

Routing Concepts

Objectives

Upon completion of this chapter, you will be able to answer the following questions:

- What are the primary functions and features of a router?
- How do you connect devices for a small routed network?
- Can you configure basic settings on a router to route between two directly connected networks?
- How can you verify connectivity between two networks that are directly connected to a router?
- How do routers encapsulate and de-encapsulate packets when switching packets between directly connected interfaces?
- How do routers determine the best path?
- How do routers build a routing table of directly connected networks?
- How do routers build a routing table using static routes?
- How do routers build a routing table using a dynamic routing protocol?

Key Terms

This chapter uses the following key terms. You can find the definitions in the Glossary.

default gateway page 3

physical topology page 4

logical topology page 4

availability page 5

scalability page 5

reliability page 5

Random Access Memory (RAM) page 6

Read-Only Memory (ROM) page 6

Non-Volatile Random Access Memory (NVRAM) page 6

Flash page 6

process switching page 9

fast switching page 10

Cisco Express Forwarding (CEF) page 11

IP address page 14

subnet mask page 14

topology diagram page 16

addressing table page 16

statically assigned IP address page 16

dynamically assigned IP address page 16

console cable page 19

terminal emulation software page 19

switched virtual interface (SVI) page 20

High-Speed WAN Interface Card (HWIC)
page 24

loopback interface page 28

directly connected network page 43

remote network page 43

Gateway of Last Resort page 43

metric page 44

best path page 44

equal cost load balancing page 45

unequal cost load balancing page 45

administrative distance page 46

routing table page 47

Introduction (1.0.1.1)

Networks allow people to communicate, collaborate, and interact in many ways. Networks are used to access web pages, talk using IP telephones, participate in video conferences, compete in interactive gaming, shop using the Internet, complete online coursework, and more.

At the core of the network is the router. A router connects one network to another network. The router is responsible for the delivery of packets across different networks. The destination of the IP packet might be a web server in another country or an email server on the local-area network.

The router uses its routing table to determine the best path to use to forward a packet. It is the responsibility of the routers to deliver those packets in a timely manner. The effectiveness of internetwork communications depends, to a large degree, on the ability of routers to forward packets in the most efficient way possible.

When a host sends a packet to a device on a different IP network, the packet is forwarded to the default gateway because a host device cannot communicate directly with devices outside of the local network. The *default gateway* is the destination that routes traffic from the local network to devices on remote networks. It is often used to connect a local network to the Internet.

This chapter will also answer the question, “What does a router do with a packet received from one network and destined for another network?” Details of the routing table will be examined, including connected, static, and dynamic routes.

Because the router can route packets between networks, devices on different networks can communicate. This chapter will introduce the router, its role in the networks, its main hardware and software components, and the routing process.



Class Activity 1.0.1.2: Do We Really Need a Map?

This modeling activity asks you to research travel directions from source to destination. Its purpose is to compare those types of directions to network routing directions.

Scenario

Using the Internet and Google Maps, located at <http://maps.google.com>, find a route between the capital city of your country and some other distant town or between two places within your own city. Pay close attention to the driving or walking directions Google Maps suggests.

Notice that in many cases, Google Maps suggests more than one route between the two locations you chose. It also allows you to put additional constraints on the route, such as avoiding highways or tolls.

Copy at least two route instructions supplied by Google Maps for this activity. Place your copies into a word processing document and save it for use with the next step.

Open the .pdf accompanying this modeling activity and complete it with a fellow student. Discuss the reflection questions listed on the .pdf and record your answers.

Be prepared to present your answers to the class.

Initial Configuration of a Router (1.1)

A router is essentially a special-purpose computer with an internetwork operating system optimized for the purpose of routing and securing networks. This section will examine the functions of a router and how a router determines the best path. It will also review the command-line interface (CLI) commands required to configure the base settings of a router.

Characteristics of a Network (1.1.1.1)

Networks have had a significant impact on our lives. They have changed the way we live, work, and play.

Networks allow us to communicate, collaborate, and interact in ways we never did before. We use the network in a variety of ways, including web applications, IP telephony, video conferencing, interactive gaming, electronic commerce, education, and more.

There are many terms, key structures, and performance-related characteristics that are referred to when discussing networks. These include:

- **Topology:** There are physical and logical topologies. The *physical topology* is the arrangement of the cables, network devices, and end systems. It describes how the network devices are actually interconnected with wires and cables. The *logical topology* is the path over which the data is transferred in a network. It describes how the network devices appear connected to network users.
- **Speed:** Speed is a measure of the data rate in bits per second (b/s) of a given link in the network.
- **Cost:** Cost indicates the general expense for purchasing of network components, and installation and maintenance of the network.
- **Security:** Security indicates how protected the network is, including the information that is transmitted over the network. The subject of security is important,

and techniques and practices are constantly evolving. Consider security whenever actions are taken that affect the network.

- **Availability:** Availability is a measure of the probability that the network is available for use when it is required.
- **Scalability:** Scalability indicates how easily the network can accommodate more users and data transmission requirements. If a network design is optimized to only meet current requirements, it can be very difficult and expensive to meet new needs when the network grows.
- **Reliability:** Reliability indicates the dependability of the components that make up the network, such as the routers, switches, PCs, and servers. Reliability is often measured as a probability of failure or as the mean time between failures (MTBF).

These characteristics and attributes provide a means to compare different networking solutions.

Note

While the term “speed” is commonly used when referring to the network bandwidth, it is not technically accurate. The actual speed that the bits are transmitted does not vary over the same medium. The difference in bandwidth is due to the number of bits transmitted per second, not how fast they travel over wire or wireless medium.

Why Routing? (1.1.1.2)

How does clicking a link in a web browser return the desired information in mere seconds? Although there are many devices and technologies collaboratively working together to enable this, the primary device is the router. Stated simply, a router connects one network to another network.

Communication between networks would not be possible without a router determining the best path to the destination and forwarding traffic to the next router along that path. The router is responsible for the routing of traffic between networks.

Video

Video 1.1.1.2: Routers Route Packets

Go to the online course and play the animation of a packet being sent through a Cisco 1841 router from sender to receiver.

When a packet arrives on a router interface, the router uses its routing table to determine how to reach the destination network. The destination of the IP packet might be a web server in another country or an email server on the local-area network. It is

the responsibility of routers to deliver those packets efficiently. The effectiveness of internetwork communications depends, to a large degree, on the ability of routers to forward packets in the most efficient way possible.

Routers Are Computers (1.1.1.3)

Most network capable devices (i.e., computers, tablets, and smartphones) require the following components to operate:

- Central processing unit (CPU)
- Operating system (OS)
- Memory and storage (RAM, ROM, NVRAM, Flash, hard drive)

A router is essentially a specialized computer. It requires a CPU and memory to temporarily and permanently store data to execute operating system instructions, such as system initialization, routing functions, and switching functions.

Note

Cisco devices use the Cisco Internetwork Operating System (IOS) as the system software.

Routers store data using:

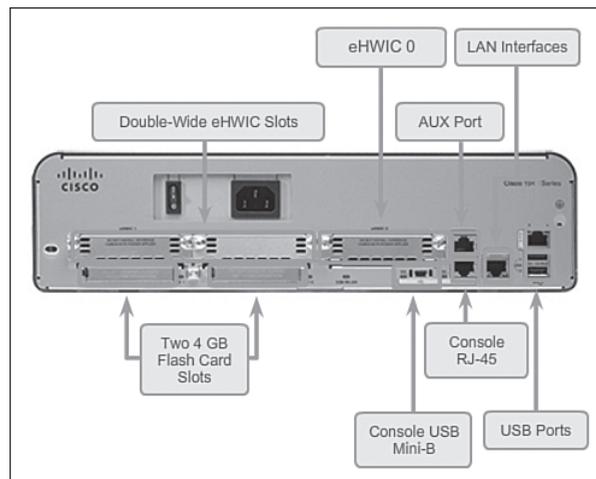
- **Random Access Memory (RAM):** Provides temporary storage for various applications and processes, including the running IOS, the running configuration file, various tables (i.e., IP routing table, Ethernet ARP table), and buffers for packet processing. RAM is referred to as volatile because it loses its contents when power is turned off.
- **Read-Only Memory (ROM):** Provides permanent storage for bootup instructions, basic diagnostic software, and a limited IOS in case the router cannot load the full featured IOS. ROM is firmware and referred to as non-volatile because it does not lose its contents when power is turned off.
- **Non-Volatile Random Access Memory (NVRAM):** Provides permanent storage for the startup configuration file (startup-config). NVRAM is non-volatile and does not lose its contents when power is turned off.
- **Flash:** Provides permanent storage for the IOS and other system-related files. The IOS is copied from flash into RAM during the bootup process. Flash is non-volatile and does not lose its contents when power is turned off.

Table 1-1 provides a summary of the types of router memory, their volatility, and examples of what is stored in each.

Table 1-1 Router Memory

Memory	Volatile/Non-Volatile	Stores
RAM	Volatile	<ul style="list-style-type: none"> ■ Running IOS ■ Running configuration file ■ IP routing and ARP tables ■ Packet buffer
ROM	Non-volatile	<ul style="list-style-type: none"> ■ Bootup instructions ■ Basic diagnostic software ■ Limited IOS
NVRAM	Non-volatile	<ul style="list-style-type: none"> ■ Startup configuration file
Flash	Non-volatile	<ul style="list-style-type: none"> ■ IOS file ■ Other system files

Unlike a computer, a router does not have video adapters or sound card adapters. Instead, routers have specialized ports and network interface cards to interconnect devices to other networks. Figure 1-1 displays the back panel of a Cisco 1941 ISR G2 and identifies those special ports and interfaces.

**Figure 1-1** Back Panel of a 1941 ISR G2

Routers Interconnect Networks (1.1.1.4)

Most users are unaware of the presence of numerous routers on their own network or on the Internet. Users expect to be able to access web pages, send emails, and download music, regardless of whether the server accessed is on their own network or on

another network. Networking professionals know that it is the router that is responsible for forwarding packets from network to network, from the original source to the final destination.

A router connects multiple networks, which means that it has multiple interfaces that each belong to a different IP network. When a router receives an IP packet on one interface, it determines which interface to use to forward the packet to the destination. The interface that the router uses to forward the packet may be the final destination, or it may be a network connected to another router that is used to reach the destination network.

Video**Video 1.1.1.4: Routers Connect**

Go to the online course and play the animation of a packet being sent through two Cisco routers. R1 and R2 are responsible for receiving the packet on one network and forwarding the packet out another network toward the destination network.

Each network that a router connects to typically requires a separate interface. These interfaces are used to connect a combination of both local-area networks (LANs) and wide-area networks (WANs). LANs are commonly Ethernet networks that contain devices, such as PCs, printers, and servers. WANs are used to connect networks over a large geographical area. For example, a WAN connection is commonly used to connect a LAN to the Internet service provider (ISP) network.

