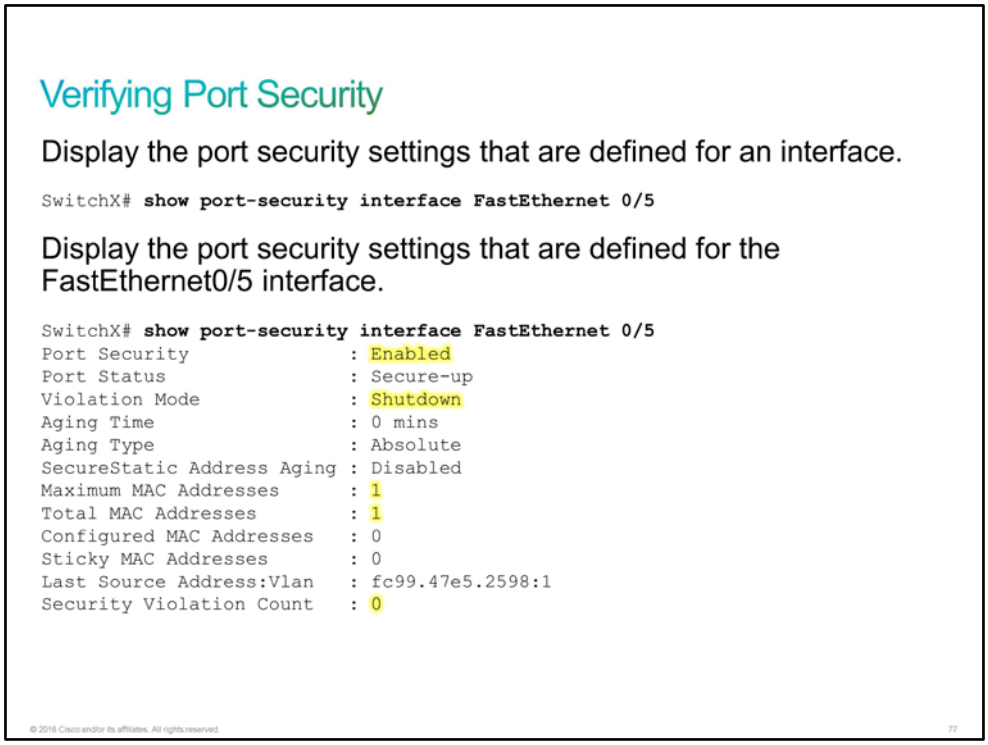

You must check each interface to verify that you have set the port security correctly. You must also verify that you have configured static MAC addresses correctly. Use the **show port-security interface** privileged EXEC command to display the port security settings that are defined for an interface.

The output displays this information (from the top down):

- Whether the port security feature is enabled
- The violation mode
- The maximum allowed number of secure MAC addresses for each interface
- The number of secure MAC addresses on the interface
- The number of security violations that have occurred

Port Security Verification (Cont.)

Display the port security violation for the FastEthernet0/5 interface.

```
SwitchX#show port-security interface FastEthernet 0/5
Port Security                       : Enabled
Port Status                         : Secure-shutdown
Violation Mode                      : Shutdown
Aging Time                          : 0 mins
Aging Type                          : Absolute
SecureStatic Address Aging          : Disabled
Maximum MAC Addresses               : 1
Total MAC Addresses                 : 1
Configured MAC Addresses            : 0
Sticky MAC Addresses                : 1
Last Source Address:Vlan            : 001a.2fe7.3089:1
Security Violation Count            : 1
```


Port Security Verification (Cont.)

Verify the status of the interface.

```
SwitchX# show interface status
Port      Name               Status       Vlan     Duplex  Speed Type
Fa0/1                        connected    1        a-full  a-100 10/100BaseTX
Fa0/2                        notconnect   1          auto   auto 10/100BaseTX
Fa0/3                        notconnect   1          auto   auto 10/100BaseTX
Fa0/4                        notconnect   1          auto   auto 10/100BaseTX
Fa0/5                        err-disabled 1           auto   auto 10/100BaseTX
<output omitted>
```


When MAC addresses are assigned to a secure port, the port does not forward frames with source MAC addresses outside the group of defined addresses. When a port that is configured with port security receives a frame, the source MAC address of the frame is compared to the list of secure source addresses that were manually configured or autoconfigured (learned) on the port. If a MAC address of a device that is attached to the port differs from the list of secure addresses, the port either shuts down until it is administratively re-enabled (default mode) or drops incoming frames from the insecure host (the **restrict** option). The behavior of the port depends on how it is configured to respond to a security violation.

The output in the figure shows that a security violation has occurred, and the port is in the secure-shutdown state.

Because the port security violation mode is set to **shutdown**, the port with the security violation (source MAC addresses outside the group of defined addresses) goes to the error-disabled state. You receive this notification on the switch:

```
Sep 20 06:44:54.966: %PM-4-ERR_DISABLE: psecure-violation error detected on Fa0/5,
putting Fa0/5 in err-disable state
Sep 20 06:44:54.966: %PORT_SECURITY-2-PSECURE_VIOLATION: Security violation occurred,
caused by MAC
address 000c.292b.4c75 on port FastEthernet0/5.
Sep 20 06:44:55.973: %LINEPROTO-5-PPDOWN: Line protocol on Interface FastEthernet0/5,
changed state
 to down
Sep 20 06:44:56.971: %LINK-3-UPDOWN: Interface FastEthernet0/5, changed state to down
```

To verify the status of the interface, use the **show interface status** command.

To make the interface operational again, you need to disable the interface administratively and then enable it again:

```
SwitchX(config)# interface FastEthernet 0/5
SwitchX(config-if)# shutdown
Sep 20 06:57:28.532: %LINK-5-CHANGED: Interface FastEthernet0/5, changed state to
administratively down
SwitchX(config-if)# no shutdown
Sep 20 06:57:48.186: %LINK-3-UPDOWN: Interface FastEthernet0/5, changed state to up
Sep 20 06:57:49.193: %LINEPROTO-5-UPDOWN: Line protocol on Interface FastEthernet0/5,
changed state
 to up
```

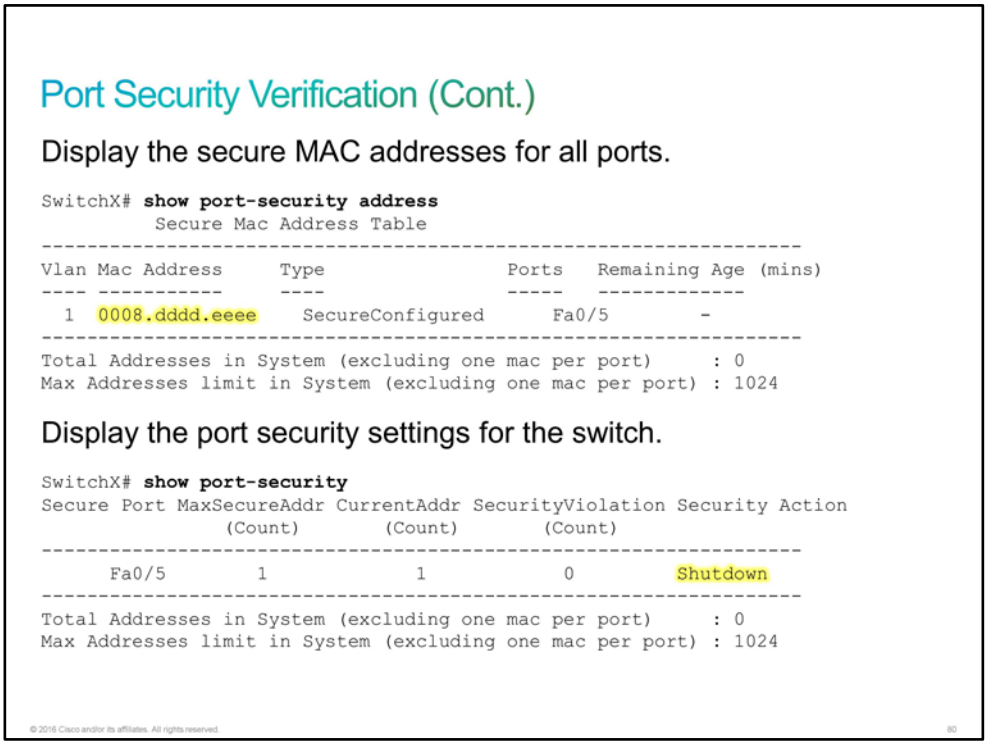

Use the **show port-security address** command to display the secure MAC addresses for all ports. Use the **show port-security** command *without* keywords to display the port security settings for the switch.

Discovery 37: Configure and Verify Port Security

Introduction

Port security restricts a switch port to a specific set of MAC addresses. You should configure it on all ports that connect to end devices.

In this discovery lab, you will configure and verify port security. You will also set error-disabled port automatic recovery.

Topology

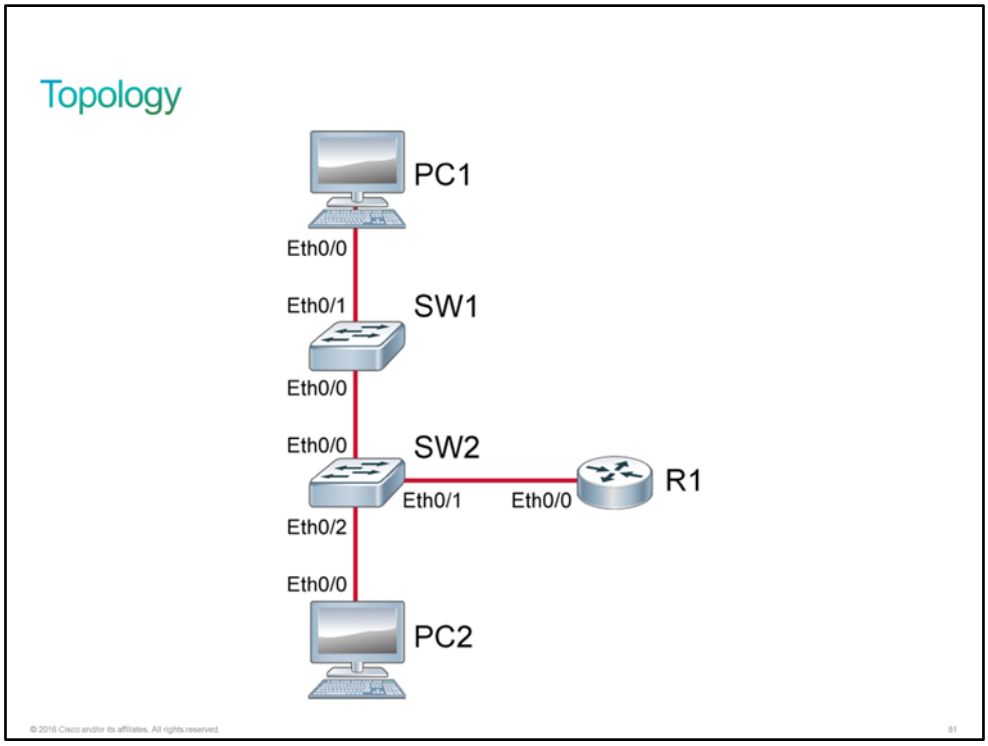

Job Aid

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

Device	Characteristic	Value
PC2	IP address	10.10.1.20/24
PC2	Default gateway	10.10.1.1
SW1	Hostname	SW1
SW1	VLAN 1 IP address	10.10.1.2/24
SW1	Default gateway	10.10.1.1
SW1	Ethernet0/0 description	Link to SW2
SW1	Ethernet0/1 description	Link to PC1
SW2	Hostname	SW2
SW2	VLAN 1 IP address	10.10.1.3/24
SW2	Default gateway	10.10.1.1
SW2	Ethernet0/0 description	Link to SW1
SW2	Ethernet0/1 description	Link to R1
SW2	Ethernet0/2 description	Link to PC2
R1	Hostname	R1
R1	Ethernet0/0 description	Link to SW2
R1	Ethernet0/0 IP address	10.10.1.1/24
R1	Loopback 0 IP	10.10.3.1/24

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 Port Security

Activity

Step 1 On SW1, configure port security with sticky learning on the Ethernet0/1 interface.

On SW1, enter the following commands:


```
SW1# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
SW1(config)# interface Ethernet 0/1
SW1(config-if)# switchport port-security mac-address sticky
SW1(config-if)# switchport mode access
SW1(config-if)# switchport port-security
SW1(config-if)# end
SW1# copy running-config startup-config
Destination filename [startup-config]? <Enter>
Building configuration...
Compressed configuration from 1075 bytes to 720 bytes[OK]
```

When SW1 learns the MAC address of the PC1, you have to save the running configuration so that the learned MAC address stays in the configuration even if the switch reboots.

Step 2 On SW1, verify the port security status.

On SW1, enter the following command:

```
SW1# show port-security
Secure Port  MaxSecureAddr  CurrentAddr  SecurityViolation  Security Action
                (Count)       (Count)          (Count)
-----------------------------------------------------------------------------
      Et0/1              1            1                  0         Shutdown
-----------------------------------------------------------------------------
Total Addresses in System (excluding one mac per port)     : 0
Max Addresses limit in System (excluding one mac per port) : 4096
```

Step 3 On SW1, verify which MAC address is learned on the Ethernet0/1.

On SW1, enter the following command:

```
SW1# show port-security interface Ethernet 0/1
Port Security              : Enabled
Port Status                : Secure-up
Violation Mode             : Shutdown
Aging Time                 : 0 mins
Aging Type                 : Absolute
SecureStatic Address Aging : Disabled
Maximum MAC Addresses      : 1
Total MAC Addresses        : 1
Configured MAC Addresses   : 0
Sticky MAC Addresses       : 1
Last Source Address:Vlan   : aabb.cc00.0200:1
Security Violation Count   : 0
```

MAC address in your output may be different.

Also verify that the same MAC address is listed in the configuration of the SW1.

```
SW1# sh run int e0/1
Building configuration...

Current configuration : 222 bytes
!
interface Ethernet0/1
 description Link to PC1
 switchport mode access
 switchport port-security
 switchport port-security mac-address sticky
 switchport port-security mac-address sticky aabb.cc00.0200
 duplex auto
end
```

Step 4 On PC1, change the MAC address on Ethernet0/0 to eeee.eeee.eeee.

On PC1, enter the following commands:

```
PC1# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
PC1(config)# int e 0/0
PC1(config-if)# mac-address eeee.eeee.eeee
PC1(config-if)# end
PC1#
```

When the MAC address on PC1 changes, SW1 will shut down the port toward PC1. You should see the following log messages on the SW1 console:

```
*Jan 29 12:38:18.353: %PM-4-ERR_DISABLE: psecure-violation error detected on
Et0/1, putting Et0/1 in err-disable state
*Jan 29 12:38:18.354: %PORT_SECURITY-2-PSECURE_VIOLATION: Security violation
occurred, caused by MAC address eeee.eeee.eeee on port Ethernet0/1.
*Jan 29 12:38:19.356: %LINEPROTO-5-UPDOWN: Line protocol on Interface
Ethernet0/1, changed state to down
*Jan 29 12:38:20.356: %LINK-3-UPDOWN: Interface Ethernet0/1, changed state to
down
```

Step 5 On SW1, verify the port security status of Ethernet0/1.

On SW1, enter the following command:

```
SW1# sh port-security int e 0/1
Port Security              : Enabled
Port Status                : Secure-shutdown
Violation Mode             : Shutdown
Aging Time                 : 0 mins
Aging Type                 : Absolute
SecureStatic Address Aging : Disabled
Maximum MAC Addresses      : 1
Total MAC Addresses        : 1
Configured MAC Addresses   : 0
Sticky MAC Addresses       : 1
Last Source Address:Vlan   : eeee.eeee.eeee:1
Security Violation Count   : 1
```

Port Ethernet0/1 is disabled by port security.

Step 6 On PC1, delete the MAC address from Ethernet0/0.

On PC1, enter the following commands:

```
PC1# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
PC1(config)# int e 0/0
PC1(config-if)# no mac-address
PC1(config-if)# end
```

PC1 should use the default MAC address as learned by SW1.

```
PC1# sh int e 0/0 | in bia
  Hardware is AmdP2, address is aabb.cc00.0200 (bia aabb.cc00.0200)
```

MAC address in your output may be different.

Step 7 On SW1, verify the port security status of Ethernet0/1.

On SW1, enter the following command:

```
SW1# show port-security int e0/1
Port Security              : Enabled
Port Status                : Secure-shutdown
Violation Mode             : Shutdown
Aging Time                 : 0 mins
Aging Type                 : Absolute
SecureStatic Address Aging : Disabled
Maximum MAC Addresses      : 1
Total MAC Addresses        : 1
Configured MAC Addresses   : 0
Sticky MAC Addresses       : 1
Last Source Address:Vlan   : eeee.eeee.eeee:1
Security Violation Count   : 1
```

Port Ethernet0/1 is still disabled by port security.

Step 8 On SW1, enable Ethernet0/1 by shut it down and then bring it back up

On SW1, enter the following command:

```
SW1# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
SW1(config)# interface e0/1
SW1(config-if)# shutdown
*Jan 29 12:48:33.655: %LINK-5-CHANGED: Interface Ethernet0/1, changed state to
administratively down
SW1(config-if)# no shutdown
*Jan 29 12:48:39.281: %LINK-3-UPDOWN: Interface Ethernet0/1, changed state to
up
*Jan 29 12:48:40.289: %LINEPROTO-5-UPDOWN: Line protocol on Interface
Ethernet0/1, changed state to up
SW1(config-if)# end
```

Step 9 On SW1, verify the port security status of Ethernet0/1 again.

On SW1, enter the following command:

```
SW1# show port-security int e0/1
Port Security              : Enabled
Port Status                : Secure-up
Violation Mode             : Shutdown
Aging Time                 : 0 mins
Aging Type                 : Absolute
SecureStatic Address Aging : Disabled
Maximum MAC Addresses      : 1
Total MAC Addresses        : 1
Configured MAC Addresses   : 0
Sticky MAC Addresses       : 1
Last Source Address:Vlan   : aabb.cc00.0200:1
Security Violation Count   : 0
```

Port Ethernet0/1 is operational now.

Error-Disabled Port Automatic Recovery

An error-disabled port will become operational after you shut it down and then bring it back up. To reduce the administrative overhead, an error-disabled port can be automatically re-enabled, after the problem that is causing the error-disabled state is fixed.

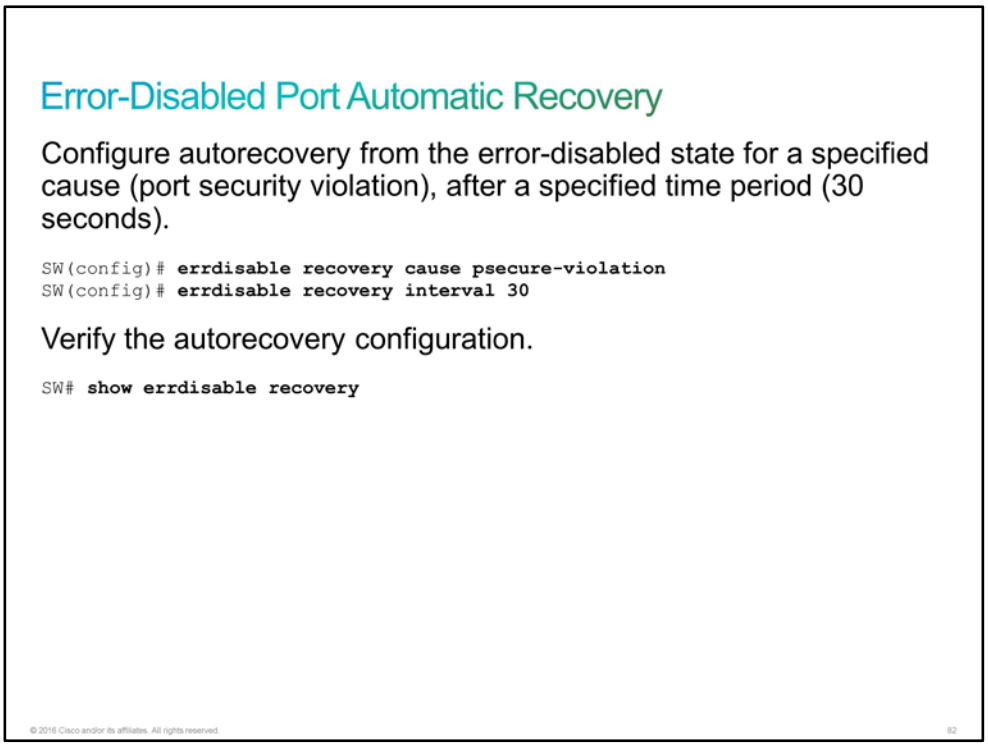

Use the **errdisable recovery** command to automatically re-enable the port after a specified time. If the problem that caused the port to change into the error-disabled state is not resolved, the port will stay in the error-disabled state.

```
SW(config)# errdisable recovery cause cause
SW(config)# errdisable recovery interval seconds
```

The default time interval is 300 seconds, and the minimum is 30 seconds.

You can verify where autorecovery is enabled by using the command **show errdisable recovery**. By default, the autorecovery feature is disabled.

Step 10 On SW1, configure the error-disabled recovery cause to **psecure-violation** and set the interval timer to 30 seconds.

On SW1, enter the following commands:

```
SW1# conf t
SW1(config)# errdisable recovery cause psecure-violation
SW1(config)# errdisable recovery interval 30
SW1(config)# end
```

Step 11 On PC1, change the MAC address on Ethernet0/0 to eeee.eeee.eeee again

On PC1, enter the following commands:

```
PC1# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
PC1(config)# int e 0/0
PC1(config-if)# mac-address eeee.eeee.eeee
PC1(config-if)# end
PC1#
```

When the MAC address on PC1 changes, SW1 will shut down the port toward PC1.

Step 12 On PC1, delete the MAC address from Ethernet0/0.

On PC1, enter the following commands:

```
PC1# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
PC1(config)# int e 0/0
PC1(config-if)# no mac-address
PC1(config-if)# end
```

After 30 seconds, Ethernet0/1 will be recovered. You will see following log message on the SW1 console:

```
*Jan 29 13:02:23.001: %PM-4-ERR_RECOVER: Attempting to recover from psecure-
violation err-disable state on Et0/1
*Jan 29 13:02:25.001: %LINK-3-UPDOWN: Interface Ethernet0/1, changed state to
up
*Jan 29 13:02:26.002: %LINEPROTO-5-UPDOWN: Line protocol on Interface
Ethernet0/1, changed state to up
```

Step 13 From PC1, ping R1 (10.10.1.1) to verify whether Ethernet0/1 on SW1 is operational.

On PC1, enter the following command:

```
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 100 percent (5/5), round-trip min/avg/max = 1/201/1004 ms
```

This is the end of the discovery lab.

Disabling Unused Services

To facilitate deployment, Cisco routers and switches start with a list of services that are turned on and considered to be appropriate for most network environments. However, because not all networks have the same requirements, some of these services may not be needed. Disabling these unnecessary services has two benefits: it helps preserve system resources, and it eliminates the potential for security exploits on the unneeded services.

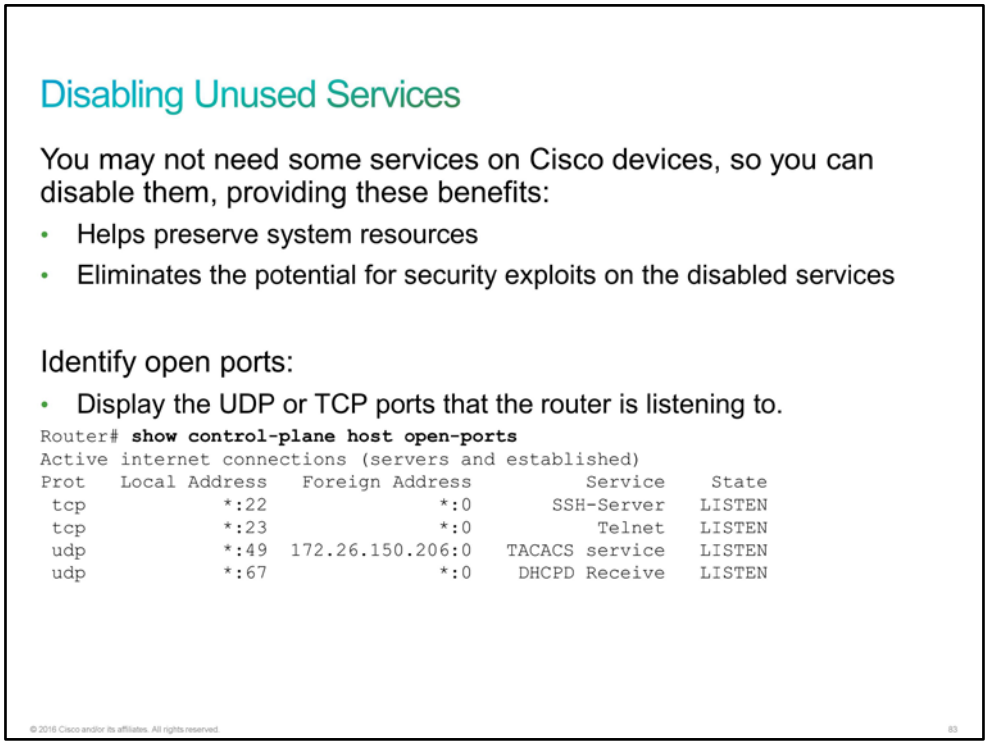

The general best practice is to identify open ports. Use the **show control-plane host open-ports** command to see which UDP or TCP ports the router is listening to and to determine which services need to be disabled.

In the example, services that are enabled on the router are SSH, Telnet, TACACS, and DHCP.

Note	As an alternative, Cisco IOS Software provides the AutoSecure function that helps disable these unnecessary services while enabling other security services.

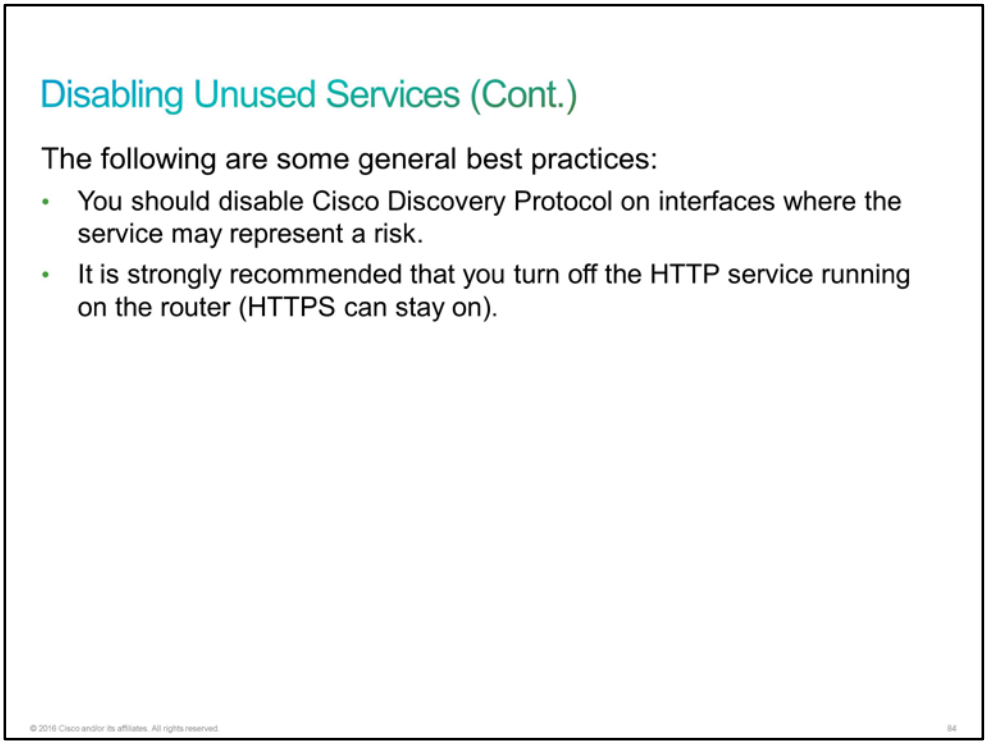

Disable Cisco Discovery Protocol on the interfaces where the service may represent a risk. Examples are external interfaces, such as those at the Internet edge, and data-only ports at the campus and branch access. Cisco Discovery Protocol is enabled by default in Cisco IOS Software Release 15.0 and later.

You can access Cisco routers via a web page, but it is strongly recommended that you turn off the HTTP service that is running on the router.

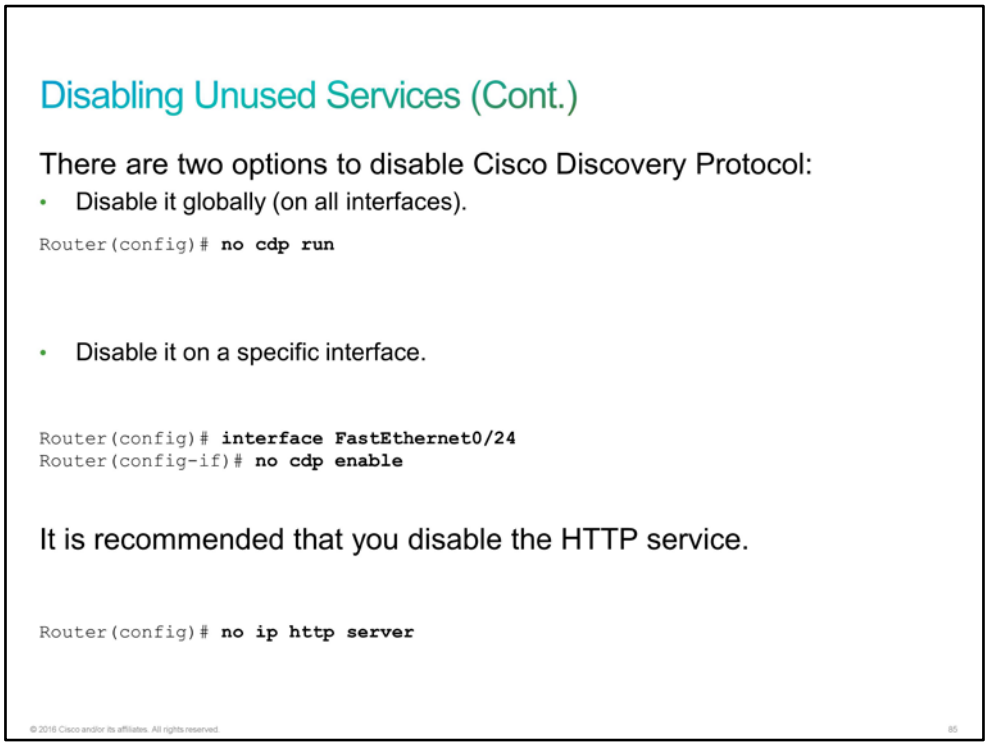

If you prefer not to use the Cisco Discovery Protocol device discovery capability, you can disable it with the **no cdp run** global configuration command. To re-enable Cisco Discovery Protocol after disabling it, use the **cdp run** command in the global configuration mode.

Cisco Discovery Protocol is enabled by default on all supported interfaces to send and receive Cisco Discovery Protocol information. Cisco Discovery Protocol is not on by default on Frame Relay interfaces. You can disable Cisco Discovery Protocol on an interface that supports it with the **no cdp enable** interface configuration command. To re-enable Cisco Discovery Protocol on an interface after disabling it, use the **cdp enable** command in the interface configuration mode.

It is strongly recommended that you turn off the HTTP service that is running on the router. You can use the **no ip http server** global configuration command to disable it. To re-enable the HTTP service after disabling it, use the **ip http server** command in global configuration mode.

Network Time Protocol

Networks use NTP to synchronize the clocks of various devices across a network. Clock synchronization within a network is critical for digital certificates and for the correct interpretation of events within syslog data. A secure method of providing clocking for the network is for network administrators to implement their own private network master clocks that are synchronized to UTC-using satellite or radio. However, if network administrators do not wish to implement their own master clocks because of cost or other reasons, other clock sources are available on the Internet.

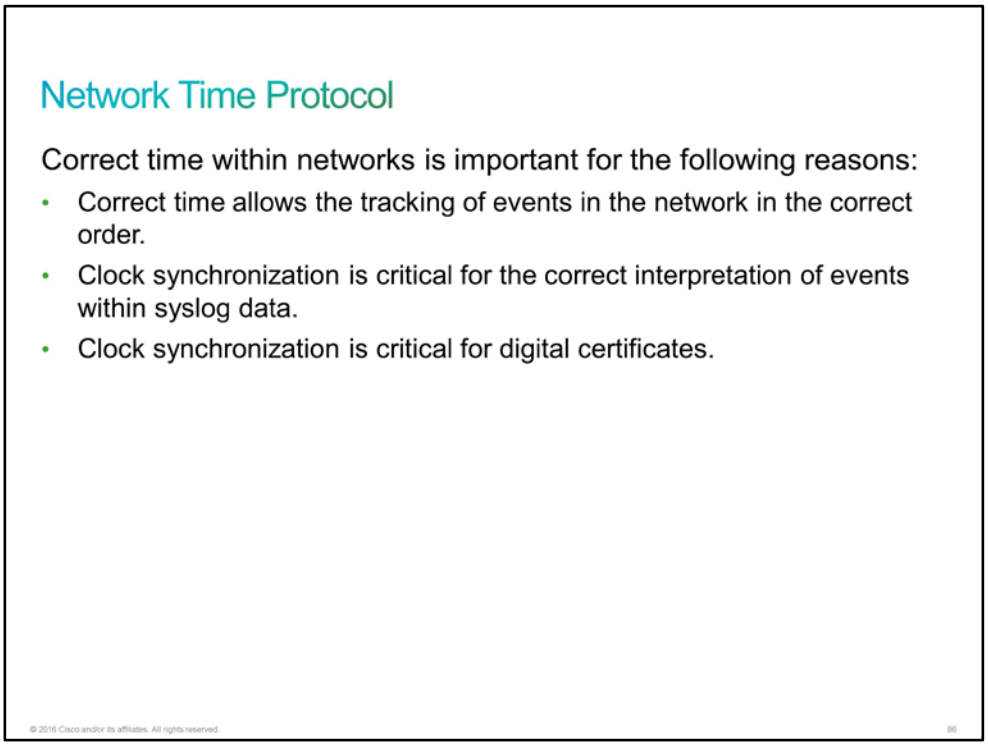

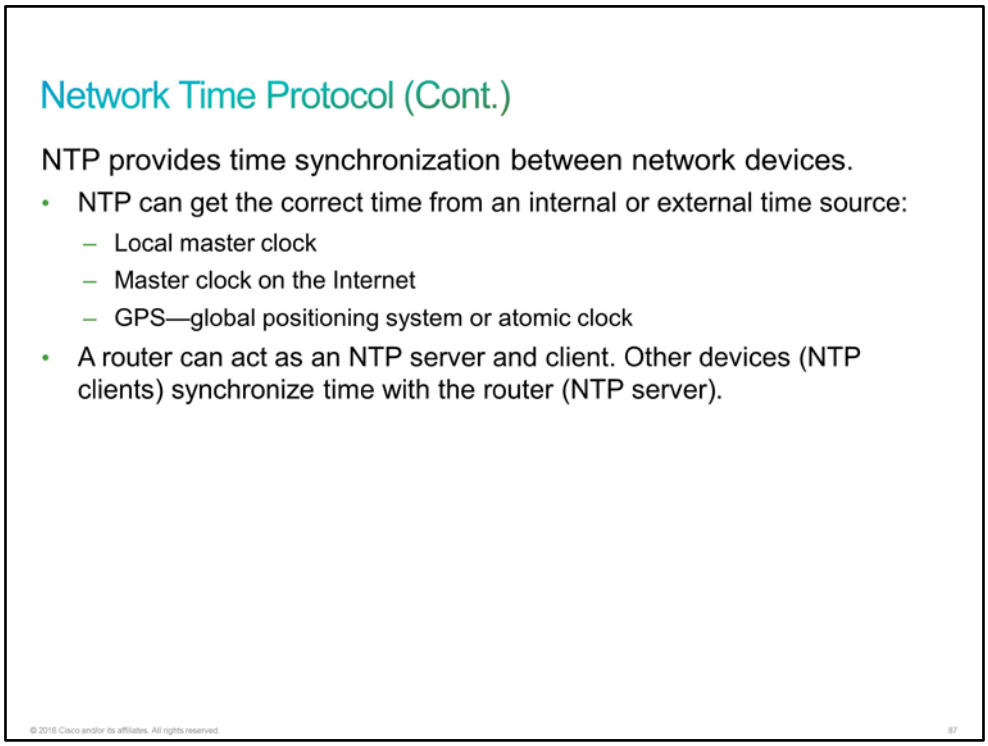

NTP is a protocol for synchronizing the clocks of computer systems over packet-switched, variable-latency data networks. NTP allows routers on your network to synchronize their time settings with an NTP server. A group of NTP clients that obtain time and date information from a single source will have more consistent time settings.

You can configure a router as an NTP server, to which other devices (NTP clients) synchronize their time settings.

Configuring NTP

To configure NTP on Cisco devices, use following commands:

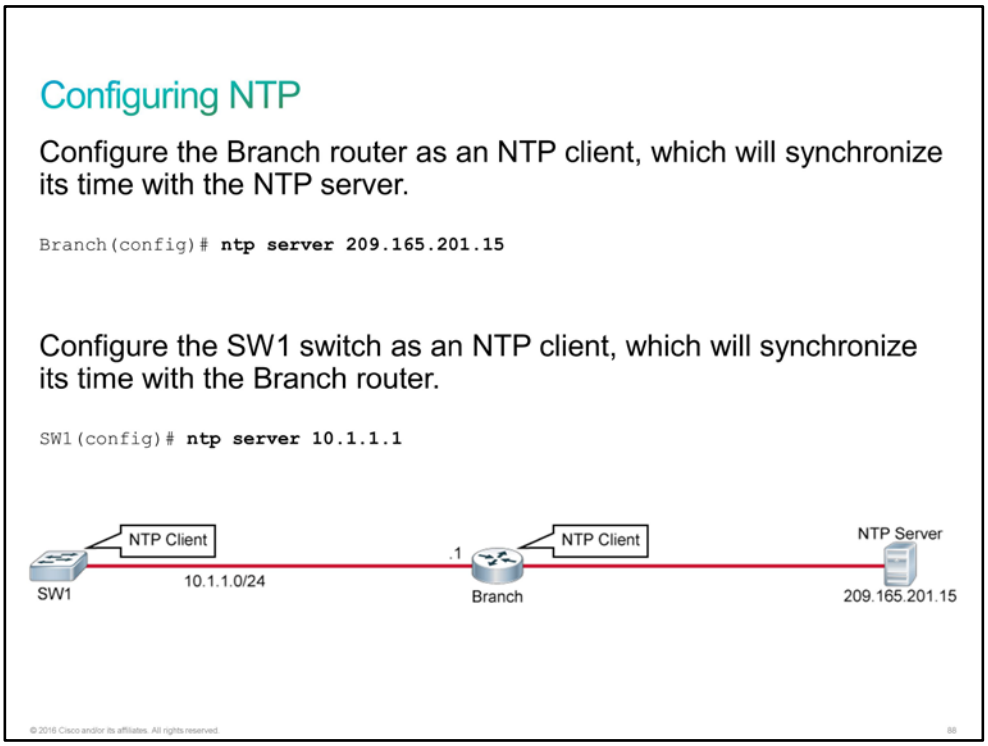

The figure shows an example configuration scenario. Both the router Branch and switch SW1 are configured as NTP clients using the **ntp server** *ip-address* global configuration command. The IP address of the NTP server is configured.

A Cisco IOS device acting as an NTP client will also respond to received time requests. This factor enables SW1 to sync directly with the router Branch and optimize traffic flows. Alternatively, you could configure the switch SW1 to sync with an external NTP server as well.

Cisco IOS devices can also act as NTP servers. To configure Cisco IOS Software as an NTP master clock to which peers synchronize themselves, use the **ntp master** command in the global configuration mode: **ntp master** [*stratum*]

Note	Use this command with caution. You can easily override valid time sources using this command, especially if a low stratum number is configured. Configuring multiple devices in the same network with the **ntp master** command can cause instability in keeping time if the devices do not agree on the time.

The *stratum* value is a number from 1 to 15. The lowest stratum value indicates a higher NTP priority. It also indicates the NTP stratum number that the system will claim.

Verifying NTP

To verify NTP on Cisco devices, use the following commands:

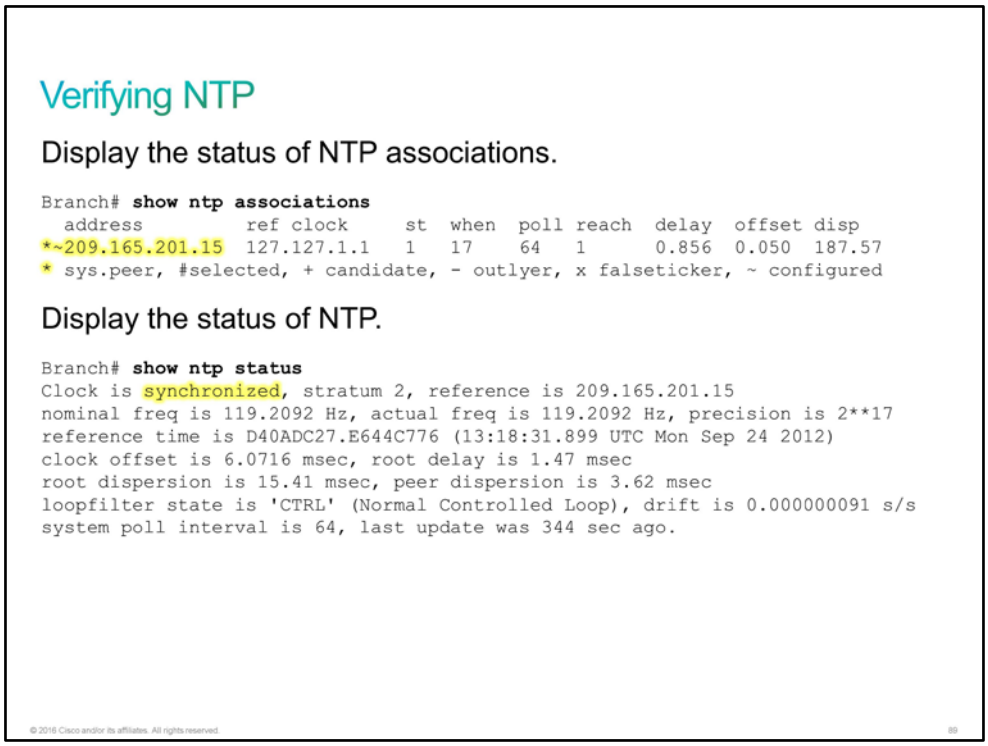

To display the status of NTP associations, use the **show ntp associations** command in the privileged EXEC mode.

The output displays these significant fields:

- ***:** The peer that is synchronized to this peer
- **~:** The peer that is is statically configured
- **address:** The address of the peer
- **st:** The stratum setting for the peer

Note	It may take several minutes for an NTP client to synchronize with the NTP server.

To display the status of NTP, use the **show ntp status** command in the user EXEC mode.

The output displays these significant fields:

- **synchronized:** The system that is synchronized to an NTP peer
- **stratum:** The NTP stratum of this system
- **reference:** The address of the peer to which a clock is synchronized

Discovery 38: Configure and Verify NTP

Introduction

Network devices generate syslog messages to convey important information about events within the network. Syslog messages have time stamps that are associated with them. For these time stamps to be of value for security analysis, the clocks on all the network devices must be in sync. NTP is the preferred method to achieve synchronisation.

This discovery lab will guide you through configuring and verifying NTP services on Cisco IOS routers. The lab is prepared as depicted in the topology diagram and the connectivity table.

Topology

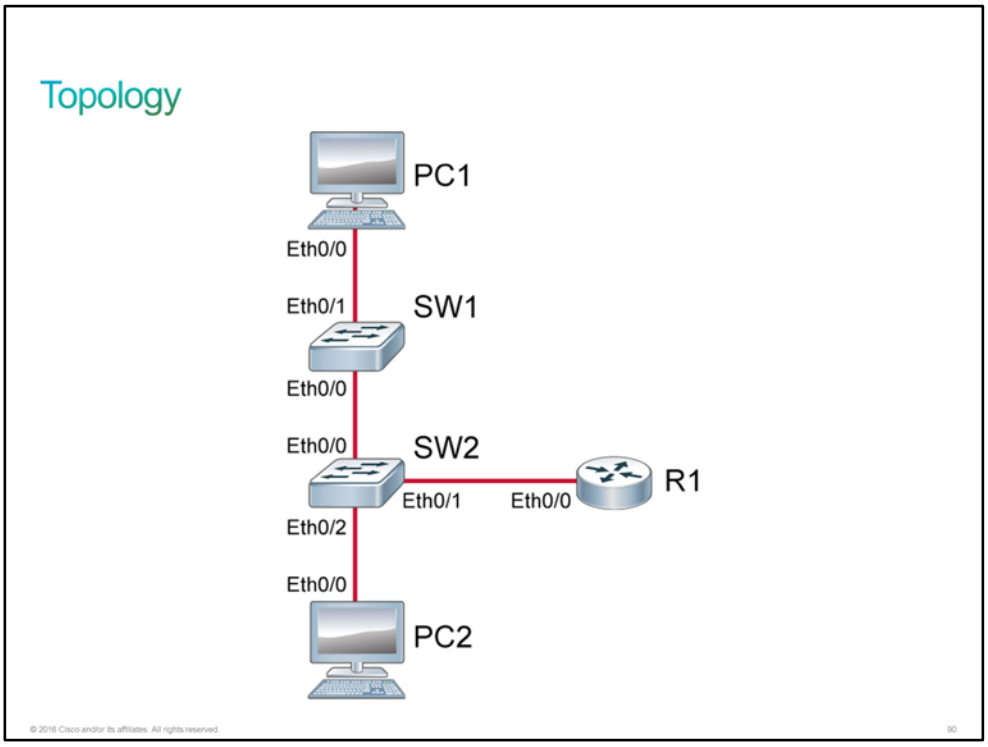

Job Aid

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	10.10.1.20/24
PC2	Default gateway	10.10.1.1
SW1	Hostname	SW1
SW1	VLAN 1 IP address	10.10.1.2/24
SW1	Default gateway	10.10.1.1
SW1	Ethernet0/0 description	Link to SW2
SW1	Ethernet0/1 description	Link to PC1
SW2	Hostname	SW2
SW2	VLAN 1 IP address	10.10.1.3/24
SW2	Default gateway	10.10.1.1
SW2	Ethernet0/0 description	Link to SW1
SW2	Ethernet0/1 description	Link to R1
SW2	Ethernet0/2 description	Link to PC2
R1	Hostname	R1
R1	Ethernet0/0 description	Link to SW2
R1	Ethernet0/0 IP address	10.10.1.1/24
R1	Loopback 0 IP	10.10.3.1/24

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 NTP

Activity

Step 1 Start by reviewing the clocks on SW1, SW2 and R1. You will find that in the emulated lab environment, the clocks are actually synchronized by default.

On SW1, enter the following command:

```
SW1# sh clock
*00:46:09.379 PST Tue Nov 24 2015
```

On SW2, enter the following command:

```
SW2# sh clock
*00:46:11.251 PST Tue Nov 24 2015
```

On R1, enter the following command:

```
R1# sh clock
*00:46:12.604 PST Tue Nov 24 2015
```

The difference in time is only the time it took you to switch from one console to the next and enter the **show clock** command.

Of course, the times that this output and the following output examples depict will differ from what you can see in the lab environment.

Step 2 Access the console of R1 and configure it as an NTP server by enabling the master clock status.

On R1, enter the following commands:

```
R1# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
R1(config)# ntp master
R1(config)# end
R1#
```

Step 3 Configure SW2 to use R1 (10.10.1.1) as its NTP server.

On SW2, enter the following commands:

```
SW2# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
SW2(config)# ntp server 10.10.1.1
SW2(config)# end
SW2#
```

Step 4 Display the current NTP associations and NTP status on SW2.

On SW2, enter the following commands:


```
SW2# show ntp associations

  address         ref clock       st   when   poll reach  delay  offset   disp
*~10.10.1.1       127.127.1.1      8     49     64     1  0.000   0.000 189.47
 * sys.peer, # selected, + candidate, - outlyer, x falseticker, ~ configured
SW2# show ntp status
Clock is synchronized, stratum 9, reference is 10.10.1.1
nominal freq is 250.0000 Hz, actual freq is 250.0000 Hz, precision is 2**10
ntp uptime is 6600 (1/100 of seconds), resolution is 4000
reference time is D9FEA2DD.74FDF4F8 (00:48:29.457 PST Tue Nov 24 2015)
clock offset is 0.0000 msec, root delay is 1.00 msec
root dispersion is 4381.02 msec, peer dispersion is 189.47 msec
loopfilter state is 'CTRL' (Normal Controlled Loop), drift is 0.000000000 s/s
system poll interval is 64, last update was 64 sec ago.
```

Step 5 One at a time, access the consoles of SW2 and R1 and display their clocks. They should be synchronized. The difference in time is due to the time that you spend switching between consoles and entering the command.

On SW2, enter the following command:

```
SW2# show clock
00:50:58.437 PST Tue Nov 24 2015
```

On R1, enter the following command:

```
R1# sh clock
00:51:00.213 PST Tue Nov 24 2015
```

Step 6 On R1, configure CET time zone.

On R1, enter the following commands:

```
R1# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
R1(config)# clock timezone CET 2
Nov 24 08:51:57.876: %SYS-6-CLOCKUPDATE: System clock has been updated from
00:51:57 PST Tue Nov 24 2015 to 10:51:57 CET Tue Nov 24 2015, configured from
console by console.
R1(config)# end
R1#
```

Step 7 Display the current time on R1 and observe that the time zone has changed.

On R1, enter the following command:

```
R1# show clock
10:53:05.222 CET Tue Nov 24 2015
```

This is the end of the discovery lab.

Challenge

1. Which command can you use to help disable multiple ports in a switch?

   A. **interface range**
   B. **interface**
   C. **shutdown range**
   D. **interface range shutdown**

2. Which of the following converts dynamically learned addresses into secure addresses by modifying the running configuration on the fly?

   A. Static Learning
   B. Dynamic Learning
   C. A combination of static and dynamic learning
   D. Sticky Learning

3. You want an interface to error-disable if traffic on the interface violates port-security parameters. Which of the following would you use?

   A. **switchport port-security shutdown**
   B. **switchport port-security violations on**
   C. **switchport port-security violation err-disabled**
   D. **switchport port-security violation shutdown**

4. Check the following command output. What state is the port in?

   SwitchX# show port-security interface FastEthernet 0/5
   Port Security : Enabled
   Port Status : Secure-up
   Violation Mode : Shutdown
   Aging Time : 0 mins
   Aging Type : Absolute
   SecureStatic Address Aging : Disabled
   Maximum MAC Addresses : 1
   Total MAC Addresses : 1
   Configured MAC Addresses : 0
   Sticky MAC Addresses : 0
   Last Source Address:Vlan : fc99.47e5.2598:1
   Security Violation Count : 0

   A. Forwarding
   B. Err-Disabled
   C. Shutdown
   D. Listening

5. Which of the following commands displays the open ports on a router?

   A. **show open-ports**
   B. **show control-plane host**
   C. **show control-plane host open-ports**
   D. **show ports open host**

6. Why is clock synchronization between network devices important?

   A. To ensure that routing protocols on devices can communicate with each other.
   B. To ensure traffic transiting network devices, do not get dropped.
   C. To ensure no security breaches happen due to an explot called 'clock attack'
   D. To ensure the correct interpretation of events within syslog data.

7. Which command will you use to configure a device as an NTP client?

   A. **ntp client**
   B. **ntp server**
   C. **ntp master**
   D. **ntp source**

Answer Key

Challenge

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

Lesson 3: Implementing Advance Security

Introduction

A CCS customer would like some advice on how to mitigate the threats at the access layer of their network. They are also considering implementing RADIUS or TACACS+ servers for authentication to their network devices. The customer also heard that using NMS in their network can help them quickly determine the operation of different network devices.

Mitigating Threats at Access Layer

The access layer is the point at which user devices connect to the network and is therefore the connection point between the network and client device. So, protecting the access layer plays an important role in protecting other users, applications, and the network itself from human errors and malicious attacks.

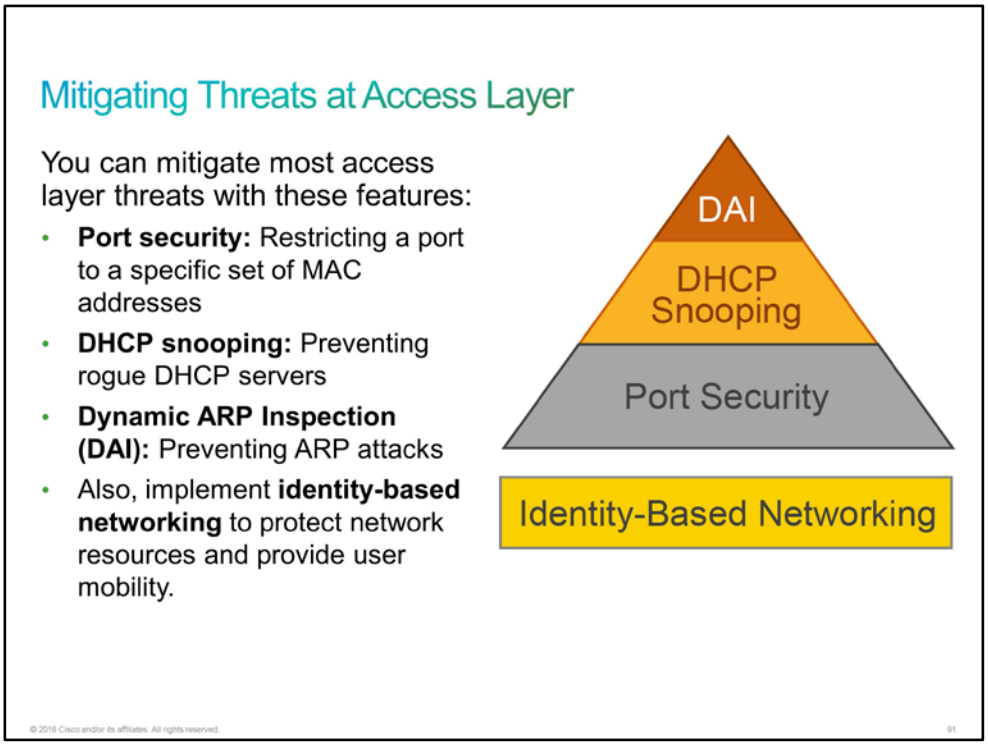

Different security features exist to protect the access layer of your network. Port security, DHCP snooping, Dynamic ARP Inspection, also known as DAI, are only some of them. Besides those features, you can configure identity-based networking, which will provide additional security and protection of your network resources even in the case of user mobility.

Note	The configuration of mentioned techniques is beyond the scope of this course (with the exception of port security, which you are already familiar with).

DHCP Snooping and DAI

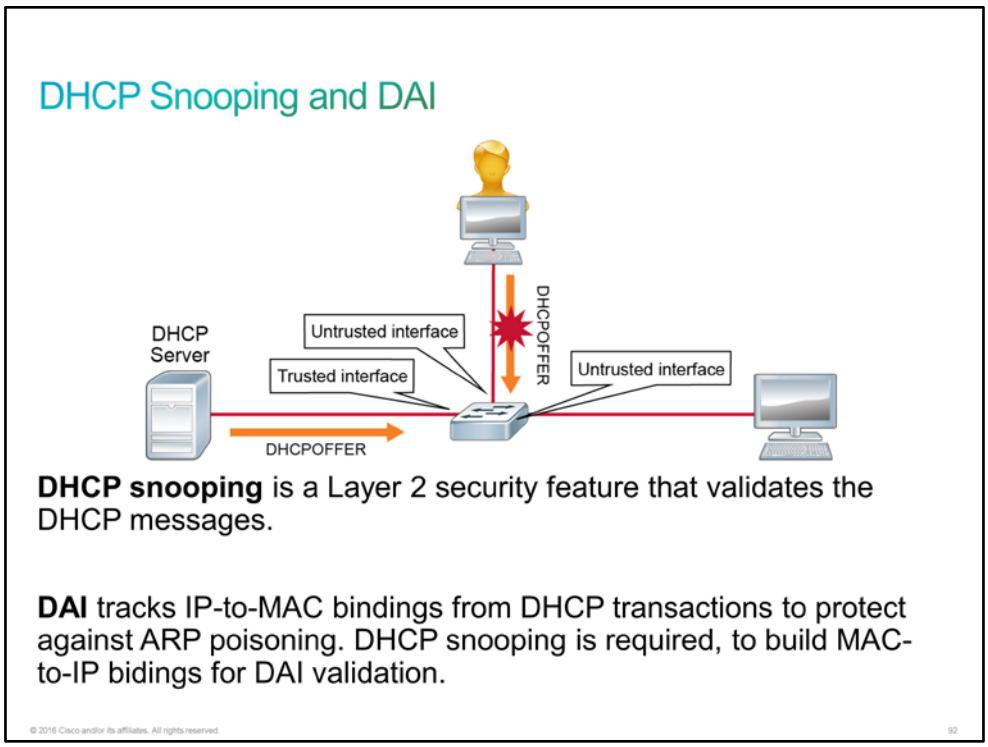

DHCP snooping is a Layer 2 security feature that acts like a firewall between untrusted hosts and trusted DHCP servers. The primary function of the DHCP snooping is to prevent rogue DHCP servers in the network. Interfaces on the switches are configured as trusted or untrusted. Trusted interfaces allow all types of DHCP messages, while untrusted interfaces allow only requests. Trusted interfaces are interfaces that connect to a DHCP server or are an uplink towards the DHCP server.

With DHCP snooping enabled, a switch also builds a DHCP snooping binding database. Each entry in the database includes the MAC address of the host, the leased IP address, the lease time, the binding type, the VLAN number, and the interface information that is associated with the host. Other security features, such as Dynamic ARP Inspection, also use this DHCP snooping binding database.

Dynamic ARP Inspection intercepts all ARP requests and all replies on the untrusted ports. It verifies each intercepted packet for a valid IP-to-MAC binding based on the database that DHCP snooping builds. The device either drops or logs ARP replies coming from invalid devices. This way, it prevents ARP poisoning attacks.

Identity-Based Networking

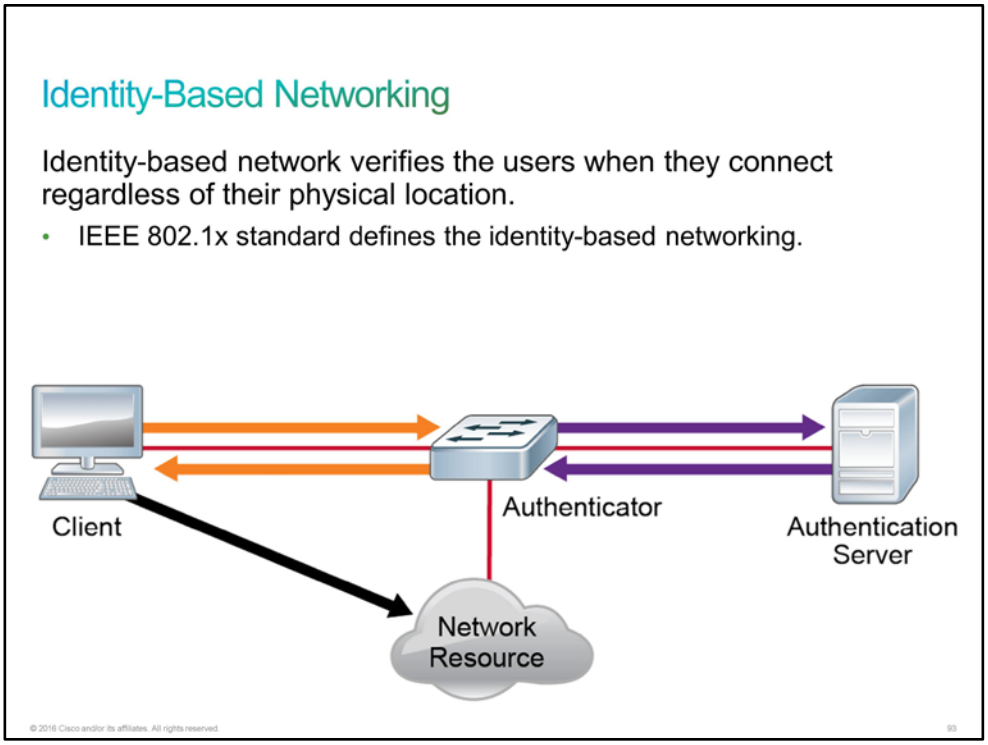

Identity-based networking is a concept that unites several authentication, access control, and user policy components with the aim to provide users with the network services that they are entitled to.

Traditional LAN security depends on physical security of the network ports. In order to gain access to the accounting VLAN, a user has to walk into the accounting department and plug the device in an Ethernet port. With user mobility as one of the core requirements of modern enterprise networks, this dependency is no longer practical, and it does not provide sufficient security.

Identity-based networking allows you to verify users when they connect to a switch port. Identity-based networking authenticates users and places them in the right VLAN based on their identity. Should any users fail to pass the authentication process, their access can be rejected, or they might be simply put in a guest VLAN.

The IEEE 802.1x standard allows you to implement the identity-based networking based on the client-server access control. These three roles are defined by the standard:

- **Client:** Also known as the supplicant. It is the workstation with 802.1x-compliant client software.
- **Authenticator:** Usually the switch, which controls the physical access to the network. It acts as a proxy between the client and authentication server.
- **Authentication server (RADIUS):** The server that authenticates each client that connects to a switch port before making available any services that the switch or the LAN behind offer.

External Authentication Options

Administrative access to a specific network device should be secured so that only authenticated users can access the device.

In a small network, local authentication is often used. When you have more than a few user accounts in a local device database, managing those user accounts becomes more complex. For example, if you have 100 network devices, adding one user account means that you have to add this user account on all 100 devices in the network. Also, when you add one network device to the network, you have to add all user accounts to the local device database to enable all users to access that device.

Because maintaining the local database for each network device for the size of the network is usually not feasible, you can use an external AAA server that will manage all user and administrative access needs for an entire network.

Note	AAA commonly stands for *authentication*, *authorization*, and *accounting*. It refers to a security architecture for distributed systems that enables control over which users are allowed access to which services and how much resources they have used.

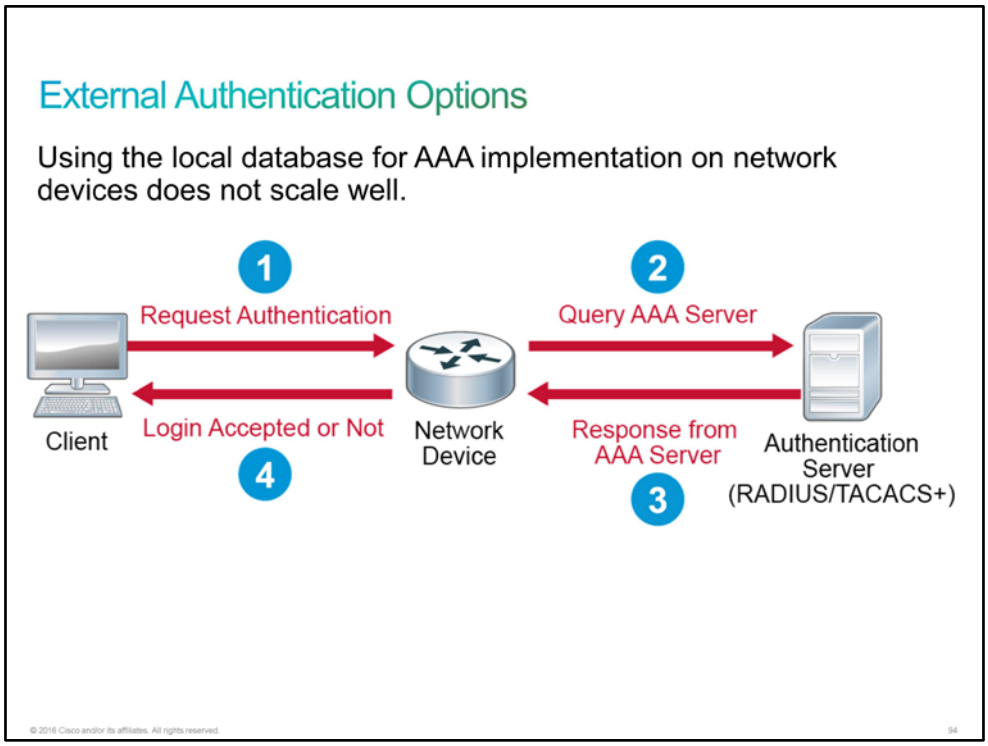

The two most popular options for external AAA are as follows:

- **RADIUS:** RADIUS is an open standard that combines authentication and authorization services as a single process—after users are authenticated, they are also authorized. It uses UDP for the authentication and authorization service.
- **TACACS+:** TACACS+ is a Cisco proprietary security mechanism that separates AAA services. Because it has separated services, you can use TACACS+ only for authorization and accounting, while using another method of authentication. It uses TCP for all three services.

By using the RADIUS or TACACS+ authentication, all authentication requests are relayed to the external server, which allows or denies the user according to its user database. The server then instructs the network device to allow or deny access.

The previous figure shows the external authentication process:

1. A host connects to the network. It can use any communication protocol, depending on the host. At this point, the host is prompted for a username and password.
2. The network device passes a RADIUS/TACACS+ access request, along with user credentials, to the authentication server.
3. The authentication server uses an identity that is stored to validate user credentials.
4. The authentication server sends a RADIUS/TACACS+ response (Access-Accept or Access-Reject) to the network device that will apply the decision.

Discovery 39: Configure External Authentication Using RADIUS and TACACS+

Introduction

This discovery will guide you through the configuration of external authentication using RADIUS and TACACS+. The live virtual lab is prepared with the router, PC, and server that are represented in the topology diagram and the connectivity table. The devices have their basic configurations in place, including hostnames and IP addresses. In the discovery, you will configure a console and vty access on the router using RADIUS and TACACS+ servers.

Topology

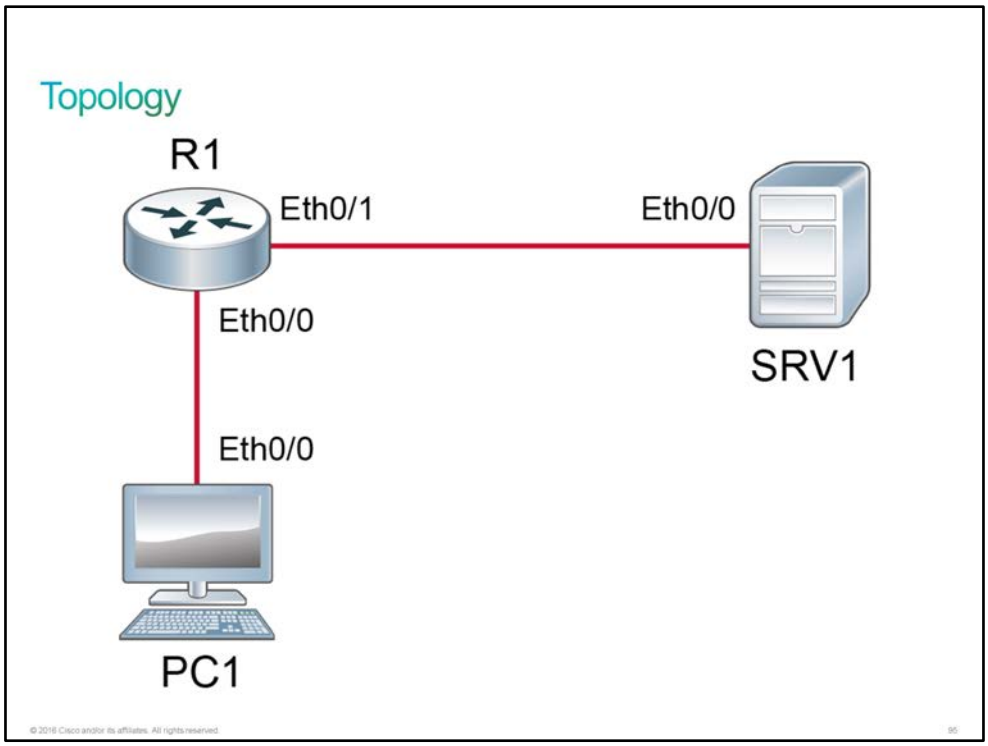

Job Aid

Device Information

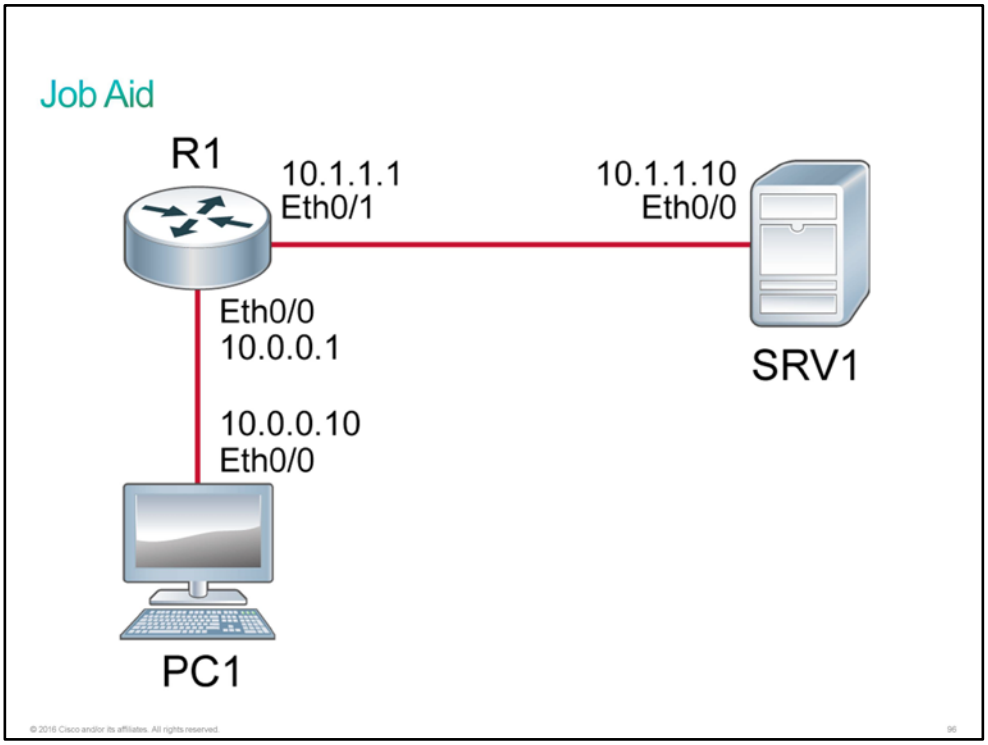

The configuration is as follows:

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

Device Details

Device	Interface	Neighbor	IP Address
PC1	Ethernet0/0	R1	10.0.0.10/24
R1	Ethernet0/0	PC1	10.0.0.1/24
R1	Ethernet0/1	SRV1	10.1.1.1/24
SRV1	Ethernet0/0	R1	10.1.1.10/24

Note	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: Configure RADIUS for Console and VTY Access

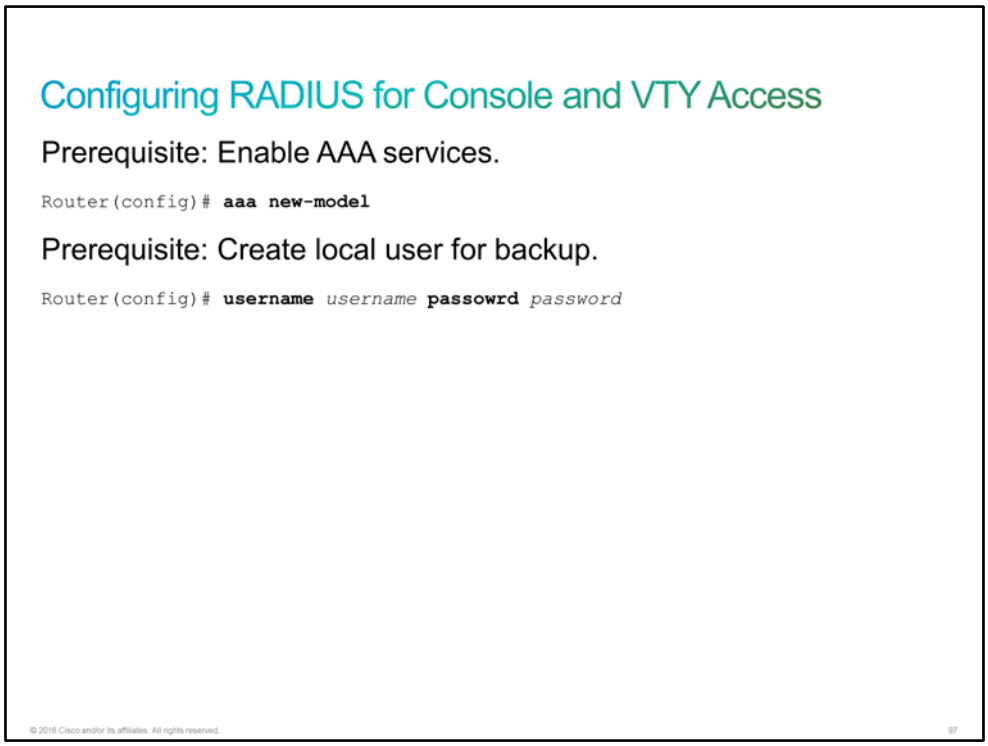

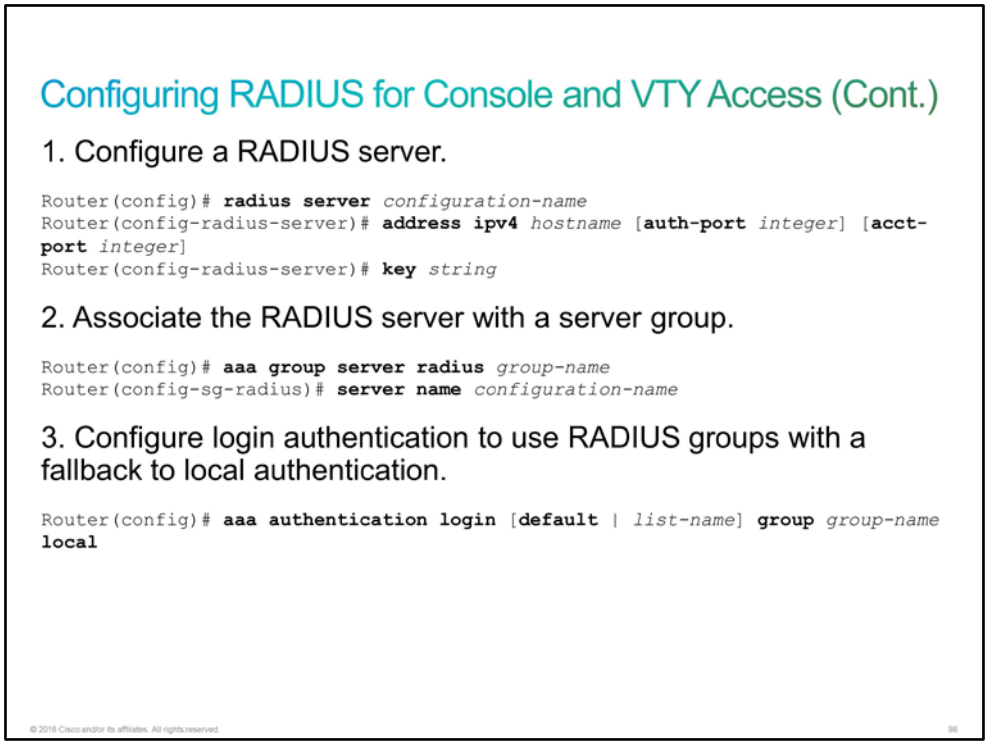

Before starting with the RADIUS configuration, you need to enable AAA services and configure local username and password to avoid being locked out.

```
Router(config)# aaa new-model
Router(config)# username username password  password
```

The RADIUS AAA configuration then starts with the configuration of a RADIUS server:

```
Router(config)# radius server configuration-name
Router(config-radius-server)# address ipv4 hostname  [auth-port integer] [ acct-port
integer]
Router(config-radius-server)# key string
```

You need to specify the *hostname*, or the IP address of the server. Optionally, you can specify a custom port number for the UDP communication, if your RADIUS server is listening on nondefault ports. Port numbers for authentication and accounting differ. The **key** string specifies the authentication and encryption key that is used between the access device and the RADIUS server. This value must match on both devices.

Next, you need to add the RADIUS server to a server group. You can add multiple RADIUS servers to a group, as long as they were previously defined using the **radius server** command.

```
Router(config)# aaa group server radius group-name
Router(config-sg-radius)# server name configuration-name
```

Then you have to configure the device to actually use RADIUS server group for login authentication. Optionally, you can also specify to fallback to local authentication.

```
Router(config)# aaa authentication login [default | list-name] group group-name local
```

The default method list is automatically applied to all interfaces, except those interfaces that have a named method list that is explicitly defined.

Note	You can also specify multiple authentication method lists, using different combinations of server groups and options of local fallback. If you decide to use method lists, you must then apply a specific list also to the console or vty lines.

Activity

Complete the following steps:

Step 1 On R1, configure local user "admin" that will have the "Cisco123" password.

```
R1# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
R1(config)# username admin password Cisco123
```

You can then use this same locally created user if the external authentication server fails.

Step 2 You need to enable AAA services to unhide all AAA commands. Access the console of R1 and configure the **aaa new-model** command in the global configuration mode.

```
R1(config)# aaa new-model
```

The **aaa new-model** command immediately applies local authentication to all lines and interfaces (except console line **line con 0**). To avoid being locked out of the router, you should define a local username and password before starting the AAA configuration.

Step 3 On R1, configure SRV1 as a RADIUS server. Use "radiusPassword" as a shared key.

The configuration name of the server can be anything, but you have to specify the SRV1 IP address as the IPv4 address of the server.

```
R1(config)# radius server myRadiusSRV1
R1(config-radius-server)# address ipv4 10.1.1.10
R1(config-radius-server)# key radiusPassword
R1(config-radius-server)# exit
```

Step 4 On R1, add this newly created RADIUS server to the group.

The configuration name of the group can be anything.

```
R1(config)# aaa group server radius MyRadiusGroup
R1(config-sg-radius)# server name myRadiusSRV1
R1(config-sg-radius)# exit
```

Step 5 Now you have to specify the router to use this RADIUS group for login authentication.

On R1, configure this newly created group to be used for AAA login authentication. If the RADIUS server fails, the fallback to local authentication should be set.

```
R1(config)# aaa authentication login default group MyRadiusGroup local
R1(config)# exit
```

Step 6 Access the console of PC1 and try to connect to R1. Use the "admin" username and "Cisco123" password for login credentials.

Remember that SRV1 is listed as a RADIUS server. Because SRV1 is a virtual server, which is simulated as router in this example, it does not have actual RADIUS capabilities. So, when you try to connect to R1, the RADIUS authentication will not work. Authentication will fall back to the local authentication, and you will be able to use local credentials that you created earlier.

```
PC1# telnet 10.0.0.1
Trying 10.0.0.1 ... Open


User Access Verification

Username: admin
Password:

R1# exit

[Connection to 10.0.0.1 closed by foreign host]
PC1#
```

Note: Because R1 first tries to authenticate you on the RADIUS server and then falls back to the local database, the authentication process may take a bit longer.

Task 2: Configure TACACS+ for Console and VTY Access

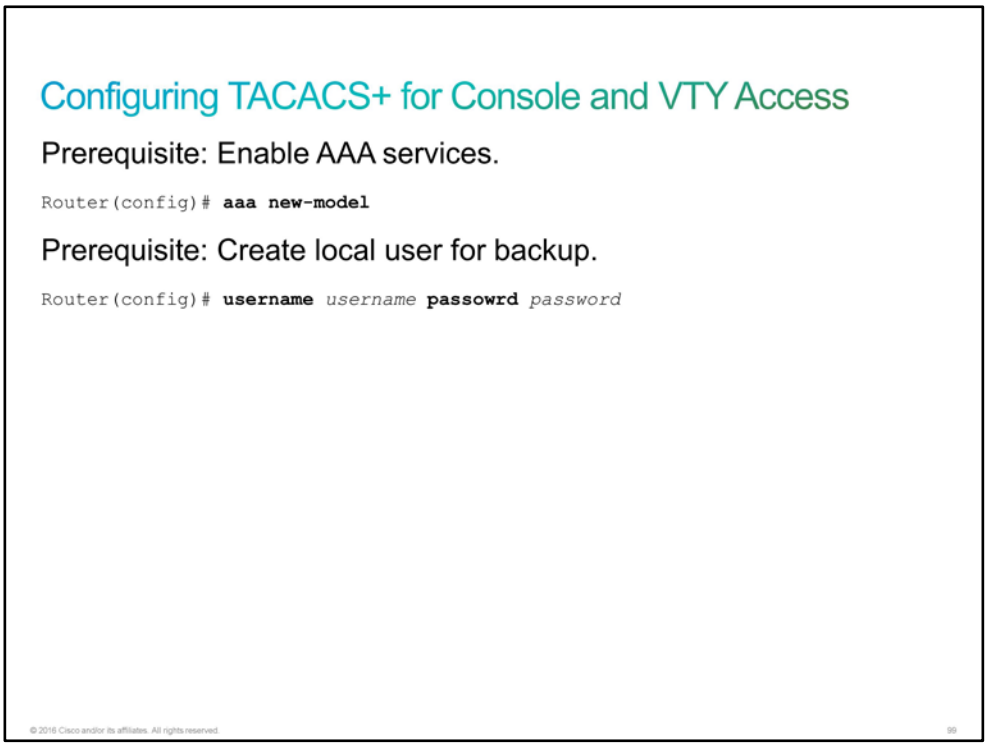

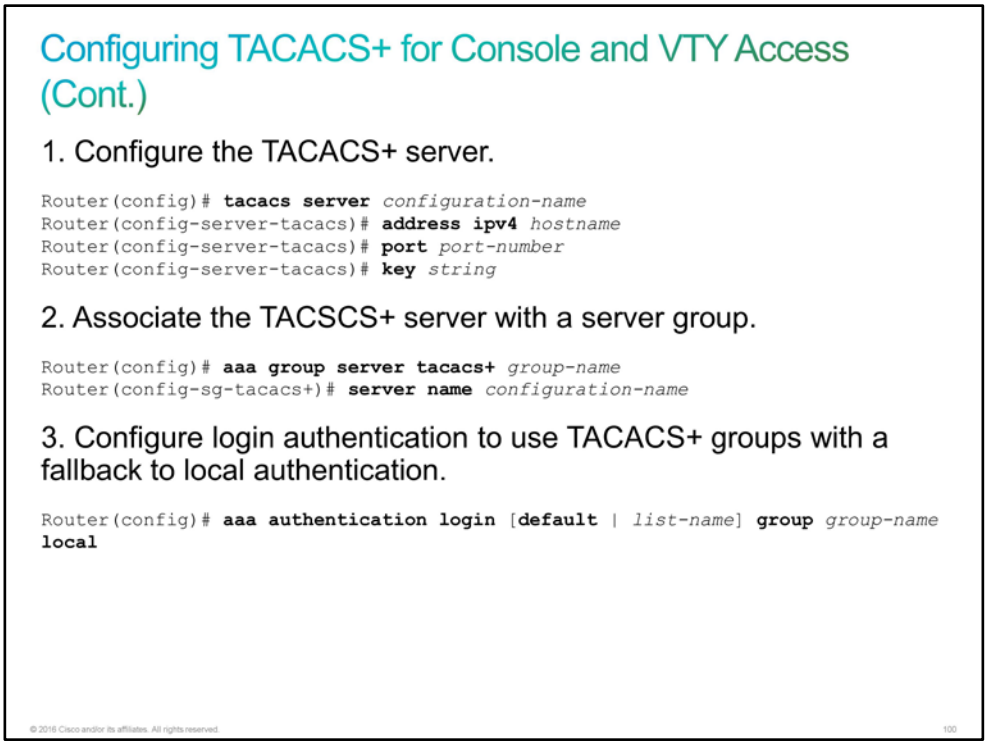

TACACS+ AAA configuration is nearly identical to the RADIUS configuration. Before starting with the TACACS+ configuration, you need to enable AAA services and configure local username and password to avoid being locked out.


```
Router(config)# aaa new-model
Router(config)# username username password password
```

Then, you can configure the TACACS+ server.

```
Router(config)# tacacs server configuration-name
Router(config-server-tacacs)# address ipv4 hostname
Router(config-server-tacacs)# port port-number
Router(config-server-tacacs)# key string
```

You need to specify the *hostname*, or the IP address of the server. Optionally, you can specify a custom port number for the TCP communication, if your TACACS+ server is listening on nondefault ports. The **key** string specifies the encryption key that is used for encrypting all traffic between the access device and TACACS+ server. This value must match on both devices.

Next, you need to add the TACACS+ server to a server group. You can add multiple TACACS+ servers to a group, as long as they were previously defined using the **tacacs server** command**.**

```
Router(config)# aaa group server tacacs+ group-name
Router(config-sg-tacacs+)# server name configuration-name
```

Then you have to configure the device to actually use the TACACS+ server group for login authentication. Optionally, you can also specify to fall back to local authentication.

```
Router(config)# aaa authentication login [default | list-name] group group-name local
```

The default method list is automatically applied to all interfaces except those interfaces that have a named method list that is explicitly defined.

Activity

Complete the following steps:

Step 1 You first need to enable AAA services and create a local user.

Because you have already configured this part in the previous procedure, you can proceed to the next step.

Step 2 Access the console of R1 and configure SRV1 as a TACACS+ server. Use "tacacsPassword" as a shared key.

The configuration name of the server can be anything, but you have to specify the SRV1 IP address as the IPv4 address of the server.

```
R1# conf t
R1(config)# tacacs server myTacacsSRV1
R1(config-server-tacacs)# address ipv4 10.1.1.10
R1(config-server-tacacs)# key tacacsPassword
R1(config-server-tacacs)# exit
```

Step 3 On R1, add this newly created TACACS+ server to the group.

The configuration name of the group can be anything.

```
R1(config)# aaa group server tacacs+ MyTacacsGroup
R1(config-sg-tacacs+)# server name myTacacsSRV1
R1(config-sg-tacacs+)# exit
```

Step 4 Now you have to specify the router to use this TACACS+ group for login authentication.

On R1, configure this newly created group to be used for AAA login authentication. If the TACACS+ server fails, the fallback to local authentication should be set.

```
R1(config)# aaa authentication login default group MyTacacsGroup local
R1(config)# exit
```

Note that this configuration will overwrite the previously specified authentication method using the RADIUS server because you can specify only one group (RADIUS or TACACS+) with the default method list.

Step 5 Access the console of PC1 and try to connect to R1. Use "admin" and "Cisco123" login credentials.

Remember that SRV1 is listed as the TACACS+ server. Because SRV1 is a virtual server, which is simulated as a router in this example, it does not have actual TACACS+ capabilities. So, when you try to connect to R1, the TACACS+ authentication will not work. Authentication will fall back to local authentication, and you will be able to use local credentials that you created earlier.

```
PC1# telnet 10.0.0.1
Trying 10.0.0.1 ... Open


User Access Verification

Username: admin
Password:

R1# exit

[Connection to 10.0.0.1 closed by foreign host]
PC1#
```

Note: Because R1 first tries to authenticate you on the TACACS+ server and then falls back to the local database, the authentication process may take a bit longer.

This is the end of the discovery lab.

Challenge

1. Which of the following will mitigate access layer threats ? (Choose two)

   A. Port Security
   B. Layer 3 IP Access Lists
   C. Dynamic ARP Inspection
   D. AAA

2. Which of the following is not true about DHCP snooping ?

   A. Validates DHCP messages received from untrusted sources and filters out invalid messages
   B. Builds and maintains the DHCP snooping binding database, which contains information about untrusted hosts with leased IP addresses.
   C. Rate-limits DHCP traffic from trusted and untrusted sources.
   D. DHCP snooping is a Layer 2 security feature that acts like a firewall between hosts.

3. Which of the following command will enable AAA on router ?

   A. **aaa enable**
   B. **enable aaa**
   C. **new-model aaa**
   D. **aaa new-model**

4. Which of the following are true about TACACS+ ? (Choose two)

   A. TACACS+ is a Cisco proprietary security mechanism.
   B. TACACS+ uses UDP.
   C. TACACS+ combines authentication and authorization services as a single process—after users are authenticated, they are also authorized.
   D. TACACS+ uses TCP.

5. Which of the following is not true about RADIUS ?

   A. RADIUS is an open standard protocol
   B. RADIUS separates AAA services.
   C. RADIUS uses UDP.
   D. RADIUS encrypts only the password in the access-request packet, from the client to the server. The remainder of the packet is unencrypted.

Answer Key

Challenge

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

Module 8: Implementing an EIGRP-Based Solution

Introduction

EIGRP is an advanced distance vector routing protocol. EIGRP was a Cisco proprietary protocol, so all routers in a network that is running EIGRP had to be Cisco routers. Partial functionality of EIGRP was converted to an open standard in 2013. EIGRP is often considered a hybrid protocol because it also sends link state updates when link states change. EIGRP is an interior gateway protocol that is suited for many different topologies and media. In a well designed network, EIGRP scales well and provides extremely quick convergence times with minimal network traffic.

In this module, you will learn how to implement basic EIGRP configuration both for IPv4 and IPv6 and how to verify the operation of this routing protocol. You will also perform basic troubleshooting steps for common EIGRP issues and configuration mistakes.

Lesson 1: Implementing EIGRP

Introduction

A new client calls CCS and reports slow network response. After speaking with the network administrator, Bob decides that the network issues can be resolved by moving this customer from RIP to a more robust routing protocol. Bob explains the benefits of EIGRP to the customer and they agree to an onsite engagement. You will need to go onsite to the new company, shut down RIP, and configure EIGRP. You should know the technology behind EIGRP before you go onsite so you can answer any customer inquiries while on the job.

Dynamic Routing Protocols

A routing protocol is a set of processes, algorithms, and messages that are used to exchange routing information. Routing information is used to populate the routing table with the best paths to destinations on the network. As routers learn of changes to network reachability, this information is dynamically passed onto other routers.

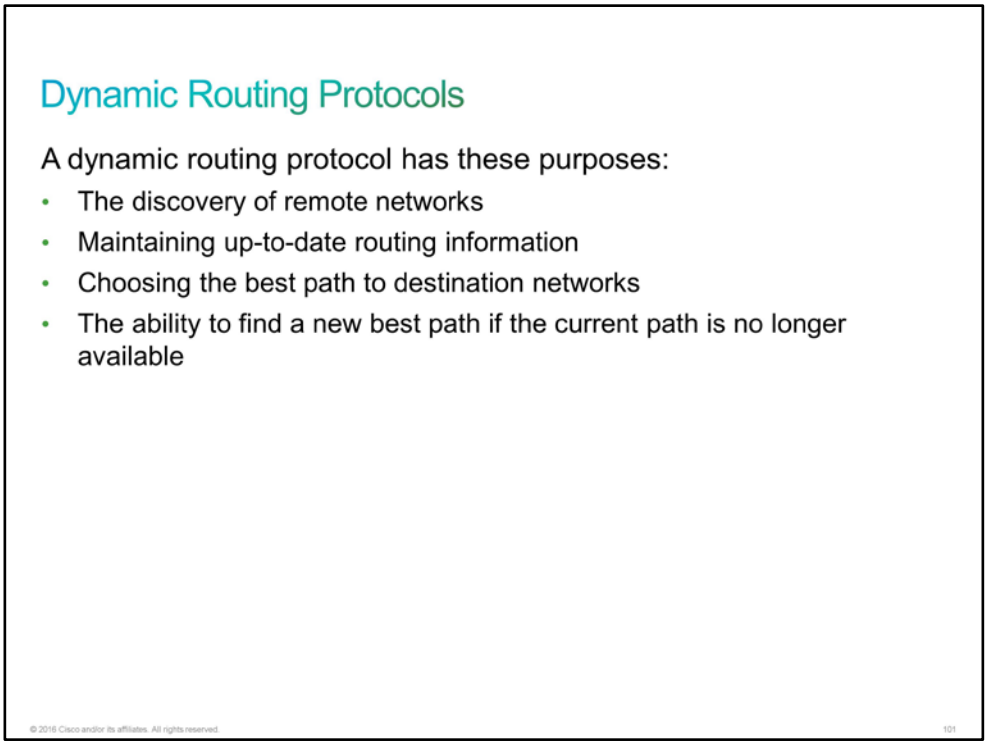

All routing protocols have the same purpose: to learn about remote networks and to quickly adapt whenever there is a change in the topology. The method that a routing protocol uses to accomplish this purpose depends upon the algorithm that it uses and the operational characteristics of this protocol. The operations of a dynamic routing protocol vary, depending on the type of routing protocol and on the routing protocol itself.

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. So, the routers that use these protocols must have sufficient processing capacity to implement the algorithms of the protocol and to perform timely packet routing and forwarding.

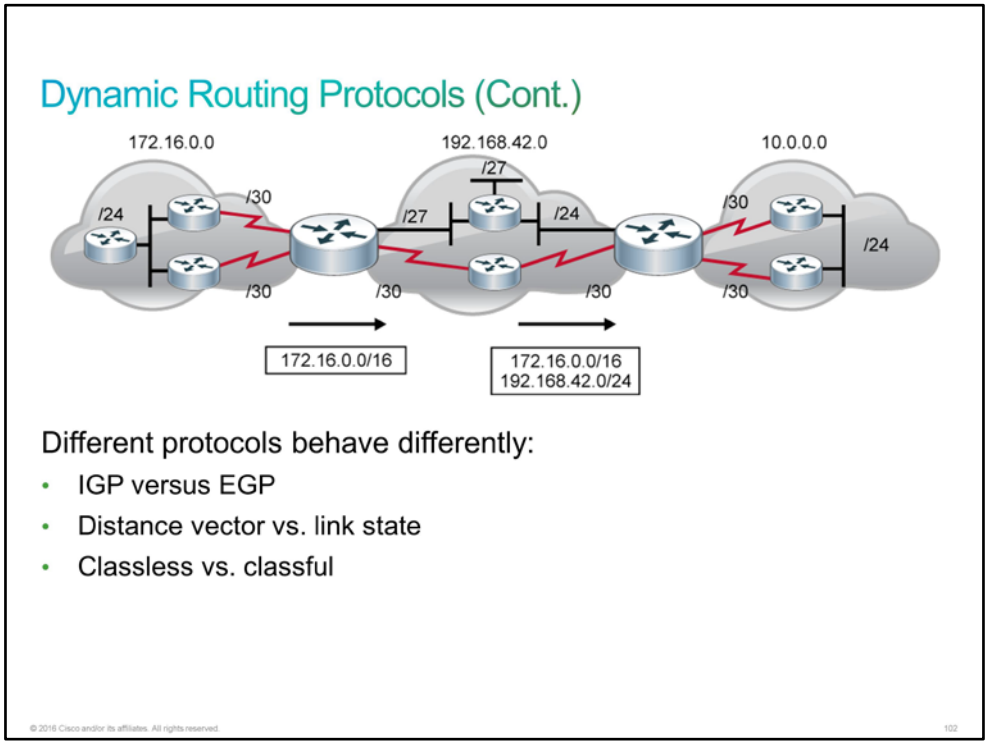

An AS, otherwise known as a routing domain, is a collection of routers under a common administration, such as an internal company network or an ISP network. Because the Internet is based on the AS concept, the following two types of routing protocols are required:

- **IGP:** The IGP routing protocol is used to exchange routing information within an AS. EIGRP, IS-IS, OSPF, and RIP are examples of IGPs.
- **EGP:** The EGP routing protocol is used to route between autonomous systems. BGP is the EGP of choice in networks today.

Within an AS, most IGP routing can be classified as distance vector or link-state routing:

- **Distance vector:** The distance vector routing approach determines the direction (vector) and distance (such as hops) to any link in the internetwork. Some distance vector protocols periodically send complete routing tables to all of the connected neighbors. In large networks, these routing updates can become very large, 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. Distance vector routing protocols do not have an actual map of the network topology. RIP is an example of a distance vector routing protocol while EIGRP is an advanced distance vector routing protocol.
- **Link state:** The link-state approach, which uses the SPF algorithm, creates an abstract of the exact topology of the entire internetwork, or at least of the partition in which the router is situated. A link-state routing protocol is like having a complete map of the network topology. 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. The OSPF and IS-IS protocols are examples of link-state routing protocols.

Also, there is classful and classless routing:

- **Classful routing protocol:** Classful routing protocol is a consequence of the fact that subnet masks are not advertised in the routing advertisements that most distance vector routing protocols generate. When a classful routing protocol is used, all subnetworks of the same major network (Class A, B, or C) must use the same subnet mask, which is not necessarily a default major class subnet mask. Routers that are running a classful routing protocol perform automatic route summarization across network boundaries. Classful routing protocols are obsolete in networks today.
- **Classless routing protocol:** Classless routing protocols can be considered second-generation protocols because they are designed to address limitations of classful routing protocols such as RIPv1 and IGRP. A prime limitation of classful routing protocols is that the subnet mask is not exchanged during the routing update process. This limitation means that the same subnet mask must be used on all subnetworks within the same major network. When you consider point-to-point serial WAN connections, using a 24-bit network prefix is very wasteful when all that is required is a 30-bit network prefix to accommodate the two endpoints.

Another limitation of the classful approach is the need to automatically summarize to the classful network number at all major network boundaries. As an example, using 172.16.0.0/16 as the classful network allows only a single, flat network. If the class B network is subnetted into /24 networks, there are now 255 subnets available. If the company connects to another network, it must advertise the 172.16.0.0/16 summary, because the classful routing protocol does not have the capability to provide subnet-specific routes.

In the classless environment, the summarization process is controlled manually and can usually be invoked at any bit position within the address. Because subnet routes are propagated throughout the routing domain, manual summarization may be required to keep the size of the routing tables manageable. Classless routing protocols include RIPv2, EIGRP, OSPF, and IS-IS.

Administrative Distance

In an enterprise network, it is not uncommon to encounter multiple dynamic routing protocols and static routes configured on Layer 3 devices. If there are several sources for routing information, such as specific routing protocols, static routes, and even directly connected networks, a method is required to rate the trustworthiness of each routing information source in order to select the best path.

Cisco IOS Software uses the concept of administrative distance to select the best path when it learns about the same destination network from two or more routing sources.

Administrative distance ranks the reliability of a routing protocol. Each routing protocol is prioritized in order of most to least reliable (believable) with the help of an administrative distance value.

The administrative distance is an integer from 0 to 255. A routing protocol with a lower administrative distance is considered more trustworthy than the one with a higher administrative distance.

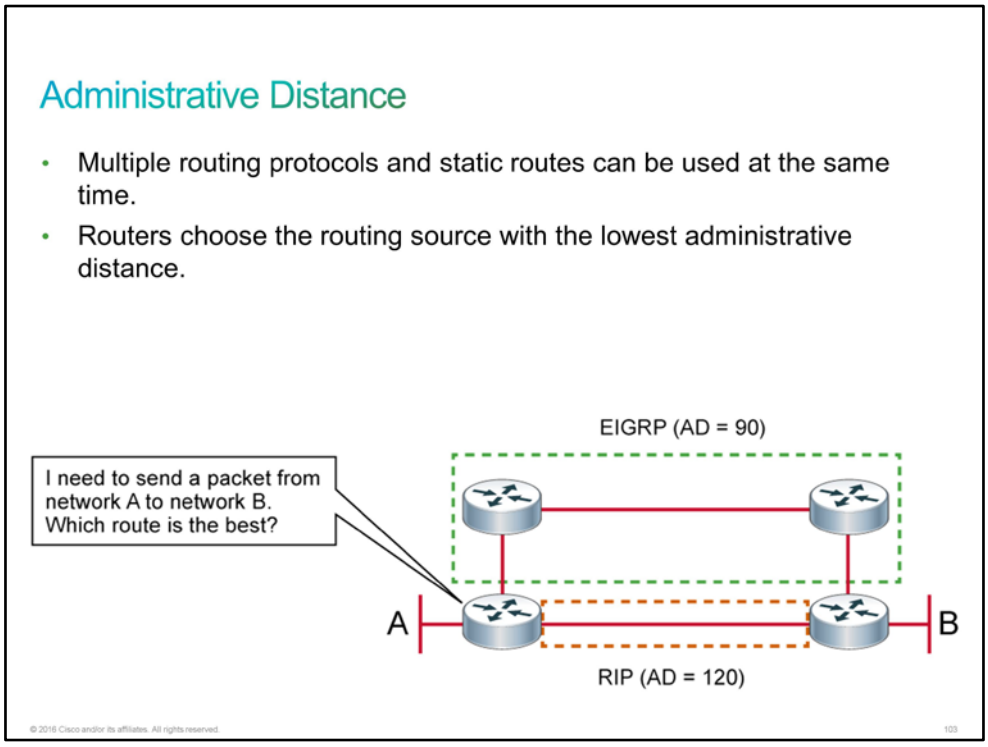

As illustrated in the figure, the router has a packet to deliver from network A to network B. The router must choose between the routes advertised by EIGRP and RIP. Given that there are fewer hops to the destination network via RIP, it appears to be the better choice. However, the EIGRP route has a lower administrative distance than RIP, so the router will choose the route that was advertised by EIGRP and install it in the routing table. If for some reason the path that was advertised by EIGRP goes down, the route that was advertised by RIP will be entered into the routing table.

The table shows the default administrative distance for selected routing information sources.

Note	The default administrative distances can be tuned for each routing protocol.

Route Source	Default Distance
Connected interface	0
Static route	1
EBGP	20
EIGRP	90
OSPF	110
IS-IS	115
RIP	120
External EIGRP	170
IBGP	200
Unreachable	255 (will not be used to pass traffic)

EIGRP Features

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 only partial updates are sent when a change occurs. However, EIGRP is based on the key distance vector routing protocol principle, in which information about the rest of the network is learned from directly connected neighbors.

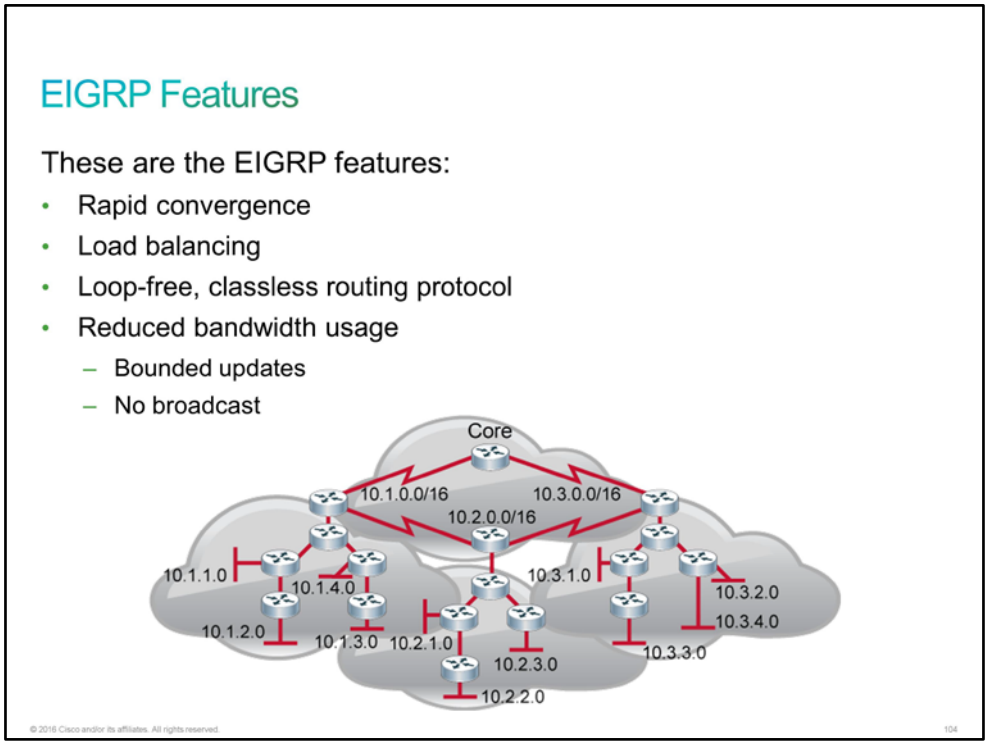

Look into the EIGRP features in more detail:

- **Rapid convergence:** EIGRP uses DUAL to achieve rapid convergence. As the computational engine that runs EIGRP, DUAL resides at the center of the routing protocol, guaranteeing loop-free paths and backup paths throughout the routing domain. A router that uses EIGRP stores all available backup routes for destinations so that it can quickly adapt to alternate routes. If the primary route in the routing table fails, the best backup route is immediately added to the routing table. If no appropriate route or backup route exists in the local routing table, EIGRP queries its neighbors to discover an alternate route.
- **Load balancing:** EIGRP supports unequal metric load balancing and equal metric load balancing, which allows administrators to better distribute traffic flow in their networks.

- **Loop-free, classless routing protocol:** Because EIGRP is a classless routing protocol, it advertises a routing mask for each destination network. The routing mask feature enables EIGRP to support discontiguous subnetworks and VLSMs.
- **Reduced bandwidth usage:** EIGRP updates can be thought of as either "partial" or "bounded." EIGRP does not make periodic updates. The term "partial" means that the update only includes information about the route changes. EIGRP sends these incremental updates when the state of a destination changes, instead of sending the entire contents of the routing table. The term "bounded" refers to the propagation of partial updates that are sent only to those routers that the changes affect. By sending only the routing information that is needed and only to those routers that need it, EIGRP minimizes the bandwidth that is required to send EIGRP updates. EIGRP uses multicast and unicast rather than broadcast. Multicast EIGRP packets use the reserved multicast address of 224.0.0.10. As a result, end stations are unaffected by routing updates and requests for topology information.

EIGRP Path Selection

In the context of dynamic IP routing protocols like EIGRP, the term *path selection* refers to the method by which the protocol determines the best path to a destination IP network.

Each EIGRP router maintains a neighbor table. This table includes a list of directly connected EIGRP routers that have formed an adjacency with this router. Neighbor relationships are used to track the status of these neighbors. EIGRP uses a lightweight Hello protocol to monitor the connection status with its neighbors.

Each EIGRP router maintains a topology table for each routed protocol configuration. The topology table includes route entries for every destination that the router learns from its directly connected EIGRP neighbors. EIGRP chooses the best routes to a destination from the topology table and places these routes in the routing table.

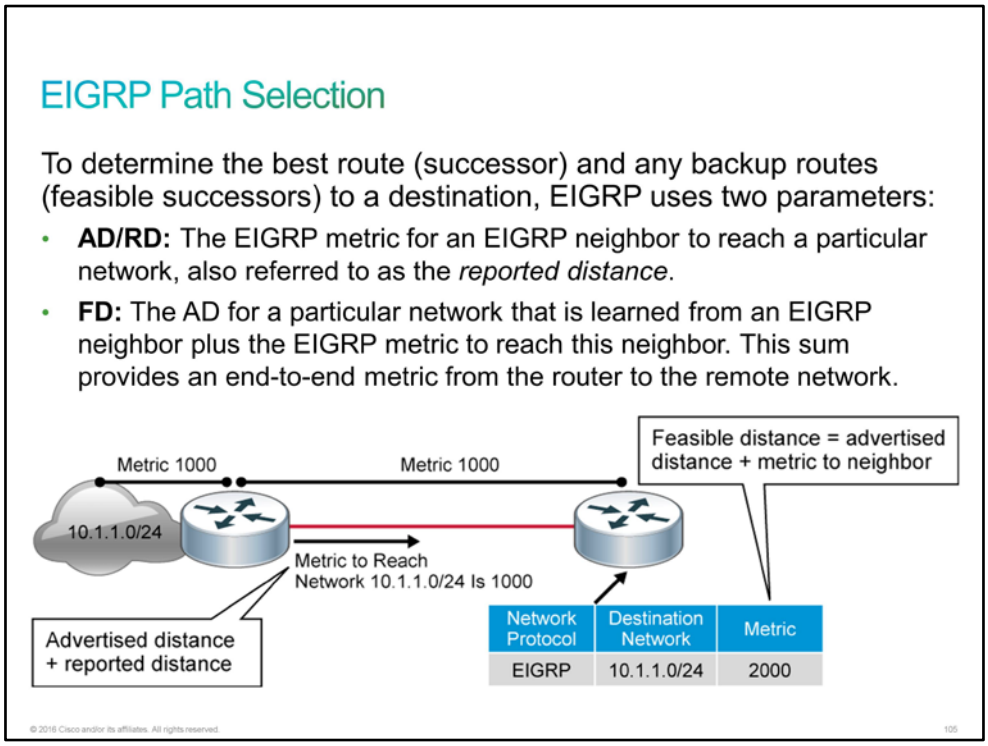

A router compares all FDs to reach a specific network and then selects the lowest FD and places it in the routing table. The FD for the chosen route becomes the EIGRP routing metric to reach this network in the routing table.

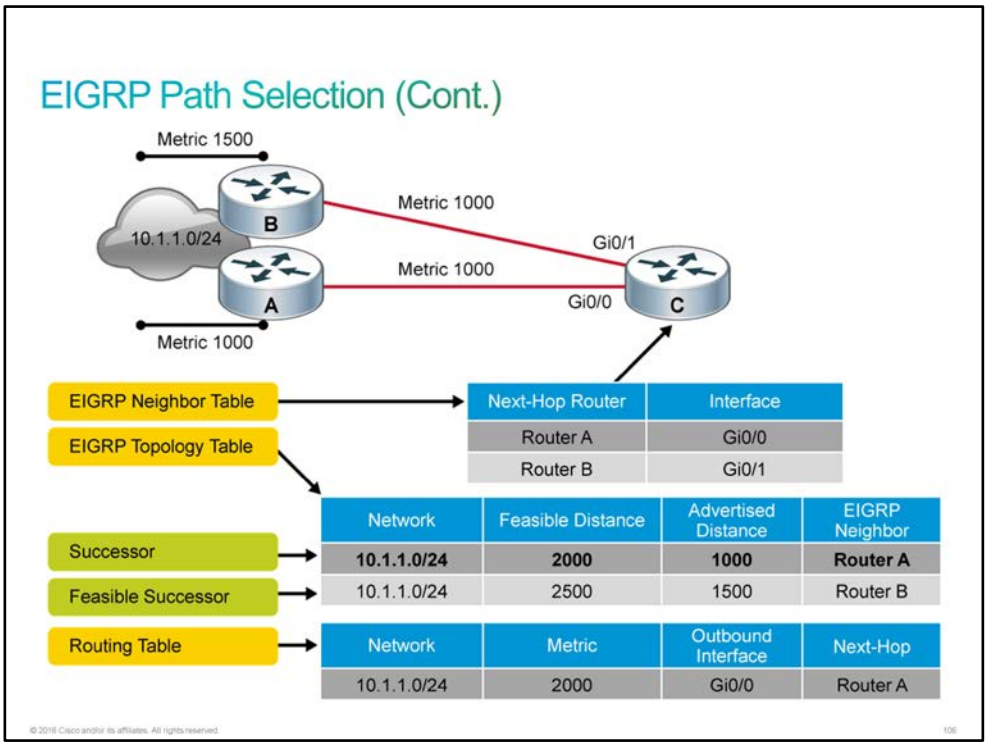

The EIGRP topology database contains all the routes that are known to each EIGRP neighbor. As shown in the example above, routers A and B sent their routing tables to router C, whose table is displayed. Both routers A and B have routes to network 10.1.1.0/24, as well as to other networks that are not shown.

Router C has two entries to reach 10.1.1.0/24 in its topology table. The EIGRP metric for router C to reach both routers A and B is 1000. Add this metric (1000) to the respective AD for each router, and the results represent the FDs that router C must travel to reach network 10.1.1.0/24.

Router C chooses the smallest FD (2000) and installs it in the IP routing table as the best route to reach 10.1.1.0/24. The route with the smallest FD that is installed in the routing table is called the "successor route."

Router C then chooses a backup route to the successor that is called a "feasible successor route," if one or more feasible successor routes exist. To become a feasible successor, a route must satisfy this feasibility condition: A next-hop router must have an AD that is less than the FD of the current successor route (therefore, the route is tagged as a feasible successor). This rule is used to ensure that the network is loop-free.

If the route via the successor becomes invalid, possibly because of a topology change, or if a neighbor changes the metric, DUAL checks for feasible successors to the destination route. If a feasible successor is found, DUAL uses it, avoiding the need to recompute the route. A route will change from a passive state to an active state if no feasible successor exists, and a recomputation must occur to determine the new successor.

Note	In this example, values for the EIGRP metric and for FDs and ADs are optimized for explanation purposes. The real metric values are much larger.

EIGRP Metric

Unlike other routing protocols (such as RIP and OSPF), EIGRP does not use a single attribute to determine the metric of its routes. EIGRP uses a combination of four different features to determine its metric. These features are all physical characteristics of an interface.

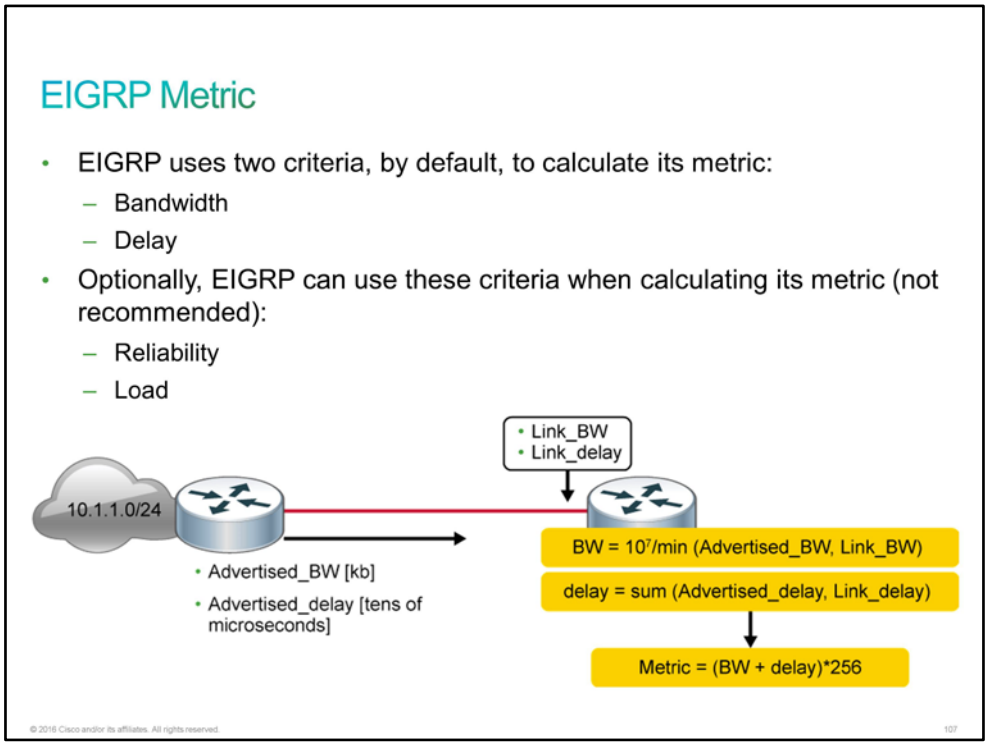

The EIGRP metric can be based on four criteria, but by default, EIGRP uses only two:

- **Bandwidth:** The smallest bandwidth of all outgoing interfaces between the source and destination, in kilobits.
- **Delay:** The cumulative (sum) of all interface delay along the path, in tens of microseconds.

Two additional criteria can be used, but are not recommended because they typically result in frequent recalculation of the topology table:

- **Reliability:** This value represents the worst reliability between the source and destination, which is based on keepalives.
- **Load:** This value represents the worst load on a link between the source and destination, which is computed based on the packet rate and the configured bandwidth of the interface.

The composite metric formula is used by EIGRP to calculate metric value. The formula consists of values K1 through K5, which are known as EIGRP metric weights. By default, K1 and K3 are set to 1, and K2, K4, and K5 are set to 0. The result is that only the bandwidth and delay values are used in the computation of the default composite metric. The metric calculation method (K values) and the EIGRP AS number must match between EIGRP neighbors.

Although an MTU is exchanged in EIGRP packets between neighbor routers, the MTU is not factored into the EIGRP metric calculation.

EIGRP uses scaled values to determine the total metric: 256 * ([K1 * bandwidth] + [K2 * bandwidth] / [256 – Load] + K3 * Delay) * (K5 / [Reliability + K4]), where if K5 = 0, the (K5 / [Reliability + K4]) part is not used (that is, equals 1). Using the default K values, the metric calculation simplifies to 256 * (bandwidth + delay).

Note	EIGRP metric K values are carried in EIGRP hello packets. Therefore a mismatched K value will cause a neighbor to be reset even if that value is unused. The values must be consistently configured throughout the network, and only changed under the recommendation of Cisco.

By using the **show interface** command, you can examine the actual values that are used for bandwidth, delay, reliability, and load in the computation of the routing metric. The output in the figure shows the values that are used in the composite metric for the Serial0/0/0 interface.

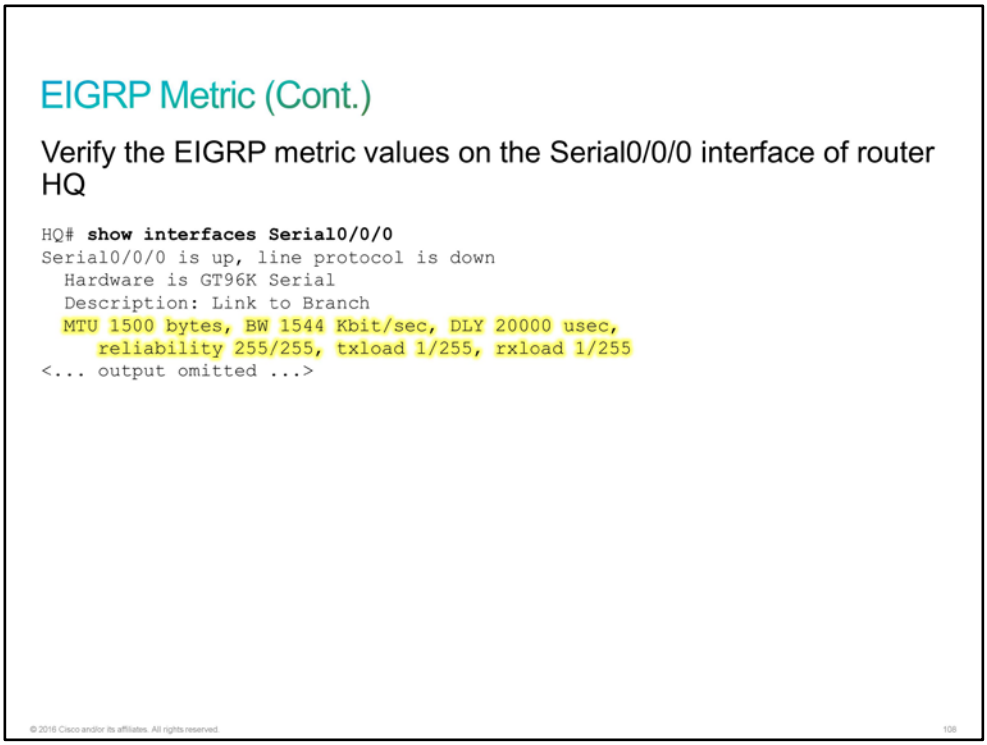

You can influence the EIGRP metric by changing bandwidth and delay on an interface, using **bandwidth** *kbps* and **delay** *microseconds* interface configuration commands.

Discovery 40: Configure and Verify EIGRP

Introduction

This discovery will guide you through the configuration and verification of EIGRP on a Cisco IOS router. The virtual lab is prepared with the devices represented in the topology diagram and the connectivity table. All devices have their basic configurations in place, including hostnames and IP addresses. R2 and R3 are also configured with EIGRP using AS number 1. In this discovery, you will configure EIGRP on R1 and verify the results.

Topology

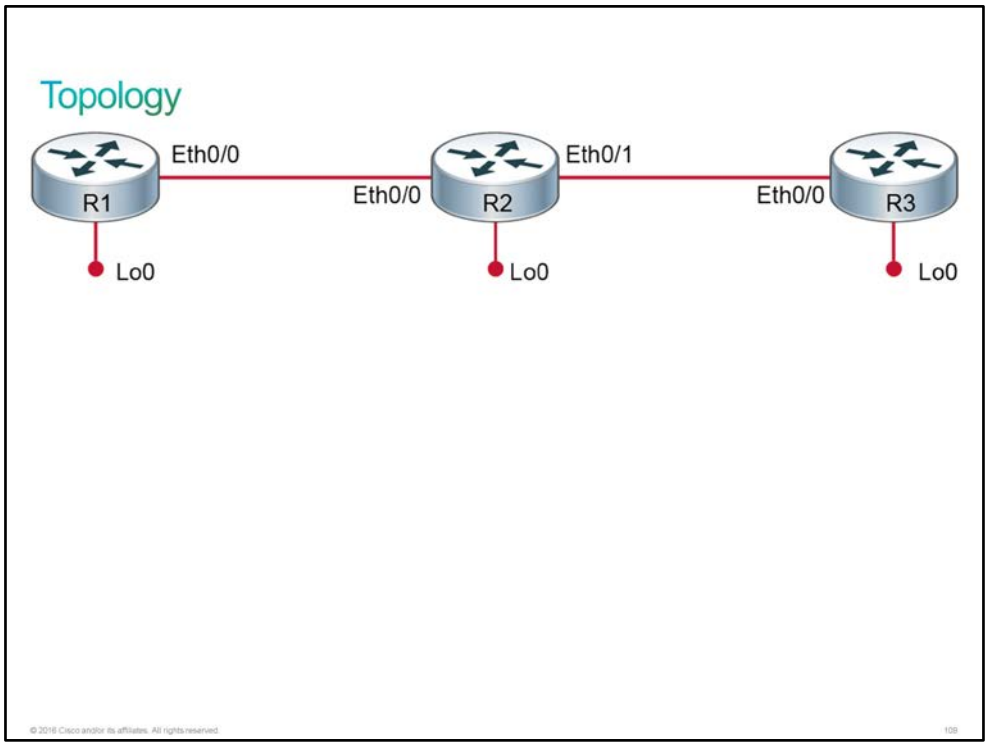

Job Aid

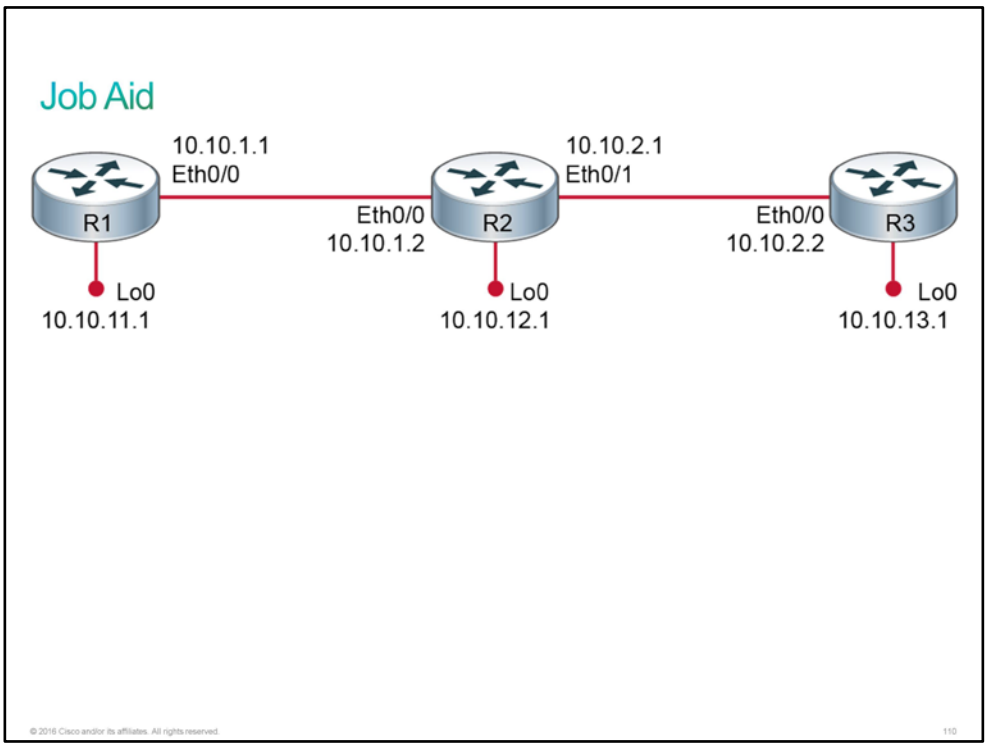

The configuration is as follows:

- All devices have their basic configurations in place, including hostnames and IP addresses.
- EIGRP is preconfigured on R2 and R3:
  - AS number 1 is used.
  - Both routers are announcing Loopback interface network.

Device Information

Device Details

Device	Interface	IP Address	Neighbor
R1	Ethernet0/0	10.10.1.1/24	R2
R1	Loopback0	10.10.11.1/24	—
R2	Ethernet0/0	10.10.1.2/24	R1
R2	Ethernet0/1	10.10.2.1/24	R3
R2	Loopback0	10.10.12.1/24	—
R3	Ethernet0/0	10.10.2.2/24	R2

Device	Interface	IP Address	Neighbor
R3	Loopback0	10.10.13.1/24	—

Task 1: Configure and Verify EIGRP

Activity

Complete the following steps:

Step 1 Access the console of R2 and display EIGRP configuration.

```
R2# show running-config | section eigrp
router eigrp 1
 network 10.0.0.0
```

You should see that EIGRP is preconfigured for AS number 1 and network 10.0.0.0/8 is included.

Step 2 Access the console of R3 and display EIGRP configuration.

```
R3# sh running-config | section eigrp
router eigrp 1
 network 10.0.0.0
```

You should see that EIGRP is preconfigured for AS number 1 and network 10.0.0.0/8 is included.

Configuring EIGRP

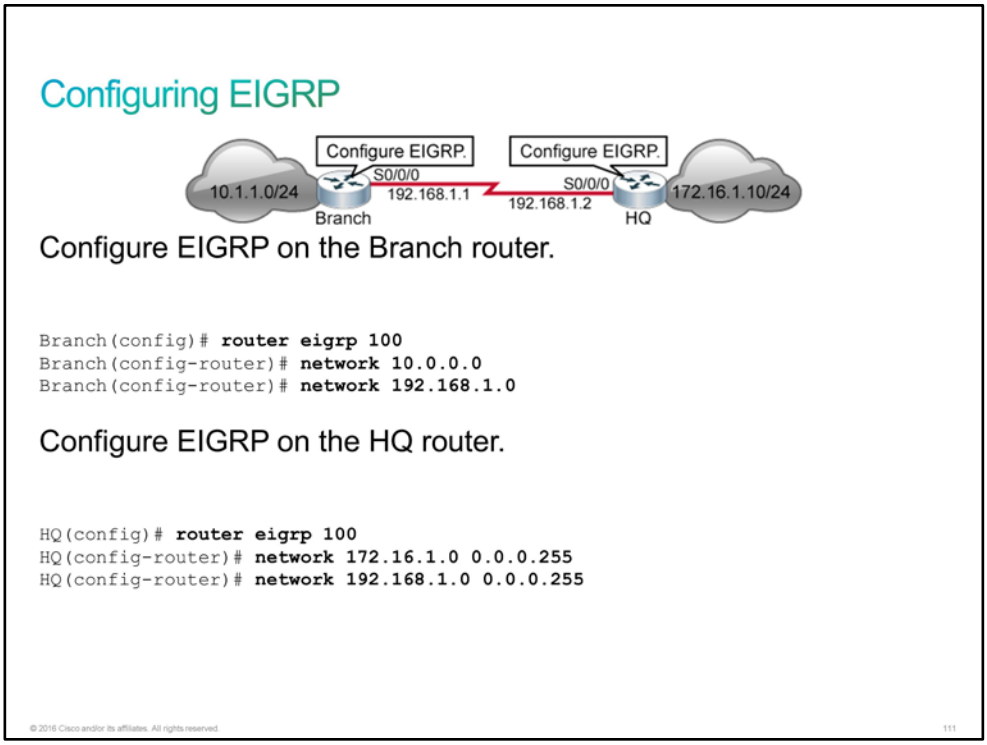

Command	Description
router eigrp *as_number*	Enables the EIGRP routing process for the AS that is specified.
network *network_number* [*wildcard_mask*]	Associates the network with the EIGRP routing process. Use of the wildcard mask is optional.
no shutdown	EIGRP has a shutdown feature. The routing process should be in the **no shutdown** mode in order to start running. The default behavior is different between Cisco IOS Software versions.

The **router eigrp** global configuration command enables EIGRP.

Use the **router eigrp** and **network** commands to create an EIGRP routing process. Note that EIGRP requires an AS number. The AS parameter is a number between 1 and 65,535 that is chosen by the network administrator.

The **network** command is used in the router configuration mode.

Note	The AS number that EIGRP refers to in the parameter can be assigned any 16-bit value. As opposed to OSPF, the AS number in EIGRP must match on all routers that are involved in the same EIGRP process.

The **network** command in EIGRP has the same function as in other IGP routing protocols:

- The **network** command defines a major network number to which the router is directly connected. Any interface on this router that matches the network address in the **network** command will be enabled to send and receive EIGRP updates. The EIGRP routing process looks for interfaces that have an IP address that belongs to the networks that are specified with the **network** command. The EIGRP process begins on these interfaces.
- This network (or subnet) will be included in EIGRP routing updates.

To configure EIGRP to advertise specific subnets only, use the *wildcard-mask* option with the **network** command. For example, for subnet 255.255.255.0 the wildcard mask will be 0.0.0.255.

Note	You can also use the subnet mask with EIGRP, however the IOS will automatically correct it to be the wildcard mask.

Step 3 Access the console of R1. Enable EIGRP AS number 1 and include the network 10.0.0.0/8 on R1.

Enter the following commands to the R1 router:

```
R1# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
R1(config)# router eigrp 1
R1(config-router)# network 10.0.0.0
*Oct  6 08:14:41.002: %DUAL-5-NBRCHANGE: EIGRP-IPv4 1: Neighbor 10.10.1.2
(Ethernet0/0) is up: new adjacency
R1(config-router)# end
R1#
```

Since 10.0.0.0 is the base address of the full Class A network 10.0.0.0/8, you do not need to include a subnet mask in the network statement.

The EIGRP neighbor relationship with R2 was established immediately after entering the network statement that included the IP address of R1's Ethernet0/0 interface.