Notice that each site in Figure 1-2 requires the use of a router to interconnect to other sites. Even the Home Office requires a router. In this topology, the router located at the Home Office is a specialized device that performs multiple services for the home network.

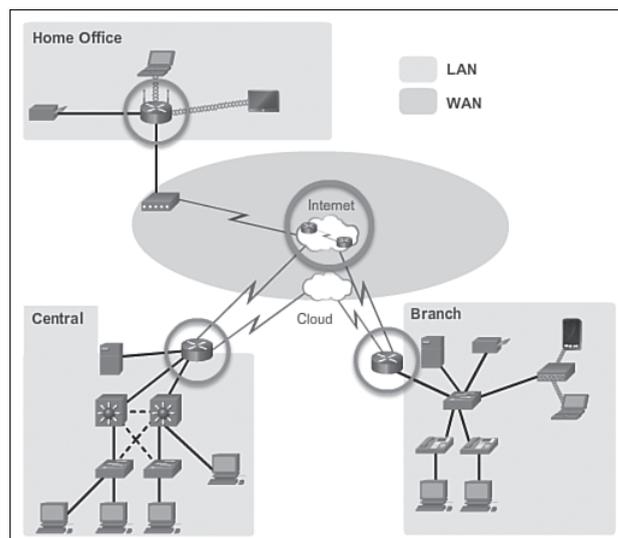


Figure 1-2 Sample Routed Topology

Routers Choose Best Paths (1.1.1.5)

The primary functions of a router are to:

- Determine the best path to send packets
- Forward packets toward their destination

The router uses its routing table to determine the best path to use to forward a packet. When the router receives a packet, it examines the destination address of the packet and uses the routing table to search for the best path to that network. The routing table also includes the interface to be used to forward packets for each known network. When a match is found, the router encapsulates the packet into the data link frame of the outgoing or exit interface, and the packet is forwarded toward its destination.

It is possible for a router to receive a packet that is encapsulated in one type of data link frame, and to forward the packet out of an interface that uses a different type of data link frame. For example, a router may receive a packet on an Ethernet interface, but must forward the packet out of an interface configured with the Point-to-Point Protocol (PPP). The data link encapsulation depends on the type of interface on the router and the type of medium to which it connects. The different data link technologies that a router can connect to include Ethernet, PPP, Frame Relay, DSL, cable, and wireless (802.11, Bluetooth).

Video

Video 1.1.1.5: How the Router Works

Go to the online course and play the animation of a packet being sent through two routers from sender to receiver.

Note

Routers use static routes and dynamic routing protocols to learn about remote networks and build their routing tables.

Packet Forwarding Mechanisms (1.1.1.6)

Routers support three packet-forwarding mechanisms:

- **Process switching:** An older packet-forwarding mechanism still available for Cisco routers. When a packet arrives on an interface, it is forwarded to the control plane, where the CPU matches the destination address with an entry in its routing table, and then determines the exit interface and forwards the packet. It is important to understand that the router does this for every packet, even if the destination is the same for a stream of packets. This process-switching mechanism is very slow and

rarely implemented in modern networks. Figure 1-3 illustrates how packets are process-switched.

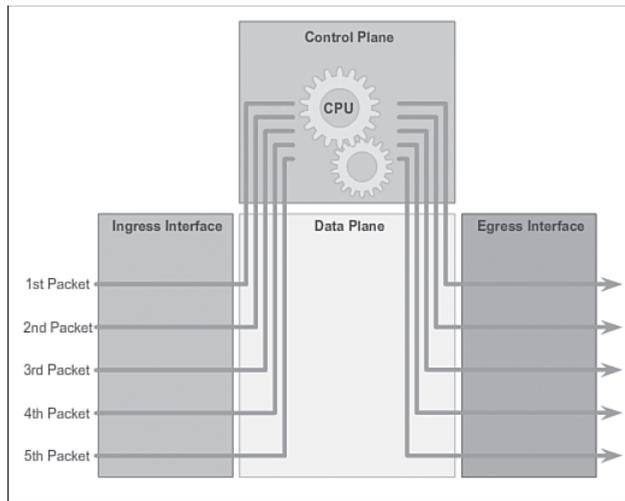


Figure 1-3 Process Switching

- **Fast switching:** This is a common packet-forwarding mechanism which uses a fast-switching cache to store next-hop information. When a packet arrives on an interface, it is forwarded to the control plane, where the CPU searches for a match in the fast-switching cache. If it is not there, it is process-switched and forwarded to the exit interface. The flow information for the packet is also stored in the fast-switching cache. If another packet going to the same destination arrives on an interface, the next-hop information in the cache is re-used without CPU intervention. Figure 1-4 illustrates how packets are fast-switched.

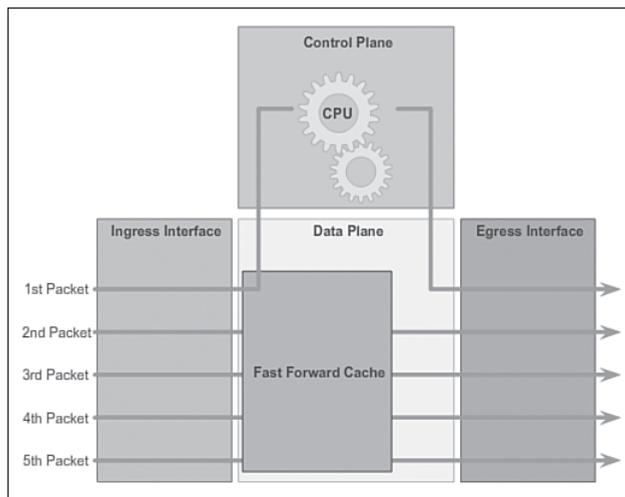


Figure 1-4 Fast Switching

- Cisco Express Forwarding (CEF):** CEF is the most recent and preferred Cisco IOS packet-forwarding mechanism. Like fast switching, CEF builds a Forwarding Information Base (FIB) and an adjacency table. However, the table entries are not packet-triggered like fast switching but change-triggered such as when something changes in the network topology. Therefore, when a network has converged, the FIB and adjacency tables contain all the information a router would have to consider when forwarding a packet. The FIB contains pre-computed reverse lookups and next-hop information for routes, including the interface and Layer 2 information. Cisco Express Forwarding is the fastest forwarding mechanism and the preferred choice on Cisco routers. Figure 1-5 illustrates how packets are forwarded using CEF.

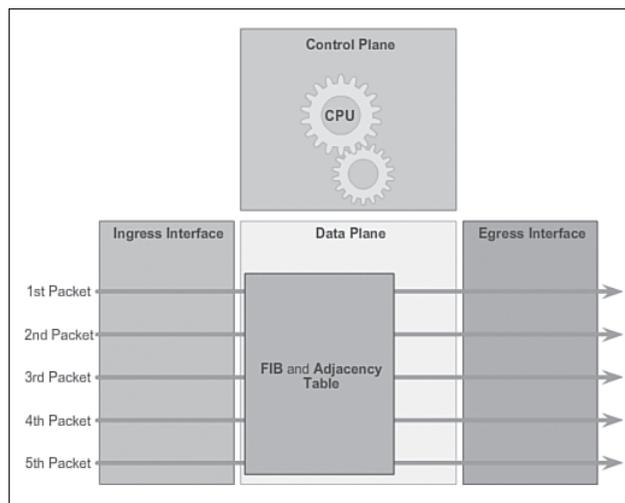


Figure 1-5 Cisco Express Forwarding

Figures 1-3 to 1-5 illustrate the differences between the three packet-forwarding mechanisms. Assume a traffic flow consisting of five packets all going to the same destination. As shown in Figure 1-3, with process switching, each packet must be processed by the CPU individually. Contrast this with fast switching, as shown in Figure 1-4. With fast switching, notice how only the first packet of a flow is process-switched and added to the fast-switching cache. The next four packets are quickly processed based on the information in the fast-switching cache. Finally, in Figure 1-5, CEF builds the FIB and adjacency tables, after the network has converged. All five packets are quickly processed in the data plane.

A common analogy used to describe the three packet-forwarding mechanisms is as follows:

- Process switching solves a problem by doing math long hand, even if it is the identical problem.
- Fast switching solves a problem by doing math long hand one time and remembering the answer for subsequent identical problems.
- CEF solves every possible problem ahead of time in a spreadsheet.

**Interactive
Graphic****Activity 1.1.1.7: Identify Router Components**

Go to the online course to perform this practice activity.

**Packet Tracer
Activity****Packet Tracer Activity 1.1.1.8: Using Traceroute to Discover the Network**

The company you work for has acquired a new branch location. You asked for a topology map of the new location, but apparently one does not exist. However, you have username and password information for the new branch's networking devices and you know the web address for the new branch's server. Therefore, you will verify connectivity and use the `tracert` command to determine the path to the location. You will connect to the edge router of the new location to determine the devices and networks attached. As a part of this process, you will use various `show` commands to gather the necessary information to finish documenting the IP addressing scheme and create a diagram of the topology.

**Lab 1.1.1.9: Mapping the Internet**

In this lab, you will complete the following objectives:

- Part 1: Determine Network Connectivity to a Destination Host
 - Part 2: Trace a Route to a Remote Server Using Tracert
-

Connect Devices (1.1.2)

In this section, you will see how accessing a network involves connecting hosts and infrastructure devices with IP addresses, subnet masks, and default gateways. This section will also introduce how to configure the initial settings of a switch.

Connect to a Network (1.1.2.1)

Network devices and end users typically connect to a network using a wired Ethernet or wireless connection. Refer to the sample reference topology in Figure 1-6. The LANs in the figure serve as an example of how users and network devices could connect to networks.

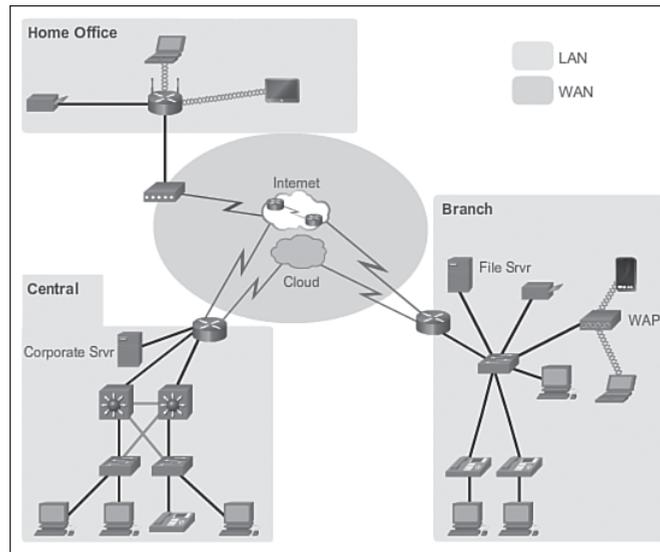


Figure 1-6 Sample LAN and WAN Connections

Home Office devices can connect as follows:

- Laptops and tablets connect wirelessly to a home router.
- A network printer connects using an Ethernet cable to the switch port on the home router.
- The home router connects to the service provider cable modem using an Ethernet cable.
- The cable modem connects to the Internet service provider (ISP) network.