Verifying EIGRP Neighbors

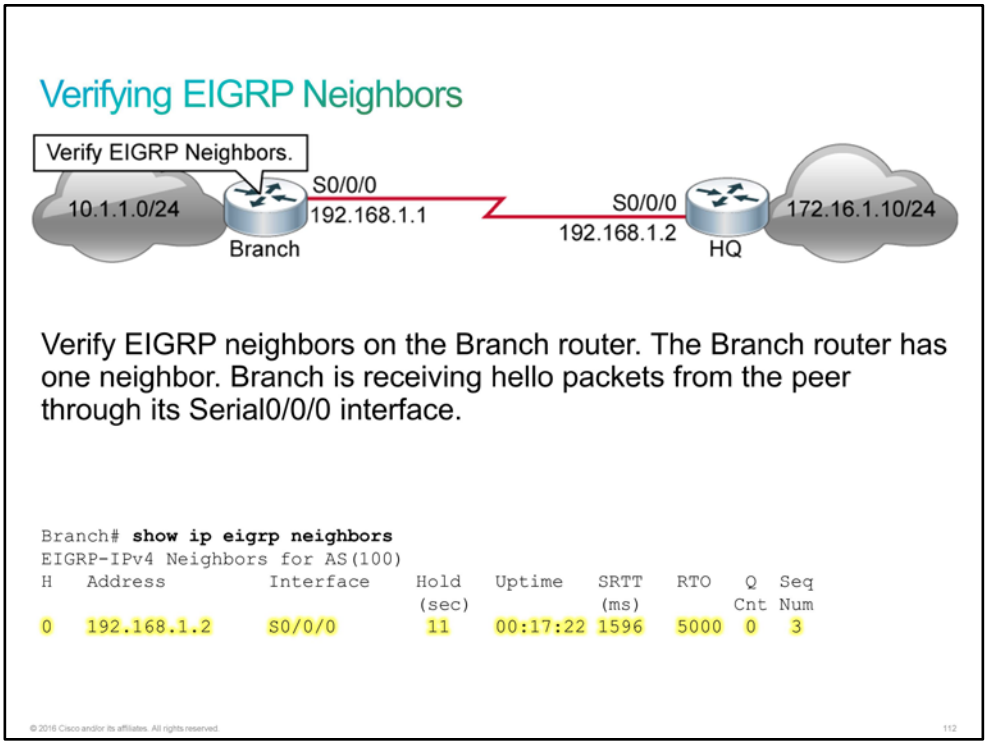

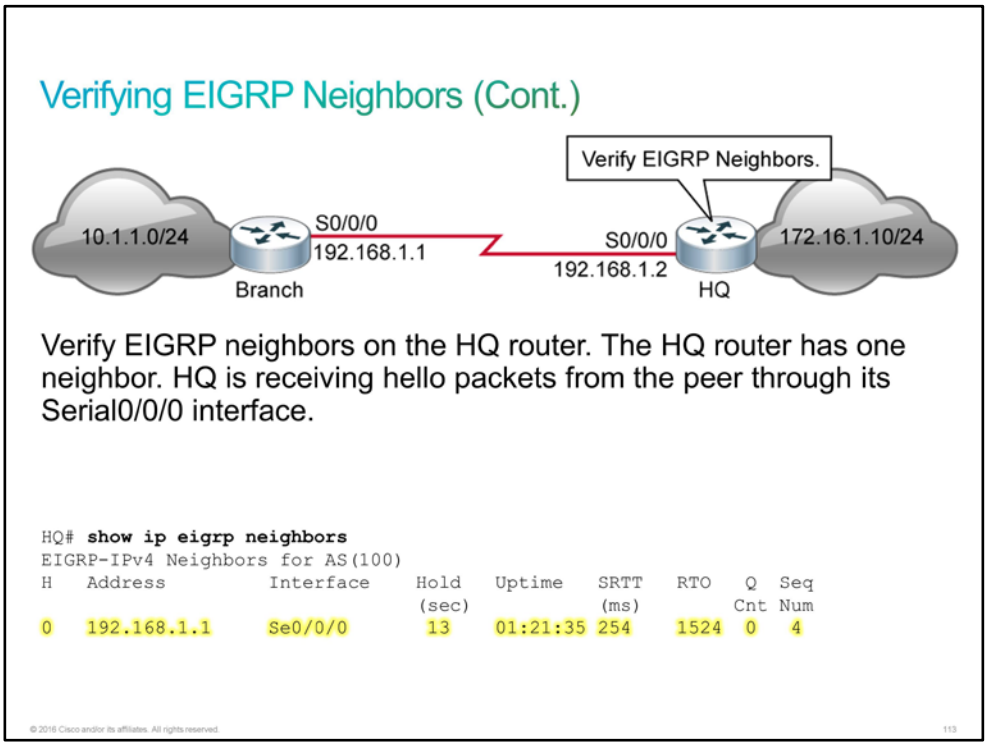

Use the **show ip eigrp neighbors** command to display the neighbors that EIGRP discovered and to determine when neighbors become active and inactive. The command is also useful for debugging transport problems.

Field	Description
AS(100)	Process number that is specified with the **router** command
Address	IP address of the EIGRP peer
Interface	Interface on which the router is receiving hello packets from the peer
Hold (sec)	Length of time (in seconds) that Cisco IOS Software waits to hear from the peer before declaring it down. If the peer is using the default hold time, this number is less than 15. If the peer configures a nondefault hold time, the nondefault hold time is displayed.
Uptime	Elapsed time (in the *hours:minutes:seconds* format) since the local router first heard from this neighbor
Q Cnt	Number of EIGRP packets (update, query, and reply) that the software is waiting to send
Seq Num	Sequence number of the last update, query, or reply packet that was received from this neighbor

Step 4 Display the EIGRP neighbor table on R1.

R2 is an EIGRP neighbor of R1.

```
R1# show ip eigrp neighbors
EIGRP-IPv4 Neighbors for AS(1)
H   Address                 Interface              Hold Uptime   SRTT   RTO  Q
Seq
                                                   (sec)         (ms)       Cnt
Num
0   10.10.1.2               Et0/0                    13 00:01:04 1599  5000  0
7
```

Verifying EIGRP Interfaces

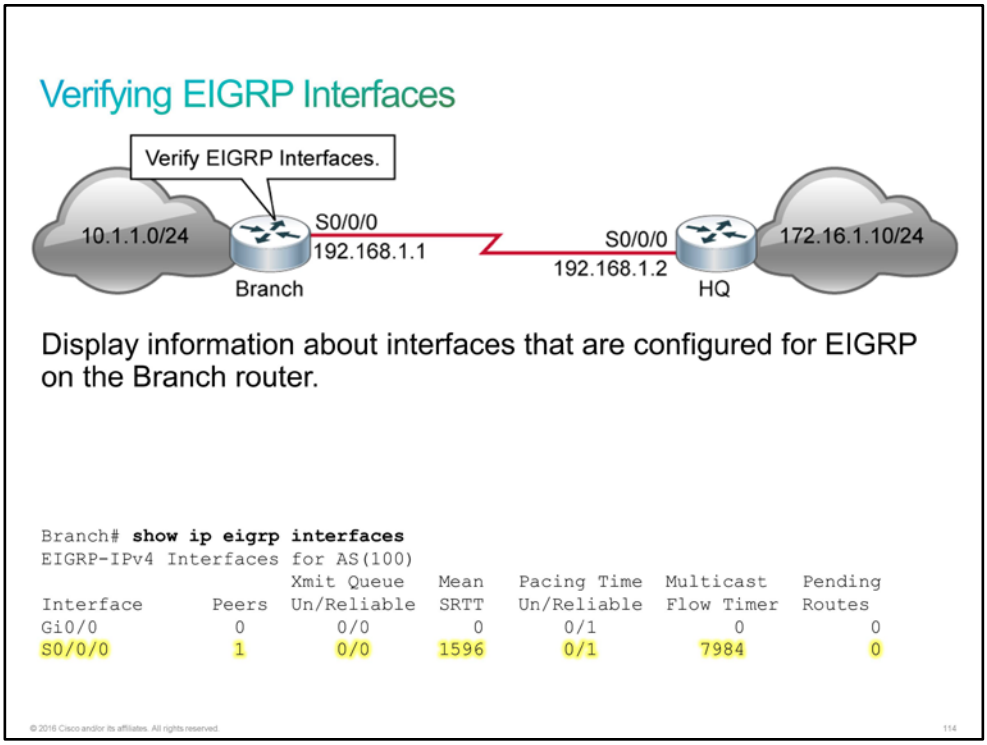

Use the **show ip eigrp interfaces** command to determine on which interfaces EIGRP is active and to learn information about EIGRP that relates to those interfaces. If you specify an interface (for example, **show ip eigrp interfaces GigabitEthernet0/0**), only this interface is displayed. Otherwise, all interfaces on which EIGRP is running are displayed. If you specify AS (for example, **show ip eigrp 100 interfaces**), only the routing process for the specified AS is displayed. Otherwise, all EIGRP processes are displayed.

Field	Description
Interface	Interface over which EIGRP is configured
Peers	Number of directly connected EIGRP neighbors on the interface
Xmit Queue Unreliable/Reliable	Number of packets remaining in the Unreliable and Reliable queues
Mean SRTT	Average SRTT interval (in milliseconds) for all neighbors on the interface
Pacing Time Unreliable/Reliable	Number of milliseconds to wait after transmitting unreliable and reliable packets
Multicast Flow Timer	Number of milliseconds to wait for acknowledgment of a multicast packet by all neighbors before transmitting the next multicast packet
Pending Routes	Number of routes in the packets in the transmit queue waiting to be sent

Step 5 Display the interfaces on R1 that are participating in EIGRP.

Both Ethernet0/0 and Loopback0 are participating in EIGRP.


```
R1# show ip eigrp interfaces
EIGRP-IPv4 Interfaces for AS(1)
                              Xmit Queue   PeerQ        Mean   Pacing Time
Multicast    Pending
Interface             Peers  Un/Reliable  Un/Reliable  SRTT   Un/Reliable
Flow Timer   Routes
Et0/0                   1        0/0          0/0        1599       0/2
7992           0
Lo0                     0        0/0          0/0           0       0/0
0              0
```

Verifying EIGRP Routes

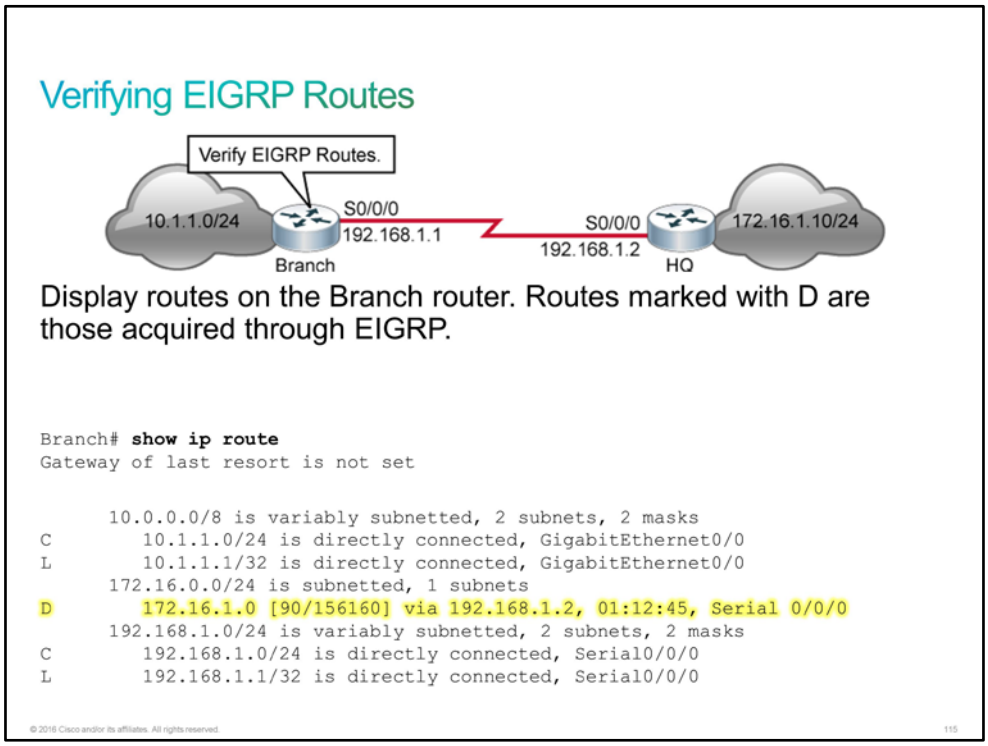

The **show ip route** command displays the current entries in the routing table. EIGRP has a default administrative distance of 90 for internal routes and 170 for routes that are imported from an external source, such as default routes. When compared to other IGPs, EIGRP is preferred by Cisco IOS Software because it has the lowest administrative distance.

Step 6 Display the routing table on R1.

"D" indicates a route that was provided by EIGRP.

```
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 not set

      10.0.0.0/8 is variably subnetted, 7 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
D        10.10.2.0/24 [90/307200] via 10.10.1.2, 00:01:25, Ethernet0/0
C        10.10.11.0/24 is directly connected, Loopback0
L        10.10.11.1/32 is directly connected, Loopback0
D        10.10.12.0/24 [90/409600] via 10.10.1.2, 00:01:25, Ethernet0/0
D        10.10.13.0/24 [90/435200] via 10.10.1.2, 00:01:25, Ethernet0/0
```

R1 has learned about the network between R2 and R3 as well as the networks of the loopback interfaces on both R2 and R3.

Verifying EIGRP Topology

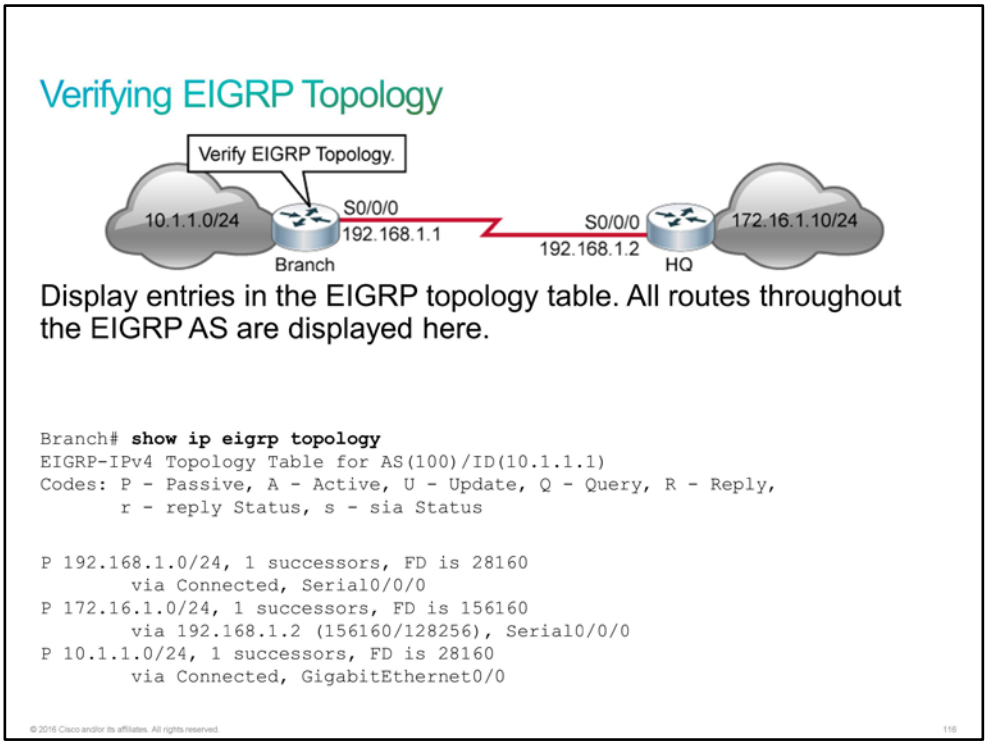

The **show ip eigrp topology** command displays the EIGRP topology table, the active or passive state of routes, the number of successors, and the FD to the destination. Use the **show ip eigrp topology all-links** command to display all paths, even the ones that are not feasible.

Field	Description
Codes	The state of this topology table entry. Passive and active refer to the EIGRP state regarding this destination; update, query, and reply refer to the type of packet that is being sent.
P – passive	Indicates that no EIGRP computations are being performed for this destination.
A – active	Indicates that EIGRP computations are being performed for this destination.
U – update	Indicates that an update packet was sent to this destination.
Q – query	Indicates that a query packet was sent to this destination.
R – reply	Indicates that a reply packet was sent to this destination.
R – reply status	A flag that is set after the software has sent a query and is waiting for a reply.
172.16.1.0	Destination IP network number.
/24	Destination subnet mask.
Successors	Number of successors. This number corresponds to the number of next hops in the IP routing table. If "successors" is capitalized, then the route or next hop is in a transition state.
FD	The FD is the best metric to reach the destination or the best metric that was known when the route went active. This value is used in the feasibility condition check. If the reported distance of the router (the metric after the slash) is less than the FD, the feasibility condition is met and this path is a feasible successor. After the software determines that it has a feasible successor, it does not need to send a query for this destination.
Replies	The number of replies that are still outstanding (have not been received) regarding this destination. This information appears only when the destination is in active state.
State	The exact EIGRP state that this destination is in. It can be 0, 1, 2, or 3. This information appears only when the destination is in the active state.
Via	The IP address of the peer that told the software about this destination. The first *n* of these entries, where *n* is the number of successors, are the current successors. The remaining entries on the list are feasible successors.
(156160/128256)	The first number is the EIGRP metric that represents the cost, or FD, to the destination. The second number is the EIGRP metric that this peer advertised.
Serial0/0/0	The interface from which this information was learned.

You can also see the router ID in the output. Each router in an EIGRP routing domain is identified by its router ID. It is used by a router each time that it communicates with its EIGRP neighbors. The EIGRP router ID is also used for validating the origin of external routes. If an external route is received with a local router ID, the route is discarded.

You can set the router ID manually using the **eigrp router-id** *router-id* command. The router ID can be configured with any IP address except 0.0.0.0 and 255.255.255.255. A unique value should be configured for each router. If the router ID is not explicitly configured, the router will select the highest address of its loopback interfaces. If there is no loopback interface on the router, it will select the highest IP address of any other active local interface. The router ID is not changed unless the EIGRP process is cleared, or if the router ID is manually configured

Step 7 Display the EIGRP topology database on R1.

You will see five networks in the topology and the router ID for the EIGRP process.

```
R1# show ip eigrp topology
EIGRP-IPv4 Topology Table for AS(1)/ID(10.10.11.1)
Codes: P - Passive, A - Active, U - Update, Q - Query, R - Reply,
       r - reply Status, s - sia Status

P 10.10.12.0/24, 1 successors, FD is 409600
        via 10.10.1.2 (409600/128256), Ethernet0/0
P 10.10.11.0/24, 1 successors, FD is 128256
        via Connected, Loopback0
P 10.10.13.0/24, 1 successors, FD is 435200
        via 10.10.1.2 (435200/409600), Ethernet0/0
P 10.10.2.0/24, 1 successors, FD is 307200
        via 10.10.1.2 (307200/281600), Ethernet0/0
P 10.10.1.0/24, 1 successors, FD is 281600
        via Connected, Ethernet0/0
```

There are five networks in the virtual lab topology: the networks associated with the Loopback interface on each router, the network between R1 and R2, and the network between R2 and R3. All five of these networks will be represented in the EIGRP topology database on each of the three routers.

Note that the current router ID is 10.10.11.1, which is the IP address of the Loopback0 interface.

Step 8 Access the console of R1. Change the EIGRP router ID to 11.11.11.11.

Enter the following commands to the R1 router:

```
R1# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
R1(config)# router eigrp 1
R1(config-router)# eigrp router-id 11.11.11.11
R1(config-router)# end
R1#
```

Step 9 Display the EIGRP topology database on R1 to verify the EIGRP router ID.

The EIGRP router ID is set to 11.11.11.11.


```
R1# show ip eigrp topology
EIGRP-IPv4 Topology Table for AS(1)/ID(11.11.11.11)
Codes: P - Passive, A - Active, U - Update, Q - Query, R - Reply,
       r - reply Status, s - sia Status

P 10.10.12.0/24, 1 successors, FD is 409600
        via 10.10.1.2 (409600/128256), Ethernet0/0
P 10.10.11.0/24, 1 successors, FD is 128256
        via Connected, Loopback0
P 10.10.13.0/24, 1 successors, FD is 435200
        via 10.10.1.2 (435200/409600), Ethernet0/0
P 10.10.2.0/24, 1 successors, FD is 307200
        via 10.10.1.2 (307200/281600), Ethernet0/0
P 10.10.1.0/24, 1 successors, FD is 281600
        via Connected, Ethernet0/0
```

Using Passive Interfaces

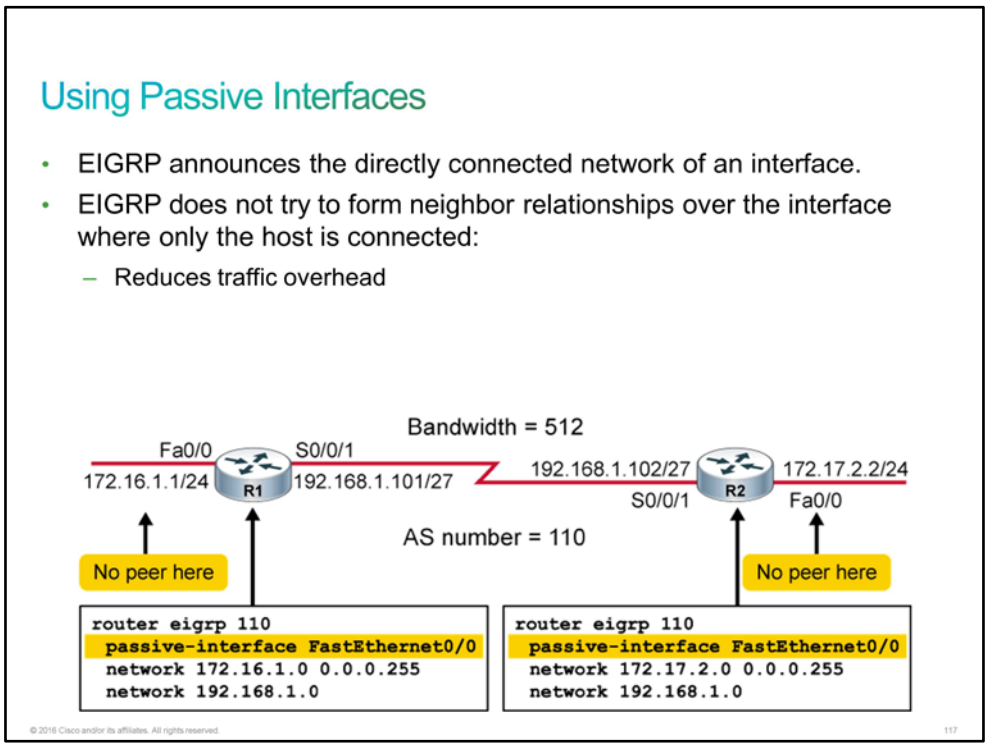

R1 and R2 have no neighbors that are available over the FastEthernet0/0 interface; therefore, there is no need to try to establish adjacency over the interfaces. Moreover, the packets that are sent are overhead to the link bandwidth and also consume CPU resources of the router. To stop sending hello packets over the interface without neighbors, use the **passive-interface** command on the specified interface. In the example, the **passive-interface** command is used in both routers for the FastEthernet0/0 interface. EIGRP will not bring up adjacencies on a passive interface.

Note	Configuring the **passive-interface** command suppresses all incoming and outgoing routing updates and hello messages.

The **passive-interface** command has the following properties:

- Prevents a neighbor relationship from being established over the passive interface
- Stops routing updates from being received or sent over the passive interface
- Allows a subnet on the passive interface to be announced in an EIGRP process

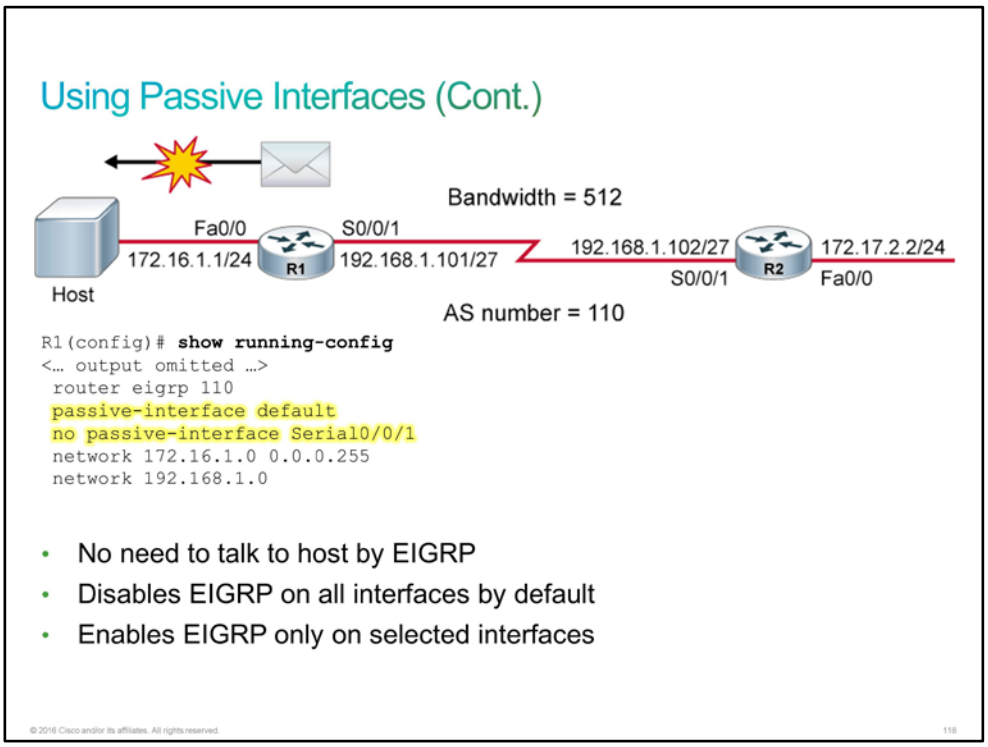

Within ISPs and large enterprise networks, distribution routers may have more than 100 interfaces, so manual configuration of the **passive-interface** command on interfaces where adjacency is not desired may create a problem. So, in some networks, you would need to enter 100 or more passive interface statements.

With the default passive interface feature, this issue is solved by allowing all interfaces to be set as passive by default using a single **passive-interface default** command. Where adjacencies are desired, the individual interfaces are configured using the **no passive-interface** command.

In the figure, R1 and R2 are configured with the **passive-interface default** command, and all interfaces are refusing the establishment of EIGRP adjacency by default. The Serial0/0/1 interface on each router is then configured to allow EIGRP adjacency, because neighbors are expected. The **passive-interface** command is disabled for these interfaces.

Step 10 Access the console of R1. Configure interface Loopback0 as EIGRP passive interface.

Enter the following commands to the R1 router:

```
R1# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
R1(config)# router eigrp 1
R1(config-router)# passive-interface Loopback0
R1(config-router)# end
R1#
```

Usually, you would configure physical interfaces that are connected to end devices as passive.

Verifying Operation with Passive Interfaces

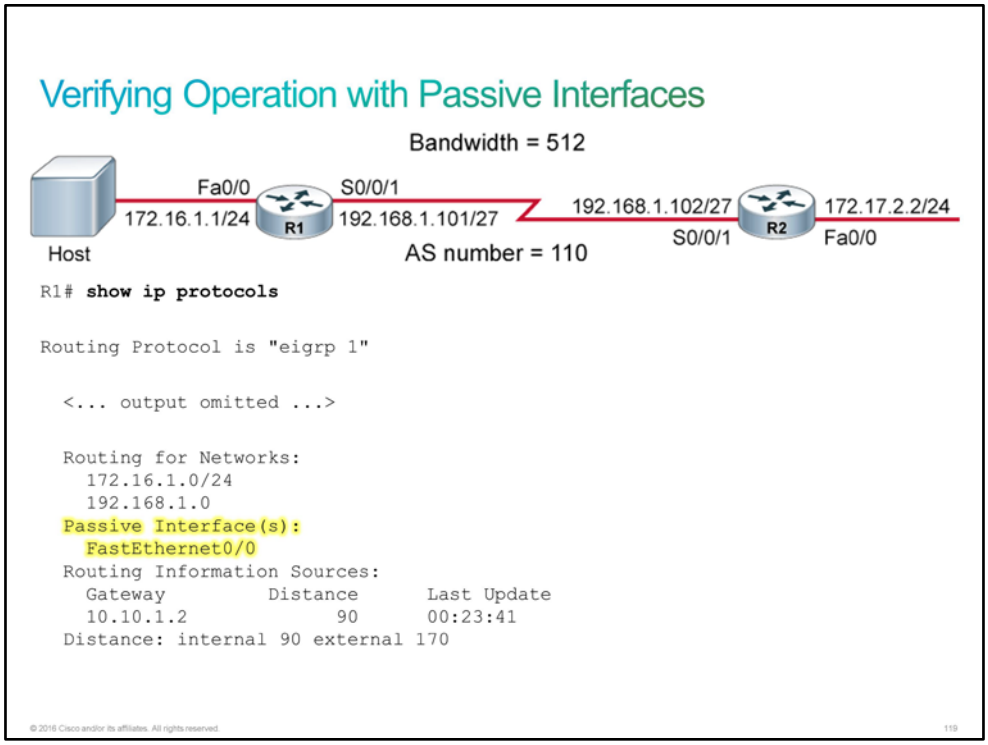

The most important questions to ask when verifying operation with passive interfaces are as follows:

- Do you see all the neighbors?
- Which interfaces in the routing process are passive?

To see all the available EIGRP neighbors, use the **show ip eigrp neighbors** command.

To see the passive interfaces in the routing protocol, use **show ip protocols** command. In the figure, the command output for R1 shows that the FastEthernet0/0 interface is defined as a passive interface.

Step 11 Use the **show ip protocols** command to verify which interfaces are configured as passive.

You should see the Loopback0 interface on the list of passive interfaces.

```
R1# show ip protocols
*** IP Routing is NSF aware ***

Routing Protocol is "eigrp 1"
  Outgoing update filter list for all interfaces is not set
  Incoming update filter list for all interfaces is not set
  Default networks flagged in outgoing updates
  Default networks accepted from incoming updates
  EIGRP-IPv4 Protocol for AS(1)
    Metric weight K1=1, K2=0, K3=1, K4=0, K5=0
    NSF-aware route hold timer is 240
    Router-ID: 11.11.11.11
    Topology : 0 (base)
      Active Timer: 3 min
      Distance: internal 90 external 170
      Maximum path: 4
      Maximum hopcount 100
      Maximum metric variance 1

  Automatic Summarization: disabled
  Maximum path: 4
  Routing for Networks:
    10.0.0.0
  Passive Interface(s):
    Loopback0
  Routing Information Sources:
    Gateway         Distance      Last Update
    10.10.1.2             90      00:23:41
  Distance: internal 90 external 170
```

You will also see that EIGRP routing is enabled for 10.0.0.0 network. Also note that only metric values K1 and K3 are enabled by default, meaning that metric value is calculated based on bandwidth and delay only.

This is the end of the discovery lab.

EIGRP Load Balancing

In general, load balancing is the capability of a router to distribute traffic over all the router network ports that are within the same distance from the destination address. Load balancing increases the utilization of network segments, and this way increases effective network bandwidth. EIGRP supports both equal and unequal cost path load balancing.

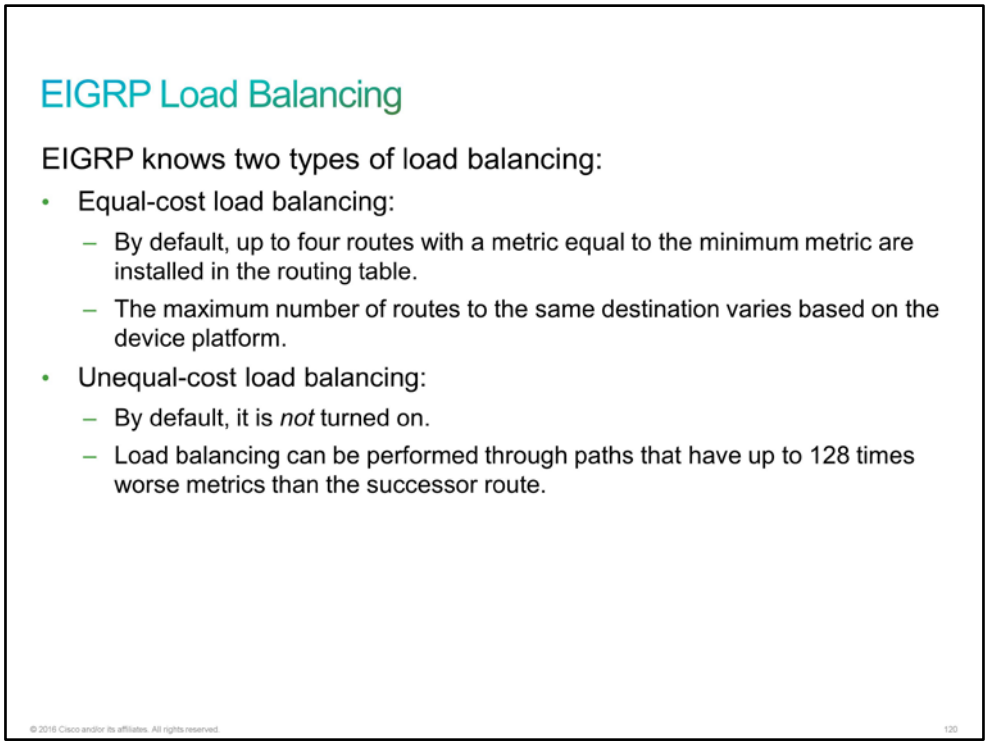

Equal-Cost Load Balancing

Given that good network design involves Layer 3 path redundancy, it is a common customer expectation that if there are multiple devices and paths to a destination, all paths should be utilized. Mostly, the paths to a destination have equal costs. In the figure, networks A and B are connected with two equal-cost paths. For this example, assume that the links are GigabitEthernet.

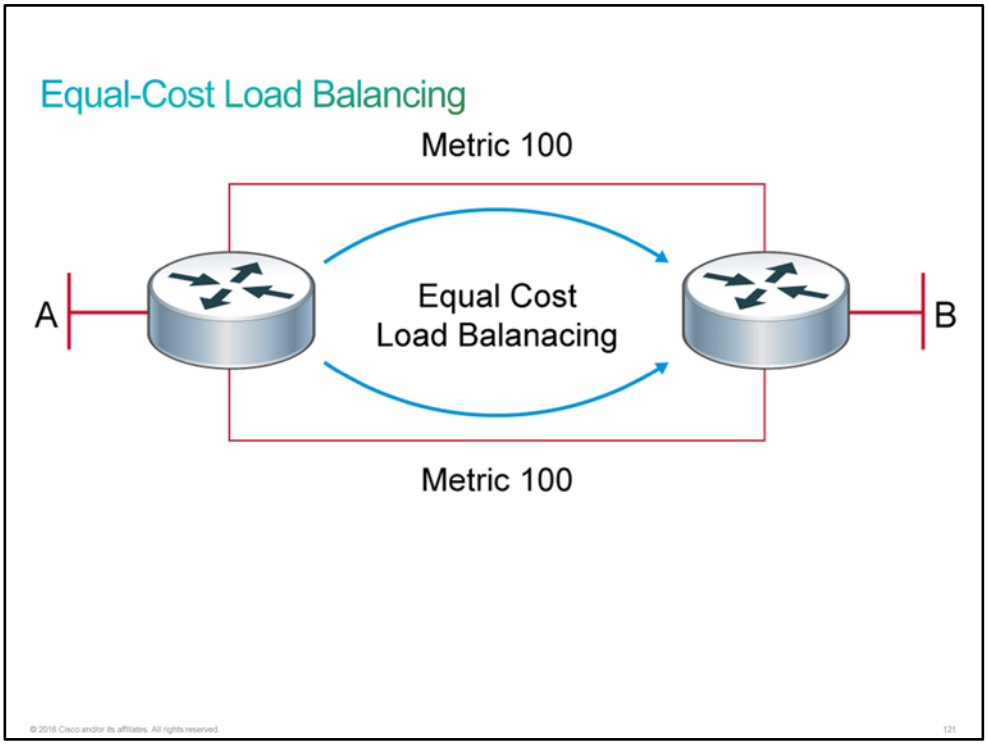

Equal-cost load balancing is the ability of a router to distribute traffic over all its network ports that are the same metric from the destination address. Load balancing increases the use of network segments and increases effective network bandwidth.

By default, Cisco IOS Software applies load balancing across up to four equal-cost paths for a certain destination IP network, if such paths exist. With the **maximum-paths** router configuration command, you can specify the number of routes that can be kept in the routing table. If you set the value to 1, you disable load balancing.

Note	The actual number of maximum-paths that can be configured varies from device to device.

The **maximum-paths** command is entered in routing protocol configuration mode. In the example, this Cisco router supports up to 16 paths.

Note	If you adjust the **maximum-paths** value, it must be the same on both sides of the path.

```
HQ(config)#router eigrp 100
HQ(config-router)#maximum-paths ?
  <1-16>  Number of paths
```


Unequal-Cost Load Balancing

EIGRP can also balance traffic across multiple routes that have different metrics. This type of balancing is called unequal-cost load balancing. In the figure, you are presented with a cost difference of almost 4:1. A real-network example of such situation is the case of a WAN connection from HQ to a branch. The primary WAN link is a 6 Mb/s MPLS link with a T1 (1.544 Mb/s) backup link.

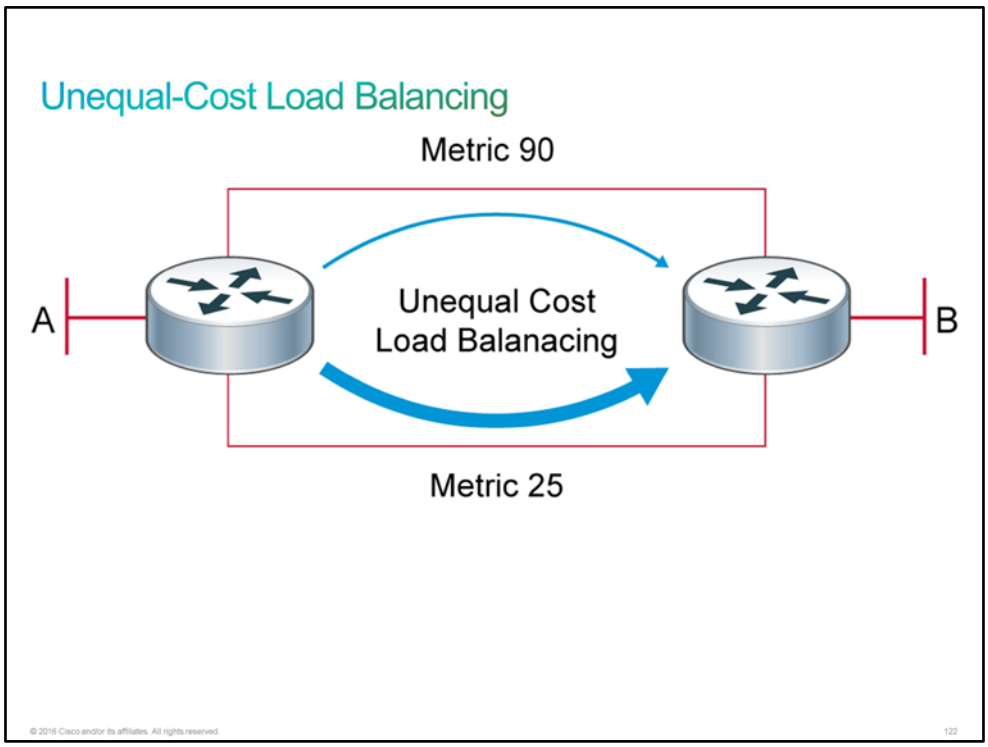

The default variance is equal to 1. EIGRP will normally only install additional routes (paths) to a destination when there is zero variance in cost. You can use the **variance** command to tell EIGRP to install routes in the routing table, as long as they are less than the current cost multiplied by the variance value. In the example in the figure, setting the variance to 4 would allow EIGRP to install the backup path and send traffic over it. The backup path is now performing work instead of just idling.

```
HQ(config-router)# router eigrp 100
HQ(config-router)# variance ?
  <1-128>  Metric variance multiplier

HQ(config-router)# variance 4
HQ(config-router)#
```

Challenge

1. Which two of the following are classified as link-state routing protocols? (Choose two.)

   A. IS-IS
   B. OSPF
   C. EIGRP
   D. RIPv2
   E. BGP

2. A router has learned three possible routes that could be used to reach a destination network. One route is from EIGRP and has a composite metric of 20584570. Another route is from OSPF with a metric of 842. The last is from RIPv2 and has a metric of 3. Which route or routes will the router install in the routing table?

   A. EIGRP route
   B. OSPF route
   C. RIPv2
   D. All three routes.
   E. None of the above.

3. Refer to the figure and the output of the **show ip protocols** command. Which two EIGRP metrics are being used to affect the calculation that selects the best path to add to the EIGRP routing table? (Choose two.)

```
R1# show ip protocols
*** IP Routing is NSF aware ***

Routing Protocol is "eigrp 65010"
  Outgoing update filter list for all interfaces is not set
  Incoming update filter list for all interfaces is not set
  Default networks flagged in outgoing updates
  Default networks accepted from incoming updates
  EIGRP-IPv4 Protocol for AS(65010)
    Metric weight K1=1, K2=0, K3=1, K4=0, K5=0
    NSF-aware route hold timer is 240
    Router-ID: 172.16.1.1
    Topology : 0 (base)
      Active Timer: 3 min
      Distance: internal 90 external 170
      Maximum path: 4
      Maximum hopcount 100
      Maximum metric variance 1
  Automatic Summarization: disabled
  Maximum path: 4
  Routing for Networks:
  Routing Information Sources:
    Gateway         Distance      Last Update
  Distance: internal 90 external 170
```

A. Bandwidth
B. Delay
C. load
D. Reliability
E. MTU

4. Which command should you use to determine whether the EIGRP router ID has been configured in the EIGRP process? (Choose two)

A. **show ip eigrp neighbor**
B. **show ip eigrp interface**
C. **show ip protocols**
D. **show ip eigrp topology**

5. What does **passive interface** command do in EIGRP?

A. Router cannot form neighbor adjacencies on that interface
B. Router cannot send routing updates on that interface.
C. Router cannot receive routing updates on that interface.
D. All the above.

6. Here is the **show ip route** command from router R1. R1 is load- balancing to 10.80.13.0/30 network.

R1>show ip route
Codes: 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
*** - candidate default, U - per-user static route**
o - ODR, P - periodic downloaded static route
Gateway of last resort is not set

10.80.0.0/8 is variably subnetted, 11 subnets, 4 masks
D 10.80.13.0/30 [90/2681856] via 10.80.234.3, 01:13:17, Serial0/0
[90/2681856] via 10.80.234.2, 01:13:17, Serial1/0
D 10.80.23.2/32 [90/2681856] via 10.80.234.2, 01:13:17, Serial0/0

<output omitted>

On R1, following change was made. How will it affect the EIGRP routing table and topology table ? (Choose two)

R1(config)# router eigrp 10
R1(config-router)# maximum-paths 1

A. There will be no change in the routing table for the route 10.80.13.0/30.
B. EIGRP topology table will have both the routes to the network 10.80.13.0/30
C. Only one route will be there is the routing table for 10.80.13.0/30.
D. EIGRP topology table will have only one route for 10.80.13.0/30.

7. Refer to the output of the show ip protocols command. According to the output, this router is configured to load-balance over unequal-cost paths. True or false ?

```
R1# show ip protocols
*** IP Routing is NSF aware ***

Routing Protocol is "eigrp 65010"
  Outgoing update filter list for all interfaces is not set
  Incoming update filter list for all interfaces is not set
  Default networks flagged in outgoing updates
  Default networks accepted from incoming updates
  EIGRP-IPv4 Protocol for AS(65010)
    Metric weight K1=1, K2=0, K3=1, K4=0, K5=0
    NSF-aware route hold timer is 240
    Router-ID: 172.16.1.1
    Topology : 0 (base)
      Active Timer: 3 min
      Distance: internal 90 external 170
      Maximum path: 4
      Maximum hopcount 100
      Maximum metric variance 1

  Automatic Summarization: disabled
  Maximum path: 4
  Routing for Networks:
  Routing Information Sources:
    Gateway         Distance      Last Update
  Distance: internal 90 external 170
```

A. True
B. False

Answer Key

Challenge

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

Lesson 2: Implementing EIGRP for IPv6

Introduction

You are about to go to a customer site to implement EIGRP for IPv6. You should know the operational theory behind EIGRP for IPv6 before you go onsite so that you can answer any customer inquiries.

Decide if you are ready to go onsite now to upgrade the network or if you first need to research how to implement EIGRP for IPv6.

EIGRP for IPv6

Although the configuration and management of EIGRP for IPv4 and EIGRP for IPv6 are similar, they are configured and managed separately.

EIGRP is inherently a multiprotocol routing protocol because it has supported non-IP IPX and AppleTalk for some time. IPv6 support is added as a separate module. IPv6 EIGRP is configured and managed separately from IPv4 EIGRP, but the mechanisms and configuration techniques will be familiar to people who are skilled with EIGRP for IPv4.

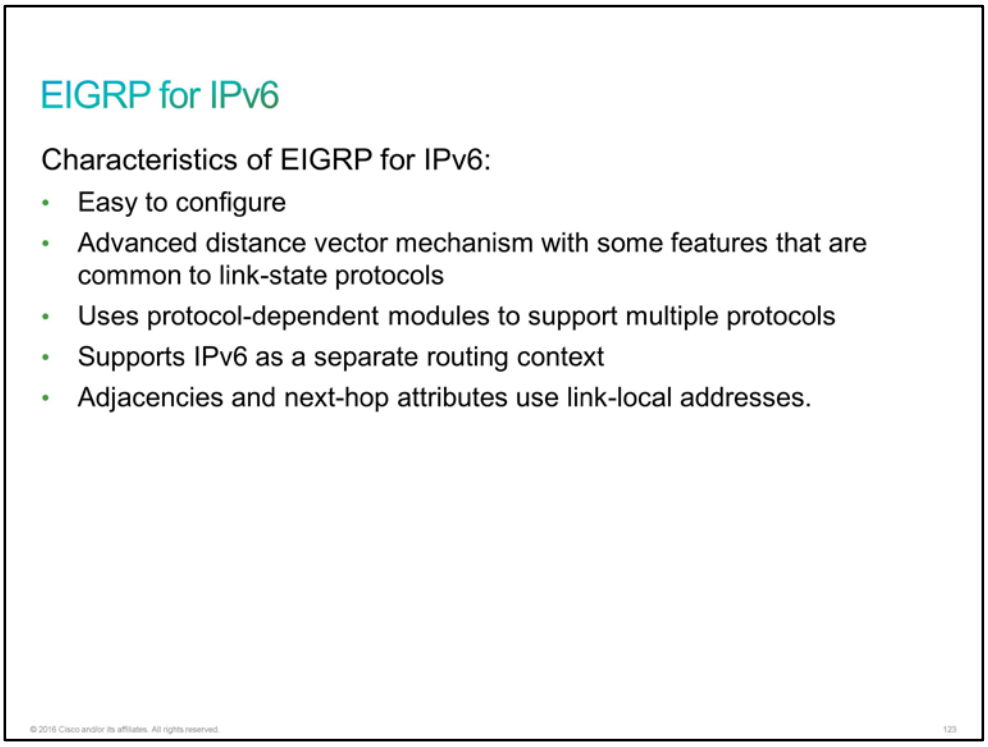

For example, both the IPv4 and IPv6 EIGRP implementations include a shutdown feature that allows the routing protocol to be configured but also easily disabled. Both use the DUAL to optimize the routing path. Both are scalable to large networks. There are also a few differences in the IPv4 and IPv6 features. For example, in contrast with IPv4 EIGRP, IPv6 EIGRP is configured over a link—there is no network statement as there is for IPv4. Also EIGRP for IPv6 adjacencies use link-local addresses to communicate and router next-hop attributes are neighboring router link-local addresses.

The basic components of EIGRP for IPv6 remain the same as the IPv4 version.

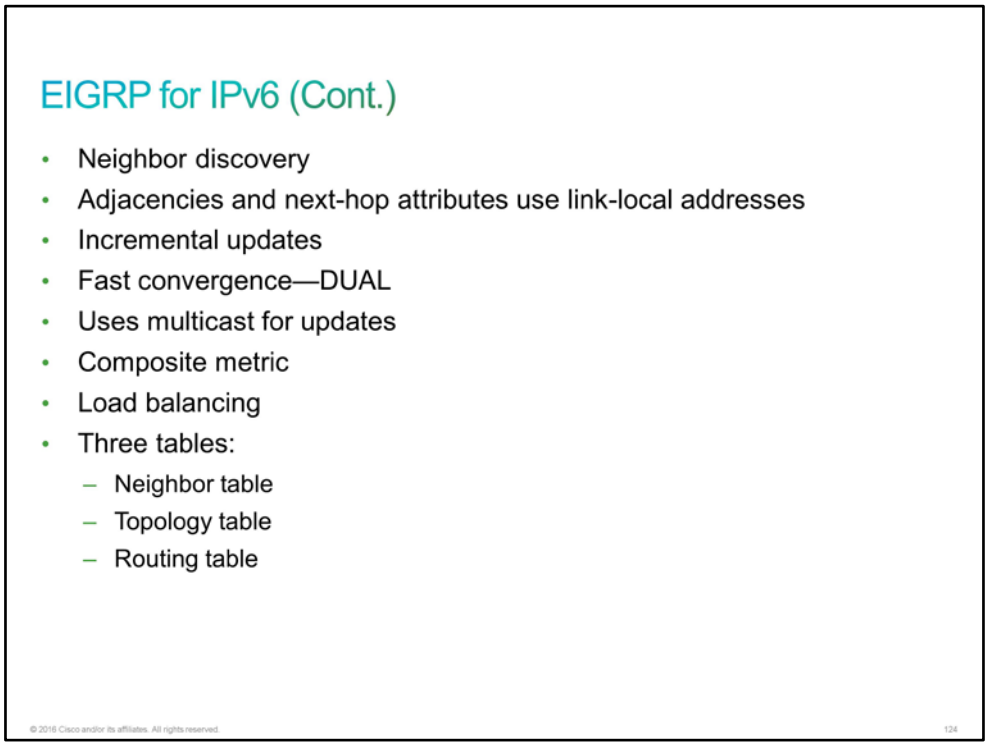

EIGRP uses a small hello packet to discover other EIGRP-capable routers on directly attached links and forms durable neighbor relationships. Updates may be acknowledged by using a reliable transport protocol, or they may be unacknowledged—depending on the specific function that is being communicated. The protocol provides the flexibility that is needed to unicast or multicast updates, whether acknowledged or unacknowledged.

Hello packets and updates are set to the well-known, link-local multicast address FF02::A, which Cisco obtained from the IANA. This multicast distribution technique is more efficient than the broadcast mechanism that is used by earlier, more primitive routing protocols, such as RIPv1. EIGRP for IPv4 also uses multicast for update distribution.

EIGRP sends incremental updates when the state of a destination changes, instead of sending the entire contents of the routing table. This feature minimizes the bandwidth that is required for EIGRP packets.

DUAL, which is an EIGRP algorithm for determining the best path through the network, uses several metrics to select efficient, loop-free paths. When multiple routes to a neighbor exist, DUAL determines which route has the lowest metric (named the FD) and enters this route into the routing table. Other possible routes to this neighbor with larger metrics are received, and DUAL determines the reported distance to this network. The reported distance is defined as the total metric that is advertised by an upstream neighbor for a path to a destination. DUAL compares the reported distance with the FD, and if the reported distance is less than the FD, DUAL considers the route to be a feasible successor and enters the route into the topology table. The feasible successor route that is reported with the lowest metric becomes the successor route to the current route if the current route fails. To avoid routing loops, DUAL ensures that the reported distance is always less than the FD for a neighbor router to reach the destination network; otherwise, the route to the neighbor may loop back through the local router.

When there are no feasible successors to a route that has failed, but there are neighbors advertising the route, a recomputation must occur. This process is where DUAL determines a new successor. The amount of time that is required to recompute the route affects the convergence time. Recomputation is processor-intensive. It is advantageous to avoid unneeded recomputation. When a topology change occurs, DUAL will test for feasible successors. If there are feasible successors, DUAL will use them in order to avoid unnecessary recomputation.

EIGRP updates contain five metrics: minimum bandwidth, delay, load, reliability, and MTU. Of these five metrics, by default, only minimum bandwidth and delay are used to compute the best path. Unlike most metrics, minimum bandwidth is set to the minimum bandwidth of the entire path, and it does not reflect how many hops or low-bandwidth links are in the path. Delay is a cumulative value that increases by the delay value of each segment in the path.

EIGRP for IPv6, like EIGRP for IPv4, is able to do load balancing. Load balancing is the capability of a router to distribute traffic over all the router network ports that are within the same distance from the destination address. Load balancing increases the utilization of network segments and this way increases effective network bandwidth. There are two types of load balancing:

- **Equal-cost path:** Applicable when different paths to a destination network report the same routing metric value.
- **Unequal-cost path:** Applicable when different paths to a destination network report different routing metric values.

When a router discovers a new neighbor, it records the neighbor address and interface as an entry in the *neighbor table*. One neighbor table exists for each protocol-dependent module. When a neighbor sends a hello packet, it advertises a hold time, which is the time that a router treats a neighbor as reachable and operational. If a hello packet is not received within the hold time, the hold time expires and DUAL is informed of the topology change.

The *topology table* contains all destinations that are advertised by the neighboring routers. Each entry in the topology table includes the destination address and a list of neighbors that have advertised the destination. For each neighbor, the entry records the advertised metric, which the neighbor stores in its routing table. An important rule that distance vector protocols must follow is that if the neighbor advertises this destination, the neighbor must use the route to forward packets.

Discovery 41: Configure and Verify EIGRP for IPv6

Introduction

This discovery will guide you through the configuration and verification of EIGRP for IPv6 on an IOS router. The 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. Both IPv4 and IPv6 are configured in this dual-stack environment. R2 and R3 are also configured with EIGRP for IPv6 using the autonomous system number 100. In this discovery, you will configure EIGRP for IPv6 on R1 and verify the results.

Topology

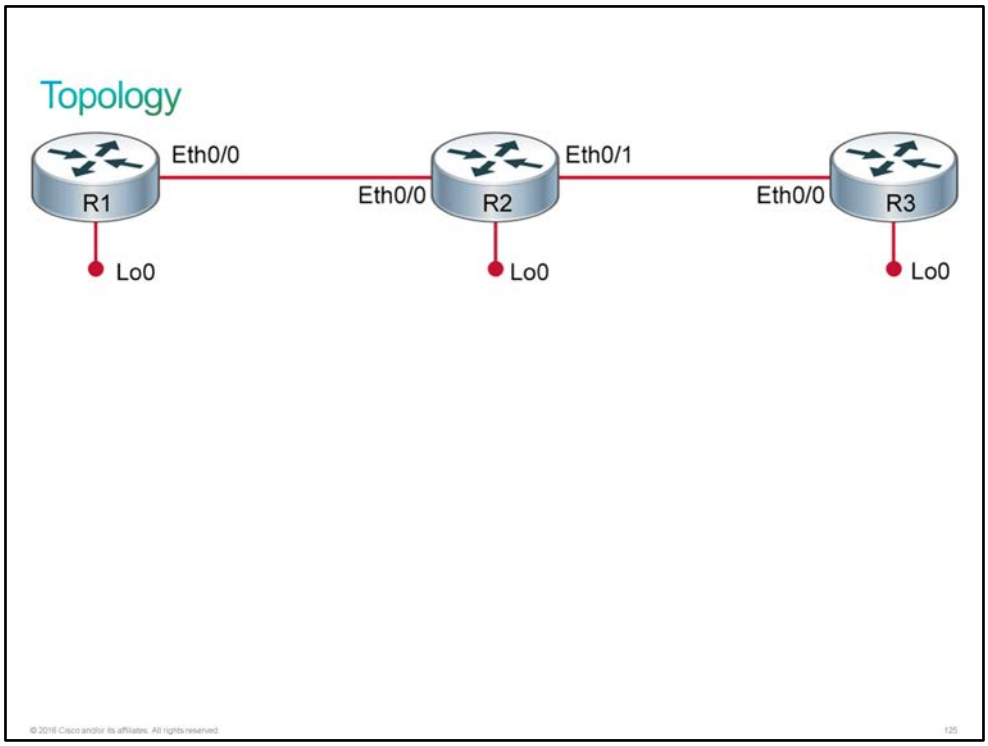

Job Aid

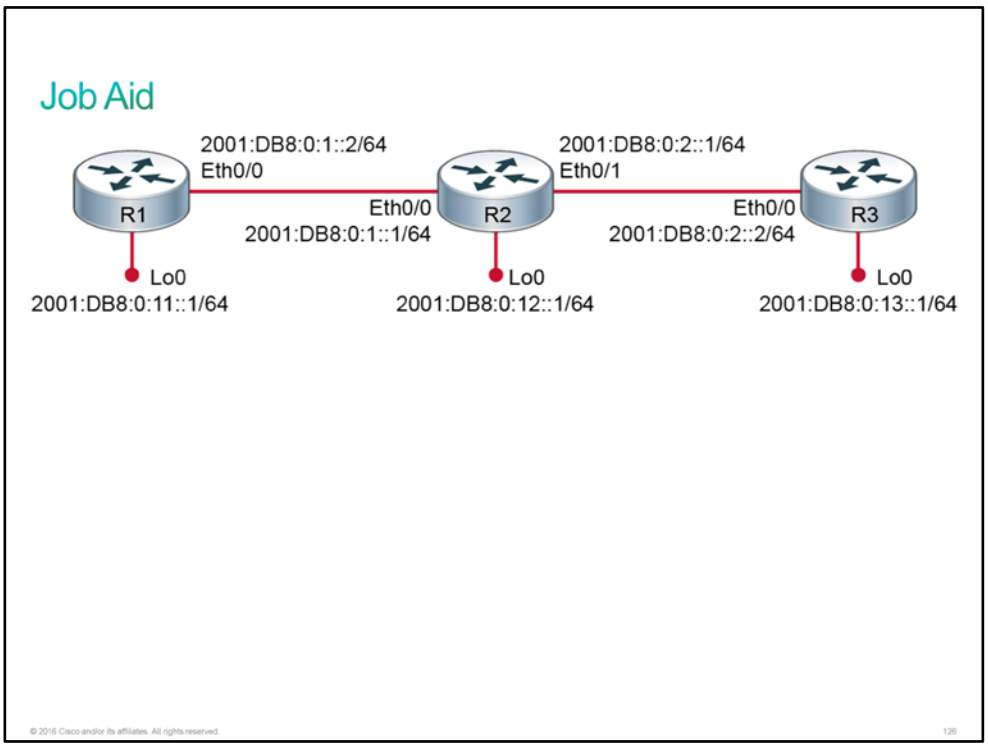

The configuration is as follows:

- All devices have their basic configurations in place including hostnames, IPv4, and IPv6 addresses.
- EIGRP for IPv6 is configured on R2 and R3:
  - AS number 100 is used.
  - Both routers are announcing Loopback interface network.

Device Information

Device Details

Device	Interface	IPv4 Address	IPv6 Address	Neighbor
R1	Ethernet0/0	10.10.1.1/24	2001:DB8:0:1::2/64	R2
R1	Loopback0	10.10.11.1/24	2001:DB8:0:11::1/64	—
R2	Ethernet0/0	10.10.1.2/24	2001:DB8:0:1::1/64	R1
R2	Ethernet0/1	10.10.2.1/24	2001:DB8:0:2::1/64	R3
R2	Loopback0	10.10.12.1/24	2001:DB8:0:12::1/64	—
R3	Ethernet0/0	10.10.2.2/24	2001:DB8:0:2::2/64	R2

Device	Interface	IPv4 Address	IPv6 Address	Neighbor
R3	Loopback0	10.10.13.1/24	2001:DB8:0:13::1/64	—

Task 1: Configure and Verify EIGRP for IPv6

Activity

Complete the following steps:

Step 1 Access the console of R2 and display EIGRP for IPv6-related configuration.

```
R2# show running-config

<... output omitted ...>
ipv6 unicast-routing
<... output omitted ...>
interface Loopback0
 ip address 10.10.12.1 255.255.255.0
 ipv6 address 2001:DB8:0:12::1/64
 ipv6 eigrp 100
!
interface Ethernet0/0
 description Link to R1
 ip address 10.10.1.2 255.255.255.0
 ipv6 address 2001:DB8:0:1::1/64
 ipv6 eigrp 100
!
!
interface Ethernet0/1
 description Link to R3
 ip address 10.10.2.1 255.255.255.0
 ipv6 address 2001:DB8:0:2::1/64
 ipv6 eigrp 100
!
<... output omitted ...>
ipv6 router eigrp 100
```

You should see that these configuration parts are preconfigured:

- IPv6 routing is globally enabled.
- EIGRP for IPv6 is enabled with AS 100 on interfaces Ethernet0/0 (towards router R1), Ethernet0/1 (towards router R3), and on Loopback0.

Step 2 Access the console of R3 and display EIGRP for IPv6-related configuration.

```
R3# show running-config

<... output omitted ...>
ipv6 unicast-routing
<... output omitted ...>
interface Loopback0
 ip address 10.10.13.1 255.255.255.0
 ipv6 address 2001:DB8:0:13::1/64
 ipv6 eigrp 100
!
interface Ethernet0/0
 description Link to R2
 ip address 10.10.2.2 255.255.255.0
 ipv6 address 2001:DB8:0:2::2/64
 ipv6 eigrp 100
!
<... output omitted ...>
ipv6 router eigrp 100
```

You should see that these configuration parts are preconfigured:

- IPv6 routing is globally enabled.
- EIGRP for IPv6 is enabled with AS 100 on interfaces Ethernet0/0 (towards router R2), and on Loopback0.

Configuring EIGRP for IPv6

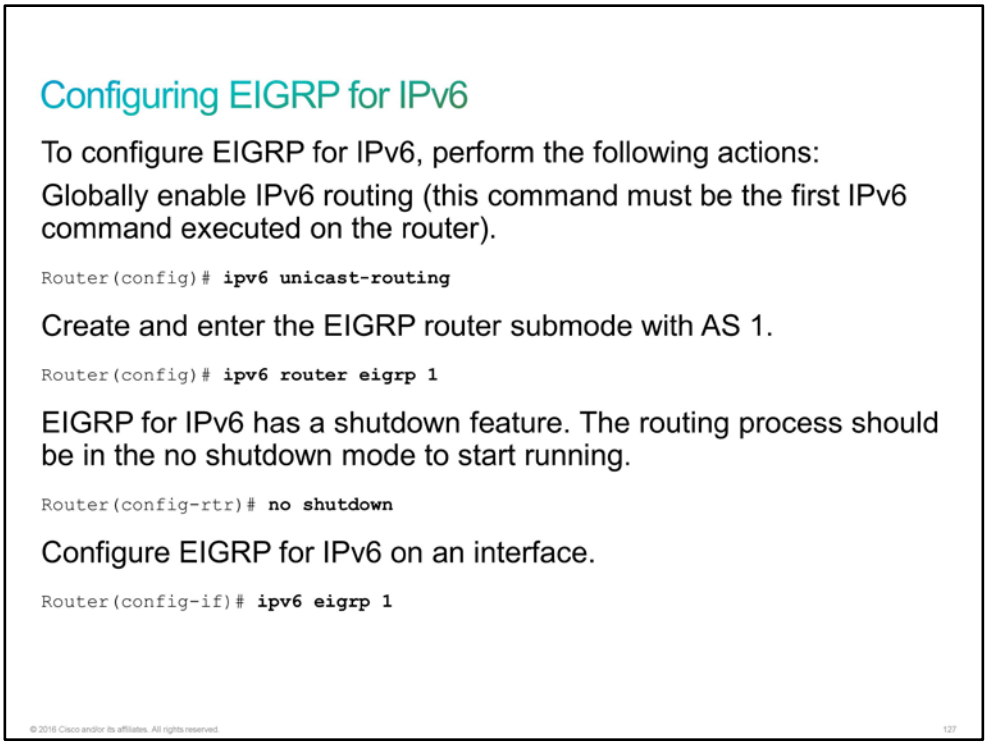

Command	Description
ipv6 unicast-routing	By default, IPv6 traffic forwarding is disabled. This command enables it.

Command	Description
ipv6 router eigrp *as-number*	To place the router in the router configuration mode, create an EIGRP routing process in IPv6, configure this process, and use the **ipv6 router eigrp** command in the global configuration mode.
no shutdown	EIGRP for IPv6 has a shutdown feature. The routing process should be in the **no shutdown** mode in order to start running. The default behavior is different between Cisco IOS Software versions.
[no] ipv6 eigrp *as-number*	To enable EIGRP for IPv6 on a specified interface, use the **ipv6 eigrp** command in the interface configuration mode. To disable EIGRP for IPv6, use the **no** form of this command.

These commands are some common configuration commands for EIGRP for IPv6. The syntax for these commands is similar, if not identical, to their IPv4 counterparts.

Step 3 Access the console of R1 and enable IPv6 unicast routing on R1.

Enter the following commands to the R1 router:

```
R1# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
R1(config)# ipv6 unicast-routing
```

Step 4 Enable EIGRP for IPv6 on AS 100 for R1.

Enter the following commands to the R1 router:

```
R1(config)# ipv6 router eigrp 100
R1(config-rtr)# exit
```

To activate AS 100, you must execute the **ipv6 router eigrp 100** configuration command. By default, this command defines the AS and enables it.

Step 5 Assign both the Ethernet0/0 and Loopback0 interfaces to AS 100.

Enter the following commands to the R1 router:

```
R1(config)# interface Ethernet0/0
R1(config-if)# ipv6 eigrp 100
R1(config-if)# exit
*Oct  8 08:43:50.642: %DUAL-5-NBRCHANGE: EIGRP-IPv6 100: Neighbor
FE80::A8BB:CCFF:FE00:3C00 (Ethernet0/0) is up: new adjacency
R1(config)# interface Loopback0
R1(config-if)# ipv6 eigrp 100
R1(config-if)# end
```

As soon as EIGRP for IPv6 was enabled on Ethernet0/0, a syslog message indicated the formation of a new neighbor relationship.

Verifying EIGRP for IPv6

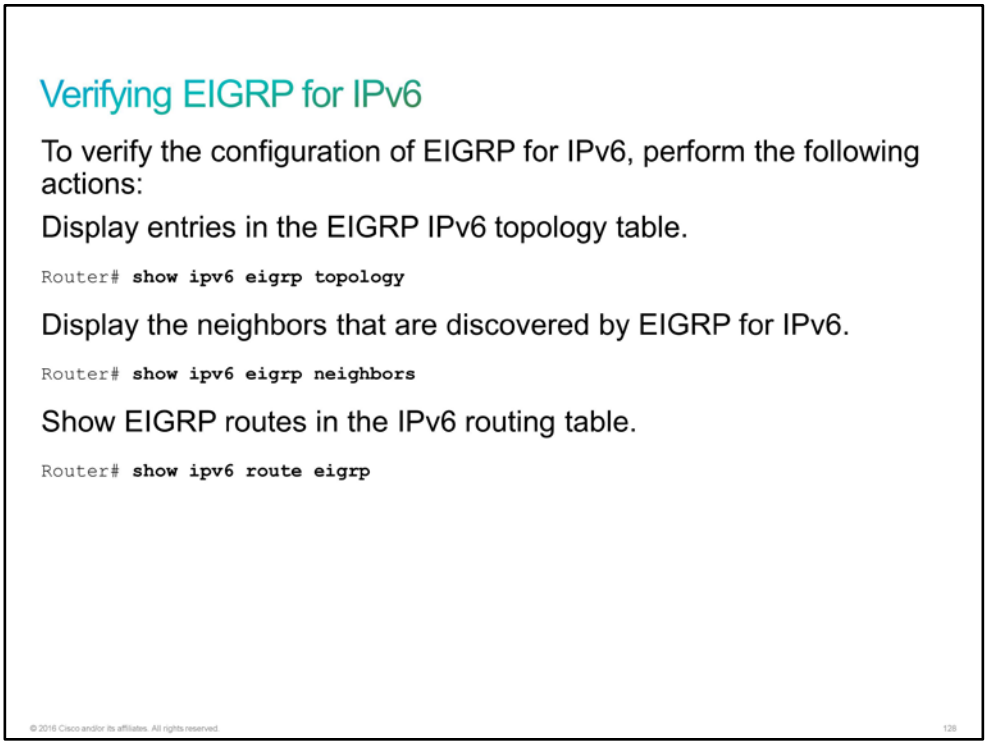

The three **show** commands that are listed have the same role that they have in EIGRP for IPv4.

To display entries in the EIGRP for IPv6 topology table, use the **show ipv6 eigrp topology** command in privileged EXEC mode.

To display the neighbors that are discovered by EIGRP for IPv6, use the **show ipv6 eigrp neighbors** command.

The **show ipv6 route eigrp** command shows the content of the IPv6 routing table that includes the routes specific to EIGRP.

Step 6 Display the IPv6 EIGRP neighbors for R1.

R2 is an EIGRP IPv6 neighbor of R1 on interface Ethernet0/0.

```
R1# show ipv6 eigrp neighbors
EIGRP-IPv6 Neighbors for AS(100)
H   Address                 Interface              Hold Uptime   SRTT   RTO  Q
Seq
                                                   (sec)         (ms)       Cnt
Num
0   Link-local address:     Et0/0                    14 00:01:39   17   102  0
8
    FE80::A8BB:CCFF:FE00:3C00
```

Note: The link local IPv6 address in your output may be different.

Step 7 Display the IPv6 routing table on R1 for networks learned through EIGRP.

R1 should learn about the network between R2 and R3 as well as the two networks that are associated with the Loopback interfaces on R2 and R3.

```
R1# show ipv6 route eigrp
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
D   2001:DB8:0:2::/64 [90/307200]
     via FE80::A8BB:CCFF:FE00:3C00, Ethernet0/0
D   2001:DB8:0:12::/64 [90/409600]
     via FE80::A8BB:CCFF:FE00:3C00, Ethernet0/0
D   2001:DB8:0:13::/64 [90/435200]
     via FE80::A8BB:CCFF:FE00:3C00, Ethernet0/0
```

Note: The link local IPv6 addresses in your output may be different.

Step 8 Display the topology table on R1 for EIGRP IPv6 and verify entries.

R1 should learn about the network between R2 and R3 as well as the two networks that are associated with the Loopback interfaces on R2 and R3.

```
R1# show ipv6 eigrp topology
EIGRP-IPv6 Topology Table for AS(100)/ID(10.10.11.1)
Codes: P - Passive, A - Active, U - Update, Q - Query, R - Reply,
       r - reply Status, s - sia Status

P 2001:DB8:0:2::/64, 1 successors, FD is 307200
        via FE80::A8BB:CCFF:FE00:3C00 (307200/281600), Ethernet0/0
P 2001:DB8:0:1::/64, 1 successors, FD is 281600
        via Connected, Ethernet0/0
P 2001:DB8:0:13::/64, 1 successors, FD is 435200
        via FE80::A8BB:CCFF:FE00:3C00 (435200/409600), Ethernet0/0
P 2001:DB8:0:11::/64, 1 successors, FD is 128256
        via Connected, Loopback0
P 2001:DB8:0:12::/64, 1 successors, FD is 409600
        via FE80::A8BB:CCFF:FE00:3C00 (409600/128256), Ethernet0/0
```

There should be five networks in the EIGRP topology table. There is the network between R1 and R2, the network between R2 and R3, and the three networks that are associated with the three loopback interfaces on the routers.

Note: The link local IPv6 addresses in your output may be different.

This is the end of the discovery lab.

Challenge

1. Which multicast address does EIGRP for IPv6 use?

   A. FF01::2
   B. FF01::10
   C. FF02::5
   D. FF02::A
   E. EIGRP for IPv6 does not use multicast addressing

2. Which of the following is a feature of IPv4 EIGRP but not IPv6 EIGRP?

   A. includes a shutdown feature
   B. uses DUAL
   C. scalable to large networks
   D. requires a network statement

3. Which command can turn off EIGRP for the IPv6 routing process?

   A. **enable**
   B. **enable router**
   C. **shutdown**
   D. **enable router shutdown**

4. Match the following:

ipv6 eigrp	enables IPv6 routing
no shutdown	places the router in router configuration mode and creates and configures an EIGRP routing process in IPv6
ipv6 unicast-routing	enables EIGRP for the IPv6 routing process
ipv6 router eigrp	enables EIGRP for IPv6 on a specified interface

5. Which of the following commands shows the content of the IPv6 routing table that includes the routes that are specific to EIGRP.

   A. **show ipv6 route**
   B. **show ipv6 eigrp topology**
   C. **show ipv6 eigrp neighbors**
   D. **show ipv6 route eigrp**

6. Which command produced the configuration output that is shown?

EIGRP-IPv6 Topology Table for AS(1)/ID(209.165.201.1)
Codes: P - Passive, A - Active, U - Update, Q - Query, R - Reply,
r - reply Status, s - sia Status

P 2001:DB8:D1A5:C900::/64, 1 successors, FD is 28160
via connected, GigabitEthernet0/165
P 2001:DB8:AC10:100::/64, 1 successors, FD is 156160
via FE80::FE99:47FF:FEE5:2671 (156160/128256), GigabitEthernet0/1

A. **show ipv6 eigrp interfaces**
B. **show ipv6 eigrp neighbors**
C. **show ipv6 eigrp topology**
D. **show ipv6 route eigrp**

7. By default, which two metrics does EIGRP use to compute the best path? (Choose two.)

A. minimum bandwidth
B. reliability
C. delay
D. load
E. MTU

Answer Key

Challenge

1. D
2. D
3. C
4.

ipv6 unicast-routing	enables IPv6 routing
ipv6 router eigrp	places the router in router configuration mode and creates and configures an EIGRP routing process in IPv6
no shutdown	enables EIGRP for the IPv6 routing process
ipv6 eigrp	enables EIGRP for IPv6 on a specified interface

5. D
6. C
7. A, C

Lesson 3: Troubleshooting EIGRP

Introduction

Two different customers have called CCS with complaints about the loss of network connectivity since EIGRP has been implemented. Trouble tickets have been created for both customers. Bob has assigned the trouble tickets to you.

Troubleshooting EIGRP Issues

You have just finished configuring EIGRP, and you tested connectivity to a remote network with a ping, but the ping failed. As you begin to troubleshoot the situation, keep in mind that EIGRP problems usually fall into one of the following categories:

- Neighbor adjacency issues
- Routing issues

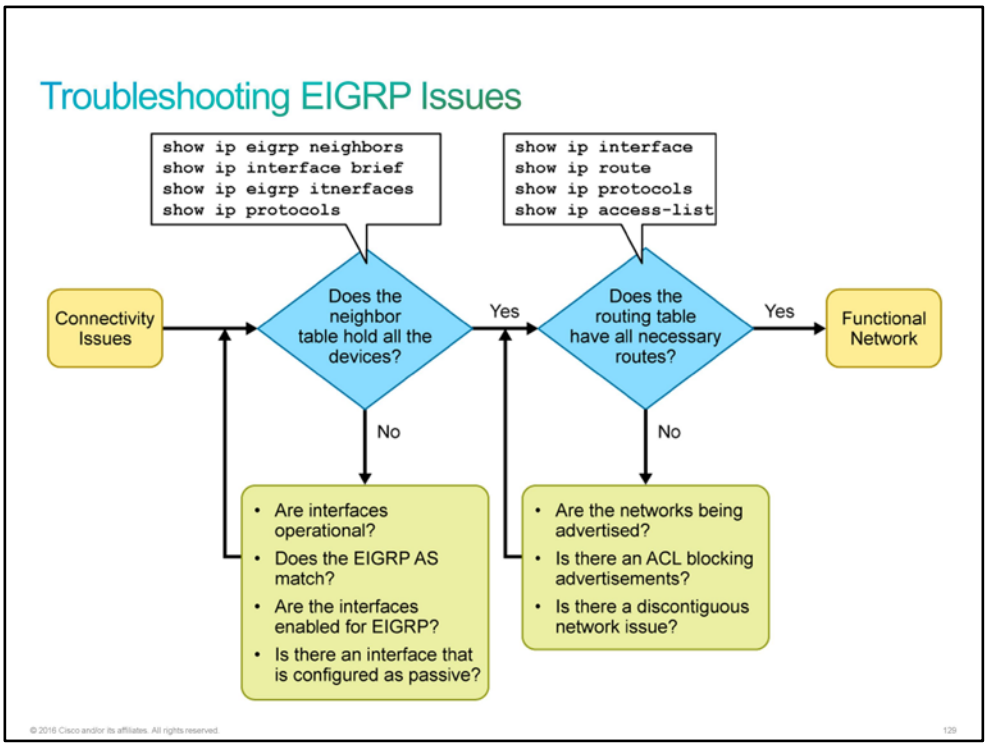

When you encounter an EIGRP problem, first use the **show ip eigrp neighbors** command to check for neighbor adjacency issues.

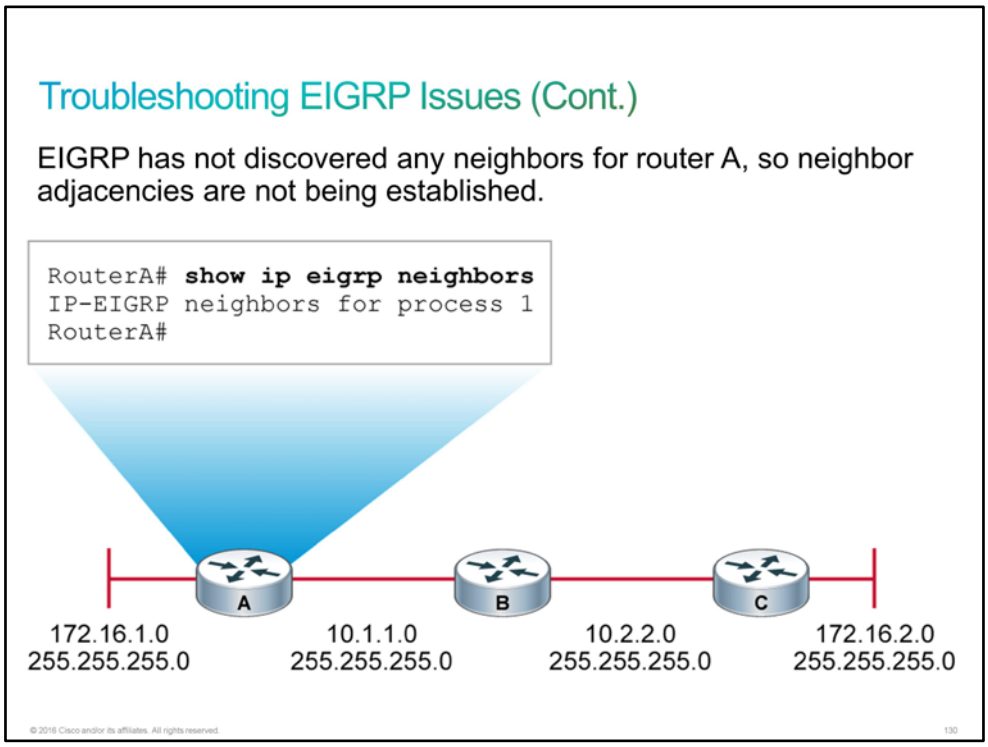

The following issues can prevent neighbor adjacencies from being established:

- The interface between the devices is down.
- The routers have mismatching EIGRP autonomous systems.
- The EIGRP process is not enabled on one of the interfaces that connects the devices.
- One of the interfaces that connects the devices is configured as a passive interface.

Aside from these issues, there are several other, more advanced issues that can cause neighbor relationships to not be formed. For example, mismatched K values can prevent neighbor relationships from being formed.

When you have eliminated EIGRP neighbor relationship issues as the cause of the problem, check to see if there is a routing problem. In the example in the next figure, the pings from the network 172.16.1.0 to the hosts on 10.2.2.0 or 172.16.2.0 will fail because router A does not have routes to those networks.

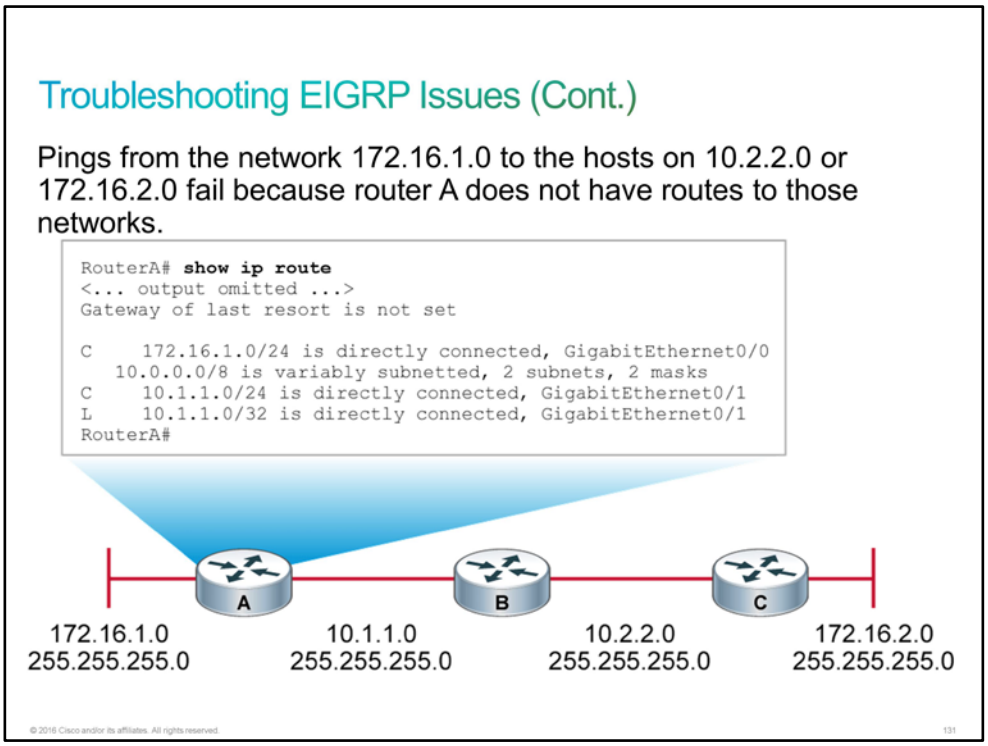

Issues that may prevent a routing table from learning the appropriate routes from EIGRP include the following:

- Networks are not being advertised on remote routers.
- An access list is blocking advertisements of remote networks.
- Automatic route summarization is causing confusion in your discontiguous network.

Note	Although there are also **debug** commands that provide excellent diagnostic information, you should use the **debug** commands with caution. In general, it is recommended that you use these commands only under the direction of your router technical support representative when you are troubleshooting specific problems.

Troubleshooting EIGRP Neighbor Issues

If the output of the **show ip eigrp neighbors** command indicates that neighbor relationships are not being formed after you configure EIGRP, your next troubleshooting step, of course, is to determine what is preventing these adjacencies from forming.

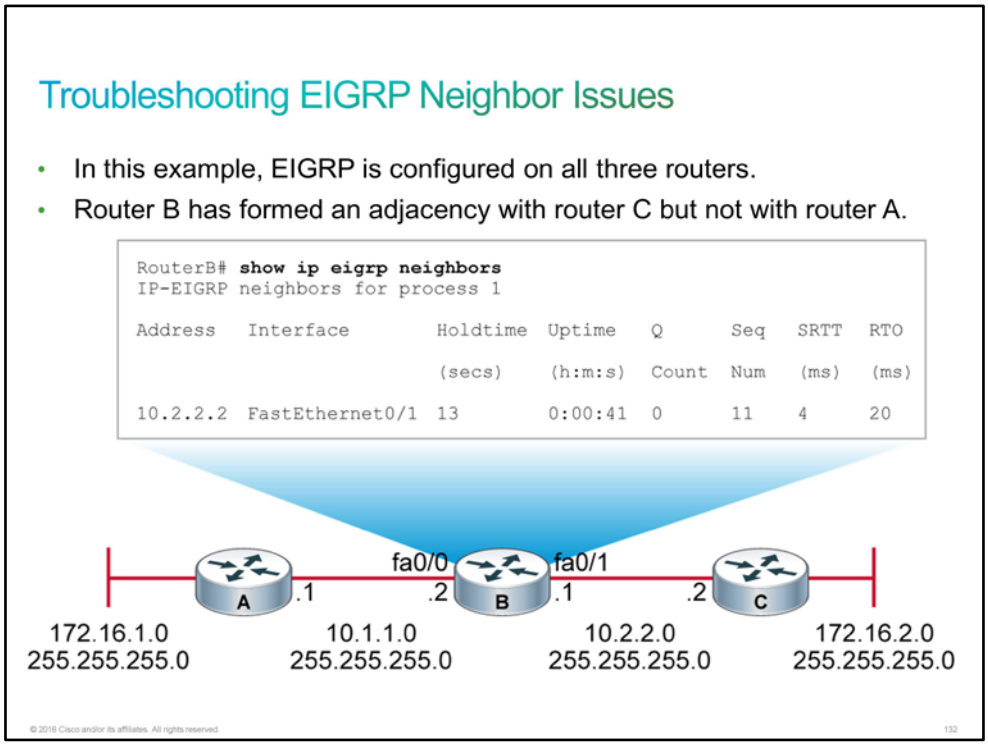

A prerequisite for the establishment of neighbor adjacencies is the OSI Layer 3 connectivity, so you want to perform basic connectivity troubleshooting steps on the link between router B and router C, starting at the physical layer.

Are All Interface Statuses "Up/Up"?

Use the **show ip interface brief** command to make sure that the interfaces that connect the two routers are "up."

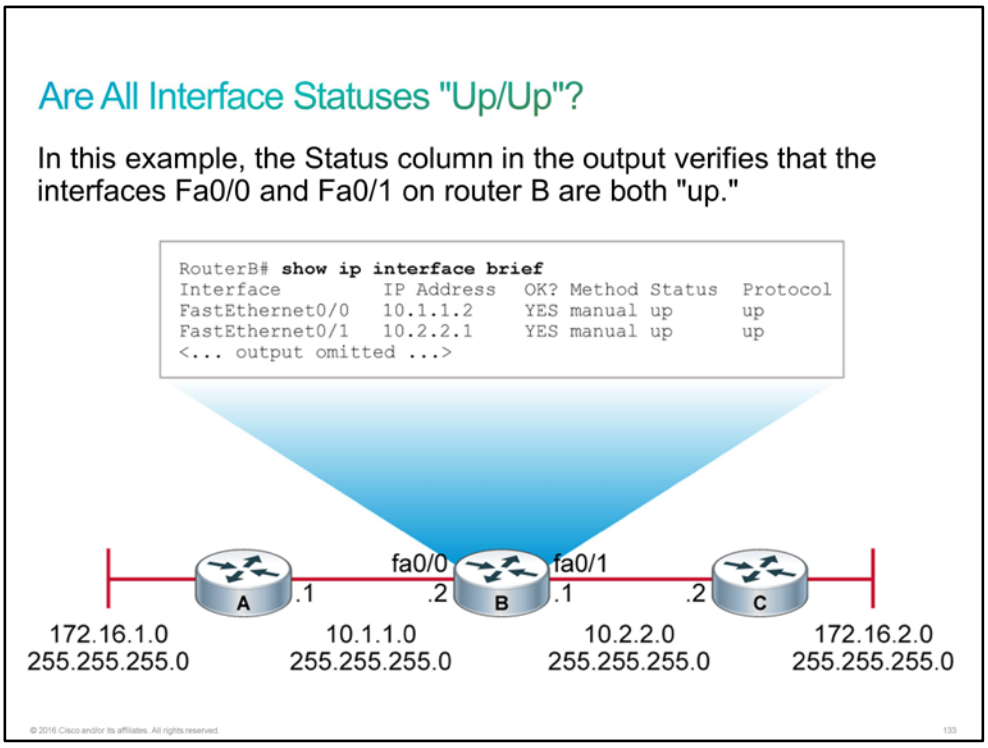

Note	When EIGRP is configured, you may receive a "not on common subnet" message on your router console. This message indicates that there is an incorrect IP address on one of the EIGRP neighbor interfaces. For a neighbor adjacency to be formed between two routers, the interfaces that connect the routers must be on the same subnet.

Is There an AS Mismatch?

Another prerequisite for the establishment of neighbor adjacencies is the matching of AS numbers. You could use the **show ip protocols** command to determine whether the routers have the same AS number. This command displays the name and AS number of the currently running routing protocol.

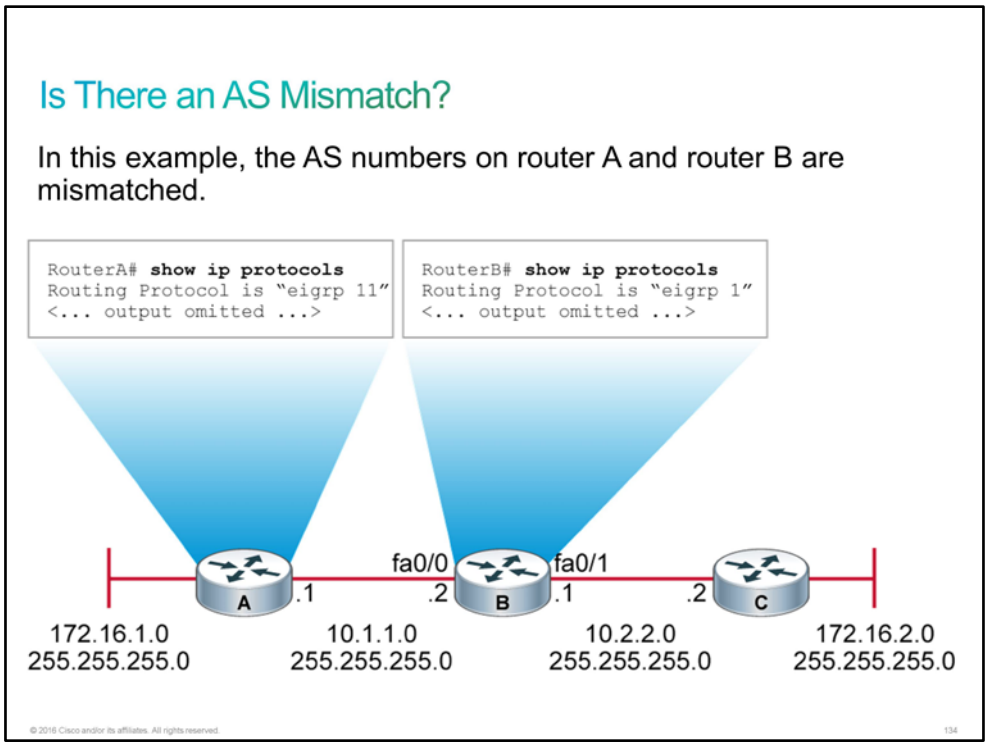

To correct a mismatched AS number situation, do the following commands to reconfigure the router that has the wrong AS number:

- Remove the old EIGRP routing process that has wrong AS number.
- Enable the EIGRP routing process with the correct AS number.
- Include the networks into the newly created EIGRP process.

The example above should be corrected in the following manner:

```
RouterA(config)# no router eigrp 11
RouterA(config)# router eigrp 1
RouterA(config-router)# network 10.1.1.0
RouterA(config-router)# network 172.16.1.0
```

Is EIGRP Enabled on the Interface?

The "Routing for Networks" section of the **show ip protocols** command output indicates which networks have been configured. Any interfaces in those networks participate in EIGRP.

```
RouterB# show ip protocols
Routing Protocol is "eigrp 11"
  Outgoing update filter list for all interfaces is not set
  Incoming update filter list for all interfaces is not set
  Default networks flagged in outgoing updates
  Default networks accepted from incoming updates
  EIGRP metric weight K1=1, K2=0, K3=1, K4=0, K5=0
  EIGRP maximum hopcount 100
  EIGRP maximum metric variance 1
  Redistributing: eigrp 11
<... output omitted ...>
  Maximum path: 4
  Routing for Networks:
    10.1.1.0
    10.2.2.0
  Routing Information Sources:
    Gateway         Distance      Last Update
    (this router)         90            00:01:08
    10.2.2.2              90            00:01:08
  Distance: internal 90 external 170
```

Note	As shown in the command output above, you can also use the **show ip protocols** command to verify the "K" values that are being used in EIGRP metric calculations.

Neighbor adjacencies will be established between two routers only if the connecting interfaces on the two routers are enabled for the EIGRP process. You can use the **show ip eigrp interfaces** command to determine on which interfaces EIGRP is enabled and to learn the information about EIGRP relating to those interfaces. If an interface is not listed in the output of this command, the router is not using EIGRP on that interface.

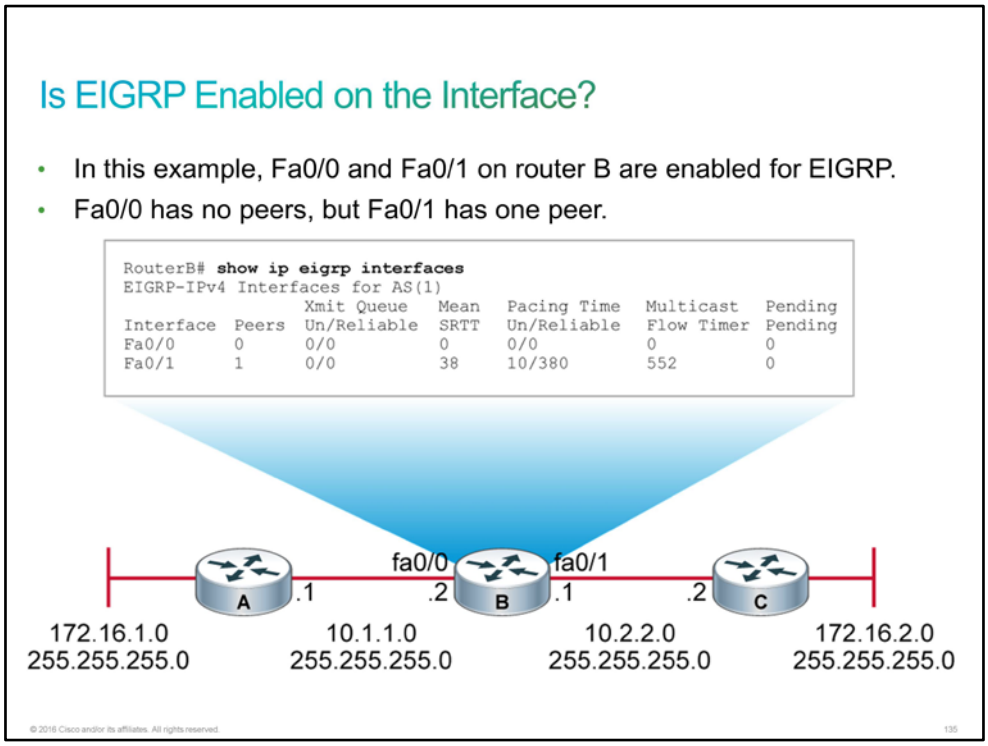

The table further explains the output of the **show ip eigrp interfaces** command.

Parameter	Description
AS(1)	The AS number that is specified with the router command
Interface	The interface over which EIGRP is configured
Peers	The number of directly connected EIGRP neighbors on the interface
Xmit Queue Unreliable and Reliable	The number of packets remaining in the Unreliable and Reliable queues
Mean SRTT	The average SRTT interval (in milliseconds) for all the neighbors on the interface
Pacing Time Unreliable and Reliable	The number of milliseconds to wait after transmitting unreliable and reliable packets
Multicast Flow Timer	The number of milliseconds to wait for acknowledgment of a multicast packet by all the neighbors before transmitting the next multicast packet
Pending Routes	The number of routes in the packets in the transmit queue that are waiting to be sent

To correct a situation in which an interface that should be enabled for EIGRP is not enabled for EIGRP, use the **network** command. This command, which is configured under the EIGRP routing process, specifies which networks will participate in the EIGRP process. Only the interfaces that fall within the range of addresses that are defined by the network entries will participate in the EIGRP process.

Is Passive Interface Configured?

Another situation in which neighbor adjacencies may fail to form occurs when an interface that should form a neighbor adjacency is configured as passive. You can use the **passive-interface** command in an EIGRP configuration to specify that certain interfaces are passive. This result means that those interfaces will not send or receive hello packets and will not form neighbor adjacencies. To determine whether any interface on a router is configured as passive, use the **show ip protocols** command.

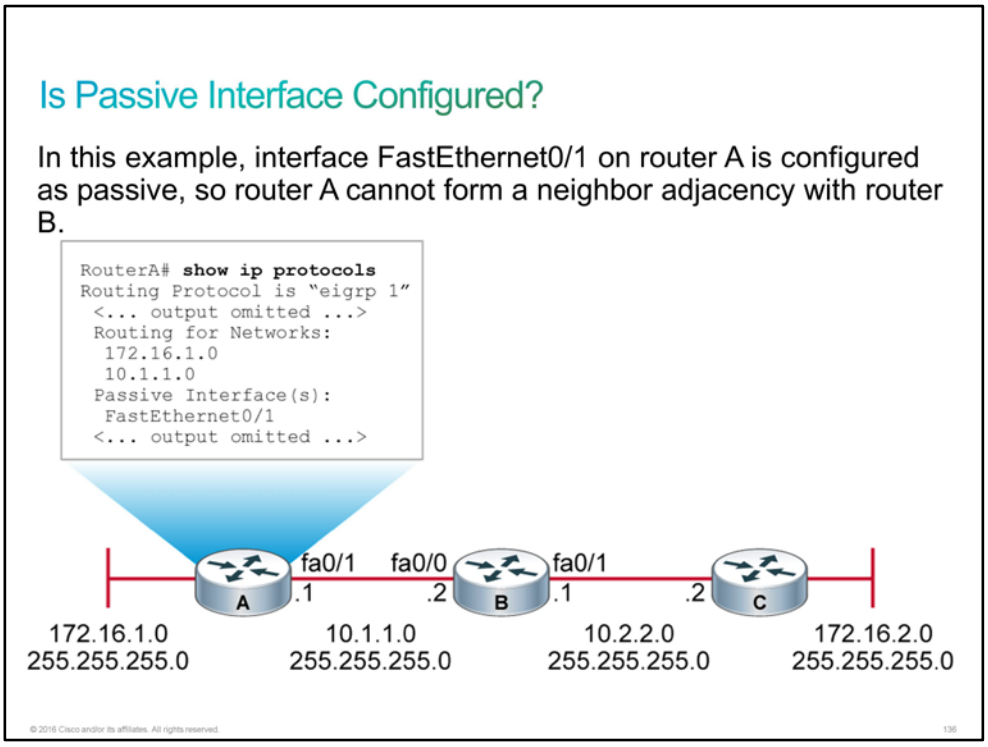

To return an interface to a nonpassive state, use the **no passive-interface** command.

Is Access List Blocking Route Advertisements?

However, if the desired routes still do not appear in the routing tables, an access-list configured on an interface could be blocking EIGRP route advertisements.

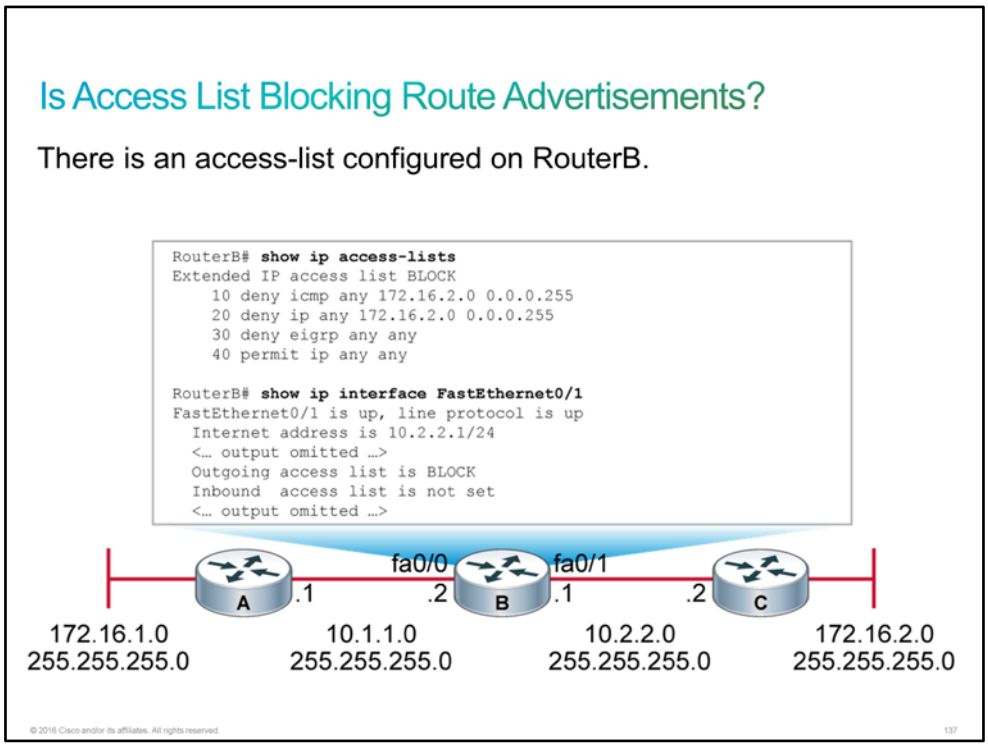

To block EIGRP traffic, you can configure an access-list, blocking this type of the traffic and apply it on the interface in desired direction. If you applied the access-list on the interface only in one direction, the neighborship will be flapping between those two routers, since EIGPR traffic will be blocked only in one direction.

The **show ip access-list** and **show ip interface** *interface slot/number* commands will give you this information. In the next figure, the output of these two commands indicates that an outgoing access-list has been set on the RouterB. The access-list is set on the interface FastEthernet0/1 in the outgoing direction and is blocking all EGIRP traffic.

Troubleshooting EIGRP Routing Table Issues

Consider this scenario. Router A and router B have established a neighbor adjacency. Router B has also established a neighbor adjacency with router C. However, a ping test from router A to a host in the 172.16.2.0/24 network is still not successful. You issue the **show ip route** command on router A and notice that router A has no route to the destination network of 172.16.2.0/24.

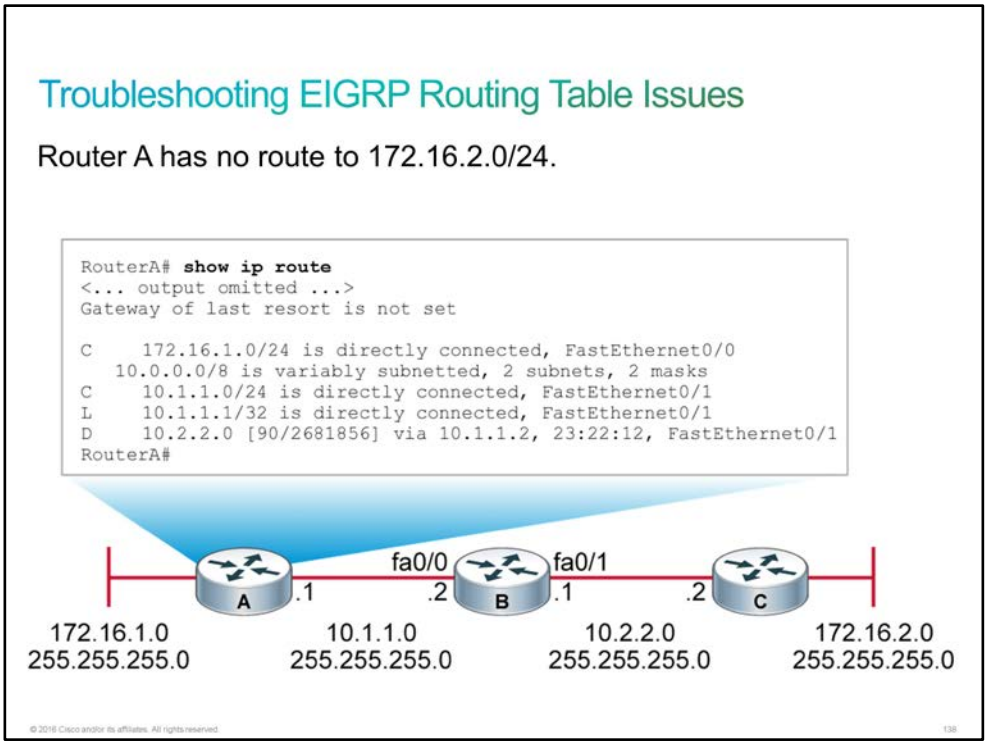

You issue the same command on router B and discover that the router is missing the route to 172.16.2.0.

Are All Required EIGRP Networks Being Advertised?

In this situation, it is a good idea to use the **show ip protocols** command to determine the reason for the missing routes. This command will tell you whether the 172.16.2.0/24 network is being advertised by its directly connected router, which is router C. The figure shows that, in this example, router C is not advertising the network 172.16.2.0/24 to its EIGRP neighbors.

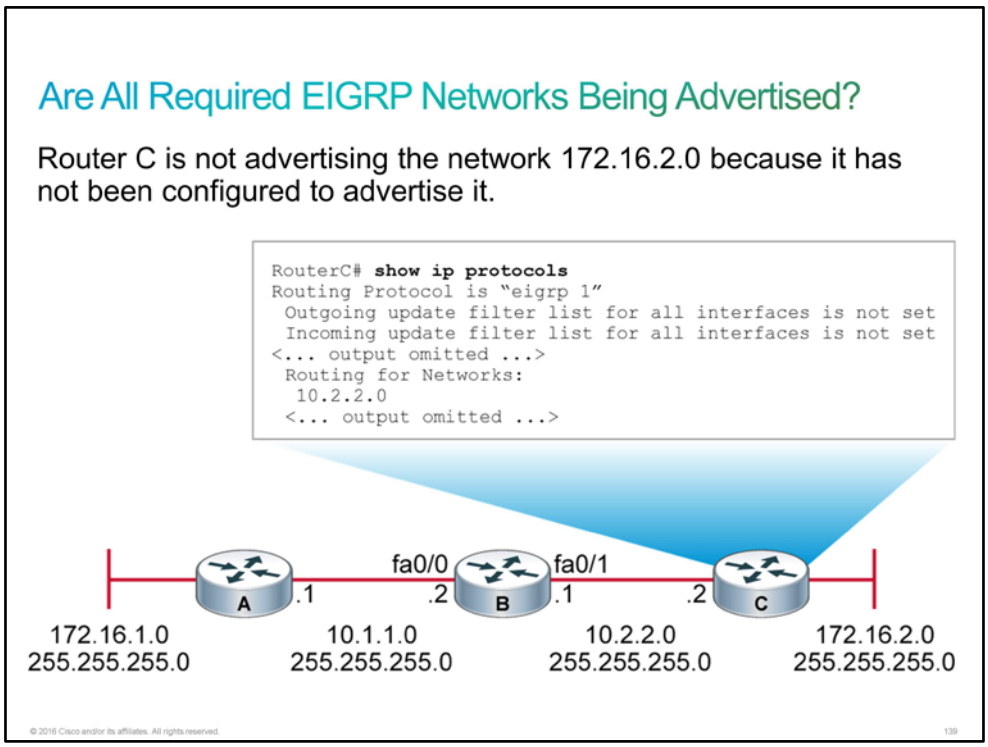

After you use the **network** command in the router configuration mode to configure router C to advertise this network, issue the **show ip protocols** command again on router C to verify that the network 172.16.2.0 is now included in the EIGRP process. Then use the **show ip route** command on router A and router B to confirm that they now have routes to the 172.16.2.0 network. If the routes exist in the routing tables on these routers, all is well.

Is There a Discontiguous Network?

Now consider a different scenario. In the figure, the networks 10.1.1.0 and 10.2.2.0 separate the subnets of the network 172.16.0.0/16: 172.16.1.0/24 and 172.16.2.0/24. The automatic route summarization feature of the EIGRP routing process summarizes the routes on the network number boundaries. Both router A and router B summarized the subnets to the 172.16.0.0/16 classful boundary. As shown in the example, router B is not receiving individual routes for the 172.16.1.0/24 and 172.16.2.0/24 subnets. The result is that router B has two routes to 172.16.0.0/16 in the routing table, which can result in an inaccurate routing and packet loss. This issue is referred to as a discontiguous network issue. A discontiguous network comprises a major network that is separated by another major network. You can address this issue by disabling the auto-summarization feature under the EIGRP process.

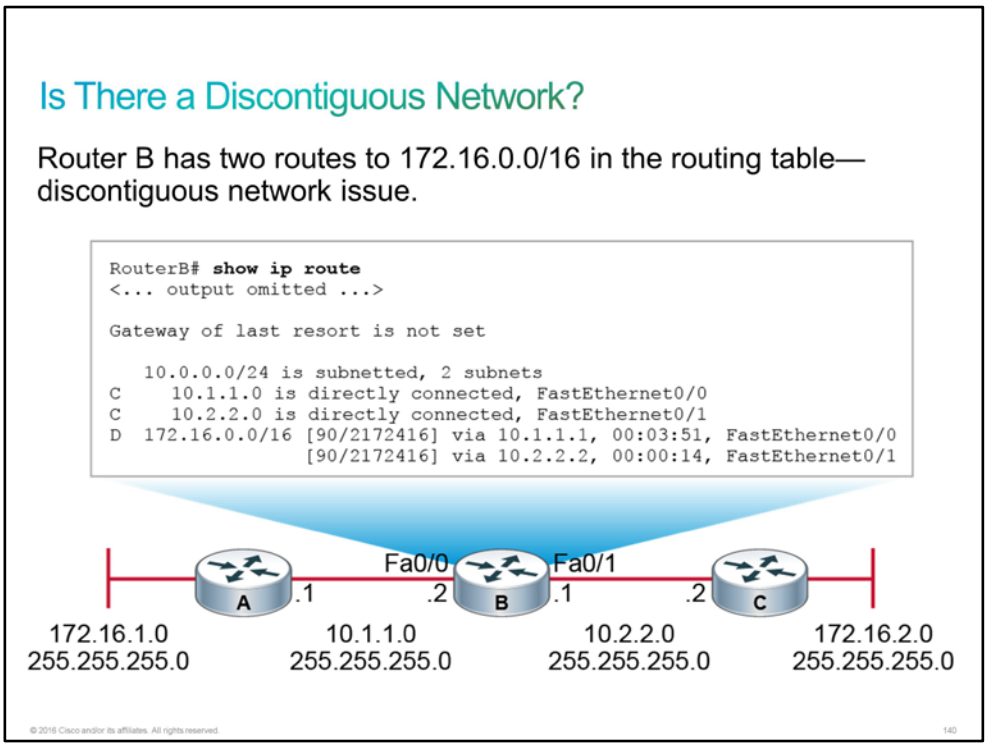

Automatic route summarization (causing a discontiguous network issue) is enabled by default in the Cisco IOS Software before Release 15 (for example, Cisco IOS Release 12). In this case, you would have to use the **no auto-summary** command. If you are using Cisco IOS Software Release 15 or later, you do not need to use the **no auto-summary** command to disable automatic route summarization. Automatic route summarization is disabled by default in Cisco IOS Software Releases 15 or later.

To solve the discontiguous network issue, make sure that all the routers in the figure, that have interfaces connected to multiple different classful networks have the automatic route summarization option disabled (the **no auto-summary** command).

You can verify, if route summarization is enabled, using the **show ip protocols** command.


```
RouterB# show ip protocols
*** IP Routing is NSF aware ***

Routing Protocol is "eigrp 1"
<... output omitted ...>

  Automatic Summarization: disabled
  Maximum path: 4
  Routing for Networks:
    10.0.0.0
  Routing Information Sources:
    Gateway         Distance      Last Update
    10.1.1.1              90      00:00:31
    10.2.2.2              90      00:00:29
  Distance: internal 90 external 170
```

Troubleshooting EIGRP for IPv6 Issues

Configuring EIGRP for IPv6 is very similar to configuring EIGRP for IPv4. The main difference is that EIGRP is enabled on the interface for IPv6 with the **ipv6 eigrp** *as-number* command. Therefore, troubleshooting EIGRP for IPv6 is very similar to troubleshooting EIGRP for IPv4.

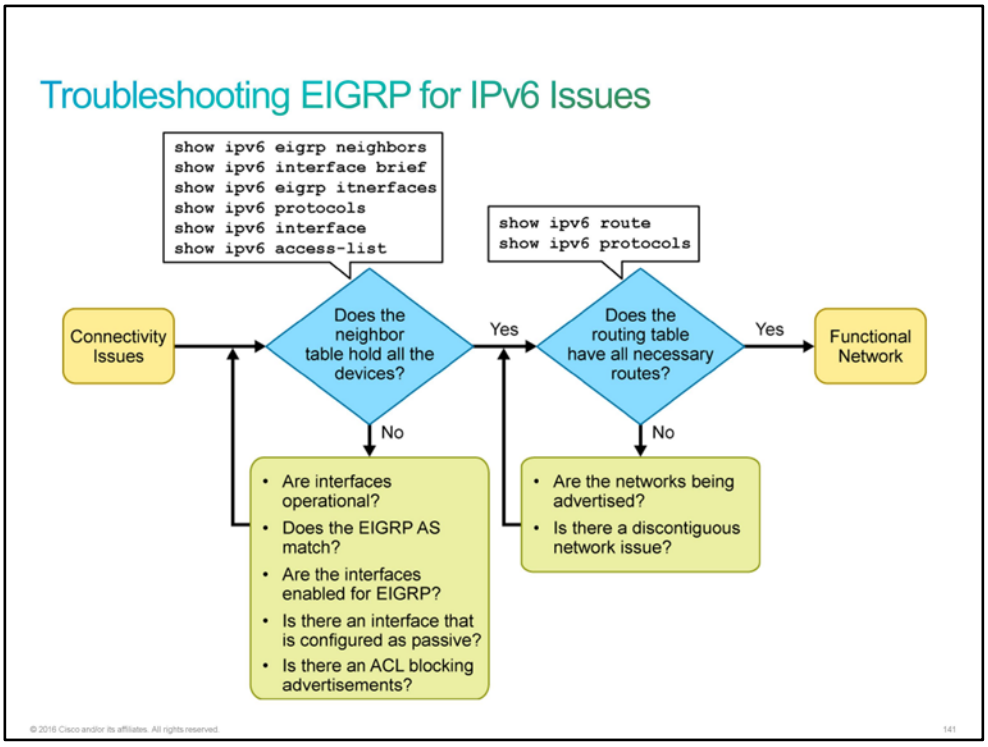

To check the IPv6 routing protocols on the router, use the **show ipv6 protocols** command. The output will show the IPv6 routing protocols that are enabled on the router. The EIGRP section shows metric weights, router ID, EIGRP interfaces, redistribution information, and so on.

To display the neighbors that are discovered by EIGRP for IPv6, use the **show ipv6 eigrp neighbors** command.

The **show ipv6 route eigrp** command shows the content of the IPv6 routing table that includes the routes that are specific to EIGRP.

To verify the topology table, use the **show ipv6 eigrp topology** command. You can see all routing updates that the router received, with AD and FD information, next-hop, and so on.

There are also other things to check, that are not directly related to EIGRP configuration. To check whether IPv6 addresses have been assigned on the interfaces, use the **show ipv6 interface brief**. To verify if there are any access-lists configured, use **show ipv6 access-list**s command.

Discovery 42: Troubleshoot EIGRP

Introduction

This discovery will guide you through the troubleshooting of various EIGRP configuration issues. The virtual lab is prepared with the devices that are represented in the topology diagram and the "Device Information" table. All devices have their basic configurations in place, including hostnames and IP addresses. EIGRP AS 10 has been configured on all seven routers, but there are problems with the router configurations. Each router has a loopback interface with the IP address 192.168.*R*.1/24 (where *R* is the router number). The routing table on R1 is missing routes to the loopback interface networks for each of its peers. In this discovery, you will troubleshoot and fix the problems that are associated with the routing of each of these networks.

You will start with the R2 loopback network and proceed one at a time, finishing with R7, which is also configured for EIGRP IPv6 routing. In each case, you will first determine the root cause. You will then fix the issue and verify that the route is properly defined in the routing table of R1.

Topology

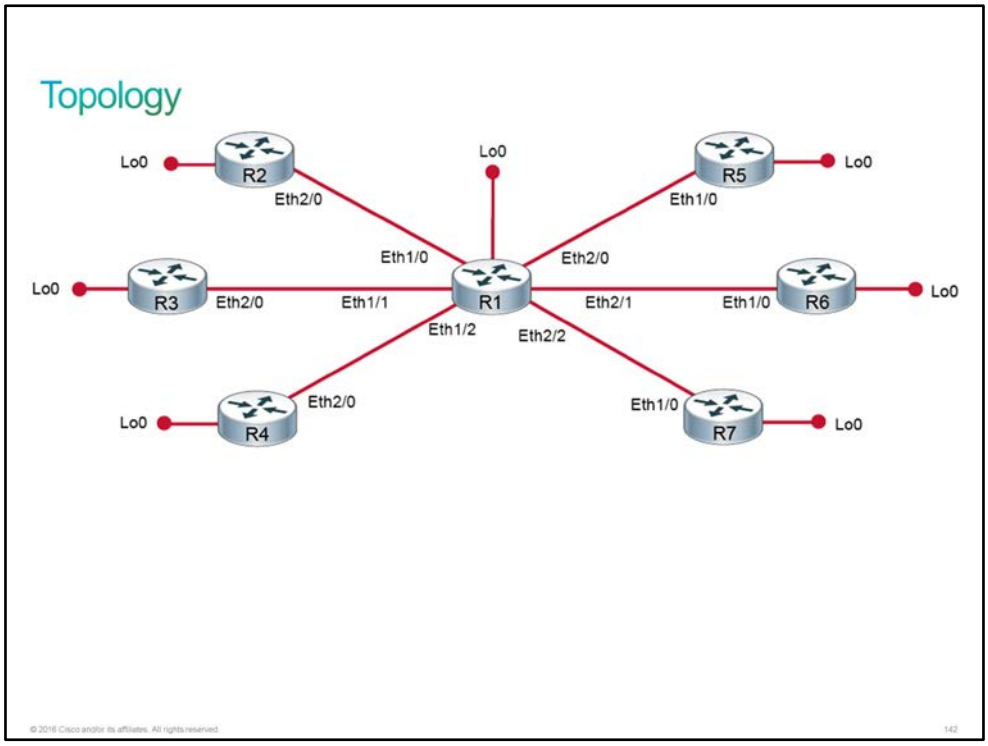

Job Aid

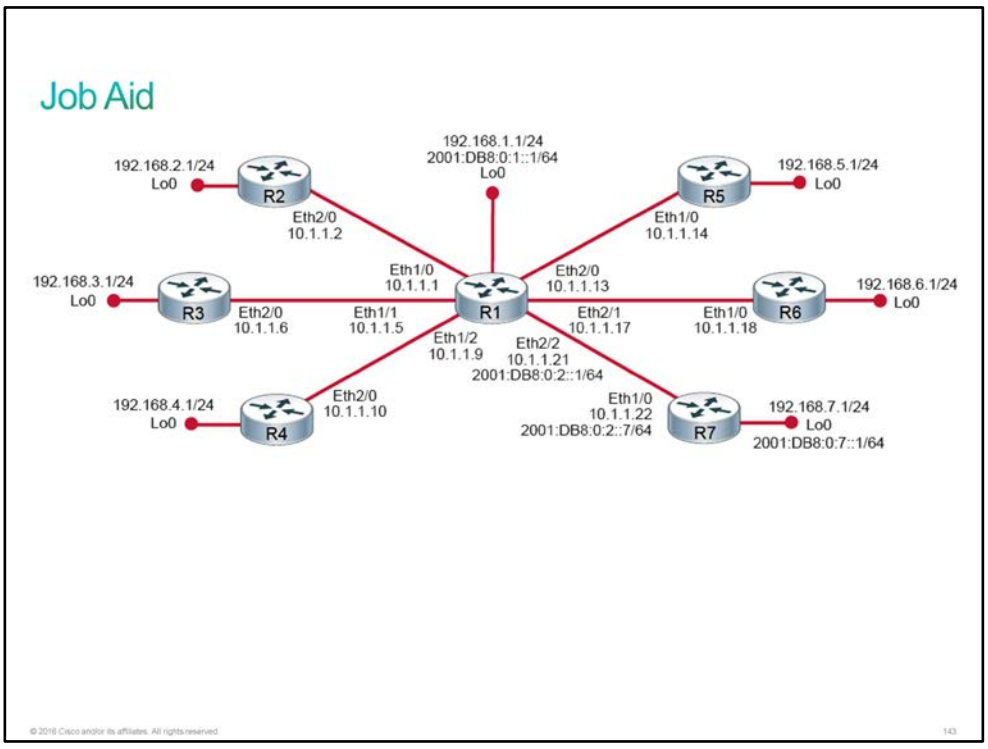

The configuration is as follows:

- All devices have their basic configurations in place including hostnames, and IP addresses. R1 and R7 also have IPv6 addresses configured.
- EIGRP AS 10 has been configured on all seven routers, but there are problems with the router configurations.
  - The routing table on R1 is missing routes to the loopback interface networks for each of its peers.
- R1 and R7 are also configured for EIGRP IPv6 routing, using AS 10.

Device Information

Device Details

Device	Interface	Neighbor	IP Address
R1	Loopback0	—	192.168.1.1/24
R1	Ethernet1/0	R2	10.1.1.1/30
R1	Loopback0	—	2001:DB8:0:1::1/64
R1	Ethernet2/2	R7	2001:DB8:0:2::1/64
R1	Ethernet1/1	R3	10.1.1.5/30

Device	Interface	Neighbor	IP Address
R1	Ethernet1/2	R4	10.1.1.9/30
R1	Ethernet2/0	R5	10.1.1.13/30
R1	Ethernet2/1	R6	10.1.1.17/30
R1	Ethernet2/2	R7	10.1.1.21/30
R2	Ethernet2/0	R1	10.1.1.2/30
R2	Loopback0	—	192.168.2.1/24
R3	Ethernet2/0	R1	10.1.1.6/30
R3	Loopback0	—	192.168.3.1/24
R4	Ethernet2/0	R1	10.1.1.10/30
R4	Loopback0	—	192.168.4.1/24
R5	Ethernet1/0	R1	10.1.1.14/30
R5	Loopback0	—	192.168.5.1/24
R6	Ethernet1/0	R1	10.1.1.18/30
R6	Loopback0	—	192.168.6.1/24
R7	Ethernet1/0	R1	10.1.1.22/30
R7	Loopback0	—	192.168.7.1/24
R7	Ethernet1/0	R1	2001:DB8:0:2::7/64
R7	Loopback0	—	2001:DB8:0:7::1/64

Task 1: Troubleshoot EIGRP

Activity

Complete the following steps:

Step 1 The network 192.168.2.0/24 does not exist in the routing table of R1. This network is associated with the Loopback0 interface of R2. Try to determine the root cause for this issue.

Note: There is no single best procedure for troubleshooting any network issue. The goal is to isolate the root cause. One strategy is to work from the application layer down. If there are aspects of the application that are working, it implies that there must be IP connectivity and link layer connectivity below the application. If the application does not function, check the IP connectivity next.

You might use the following commands on R1 and observe these results:

- **show ip eigrp neighbor—**R2 (10.1.1.2) is not in the EIGRP neighbor table.
- **show ip interface brief—**The interface Ethernet1/0 is "up/up."
- **ping 10.1.1.2—**R2 responds to the ping.

These results indicate that there is Layer 2 and Layer 3 connectivity, but there is an issue with EIGRP communication.

On R2, virtually all show commands that are associated with EIGRP will provide the hint. The AS number 10 is used across the network, but R2 is configured with AS 100.

```
R2# show ip protocols
*** IP Routing is NSF aware ***
Routing Protocol is "eigrp 100"
<... output omitted ...>

R2# show ip eigrp interfaces
EIGRP-IPv4 Interfaces for AS(100)
<... output omitted ...>

R2# show ip eigrp neighbors
EIGRP-IPv4 Neighbors for AS(100)

R2# show ip eigrp topology
EIGRP-IPv4 Topology Table for AS(100)/ID(192.168.2.1)
<... output omitted ...>
```

You can see the root cause in the configuration.

```
R2# show running-config | section eigrp
router eigrp 100
 network 10.0.0.0
 network 192.168.2.0
```

Step 2 With the root cause determined, fix the problem and verify that the route to 192.168.2.0/24 now exists in the routing table of R1.

Enter the following configuration on the R2 router:

```
R2# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
R2(config)# no router eigrp 100
R2(config)# router eigrp 10
R2(config-router)# network 10.0.0.0
*Oct 13 13:34:45.096: %DUAL-5-NBRCHANGE: EIGRP-IPv4 10: Neighbor 10.1.1.1
(Ethernet2/0) is up: new adjacency
R2(config-router)# network 192.168.2.0
R2(config-router)# end
```

The neighbor adjacency is initiated when the network 10.0.0.0 is enabled under the EIGRP AS 10.


```
R1# show ip route 192.168.2.0
Routing entry for 192.168.2.0/24
  Known via "eigrp 10", distance 90, metric 409600, type internal
  Redistributing via eigrp 10
  Last update from 10.1.1.2 on Ethernet1/0, 00:03:31 ago
  Routing Descriptor Blocks:
  * 10.1.1.2, from 10.1.1.2, 00:03:31 ago, via Ethernet1/0
      Route metric is 409600, traffic share count is 1
      Total delay is 6000 microseconds, minimum bandwidth is 10000 Kbit
      Reliability 255/255, minimum MTU 1500 bytes
      Loading 1/255, Hops 1

R1# show ip route eigrp
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

D     192.168.2.0/24 [90/409600] via 10.1.1.2, 00:03:49, Ethernet1/0
```

Step 3 The network 192.168.3.0/24 does not exist in the routing table of R1. This network is associated with the Loopback0 interface of R3. Try to determine the root cause for this issue.

You might use the following commands on R1 and observe these results:

- **show ip eigrp neighbor**—R3 (10.1.1.6) is not in the EIGRP neighbor table.
- **show ip interface brief**—The interface Ethernet1/1 is "up/up."
- **ping 10.1.1.6**—R3 responds to a ping from R1.

These results indicate that there is Layer 2 and Layer 3 connectivity, but there is an issue with EIGRP communication.

You might use the following commands on R3 and observe these results:

- **show ip eigrp neighbor**—R1 (10.1.1.5) is not in the EIGRP neighbor table.
- **show ip eigrp interfaces**—Only Loopback0 is included (Ethernet2/0 is missing).
- **show ip protocols**—Only the network for Loopback0 is included (10.0.0.0 is missing).

```
R3# show ip eigrp neighbors
EIGRP-IPv4 Neighbors for AS(10)

R3# show ip eigrp interfaces
EIGRP-IPv4 Interfaces for AS(10)
                              Xmit Queue   PeerQ        Mean   Pacing Time
ticast    Pending
Interface              Peers  Un/Reliable  Un/Reliable  SRTT   Un/Reliable  w
Timer   Routes
Lo0                      0        0/0       0/0           0       0/0        0
0

R3# show ip protocols
*** IP Routing is NSF aware ***

Routing Protocol is "eigrp 10"
  Outgoing update filter list for all interfaces is not set
  Incoming update filter list for all interfaces is not set
  Default networks flagged in outgoing updates
  Default networks accepted from incoming updates
  EIGRP-IPv4 Protocol for AS(10)
    Metric weight K1=1, K2=0, K3=1, K4=0, K5=0
    NSF-aware route hold timer is 240
    Router-ID: 192.168.3.1
    Topology : 0 (base)
      Active Timer: 3 min
      Distance: internal 90 external 170
      Maximum path: 4
      Maximum hopcount 100
      Maximum metric variance 1

  Automatic Summarization: disabled
  Maximum path: 4
  Routing for Networks:
    192.168.3.0
  Routing Information Sources:
    Gateway         Distance      Last Update
  Distance: internal 90 external 170
```

EIGRP must include 10.1.1.6 (the IP address for Ethernet2/0) among its routed networks.

You can see the root cause in the configuration.

```
R3# show running-config | section eigrp
router eigrp 10
 network 192.168.3.0
```

Step 4 With the root cause determined, fix the problem and verify that the route to 192.168.3.0/24 now exists in the routing table of R1.

Enter the following configuration on the R3 router:

```
R3# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
R3(config)# router eigrp 10
R3(config-router)# network 10.0.0.0
*Oct 14 07:10:55.227: %DUAL-5-NBRCHANGE: EIGRP-IPv4 10: Neighbor 10.1.1.5
(Ethernet2/0) is up: new adjacency
R3(config-router)# end
```


The neighbor relationship is initiated between R3 and R1 when the missing **network** statement is added to the configuration.

```
R1# show ip route 192.168.3.0
Routing entry for 192.168.3.0/24
  Known via "eigrp 10", distance 90, metric 409600, type internal
  Redistributing via eigrp 10
  Last update from 10.1.1.6 on Ethernet1/1, 00:01:24 ago
  Routing Descriptor Blocks:
  * 10.1.1.6, from 10.1.1.6, 00:01:24 ago, via Ethernet1/1
      Route metric is 409600, traffic share count is 1
      Total delay is 6000 microseconds, minimum bandwidth is 10000 Kbit
      Reliability 255/255, minimum MTU 1500 bytes
      Loading 1/255, Hops 1

R1# show ip route eigrp
<... output omitted ...>
D     192.168.3.0/24 [90/409600] via 10.1.1.6, 00:01:38, Ethernet1/1
<... output omitted ...>
```

Step 5 The network 192.168.4.0/24 does not exist in the routing table of R1. This network is associated with the Loopback0 interface of R4. Try to determine the root cause for this issue.

You might use the following commands on R1 and observe these results:

- **show ip eigrp neighbor**—R4 (10.1.1.10) is not in the EIGRP neighbor table.
- **ping 10.1.1.10**—R4 does not respond to the ping.

These results indicate that there is no connectivity between R1 and R4, leading to problems with EIGRP on top of IP.

If your exploration on R1 indicates a Layer 2 problem, you might start with looking at Layer 2 on R4 with the **show ip interface brief** command. The output shows that the interface Ethernet2/0 is administratively down. "Administratively down" indicates that the interface is shut down in the running configuration. In real-life situations, Layer 2 problems are more likely caused by a cable issue, a faulty interface, or a faulty piece of equipment in the path to the service provider. In the virtual lab environment, administratively shutting down interfaces is the only reliable way to implement a Layer 2 problem.

```
R4# show ip interface brief
Interface                  IP-Address      OK? Method Status
Protocol
<... output omitted ...>
Ethernet2/0                10.1.1.10       YES manual administratively down
down
<... output omitted ...>
```

You can see the root cause can be seen in the configuration.

```
R4# show running-config interface Ethernet2/0
Building configuration...

Current configuration : 101 bytes
!
interface Ethernet2/0
 description Link to R1
 ip address 10.1.1.10 255.255.255.252
 shutdown
end
```

Step 6 With the root cause determined, fix the problem and verify that the route to 192.168.4.0/24 now exists in the routing table of R1.