The Branch site devices connect as follows:

- Corporate resources (i.e., file servers and printers) connect to Layer 2 switches using Ethernet cables.
- Desktop PCs and voice over IP (VoIP) phones connect to Layer 2 switches using Ethernet cables.
- Laptops and smartphones connect wirelessly to wireless access points (WAPs).
- The WAPs connect to switches using Ethernet cables.

- Layer 2 switches connect to an Ethernet interface on the edge router using Ethernet cables. An edge router is a device that sits at the edge or boundary of a network and routes between that network and another, such as between a LAN and a WAN.
- The edge router connects to a WAN service provider (SP).
- The edge router also connects to an ISP for backup purposes.

The Central site devices connect as follows:

- Desktop PCs and VoIP phones connect to Layer 2 switches using Ethernet cables.
- Layer 2 switches connect redundantly to multilayer Layer 3 switches using Ethernet fiber-optic cables (orange connections).
- Layer 3 multilayer switches connect to an Ethernet interface on the edge router using Ethernet cables.
- The corporate website server is connected using an Ethernet cable to the edge router interface.
- The edge router connects to a WAN SP.
- The edge router also connects to an ISP for backup purposes.

In the Branch and Central LANs, hosts are connected either directly or indirectly (via WAPs) to the network infrastructure using a Layer 2 switch.

Default Gateways (1.1.2.2)

To enable network access, devices must be configured with IP address information to identify the appropriate:

- **IP address:** Identifies a unique host on a local network
- **Subnet mask:** Identifies with which network subnet the host can communicate
- **Default gateway:** Identifies the router to send a packet to when the destination is not on the same local network subnet

When a host sends a packet to a device that is on the same IP network, the packet is simply forwarded out of the host interface to the destination device.

When a host sends a packet to a device on a different IP network, then the packet is forwarded to the default gateway, because a host device cannot communicate directly with devices outside of the local network. The default gateway is the destination that routes traffic from the local network to devices on remote networks. It is often used to connect a local network to the Internet.

The default gateway is usually the address of the interface on the router connected to the local network. The router maintains routing table entries of all connected networks as well as entries of remote networks, and determines the best path to reach those destinations.

For example, if PC1 sends a packet to the Web Server located at 172.16.1.99, it would discover that the Web Server is not on the local network and it, therefore, must send the packet to the Media Access Control (MAC) address of its default gateway. The packet protocol data unit (PDU) in Figure 1-7 identifies the source and destination IP and MAC addresses.

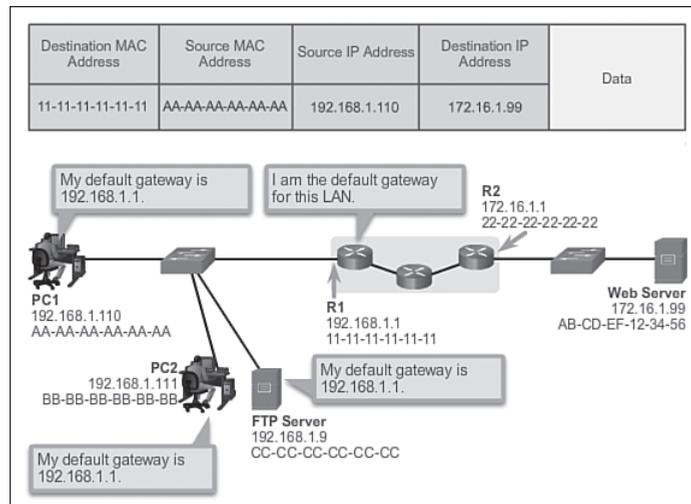


Figure 1-7 Getting the Pieces to the Correct Network

Note

A router is also usually configured with its own default gateway. This is sometimes known as the Gateway of Last Resort.

Document Network Addressing (1.1.2.3)

When designing a new network or mapping an existing network, document the network. At a minimum, the documentation should identify:

- Device names
- Interfaces used in the design
- IP addresses and subnet masks
- Default gateway addresses

This information is captured by creating two useful network documents:

- **Topology diagram:** Provides a visual reference that indicates the physical connectivity and logical Layer 3 addressing. Often created using software, such as Microsoft Visio.
- **Addressing table:** A table that captures device names, interfaces, IPv4 addresses, subnet masks, and default gateway addresses.

Figure 1-8 displays the sample topology diagram, while Table 1-2 provides a sample addressing table for the topology.

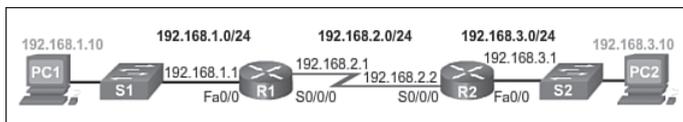


Figure 1-8 Documenting Network Addressing

Table 1-2 Addressing Table

Device	Interface	IP Address	Subnet Mask	Default Gateway
R1	Fa0/0	192.168.1.1	255.255.255.0	N/A
	S0/0/0	192.168.2.1	255.255.255.0	N/A
R2	Fa0/0	192.168.3.1	255.255.255.0	N/A
	S0/0/0	192.168.2.2	255.255.255.0	N/A
PC1	N/A	192.168.1.10	255.255.255.0	192.168.1.1
PC2	N/A	192.168.3.10	255.255.255.0	192.168.3.1

Enable IP on a Host (1.1.2.4)

A host can be assigned its IP address information in one of two ways. A host can get a:

- **Statically Assigned IP Address:** The host is manually assigned the correct IP address, subnet mask, and default gateway. The DNS server IP address can also be configured.
- **Dynamically Assigned IP Address:** IP address information is provided by a server using the Dynamic Host Configuration Protocol (DHCP). The DHCP server provides a valid IP address, subnet mask, and default gateway for end devices. Other information may be provided by the server.

Figures 1-9 and 1-10 provide static and dynamic IPv4 address configuration examples.

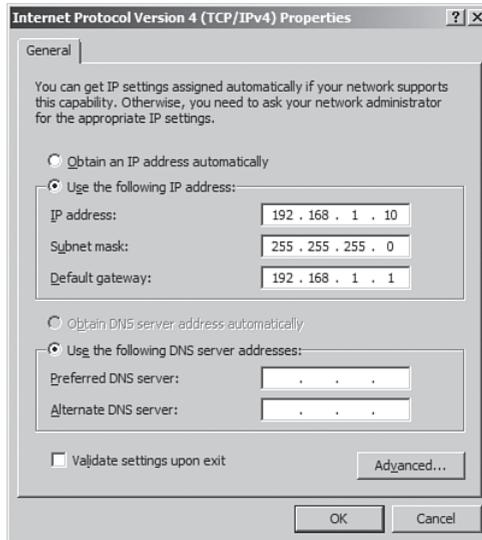


Figure 1-9 Statically Assigning an IP Address

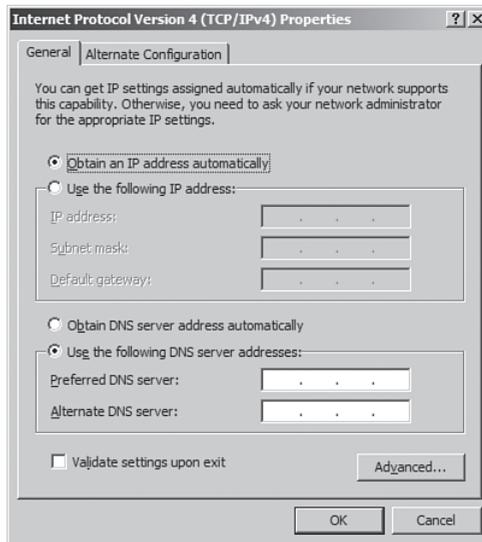


Figure 1-10 Dynamically Assigning an IP Address

Statically assigned addresses are commonly used to identify specific network resources, such as network servers and printers. They can also be used in smaller networks with few hosts. However, most host devices acquire their IPv4 address information by accessing a DHCP server. In large enterprises, dedicated DHCP servers providing services to many LANs are implemented. In a smaller branch or small office setting, DHCP services can be provided by a Cisco Catalyst switch or a Cisco ISR.

Device LEDs (1.1.2.5)

Host computers connect to a wired network using a network interface and RJ-45 Ethernet cable. Most network interfaces have one or two LED link indicators next to the interface. Typically, a green LED means a good connection while a blinking green LED indicates network activity.

If the link light is not on, then there may be a problem with either the network cable or the network itself. The switch port where the connection terminates would also have an LED indicator lit. If one or both ends are not lit, try a different network cable.

Note

The actual function of the LEDs varies between computer manufacturers.

Similarly, network infrastructure devices commonly use multiple LED indicators to provide a quick status view. For example, a Cisco Catalyst 2960 switch has several status LEDs to help monitor system activity and performance. These LEDs are generally lit green when the switch is functioning normally and lit amber when there is a malfunction.

Cisco ISRs use various LED indicators to provide status information. The LEDs on the router help the network administrator conduct some basic troubleshooting. Each device has a unique set of LEDs. Consult the device-specific documentation for an accurate description of the LEDs.

The LEDs of the Cisco 1941 router shown in Figure 1-11 are explained in Table 1-3.

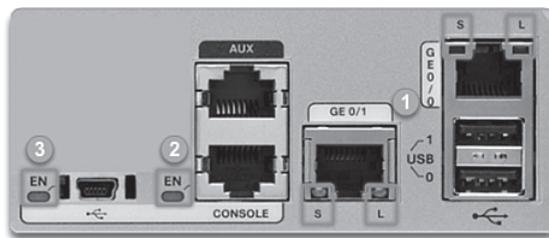


Figure 1-11 Cisco 1941 LEDs

Table 1-3 Description of the Cisco 1941 LEDs

#	Port	LED	Color	Description
1	GE0/0 and GE0/1	S (Speed)	1 blink + pause	Port operating at 10 Mb/s
			2 blink + pause	Port operating at 100 Mb/s
			3 blink + pause	Port operating at 1000 Mb/s
		L (Link)	Green	Link is active
			Off	Link is inactive
2	Console	EN	Green	Port is active
			Off	Port is inactive
3	USB	EN	Green	Port is active
			Off	Port is inactive

Console Access (1.1.2.6)

In a production environment, infrastructure devices are commonly accessed remotely using Secure Shell (SSH) or HyperText Transfer Protocol Secure (HTTPS). Console access is really only required when initially configuring a device, or if remote access fails.