Enter the following configuration on the R4 router:

```
R4# conf t
R4(config)# interface Ethernet2/0
R4(config-if)# no shutdown
R4(config-if)#
*Oct 14 07:49:23.416: %LINK-3-UPDOWN: Interface Ethernet2/0, changed state to
up
*Oct 14 07:49:24.420: %LINEPROTO-5-UPDOWN: Line protocol on Interface
Ethernet2/0, changed state to up
*Oct 14 07:49:24.431: %DUAL-5-NBRCHANGE: EIGRP-IPv4 10: Neighbor 10.1.1.9
(Ethernet2/0) is up: new adjacency
R4(config-if)# end
```

After the interface is enabled, an EIGRP neighbor relationship gets formed between R4 and R1.

```
R1# show ip route 192.168.4.0
Routing entry for 192.168.4.0/24
  Known via "eigrp 10", distance 90, metric 409600, type internal
  Redistributing via eigrp 10
  Last update from 10.1.1.10 on Ethernet1/2, 00:00:28 ago
  Routing Descriptor Blocks:
  * 10.1.1.10, from 10.1.1.10, 00:00:28 ago, via Ethernet1/2
      Route metric is 409600, traffic share count is 1
      Total delay is 6000 microseconds, minimum bandwidth is 10000 Kbit
      Reliability 255/255, minimum MTU 1500 bytes
      Loading 1/255, Hops 1

R1# show ip route eigrp

<... output omitted ...>
D     192.168.4.0/24 [90/409600] via 10.1.1.10, 00:00:44, Ethernet1/2
  <... output omitted ...>
```

Step 7 The network 192.168.5.0/24 does not exist in the routing table of R1. This network is associated with the Loopback0 interface with R5. Try to determine the root cause for this issue.

You might use the following commands on R1 and observe these results:

- **show ip eigrp neighbor—**R5 (10.1.1.14) is in the EIGRP neighbor table.
- **ping 10.1.1.14—**R5 responds to the ping.

From this situation, apparently EIGRP is working between R1 and R5, so IP and the data link layers must be working as well. Why is R5 not advertising 192.168.5.0/24?

You might use the following commands on R5 and observe these results:

- **show ip eigrp neighbor—R1 (10.1.1.13)** is in the neighbor table.
- **show ip eigrp interfaces—**Only Ethernet1/0 is in the interface table.
- **show ip protocols—**EIGRP is routing only for network 10.0.0.0, not the network of Loopback0 interface.


```
R5# show ip eigrp neighbors
EIGRP-IPv4 Neighbors for AS(10)
H   Address                 Interface              Hold Uptime   SRTT   RTO  Q
Seq
                                                   (sec)         (ms)       Cnt
Num
0   10.1.1.13               Et1/0                    13 21:03:06  421  2526  0
17

R5# show ip eigrp interfaces
EIGRP-IPv4 Interfaces for AS(10)
                              Xmit Queue   PeerQ        Mean   Pacing Time
Multicast    Pending
Interface             Peers  Un/Reliable  Un/Reliable  SRTT   Un/Reliable
Flow Timer   Routes
Et1/0                   1        0/0       0/0         421       0/2
2616           0

R5# show ip protocols
*** IP Routing is NSF aware ***

Routing Protocol is "eigrp 10"
  Outgoing update filter list for all interfaces is not set
  Incoming update filter list for all interfaces is not set
  Default networks flagged in outgoing updates
  Default networks accepted from incoming updates
  EIGRP-IPv4 Protocol for AS(10)
    Metric weight K1=1, K2=0, K3=1, K4=0, K5=0
    NSF-aware route hold timer is 240
    Router-ID: 192.168.5.1
    Topology : 0 (base)
      Active Timer: 3 min
      Distance: internal 90 external 170
      Maximum path: 4
      Maximum hopcount 100
      Maximum metric variance 1

  Automatic Summarization: disabled
  Maximum path: 4
  Routing for Networks:
    10.0.0.0
  Routing Information Sources:
    Gateway         Distance      Last Update
    10.1.1.13             90      00:14:59
  Distance: internal 90 external 170
```

You can see the root cause in the configuration. Including the network 10.0.0.0 will allow EIGRP to run on Ethernet1/0. Hence R1 and R5 are neighbors, and R5 can learn routes from R1. But missing the network 192.168.5.0 prevents R5 from advertising that route to R1.

```
R5# show running-config | section eigrp
router eigrp 10
 network 10.0.0.0
```

Step 8 With the root cause determined, fix the problem and verify that the route to 192.168.5.0/24 now exists in the routing table of R1.

Enter the following configuration on the R5 router:

```
R5# conf t
R5(config)# router eigrp 10
R5(config-router)# network 192.168.5.0
R5(config-router)# end
R5#
```

The neighbor relationship was already functional between R1 and R5, so there was no syslog message to indicate any changes in EIGRP. Before verifying routes on R1, it makes sense to verify that the Loopback0 interface is now included among the EIGRP interfaces on R5.

```
R5# show ip eigrp interfaces
EIGRP-IPv4 Interfaces for AS(10)
                              Xmit Queue   PeerQ        Mean   Pacing Time
Multicast    Pending
Interface             Peers  Un/Reliable  Un/Reliable  SRTT   Un/Reliable
Flow Timer   Routes
Et1/0                   1        0/0       0/0         336       0/2
1684           0
Lo0

R1# show ip route 192.168.5.0
Routing entry for 192.168.5.0/24
  Known via "eigrp 10", distance 90, metric 409600, type internal
  Redistributing via eigrp 10
  Last update from 10.1.1.14 on Ethernet2/0, 00:02:58 ago
  Routing Descriptor Blocks:
  * 10.1.1.14, from 10.1.1.14, 00:02:58 ago, via Ethernet2/0
      Route metric is 409600, traffic share count is 1
      Total delay is 6000 microseconds, minimum bandwidth is 10000 Kbit
      Reliability 255/255, minimum MTU 1500 bytes
      Loading 1/255, Hops 1

R1# show ip route eigrp

<... output omitted ...>
D     192.168.5.0/24 [90/409600] via 10.1.1.14, 00:03:10, Ethernet2/0
<... output omitted ...>
```

Step 9 The network 192.168.6.0/24 does not exist in the routing table of R1. This network is associated with the Loopback0 interface of R6. Try to determine the root cause for this issue.

You might use the following commands on R1 and observe these results:

- **show ip eigrp neighbor**—R6 (10.1.1.18) is not in the EIGRP neighbor table.
- **show ip interface brief**—The interface Ethernet2/1 is "up/up."
- **ping 10.1.1.18**—R6 responds to the ping.

You might use the following commands on R6 and observe these results:

- **show ip eigrp neighbor**—R1 is not in the EIGRP neighbor table of R6.
- **show ip eigrp interfaces**—Only Loopback0 is in the interface table.
- **show ip protocols**—The interface that is linking R6 to R1 (Ethernet1/0) is configured as a passive interface.


```
R6# show ip eigrp neighbors
EIGRP-IPv4 Neighbors for AS(10)

R6# show ip eigrp interfaces
EIGRP-IPv4 Interfaces for AS(10)
                              Xmit Queue   PeerQ        Mean   Pacing Time
Multicast    Pending
Interface             Peers  Un/Reliable  Un/Reliable  SRTT   Un/Reliable
Flow Timer   Routes
Lo0                     0        0/0       0/0           0       0/0
0           0

R6# show ip protocols
*** IP Routing is NSF aware ***

Routing Protocol is "eigrp 10"
  Outgoing update filter list for all interfaces is not set
  Incoming update filter list for all interfaces is not set
  Default networks flagged in outgoing updates
  Default networks accepted from incoming updates
  EIGRP-IPv4 Protocol for AS(10)
    Metric weight K1=1, K2=0, K3=1, K4=0, K5=0
    NSF-aware route hold timer is 240
    Router-ID: 192.168.6.1
    Topology : 0 (base)
      Active Timer: 3 min
      Distance: internal 90 external 170
      Maximum path: 4
      Maximum hopcount 100
      Maximum metric variance 1

  Automatic Summarization: disabled
  Maximum path: 4
  Routing for Networks:
    10.0.0.0
    192.168.6.0
  Passive Interface(s):
    Ethernet1/0
  Routing Information Sources:
    Gateway         Distance      Last Update
  Distance: internal 90 external 170
```

There is a lot of information in the output of the **show ip protocols** command. It is probably the single best command for providing troubleshooting information for dynamic routing protocols. The downside is that it can be difficult to pick out the one piece of inconsistency in the large amount of output. You develop this skill through experience.

You can see the root cause in the configuration.

```
R6# show running-config | section eigrp
router eigrp 10
 network 10.0.0.0
 network 192.168.6.0
 passive-interface Ethernet1/0
```

Step 10 With the root cause determined, fix the problem and verify that the route to 192.168.6.0/24 now exists in the routing table of R1.

Enter the following configuration on the R6 router:

```
R6# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
R6(config)# router eigrp 10
R6(config-router)# no passive-interface Ethernet1/0
R6(config-router)#
*Oct 14 09:12:47.497: %DUAL-5-NBRCHANGE: EIGRP-IPv4 10: Neighbor 10.1.1.17
(Ethernet1/0) is up: new adjacency
R6(config-router)# end
R6#
```

The neighbor relationship between R6 and R1 was initiated when the **passive-interface** statement was removed from the running configuration.

```
R1# show ip route 192.168.6.0
Routing entry for 192.168.6.0/24
  Known via "eigrp 10", distance 90, metric 409600, type internal
  Redistributing via eigrp 10
  Last update from 10.1.1.18 on Ethernet2/1, 00:03:23 ago
  Routing Descriptor Blocks:
  * 10.1.1.18, from 10.1.1.18, 00:03:23 ago, via Ethernet2/1
      Route metric is 409600, traffic share count is 1
      Total delay is 6000 microseconds, minimum bandwidth is 10000 Kbit
      Reliability 255/255, minimum MTU 1500 bytes
      Loading 1/255, Hops 1

R1# show ip route eigrp

<... output omitted ...>
D     192.168.6.0/24 [90/409600] via 10.1.1.18, 00:03:31, Ethernet2/1
<... output omitted ...>
```

Step 11 The EIGRP IPv6 neighbor relationship is not established between R1 and R7. Try to determine the root cause for this issue.

You might use the following commands on R1 and observe these results:

- **show ipv6 eigrp neighbor—**R7 (outgoing interface Ethernet2/2) is not in the EIGRP neighbor table.
- **show ipv6 interface brief—**The interface Ethernet2/2 is "up/up."
- **ping 2001:DB8:0:2::7—**R7 responds to the ping.

These results indicate that there is Layer 2 and Layer 3 connectivity, but there is an issue with EIGRP communication.

On R7, virtually all the show commands that are associated with EIGRP will provide the hint. The AS number 10 is used across the network, but R7 is configured with AS 100.


```
R7# show ipv6 protocols
IPv6 Routing Protocol is "connected"
IPv6 Routing Protocol is "ND"
IPv6 Routing Protocol is "eigrp 100"
EIGRP-IPv6 Protocol for AS(100)
<... output omitted ...>

R7# show ipv6 eigrp interfaces
EIGRP-IPv6 Interfaces for AS(100)
                              Xmit Queue   PeerQ        Mean   Pacing Time
Multicast    Pending
Interface             Peers  Un/Reliable  Un/Reliable  SRTT   Un/Reliable
Flow Timer   Routes
Et1/0                   0        0/0       0/0           0       0/0
0           0
<... output omitted ...>

R7# show ipv6 eigrp neighbors
EIGRP-IPv6 Neighbors for AS(100)
```

You can see the root cause in the configuration.

```
R7# show running-config | section ipv6 eigrp
 ipv6 eigrp 100

R7# show running-config interface Ethernet1/0
!
interface Ethernet1/0
 description Link to R1
 ip address 10.1.1.22 255.255.255.252
 ipv6 address 2001:DB8:0:2::7/64
 ipv6 eigrp 100
end
```

Step 12 With the root cause determined, fix the problem and verify that EIGRP for IPv6 neighbor relationship now exists between R1 and R7 routers.

Enter the following configuration on the R7 router:

```
R7# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
R7(config)# no ipv6 router eigrp 100
R7(config)# ipv6 router eigrp 10
R7(config-rtr)# exit
R7(config)# interface Ethernet1/0
R7(config-if)# no ipv6 eigrp 100
R7(config-if)# ipv6 eigrp 10
*Oct 15 08:10:09.874: %DUAL-5-NBRCHANGE: EIGRP-IPv6 10: Neighbor
FE80::A8BB:CCFF:FE00:2322 (Ethernet1/0) is up: new adjacency
R2(config-if)# end
```

The neighbor adjacency is initiated when EIGRP AS gets changed from 100 to 10 on R7.

```
R1# show ipv6 eigrp neighbors
EIGRP-IPv6 Neighbors for AS(10)
H   Address                 Interface              Hold Uptime   SRTT   RTO  Q
Seq
                                                   (sec)         (ms)       Cnt
Num
0   Link-local address:     Et2/2                    12 00:03:26   14   100  0
2
    FE80::A8BB:CCFF:FE00:2E01
```

Note: The link local IPv6 address may be different in your output.

Step 13 The network 2001:DB8:0:7::/64 still does not exist in the IPv6 routing table of R1. This network is associated with the Loopback0 interface of R7. Try to determine the root cause for this issue.

There is no single best procedure for troubleshooting any network issues. The goal is to isolate the root cause.

You might use the following commands on R7 and observe these results:

- **show ipv6 eigrp interfaces**—Only the interface Ethernet1/0 is included (Loopback0 is missing).
- **show ipv6 protocols**—Only the interface Ethernet1/0 is included (Loopback0 is missing).

```
R7# show ipv6 eigrp interfaces
EIGRP-IPv6 Interfaces for AS(10)
                              Xmit Queue   PeerQ        Mean   Pacing Time
Multicast    Pending
Interface             Peers  Un/Reliable  Un/Reliable  SRTT   Un/Reliable
Flow Timer   Routes
Et1/0                   1        0/0       0/0          8       0/2
50           0
```

```
R7# show ipv6 protocols
IPv6 Routing Protocol is "connected"
IPv6 Routing Protocol is "ND"
IPv6 Routing Protocol is "eigrp 10"
EIGRP-IPv6 Protocol for AS(10)
  Metric weight K1=1, K2=0, K3=1, K4=0, K5=0
  NSF-aware route hold timer is 240
  Router-ID: 192.168.7.1
  Topology : 0 (base)
    Active Timer: 3 min
    Distance: internal 90 external 170
    Maximum path: 16
    Maximum hopcount 100
    Maximum metric variance 1

  Interfaces:
    Ethernet1/0
  Redistribution:
    None
```

You need to add interface Loopback0 to EIGRP IPv6 process.

You can see the root cause in the configuration.


```
R7# show running-config interface Loopback0
interface Loopback0
 description Logical loopback interface
 ip address 192.168.7.1 255.255.255.0
 ipv6 address 2001:DB8:0:7::1/64
end
```

Step 14 With the root cause determined, fix the problem and verify that the route to 2001:DB8:0:7::/64 now exists in the IPv6 routing table of R1.

Enter the following configuration on the R7 router:

```
R7# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
R7(config)# interface Loopback0
R7(config-if)# ipv6 eigrp 10
```

The route to 2001:DB8:0:7::/64 now exists in the IPv6 routing table of R1.

```
R1# show ipv6 route eigrp
IPv6 Routing Table - default - 6 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
D   2001:DB8:0:7::/64 [90/409600]
     via FE80::A8BB:CCFF:FE00:2E01, Ethernet2/2
```

Note: The link local IPv6 address may be different in your output.

This is the end of the discovery lab.

Challenge

1. Which command can determine whether two routers have formed an EIGRP IPv6 neighbor adjacency?

   A. **show ipv6 eigrp interface**
   B. **show ipv6 eigrp neighbor**
   C. **show cdp neighbor**
   D. **show ipv6 eigrp traffic**

2. Refer to the figure. PC1 is unable to communicate with PC2. Which action could correct this problem?

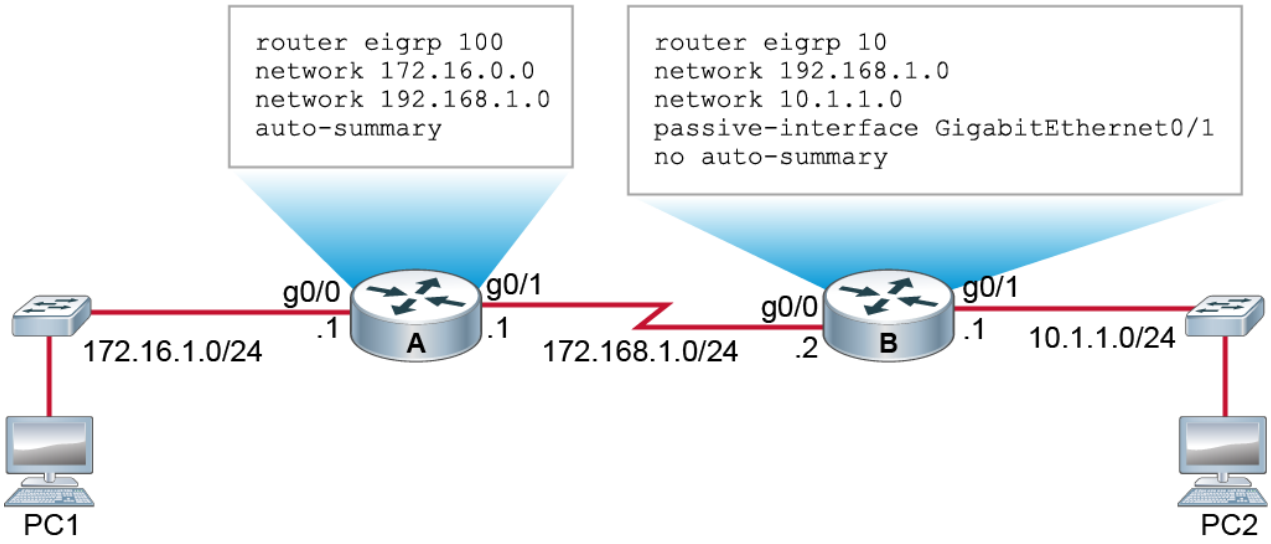

   A. Issue the **no auto-summary** command on router A.
   B. Add the network statement network 172.16.1.0 0.0.0.255 to router A.
   C. Remove the passive-interface command from router B.
   D. Change the AS number on router A to 10.

3. Refer to the figure. You configured EIGRP on your network, and you determined that all routers have formed neighbor adjacencies as expected. However, you still cannot access PC2 from PC1. Which two commands should you use next to troubleshoot this issue? (Choose two.)

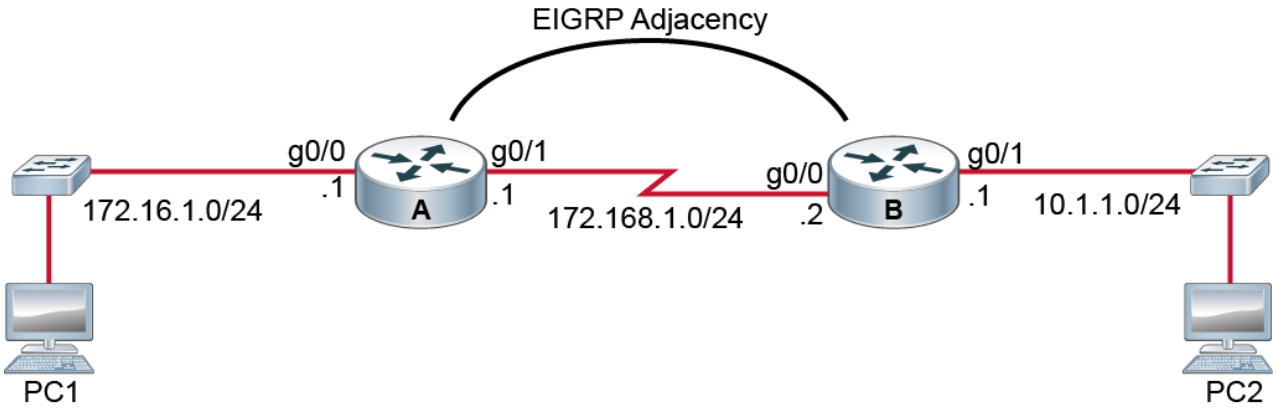

A. **show ip eigrp neighbors**
B. **show ip route**
C. **show ip protocols**
D. **show ip eigrp traffic**

4. Which of the following will not allow EIGRP neighbors to be formed ? (Choose two)

A. K value mismatch
B. Auto summary enabled
C. Access-list denying multicast on interface forming EIGRP neighbors
D. Enabling load-balancing using **variance** command in EIGRP process.

5. Which command should you use to determine the hello and hold timers for the EIGRP neighbors?

A. **Show ip eigrp interface**
B. **Show ip eigrp interface detail**
C. **show ip eigrp neighbor**
D. **show ip eigrp neighbor detail**

6. What could be the reason for the following error message in EIGRP ?

IP-EIGRP: Neighbor ip address not on common subnet for interface

A. The **network** command is misconfigured in EIGRP process.
B. The AS numbers in EIGRP don't match.
C. The interface has been made passive in EIGRP.
D. The IP address has been misconfigured on interfaces.

7. From the show output below, is auto summarization disabled ?

R1# show ip protocols
***** IP Routing is NSF aware *****

Routing Protocol is "eigrp 65010"
Outgoing update filter list for all interfaces is not set
Incoming update filter list for all interfaces is not set
Default networks flagged in outgoing updates
Default networks accepted from incoming updates
EIGRP-IPv4 Protocol for AS(65010)
Metric weight K1=1, K2=0, K3=1, K4=0, K5=0
NSF-aware route hold timer is 240
Router-ID: 172.16.1.1
Topology : 0 (base)
Active Timer: 3 min
Distance: internal 90 external 170
Maximum path: 4
Maximum hopcount 100
Maximum metric variance 1

Automatic Summarization: disabled
Maximum path: 4
Routing for Networks:
Routing Information Sources:
Gateway Distance Last Update
Distance: internal 90 external 170

A. Yes
B. No

Answer Key

Challenge

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

Module 9: Summary Challenge

Introduction

This module challenges you to use the knowledge and skill that you have obtained related in the *previous* modules.

Lesson 1: Troubleshooting a Medium-Sized Network

Introduction

In this lesson, there is a Proof Of Concept lab that has been designed to test your skills on implementation and troubleshooting of a medium-sized network. You need to implement and troubleshoot the issues in the given lab.

Challenge

1. Which of the following commands will you use to check if the interface has been configured as trunk? (Choose two)

   A. **show ip interface brief**
   B. **show vlan**
   C. **show interface** *interface-ID* **switchport**
   D. **show interface trunk**

2. DTP is Cisco proprietary protocol. True or False ?

   A. True
   B. False

3. Which of the following VTP modes allows VLAN configuration changes but does not propagate it to other switches in the VTP domain ?

   A. Transparent
   B. Client
   C. Server
   D. None of the above.

4. The IPv4 access-list that you create end with what is called the "implicit" deny all. Is this statement true or false.

   A. True
   B. False

5. Which of the following is an Etherchannel feature ?

   A. Load sharing across links.
   B. Redundancy.
   C. Higher bandwidth
   D. All the above.

6. You see the following error message when you enable IPv6 EIGRP routing. Which of the following will fix the issue?

   R1(config)# ipv6 router eigrp 1
   % IPv6 routing not enabled
   R1(config)#

   A. Enable IPv6 EIGRP on the interface.
   B. Enable IPv6 CEF.
   C. Enable IPv6 unicast routing first.
   D. It is an IOS-related issue and you need to upgrade the IOS.

7. EIGRP does not support unequal-cost load balancing by default. True or False.

   A. True
   B. False

Answer Key

Challenge

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

Lesson 2: Troubleshooting Scalable Medium-Sized Network

Introduction

You work for DENTIC Networking. Your colleague, Andy did some maintenance on the network over the weekend and now, they are seeing some issues. You need to troubleshoot and resolve the network issues.

Challenge

1. Which of the following table is not used for route selection in EIGRP ?

   A. EIGRP Neighbor Table
   B. EIGRP Topology Table
   C. EIGRP Interface Table
   D. EIGRP Route Table

2. You are assigning VLANs to the ports of a switch. What VLAN number value is assigned to a port by default ?

   A. VLAN 1023
   B. VLAN 99
   C. No VLANs
   D. VLAN 1

3. Which of the following two statements are true about default HSRP configuration? (Choose two)

   A. The Standby priority is 100.
   B. The Standby hello time is 2 seconds.
   C. Two HSRP groups are configured.
   D. Standby group number is 1.
   E. The Standby hold time is 10 seconds

4. Which of the following is true about VLANs in a network ?

   A. End hosts use DHCP to request their VLAN.
   B. End hosts are unaware of any VLANs.
   C. End hosts are assigned to VLANs based on their MAC addresses.
   D. End hosts are all in the same VLAN regardless of which port they attach to.

5. Which two states are the port states when RSTP has converged? (Choose two)

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

6. Which of the following is true about source port used in SPAN ? (Choose two)

   A. It can be a destination port as well for SPAN.
   B. It can be any port type, such as EtherChannel, Fast Ethernet, Gigabit Ethernet and so on.
   C. All source ports can only be configured for ingress traffic.
   D. Source ports can be in the same or different VLANs.

7. Which three components are combined to form STP bridge ID?

   A. Bridge Priority
   B. MAC address
   C. Port Cost
   D. Extended System ID
   E. Port ID

Answer Key

Challenge

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

Module 10: Implementing a Scalable OSPF-Based Solution

Introduction

This module examines OSPF, which is one of the most commonly used IGPs in IP networking. OSPF is a complex protocol, and therefore configuration and verification of OSPF on a Cisco router is a primary learning objective.

The module discusses the primary configuration commands for a multiarea OSPF and explains the benefits of a multiarea OSPF solution compared to a single-area solution. Specifically, it covers link-state protocols, OSPF components, the OSPF metric, the way in which OSPF operates, and how to configure multiarea OSPF. Several OSPF show commands are also described in this module for verification purposes.

Lesson 1: Understanding OSPF

Introduction

Your first deployment at CCS was at the Small Law Firm (in further text—SLF). Since the multinational law firm Big Law Firm (in further text—BLF) purchased SLF, it has standardized on OSPF as its IGP at the corporate office and at all branches. You and your team leader Bob successfully implemented OSPF at SLF. BLF has been very satisfied with the results that CCS has delivered. BLF corporate has decided to award CCS the contract to provide network infrastructure, management, and security services for all BLF corporate and its branch offices.

The strong growth at BLF has put a strain on the existing flat network design. You and Bob have decided to implement a hierarchical design to optimize routing using multiarea OSPF. You will be deployed soon to implement the change, but before you go, you should have a solid understanding of OSPF functions, packet types, and the LSDB. As before, you can take the training to gain the OSPF knowledge or take the test instead.

Introduction to Link-State Routing Protocol

The two basic types of routing protocols are distance vector and link state. OSPF is an example of a link-state routing protocol.

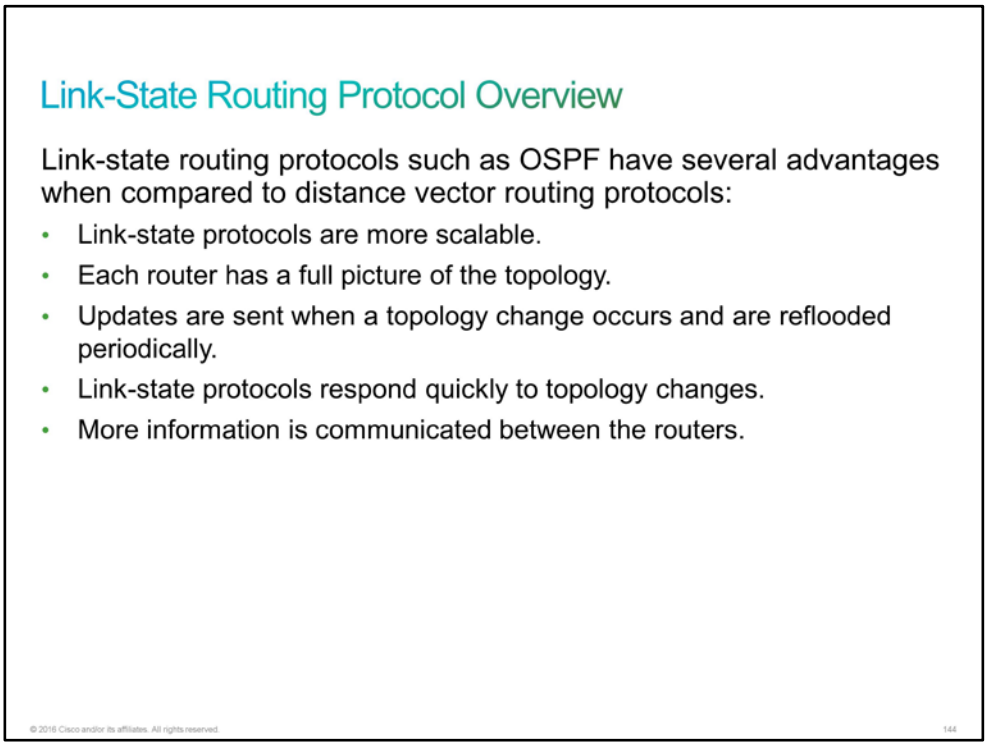

When a failure occurs in a network, routing protocols should detect the failure as soon as possible and find another path across the network. Only link-state protocols support fast convergence with support for scalability and multivendor environments, so they are the only type of IGP that is found in large network environments.

Link-state protocols have the following advantages when compared to distance vector routing protocols:

- **They are more scalable:** Link-state protocols use a hierarchical design and can scale to very large networks, if properly designed.
- **Each router has a full picture of the topology:** Because each router contains full information about the routers and links in a network, each router is able to independently select a loop-free and efficient pathway, which is based on cost, to reach every neighbor in the network.
- **Updates are sent when a topology change occurs and are reflooded periodically:** Link-state protocols send updates of a topology change by using triggered updates. Also, updates are made periodically—by default every 30 minutes.
- **They respond quickly to topology changes:** Link-state protocols establish neighbor relations with the adjacent routers. The failure of a neighbor is detected quickly, and this failure is communicated by using triggered updates to all routers in the network. This immediate reporting generally leads to fast convergence times.
- **More information is communicated between routers:** Routers that are running a link-state protocol have a common view on the network. This means that each router has full information about other routers and links between them, including the metric on each link.

Link-State Routing Protocol Data Structures

A router that is running a link-state routing protocol must first recognize other routers and establish a neighbor adjacency with its neighboring routers. A router achieves this neighbor adjacency by exchanging hello packets with the neighboring routers. After a router establishes a neighbor adjacency by using the hello packets, a neighbor is put into the neighbor database.

In the example, router A recognizes routers B and D as neighbors.

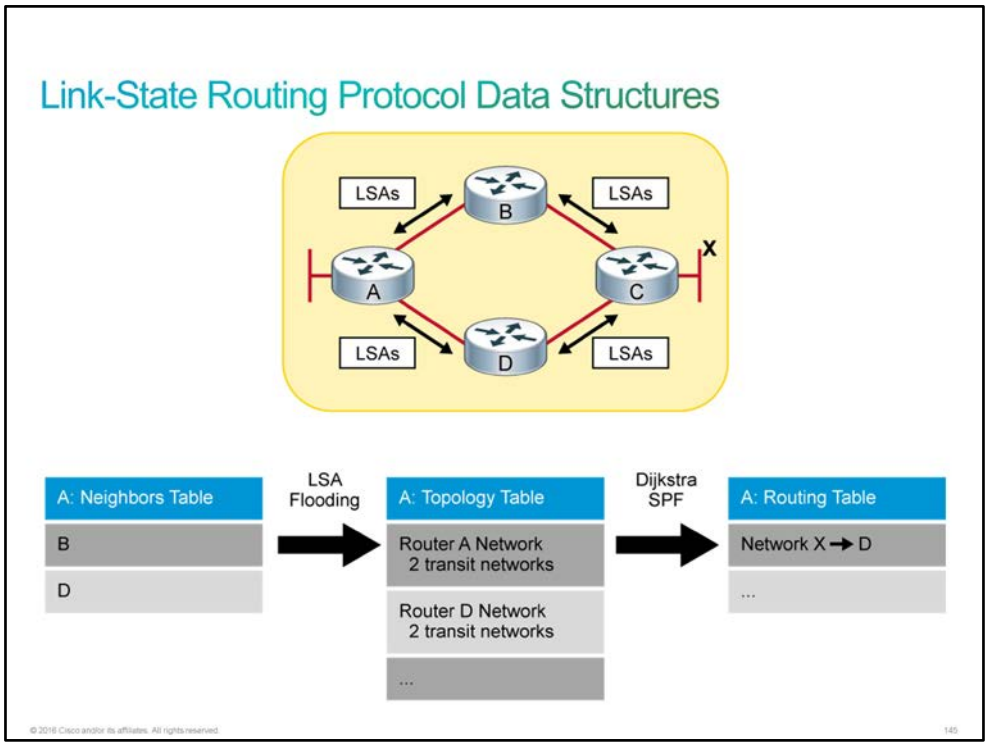

After a neighbor relationship is established between routers, the routers synchronize their topology databases (LSDBs) by reliably exchanging LSAs. An LSA describes a router and networks that are connected to this router. LSAs are stored in the LSDB. By exchanging all LSAs, routers learn the complete topology of the network. Each router will have the same topology database within an area.

After the topology database is built, each router applies the SPF algorithm to the topology map. The SPF algorithm uses the Dijkstra algorithm to calculate the shortest path to each destination.

The best (shortest) paths to destinations are then put into the routing table. The routing table includes a destination network and the next-hop IP address. In the example, the routing table on router A states that a packet should be sent to router D to reach network X.

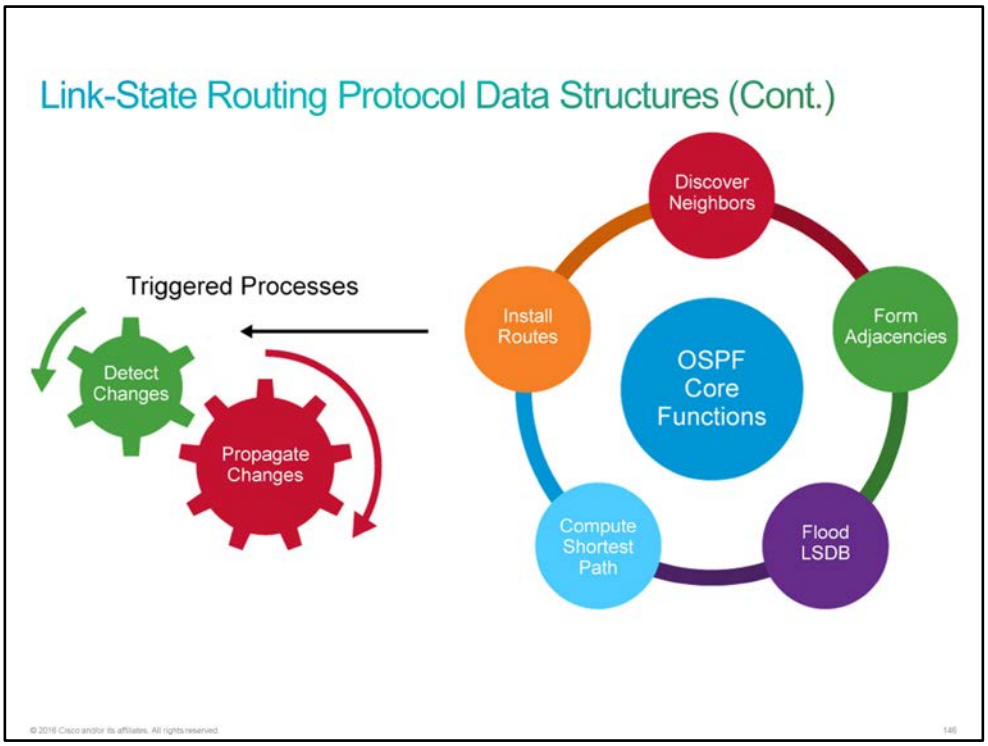

Whenever there is a change in a topology, new LSAs are created and sent throughout the network. All routers change their LSDB when they receive the new LSA, and the SPF algorithm is run again on the updated LSDB to verify new paths to destinations.

Introducing OSPF

OSPF is a link-state routing protocol. You can think of a link as an interface on a router. The state of the link is a description of that interface and of its relationship to its neighboring routers. A description of the interface would include, for example, the IP address of the interface, the subnet mask, the type of network to which it is connected, the routers that are connected to that network, and so on. The collection of all these link states forms a link-state database.

Note	OSPF was developed based on an open standard and is supported by several router manufacturers. OSPF is widely used as an IGP, especially in large network environments. OSPF was developed as a replacement for the distance vector routing protocol RIP. The major advantages of OSPF over RIP are its fast convergence and its ability to scale to much larger networks.

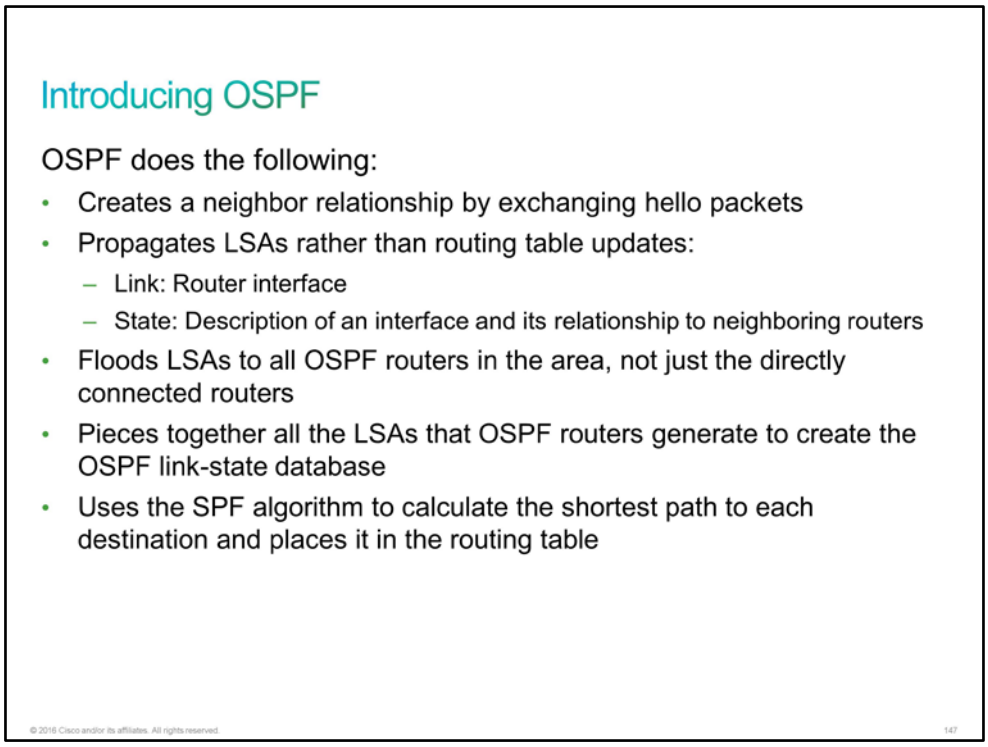

A router sends LSA packets immediately to advertise its state when there are state changes. The router sends the packets periodically as well (every 30 minutes by default). The information about the attached interfaces, the metrics that are used, and other variables are included in OSPF LSAs. As OSPF routers accumulate link-state information, they use the SPF algorithm to calculate the shortest path to each node.

A topological (link-state) database is, essentially, an overall picture of the networks in relation to the routers. The topological database contains the collection of LSAs that all routers in the same area sent. Because the routers within the same area share the same information, they have identical topological databases.

OSPF can operate within a hierarchy. The largest entity within the hierarchy is the AS, which is a collection of networks under a common administration that share a common routing strategy. An AS can be divided into several areas, which are groups of contiguous networks and attached hosts. The figure shows an example of an OSPF hierarchy.

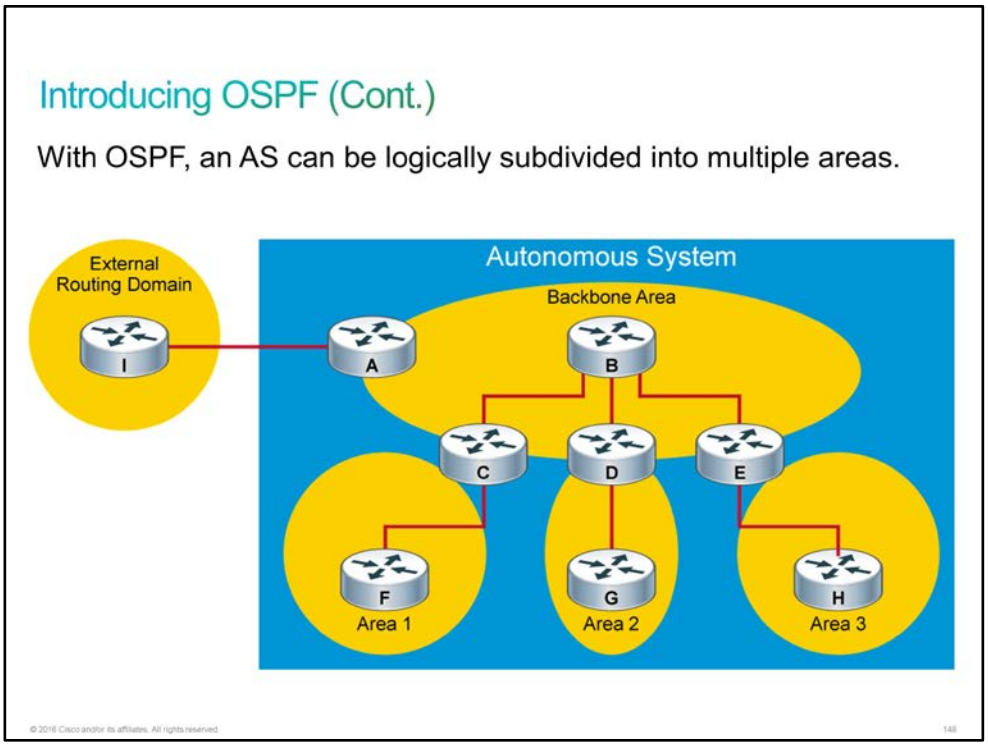

OSPF uses a two-layer network hierarchy that has two primary elements:

- **AS:** An AS consists of a collection of networks under a common administration that share a common routing strategy. An AS, which is sometimes called a *domain*, can be logically subdivided into multiple areas.
- **Area:** An *area* is a grouping of contiguous networks. Areas are logical subdivisions of the AS.

Within each AS, a contiguous backbone area must be defined. In the multiarea design, all other nonbackbone areas are connected off the backbone area.

Multiarea design is more effective since the network is segmented to limit the propagation of LSAs inside an area. It is especially useful for large networks.

Establishing OSPF Neighbor Adjacencies

Neighbor OSPF routers must recognize each other on the network before they can share information because OSPF routing depends on the status of the link between two routers. This process is done using the Hello protocol. OSPF routers send hello packets on all OSPF-enabled interfaces to determine if there are any neighbors on those links.

The Hello protocol establishes and maintains neighbor relationships by ensuring bidirectional (two-way) communication between neighbors.

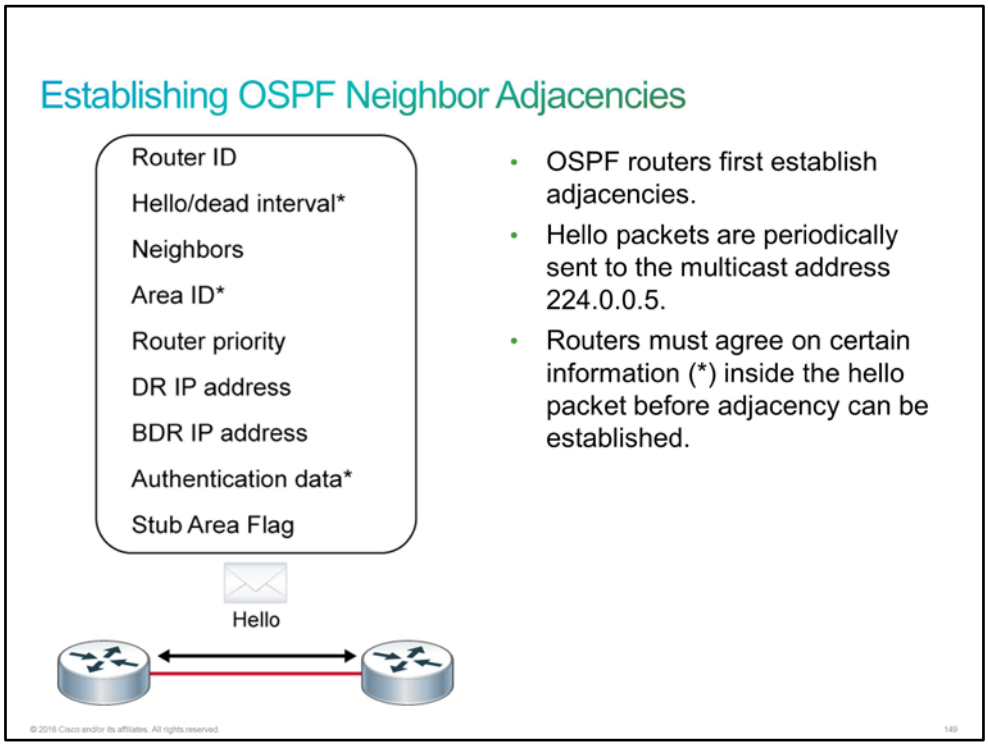

An OSPF neighbor relationship, or adjacency, is formed between two routers if they both agree on the area ID, hello and dead intervals, and authentication. Of course, the routers must be on the same IP subnet. Bidirectional communication occurs when a router recognizes itself in the neighbors list that is held in the hello packet that it receives from a neighbor.

Each interface that is participating in OSPF uses the multicast address 224.0.0.5 to periodically send hello packets. A hello packet contains the following information:

- **Router ID**: The router ID is a 32-bit number that uniquely identifies the router. The router ID is, by default, the highest IP address on a loopback interface, if configured. If the router ID is not configured, it is the highest IP address on any interface. You can also manually configure the router ID using the **router-id** command. It is recommended that you always use a loopback IP address for the router ID or to set the router ID manually. In this way, the router ID is stable and will not change.
- **Hello and dead intervals**: The hello interval specifies the frequency in seconds at which a router sends hello packets. The default hello interval on multiaccess networks is 10 seconds. The dead interval is the time in seconds that a router waits to hear from a neighbor before declaring the neighboring router out of service. By default, the dead interval is four times the hello interval. These timers must be the same on neighboring routers; otherwise, an adjacency will not be established.

- **Neighbors**: The Neighbors field lists the adjacent routers with an established bidirectional communication. This bidirectional communication is indicated when the router recognizes itself as it is listed in the Neighbors field of the hello packet from the neighbor.
- **Area ID**: To communicate, two routers must share a common segment and their interfaces must belong to the same OSPF area on this segment. The neighbors must also share the same subnet and mask. These routers in the same area will all have the same link-state information for that area.
- **Router priority**: The router priority is an 8-bit number that indicates the priority of a router. OSPF uses the priority to select a DR and BDR. In certain types of networks, OSPF elects DRs and BDRs. The DR acts as a hub to reduce traffic between routers.
- **DR and BDR IP addresses**: These addresses are the IP addresses of the DR and BDR for the specific network, if they are known.
- **Authentication data:** If router authentication is enabled, two routers must exchange the same authentication data. Authentication is not required, but if it is enabled, all peer routers must have the same key configured.
- **Stub area flag:** A stub area is a special area. Designating a stub area is a technique that reduces routing updates by replacing them with a default route. Two routers have to also agree on the stub area flag in the hello packets in order to become neighbors.

OSPF Neighbor States

When routers that are running OSPF are initialized, an exchange process using the Hello protocol is the first procedure.

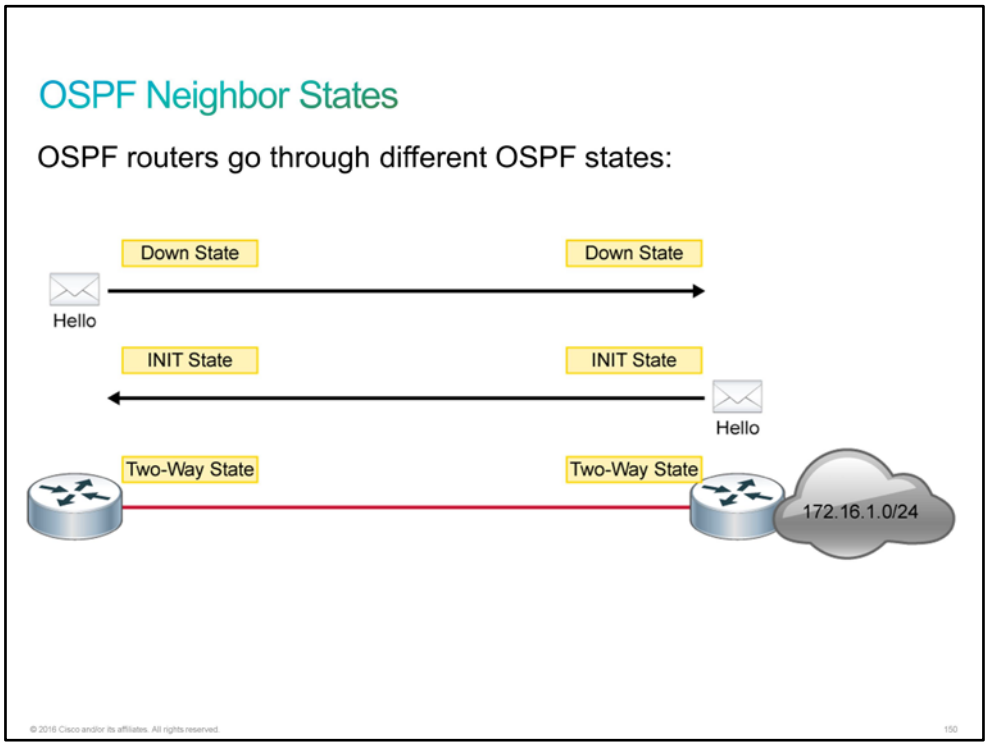

The figure illustrates the exchange process that happens when routers appear on the network:

1. A router is enabled on the LAN and is in a down state because it has not exchanged information with any other router. The router begins by sending a hello packet through each of its interfaces that are participating in OSPF, although it does not know the identity of any other routers.
2. All directly connected routers that are running OSPF receive the hello packet from the first router and add the router to their lists of neighbors. After adding the router to the list, other routers are in the INIT state.
3. Each router that received the hello packet sends a unicast reply hello packet to the first router with its corresponding information. The neighbor field in the hello packet includes all neighboring routers and the first router.
4. When the first router receives the hello packets, it adds all the routers that had its router ID in their hello packets to its own neighbor relationship database. After this process, the first router is in the two-way state. At this point, all routers that have each other in their lists of neighbors have established a bidirectional communication.

If the link type is a broadcast network (for example, a LAN link like Ethernet), a DR and BDR must first be selected. The DR acts as a central exchange point for routing information and reduces the amount of routing information that the routers have to exchange. The DR and BDR are selected after routers are in the two-way state. The router with the highest priority will become the DR. If there is a tie, the router with the highest router ID will become the DR. Among the routers on a LAN that were not elected as the DR or BDR, the exchange process stops at this point, and the routers remain in the two-way state. Routers then communicate only with DR (or BDR) router using multicast IP address 224.0.0.6. The DR router uses 224.0.0.5 multicast IP address to communicate with all other non-DR routers.

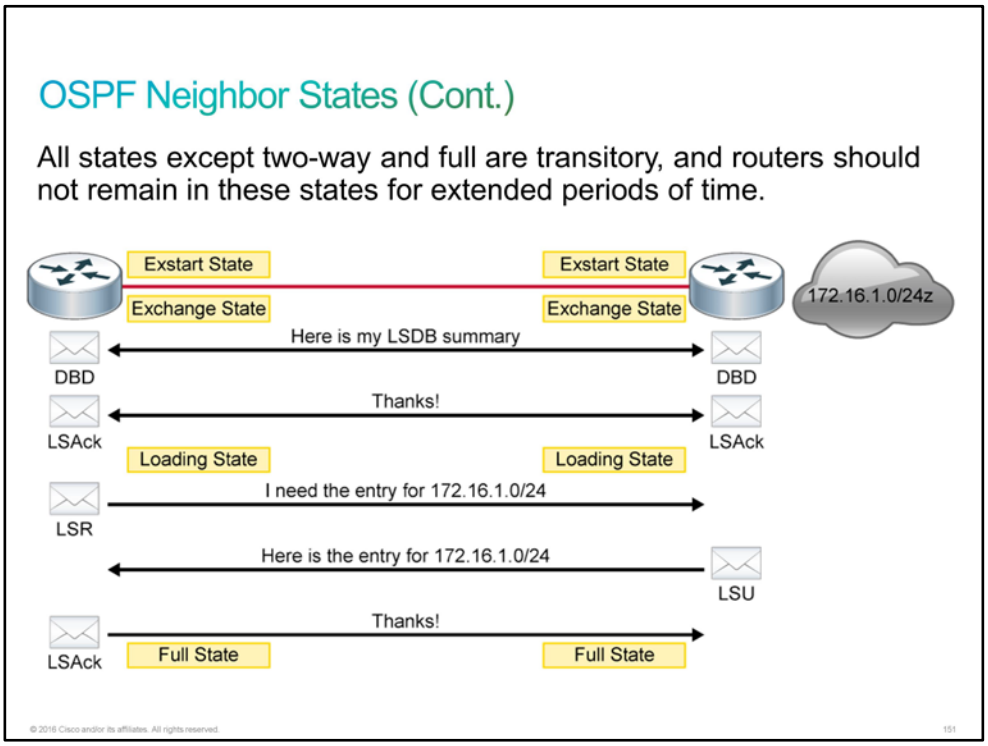

After the DR and BDR have been selected, the routers are considered to be in the exstart state. The routers are then ready to discover the link-state information about the internetwork and create their LSDBs. The exchange protocol is used to discover the network routes, and it brings all the routers from the exchange state to a full state of communication. The first step in this process is for the DR and BDR to establish adjacencies with each of the other routers.

As shown in the figure, the exchange protocol continues as follows:

1. In the exstart state, the DR and BDR establish adjacencies with each router in the network. During this process, a primary-subordinate relationship is created between each router and its adjacent DR and BDR. The router with the higher router ID acts as the primary during the exchange process. The primary-subordinate election dictates which router will start the exchange of routing information. This step is not shown in the figure.
2. The primary and subordinate routers exchange one or more DBD packets. The routers are in the exchange state.
3. A router compares the DBD that it received with the LSAs that it has. If the DBD has a more up-to-date link-state entry, the router sends an LSR to the other router. When routers start sending LSRs, they are in the loading state.
4. When all LSRs have been satisfied for a given router, the adjacent routers are considered synchronized and are in the full state.

You should be aware that all states except two-way and full are transitory, and routers should not remain in these states for extended periods of time.

SPF Algorithm

The SPF algorithm places each router at the root of a tree and calculates the shortest path to each node. The path calculation is based on the cumulative cost that is required to reach that destination. LSAs are flooded throughout the area by using a reliable algorithm, which ensures that all the routers in an area have the same topological database. Each router uses the information in its topological database to calculate a shortest path tree, with itself as the root. The router then uses this tree to route network traffic.

The figure represents the R1 view of the network, where R1 is the root and calculates the pathways by assuming this view.

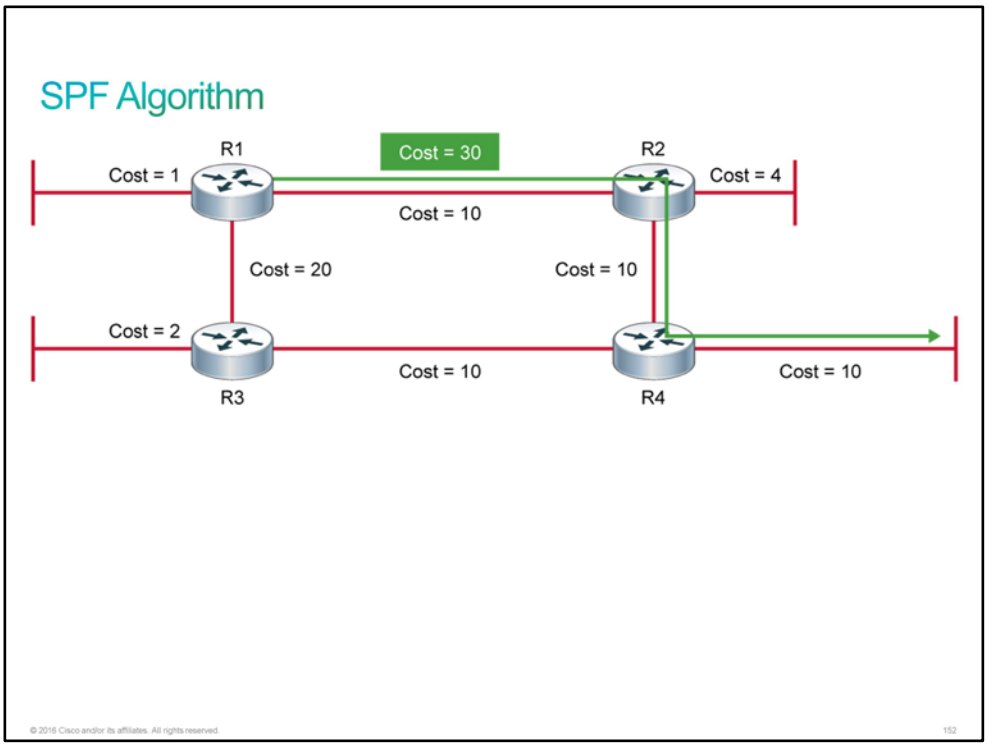

Each router has its own view of the topology, even though all the routers build the shortest path trees by using the same link-state database.

A metric is an indication of the overhead that is required to send packets across a certain interface. OSPF uses cost as a metric. A smaller cost indicates a better path than a higher cost. By default on Cisco devices, the cost of an interface is inversely proportional to the bandwidth of this interface, so a higher bandwidth indicates a lower cost. There is more overhead, a higher cost, and more time delays that are involved in crossing a 10-Mbps Ethernet line than in crossing a 100-Mbps Ethernet line.

The formula that you use to calculate OSPF cost is **cost = reference bandwidth / interface bandwidth (in bits per second)**.

The default reference bandwidth is 10^8, which is 100,000,000 or the equivalent of the bandwidth of FastEthernet. Therefore, the default cost of a 10-Mbps Ethernet link will be $10^8 / 10^7 = 10$, and the cost of a 100-Mbps link will be $10^8 / 10^8 = 1$. The problem arises with links that are faster than 100 Mbps. Because the OSPF cost has to be an integer, all links that are faster than FastEthernet will have an OSPF cost of 1.

The cost to reach a distant network from a router is the cumulative cost of all links on the path from the router to the network. In the example, the cost from router R1 to the destination network via R3 is 40 (20 + 10 + 10), and the cost via router R2 is 30 (10 + 10 + 10). The path via R2 is better because it has a lower cost.

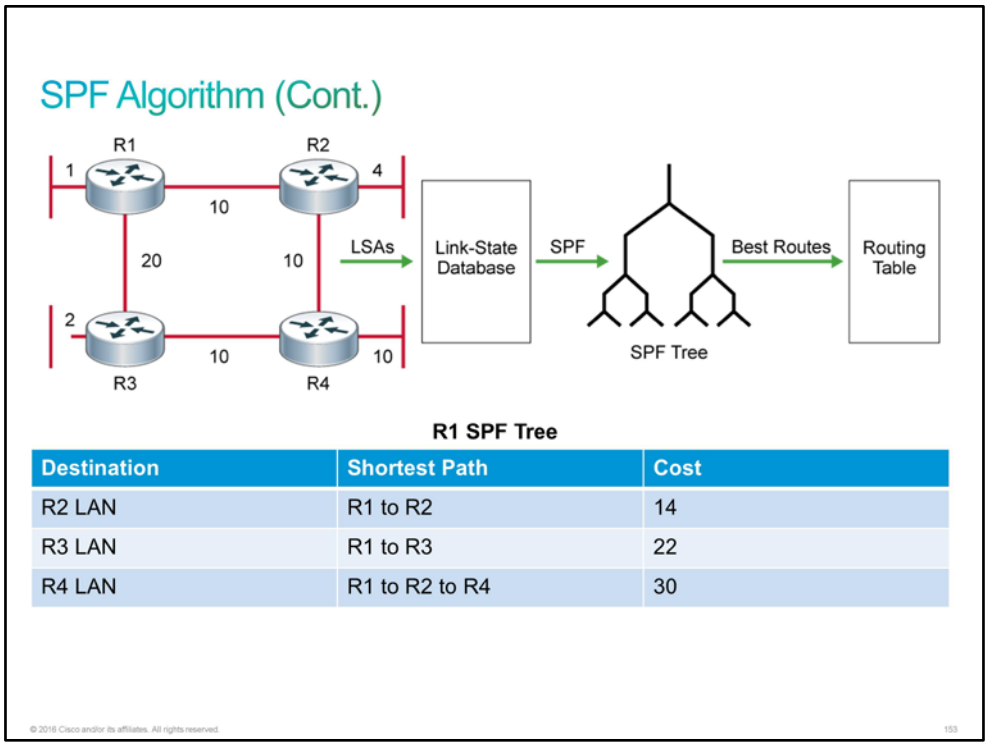

LSAs are flooded through the area by using a reliable algorithm, which ensures that all routers in an area have the same topological database. Because of the flooding process, R1 has learned the link-state information for each router in its routing area. Each router uses the information in its topological database to calculate a shortest path tree, with itself as the root. The tree is then used to populate the IP routing table with the best paths to each network.

For R1, the shortest path to each LAN and its cost are shown in the table. The shortest path is not necessarily the best path. Each router has its own view of the topology, even though the routers build shortest path trees by using the same link-state database.

Building a Link-State Database

When two routers discover each other and establish adjacency using hello packets, they use the exchange protocol to exchange information about the LSAs.

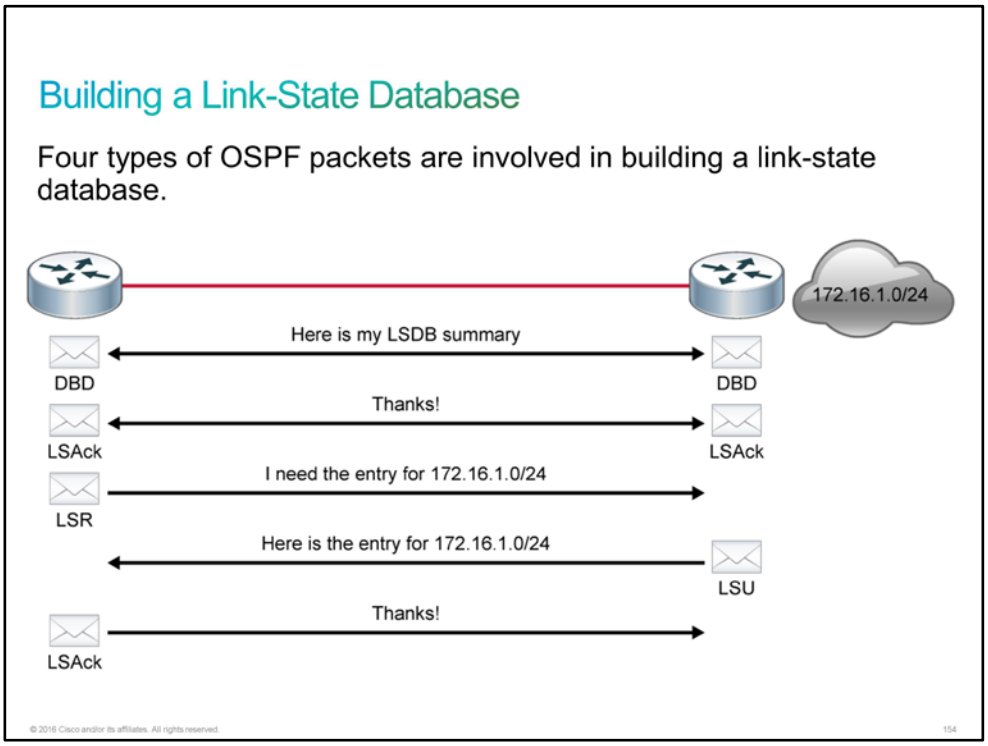

As shown in the figure, the exchange protocol operates as follows:

1. The routers exchange one or more DBD packets. A DBD includes information about the LSA entry header that appears in the LSDB of the router. Each LSA entry header includes information about the link-state type, the address of the advertising router, the cost of the link, and the sequence number. The router uses the sequence number to determine the "newness" of the received link-state information.
2. When the router receives the DBD, it acknowledges the receipt of the DBD that is using the LSAck packet.
3. The routers compare the information that they receive with the information that they have. If the received DBD has a more up-to-date link-state entry, the router sends an LSR to the other router to request the updated link-state entry.
4. The other router responds with complete information about the requested entry in an LSU packet. The other router adds the new link-state entries to its LSDB.
5. When the router receives an LSU, it sends an LSAck.

Four types of update packets are used when building and synchronizing LSDBs:

- **DBD packet**: A DBD packet is used to describe the network routes of each neighbor.
- **LSR packet**: After DBD packets are exchanged, the routers request the missing information by using LSR packets.

- **LSU packet**: All missing information is sent to the neighbors by sending LSU packets that contain different LSAs.
- **LSAck packet**: Every packet receives an LSAck to ensure a reliable transport and a reliable exchange of information.

OSPF Packet Types

OSPF uses five types of routing protocol packets that share a common protocol header. The Protocol field in the IP header is set to 89. All five packet types are used in a normal operation of OSPF.

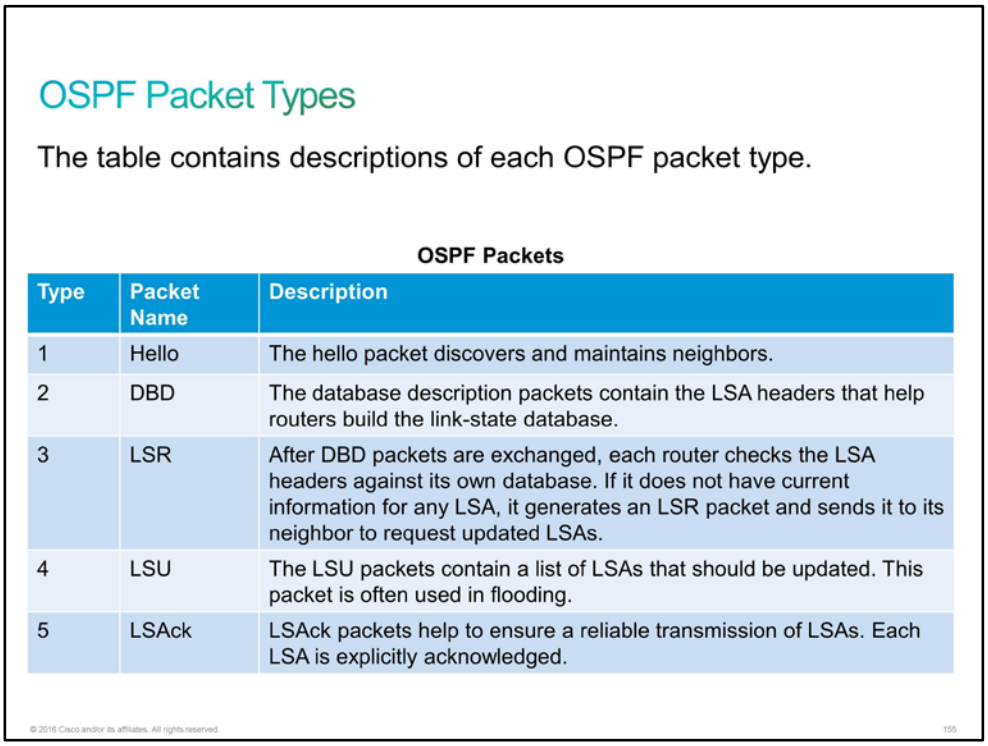

OSPF Packet Types

The table contains descriptions of each OSPF packet type.

OSPF Packets

Type	Packet Name	Description
1	Hello	The hello packet discovers and maintains neighbors.
2	DBD	The database description packets contain the LSA headers that help routers build the link-state database.
3	LSR	After DBD packets are exchanged, each router checks the LSA headers against its own database. If it does not have current information for any LSA, it generates an LSR packet and sends it to its neighbor to request updated LSAs.
4	LSU	The LSU packets contain a list of LSAs that should be updated. This packet is often used in flooding.
5	LSAck	LSAck packets help to ensure a reliable transmission of LSAs. Each LSA is explicitly acknowledged.

OSPF Packet Header Format

All five OSPF packet types are encapsulated directly into an IP payload, as shown in the figure. OSPF packets do not use TCP or UDP. OSPF requires a reliable packet transport, but because it does not use TCP, OSPF defines an acknowledgment packet (OSPF packet type 5) to ensure reliability.

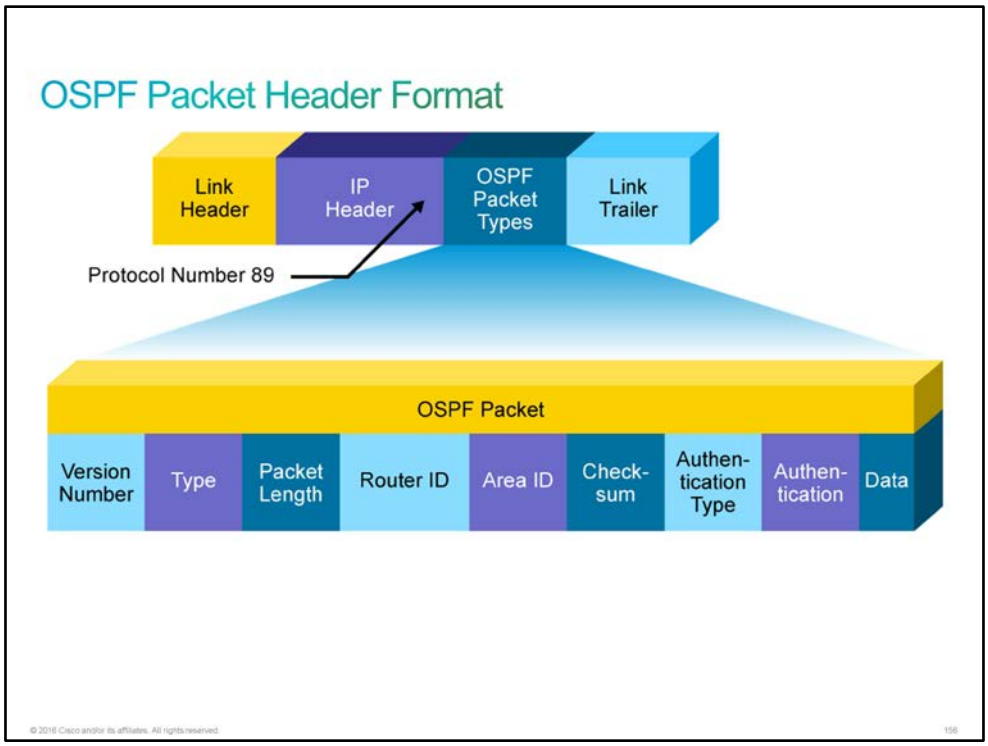

In the Protocol field of the IP header, the value of 89 is set for all OSPF packet types. Each of the five OSPF packet types begins with the same header format. This header includes the following fields:

- **Version number:** Version 2 for OSPF with IPv4 and version 3 for OSPF with IPv6
- **Type:** Differentiates the five OSPF packet types
- **Packet length:** The length of the OSPF packet in bytes
- **Router ID:** Defines which router is the source of the packet
- **Area ID:** Defines the area where the packet originated
- **Checksum:** Used for packet-header error detection to ensure that the OSPF packet was not corrupted during transmission
- **Authentication type:** An OSPF option that describes either the *no authentication*, *cleartext passwords*, or *passwords protected by an* MD5 *hash* formats for router authentication
- **Authentication:** Used in the authentication scheme
- **Data:** Each of the five packet types includes different data:
  - **Hello packets:** Contains a list of known neighbors
  - **DBD packet:** Contains a summary of the LSDB, which includes all known router IDs and their last sequence numbers, among several other fields
  - **LSR packet:** Contains the type of LSU that is needed and the router ID that has the needed LSU

- **LSU packet:** Contains the complete LSA entries. Multiple LSA entries can fit in one OSPF update packet.
- **LSAck packet:** Empty

Discovery 43: Configure and Verify Single-Area OSPF

Introduction

This discovery will guide you through the configuration and verification of OSPF for IPv4 on a Cisco IOS router. The 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. R2 and R3 are also configured with OSPF. You will configure OSPF on R1 and verify the results.

Topology

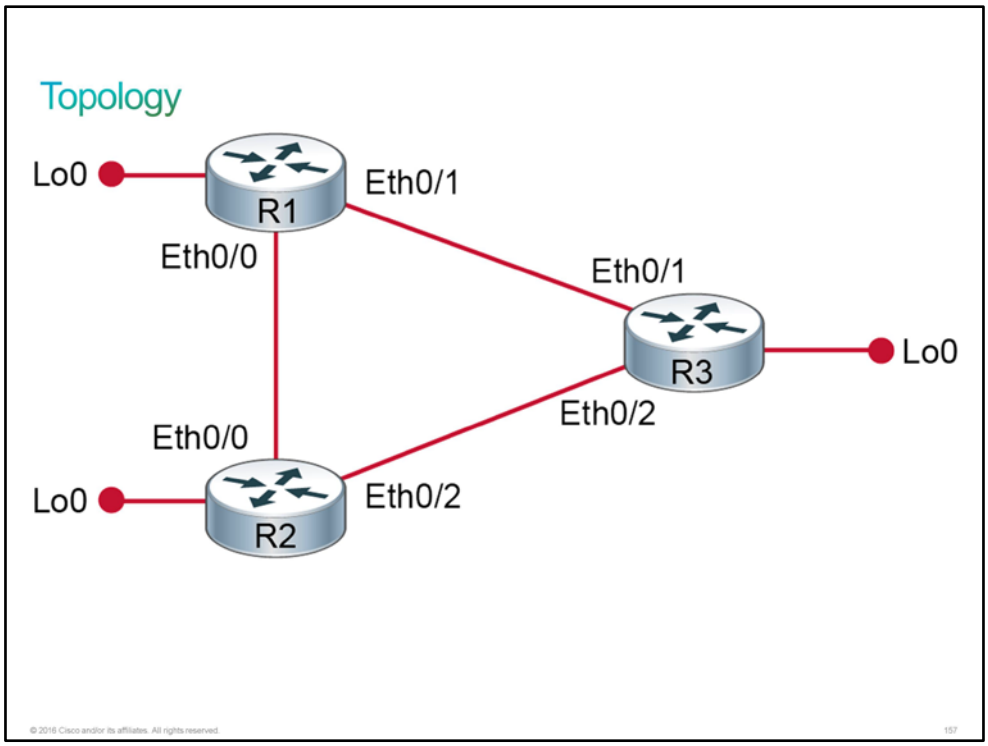

Job Aid

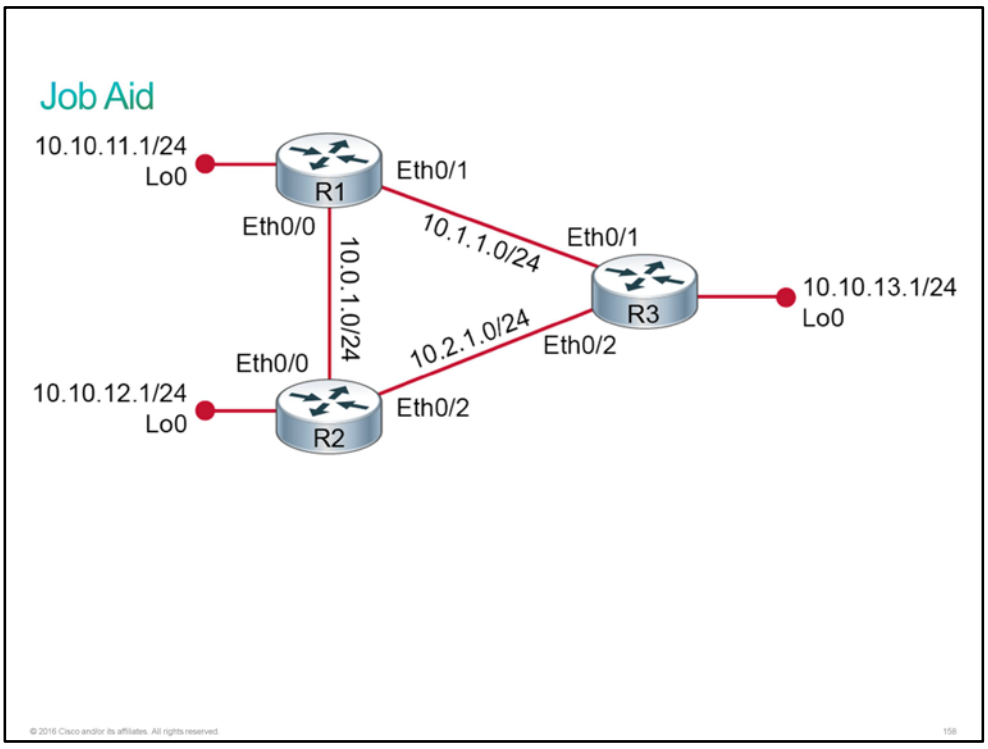

The configuration is as follows:

- All devices have their basic configurations in place, including hostnames and IP addresses.
- OSPF is preconfigured on R2 and R3:
  - AS number 1 is used.
  - Both routers are announcing Loopback interface network.

Device Information

Device Details

Device	Interface	IP Address	Neighbor
R1	Ethernet0/0	10.0.1.1/24	R2
R1	Ethernet0/1	10.1.1.1/24	R3
R1	Loopback0	10.10.11.1/24	—
R2	Ethernet0/0	10.1.1.2/24	R1
R2	Ethernet0/2	10.2.1.2/24	R3
R2	Loopback0	10.10.12.1/24	—

Device	Interface	IP Address	Neighbor
R3	Ethernet0/1	10.1.1.3/24	R1
R3	Ethernet0/2	10.2.1.3/24	R2
R3	Loopback0	10.10.13.1/24	—

Task 1: Configure and Verify Single-Area OSPF

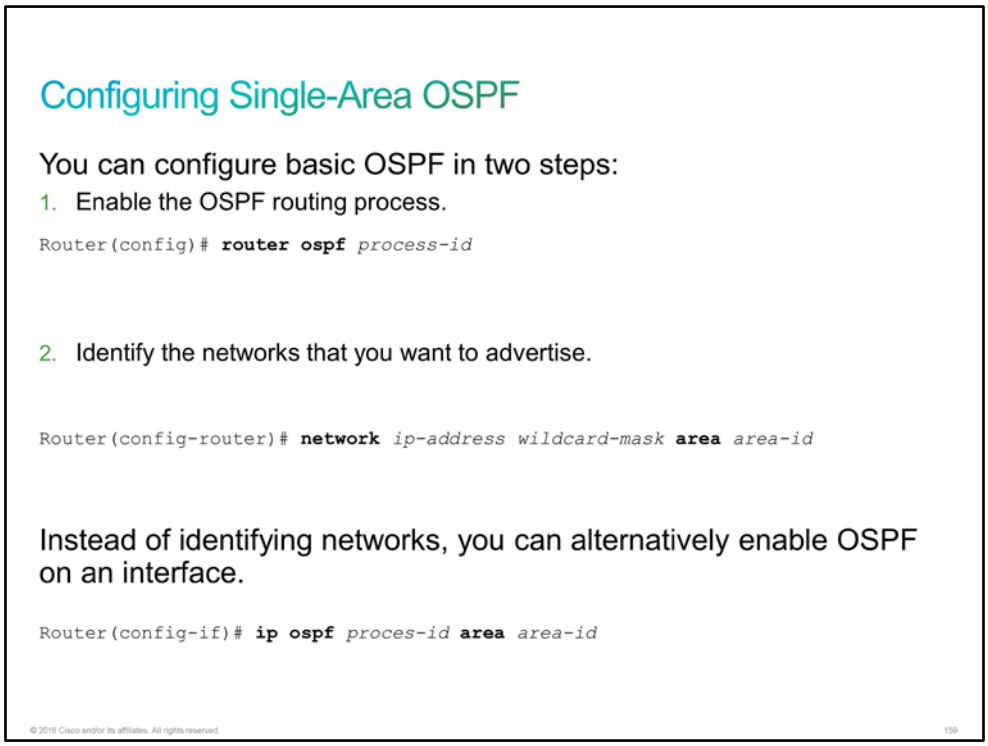

The **router ospf** command uses a process identifier as an argument. The process ID is a unique, arbitrary number that you select to identify the routing process. The process ID is locally significant and does not need to match the OSPF process ID on other OSPF routers.

The **network** command identifies which IP networks on the router are part of the OSPF network. For each network, you must also identify the OSPF area to which the networks belong. The network that is identified in the **network** command does not tell the router which network to advertise; instead, it indicates the interfaces on which OSPF will be enabled.

The table defines the commands that you use to configure OSPF.

Command and Variable	Description
router ospf *process_id*	Enters into the OSPF routing configuration mode. The network administrator chooses the process ID, which is a number between 1 and 65,535. The process ID is locally significant, which means that it does not have to match other OSPF routers to establish adjacencies with those neighbors.
network *ip-address wildcard_mask* **area** *area_id*	Uses a combination of the network address and wildcard mask and serves as the criteria to match when identifying the interfaces that can send and receive OSPF packets. The network address, along with the wildcard mask, identifies which IP networks are part of the OSPF network and are included in OSPF routing updates. The area ID identifies the OSPF area to which the network belongs. When all the routers are within the same OSPF area, the **network** commands must be configured with the same area ID on all routers. Even if no areas are specified, there must be an area 0. In a single-area OSPF environment, the area is always 0.
ip ospf *process_id* **area** *area_id*	As an alternative to a **network** command, you can use this interface configuration mode command that enables OSPF explicitly on the selected interface.

To be able to perform routing toward external networks or toward the Internet, the router must either know all the destination networks or have a default route. You can statically configure a default route, but it can also be learned dynamically via the OSPF routing protocol. The router that announces the default route needs to be configured with the **default-information originate** command in the routing protocols configuration mode. You can also add the **always** keyword at the end of the command (**default-information originate always**) to always advertise the default route regardless of whether the route table has a default route.

Activity

Complete the following steps:

Step 1 Access the console of R2 and display the OSPF configuration.

You can verify the OSPF configuration using the **show running-config** command.

```
R2# show running-config | section ospf
router ospf 1
 router-id 2.2.2.2
 network 10.0.1.0 0.0.0.255 area 0
 network 10.2.1.0 0.0.0.255 area 0
 network 10.10.12.0 0.0.0.255 area 0
```

You should see that the OSPF with the process ID 1 is preconfigured with the following networks:

- 10.0.1.0/24
- 10.2.1.0/24
- 10.10.12.0/24

If you refer to the Job Aids section, you can quickly determine that the configured OSPF networks are associated with each of the active interfaces on R2. All networks, meaning all active interfaces on the router, belong to the same area—area 0.

At this point, note that the router is configured with the router ID 2.2.2.2. The router ID will be discussed later on.

Step 2 Access the console of R3 and display the OSPF configuration.

Another way to verify the OSPF configuration is by using the **show ip protocols** command. This command will display the status of the configured dynamic routing protocols.

```
R3# show ip protocols
*** IP Routing is NSF aware ***

Routing Protocol is "ospf 1"
  Outgoing update filter list for all interfaces is not set
  Incoming update filter list for all interfaces is not set
  Router ID 3.3.3.3
  Number of areas in this router is 1. 1 normal 0 stub 0 nssa
  Maximum path: 4
  Routing for Networks:
    10.1.1.0 0.0.0.255 area 0
    10.2.1.0 0.0.0.255 area 0
    10.10.13.0 0.0.0.255 area 0
  Routing Information Sources:
    Gateway         Distance      Last Update
    2.2.2.2              110      17:04:48
  Distance: (default is 110)
```

You should see that only the OSPF routing protocol is configured on R3. OSPF is using the process ID 1 and is preconfigured with the following networks:

- 10.1.1.0/24
- 10.2.1.0/24
- 10.10.13.0/24

If you refer to the Job Aids section, you can quickly determine that OSPF on R3 is enabled on all enabled interfaces. All networks, meaning all interfaces on the router, belong to the same area—area 0.

The Router is preconfigured with the router ID 3.3.3.3.

Router ID

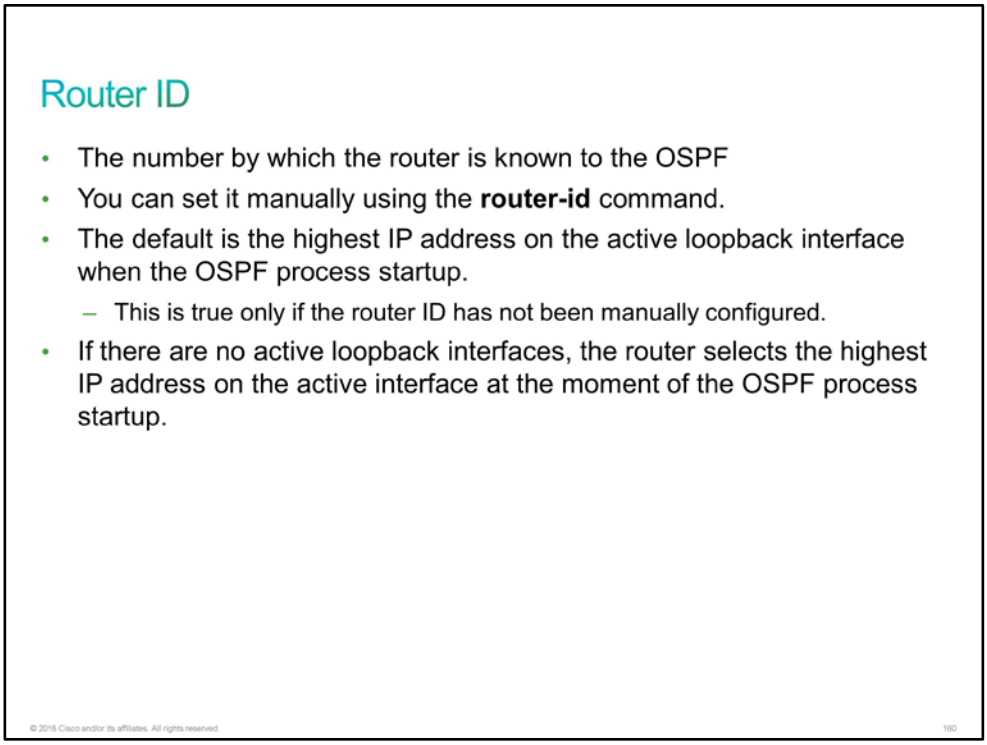

The OSPF router ID is used to uniquely identify each router in the OSPF routing domain. A router ID is simply a label and is expressed as an IP address. Cisco routers derive the router ID based on three criteria and with this precedence:

1. The router uses the IP address (or dotted decimal number) that is configured with the OSPF **router-id** command.
2. If the router ID is not configured, the router chooses the highest IP address of its loopback interfaces.
3. If no loopback interfaces are configured, the router chooses the highest active IP address of its physical interfaces.

Note	The router ID looks like an IP address, but it is not routable and therefore not included in the routing table, unless the OSPF routing process chooses an interface (physical or loopback) that is appropriately defined by a **network** command or **ip ospf**o interface command.

If an OSPF router is not configured with an OSPF **router-id** command and no loopback interfaces are configured, the OSPF router ID will be the highest active IP address on any of its interfaces. The interface does not need to be enabled for OSPF, meaning that it does not need to be included in one of the OSPF network commands. However, the interface must be active—it must be in the "up" state.

Step 3 Access the console of R1 and configure the OSPF process ID 1. Include all the networks that are associated with each of the three active interfaces for R1 in area 0. Also configure the router ID to 1.1.1.1.


```
R1# conf t
R1(config)# router ospf 1
R1(config-router)# router-id 1.1.1.1
R1(config-router)# network 10.0.1.0 0.0.0.255 area 0
R1(config-router)# network 10.1.1.0 0.0.0.255 area 0
R1(config-router)# network 10.10.11.0 0.0.0.255 area 0
R1(config-router)# end
```

During configuration, a syslog message will indicate that new adjacencies have been initiated with two neighbors, R2 and R3. Note that the R2 and R3 routers are represented using the preconfigured router ID.

```
*Oct 13 07:24:35.278: %OSPF-5-ADJCHG: Process 1, Nbr 2.2.2.2 on Ethernet0/0
from LOADING to FULL, Loading Done
*Oct 13 07:24:46.037: %OSPF-5-ADJCHG: Process 1, Nbr 3.3.3.3 on Ethernet0/1
from LOADING to FULL, Loading Done.
```

Verifying Single-Area OSPF

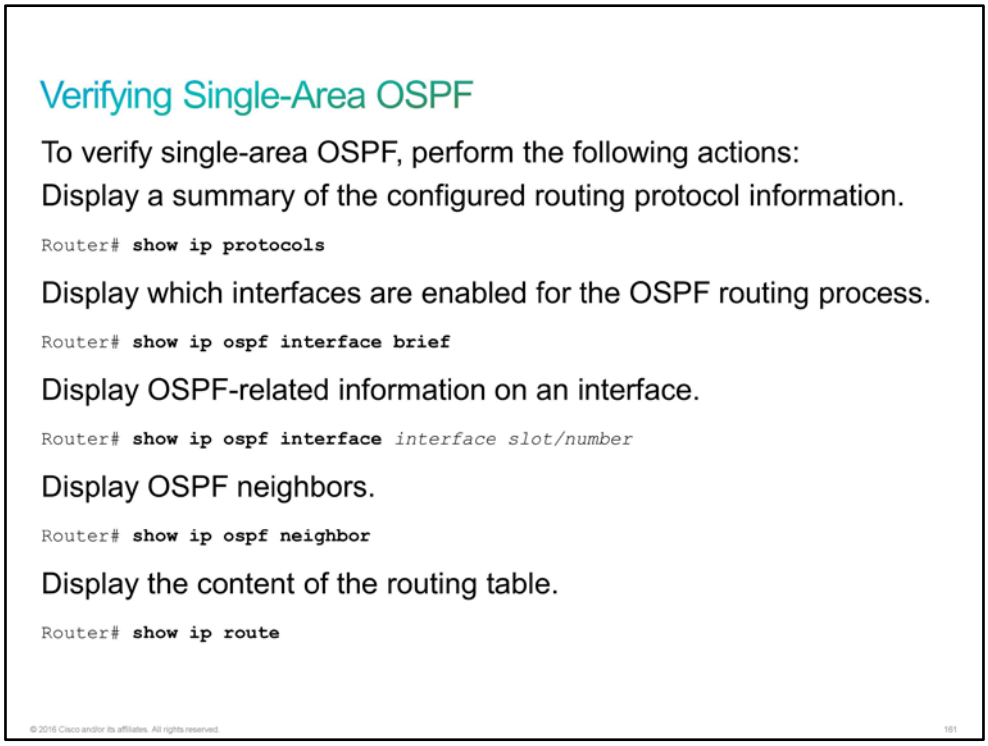

You can use several commands to verify the configuration of single-area OSPF:

- The **show ip protocols** command shows a summary of the configured routing protocol information. You can see which protocols are enabled and which networks these protocols are routing for. You can also see on which interfaces the routing protocols were enabled explicitly.
- Using the **show ip ospf interface** *interface slot/number* you can verify all OSPF-related configuration on an interface.
- Using the **show ip ospf interface brief** command, you can verify which interfaces are enabled for OSPF. It is useful to determine if your network statements were correctly composed.

- Using the **show ip ospf neighbor** you can display the OSPF neighbor information on a per-interface basis.
- The **show ip route** command displays the routes that are known to the router and how they were learned. This command is one of the best ways to determine connectivity between the local router and the rest of the internetwork.

Step 4 Display the interfaces on R1 that are participating in OSPF.

```
R1# show ip ospf interface brief
Interface    PID   Area            IP Address/Mask    Cost  State Nbrs F/C
Lo0          1     0               10.10.11.1/24      1     LOOP  0/0
Et0/1        1     0               10.1.1.1/24        10    BDR   1/1
Et0/0        1     0               10.0.1.1/24        10    BDR   1/1
```

Ethernet0/0, Ethernet0/1, and Loopback0 are participating in OSPF in area 0, under the process ID 1.

Step 5 Display the list of the OSPF neighbors for R1.

```
R1# show ip ospf neighbor

Neighbor ID     Pri   State           Dead Time   Address         Interface
3.3.3.3           1   FULL/DR         00:00:38    10.1.1.3        Ethernet0/1
2.2.2.2           1   FULL/DR         00:00:34    10.0.1.2        Ethernet0/0
```

R1 has two neighbors:

- 3.3.3.3 (R3) that can be reached via the Ethernet0/1 interface. The interface IP address of the interface to which the neighbor is directly connected is 10.1.1.3.
- 2.2.2.2 (R2) that can be reached via the Ethernet0/0 interface. The interface IP address of the interface to which the neighbor is directly connected is 10.0.1.2.

Notice that the neighbor state is "FULL/DR", indicating that the OSPF adjacency is established and both of the neighbors are DR routers. Instead of DR, you could also see BDR state, indicating that the router is BDR, or DROTHER. DROTHER would indicate that the router has priority set to 0 and cannot become DR or BDR.

Step 6 Display the routing table on R1.

The routes that the router has learned via OSPF are tagged with an "O."


```
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 not set

      10.0.0.0/8 is variably subnetted, 9 subnets, 2 masks
C        10.0.1.0/24 is directly connected, Ethernet0/0
L        10.0.1.1/32 is directly connected, Ethernet0/0
C        10.1.1.0/24 is directly connected, Ethernet0/1
L        10.1.1.1/32 is directly connected, Ethernet0/1
O        10.2.1.0/24 [110/20] via 10.1.1.3, 00:04:50, Ethernet0/1
                     [110/20] via 10.0.1.2, 00:05:00, Ethernet0/0
C        10.10.11.0/24 is directly connected, Loopback0
L        10.10.11.1/32 is directly connected, Loopback0
O        10.10.12.0/24 [110/11] via 10.0.1.2, 00:05:00, Ethernet0/0
O        10.10.13.0/24 [110/11] via 10.1.1.3, 00:04:50, Ethernet0/1
```

R1 has learned about the following networks:

- Network between R2 and R3, which has two equal paths—via the Ethernet0/0 and Ethernet0/1 interfaces
- Network of the Loopback interface on R2
- Network of the Loopback interface on R3

OSPF Metric

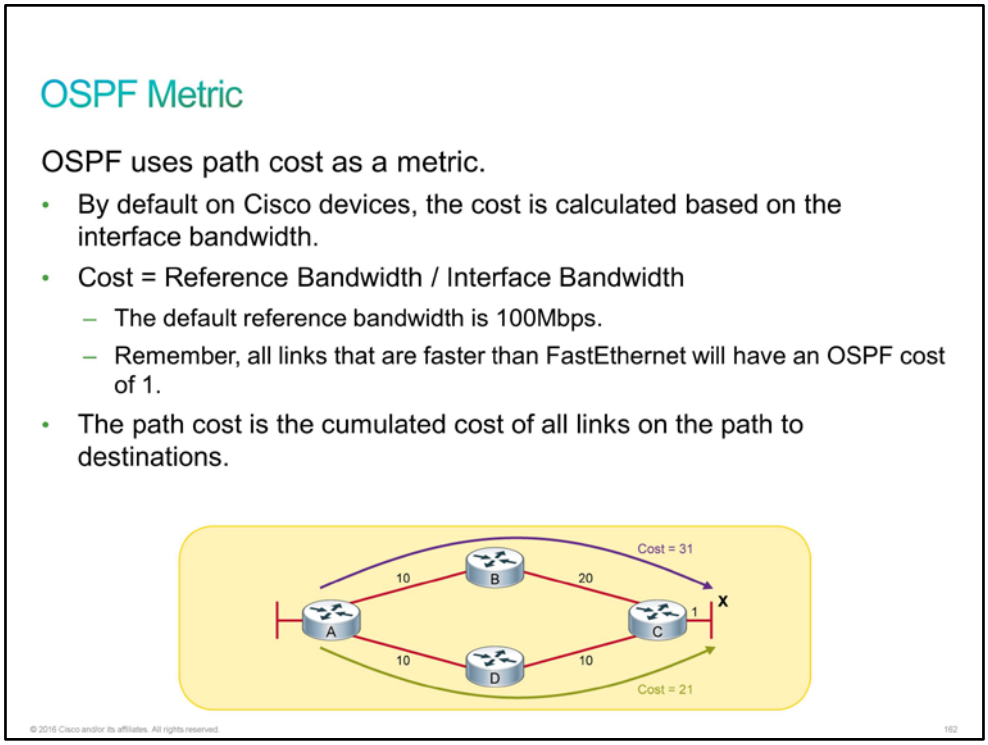

You can apply realistic cost, influencing the following:

- Reference bandwidth, using the **ospf auto-cost reference bandwidth** *reference-bandwidth* command
- Interface cost, using the **ip ospf cost** *cost* command
- Interface bandwidth, using the **bandwidth** command

OSPF uses path costs as a metric. The limitation is that all links that are faster than FastEthernet will have an OSPF cost of 1. There are three approaches you can take to apply realistic costs to your high-speed links.

- **Reference Bandwidth:** You can set the reference bandwidth on the router globally to provide granular link costs.

By default, reference bandwidth is 100 Mbps. To adjust the reference bandwidth for a link, use the **ospf auto-cost reference-bandwidth** *reference-bandwith* command that is configured in the OSPF routing process configuration mode.

- **Interface Cost**: You can choose to use arbitrary cost numbers on every interface.

To override the cost that is calculated for an interface for OSPF routing process, use the **ip ospf cost** *cost* interface configuration command.

- **Interface Bandwidth**: You can configure the **bandwidth** *kilobits-per-second* command on an interface to override the default bandwidth. This has the effect of adjusting the cost of the link regarding routing protocols.

Note	Whether you choose the reference bandwidth method, interface cost method, or interface bandwidth method for adjusting OSPF link costs, it is imperative that you consistently configure adjustments on every router in the OSPF network. Inconsistent application of OSPF link costs can lead to suboptimal path selection.

Step 7 R1 has two paths to the 10.2.1.0/24 network, because both paths have equal cost. Influence the interface cost on R1, so that only the path via Ethernet0/0 will be chosen as the best one.

First, verify the cost of the Ethernet0/0 and Ethernet0/1 interfaces.

```
R1# show ip ospf interface brief
Interface    PID   Area            IP Address/Mask    Cost  State Nbrs F/C
Lo0          1     0               10.10.11.1/24      1     LOOP  0/0
Et0/1        1     0               10.1.1.1/24        10    BDR   1/1
Et0/0        1     0               10.0.1.1/24        10    BDR   1/1
```

Both interfaces have the same cost—10. If you want the path via interface Ethernet0/0 to be chosen, you have to change its cost to a lower value.

```
R1# conf t
R1(config)# interface Ethernet0/0
R1(config-if)# ip ospf cost 1
R1(config-if)# end
```


Alternatively, you could also change the cost of Ethernet0/1 to a higher value.

Step 8 Again, display the routing table of R1.

Verify that there is only one path, the path via Ethernet0/0, to reach the 10.2.1.0/24 network.

```
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 not set

      10.0.0.0/8 is variably subnetted, 9 subnets, 2 masks
C        10.0.1.0/24 is directly connected, Ethernet0/0
L        10.0.1.1/32 is directly connected, Ethernet0/0
C        10.1.1.0/24 is directly connected, Ethernet0/1
L        10.1.1.1/32 is directly connected, Ethernet0/1
O        10.2.1.0/24 [110/11] via 10.0.1.2, 00:00:02, Ethernet0/0
C        10.10.11.0/24 is directly connected, Loopback0
L        10.10.11.1/32 is directly connected, Loopback0
O        10.10.12.0/24 [110/2] via 10.0.1.2, 00:00:02, Ethernet0/0
O        10.10.13.0/24 [110/11] via 10.1.1.3, 00:06:44, Ethernet0/1
```

The total cost to reach the 10.2.1.0/24 network is 11, which is the sum of all the links to reach the network (including the cost of the links that R2 needs to reach this network).

- R1 can reach this network via R2. The cost of the link for R1 to reach R2 is 1. (This is the cost that you configured.)
- The cost of the link for R2 to reach the network is 10.

Passive Interfaces in OSPF

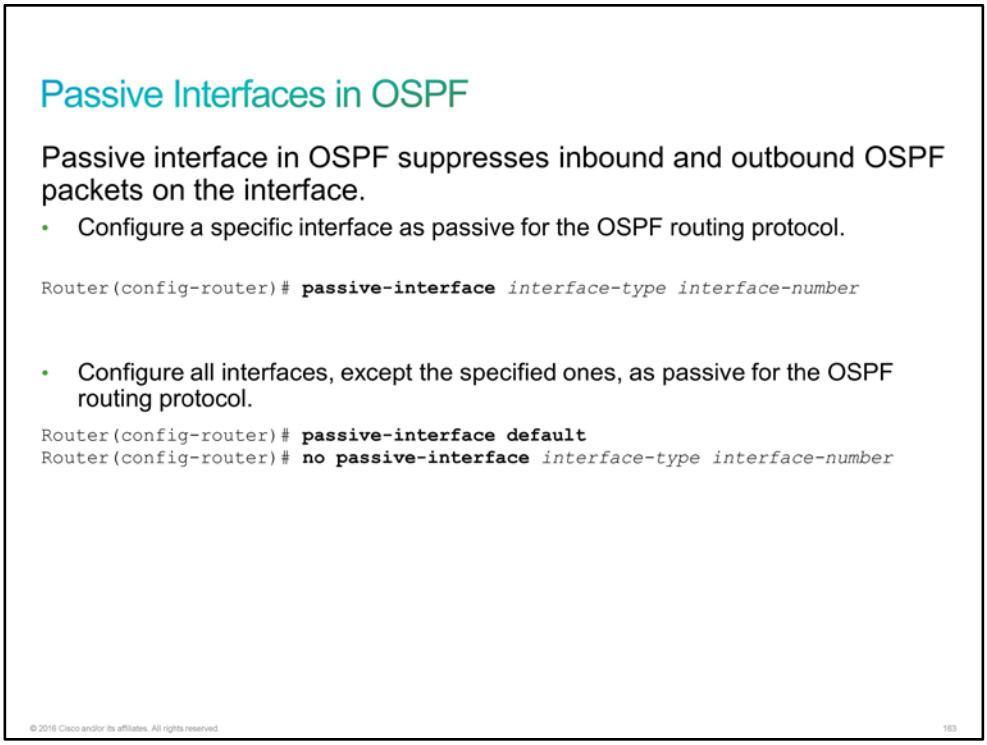

With OSPF running on a network, the **passive-interface** command stops both outgoing and incoming routing updates because the effect of the command causes the router to stop sending and receiving hello packets over an interface. For this reason, the routers will not become neighbors. Use the passive interface configuration only on the interfaces where you do not expect the router to form any OSPF neighbor adjacency.

You can configure either a specific interface as passive, or turn on a passive interface setting as default. Then mark the interfaces which should not be configured as passive with the **no passive-interface** configuration command.

Step 9 On R1, set all interfaces as passive, except the interface connecting to R3.

The easiest way is to use the **passive-interface default** command.

```
R1# conf t
R1(config)# router ospf 1
R1(config-router)# passive-interface default
*Oct 13 11:30:01.326: %OSPF-5-ADJCHG: Process 1, Nbr 2.2.2.2 on Ethernet0/0
from FULL to DOWN, Neighbor Down: Interface down or detached
*Oct 13 11:30:01.326: %OSPF-5-ADJCHG: Process 1, Nbr 3.3.3.3 on Ethernet0/1
from FULL to DOWN, Neighbor Down: Interface down or detached
R1(config-router)# no passive-interface Ethernet0/1
*Oct 13 11:31:07.174: %OSPF-5-ADJCHG: Process 1, Nbr 3.3.3.3 on Ethernet0/1
from LOADING to FULL, Loading Done
R1(config-router)# end
```

Note that during configuration, a syslog message will indicate that existing adjacencies have been terminated. After you specify that Ethernet0/1 should not be configured as passive, the new adjacency is initiated with R3.

Step 10 Display the list of the OSPF neighbors on R1.

```
R1# show ip ospf neighbor

Neighbor ID     Pri   State           Dead Time   Address         Interface
3.3.3.3           1   FULL/DR         00:00:31    10.1.1.3        Ethernet0/1
```

Because only Ethernet0/1 has been excluded from the passive-interface configuration, R1 has formed adjacency only with R3.

This is the end of the discovery lab.

Challenge

1. How can link-state protocols limit the scope of route changes?

   A. by supporting classless addressing
   B. by sending the mask along with the address
   C. by sending only updates of a topology change
   D. by segmenting the network into area hierarchies

2. Which two data structures do the link-state routing protocols use? (Choose two.)

   A. the LSU database
   B. the neighbors database
   C. the link-state interfaces database
   D. the topology database
   E. the next-hop database

3. What is the purpose of link-state advertisements?

   A. constructing a topological database
   B. specifying the cost to reach a destination
   C. determining the best path to a destination
   D. verifying that a neighbor is still functioning

4. Match the OSPF packet type with the correct description.

LSR	sends specifically requested link-state records
DBD	acknowledges the other packet types
LSAck	checks for database synchronization
LSU	requests specific link-state records from another router
Hello	discovers neighbors

5. The SPF algorithm uses a value that is inversely proportional to the bandwidth. What is this value called?

   A. link cost
   B. hop count
   C. link state
   D. MTU

6. Match the OSPF neighbor states with their correct positions in the right order.

4	down
5	init
3	two-way
6	exstart
2	exchange
1	loading
7	full

7. Refer to the figure. Which path will OSPF install in the routing table to reach the 10.10.1.0/24 network from router A?

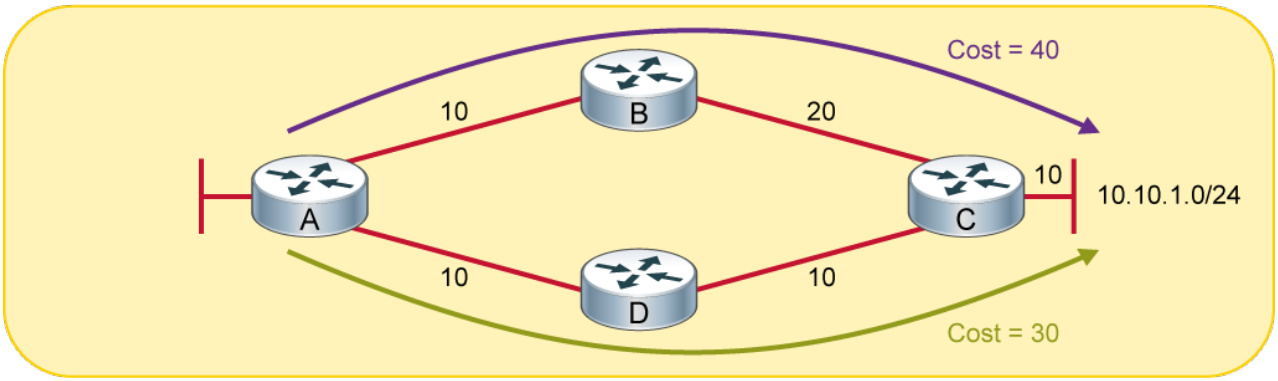

A. the path A-B-C
B. the path A-D-C

8. What is a concern of the default OSPF metric?

A. Link with speeds greater than 10 Gbps will not be supported until the release of OSPF v4.
B. Link with speeds greater than 1 Gbps are converted to 10 Mbps for the purposes of OSPF cost calculation.
C. Link with speeds greater than 100 Mbps have a cost of 1.
D. Link with speeds greater than 1 Gbps require additional memory in the OSPF router or switch to calculate the larger costs.

Answer Key

Challenge

1. D
2. B, D
3. A
4.

LSU	sends specifically requested link-state records
LSAck	acknowledges the other packet types
DBD	checks for database synchronization
LSR	requests specific link-state records from another router
Hello	discovers neighbors

5. A
6.

1	down
2	init
3	two-way
4	exstart
5	exchange
6	loading
7	full

7. B
8. C

Lesson 2: Multiarea OSPF IPv4 Implementation

Introduction

CCS has recently started providing IT services for a startup company that is growing at an unusually rapid pace. The company was so small when they hired CCS that it was still using static routing on its network. Shortly after taking over the IT services for this company, Bob explained to its management that a routing protocol was absolutely necessary. After much discussion, OSPF was selected as the routing protocol. Also, because of the rapid growth of the company, they made the decision to implement multiarea OSPF.

As you are leaving for the customer site to perform the implementation, Bob stops you. He reminds you that not only must you be able to perform the configuration, you must also be able to answer any questions about OSPF that arise from the customer who knows little about it.

OSPF Area Structure

In small networks, the web of router links is not complex, and paths to individual destinations are easily deduced. However, in large networks, the web is highly complex, and the number of potential paths to each destination is large. Therefore, the Dijkstra calculations that compare all these possible routes can be very complex and can take a significant amount of time to complete.

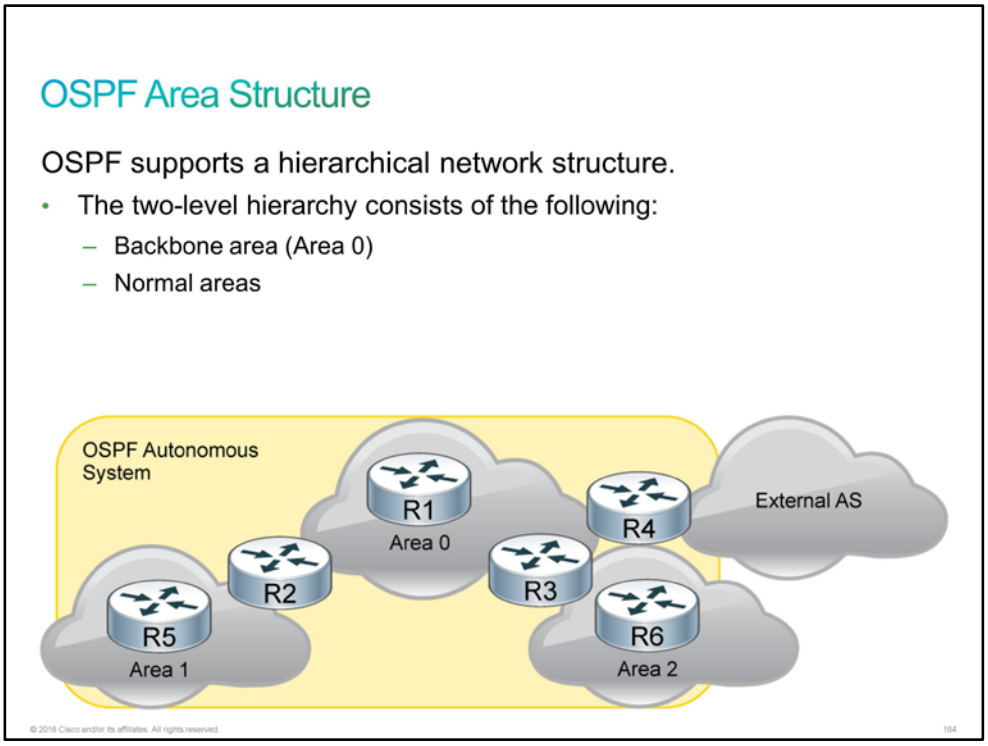

Link-state routing protocols usually reduce the size of the Dijkstra calculations by partitioning the network into areas. The number of routers in an area and the number of LSAs that flood within the area are small, which means that the link-state or topology database for an area is small. So, the Dijkstra calculation is easier and takes less time. The routers that are inside an area maintain detailed information about the links and only general or summary information about the routers and links in other areas. However, summarization is not done by default; it must be configured. Another advantage of using a multiarea OSPF design is that a topology change in an area causes LSA flooding only within the area. SPF recalculations therefore occur only in an area where a topology change has happened.

Link-state routing protocols use a two-layer area hierarchy:

- **Backbone area:** The primary function of this OSPF area is to quickly and efficiently move IP packets. Backbone areas interconnect with other OSPF area types. The OSPF hierarchical area structure requires that all areas connect directly to the backbone area. In the figure, the links between Area 1 and Area 2 routers are not allowed. Generally, end users are not found within a backbone area, which is also known as OSPF Area 0.
- **Normal or nonbackbone area:** The primary function of this OSPF area is to connect users and resources. Normal areas are usually set up according to functional or geographical groupings. By default, a normal area does not allow traffic from another area to use its links to reach other areas. All traffic from other areas must cross a transit area such as Area 0. Normal areas can be of different types. Normal area types affect the amount of routing information that is propagated into the normal area. For example, instead of propagating all routes from the backbone area into a normal area, you could propagate only a default route.

OSPF has special restrictions when multiple areas are involved. One of the areas has to be Area 0, the backbone. All other areas have to be connected to the backbone, which is responsible for distributing routing information between nonbackbone areas.

Note	An OSPF area is identified using a 32-bit Area ID. It can be expressed as either a decimal number or a dotted decimal. You can use both formats at the same time. For example, Area 0 and Area 0.0.0.0 are equivalent. The same goes for Area 14 and Area 0.0.0.14. Area 300 would be the same as 0.0.1.44.

The maximum number of routers per area depends on several factors, however in general it is recommended to minimize the number of routers in one area. Also, you should consider the number of neighbors. Areas with unstable links should be smaller. In general, to maximize stability, one router should not be in more than three areas.

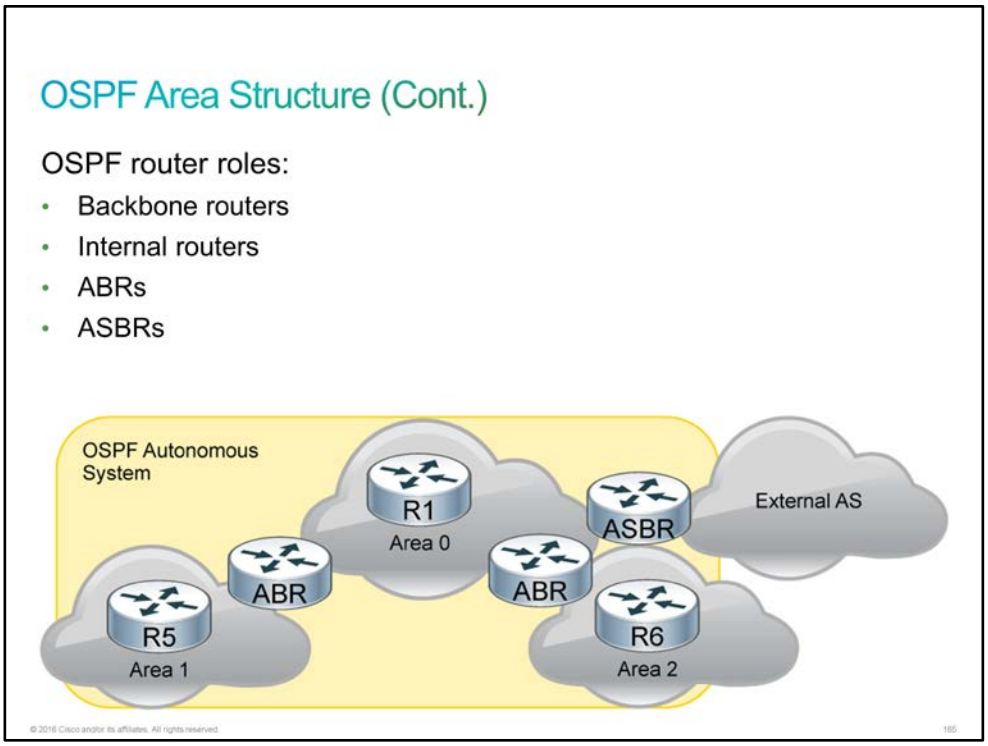

All OSPF areas and routers that are running the OSPF routing protocol comprise the OSPF AS.

Within each AS, a contiguous backbone area, Area 0, must be defined. OSPF hierarchical networking defines Area 0 as the core. All other areas connect directly to backbone. The backbone area is the transition area because all other areas communicate through it.

The routers that make up Area 0 are known as backbone routers. The routers that make up nonbackbone (normal) areas are known as internal routers; they have all interfaces only in one area.

An ABR connects Area 0 to the nonbackbone areas. An OSPF ABR plays a very important role in the network design and has interfaces in more than one area. An ABR has the following characteristics:

- It separates LSA flooding zones.
- It becomes the primary point for area address summarization.
- It functions regularly as the source for default routes.
- It maintains the LSDB for each area with which it is connected.

The ideal design is to have each ABR connected to only two areas, the backbone and another area, with three areas being the upper limit.

An ASBR connects any OSPF area to a different routing administration. The ASBR is the point where external routes can be introduced into the OSPF AS.

In the example, R1 is the backbone router, R2 is an ABR between Areas 0 and 1. R4 acts as the ASBR between the OSPF routing domain and an external domain.

Single-Area vs. Multiarea OSPF

The single-area OSPF design puts all routers into a single OSPF area. This design results in many LSAs being processed on every router and in larger routing tables. The OSPF configuration follows a single-area design in which all the routers are treated as being internal routers to the area and all the interfaces are members of this single area.

As you know, OSPF uses flooding to exchange link-state updates between routers. Any change in the routing information is flooded to all routers in an area. For this reason, the single-area OSPF design can become undesirable as the network grows. The number of LSAs that are processed on every router will increase, and the routing tables may grow quite large.

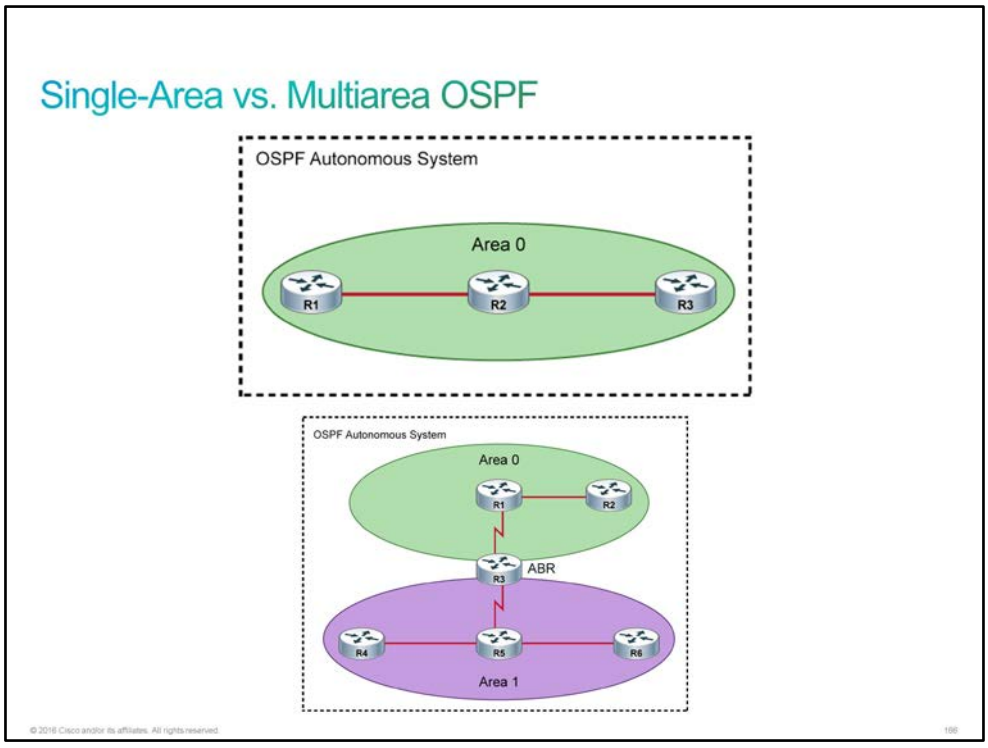

For large or growing networks especially, a multiarea design is a better solution than a single-area design. In a multiarea design, the network is segmented to limit the propagation of LSAs inside an area and to make the routing tables smaller by utilizing summarization.

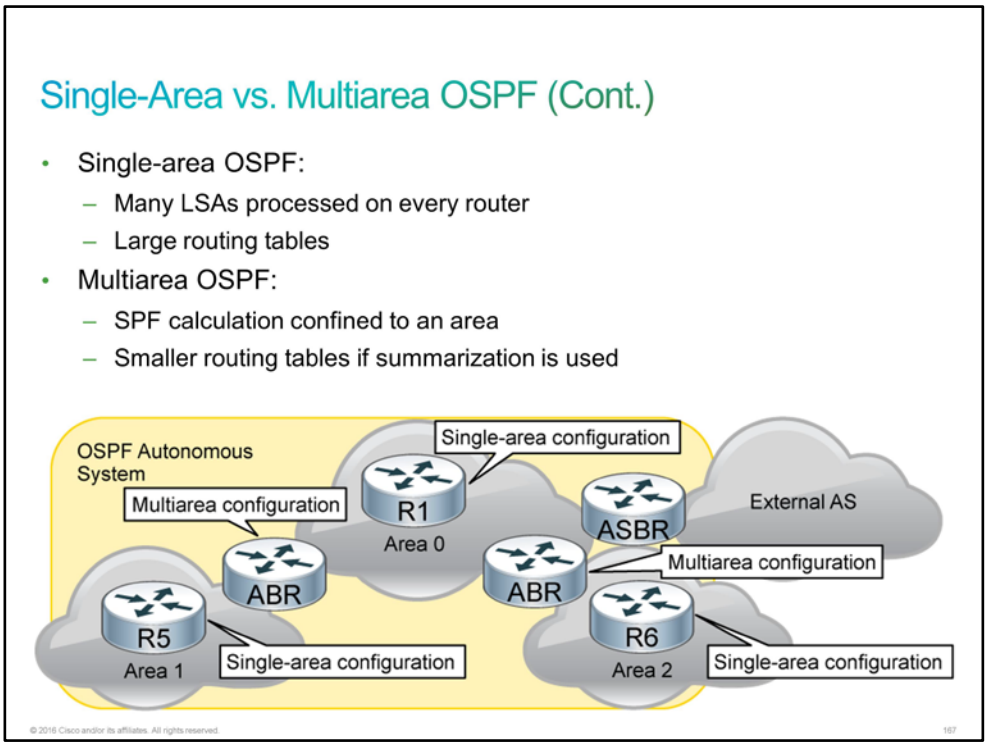

There are two types of routers from the configuration point of view:

- **Routers with single-area configuration:** Internal routers, backbone routers, and ASBRs that are residing in one area
- **Routers with a multiarea configuration:** ABRs and ASBRs that are residing in more than one area

While multiarea OSPF is a scalable and powerful routing protocol, it requires much knowledge to properly design, implement, or troubleshoot.

Multiarea OSPF offers the following advantages over single-area OSPF:

- It can make routing tables smaller if you use route summarization.
- It divides routers into separate areas to limit the propagation and processing of LSAs.

Discovery 44: Configure and Verify Multiarea OSPF

Introduction

This discovery will guide you through the configuration of an ABR in a multiarea OSPF environment. The 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. R1 has also been configured as an internal router in Area 0, while R3 has been configured as an internal router in Area 1. Area 0 spans subnets of 10.0.0.0/16, while Area 1 spans subnets of 10.1.0.0/16. Your job in this discovery is to configure R2 as an ABR between Area 0 and Area 1. After R2 is configured, you will verify the results.

Topology

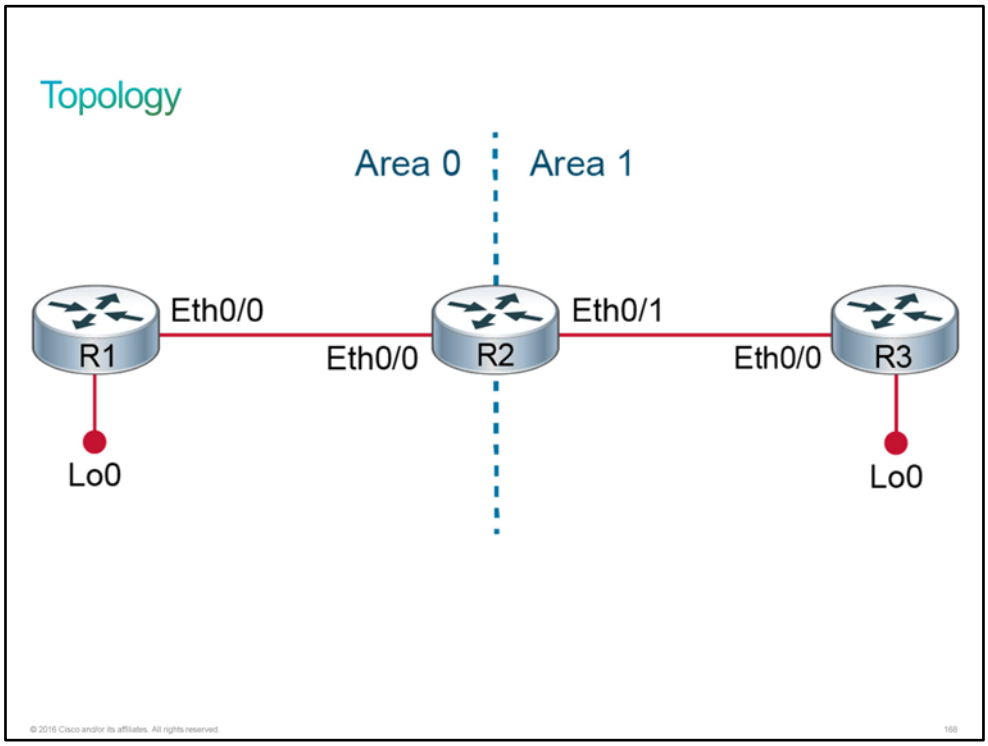

Job Aid

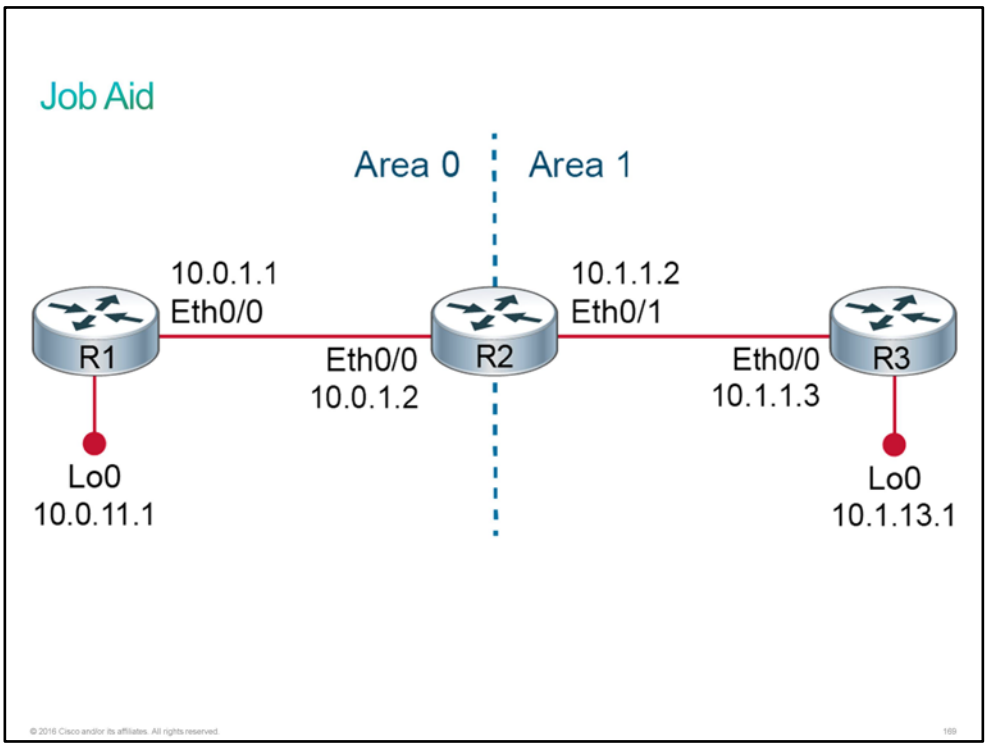

The configuration is as follows:

- All devices have their basic configurations in place, including hostnames and IP addresses.
- OSPF is preconfigured on R1 and R3:
  - R1 has been configured as an internal router in Area 0.
  - R3 has been configured as an internal router in Area 1
- You will configure R2 as an ABR between Area 0 and Area 1.

Device Information

Device Details

Device	Interface	IP Address	Neighbor
R1	Ethernet0/0	10.0.1.1/24	R2
R1	Loopback0	10.0.11.1/24	—
R2	Ethernet0/0	10.0.1.2/24	R1
R2	Ethernet0/1	10.1.1.2/24	R3
R3	Ethernet0/0	10.1.1.3/24	R2

Device	Interface	IP Address	Neighbor
R3	Loopback0	10.1.13.1/24	—

Task 1: Configure and Verify Multiarea OSPF

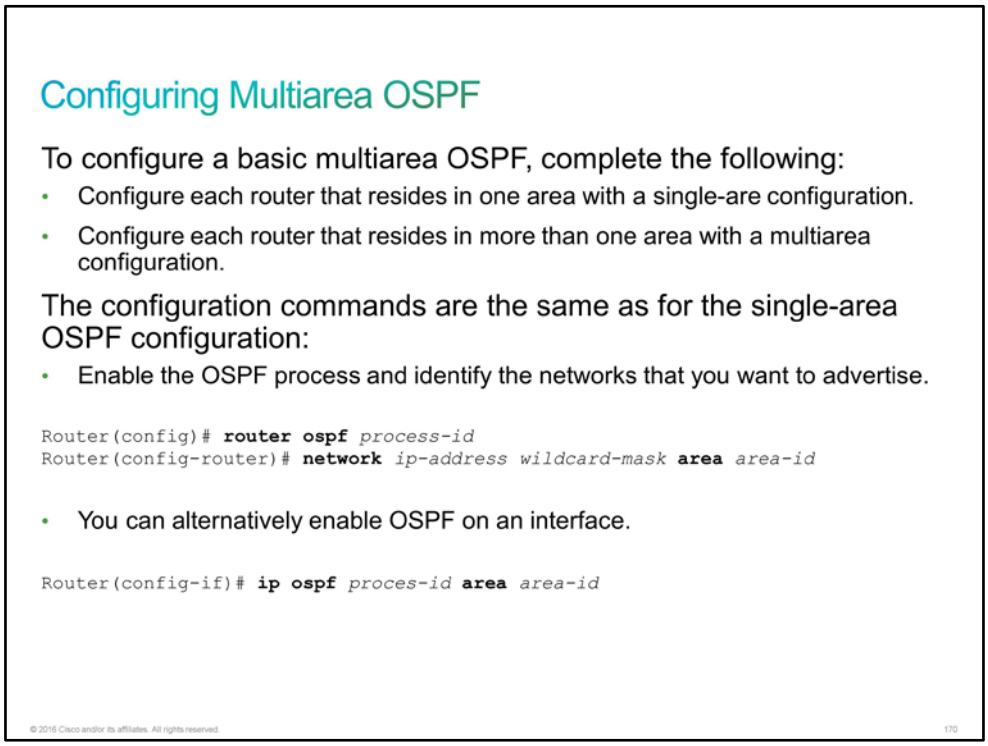

To configure a multiarea OSPF, use the same commands that you would use to configure a single-area OSPF.