Console access requires:

- *Console cable*: RJ-45-to-DB-9 console cable
- *Terminal emulation software*: Tera Term, PuTTY, HyperTerminal

The cable is connected between the serial port of the host and the console port on the device. Most computers and notebooks no longer include built-in serial ports. If the host does not have a serial port, the USB port can be used to establish a console connection. A special USB-to-RS-232 compatible serial port adapter is required when using the USB port.

The Cisco ISR G2 supports a USB serial console connection. To establish connectivity, a USB Type-A to USB Type-B (mini-B USB) is required, as well as an operating system device driver. This device driver is available from <http://www.cisco.com>. Although these routers have two console ports, only one console port can be active at a time. When a cable is plugged into the USB console port, the RJ-45 port becomes inactive. When the USB cable is removed from the USB port, the RJ-45 port becomes active.

Table 1-4 summarizes the console connection requirements, while Figure 1-12 displays the various ports and cables required.

Table 1-4 Console Connection Requirements

Port on Computer	Cable Required	Port on ISR	Terminal Emulation
Serial Port	RJ-45 to DB-9 Console Cable		
USB Type-A Port	USB to RS-232 compatible serial port adapter <ul style="list-style-type: none"> Adapter may require a software driver 	RJ-45 Console Port	Tera Term PuTTY
	RJ-45 to DB-9 Console Cable		
	USB Type-A to USB Type-B (Mini-B USB) <ul style="list-style-type: none"> A device driver is required and available from Cisco.com 	USB Type-B (Mini-B USB)	

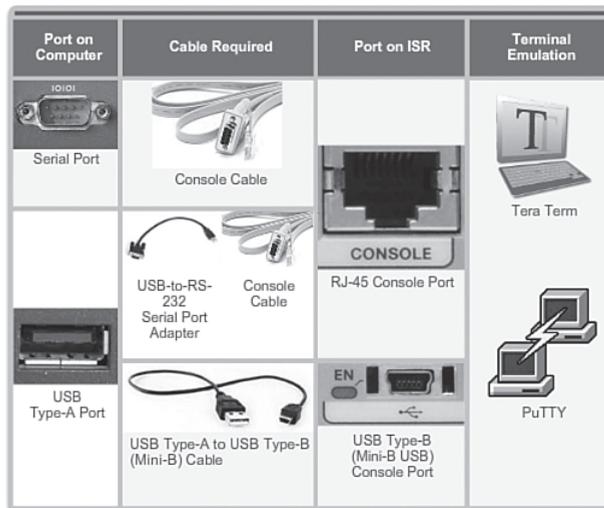


Figure 1-12 Ports and Cables

Enable IP on a Switch (1.1.2.7)

Network infrastructure devices require IP addresses to enable remote management. Using the device IP address, the network administrator can remotely connect to the device using Telnet, SSH, HTTP, or HTTPS.

A switch does not have a dedicated interface to which an IP address can be assigned. Instead, the IP address information is configured on a virtual interface called a *switched virtual interface (SVI)*.



The steps to configure the basic settings on a switch are as follows:

- Step 1.** Name the device.
- Step 2.** Configure the SVI. This makes the switch accessible for network management.
- Step 3.** Enable the SVI.
- Step 4.** Configure the default gateway for the switch. Packets generated by the switch and destined for an address other than its management network segment will be forwarded to this address. This default gateway is used by the switch only for the packets it generates, not any hosts connected to the switch.

For example, the following commands would configure the management VLAN interface and default gateway of switch S1 shown in Figure 1-13.

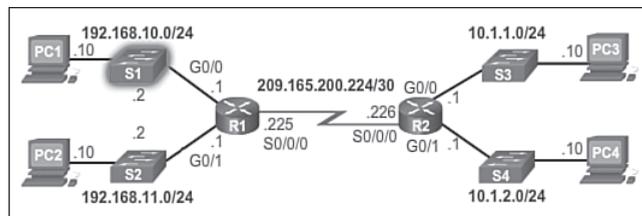


Figure 1-13 Configuring the SVI of S1

```
S1(config)# interface vlan 1
S1(config-if)# ip address 192.168.10.2 255.255.255.0
S1(config-if)# no shutdown
%LINK-5-CHANGED: Interface Vlan1, changed state to up
S1(config-if)# exit
S1(config)#
S1(config)# ip default-gateway 192.168.10.1
S1(config)#
```

In the example, the switch SVI is configured and enabled with the IP address 192.168.10.2/24 and a default gateway of the router located at 192.168.10.1. Packets generated by the switch and destined for an address outside of the 192.168.10/24 network segment will be forwarded to this address. In the example, the address is that of the G0/0 interface of R1.

Interactive
Graphic

Activity 1.1.2.7: Configure the Management SVI on S2

Go to the online course to use the Syntax Checker in the second graphic to configure the S2 Layer 2 switch.

**Interactive
Graphic****Activity 1.1.2.8: Document an Addressing Scheme**

Go to the online course to perform this practice activity.

**Packet Tracer
Activity****Packet Tracer Activity 1.1.2.9: Documenting the Network**

Your job is to document the addressing scheme and connections used in the Central portion of the network. You will need to use a variety of commands to gather the required information.

Basic Settings on a Router (1.1.3)

The basic addressing and configuration of Cisco devices was covered in either the Introduction to Networks or Network Basics course. However, we will spend some time reviewing these topics as well as preparing you for the hands-on lab experience in this course.

Configure Basic Router Settings (1.1.3.1)

Cisco routers and Cisco switches have many similarities. They support a similar modal operating system, similar command structures, and many of the same commands. In addition, both devices have similar initial configuration steps.

When initially configuring a Cisco switch or router, the following steps should be executed:



- Step 1.** Name the device. This changes the router prompt and helps distinguish the device from others.
- Step 2.** Secure management access. Specifically, secure the privileged EXEC, user EXEC, and Telnet access, and encrypt passwords to their highest level.
- Step 3.** Configure a banner. Although optional, this is a recommended step to provide legal notice to anyone attempting to access the device.
- Step 4.** Save the configuration.

For example, the following commands would configure the basic settings for router R1 shown in Figure 1-14.

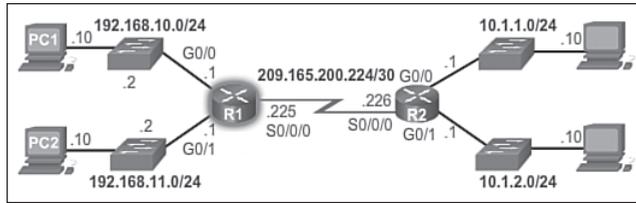


Figure 1-14 Configuring the Basic Settings of R1

```

Router# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Router(config)# hostname R1
R1(config)#
R1(config)# enable secret class
R1(config)#
R1(config)# line console 0
R1(config-line)# password cisco
R1(config-line)# login
R1(config-line)# exit
R1(config)#
R1(config)# line vty 0 4
R1(config-line)# password cisco
R1(config-line)# login
R1(config-line)# exit
R1(config)#
R1(config)# service password-encryption
R1(config)#
R1(config)# banner motd $ Authorized Access Only! $
R1(config)# end
R1#
R1# copy running-config startup-config
Destination filename [startup-config]?
Building configuration...
[OK]
R1#

```

Interactive
Graphic

Activity 1.1.3.1: Configure Basic Settings on R2

Go to the online course to use the Syntax Checker in the fifth graphic to configure basic settings on R2.

Configure an IPv4 Router Interface (1.1.3.2)

One distinguishing feature between switches and routers is the type of interfaces supported by each. For example, Layer 2 switches support LANs and, therefore, have multiple FastEthernet or Gigabit Ethernet ports.

Routers support LANs and WANs and can interconnect different types of networks; therefore, they support many types of interfaces. For example, G2 ISRs have one or two integrated Gigabit Ethernet interfaces and *High-Speed WAN Interface Card (HWIC)* slots to accommodate other types of network interfaces, including serial, DSL, and cable interfaces.

To be available, an interface must be:

- **If using IPv4, configured with an address and a subnet mask:** Use the **ip address** *ip-address subnet-mask* interface configuration command.
- **Activated:** By default, LAN and WAN interfaces are not activated (**shutdown**). To enable an interface, it must be activated using the **no shutdown** command. (This is similar to powering on the interface.) The interface must also be connected to another device (a hub, a switch, or another router) for the physical layer to be active.

Optionally, the interface could also be configured with a short description. It is good practice to configure a description on each interface. The description text is limited to 240 characters. On production networks, a description can be helpful in troubleshooting by providing information about the type of network to which the interface is connected. If the interface connects to an ISP or service carrier, it is helpful to enter the third-party connection and contact information.

Depending on the type of interface, additional parameters may be required. For example, in the lab environment, the serial interface connecting to the serial cable end labeled DCE must be configured with the **clock rate** command.

Note

Accidentally using the **clock rate** command on a DTE interface generates a “%Error: This command applies only to DCE interface” message.

How To

The steps to configure an IPv4 interface on a router are:

- Step 1.** Add a description. Although optional, it is a necessary component for documenting a network.
- Step 2.** Configure the IPv4 address.
- Step 3.** Configure a clock rate on Serial interfaces. This is only necessary on the DCE device in our lab environment and does not apply to Ethernet interfaces.
- Step 4.** Enable the interface.

For example, the following commands would configure the three directly connected interfaces of router R1 shown in Figure 1-14 (in the previous section):

```
R1(config)# interface gigabitethernet 0/0
R1(config-if)# description Link to LAN 1
R1(config-if)# ip address 192.168.10.1 255.255.255.0
R1(config-if)# no shutdown
R1(config-if)# exit
R1(config)#
R1(config)# interface gigabitethernet 0/1
R1(config-if)# description Link to LAN 2
R1(config-if)# ip address 192.168.11.1 255.255.255.0
R1(config-if)# no shutdown
R1(config-if)# exit
R1(config)#
R1(config)# interface serial 0/0/0
R1(config-if)# description Link to R2
R1(config-if)# ip address 209.165.200.225 255.255.255.252
R1(config-if)# clock rate 128000
R1(config-if)# no shutdown
R1(config-if)# exit
R1(config)#
```

Interactive Graphic

Activity 1.1.3.2: Configure the R2 Interfaces

Go to the online course to use the Syntax Checker in the fourth graphic to configure the R2 interfaces.

Configure an IPv6 Router Interface (1.1.3.3)

Configuring an IPv6 interface is similar to configuring an interface for IPv4. Most IPv6 configuration and verification commands in the Cisco IOS are very similar to their IPv4 counterparts. In many cases, the only difference uses `ipv6` in place of `ip` in commands.