Command	Description
ip ospf cost *cost*	Specifies the OSPF cost of sending a packet on an interface. The cost can be a value in the range from 1 to 65,535.
router ospf *process_id*	Configures an OSPF routing process. The *process-id* parameter is an internally used identification parameter for the OSPF routing process. It is locally assigned and can be any positive integer. A unique value is assigned for each OSPF routing process in the router.
network *network wildcard_mask* **area** *area_id*	Defines the interfaces on which OSPF runs and defines the area IDs for those interfaces. The *wildcard_mask* parameter determines how to interpret the IP address. The mask has wildcard bits in which 0 is a match and 1 indicates that the value is not significant. For example, 0.0.255.255 indicates a match in the first two octets.
ip ospf *process-id* **area** *area-id*	Used in the interface configuration mode to enable OSPFv2 on an interface. The *process-id* parameter is a decimal value in the range from 1 to 65535 that identifies the process ID. The *area-id* parameter is a decimal value in the range from 0 to 4,294,967,295, or an IP address.

Activity

Complete the following steps:

Step 1 Access the console of R2. Initialize the OSPF process number 10 and set R2's OSPF router ID to 0.0.0.2.

```
R2# conf t
R2(config)# router ospf 10
R2(config-router)# router-id 0.0.0.2
```

Step 2 Include the interfaces with IP addresses in the 10.0.0.0/16 address range in Area 0. Wait for a syslog message that indicates that a neighbor relationship with R1 has been established.

```
R2(config-router)# network 10.0.0.0 0.0.255.255 area 0
*Oct 15 09:10:26.381: %OSPF-5-ADJCHG: Process 10, Nbr 0.0.0.1 on Ethernet0/0
from LOADING to FULL, Loading Done
```

Step 3 Include the interfaces with IP addresses in the 10.1.0.0/16 address range in Area 1. Wait for a syslog message that indicates that a neighbor relationship with R3 has been established.


```
R2(config-router)# network 10.1.0.0 0.0.255.255 area 1
R2(config-router)#
*Oct 15 09:11:58.769: %OSPF-5-ADJCHG: Process 10, Nbr 0.0.0.3 on Ethernet0/1
from LOADING to FULL, Loading Done
R2(config-router)# end
R2#
```

Verifying Multiarea OSPF

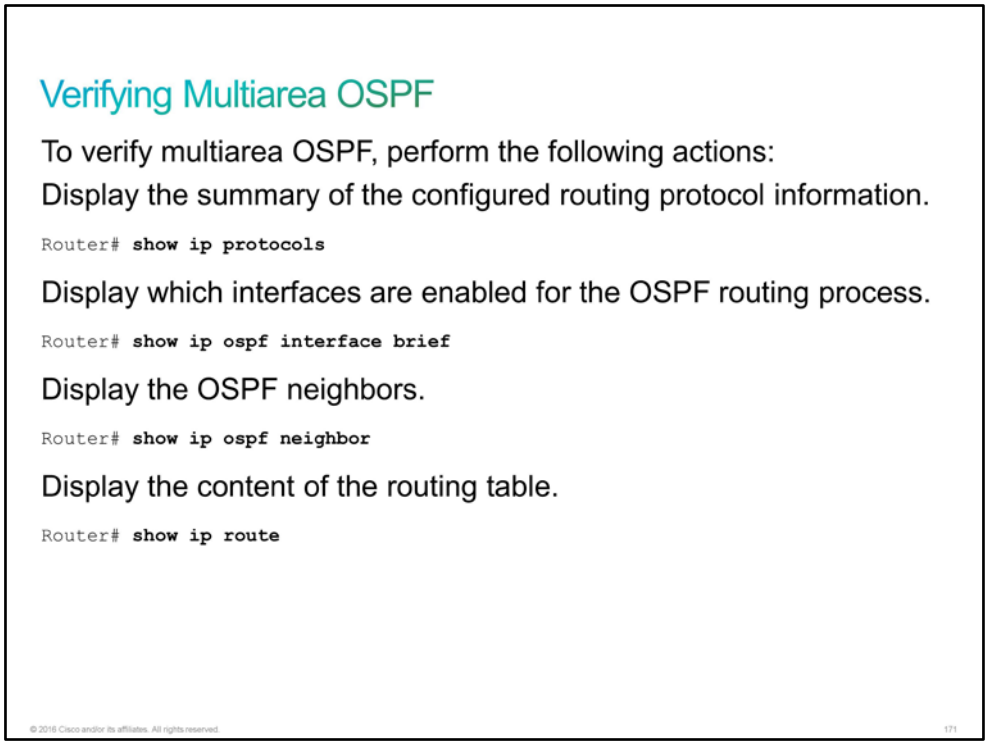

To verify multiarea OSPF configuration, use the same commands that you would use to verify the single-area OSPF.

- **show ip protocols**—to verify the OSPF status, router ID, number of areas in the router, and the networks for which the router routes.
- **show ip ospf interface**—to display OSPF-related information on OSPF-enabled interface. This command will reveal the OSPF process ID to which the interface is assigned, the area that the interfaces are in, and the cost of the interfaces
- **show ip ospf neighbor**—to verify the OSPF neighbors.
- **show ip route ospf**—to verify the OSPF routes in the IP routing table.

Step 4 Display the OSPF neighbors of R2.

```
R2# show ip ospf neighbor

Neighbor ID     Pri   State           Dead Time   Address         Interface
0.0.0.1           1   FULL/BDR        00:00:38    10.0.1.1        Ethernet0/0
0.0.0.3           1   FULL/DR         00:00:37    10.1.1.3        Ethernet0/1
```

R2 has two neighbors:

- 0.0.0.1 (R1) that can be reached via the Ethernet0/0 interface. The IP address of the interface to which the neighbor is directly connected is 10.0.1.1.
- 0.0.0.3 (R3) that can be reached via the Ethernet0/1 interface. The IP address of the interface to which the neighbor is directly connected is 10.1.1.3.

Note: Your router may show a different DR/BDR neighbor state.

Step 5 Display the status of the dynamic routing protocols that are running on R2.

You will find the **show ip protocols** command useful when verifying the configuration and status of all IPv4 dynamic routing protocols.

```
R2# show ip protocols
*** IP Routing is NSF aware ***

Routing Protocol is "ospf 10"
  Outgoing update filter list for all interfaces is not set
  Incoming update filter list for all interfaces is not set
  Router ID 0.0.0.2
  It is an area border router
  Number of areas in this router is 2. 2 normal 0 stub 0 nssa
  Maximum path: 4
  Routing for Networks:
    10.0.0.0 0.0.255.255 area 0
    10.1.0.0 0.0.255.255 area 1
  Routing Information Sources:
    Gateway         Distance      Last Update
    0.0.0.1              110      00:54:21
    0.0.0.3              110      00:54:11
  Distance: (default is 110)
```

All OSPF settings that you configured on R2 are apparent in this output: the process ID, the router ID, and the network definitions. The dynamic status information, such as its ABR status and its OSPF peers, are also apparent in the output.

R2 is configured for OSPF with the process ID 10. Its router ID is 0.0.0.2 and it is an area border router. R2 is routing for the 10.0.0.0/16 networks in area 0 and for 10.1.0.0./16 networks in area 1. It has two neighbors, 0.0.0.1 and 0.0.0.3.

Step 6 Display the summary of information about the OSPF configuration and status on the interfaces of R2.

```
R2# show ip ospf interface brief
Interface    PID   Area            IP Address/Mask    Cost  State Nbrs F/C
Et0/0        10    0               10.0.1.2/24        10    DR    1/1
Et0/1        10    1               10.1.1.2/24        10    BDR   1/1
```

This command makes it clear which interfaces are participating in the OSPF process. If any expected interfaces are missing in the output, it can provide a direction for further troubleshooting.

Besides interface participation, other pertinent information is displayed in the output, including area designation, cost, and neighbor count.

R2 has two interfaces that are participating in OSPF. Ethernet0/0 is in area 0 nad Ethernet0/1 is in area 1.

Step 7 Display the routes from the R2 routing table that were populated via OSPF.

```
R2# show ip route ospf
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, 2 masks
O        10.0.11.0/24 [110/11] via 10.0.1.1, 01:15:59, Ethernet0/0
O        10.1.13.0/24 [110/11] via 10.1.1.3, 01:15:49, Ethernet0/1
```

R2 has two entries for OSPF in its routing table. These two entries are associated with the Loopback interfaces on R1 and R3. Note the code in front of the routes is "O," which indicates OSPF routes.

The presence of OSPF routes indicates that OSPF is operational. If there are expected OSPF routes that are missing from the route table, it can provide a direction for further troubleshooting.

Step 8 Access the console of R1 and display the routes in the routing table that were populated via OSPF.

```
R1# show ip route ospf
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, 2 masks
O IA     10.1.1.0/24 [110/20] via 10.0.1.2, 00:00:09, Ethernet0/0
O IA     10.1.13.0/24 [110/21] via 10.0.1.2, 00:00:09, Ethernet0/0
```

R1 has two entries for OSPF in its routing table. These two entries are associated with the Loopback interface on R3 and the link between R2 and R3. Note that the code in front of the routes is "O IA," which indicates that the routes are interarea routes. This means that the routes originated in another area. Recall that R1 is in Area 0 and R3 is in Area 1, so in this case, both routes originated in Area 1.

The **show ip protocols**, **show ip ospf interface brief**, and **show ip route ospf** commands are three useful commands to quickly summarize the OSPF status on a Cisco IOS router. The most important information is available in an easy-to-read format.

This is the end of the discovery lab.

Challenge

1. Which two statements about single-area and multiarea OSPF are true? (Choose two.)

   A. Single-area OSPF has one advantage over multiarea OSPF: smaller routing tables.
   B. When multiarea design is used, one of the areas should be Area 0.
   C. In multiarea OSPF, Area 1 must be physically connected to the backbone, and all other areas must be connected to area 1.
   D. Multiarea OSPF is more scalable than single-area OSPF, and easier to implement.
   E. In single-area OSPF, all routers inject routing information into the backbone router, and in turn the backbone router disseminates that information to other routers.
   F. Multiarea OSPF can be used to limit the propagation and processing of LSAs.

2. Which command is most efficient for determining the number of areas that are configured on a router?

   A. **show ip protocols**
   B. **show ip ospf interface**
   C. **show ip ospf neighbor**
   D. **show ip route ospf**

3. What does "[110/11]" represent in the command output below?

   R1# show ip route ospf

   <output omitted>

   Gateway of last resort is not set

   172.19.0.0/32 is subnetted, 3 subnets
   C 172.19.0.1 is directly connected, Loopback0
   O IA 172.19.0.3 [110/129] via 192.168.44.2, 00:05:00, Serial1/0
   O IA 172.19.0.2 [110/11] via 192.168.44.2, 00:05:00, Serial1/0

   <Output omitted>

   A. the Administrative Distance that is assigned to OSPF (110) and the hop count to subnet 172.19.0.2 (hop count of 11)
   B. the hop count to subnet 172.19.0.2 (hop count of 110) and the Administrative Distance that is assigned to OSPF (11)
   C. the Administrative Distance that is assigned to OSPF (110) and the total cost of the route to subnet 172.19.0.2 (cost of 11)
   D. the total cost of the route to subnet 172.19.0.2 (cost of 110) and the Administrative Distance that is assigned to OSPF (11)

4. Refer to the output of show ip route below. What does "IA" indicate regarding the destination network?

R2>show ip route
O IA 141.108.1.128/25 [110/846] via 141.108.10.2, 00:08:05, Serial1/0
O IA 141.108.9.128/25 [110/782] via 141.108.10.2, 00:26:20, Serial1/0
O IA 141.108.1.0/25 [110/846] via 141.108.10.2, 00:08:15, Serial1/0
O IA 141.108.9.0/25 [110/782] via 141.108.10.2, 00:26:20, Serial1/0
C 141.108.10.0/30 is directly connected, Serial1/0
O IA 141.108.12.0/24 [110/782] via 141.108.10.2, 00:26:20, Serial1/0
O IA 141.108.10.4/30 [110/845] via 141.108.10.2, 00:26:20, Serial1/0
131.108.0.0/16 is variably subnetted, 8 subnets, 3 masks
O 131.108.4.129/32 [110/11] via 131.108.1.1, 00:46:09, Ethernet0/0
C 131.108.5.32/27 is directly connected, Loopback0

A. It's in the same area as the local router.
B. It's in another area.
C. It will be reached via a default route.
D. The route was learned via another routing protocol.

5. Which type of router is specific to multiarea OSPF design?

A. Backbone router
B. Internal Router
C. ASBR
D. ABR

6. Refer to the output of **show ip route** command from R1. Which of the network from output below is being load balanced ?

R1# show ip route
<output omitted>

Gateway of last resort is not set

10.0.0.0/8 is variably subnetted, 9 subnets, 4 masks
O IA 10.90.50.1/32 [110/64767] via 10.90.245.5, 01:12:43, Serial0/0
O IA 10.90.145.0/24 [110/65766] via 10.90.245.4, 00:18:43, Serial0/0
O IA 10.90.45.0/30 [110/129532] via 10.90.245.5, 01:12:32, Serial0/0
[110/129532] via 10.90.245.4, 01:12:33, Serial0/0
O IA 10.90.20.1/32 [110/64767] via 10.90.245.2, 01:12:43, Serial0/0
C 10.90.10.0/24 is directly connected, Loopback0
C 10.90.245.0/29 is directly connected, Serial0/0

A. 10.90.50.1/32
B. 10.90.145.0/24
C. 10.90.45.0/30
D. 10.90.20.1/32

7. Which of the following command will tell you if the router is an ABR ?

   A. **show ip ospf neighbor**
   B. **show ip ospf interface brief**
   C. **show ip protocol**
   D. **show ip route ospf**

Answer Key

Challenge

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

Lesson 3: Implementing OSPFv3 for IPv6

Introduction

A customer has contacted CCS inquiring about OSPFv3 implementations. The customer wants to implement OSPFv3 but will more than likely want to ask you some questions about OSPFv3 before you configure it. In addition to being able to perform the configuration, you must be able to explain the OSPFv3 enhancements and differences between OSPF for IPv4 and IPv6.

Decide if you are ready to go on site to solve the problem or if you want to do some research first.

OSPFv3 for IPv6

Many concepts of OSPF version 3 are the same as in OSPF version 2. The foundation mostly remains the same as in OSPFv2. OSPFv3 only expands on OSPFv2 to provide support for IPv6 routing prefixes and 128-bit IPv6 addresses.

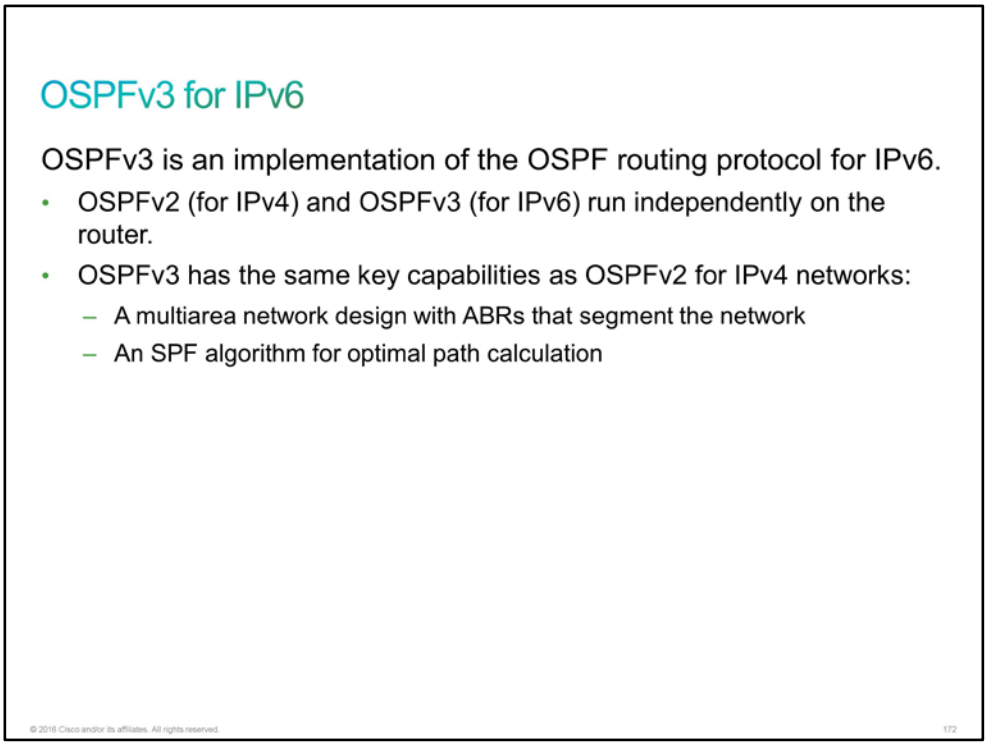

Note	OSPFv3 and OSPFv2 can coexist on the same router, but they run independently in separate processes.

As in OSPFv2, the OSPFv3 metric is still based on interface costing. The default metric remains 100 Mbps.

The packet types and neighbor discovery mechanisms are the same in OSPFv3 as they are in OSPFv2.

LSAs are still flooded throughout an OSPF domain.

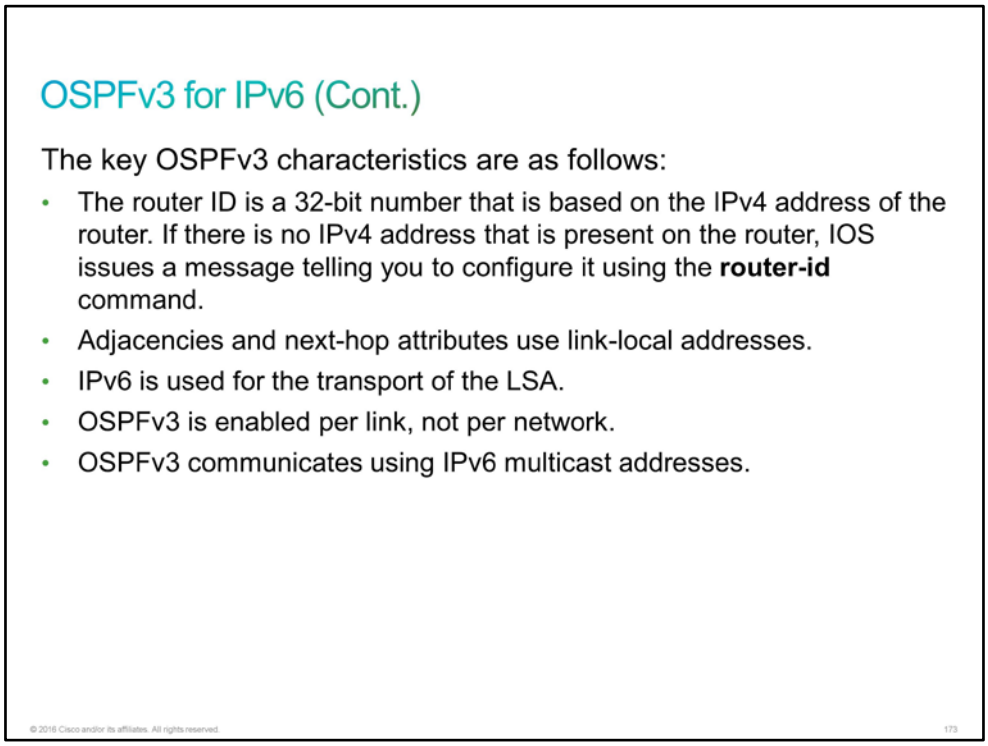

In OSPFv2, the router ID is a 32-bit number, which is derived from the "highest" IPv4 address of an existing router. It is a general practice to set a loopback interface on the router for maintaining the router ID or setting it administratively in the routing process configuration.

In OSPFv3, the OSPF process still requires a 32-bit number to be set. However, if you do not have any IPv4 configuration on the router, you have to enter this 32-bit number manually. This 32-bit number has the same form, as in OSPFv2—four octets that are separated by dots [.]. You set the router ID using the **router-id** *router_id* command. If you don't set it manually, and you have IPv4 configuration on the router, the router ID will be the same as the highest configured loopback IPv4 address. If there is also no loopback configured on the device, then it will use the highest IPv4 address of a physical interface.

OSPFv3 adjacencies use link-local addresses to communicate. Router next-hop attributes are neighboring router link-local addresses. Because link-local addresses have the same prefix, OSPF needs to store the information about the outgoing interface.

OSPFv3 uses IPv6 for the transport of LSAs. The IPv6 protocol number 89 is used. OSPFv3 takes advantage of IPv6 multicasting by using FF02::5 for all OSPF routers and FF02::6 for the OSPF DR and OSPF BDR.

OSPFv3 is enabled per link and identifies which networks (prefixes) are attached to this link for determining prefix reachability propagation and the OSPF area. This feature is different from OSPFv2, in which you can indirectly enable interfaces using the device configuration mode.

Note	One of the most noticeable changes in OSPFv3 is that you don't have to explicitly create a routing process. If you enable OSPFv3 on an interface, a routing process, and its associated configuration, will be created. However, note that you also have to have IPv4 configuration on a router so that the router can obtain the router ID. Otherwise, you have to manually create an OSPV3 routing process.

Discovery 45: Configure and Verify OSPFv3

Introduction

This discovery will guide you through the configuration and verification of OSPFv3 on a Cisco IOS router. The 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 the hostnames and IP addresses. R2 and R3 are also configured with OSPFv3. In this discovery, you will configure OSPFv3 on R1 and verify the results.

Topology

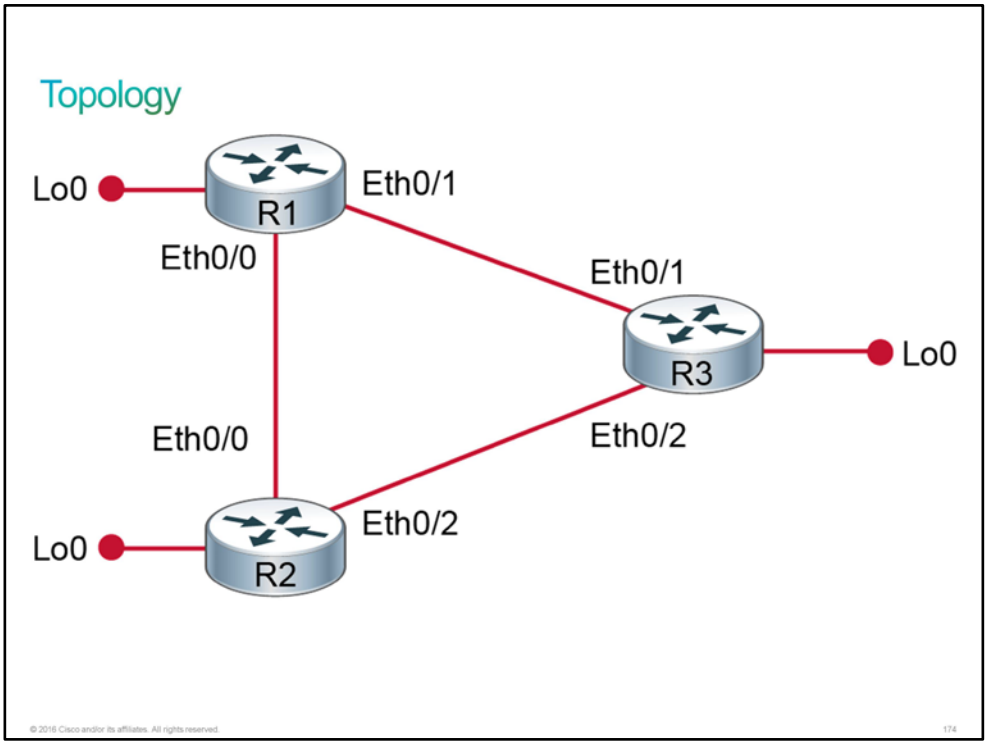

Job Aid

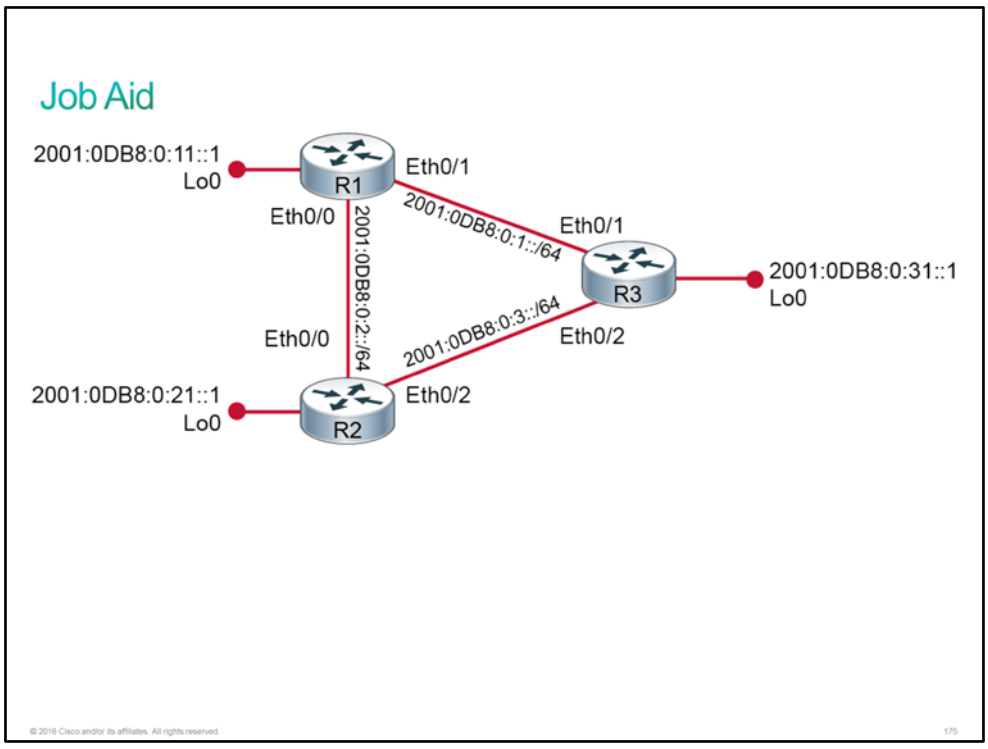

The configuration is as follows:

- All devices have their basic configurations in place including hostnames, IPv4, and IPv6 addresses.
- OSPFv3 is configured on R2 and R3:
  - AS number 10 is used.
  - Both routers are announcing Loopback interface network.

Device Information

Device Details

Device	Interface	IP Address	Neighbor
R1	Ethernet0/0	2001:0DB8:0:2::1	R2
R1	Ethernet0/1	2001:0DB8:0:1::1	R3
R1	Loopback0	2001:0DB8:0:11::1	—
R2	Ethernet0/0	2001:0DB8:0:2::2	R1
R2	Ethernet0/2	2001:0DB8:0:3::2	R3
R2	Loopback0	2001:0DB8:0:21::1	—

Device	Interface	IP Address	Neighbor
R3	Ethernet0/1	2001:0DB8:0:1::3	R1
R3	Ethernet0/2	2001:0DB8:0:3::3	R2
R3	Loopback0	2001:0DB8:0:31::1	—

Task 1: Configure and Verify OSPFv3

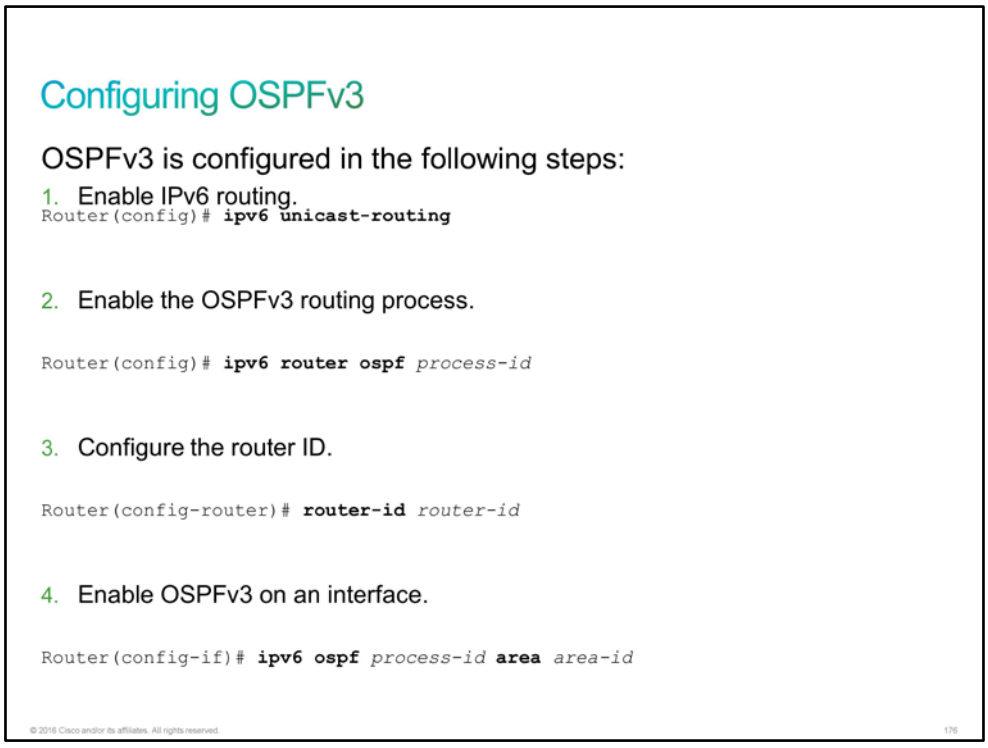

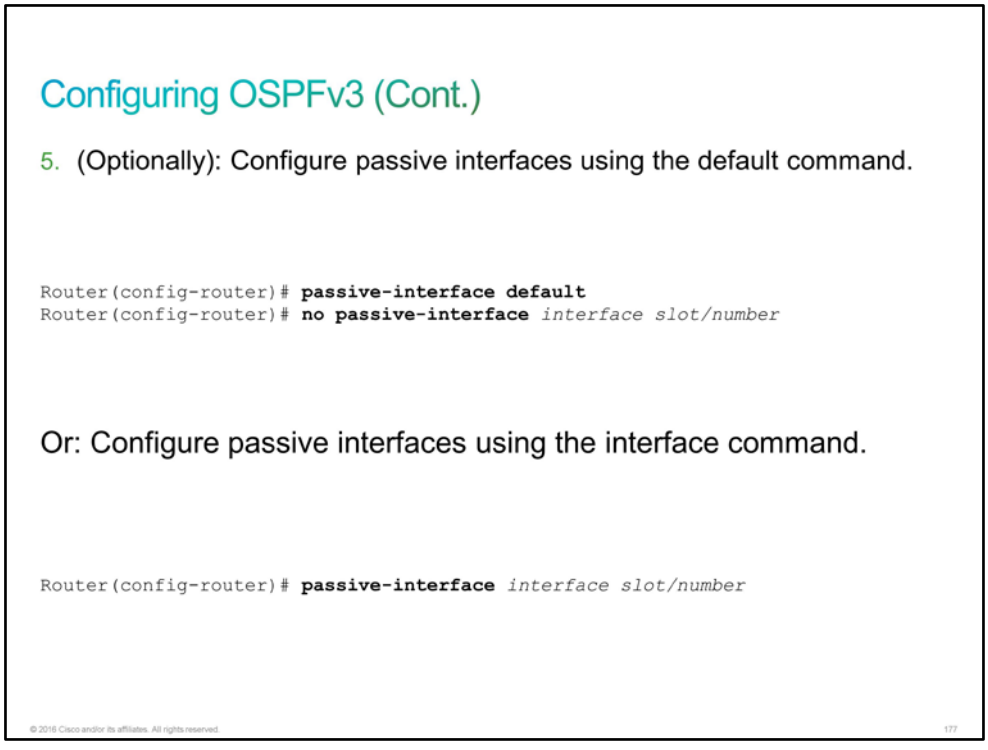

You need to take the following steps to configure OSPFv3 on a router:

1. Because IPv6 routing is not enabled by default, you must first enable it using the **ipv6 unicast-routing** command.
2. Next you have to enable the OSPFv3 routing process with the selected *process-id* parameter.

3. If there are no IPv4 addresses configured on the router, the OSPFv3 routing process requires that you manually configure the router ID.
4. Optionally, you can configure passive-interfaces.
5. All that remains then is for you to enable OSPFv3 routing on the interface. Of course, the interface must have an IPv6 address assigned and it has to be administratively enabled.

The table defines the commands that you use to configure OSPFv3.

Configuring OSPFv3 Commands

Command	Description
ipv6 unicast-routing	Enables the forwarding of IPv6 unicast datagrams and is used in the global configuration mode. To disable the forwarding of IPv6 unicast datagrams, use the **no** form of this command.
ipv6 router ospf *process-id*	Enables OSPF for the IPv6 router configuration mode. The *process-id* value is an internal identification. It is locally assigned and can be a positive integer from 1 to 65,535.
router-id *router-id*	Executed in the OSPF router configuration mode to statically configure a router ID, which is the name for the router within the OSPFv3 process.
ipv6 ospf *process-id* **area** *area-id*	Enables OSPFv3 on an interface and assigns it to the specified area.
passive-interface default	Configures all interfaces as passive for OSPFv3 process.
passive-interface *interface slot/number*	Configures specified interface as passive for OSPFv3 process.

Activity

Complete the following steps:

Step 1 Access the console of R2 and display the dynamic IPv6 routing protocols that are running on it.

Use the **show ipv6 protocols** command to display the status of the configured dynamic IPv6 protocols.

```
R2# show ipv6 protocols
IPv6 Routing Protocol is "connected"
IPv6 Routing Protocol is "ND"
IPv6 Routing Protocol is "ospf 10"
  Router ID 2.2.2.2
  Number of areas: 1 normal, 0 stub, 0 nssa
  Interfaces (Area 0):
    Loopback0
    Ethernet0/2
    Ethernet0/0
  Redistribution:
    None
```


This command verifies the OSPFv3 process ID and the interfaces that are configured, along with the area ID to which they were assigned.

R2 is running OSPFv3 with the process ID 10. Its router ID is 2.2.2.2, and it has three interfaces assigned to Area 0—Loopback0, Ethernet0/0, and Ethernet0/2.

Step 2 Access the console of R3 and display the dynamic IPv6 routing protocols that are running on it.

```
R3# show ipv6 protocols
IPv6 Routing Protocol is "connected"
IPv6 Routing Protocol is "ND"
IPv6 Routing Protocol is "ospf 10"
  Router ID 3.3.3.3
  Number of areas: 1 normal, 0 stub, 0 nssa
  Interfaces (Area 0):
    Loopback0
    Ethernet0/2
    Ethernet0/1
  Redistribution:
    None
```

R3 is running OSPFv3 with the process ID 10. Its router ID is 3.3.3.3, and it has three interfaces assigned to Area 0—Loopback0, Ethernet0/1, and Ethernet0/2.

Step 3 Access the console of R1 and first enable IPv6 routing on it.

```
R1# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
R1(config)# ipv6 unicast-routing
```

Step 4 Define the OSPv3 process ID 10 on R1, and assign to it the router ID 1.1.1.1.

```
R1(config)# ipv6 router ospf 10
R1(config-rtr)# router-id 1.1.1.1
```

The OSPF process ID doesn't have to match among peers in an OSPFv3 network, but standardizing on a process ID across a deployment minimizes potential administrative confusion.

Step 5 Enable OSPFv3 on the R1 Ethernet0/0, Ethernet0/1, and Loopback0 interfaces. All interfaces should be assigned to Area 0.

```
R1(config-rtr)# interface Loopback0
R1(config-if)# ipv6 ospf 10 area 0
R1(config-if)# interface Ethernet0/0
R1(config-if)# ipv6 ospf 10 area 0
R1(config-if)#
*Oct 16 11:25:05.346: %OSPFv3-5-ADJCHG: Process 10, Nbr 2.2.2.2 on Ethernet0/0
from LOADING to FULL, Loading Done
R1(config-if)# interface Ethernet0/1
R1(config-if)# ipv6 ospf 10 area 0
R1(config-if)# end
*Oct 16 11:25:15.912: %OSPFv3-5-ADJCHG: Process 10, Nbr 3.3.3.3 on Ethernet0/1
from LOADING to FULL, Loading Done
```

Note that when you enabled OSPFv3 on Ethernet0/0, the neighbor relationship between R1 and R2 was initiated. The same goes for enabling OSPFv3 on Ethernet0/1 and the neighbor relationship between R1 and R3.

Verifying OSPFv3

There are some important commands that you should be familiar with in order to validate your OSPFv3 configurations. These commands are similar to the ones that you used to verify OSPF for IPv4.

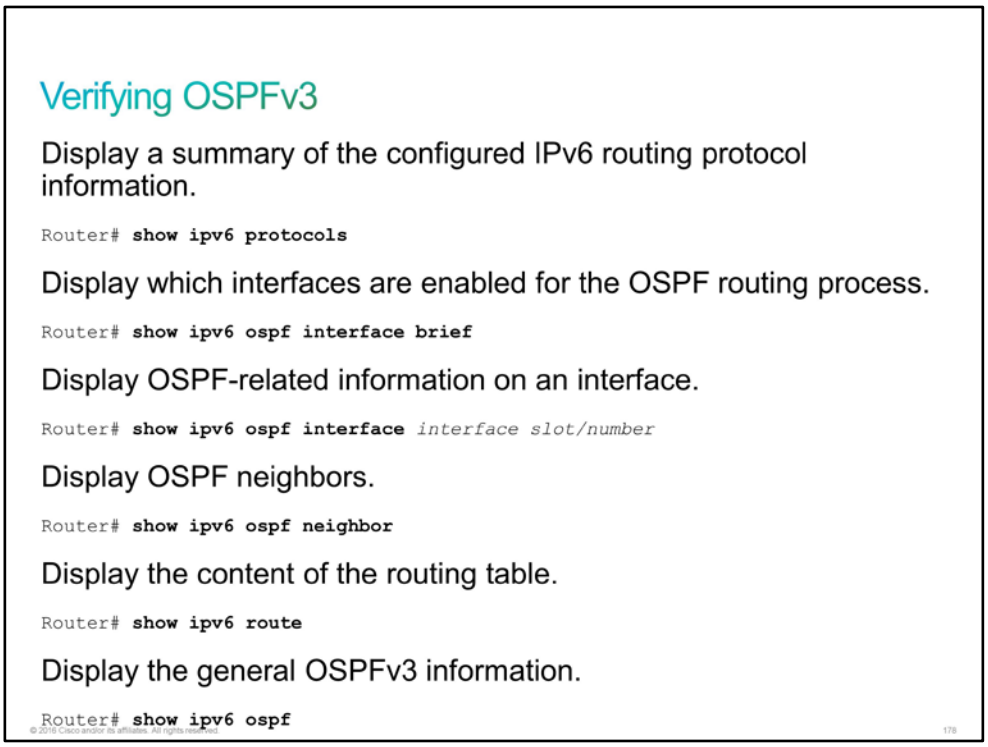

- **show ipv6 protocols**—shows a summary of the configured IPv6 routing protocol information. You can see which protocols are enabled and which networks these protocols are routing for.
- **show ipv6 ospf interface brief**—shows the interfaces that are enabled with OSPFv3 and OSPF-related interface information.
- **show ipv6 ospf interface** *interface slot/number*—shows the all OSPFv3 information on an interface. You can also see if the interface is configured as passive.
- **show ipv6 ospf neighbor**—shows OSPFv3 neighbors on per-interface basis.
- **show ipv6 route**—shows IPv6 routers that this router learns.
- **show ipv6 ospf**—shows general information about the OSPF routing process, such as OSPFv3 process ID, router ID, timers, areas that are configured, and reference bandwidth.

Step 6 Display the summary of the interface status about OSPFv3 on R1.


```
R1# show ipv6 ospf interface brief
Interface    PID   Area            Intf ID    Cost  State Nbrs F/C
Lo0          10    0               10         1     LOOP  0/0
Et0/1        10    0               4          10    BDR   1/1
Et0/0        10    0               3          10    BDR   1/1
```

The output verifies that OSPFv3 is running on the three interfaces, as expected. They are also all associated with the process ID 10 and with Area 0.

Ethernet0/0 and Ethernet0/1 have a neighbor count of 1. R2 and R3 are peers on those interfaces.

Step 7 Display the IPv6 routing table on R1.

The routes that have been learned via OSPF are tagged with an "O."

```
R1# show ipv6 route
IPv6 Routing Table - default - 10 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
C   2001:DB8:0:1::/64 [0/0]
     via Ethernet0/1, directly connected
L   2001:DB8:0:1::1/128 [0/0]
     via Ethernet0/1, receive
C   2001:DB8:0:2::/64 [0/0]
     via Ethernet0/0, directly connected
L   2001:DB8:0:2::1/128 [0/0]
     via Ethernet0/0, receive
O   2001:DB8:0:3::/64 [110/20]
     via FE80::A8BB:CCFF:FE00:1C00, Ethernet0/0
     via FE80::A8BB:CCFF:FE00:1D10, Ethernet0/1
C   2001:DB8:0:11::/64 [0/0]
     via Loopback0, directly connected
L   2001:DB8:0:11::1/128 [0/0]
     via Loopback0, receive
O   2001:DB8:0:21::/64 [110/11]
     via FE80::A8BB:CCFF:FE00:1C00, Ethernet0/0
O   2001:DB8:0:31::/64 [110/11]
     via FE80::A8BB:CCFF:FE00:1D10, Ethernet0/1
L   FF00::/8 [0/0]
     via Null0, receive
```

R1 has learned, via OSPFv3, the routes to the networks that are associated with the loopback interfaces on R2 and R3 and also the networks that connect R1 to R2, and R1 to R3.

The next-hop IPv6 addresses are link-local addresses and not global unicast addresses. Note that the IPv6 addresses in your output may be different.

Step 8 Display the list of OSPFv3 neighbors for R1.

```
R1# show ipv6 ospf neighbor

            OSPFv3 Router with ID (1.1.1.1) (Process ID 10)

Neighbor ID     Pri   State           Dead Time   Interface ID    Interface
3.3.3.3           1   FULL/DR         00:00:35    4               Ethernet0/1
2.2.2.2           1   FULL/DR         00:00:30    3               Ethernet0/0
```

This command verifies that R2 is a neighbor on the interface Ethernet0/0 and that R3 is a neighbor on the interface Ethernet0/1. Both neighbors are represented using the configured router ID.

Step 9 Display the detailed information that is available about the R1 OSPFv3 neighbors.

```
R1# show ipv6 ospf neighbor detail

            OSPFv3 Router with ID (1.1.1.1) (Process ID 10)

 Neighbor 3.3.3.3
    In the area 0 via interface Ethernet0/1
    Neighbor: interface-id 4, link-local address FE80::A8BB:CCFF:FE00:2F10
    Neighbor priority is 1, State is FULL, 6 state changes
    DR is 3.3.3.3 BDR is 1.1.1.1
    Options is 0x000013 in Hello (V6-Bit, E-Bit, R-bit)
    Options is 0x000013 in DBD (V6-Bit, E-Bit, R-bit)
    Dead timer due in 00:00:31
    Neighbor is up for 00:07:05
    Index 1/2/2, retransmission queue length 0, number of retransmission 0
    First 0x0(0)/0x0(0)/0x0(0) Next 0x0(0)/0x0(0)/0x0(0)
    Last retransmission scan length is 0, maximum is 0
    Last retransmission scan time is 0 msec, maximum is 0 msec

 Neighbor 2.2.2.2
    In the area 0 via interface Ethernet0/0
    Neighbor: interface-id 3, link-local address FE80::A8BB:CCFF:FE00:2E00
    Neighbor priority is 1, State is FULL, 6 state changes
    DR is 2.2.2.2 BDR is 1.1.1.1
    Options is 0x000013 in Hello (V6-Bit, E-Bit, R-bit)
    Options is 0x000013 in DBD (V6-Bit, E-Bit, R-bit)
    Dead timer due in 00:00:39
    Neighbor is up for 00:07:05
    Index 1/1/1, retransmission queue length 0, number of retransmission 0
    First 0x0(0)/0x0(0)/0x0(0) Next 0x0(0)/0x0(0)/0x0(0)
    Last retransmission scan length is 0, maximum is 0
    Last retransmission scan time is 0 msec, maximum is 0 msec
```

OSPFv3 neighbors are identified by their IPv6 link-local addresses, not by their IPv6 global unicast addresses. Again, note that the IPv6 addresses in your output may be different.

Step 10 Display a robust set of data that is associated with the global status and configuration of OSPFv3 on R1.

Use the **show ipv6 ospf** command.


```
R1# show ipv6 ospf
 Routing Process "ospfv3 10" with ID 1.1.1.1
 Event-log enabled, Maximum number of events: 1000, Mode: cyclic
 Router is not originating router-LSAs with maximum metric
 Initial SPF schedule delay 5000 msecs
 Minimum hold time between two consecutive SPFs 10000 msecs
 Maximum wait time between two consecutive SPFs 10000 msecs
 Minimum LSA interval 5 secs
 Minimum LSA arrival 1000 msecs
 LSA group pacing timer 240 secs
 Interface flood pacing timer 33 msecs
 Retransmission pacing timer 66 msecs
 Number of external LSA 0. Checksum Sum 0x000000
 Number of areas in this router is 1. 1 normal 0 stub 0 nssa
 Graceful restart helper support enabled
 Reference bandwidth unit is 100 mbps
 RFC1583 compatibility enabled
    Area BACKBONE(0)
        Number of interfaces in this area is 3
        SPF algorithm executed 1 times
        Number of LSA 16. Checksum Sum 0x05FF18
        Number of DCbitless LSA 0
        Number of indication LSA 0
        Number of DoNotAge LSA 0
        Flood list length 0
```

Note that you can also see the reference bandwidth in the output.

This is the end of the discovery lab.

Challenge

1. Which of the following is a characteristic that is unique to OSPFv3 compared to OSPFv2 ?

   A. LSA flooding distributes link information.
   B. Adjacencies are formed with link-local addresses.
   C. The IP protocol is 89.
   D. OSPF process IDs must be uniform on all routers to allow neighbor discovery.

2. You are configuring OSPFv3 on a new router and you are presented with the following logging message. What does this message indicate?

   ***Apr 3 08:14:59.727: %OSPFv3-4-NORTRID: OSPFv3 process 99 could not pick a router-id, please configure manually**

   A. OSPF process 99 is not within the valid range of 1 to 64.
   B. No IPv4 addresses are configured on this router.
   C. No loopback interfaces are configured on this router.
   D. IPv6 routing is not enabled on this router.

3. Which interface speed does OSPFv3 use in the default cost metric calculation?

   A. 100 Mbps
   B. 1000 Mbps
   C. 10 Gbps
   D. 100 Gbps

4. Which command includes the router IDs of all peers and the link-local addresses of the local router and its peers in its output?

   A. **show ipv6 ospf**
   B. **show ipv6 route ospf**
   C. **show ipv6 ospf neighbor**
   D. **show ipv6 ospf interface**

5. When implementing OSPFv3 for IPv6, which statement describes the configuration of OSPF areas?

   A. In interface configuration mode, the OSPFv3 area ID combination assigns interfaces to OSPFv3 areas.
   B. In router configuration mode, the network wildcard area ID combination assigns networks to OSPFv3 areas.
   C. In interface configuration mode, the IPv6 OSPF process area ID combination assigns interfaces to OSPFv3 areas.
   D. In router configuration mode, the IPv6 OSPF interface area ID combination assigns interfaces to OSPFv3 areas.

6. In OSPFv3, you don't have to explicitly create a routing process if you have IPv4 configuration on the router. True or false ?

   A. False
   B. True

7. You are configuring OSPFv3 on a router and you see the following logging message. What needs to be configured before these commands are entered on the router ?

   Router(config)# interface ethernet0/0
   Router(config-if)# ipv6 ospf 1 area 2
   % OSPFv3: IPv6 routing not enabled

   A. **Router(config)# ipv6 router ospf 1**
   B. **Router(config)# ipv6 unicast-routing**
   C. **Router(config)# interface e0/0**
      Router(config-if)# ipv6 enable
   D. **Router(config)# ipv6 multicast-routing**

Answer Key

Challenge

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

Lesson 4: Troubleshooting Multiarea OSPF

Introduction

Two different customers have called CCS with complaints involving OSPF routing issues. Trouble tickets have been issued for both complaints. Bob is reviewing the trouble tickets and trying to decide which ones to dispatch you on.

Components of Troubleshooting OSPF

Troubleshooting OSPF requires an understanding of the operation of the protocol and also of a specific approach methodology. The figure shows the major components of OSPF troubleshooting and the order in which the process flows.

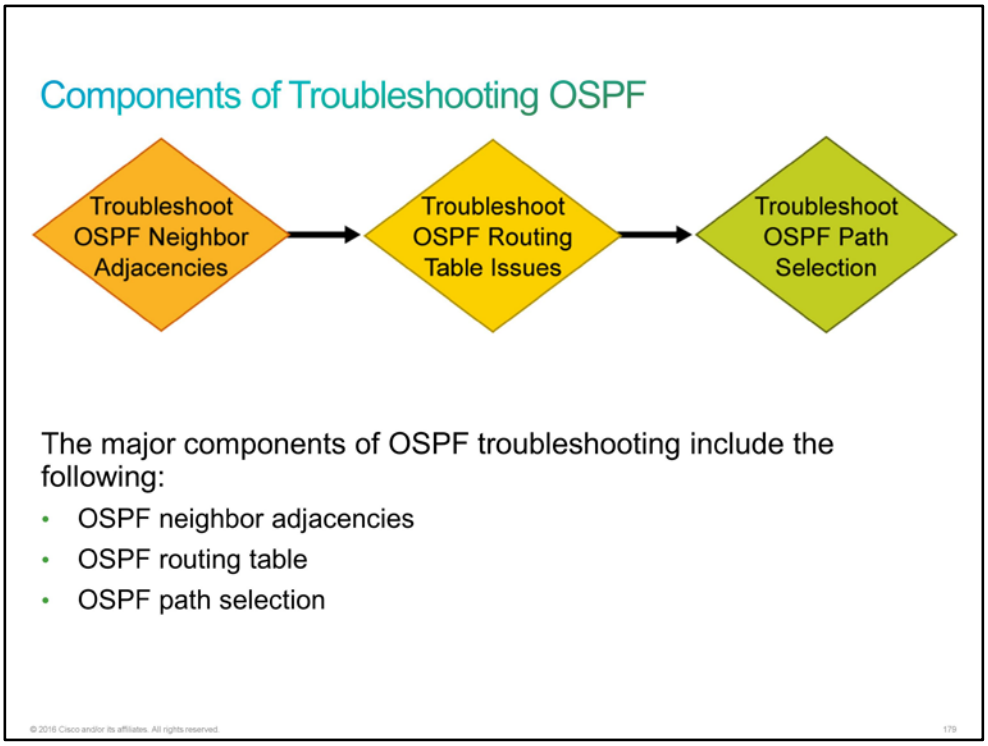

Troubleshooting OSPF Issues

When you are notified that there are connectivity issues in your network, you should first test connectivity by using the **ping** and **traceroute** commands. If there are connectivity issues and your network uses OSPF as the routing protocol, then follow these high-level steps to troubleshoot it.

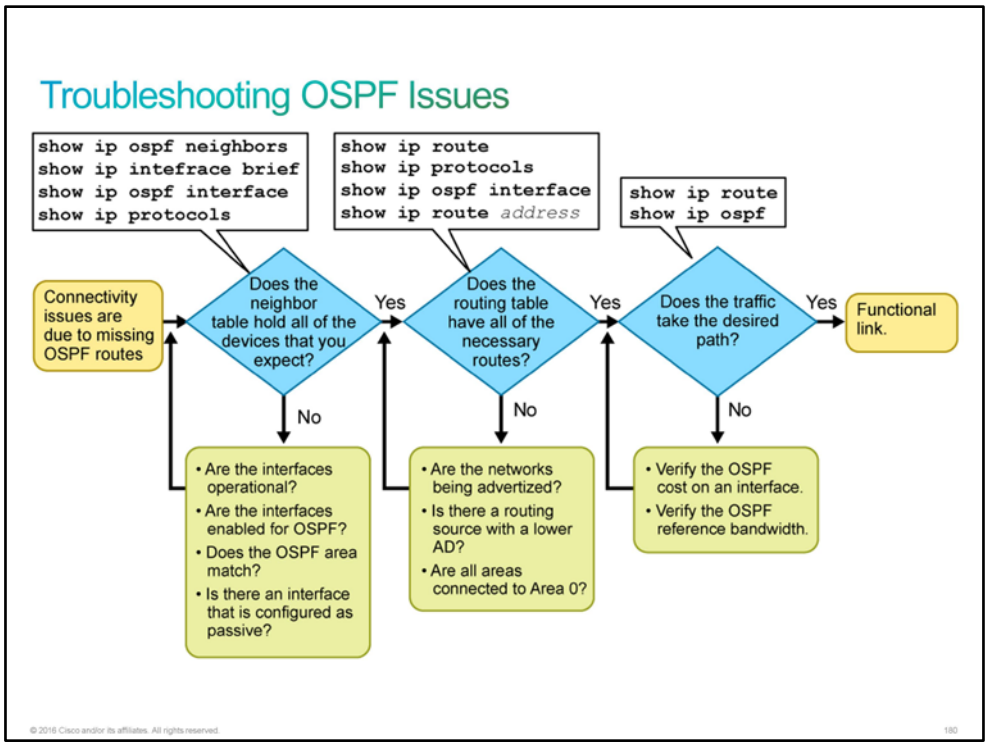

1. Verify that your router established an adjacency with a neighboring router by using the **show ip ospf neighbors** command. If an adjacency between two routers is not established, the routers cannot exchange routes. If the adjacency is not established, you should first verify that the interfaces are operational and enabled for OSPF. If the interfaces are operational and enabled for OSPF, you should also make sure that the interfaces on both routers are configured for the same OSPF area and the interfaces are not configured as passive interfaces.
2. If an adjacency between two routers is established, but you see no routes in the routing table when you use the **show ip route** command, you should first verify if there is another routing protocol with a lower administrative distance running in the network. In this case, OSPF routes would not be considered and will not be placed into the routing table. If no other routing protocols are configured, verify that all the required networks are advertised into OSPF. In the case of multiarea OSPF, you should also verify whether all regular nonbackbone areas are connected directly to area 0 or to the backbone area. If a regular area is not connected to the backbone area, routers in this area will not be able to send and receive updates to and from other areas.
3. If you see all the required routes in the routing table but the path that the traffic takes is not correct, you should verify the OSPF cost on the interfaces on the path. You should also be careful in cases where you have interfaces that are faster than 100 Mbps, because all interfaces above this bandwidth will have the same OSPF cost, by default.

Troubleshooting OSPF Neighbor Issues

The first component to troubleshoot and verify is the OSPF neighbor adjacency. The troubleshooting and verification components for neighbor adjacencies are as follows:

1. Verify that links on routers are Layer 2 operational.
2. Verify Layer 3 connectivity between routers.
3. Verify that the interfaces on both routers are enabled for OSPF.
4. Verify that the OSPF area and other required parameters matches on both ends.
5. Verify that the interfaces on both routers are not configured as passive.

In the example, you will investigate a neighbor issue between the Branch and HQ router.

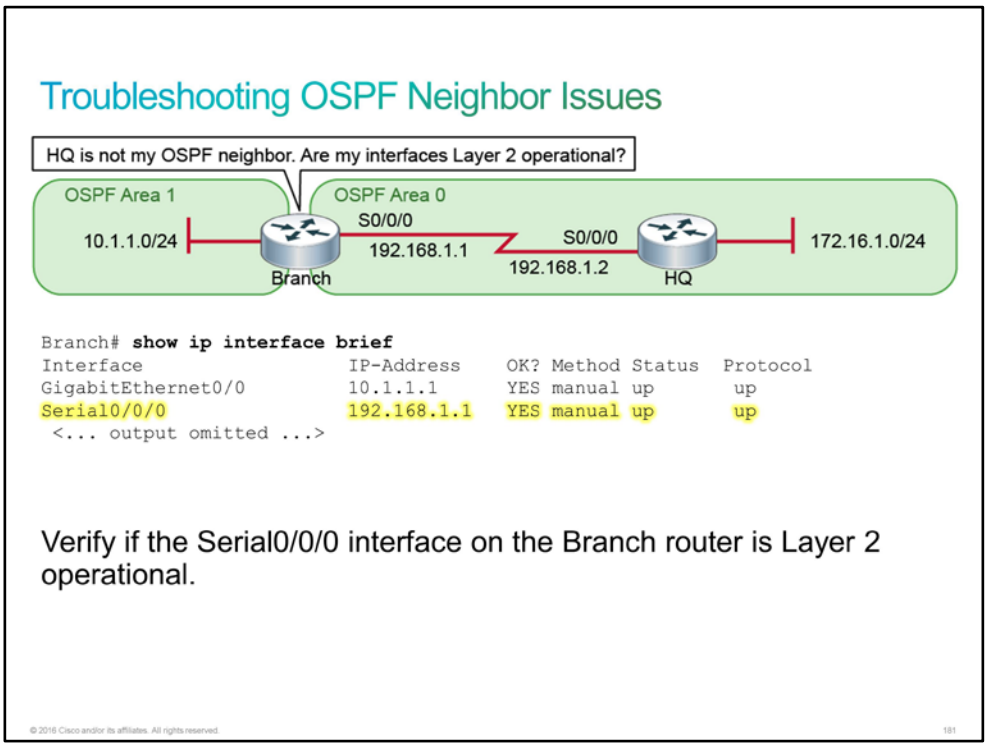

A prerequisite for the neighbor relationship to form between the Branch and HQ routers is the OSI Layer 3 connectivity. By investigating the **show ip interface brief** output, you can verify that the status and protocol are both "up" for the Serial0/0/0 interface that is connected to the Branch router. This instance confirms that the link is operational on Layer 2.

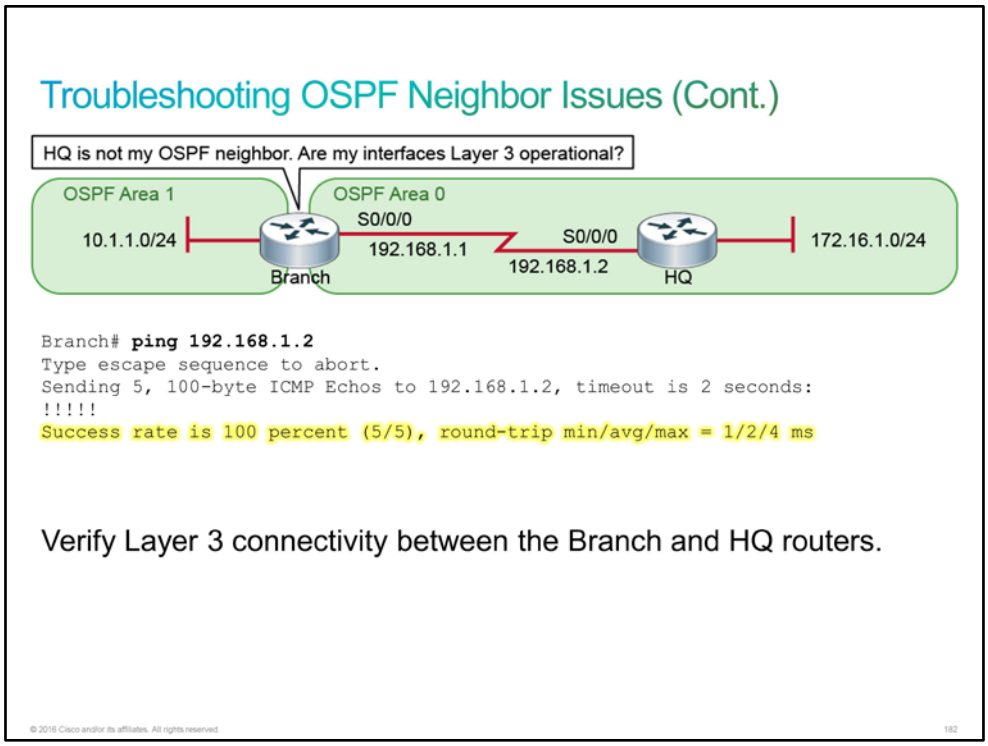

A **ping** from the Branch to the HQ router will confirm IP connectivity between the devices. If the **ping** is not successful, check the cabling and verify that the interfaces on connected devices are operational and that they are on a common subnet with the same subnet mask.

In the example, the Serial0/0/0 interface is enabled on both routers and there is connectivity between the Branch and HQ routers.

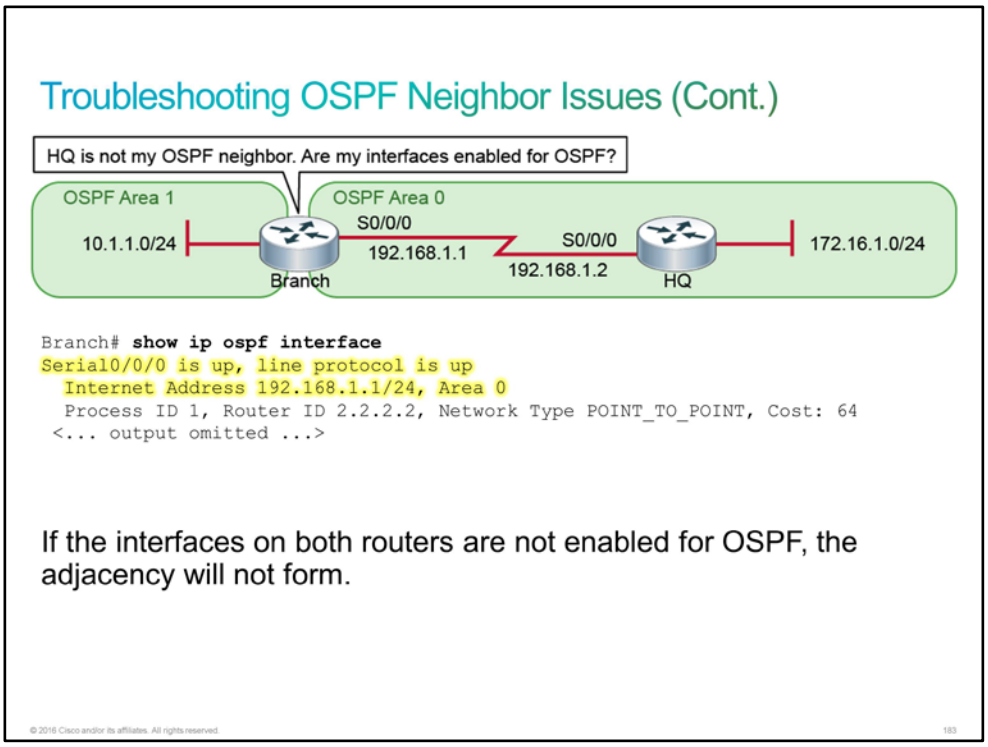

Note	The OSPF router ID for Branch is 2.2.2.2 and the OSPF router ID for HQ is 1.1.1.1.

If interfaces are operational, and there is IP connectivity between the devices, you have to verify that the interfaces on both routers are enabled for OSPF. If the interfaces on both router are not enabled for OSPF, the adjacency will not form. The **network** command that you configure under the OSPF routing process indicates which router interfaces will participate in OSPF. You can use the **show ip ospf interface** command to verify which interfaces are enabled for OSPF. The output will also show you which interface is functional and the OSPF-related parameters. If connected interfaces on two routers are not enabled for OSPF, the neighbors will not form an adjacency.

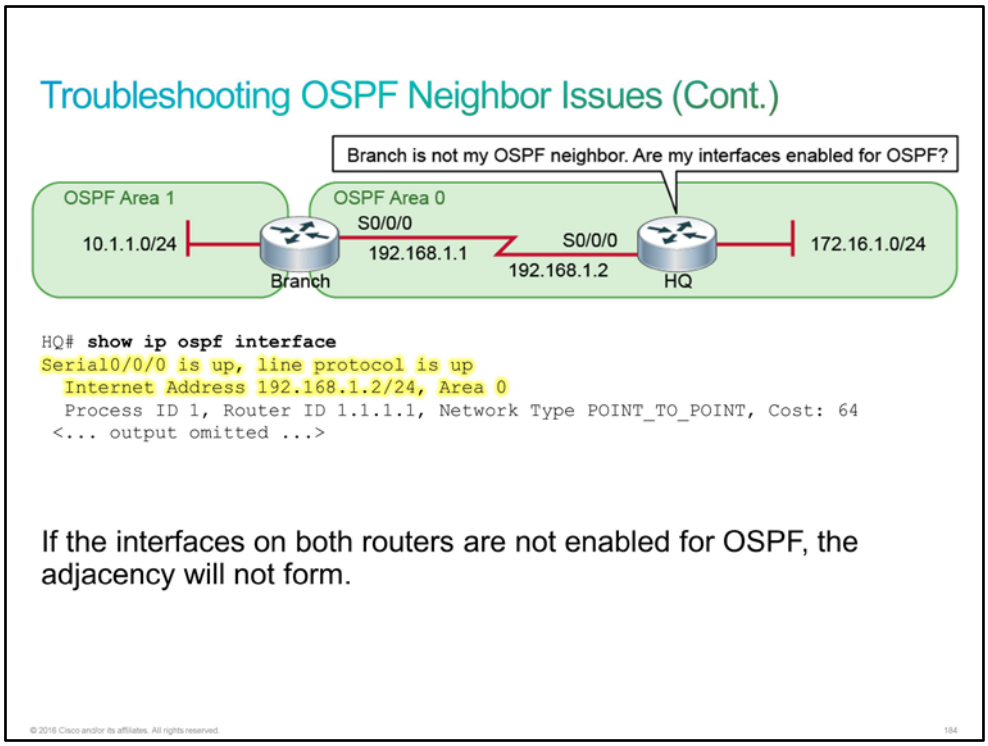

You can also use the **show ip protocols** command to verify which interfaces are configured for OSPF. The output will show you IP addresses or networks that are enabled using the **network** command. If an IP address on an interface falls within a network that has been enabled for OSPF, the interface will be enabled for OSPF. The output of this command will also show you if OSPF is enabled on an interface, using the **ip ospf area** command. The following is an example of the **show ip protocols** command:

```
HQ# show ip protocols
*** IP Routing is NSF aware ***

Routing Protocol is "ospf 1"
  Outgoing update filter list for all interfaces is not set
  Incoming update filter list for all interfaces is not set
  Router ID 1.1.1.1
  Number of areas in this router is 1. 1 normal 0 stub 0 nssa
  Maximum path: 4
  Routing for Networks:
    172.16.1.0 0.0.0.255 area 0
    192.168.1.0 0.0.0.255 area 0
  Routing on Interfaces Configured Explicitly (Area 0):
    Loopback0
 <... output omitted ...>
```

In the example, OSPF is enabled on the Serial0/0/0 interfaces on both routers.