An IPv6 interface must be:

- **Configured with IPv6 address and subnet mask:** Use the `ipv6 address` *ipv6-address/prefix-length* [`link-local` | `eui-64`] interface configuration command.
- **Activated:** The interface must be activated using the `no shutdown` command.

Note

An interface can generate its own IPv6 link-local address without having a global unicast address by using the `ipv6 enable` interface configuration command.

Unlike IPv4, IPv6 interfaces will typically have more than one IPv6 address. At a minimum, an IPv6 device must have an IPv6 link-local address but will most likely also have an IPv6 global unicast address. IPv6 also supports the ability for an interface to have multiple IPv6 global unicast addresses from the same subnet. The following commands can be used to statically create a global unicast or link-local IPv6 address:

- **ipv6 address *ipv6-address/prefix-length***: Creates a global unicast IPv6 address as specified.
- **ipv6 address *ipv6-address/prefix-length eui-64***: Configures a global unicast IPv6 address with an interface identifier (ID) in the low-order 64 bits of the IPv6 address using the EUI-64 process.
- **ipv6 address *ipv6-address/prefix-length link-local***: Configures a static link-local address on the interface that is used instead of the link-local address that is automatically configured when the global unicast IPv6 address is assigned to the interface or enabled using the **ipv6 enable** interface command. Recall, the **ipv6 enable** interface command is used to automatically create an IPv6 link-local address whether or not an IPv6 global unicast address has been assigned.



The steps to configure an IPv6 interface on a router are:

- Step 1.** Add a description. Although optional, it is a necessary component for documenting a network.
- Step 2.** Configure the IPv6 global unicast address. Configuring a global unicast address automatically creates a link-local IPv6 address.
- Step 3.** Configure a link-local unicast address which automatically assigns a link-local IPv6 address and overrides any previously assigned address.
- Step 4.** Configure a clock rate on Serial interfaces. This is only necessary on the DCE device in our lab environment and does not apply to Ethernet interfaces.
- Step 5.** Enable the interface.

In the example topology shown in Figure 1-15, R1 must be configured to support the following IPv6 global network addresses:

- 2001:0DB8:ACAD:0001:/64 (2001:DB8:ACAD:1::/64)
- 2001:0DB8:ACAD:0002:/64 (2001:DB8:ACAD:2::/64)
- 2001:0DB8:ACAD:0003:/64 (2001:DB8:ACAD:3::/64)

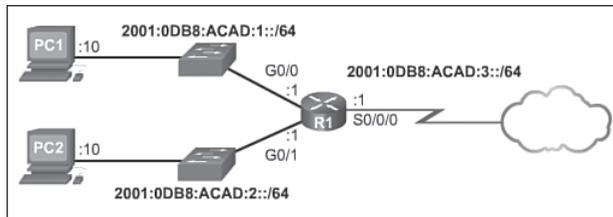


Figure 1-15 IPv6 Topology

When the router is configured using the **ipv6 unicast-routing** global configuration command, the router begins sending ICMPv6 Router Advertisement messages out the interface. This enables a PC connected to the interface to automatically configure an IPv6 address and to set a default gateway without needing the services of a DHCPv6 server. Alternatively, a PC connected to the IPv6 network can get its IPv6 address statically assigned, as shown in Figure 1-16. Notice that the default gateway address configured for PC1 is the IPv6 global unicast address of the R1 Gigabit Ethernet 0/0 interface.

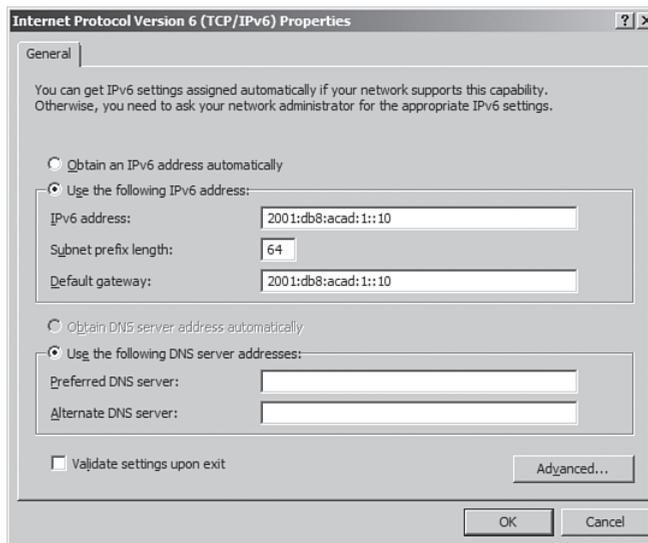


Figure 1-16 Statically Assign an IPv6 Address to PC1

For example, the following commands would configure the IPv6 global unicast addresses of the three directly connected interfaces of the R1 router shown in Figure 1-15:

```
R1# configure terminal
R1(config)# interface gigabitethernet 0/0
R1(config-if)# description Link to LAN 1
R1(config-if)# ipv6 address 2001:db8:acad:1::1/64
R1(config-if)# no shutdown
R1(config-if)# exit
R1(config)#
```

```
R1(config)# interface gigabitethernet 0/1
R1(config-if)# description Link to LAN 2
R1(config-if)# ipv6 address 2001:db8:acad:2::1/64
R1(config-if)# no shutdown
R1(config-if)# exit
R1(config)#
R1(config)# interface serial 0/0/0
R1(config-if)# description Link to R2
R1(config-if)# ipv6 address 2001:db8:acad:3::1/64
R1(config-if)# clock rate 128000
R1(config-if)# no shutdown
R1(config-if)#
```

**Interactive
Graphic****Activity 1.1.3.3: Configure the R2 Interfaces**

Go to the online course to use the Syntax Checker in the sixth graphic to configure the IPv6 global unicast addresses on the R2 router.

Configure an IPv4 Loopback Interface (1.1.3.4)

Another common configuration of Cisco IOS routers is enabling a loopback interface.

The *loopback interface* is a logical interface internal to the router. It is not assigned to a physical port and can therefore never be connected to any other device. It is considered a software interface that is automatically placed in an “up/up” state, as long as the router is functioning.

The loopback interface is useful in testing and managing a Cisco IOS device because it ensures that at least one interface will always be available. For example, it can be used for testing purposes, such as testing internal routing processes, by emulating networks behind the router.

Additionally, the IPv4 address assigned to the loopback interface can be significant to processes on the router that use an interface IPv4 address for identification purposes, such as the Open Shortest Path First (OSPF) routing process. By enabling a loopback interface, the router will use the always available loopback interface address for identification, rather than an IP address assigned to a physical port that may go down.

How To 

The steps to configure a loopback interface on a router are:

- Step 1.** Create the loopback interface using the `interface loopback number` global configuration command.
- Step 2.** Add a description. Although optional, it is a necessary component for documenting a network.
- Step 3.** Configure the IP address.

For example, the following commands configure a loopback interface of the R1 router shown in Figure 1-14 (shown earlier in the chapter):

```
R1# configure terminal
R1(config)# interface loopback 0
R1(config-if)# ip address 10.0.0.1 255.255.255.0
R1(config-if)# exit
R1(config)#
```

A loopback interface is always enabled and therefore does not require a **no shutdown** command. Multiple loopback interfaces can be enabled on a router. The IPv4 address for each loopback interface must be unique and unused by any other interface.

Packet Tracer
Activity

Packet Tracer Activity 1.1.3.5: Configuring IPv4 and IPv6 Interfaces

Routers R1 and R2 each have two LANs. Your task is to configure the appropriate addressing on each device and verify connectivity between the LANs.

Verify Connectivity of Directly Connected Networks (1.1.4)

The first task to undertake once the basic settings and interfaces are configured is to verify and validate the configured settings. This is an important step and should be done before any other configurations are added to the router.

Verify Interface Settings (1.1.4.1)

There are several **show** commands that can be used to verify the operation and configuration of an interface. The following three commands are especially useful to quickly identify an interface status:

- **show ip interface brief:** Displays a summary for all interfaces, including the IPv4 address of the interface and current operational status.
- **show ip route:** Displays the contents of the IPv4 routing table stored in RAM. In Cisco IOS 15, active interfaces should appear in the routing table with two related entries identified by the code 'C' (Connected) or 'L' (Local). In previous IOS versions, only a single entry with the code 'C' will appear.
- **show running-config interface *interface-id*:** Displays the commands configured on the specified interface.

Figure 1-17 displays the output of the **show ip interface brief** command.

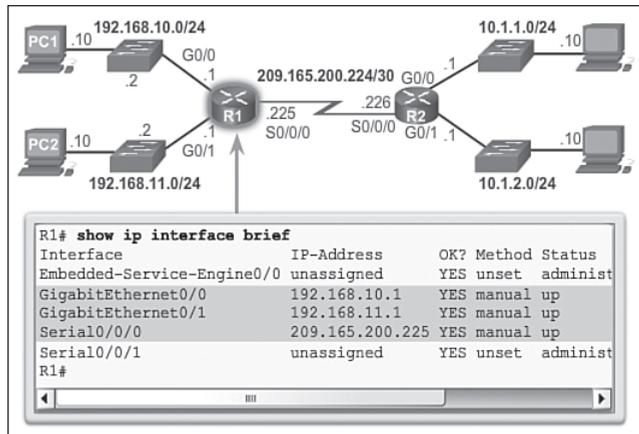


Figure 1-17 Display Interface Summaries

The output reveals that the LAN interfaces and the WAN link are all activated and operational as indicated by the Status of “up” and Protocol of “up.” A different output would indicate a problem with either the configuration or the cabling.

Note

The entire output of the **show ip interface brief** command in Figure 1-17 can be viewed in the online course on page 1.1.4.1 graphic number 1.

Note

In Figure 1-17, the Embedded-Service-Engine0/0 interface is displayed because Cisco ISRs G2 have dual-core CPUs on the motherboard. The Embedded-Service-Engine0/0 interface is outside the scope of this course.

Figure 1-18 displays the output of the **show ip route** command.

```

R1# show ip route
Codes: L - local, C - connected, S - static, R - RIP, M - m
<output omitted>
Gateway of last resort is not set

  192.168.10.0/24 is variably subnetted, 2 subnets, 2 ma
C    192.168.10.0/24 is directly connected, GigabitEther
L    192.168.10.1/32 is directly connected, GigabitEther
C    192.168.11.0/24 is variably subnetted, 2 subnets, 2 ma
L    192.168.11.0/24 is directly connected, GigabitEther
L    192.168.11.1/32 is directly connected, GigabitEther
  209.165.200.0/24 is variably subnetted, 2 subnets, 2 m

```

Figure 1-18 Verify the IPv4 Routing Table

Note

The entire output of the **show ip route** command in Figure 1-18 can be viewed in the online course on page 1.1.4.1 graphic number 2.

Notice the three directly connected network entries and the three local host route interface entries. A local host route has an administrative distance of 0. It also has a /32 mask for IPv4, and a /128 mask for IPv6. The local host route is for routes on the router owning the IP address. It is used to allow the router to process packets destined to that IP.

Figure 1-19 displays the output of the **show running-config interface** command. The output displays the current commands configured on the specified interface.

```

R1# show running-config interface gigabitEthernet 0/0
Building configuration...

Current configuration : 128 bytes
!
interface GigabitEthernet0/0
  description Link to LAN 1
  ip address 192.168.10.1 255.255.255.0
  duplex auto
  speed auto
end

R1#

```

Figure 1-19 Verify an Interface Configuration

The following two commands are used to gather more detailed interface information:

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

**Interactive
Graphic**

Activity 1.1.4.1: Verify Router Interfaces

Go to the online course to use the Syntax Checker in the fourth and fifth graphics to verify the interfaces of the R2 router.

Verify IPv6 Interface Settings (1.1.4.2)

The commands to verify the IPv6 interface configuration are similar to the commands used for IPv4.

The `show ipv6 interface brief` command in Figure 1-20 displays a summary for each of the interfaces.

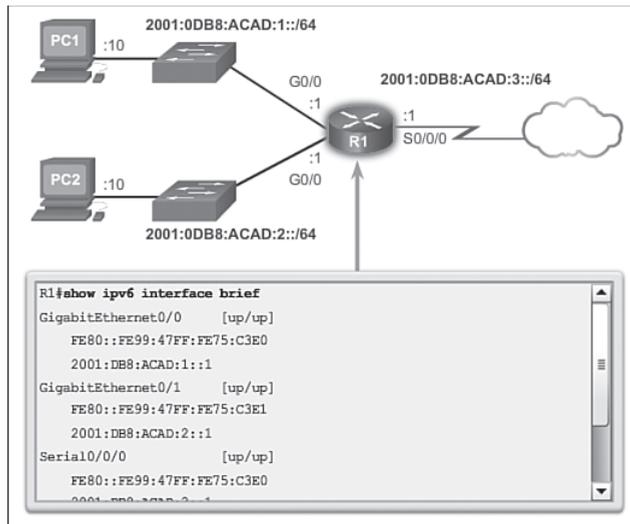


Figure 1-20 Verify the R1 IPv6 Interface Status

Note

The entire output of the `show ipv6 interface brief` command in Figure 1-20 can be viewed in the online course on page 1.1.4.2 graphic number 1.

The “up/up” output on the same line as the interface name indicates the Layer 1/ Layer 2 interface state. This is the same as the Status and Protocol columns in the equivalent IPv4 command.

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

The `show ipv6 interface gigabitethernet 0/0` command output shown in Figure 1-21 displays the interface status and all of the IPv6 addresses belonging to the interface. Along with the link-local address and global unicast address, the output includes the multicast addresses assigned to the interface, beginning with prefix FF02.

```

R1#show ipv6 interface gigabitEthernet 0/0
GigabitEthernet0/0 is up, line protocol is up
IPv6 is enabled, link-local address is FE80::32F7:DFF:FEA3:DA0
No Virtual link-local address(es):
Global unicast address(es):
  2001:DB8:ACAD:1::1, subnet is 2001:DB8:ACAD:1::/64
Joined group address(es):
  FF02::1
  FF02::1:FF00:1
  FF02::1:FF00:1

```

Figure 1-21 Verify the IPv6 Configuration on R1 G0/0

Note

The entire output of the **show ipv6 interface** command in Figure 1-21 can be viewed in the online course on page 1.1.4.2 graphic number 2.

The **show ipv6 route** command shown in Figure 1-22 can be used to verify that IPv6 networks and specific IPv6 interface addresses have been installed in the IPv6 routing table. The **show ipv6 route** command will only display IPv6 networks, not IPv4 networks.

```

<output omitted>
C 2001:DB8:ACAD:1::/64 [0/0]
  via GigabitEthernet0/0, directly connected
L 2001:DB8:ACAD:1::1/128 [0/0]
  via GigabitEthernet0/0, receive
C 2001:DB8:ACAD:2::/64 [0/0]
  via GigabitEthernet0/1, directly connected
L 2001:DB8:ACAD:2::1/128 [0/0]

```

Figure 1-22 Verify the R1 IPv6 Routing Table

Note

The entire output of the **show ipv6 route** command in Figure 1-22 can be viewed in the online course on page 1.1.4.2 graphic number 3.

Within the routing table, a ‘c’ next to a route indicates that this is a directly connected network. When the router interface is configured with a global unicast address and is in the “up/up” state, the IPv6 prefix and prefix length is added to the IPv6 routing table as a connected route.

The IPv6 global unicast address configured on the interface is also installed in the routing table as a local route, as indicated with an ‘L’ next to the route entry. The local route has a /128 prefix. Local routes are used by the routing table to efficiently process packets with the interface address of the router as the destination.

The **ping** command for IPv6 is identical to the command used with IPv4 except that an IPv6 address is used. As shown in Figure 1-23, the **ping** command is used to verify Layer 3 connectivity between R1 and PC1.

```
R1#ping 2001:db8:acad:1::10
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2001:DB8:ACAD:1::10, timeout is 2
seconds:
!!!!!!
Success rate is 100 percent (5/5)
R1#
```

Figure 1-23 Verify Connectivity on R1

Other useful IPv6 verification commands include:

- **show interface**
- **show ipv6 routers**

Filter Show Command Output (1.1.4.3)

Commands that generate multiple screens of output are, by default, paused after 24 lines. At the end of the paused output, the **--More--** text displays. Pressing Enter displays the next line and pressing the spacebar displays the next set of lines. Use the **terminal length number** command to specify the number of lines to be displayed. A value of 0 (zero) prevents the router from pausing between screens of output.

Another very useful feature that improves the user experience in the command-line interface (CLI) is the filtering of **show** output. Filtering commands can be used to display specific sections of output. To enable the filtering command, enter a pipe (**|**) character after the **show** command and then enter a filtering parameter and a filtering expression.

The filtering parameters that can be configured after the pipe include:

- **section:** Shows entire section that starts with the filtering expression
- **include:** Includes all output lines that match the filtering expression
- **exclude:** Excludes all output lines that match the filtering expression
- **begin:** Shows all the output lines from a certain point, starting with the line that matches the filtering expression

Note

Output filters can be used in combination with any **show** command.

Figures 1-24 through 1-27 provide examples of the various output filters. The example in Figure 1-24 uses the pipe character and the **section** keyword.

```
R1# show running-config | section line vty
line vty 0 4
  password 7 030752180500
  login
  transport input all
R1#
```

Figure 1-24 Filter **show** Commands by Section

The example in Figure 1-25 uses the pipe character and the **include** keyword.

```
R1# show ip interface brief
Interface          IP-Address      OK? Method Status
Embedded-Service-Engine0/0 unassigned      YES unset  administr
GigabitEthernet0/0 192.168.10.1    YES manual  up
GigabitEthernet0/1 192.168.11.1    YES manual  up
Serial0/0/0        209.165.200.225 YES manual  up
Serial0/0/1        unassigned      YES unset  administr
R1#
R1# show ip interface brief | include up
GigabitEthernet0/0 192.168.10.1    YES manual  up
GigabitEthernet0/1 192.168.11.1    YES manual  up
Serial0/0/0        209.165.200.225 YES manual  up
R1#
```

Figure 1-25 Filter **show** Commands by Common Keyword

Note

The entire output of the **show ip interface brief** command in Figure 1-25 can be viewed in the online course on page 1.1.4.3 graphic number 2.

The example in Figure 1-26 uses the pipe character and the **exclude** keyword.

```
R1# show ip interface brief
Interface          IP-Address      OK? Method Status
Embedded-Service-Engine0/0 unassigned      YES unset  administr
GigabitEthernet0/0 192.168.10.1    YES manual  up
GigabitEthernet0/1 192.168.11.1    YES manual  up
Serial0/0/0        209.165.200.225 YES manual  up
Serial0/0/1        unassigned      YES unset  administr
R1#
R1# show ip interface brief | exclude unassigned
Interface          IP-Address      OK? Method Status
GigabitEthernet0/0 192.168.10.1    YES manual  up
GigabitEthernet0/1 192.168.11.1    YES manual  up
Serial0/0/0        209.165.200.225 YES manual  up
R1#
```

Figure 1-26 Filter **show** Commands to Exclude Rows of Output

Note

The entire output of the **show ip interface brief** command in Figure 1-26 can be viewed in the online course on page 1.1.4.3 graphic number 3.

The example in Figure 1-27 uses the pipe character and the **begin** keyword.

```

R1# show ip route | begin Gateway
Gateway of last resort is not set

      192.168.10.0/24 is variably subnetted, 2 subnets, 2 masks
C       192.168.10.0/24 is directly connected, GigabitEthernet0/0
L       192.168.10.1/32 is directly connected, GigabitEthernet0/0
      192.168.11.0/24 is variably subnetted, 2 subnets, 2 masks
C       192.168.11.0/24 is directly connected, GigabitEthernet0/1
L       192.168.11.1/32 is directly connected, GigabitEthernet0/1
      209.165.200.0/24 is variably subnetted, 2 subnets, 2 masks
C       209.165.200.224/30 is directly connected, Serial0/0/0
L       209.165.200.225/32 is directly connected, Serial0/0/0
R1#

```

Figure 1-27 Filter **show** Commands Beginning from a Keyword

**Interactive
Graphic****Activity 1.1.4.3: Filter Command Output**

Go to the online course to use the Syntax Checker in the fifth graphic to practice how to filter command output on the R1 router.

Command History Feature (1.1.4.4)

The command history feature is useful, because it temporarily stores the list of executed commands to be recalled.

To recall commands in the history buffer, press **Ctrl+P** or the **Up Arrow** key. The command output begins with the most recent command. Repeat the key sequence to recall successively older commands. To return to more recent commands in the history buffer, press **Ctrl+N** or the **Down Arrow** key. Repeat the key sequence to recall successively more recent commands.

By default, command history is enabled and the system captures the last 10 command lines in its history buffer. Use the **show history** privileged EXEC command to display the contents of the buffer.

It is also practical to increase the number of command lines that the history buffer records during the current terminal session only. Use the **terminal history size** user EXEC command to increase or decrease the size of the buffer.

For example, the following displays a sample of the **terminal history size** and **show history** commands:

```
R1# terminal history size 200
R1#
R1# show history
  show ip interface brief
  show interface g0/0
  show ip interface g0/1
  show ip route
  show ip route 209.165.200.224
  show running-config interface s0/0/0
  terminal history size 200
  show history
R1#
```

**Interactive
Graphic**

Activity 1.1.4.4: Adjusting the Command History

Go to the online course to use the Syntax Checker in the second graphic to adjust and list the command history output on the R1 router.