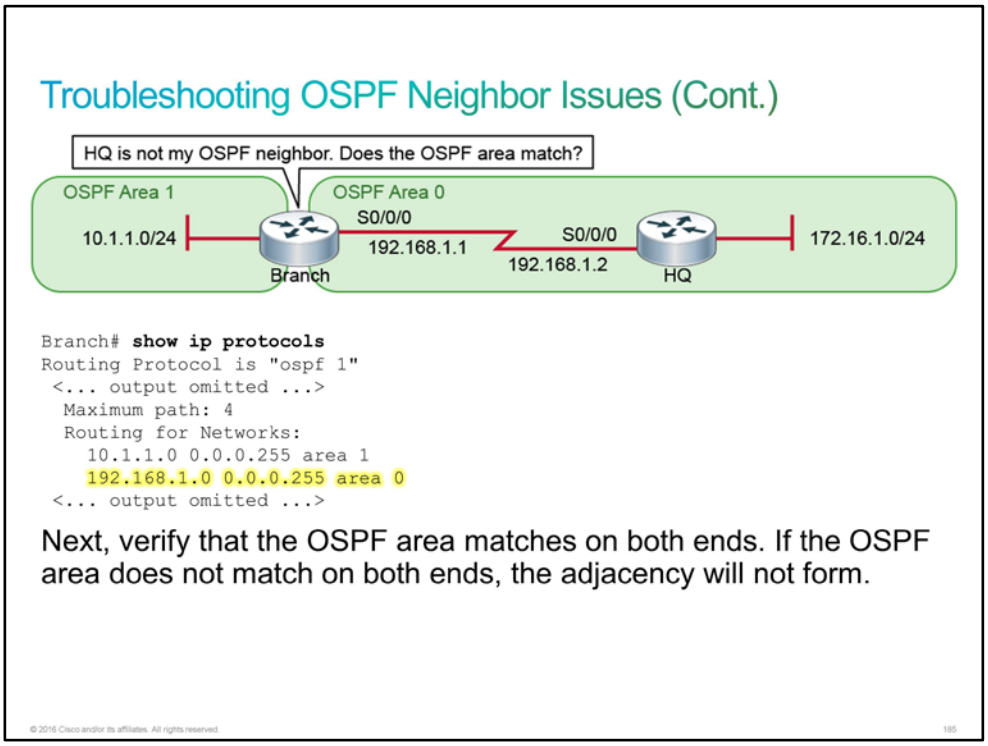

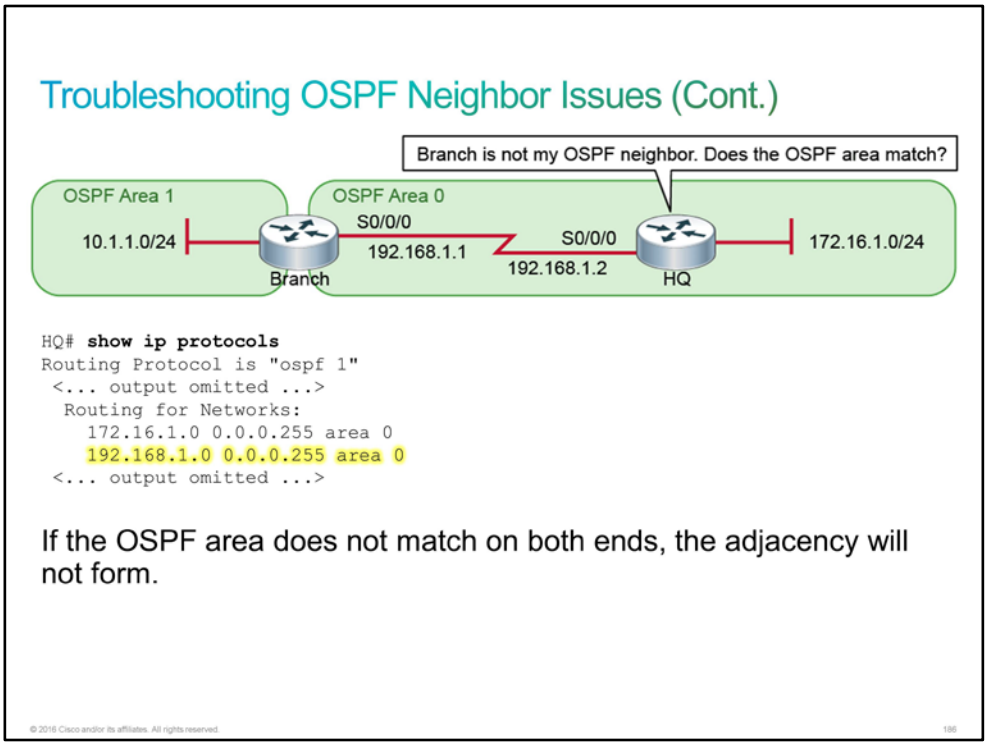

When you specify networks that will be advertised using OSPF, you have to provide the OSPF area number. The OSPF area numbers on two directly connected interfaces have to be the same, or the adjacency will not form. You can verify the area that an interface has been enabled for by using the **show ip protocols** command.

In the example, OSPF is enabled for the same area on both routers.

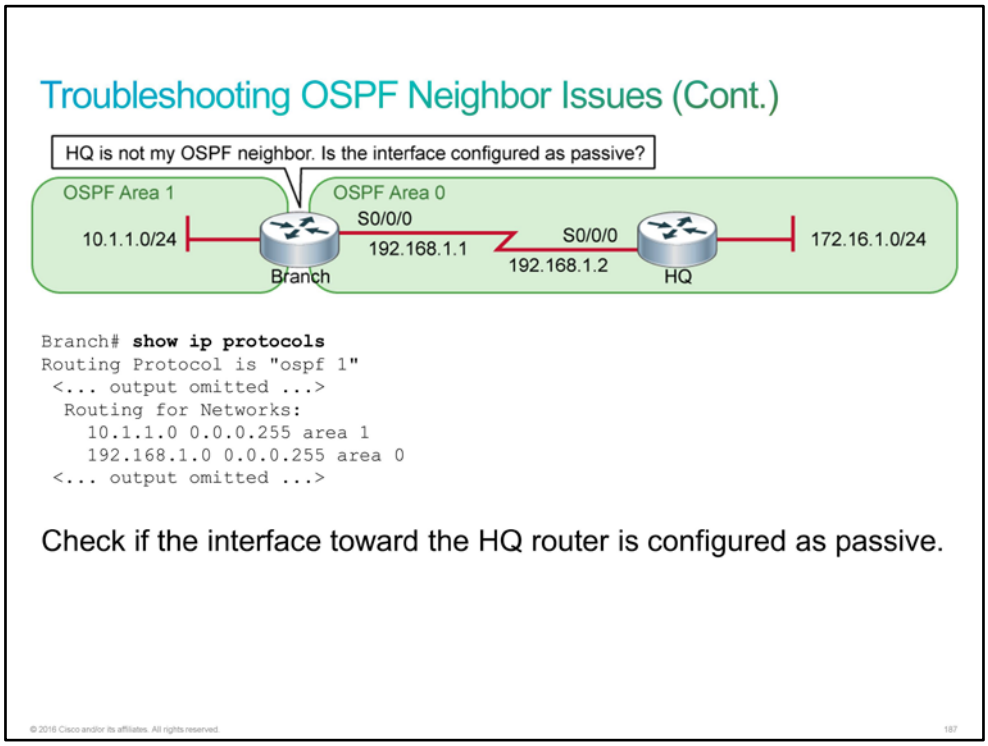

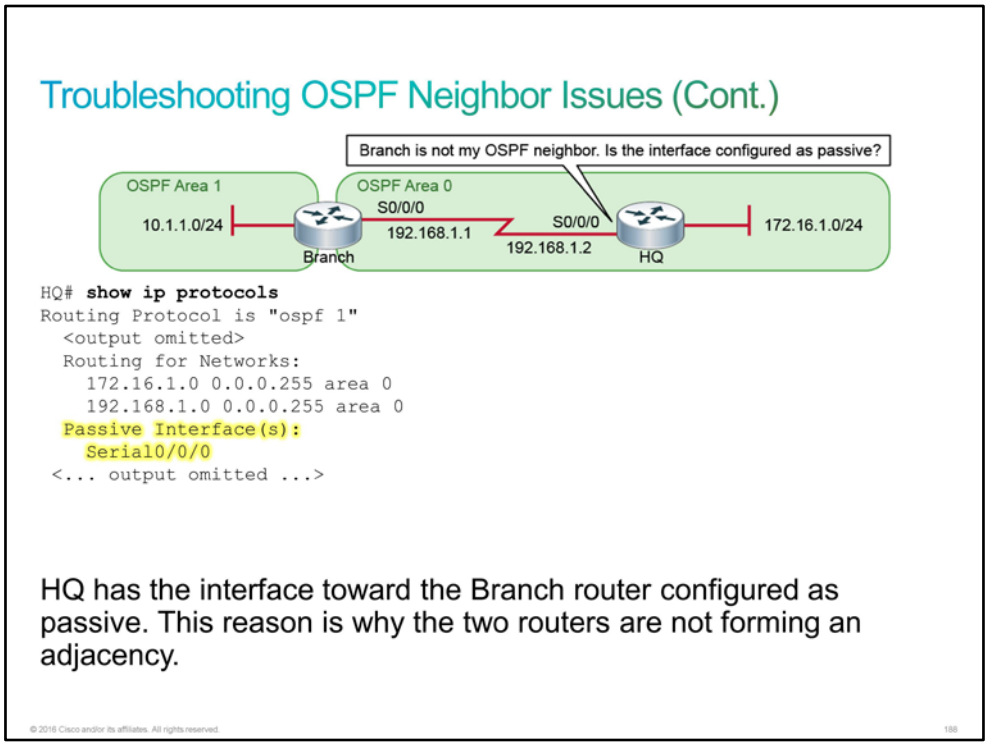

With OSPF running on a network, the **passive-interface** command stops both outgoing and incoming routing updates because the effect of the command causes the router to stop sending and receiving hello packets over an interface. For this reason, the routers will not become neighbors.

To verify if any interface on a router is configured as passive, use the **show ip protocols** command in the privileged mode.

An example in which you want to configure the interface as passive is handing off a link to a third-party organization that you have no control over (for example, an ISP). In this case, you would need to advertise this particular link through your own network but not allow the third party to receive hellos or send hellos to your device. This would be a security risk.

To configure an interface as a passive interface in OSPF, you will use the **passive-interface** *interface* command in the OSPF router configuration mode. To disable the interface as passive, use the **no passive-interface** *interface* command.

When you disable the passive interface, the routers should become adjacent, as indicated by the **show ip ospf neighbor** command output. Recall that two routers should be in the FULL state in order to exchange LSAs.

```
HQ# show ip ospf neighbor

Neighbor ID     Pri   State           Dead Time   Address         Interface
 2.2.2.2          0   FULL/  -        00:00:31    192.168.1.1     Serial0/0/0
```

Note	Routers will establish the FULL state only with the DR and BDR, while the established state will be two-way with other routers.

Troubleshooting OSPF Routing Table Issues

After you have verified that the adjacencies are correct, the next step is to verify the routing tables. The troubleshooting and verification components are as follows:

- If there is no route to a destination network in the routing table, verify that the OSPF neighbor is advertising the correct networks.
- If there is no OSPF route to a destination network in the routing table, verify if there is a routing protocol with a lower administrative distance configured in the network.

In the example, you will investigate why the Branch router cannot reach the networks that the HQ router is advertising.

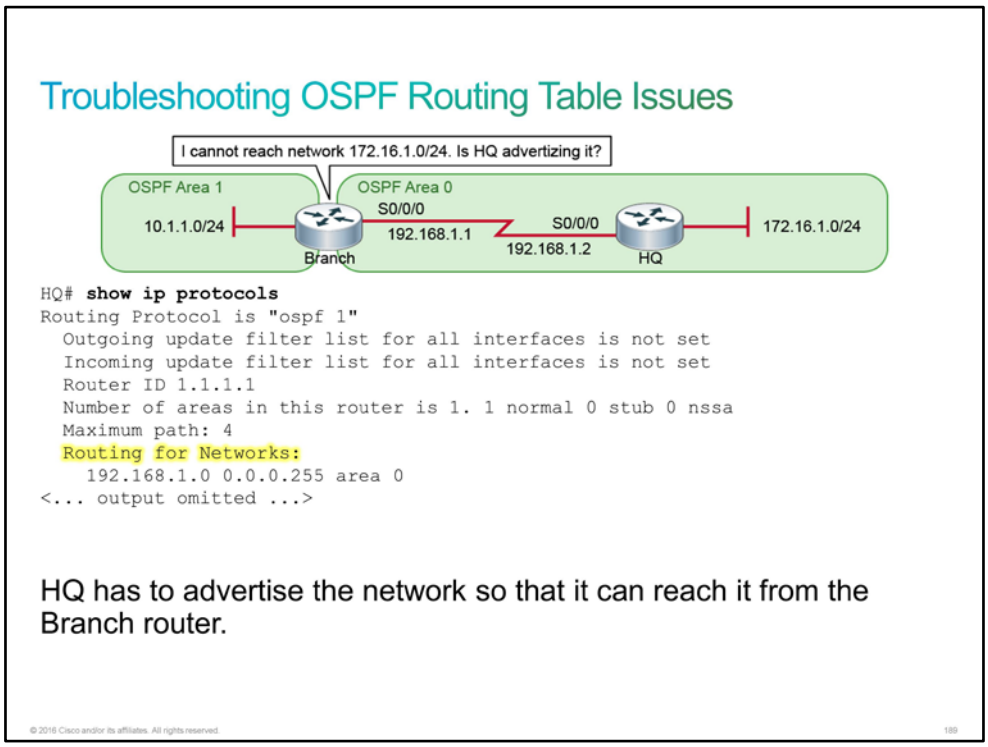

The Branch and HQ routers have their neighbor adjacency set up, but a **ping** test from the Branch router to a host in the 172.16.1.0/24 network is not successful. Checking the routing table of the Branch router leads you to the conclusion that there is a route missing to the destination network of 172.16.1.0/24.

You can use the **show ip protocols** command on the HQ router to verify if the 172.16.1.0/24 network is being advertised to the OSPF neighbors.

In the example, the HQ router is not configured to advertise the 172.16.1.0/24 network to the neighbor. To solve this issue, you have to start advertising this network on the HQ router.

Now, consider another scenario where several routing protocols are configured on a router. If several routing protocols are configured on routers, the administrative distance will decide which protocol the router will use.

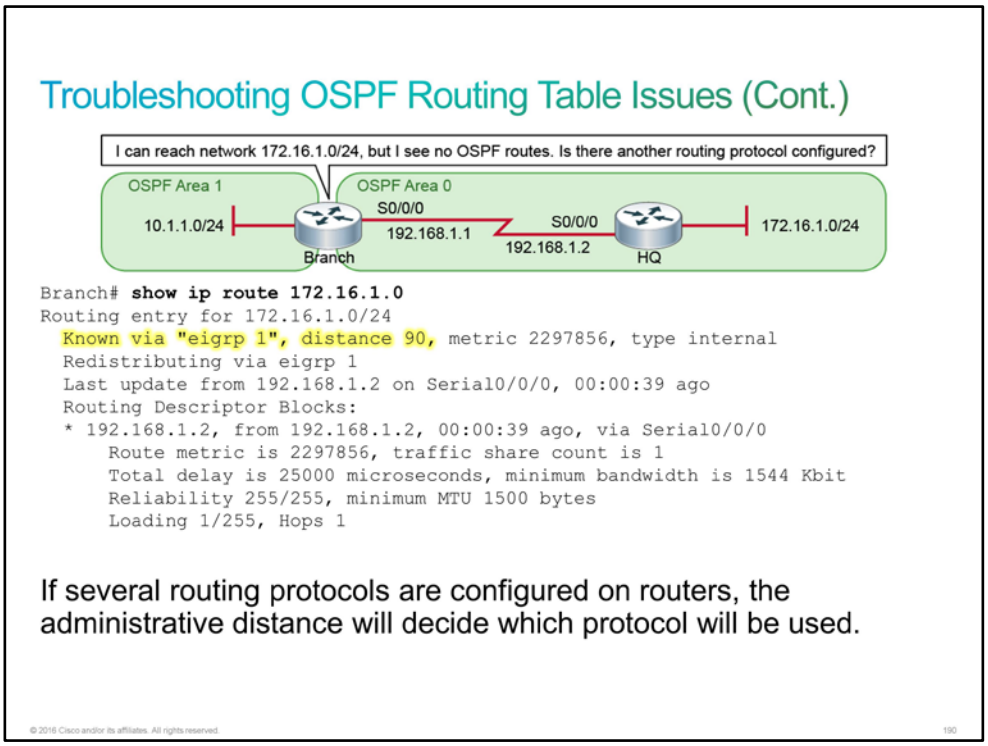

When you have more than one routing protocol configured in a network, you may receive routing information about a network through an undesired routing protocol. Recall that the routing protocol administrative distance influences which routes will be installed in the routing table. Although it does not affect connectivity, you may want to receive all routing information through the same routing protocol for the sake of easier troubleshooting and management. To verify which routing protocols are configured and their administrative distances, use the **show ip protocols** command:


```
Branch# show ip protocols
Routing Protocol is "ospf 1"
  Outgoing update filter list for all interfaces is not set
  Incoming update filter list for all interfaces is not set
  Router ID 2.2.2.2
  It is an area border router
  Number of areas in this router is 2. 2 normal 0 stub 0 nssa
  Maximum path: 4
  Routing for Networks:
    10.1.1.0 0.0.0.255 area 1
    192.168.1.0 0.0.0.255 area 0
 Reference bandwidth unit is 100 mbps
  Routing Information Sources:
    Gateway         Distance      Last Update
    1.1.1.1              80      00:02:37
  Distance: (default is 110)

Routing Protocol is "eigrp 1"
  Outgoing update filter list for all interfaces is not set
  Incoming update filter list for all interfaces is not set
  Default networks flagged in outgoing updates
  Default networks accepted from incoming updates
  EIGRP metric weight K1=1, K2=0, K3=1, K4=0, K5=0
  EIGRP maximum hopcount 100
  EIGRP maximum metric variance 1
  Redistributing: eigrp 1
  EIGRP NSF-aware route hold timer is 240s
  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
    (this router)         90      00:12:09
    192.168.1.2           90      00:02:39
  Distance: internal 90 external 170
```

In the example, the device received the route for 172.16.1.0/24 through EIGRP and OSPF. However, because EIGRP with the administrative distance of 90 is more trustworthy than OSPF with the administrative distance of 110, the device will install the EIGRP route in the routing table.

Troubleshooting OSPF Path Selection

Incorrect path selection doesn't usually lead to a loss of connectivity. However, certain links in a network should not be used, if possible. This case applies, for example, to backup WAN links, which can be charged by the amount of transferred data and can be expensive.

When you have redundant paths that are available in a network, you have to make sure that traffic takes the desired path through the network. For example, you could have two locations that are connected via the primary, high-speed link and via the dial-up, low-speed link for backup purposes. In this case, you have to make sure that the devices use the backup link only when the primary link fails.

Troubleshooting and verification components of OSPF path selection are as follows:

- If there are two OSPF paths to the destination network, verify the OSPF cost on both interfaces.

In the example, you will investigate the Branch router having two paths to the destination network on the HQ site.

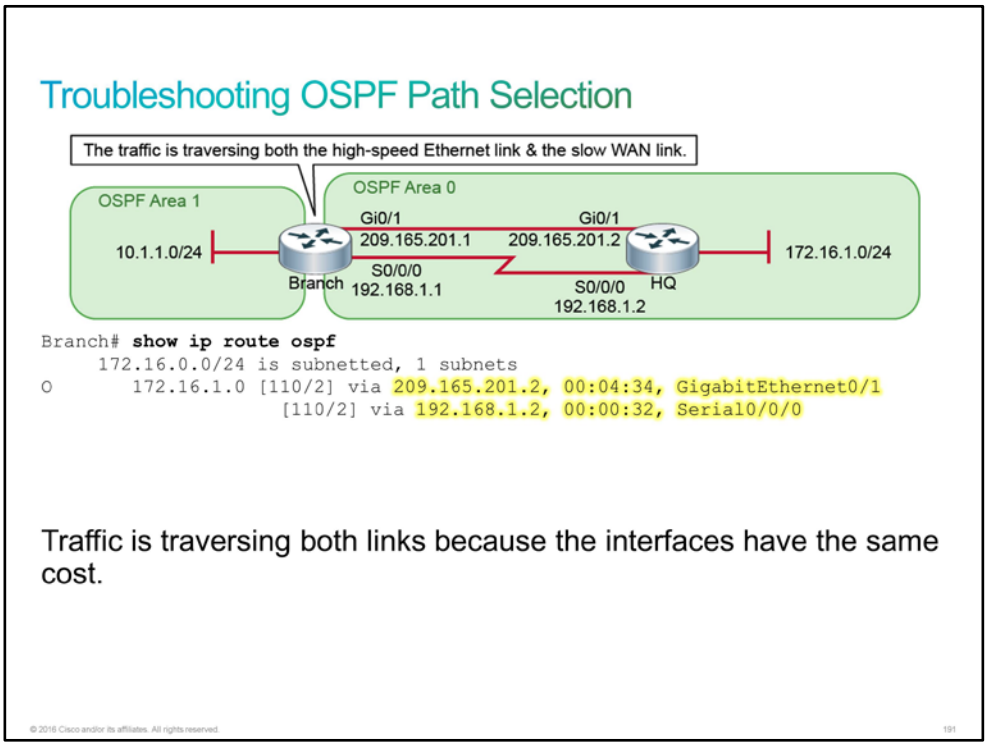

In the example, network 172.16.0.0/24 is reachable from the Branch router via the GigabitEthernet0/1 interface and the Serial0/0/0 interface. Because both interfaces have the same OSPF cost, load balancing across both links will be used. The reason for the same OSPF cost on both interfaces could be that someone manually changed the cost on the interfaces, or that there is incorrect reference bandwidth when managing interfaces that are faster than 100 Mbps. Recall that the OSPF cost is calculated as the interface bandwidth divided by the reference bandwidth, which is 100 Mbps by default. For example, with two interfaces—a 1000-Mbps and a 100-Mbps interface, both will have the same OSPF cost with a value of 1. In this case, you either need to increase the reference bandwidth to 1000 Mbps or manually change the OSPF cost on an interface to reflect the actual bandwidth of the interface.

Use the **show ip ospf interface** command to verify the OSPF cost on an interface:


```
Branch# show ip ospf interface
GigabitEthernet0/1 is up, line protocol is up
  Internet Address 209.165.201.2/27, Area 0
  Process ID 1, Router ID 1.1.1.1, Network Type BROADCAST, Cost: 1
<... output omitted ...>
Serial0/0/0 is up, line protocol is up
  Internet Address 192.168.1.2/24, Area 0
  Process ID 1, Router ID 1.1.1.1, Network Type POINT_TO_POINT, Cost: 1
```

When you increase the OSPF cost on the Serial0/0/0 interfaces on both routers, only the preferred route will be installed in the routing table:

```
Branch(config)# interface Serial0/0/0
Branch(config-if)# ip ospf cost 10

Branch# show ip route ospf
     172.16.0.0/24 is subnetted, 1 subnets
O       172.16.1.0 [110/2] via 209.165.201.2, 00:14:31, GigabitEthernet0/1
```

The reference bandwidth can be changed from the default of 100 Mbps. You can verify what the reference bandwidth is by using the **show ip ospf** command:

```
Branch# show ip ospf
<... output omitted ...>
 Cisco NSF helper support enabled
 Reference bandwidth unit is 100 mbps
    Area BACKBONE(0)
        Number of interfaces in this area is 4
        Area has no authentication
<... output omitted ...>
```

A changed reference bandwidth means a changed cost on the link. Make sure that all the routers within the OSPF AS have the same reference bandwidth. Change it by using the **auto-cost reference-bandwidth** *bandwidth_in_Mbits_per_second* command from the router OSPF configuration mode.

Troubleshooting OSPFv3 Issues

The OSPFv3 runs on IPv6 and it uses IPv6 link-local addresses as the source of hello packets and next-hop calculations. Configuring OSPFv3 is very similar to configuring OSPF for IPv4. The main difference is that OSPFv3 is enabled on the interface for IPv6 with the **ipv6 ospf** *process-id* **area** *area-id* command. Therefore, troubleshooting OSPFv3 is very similar to troubleshooting OSPF for IPv4.

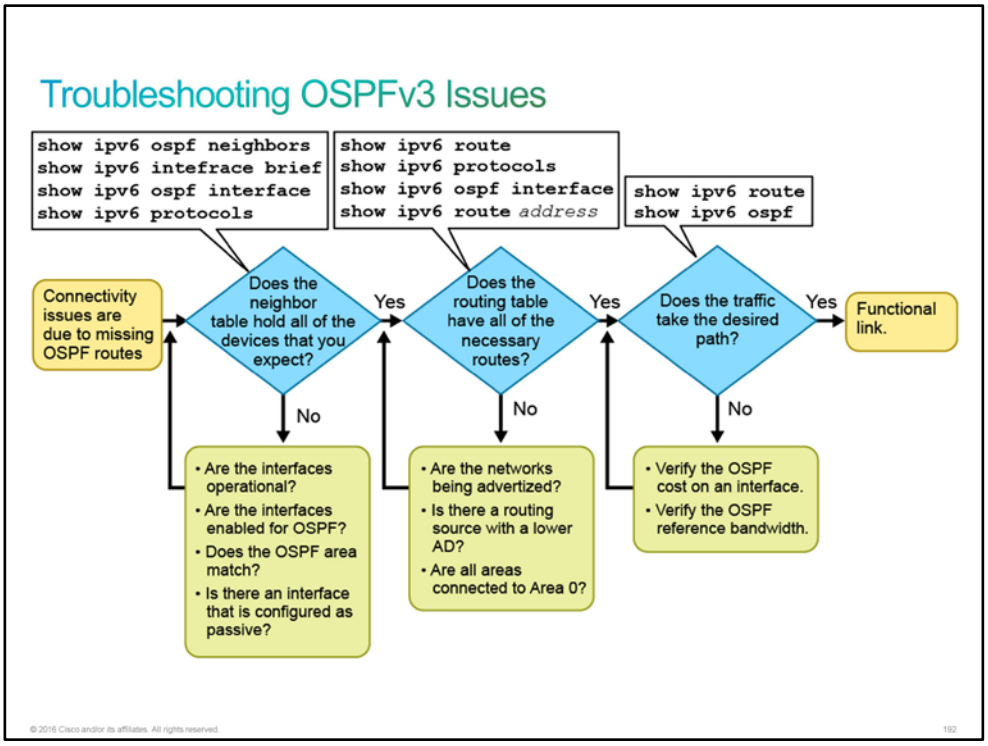

To check the IPv6 routing protocols on the router, use the **show ipv6 protocols** command. The output will show the IPv6 routing protocols that are enabled on the router. The OSPF section shows the router ID, OSPF interfaces, and so on.

To display the neighbors that OSPFv3 discovers, use the **show ipv6 ospf neighbors** command. If you want to display the interfaces that are enabled for OSPFv3 and their costs, issue the **show ipv6 ospf interfaces** command.

The **show ipv6 route ospf** command will show you the content of the IPv6 routing table, which includes the routes that are specific to OSPF.

Note	Remember that for OSPFv3 to work, IPv6 routing must be enabled.

Note	Remember that if no IPv4 is configured configured on the router, you need to manually configure the router ID for the OSPFv3 routing process.

Discovery 46: Troubleshoot Multiarea OSPF

Introduction

This discovery will guide you through the troubleshooting of various OSPF configuration issues. The 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 their hostnames and IP addresses. OSPF has been configured on all seven routers, but there are problems with the router configurations. Each router has a loopback interface with the IP address 192.168.R.1/24 (where R indicates the router number). The routing table on R1 is missing routes to the loopback interface networks for each of its peers. In this discovery, you will troubleshoot and fix the problem that is associated with the routing of each of these networks.

You will start with the R2 loopback network, and then proceed one at a time, finishing with R7, which is also configured for OSPFv3. In each case, you will first determine the root cause and then you will fix the issue and verify that the route is properly defined in the routing table of R1.

Topology

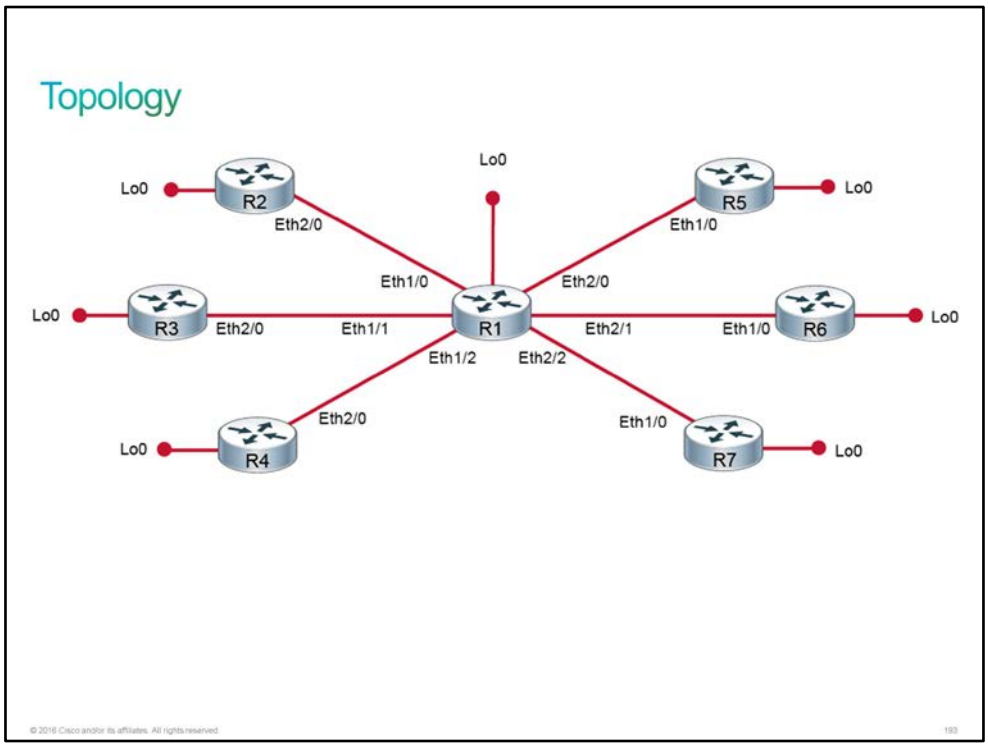

Job Aid

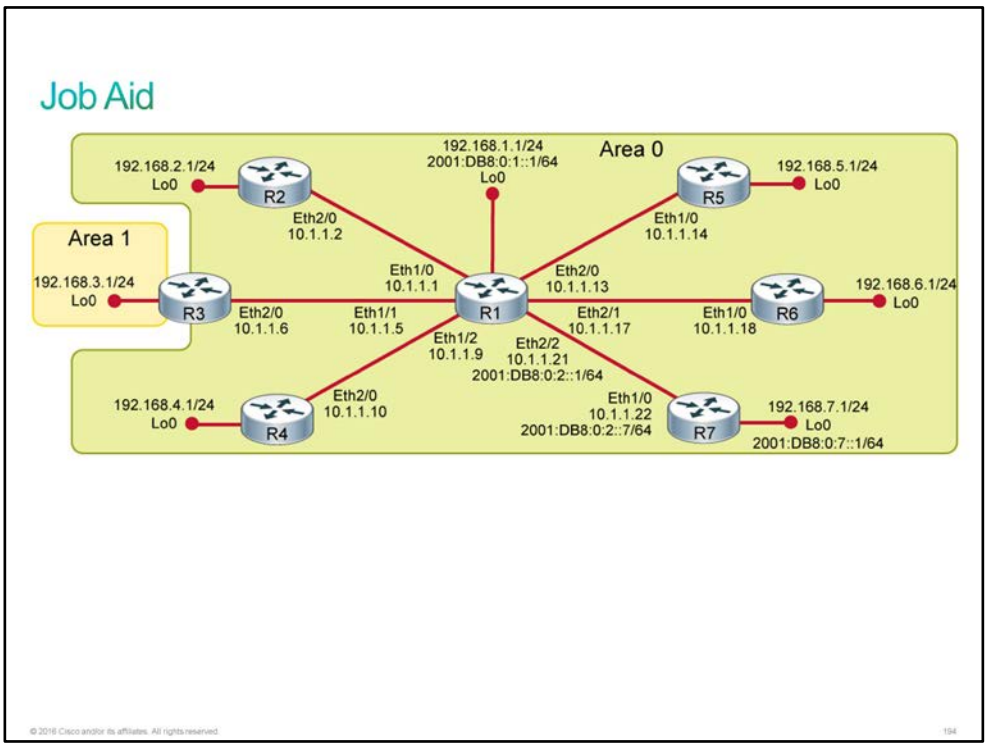

The configuration is as follows:

- All devices have their basic configurations in place including hostnames, and IP addresses. R1 and R7 also have IPv6 addresses configured.
- OSPF AS 100 has been configured on all seven routers, but there are problems with the router configurations.
  - The routing table on R1 is missing routes to the loopback interface networks for each of its peers.
- R1 and R7 are also configured for OSPv3 routing, using AS 100.

Device Information

Device Details

Device	Interface	Neighbor	IP Address
R1	Loopback0	—	192.168.1.1/24
R1	Ethernet1/0	R2	10.1.1.1/30
R1	Loopback0	—	2001:DB8:0:1::1/64
R1	Ethernet2/2	R7	2001:DB8:0:2::1/64
R1	Ethernet1/1	R3	10.1.1.5/30

Device	Interface	Neighbor	IP Address
R1	Ethernet1/2	R4	10.1.1.9/30
R1	Ethernet2/0	R5	10.1.1.13/30
R1	Ethernet2/1	R6	10.1.1.17/30
R1	Ethernet2/2	R7	10.1.1.21/30
R2	Ethernet2/0	R1	10.1.1.2/30
R2	Loopback0	—	192.168.2.1/24
R3	Ethernet2/0	R1	10.1.1.6/30
R3	Loopback0	—	192.168.3.1/24
R4	Ethernet2/0	R1	10.1.1.10/30
R4	Loopback0	—	192.168.4.1/24
R5	Ethernet1/0	R1	10.1.1.14/30
R5	Loopback0	—	192.168.5.1/24
R6	Ethernet1/0	R1	10.1.1.18/30
R6	Loopback0	—	192.168.6.1/24
R7	Ethernet1/0	R1	10.1.1.22/30
R7	Loopback0	—	192.168.7.1/24
R7	Ethernet1/0	R1	2001:DB8:0:2::7/64
R7	Loopback0	—	2001:DB8:0:7::1/64

Task 1: Troubleshoot Multiarea OSPF

Activity

Complete the following steps:

Step 1 The network 192.168.2.0/24 does not exist in the routing table of R1. This network is associated with the Loopback0 interface of R2. Try to determine the root cause of this issue.

There is no single best process to troubleshoot the problem. Use the **show running-config** command only after you have a good idea of where the problem is.

Note: There is no single best procedure for troubleshooting any network issue. The goal is to isolate the root cause. One strategy is to work from the application layer down. If there are aspects of the application that are working, it implies that there must be an IP connectivity and link layer connectivity below the application. If the application does not function, check the IP connectivity next. If the IP connectivity appears to be working, look at the link layer connectivity.

You might use the following commands on R1 and observe these results:

- **show ip ospf neighbor**—R2 (0.0.0.2) is in the OSPF neighbor table. Given that the OSPF neighbor relationship is established, you can assume that both Layer 2 and Layer 3 connectivity is functioning properly below OSPF.

Given the results on R1, you might skip troubleshooting Layer 2 (with the **show cdp neighbor** and **show ip interface brief** commands, and so on) and Layer 3 connectivity (with the **ping** and **show ip interface brief** commands, and so on) by going straight to looking more closely at the OSPF configuration.

- **show ip protocols**—Everything looks like it is appropriately configured.

```
R2# show ip ospf interface brief
Interface    PID   Area            IP Address/Mask    Cost  State Nbrs F/C
Lo0          100   0               192.168.2.1/24     1     DOWN  0/0
Et2/0        100   0               10.1.1.2/30        10    BDR   1/1

 R2# show ip interface brief
Interface                  IP-Address      OK? Method Status
Protocol
<... output omitted ...>
Loopback0                  192.168.2.1     YES manual administratively down
down
```

There was no Layer 3 issue between R1 and R2, but a Layer 2 issue is associated with the Loopback0 interface.

- The **show running-configuration interface Loopback0** command will show you the root cause.

```
R2# show running-config interface Loopback0
Building configuration...

Current configuration : 147 bytes
!
interface Loopback0
 description Logical loopback interface
 ip address 192.168.2.1 255.255.255.0
 ip ospf network point-to-point
 shutdown
end
```

Step 2 With the root cause determined, fix the problem and verify that the route to 192.168.2.0/24 now exists in the routing table of R1.

Enter the following configuration on the R2 router:


```
R2# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
R2(config)# interface Loopback0
R2(config-if)# no shut
R2(config-if)#
*Oct 26 13:54:34.236: %LINK-3-UPDOWN: Interface Loopback0, changed state to up
*Oct 26 13:54:35.240: %LINEPROTO-5-UPDOWN: Line protocol on Interface
Loopback0, changed state to up
```

The network is present in the routing table of R1 after you enable the loopback interface.

```
R1# show ip route 192.168.2.0
Routing entry for 192.168.2.0/24
  Known via "ospf 100", distance 110, metric 11, type intra area
  Last update from 10.1.1.2 on Ethernet1/0, 00:02:40 ago
  Routing Descriptor Blocks:
  * 10.1.1.2, from 0.0.0.2, 00:02:40 ago, via Ethernet1/0
      Route metric is 11, traffic share count is 1
```

Step 3 The network 192.168.3.0/24 does not exist in the routing table of R1. This network is associated with the Loopback0 interface of R3. Try to determine the root cause of this issue.

You might use the following commands on R1 and observe these results:

- **show ip ospf neighbor**—R3 is not an OSPF neighbor of R1.
- **show ip ospf interface brief**—Ethernet1/1 is running OSPF in area 0.
- **show cdp neighbor**—R3 is a neighbor on Ethernet1/1 (there is Layer 2 connectivity).
- **ping 10.1.1.6**—There is Layer 3 connectivity between R1 and R3.

From this information, it is apparent that there are no problems with the link layer or the IP layer, but that something is not working with OSPF.

When your investigation moves to R3, the syslog messages that are waiting for you make the root cause quite apparent.

```
*Oct 26 14:06:20.811: %OSPF-4-ERRRCV: Received invalid packet: mismatched area
ID, from backbone area must be virtual-link but not found from 10.1.1.5,
Ethernet2/0
```

All R1 interfaces are configured for area 0. All the directly neighboring interfaces on the peer routers must also be configured for area 0.

```
R3# show ip protocols
*** IP Routing is NSF aware ***

Routing Protocol is "ospf 100"
  Outgoing update filter list for all interfaces is not set
  Incoming update filter list for all interfaces is not set
  Router ID 0.0.0.3
  Number of areas in this router is 1. 1 normal 0 stub 0 nssa
  Maximum path: 4
  Routing for Networks:
    10.0.0.0 0.255.255.255 area 1
    192.168.3.0 0.0.0.255 area 1
  Routing Information Sources:
    Gateway         Distance      Last Update
  Distance: (default is 110)


R3# show running-config | section router ospf
router ospf 100
 router-id 0.0.0.3
 network 10.0.0.0 0.255.255.255 area 1
 network 192.168.3.0 0.0.0.255 area 1
```

Step 4 With the root cause determined, fix the problem and verify that the route to 192.168.3.0/24 now exists in the routing table of R1.

If you want the route to 192.168.3.0/24 to be distributed through the entire virtual lab topology, you can either move both network statements to area 0 (making R3 a backbone router), or you can move the network 10.0.0.0 to area 0 (making R3 an area border router). The following output shows the second option:

```
R3# conf t
R3(config)# router ospf 100
R3(config-router)# network 10.0.0.0 0.255.255.255 area 0
R3(config-router)#
*Oct 26 14:09:48.274: %OSPF-6-AREACHG: 10.0.0.0/8 changed from area 1 to area 0
R3(config-router)#
*Oct 26 14:09:56.320: %OSPF-5-ADJCHG: Process 100, Nbr 0.0.0.1 on Ethernet2/0
from LOADING to FULL, Loading Done
```

The neighbor relationship is initiated between R3 and R1 when the OSPF area is configured to match.


```
R1# show ip route 192.168.3.0
Routing entry for 192.168.3.0/24
  Known via "ospf 100", distance 110, metric 11, type inter area
  Last update from 10.1.1.6 on Ethernet1/1, 00:01:28 ago
  Routing Descriptor Blocks:
  * 10.1.1.6, from 0.0.0.3, 00:01:28 ago, via Ethernet1/1
      Route metric is 11, traffic share count is 1

R1# show ip route ospf
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

O     192.168.2.0/24 [110/11] via 10.1.1.2, 00:17:22, Ethernet1/0
O IA  192.168.3.0/24 [110/11] via 10.1.1.6, 00:02:01, Ethernet1/1
O     192.168.7.0/24 [110/11] via 10.1.1.22, 00:33:47, Ethernet2/2
```

Because 192.168.3.0/24 was left in area 1 on R3, the route in the routing table of R1 is an OSPF interarea route.

Step 5 The network 192.168.4.0/24 does not exist in the R1 routing table. This network is associated with the Loopback0 interface of R4. Try to determine the root cause of this issue.

You might use the following commands on R1 and observe the results:

- **show ip ospf neighbor**—R4 is not an OSPF neighbor.
- **show ip ospf interface brief**—Ethernet1/2 is running OSPF in area 0.
- **show cdp neighbor**—R4 is a Cisco Discovery Protocol neighbor on Ethernet1/2 (there is Layer 2 connectivity).
- **ping 10.1.1.10**—There is Layer 3 connectivity between R1 and R4.

From these results, you might conclude that there are no issues at the IP layer or the data link layer and that R1 is properly configured for OSPF. You might then investigate the OSPF configuration on R4.

```
R4# show ip protocols
*** IP Routing is NSF aware ***

Routing Protocol is "ospf 100"
  Outgoing update filter list for all interfaces is not set
  Incoming update filter list for all interfaces is not set
  Router ID 0.0.0.4
  Number of areas in this router is 1. 1 normal 0 stub 0 nssa
  Maximum path: 4
  Routing for Networks:
    10.1.1.0 0.0.0.3 area 0
    192.168.4.0 0.0.0.255 area 0
  Routing Information Sources:
    Gateway         Distance      Last Update
  Distance: (default is 110)
```

It may not be immediately obvious, but there is a problem with the networks that are included under OSPF.

```
R4# show ip ospf interface brief
Interface    PID   Area            IP Address/Mask    Cost  State Nbrs F/C
Lo0          100   0               192.168.4.1/24     1     P2P   0/0
R4# show ip interface Ethernet2/0
Ethernet2/0 is up, line protocol is up
  Internet address is 10.1.1.10/30
<... output omitted ...>
```

Note: The IP address of Ethernet2/0 is not included in the 10.1.1.0/30 subnet.

Step 6 With the root cause determined, fix the problem and verify that the route to 192.168.4.0/24 now exists in the routing table of R1.

Enter the following configuration on the R4 router:

```
R4# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
R4(config)# router ospf 100
R4(config-router)# no network 10.1.1.0 0.0.0.3 area 0
R4(config-router)# network 10.1.1.8 0.0.0.3 area 0
R4(config-router)#
*Oct 26 14:44:47.937: %OSPF-5-ADJCHG: Process 100, Nbr 0.0.0.1 on Ethernet2/0
from LOADING to FULL, Loading Done
```

When the network statements are updated to include 10.1.1.10, the neighbor relationship with R1 is established.

```
R1# show ip route 192.168.4.0
Routing entry for 192.168.4.0/24
  Known via "ospf 100", distance 110, metric 11, type intra area
  Last update from 10.1.1.10 on Ethernet1/2, 00:03:52 ago
  Routing Descriptor Blocks:
  * 10.1.1.10, from 0.0.0.4, 00:03:52 ago, via Ethernet1/2
      Route metric is 11, traffic share count is 1
```

Step 7 The network 192.168.5.0/24 does not exist in the routing table of R1. This network is associated with the Loopback0 interface of R5. Try to determine the root cause of this issue.

You might use the following commands on R1 and observe these results:

- **show ip ospf neighbor**—R5 is not an OSPF neighbor.
- **show ip ospf interface brief**—Ethernet2/0 is running OSPF in area 0.
- **show cdp neighbor**—R5 is a neighbor on Ethernet2/0 (there is Layer 2 connectivity).
- **ping 10.1.1.1**4—There is Layer 3 connectivity between R1 and R5.

From these results, you might conclude that there are no issues at the IP layer or the data link layer and that R1 is properly configured for OSPF. You might then move to investigate the OSPF configuration on R5.

You might use the following command on R5 and observe these results:

- **show ip ospf interface** (potentially with the **brief** argument added)—Both Ethernet1/0 and Loopback0 are participating in OSPF.


```
R5# show ip protocols
*** IP Routing is NSF aware ***

Routing Protocol is "ospf 100"
  Outgoing update filter list for all interfaces is not set
  Incoming update filter list for all interfaces is not set
  Router ID 0.0.0.5
  Number of areas in this router is 1. 1 normal 0 stub 0 nssa
  Maximum path: 4
  Routing for Networks:
    10.0.0.0 0.255.255.255 area 0
    192.168.5.0 0.0.0.255 area 0
  Passive Interface(s):
    Ethernet1/0
  Routing Information Sources:
    Gateway         Distance      Last Update
  Distance: (default is 110)
```

The root cause of the issue is that Ethernet1/0 is configured as a passive interface, preventing the OSPF neighbor relationships from forming on that interface.

```
R5# show running-config | section router ospf
router ospf 100
 router-id 0.0.0.5
 passive-interface Ethernet1/0
 network 10.0.0.0 0.255.255.255 area 0
 network 192.168.5.0 0.0.0.255 area 0
```

Step 8 With the root cause determined, fix the problem and verify that the route to 192.168.5.0/24 now exists in the routing table of R1.

Enter the following configuration on the R5 router:

```
R5# conf t
R5(config)# router ospf 100
R5(config-router)# no passive-interface Ethernet1/0
R5(config-router)#
*Oct 26 14:56:41.387: %OSPF-5-ADJCHG: Process 100, Nbr 0.0.0.1 on Ethernet1/0
from LOADING to FULL, Loading Done
```

Immediately after you remove the passive interface restriction, the neighbor relationship is established between R5 and R1.

```
R1# show ip route 192.168.5.0
Routing entry for 192.168.5.0/24
  Known via "ospf 100", distance 110, metric 11, type intra area
  Last update from 10.1.1.14 on Ethernet2/0, 00:01:03 ago
  Routing Descriptor Blocks:
  * 10.1.1.14, from 0.0.0.5, 00:01:03 ago, via Ethernet2/0
      Route metric is 11, traffic share count is 1
```

Step 9 The network 192.168.6.0/24 does not exist in the routing table of R1. This network is associated with the Loopback0 interface of R6. Try to determine the root cause of this issue.

You might use the following commands on R1 and observe these results:

- **show ip ospf neighbor**—R6 is an OSPF neighbor.

From this result, you can see that there is some OSPF functionality, which implies that Layer 3 and Layer 2 are working to support the OSPF communication. You might discount further investigation on R1 and focus on R6.

```
R6# show ip protocols
*** IP Routing is NSF aware ***

Routing Protocol is "ospf 100"
  Outgoing update filter list for all interfaces is not set
  Incoming update filter list for all interfaces is not set
  Router ID 0.0.0.6
  Number of areas in this router is 1. 1 normal 0 stub 0 nssa
  Maximum path: 4
  Routing for Networks:
    10.0.0.0 0.255.255.255 area 0
  Routing Information Sources:
    Gateway         Distance      Last Update
    0.0.0.1              110      03:25:41
    0.0.0.2              110      03:08:11
    0.0.0.3              110      02:52:40
    0.0.0.4              110      02:17:48
    0.0.0.5              110      02:05:55
    0.0.0.7              110      03:24:26
  Distance: (default is 110)
```

OSPF is not routing for the network 192.168.6.0/24. So, it is not willing to advertise this network to the OSPF area.

```
R6# show ip ospf interface brief
Interface    PID   Area            IP Address/Mask    Cost  State Nbrs F/C
Et1/0        100   0               10.1.1.18/30       10    BDR   1/1

R6# show running-config | section router ospf
router ospf 100
 router-id 0.0.0.6
 network 10.0.0.0 0.255.255.255 area 0
```

Step 10 With the root cause determined, fix the problem and verify that the route to 192.168.6.0/24 now exists in the routing table of R1.

```
R6# conf t
R6(config)# router ospf 100
R6(config-router)# network 192.168.6.0 0.0.0.255 area 0
R6(config-router)# end
```

Because the neighbor relationship is already in place, there were no syslog messages to indicate an OSPF state change. Before examining the R1 routing table, you may want to verify that the OSPF interface list on R6 now includes Loopback0.

```
R6# show ip ospf interface brief
Interface    PID   Area            IP Address/Mask    Cost  State Nbrs F/C
Lo0          100   0               192.168.6.1/24     1     P2P   0/0
Et1/0        100   0               10.1.1.18/30       10    BDR   1/1
```



```
R1# show ip route 192.168.6.0
Routing entry for 192.168.6.0/24
  Known via "ospf 100", distance 110, metric 11, type intra area
  Last update from 10.1.1.18 on Ethernet2/1, 00:01:55 ago
  Routing Descriptor Blocks:
  * 10.1.1.18, from 0.0.0.6, 00:01:55 ago, via Ethernet2/1
      Route metric is 11, traffic share count is 1
```

Step 11 The OSPFv3 neighbor relationship is not established between R1 and R7. Try to determine the root cause of this issue.

You might use the following commands on R1 and observe these results:

- **show ipv6 ospf neighbor**—R7 is not an OSPFv3 neighbor.
- **show ipv6 ospf interface brief**—Ethernet2/2 is running OSPFv3 100 in area 0.
- **show cdp neighbor**—R7 is a neighbor on Ethernet2/2 (there is Layer 2 connectivity).
- **ping 2001:db8:0:2::7**—There is Layer 3 connectivity between R1 and R7.

From these results, you might conclude that there are no issues at the IP layer or the data link layer and that R1 is properly configured for OSPFv3. You might then move to investigate the OSPFv3 configuration on R7.

You might use the following command on R7 and observe these results:

- **show ipv6 ospf interface** (potentially with the **brief** argument added)—No interfaces are participating in OSPFv3.

```
R7# show ipv6 protocols
IPv6 Routing Protocol is "connected"
IPv6 Routing Protocol is "ND"
R7#
```

The root cause of the issue is that R7 is missing OSPFv3 configuration.

```
R7# show running-config | include ipv6 unicast-routing
R7#
```

IPv6 routing is not enabled on R7. Someone could have configured IPv6 routing and OSPFv3, but when the IPv6 routing was disabled, all OSPFv3 configuration was removed.

Step 12 With the root cause determined, fix the problem and verify that OSPFv3 neighbor relationship now exists between R1 and R7 routers.

Enter the following configuration on the R7 router:

```
R7# conf t
R7(config)# ipv6 unicast-routing
R7(config)# ipv6 router ospf 100
R7(config-rtr)# router-id 0.0.0.7
R7(config-rtr)# exit
R7(config)# interface Ethernet1/0
R7(config-if)# ipv6 ospf 100 area 0
R7(config-if)# end
*Oct 27 08:20:10.544: %OSPFv3-5-ADJCHG: Process 100, Nbr 0.0.0.1 on Ethernet1/0
from LOADING to FULL, Loading Done
```

The neighbor adjacency is initiated when you enable IPv6 routing on R7 and configure OSPFv3 on the interface facing R1 (Ethernet1/0).

```
R1# show ipv6 ospf neighbor

            OSPFv3 Router with ID (0.0.0.1) (Process ID 100)

Neighbor ID     Pri   State           Dead Time   Interface ID    Interface
0.0.0.7           1   FULL/BDR        00:00:37    7               Ethernet2/2
```

Step 13 The network 2001:DB8:0:7::/64 still does not exist in the IPv6 routing table of R1. This network is associated with the Loopback0 interface of R7. Try to determine the root cause of this issue.

You might use the following commands on R7 and observe these results:

- **show ipv6 ospf interface brief**—Only Ethernet1/0 is included (Loopback0 is missing).

```
R7# show ipv6 ospf interface brief
Interface    PID   Area            Intf ID    Cost  State Nbrs F/C
Et1/0        100   0               7          10    BDR   1/1
```

This part is the root cause of the issue. However, it is also expected because you configured only Ethernet1/0 on R7 for OSPFv3. You can also verify this part in the running configuration.

```
R7# show running-config interface Loopback0
Building configuration...

Current configuration : 170 bytes
!
interface Loopback0
 description Logical loopback interface
 ip address 192.168.7.1 255.255.255.0
 ip ospf network point-to-point
 ipv6 address 2001:DB8:0:7::1/64
end

R7# show running-config interface Ethernet1/0
Building configuration...

Current configuration : 146 bytes
!
interface Ethernet1/0
 description Link to R1
 ip address 10.1.1.22 255.255.255.252
 ipv6 address 2001:DB8:0:2::7/64
 ipv6 ospf 100 area 0
end
```

Step 14 With the root cause determined, fix the problem and verify that the route to 2001:DB8:0:7::/64 now exists in the IPv6 routing table of R1.

Enter the following configuration on the R7 router:

```
R7# conf t
R7(config)# interface Loopback0
R7(config-if)# ipv6 ospf 100 area 0
```

The route to 2001:DB8:0:7::/64 now exists in the IPv6 routing table of R1.


```
R1# show ipv6 route ospf
IPv6 Routing Table - default - 6 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
O   2001:DB8:0:7::1/128 [110/10]
     via FE80::A8BB:CCFF:FE00:4701, Ethernet2/2
```

Note: The link-local IPv6 address in your output may be different.

This is the end of the discovery lab.

Challenge

1. Two OSPF neighbors are stuck in EXCHANGE/EXSTART state. What could be possible reason for it ?

   A. Access-list blocking OSPF hellos on the interface.
   B. MTU issue
   C. OSPF interface is made passive on one router.
   D. Multicast is broken on the link connecting two routers.

2. Which of the following in not a parameter that must be matched for routers to become OSPF neighbors ?

   A. Hello and Dead interval
   B. Area ID
   C. Stub Area flag
   D. OSPF cost
   E. Subnet ID and Subnet mask

3. Which OSPF neighbor state indicates that two neighbors have exchanged routes?

   A. INIT
   B. EXCHANGE
   C. LOADING
   D. FULL

4. You need to check if there is mismatch in hello and dead intervals on two connected routers. Which command will you use ?

   A. **show ipv6 ospf neighbor**
   B. **show ipv6 protocols**
   C. **show ipv6 interface brief**
   D. **show ipv6 ospf interface**

5. Below is the **show ip ospf neighbor** output from R1.

   R1# show ip ospf neighbor

Neighbor ID	Pri	State	Dead Time	Address	Interface
170.170.3.4	1	2WAY/DROTHER	00:00:34	170.170.3.4	Ethernet0
170.170.3.3	1	2WAY/DROTHER	00:00:34	170.170.3.3	Ethernet0
170.170.3.8	1	FULL/DR	00:00:32	170.170.3.8	Ethernet0
170.170.3.2	1	FULL/BDR	00:00:39	170.170.3.2	Ethernet0

   Notice that R1 establishes full adjacency only with the Designated Router (DR) and the Backup Designated Router (BDR). All other routers have a two-way adjacency established. These routers are in broadcast network. Is this normal expected OSPF behavior ?

   A. Yes
   B. No

6. Which command is used to verify the OSPF cost on an interface?

   A. **show ip protocols**
   B. **show ip route**
   C. **show ip ospf**
   D. **show ip ospf interface**

7. Susan is troubleshooting an OSPF neighbor issue between two routers. OSPF has been enabled on the interfaces and the configuration looks good. She enters the following command on one of the routers.

   Router# ping 224.0.0.5

   Type escape sequence to abort.
   Sending 1, 100-byte ICMP Echos to 224.0.0.5, timeout is 2 seconds:
   .
   Router#

   From the output, what could be the reason for OSPF not forming neighbor.

   A. The unicast is broken on the link.
   B. Multicast is broken.
   C. OSPF hello and dead timers do not match
   D. None of the above

Answer Key

Challenge

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

Glossary

AAA

authentication, authorization, and accounting. Pronounced "triple a."

ABR

Area Border Router. Router located on the border of one or more OSPF areas that connects those areas to the backbone network. ABRs are considered members of both the OSPF backbone and the attached areas. They therefore, maintain routing tables describing both the backbone topology and the topology of the other areas.

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).

AD

administrative distance. Rating of the trustworthiness of a routing information source. Administrative distance often is expressed as a numerical value between 0 and 255. The higher the value, the lower the trustworthiness rating.

AD

advertised distance.

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.

ASBR

Autonomous System Boundary Router. An ABR located between an OSPF autonomous system and a non-OSPF network. ASBRs run both OSPF and another routing protocol, such as RIP. ASBRs must reside in a nonstub OSPF area.

BDR

backup designated router.

BGP

Border Gateway Protocol. Interdomain routing protocol that replaces EGP. BGP exchanges reachability information with other BGP systems. It is defined by RFC 1163.

CLI

Command Language Interpreter. The basic Cisco IOS configuration and management interface.

connected interface

The network prefix associated with an interface on the router is considered 100% reliable and is represented with an administrative distance of zero.

CRC

cyclic redundancy check. Error-checking technique in which the frame recipient calculates a remainder by dividing frame contents by a prime binary divisor and compares the calculated remainder to a value stored in the frame by the sending node.

DAI

Dynamic ARP Inspection.

DBD

database description.

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.

DR

designated router.

DUAL

Diffusing Update Algorithm. Convergence algorithm used in EIGRP that provides loop-free operation at every instant throughout a route computation. Allows routers involved in a topology change to synchronize at the same time, while not involving routers that are unaffected by the change.

EBGP

Exterior Border Gateway Protocol

EGP

exterior gateway protocol.

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.

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

external EIGRP

EIGRP marks routes redistributed into an AS with an administrative distance of 170. External routes are less authoritative than internal EIGRP routes.

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.

FD

feasible distance.

FLP

Fast Link Pulse. A type of link pulse that encodes information used in autonegotiation.

FTP

File Transfer Protocol. Protocol for exchanging files over the Internet.

Gigabit Ethernet

Standard for a high-speed Ethernet, approved by the IEEE (Institute of Electrical and Electronics Engineers) 802.3z standards committee in 1996.

Hello protocol

Protocol used by OSPF systems for establishing and maintaining neighbor relationships.

HTTP

Hypertext Transfer Protocol *(common term)*. The protocol that is used by web browsers and web servers to transfer files, such as text and graphic files.

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.

IBGP

Internal Border Gateway Protocol.

ICANN

Internet Corporation for Assigned Names and Numbers. Nonprofit, private corporation that assumed responsibility for IP address-space allocation, protocol parameter assignment, domain name system management, and root server system management functions that formerly were performed under a U.S. government contract by IANA and other entities.

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.

ICMPv6

Internet Control Message Protocol version 6 is the implementation of ICMP for IPv6. It is defined in RFC 4443.

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 802.1X

An IEEE standard for port-based network access control.

IEEE 802.3ab

IEEE 802.3ab is a IEEE standard, that defines physical layer and data link layer's media access control of wired 1000BASE-T (GigabitEthernet transmission over UTP cat 5, 5e, or 6 cabling).

IETF

Internet Engineering Task Force. Task force consisting of over 80 working groups responsible for developing Internet standards. The IETF operates under the auspices of ISOC.

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.

IGRP

Interior Gateway Routing Protocol. Developed by Cisco to address the issues associated with routing in large, heterogeneous networks.

INIT state

initial state.

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.

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).

IS-IS

Intermediate System-to-Intermediate System. OSI link-state hierarchical routing protocol based on DECnet Phase V routing, whereby ISs (routers) exchange routing information based on a single metric to determine network topology.

ISP

Internet service provider. Company that provides Internet access to other companies and individuals.

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.

LED

light emitting diode. A semiconductor device that emits light produced by converting electrical energy. Status lights on hardware devices are typically LEDs.

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.

LSAck

link-state acknowledgment.

LSDB

link-state database.

LSR

link-state request.

LSR

label switch router.

LSU

link-state update.

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.

MD5

Message Digest 5. A one-way hashing algorithm that produces a 128-bit hash. Both MD5 and Secure Hash Algorithm (SHA) are variations on MD4 and are designed to strengthen the security of the MD4 hashing algorithm. Cisco uses hashes for authentication within the IPsec framework. Also used for message authentication in SNMP v.2. MD5 verifies the integrity of the communication, authenticates the origin, and checks for timeliness.

MOTD banner

message-of-the-day banner.

MTU

maximum transmission unit. The maximum packet size, in bytes, that a particular interface can handle.

NIC

network interface card. A board that provides network communication capabilities to and from a computer system. A NIC is also called an adapter.

NTP

Network Time Protocol. A protocol that is built on top of TCP that ensures accurate local timekeeping with reference to radio and atomic clocks that are located on the Internet. This protocol is capable of synchronizing distributed clocks within milliseconds over long time periods.

NVRAM

nonvolatile RAM. RAM that retains its contents when a unit is powered off.

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.

OUI

Organizational Unique Identifier. Three octets that are assigned by the IEEE in a block of 48-bit LAN addresses.

PuTTY

An SSH and Telnet client for Windows and Unix platforms.

QoS

quality of service. Measure of performance for a transmission system that reflects its transmission quality and service availability.

RADIUS

Remote Authentication Dial-In User Service. Database for authenticating modem and ISDN connections and for tracking connection time.

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.

RIP

Routing Information Protocol. A distance-vector routing protocol that uses hop count as a routing metric.

RIPv1

Routing Information Protocol version 1.

RIPv2

Routing Information Protocol version 2.

RIR

regional Internet registry.

RSA

Acronym stands for Rivest, Shamir, and Adleman, the inventors of the technique. Public-key cryptographic system that can be used for encryption and authentication.

SHA-256

Secure Hash Algorithm. SHA-256 is part of the SHA-2 set of cryptographic hash functions.

SLA

service level agreement.

SMTP

Simple Mail Transfer Protocol. Internet protocol providing email services.

SNMP

Simple Network Management Protocol. Network management protocol used almost exclusively in TCP/IP networks. SNMP provides a means to monitor and control network devices, and to manage configurations, statistics collection, performance, and security.

SPAN

Switched Port Analyzer. SPAN is a feature that is available on switches based on Cisco IOS and NX-OS Software that allows traffic received on a port or VLAN to be copied to another port for analysis. It is also referred to as "port mirroring."

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.

SRTT

smoothed round-trip time.

SSH

Secure Shell Protocol. Protocol that provides a secure remote connection to a route through a TCP application.

SSHv1

Secure Shell Protocol version 1.

SSHv2

Secure Shell Protocol version 2.

static route

Route that is explicitly configured and entered into the routing table. Static routes take precedence over routes chosen by dynamic routing protocols.

syslog

system logging.

TACACS

Terminal Access Controller Access Control System. Authentication protocol, developed by the DDN community, that provides remote access authentication and related services, such as event logging. User passwords are administered in a central database rather than in individual routers, providing an easily scalable network security solution.

TACACS+

Terminal Access Controller Access Control System Plus. Proprietary Cisco enhancement to TACACS. Provides additional support for authentication, authorization, and accounting.

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.

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.

TFTP

Trivial File Transfer Protocol. Simplified version of FTP that allows files to be transferred from one computer to another over a network, usually without the use of client authentication (for example, username and password).

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.

unreachable

If a route is received with an administrative distance of 255, it is considered unreachable.

UPS

uninterruptible power supply.

UTC

Coordinated Universal Time. Time zone at zero degrees longitude. Formerly called Greenwich Mean Time (GMT) and Zulu time.

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.

VoIP

Voice over IP. The capability to carry normal telephony-style voice over an IP-based internet with POTS-like functionality, reliability, and voice quality. VoIP enables a router to carry voice traffic (for example, telephone calls and faxes) over an IP network. In VoIP, the DSP segments the voice signal into frames, which then are coupled in groups of two and stored in voice packets. These voice packets are transported using IP in compliance with ITU-T specification H.323. A primary attraction of VoIP is its ability to reduce expenses, because phone calls travel over the data network rather than over the phone company network.

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.

vty

virtual type terminal. Commonly used as virtual terminal lines.

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.