**Packet Tracer
Activity**

Packet Tracer Activity 1.1.4.5: Configuring and Verifying a Small Network

In this activity, you will configure a router with basic settings including IP addressing. You will also configure a switch for remote management and configure the PCs. After you have successfully verified connectivity, you will use **show** commands to gather information about the network.



Lab 1.1.4.6: Configuring Basic Router Settings with IOS CLI

In this lab, you will complete the following objectives:

- Part 1: Set Up the Topology and Initialize Devices
 - Part 2: Configure Devices and Verify Connectivity
 - Part 3: Display Router Information
 - Part 4: Configure IPv6 and Verify Connectivity
-



Lab 1.1.4.7: Configuring Basic Router Settings with CCP

In this lab, you will complete the following objectives:

- Part 1: Set Up the Topology and Initialize Devices
 - Part 2: Configure Devices and Verify Connectivity
 - Part 3: Configure Router to Allow CCP Access
 - Part 4: (Optional) Install and Set Up CCP on PC-A
 - Part 5: Configure R1 Settings Using CCP
 - Part 6: Use CCP Utilities
-

Routing Decisions (1.2)

The key to understanding the role of a router in the network is to understand that a router is a Layer 3 device responsible for forwarding packets. However, a router also operates at Layers 1 and 2.

Router Switching Function (1.2.1.1)

A primary function of a router is to forward packets toward their destination. This is accomplished by using a switching function, which is the process used by a router to accept a packet on one interface and forward it out of another interface. A key responsibility of the switching function is to encapsulate packets in the appropriate data link frame type for the outgoing data link.

Note

In this context, the term “switching” literally means moving packets from source to destination and should not be confused with the function of a Layer 2 switch.

After the router has determined the exit interface using the path determination function, the router must encapsulate the packet into the data link frame of the outgoing interface.

What does a router do with a packet received from one network and destined for another network? The router performs the following three major steps:

- Step 1.** De-encapsulates the Layer 3 packet by removing the Layer 2 frame header and trailer.
- Step 2.** Examines the destination IP address of the IP packet to find the best path in the routing table.

Step 3. If the router finds a path to the destination, it encapsulates the Layer 3 packet into a new Layer 2 frame and forwards the frame out the exit interface.

As shown in Figure 1-28, devices have Layer 3 IPv4 addresses and Ethernet interfaces have Layer 2 data link addresses. For example, PC1 is configured with IPv4 address 192.168.1.10 and an example MAC address of 0A-10. As a packet travels from the source device to the final destination device, the Layer 3 IP addresses do not change. However, the Layer 2 data link addresses change at every hop as the packet is de-encapsulated and re-encapsulated in a new frame by each router. It is very likely that the packet is encapsulated in a different type of Layer 2 frame than the one in which it was received. For example, an Ethernet encapsulated frame might be received by the router on a FastEthernet interface, and then processed to be forwarded out of a serial interface as a Point-to-Point Protocol (PPP) encapsulated frame.

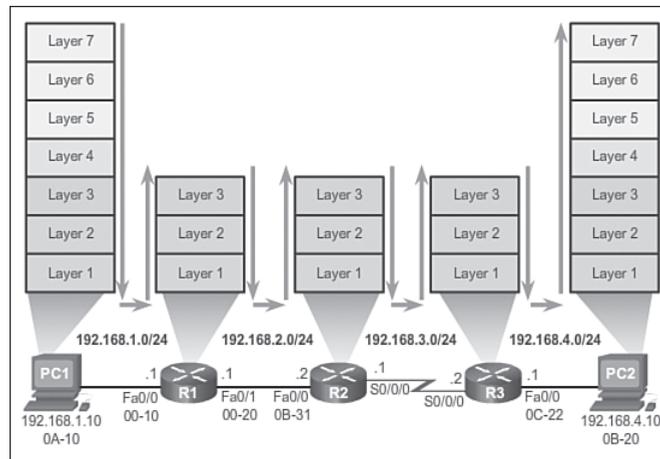


Figure 1-28 Encapsulating and De-Encapsulating Packets

Send a Packet (1.2.1.2)

In the animation in the online course, PC1 is sending a packet to PC2.

Video

Video 1.2.1.2: PC1 Sends a Packet to PC2

Go to the online course and play the animation of a packet being sent from PC1 to PC2.

PC1 must determine if the destination IPv4 address is on the same network. PC1 determines its own subnet by doing an **AND** operation on its own IPv4 address and subnet mask. This produces the network address that PC1 belongs to. Next, PC1

does this same **AND** operation using the packet destination IPv4 address and the PC1 subnet mask.

If the destination network address is the same network as PC1, then PC1 does not use the default gateway. Instead, PC1 refers to its ARP cache for the MAC address of the device with that destination IPv4 address. If the MAC address is not in the cache, then PC1 generates an ARP request to acquire the address to complete the packet and send it to the destination. If the destination network address is on a different network, then PC1 forwards the packet to its default gateway.

To determine the MAC address of the default gateway, PC1 checks its ARP table for the IPv4 address of the default gateway and its associated MAC address.

If an ARP entry does not exist in the ARP table for the default gateway, PC1 sends an ARP request. Router R1 sends back an ARP reply. PC1 can then forward the packet to the MAC address of the default gateway, the Fa0/0 interface of router R1.

A similar process is used for IPv6 packets. Instead of the ARP process, IPv6 address resolution uses ICMPv6 Neighbor Solicitation and Neighbor Advertisement messages. IPv6-to-MAC address mappings are kept in a table similar to the ARP cache, called the neighbor cache.

Forward to the Next Hop (1.2.1.3)

The following processes take place when R1 receives the Ethernet frame from PC1:

- 1.** R1 examines the destination MAC address, which matches the MAC address of the receiving interface, FastEthernet 0/0. R1, therefore, copies the frame into its buffer.
- 2.** R1 identifies the Ethernet Type field as 0x800, which means that the Ethernet frame contains an IPv4 packet in the data portion of the frame.
- 3.** R1 de-encapsulates the Ethernet frame.
- 4.** Because the destination IPv4 address of the packet does not match any of the directly connected networks of R1, R1 consults its routing table to route this packet. R1 searches the routing table for a network address that would include the destination IPv4 address of the packet as a host address within that network. In this example, the routing table has a route for the 192.168.4.0/24 network. The destination IPv4 address of the packet is 192.168.4.10, which is a host IPv4 address on that network.

The route that R1 finds to the 192.168.4.0/24 network has a next-hop IPv4 address of 192.168.2.2 and an exit interface of FastEthernet 0/1. This means that the IPv4 packet is encapsulated in a new Ethernet frame with the destination MAC address of the IPv4 address of the next-hop router.

Because the exit interface is on an Ethernet network, R1 must resolve the next-hop IPv4 address with a destination MAC address using ARP:

1. R1 looks up the next-hop IPv4 address of 192.168.2.2 in its ARP cache. If the entry is not in the ARP cache, R1 would send an ARP request out of its FastEthernet 0/1 interface and R2 would send back an ARP reply. R1 would then update its ARP cache with an entry for 192.168.2.2 and the associated MAC address.
2. The IPv4 packet is now encapsulated into a new Ethernet frame and forwarded out the FastEthernet 0/1 interface of R1.

The animation in the online course illustrates how R1 forwards the packet to R2.

Video**Video 1.2.1.3: R1 Forwards the Packet to R2**

Go to the online course and play the animation of a packet being sent through three routers from sender to receiver.

Packet Routing (1.2.1.4)

The following processes take place when R2 receives the frame on its Fa0/0 interface:

1. R2 examines the destination MAC address, which matches the MAC address of the receiving interface, FastEthernet 0/0. R2, therefore, copies the frame into its buffer.
2. R2 identifies the Ethernet Type field as 0x800, which means that the Ethernet frame contains an IPv4 packet in the data portion of the frame.
3. R2 de-encapsulates the Ethernet frame.
4. Because the destination IPv4 address of the packet does not match any of the interface addresses of R2, R2 consults its routing table to route this packet. R2 searches the routing table for the destination IPv4 address of the packet using the same process R1 used.
5. The routing table of R2 has a route to the 192.168.4.0/24 network, with a next-hop IPv4 address of 192.168.3.2 and an exit interface of Serial 0/0/0. Because the exit interface is not an Ethernet network, R2 does not have to resolve the next-hop IPv4 address with a destination MAC address.
6. The IPv4 packet is now encapsulated into a new data link frame and sent out the Serial 0/0/0 exit interface.

When the interface is a point-to-point (P2P) serial connection, the router encapsulates the IPv4 packet into the proper data link frame format used by the exit interface (HDLC, PPP, etc.). Because there are no MAC addresses on serial interfaces, R2 sets the data link destination address to an equivalent of a broadcast (MAC address: FF:FF:FF:FF:FF:FF).

The animation in the online course illustrates how R2 forwards the packet to R3.

Video**Video 1.2.1.4: R2 Forwards the Packet to R3**

Go to the online course and play the animation of a packet being sent from R2 to R3.

Reach the Destination (1.2.1.5)

The following processes take place when the frame arrives at R3:

1. R3 copies the data link PPP frame into its buffer.
2. R3 de-encapsulates the data link PPP frame.
3. R3 searches the routing table for the destination IPv4 address of the packet. The routing table has a route to a directly connected network on R3. This means that the packet can be sent directly to the destination device and does not need to be sent to another router.

Because the exit interface is a directly connected Ethernet network, R3 must resolve the destination IPv4 address of the packet with a destination MAC address:

1. R3 searches for the destination IPv4 address of the packet in its Address Resolution Protocol (ARP) cache. If the entry is not in the ARP cache, R3 sends an ARP request out of its FastEthernet 0/0 interface. PC2 sends back an ARP reply with its MAC address. R3 then updates its ARP cache with an entry for 192.168.4.10 and the MAC address that is returned in the ARP reply.
2. The IPv4 packet is encapsulated into a new Ethernet data link frame and sent out the FastEthernet 0/0 interface of R3.
3. When PC2 receives the frame, it examines the destination MAC address, which matches the MAC address of the receiving interface, its Ethernet network interface card (NIC). PC2, therefore, copies the rest of the frame into its buffer.
4. PC2 identifies the Ethernet Type field as 0x800, which means that the Ethernet frame contains an IPv4 packet in the data portion of the frame.
5. PC2 de-encapsulates the Ethernet frame and passes the IPv4 packet to the IPv4 process of its operating system.

The animation in the online course illustrates how R3 forwards the packet to PC2.

Video

Video 1.2.1.5: R3 Forwards the Packet to PC2

Go to the online course and play the animation of a packet being sent from R3 to PC2.

Interactive Graphic

Activity 1.2.1.6: Match Layer 2 and Layer 3 Addressing

Go to the online course to perform this practice activity.

Path Determination (1.2.2)

This section discusses the best path to send packets, load balancing, and the concept of administrative distance.

Routing Decisions (1.2.2.1)

A primary function of a router is to determine the best path to use to send packets. To determine the best path, the router searches its routing table for a network address that matches the destination IP address of the packet.

The routing table search results in one of three path determinations:

- **Directly connected network:** If the destination IP address of the packet belongs to a device on a network that is directly connected to one of the interfaces of the router, that packet is forwarded directly to the destination device. This means that the destination IP address of the packet is a host address on the same network as the interface of the router.
- **Remote network:** If the destination IP address of the packet belongs to a remote network, then the packet is forwarded to another router. Remote networks can only be reached by forwarding packets to another router.
- **No route determined:** If the destination IP address of the packet does not belong to either a connected or remote network, the router determines if there is a Gateway of Last Resort available. A *Gateway of Last Resort* is set when a default route is configured on a router. If there is a default route, the packet is forwarded to the Gateway of Last Resort. If the router does not have a default route, then the packet is discarded. If the packet is discarded, the router sends an ICMP Unreachable message to the source IP address of the packet.

The logic flowchart in Figure 1-29 illustrates the router packet-forwarding decision process.

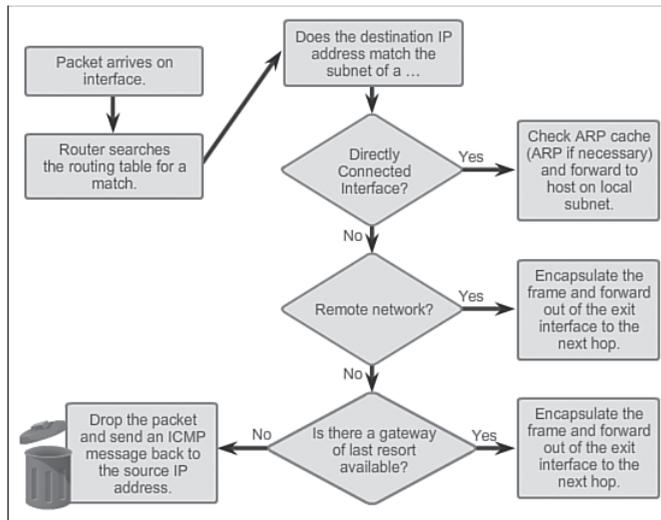


Figure 1-29 Packet Forwarding Decision Process

Best Path (1.2.2.2)

Determining the best path involves the evaluation of multiple paths to the same destination network and selecting the optimum or shortest path to reach that network. Whenever multiple paths to the same network exist, each path uses a different exit interface on the router to reach that network.

The best path is selected by a routing protocol based on the value or metric it uses to determine the distance to reach a network. A *metric* is the quantitative value used to measure the distance to a given network. The *best path* to a network is the path with the lowest metric.

Dynamic routing protocols typically use their own rules and metrics to build and update routing tables. The routing algorithm generates a value, or a metric, for each path through the network. Metrics can be based on either a single characteristic or several characteristics of a path. Some routing protocols can base route selection on multiple metrics, combining them into a single metric.

The following lists some dynamic protocols and the metrics they use:

- **Routing Information Protocol (RIP)** : Hop count
- **Open Shortest Path First (OSPF)**: Cisco routers use a cost based on cumulative bandwidth from source to destination

- **Enhanced Interior Gateway Routing Protocol (EIGRP):** Bandwidth, delay, load, reliability

The animation in the online course highlights how the path may be different depending on the metric being used.

Video**Video 1.2.2.2: Hop Count vs. Bandwidth as a Metric**

Go to the online course and play the animation showing how a network path may be different depending on the metric being used.

Load Balancing (1.2.2.3)

What happens if a routing table has two or more paths with identical metrics to the same destination network?

When a router has two or more paths to a destination with equal cost metrics, then the router forwards the packets using both paths equally. This is called *equal cost load balancing*. The routing table contains the single destination network, but has multiple exit interfaces, one for each equal cost path. The router forwards packets using the multiple exit interfaces listed in the routing table.

If configured correctly, load balancing can increase the effectiveness and performance of the network. Equal cost load balancing can be configured to use both dynamic routing protocols and static routes.

By default, Cisco routers can load balance up to four equal cost paths. The maximum number of equal cost paths depends on the routing protocol and IOS version.

EIGRP supports equal cost load balancing and is also the only routing protocol to support *unequal cost load balancing*. Unequal cost load balancing is when a router distributes traffic over network interfaces, even those that are different distances from the destination address.

Note

EIGRP supports unequal cost load balancing by using the `variance` command.

The animation in the online course provides an example of equal cost load balancing.

Video**Video 1.2.2.3: Equal Cost Load Balancing**

Go to the online course and play the animation showing an example of equal cost load balancing.

Administrative Distance (1.2.2.4)

It is possible for a router to be configured with multiple routing protocols and static routes. If this occurs, the routing table may have more than one route source for the same destination network. For example, if both RIP and EIGRP are configured on a router, both routing protocols may learn of the same destination network. However, each routing protocol may decide on a different path to reach the destination based on that routing protocol's metrics. RIP chooses a path based on hop count, whereas EIGRP chooses a path based on its composite metric. How does the router know which route to use?

Cisco IOS uses what is known as the *administrative distance* (AD) to determine the route to install into the IP routing table. The AD represents the “trustworthiness” of the route; the lower the AD, the more trustworthy the route source. For example, a static route has an AD of 1, whereas an EIGRP-discovered route has an AD of 90. Given two separate routes to the same destination, the router chooses the route with the lowest AD. When a router has the choice of a static route and an EIGRP route, the static route takes precedence. Similarly, a directly connected route with an AD of 0 takes precedence over a static route with an AD of 1.

Table 1-5 lists various routing protocols and their associated ADs.

Table 1-5 Default Administrative Distances

Route Source	Administrative Distance
Connected	0
Static	1
EIGRP summary route	5
External BGP	20
Internal EIGRP	90
IGRP	100
OSPF	110
IS-IS	115
RIP	120
External EIGRP	170
Internal BGP	200
Unknown	255

**Interactive
Graphic****Activity 1.2.2.5: Order the Steps in the Packet Forwarding Process**

Go to the online course to perform these four practice activities.

Router Operation (1.3)

The primary function of a router is to forward packets toward their destination network, the destination IP address of the packet. To do this, a router needs to search the routing information stored in its routing table. In the following sections, you will learn how a router builds the routing table. Then, you will learn the three basic routing principles.

Analyze the Routing Table (1.3.1)

A good understanding of routing tables is crucial for any network administrator.

The Routing Table (1.3.1.1)

The *routing table* of a router stores information about:

- **Directly connected routes:** These routes come from the active router interfaces. Routers add a directly connected route when an interface is configured with an IP address and is activated.
- **Remote routes:** These are remote networks connected to other routers. Routes to these networks can be either statically configured or dynamically configured using dynamic routing protocols.

Specifically, a routing table is a data file in RAM that is used to store route information about directly connected and remote networks. The routing table contains network or next-hop associations. These associations tell a router that a particular destination can be optimally reached by sending the packet to a specific router that represents the next hop on the way to the final destination. The next-hop association can also be the outgoing or exit interface to the next destination.

Figure 1-30 identifies the directly connected networks and remote networks of router R1.

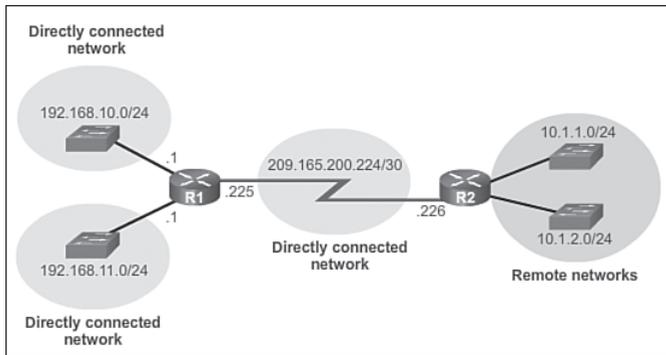


Figure 1-30 Directly Connected and Remote Network Routes

Routing Table Sources (1.3.1.2)

On a Cisco IOS router, the `show ip route` command can be used to display the IPv4 routing table of a router. A router provides additional route information, including how the route was learned, how long the route has been in the table, and which specific interface to use to get to a predefined destination.

Entries in the routing table can be added as:

- **Local Route interfaces:** Added when an interface is configured and active. This entry is only displayed in IOS 15 or newer for IPv4 routes and all IOS releases for IPv6 routes.
- **Directly connected interfaces:** Added to the routing table when an interface is configured and active.
- **Static routes:** Added when a route is manually configured and the exit interface is active.
- **Dynamic routing protocol:** Added when routing protocols that dynamically learn about the network, such as EIGRP or OSPF, are implemented and networks are identified.

The sources of the routing table entries are identified by a code. The code identifies how the route was learned. For instance, common codes include:

- **L:** Identifies the address assigned to a router's interface. This allows the router to efficiently determine when it receives a packet for the interface instead of being forwarded.
- **C:** Identifies a directly connected network.

- **S:** Identifies a static route created to reach a specific network.
- **D:** Identifies a dynamically learned network from another router using EIGRP.
- **O:** Identifies a dynamically learned network from another router using the OSPF routing protocol.

Note

Other codes are beyond the scope of this chapter.

Figure 1-31 shows a sample routing table of R1.

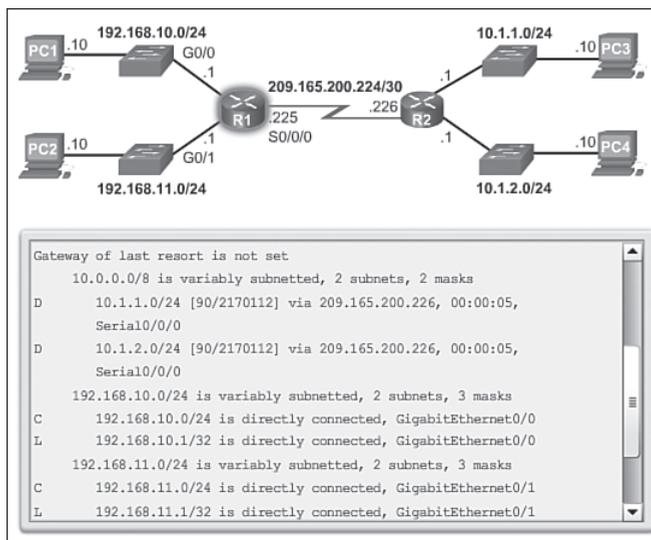


Figure 1-31 Routing Table of R1

Note

The entire output of the `show ip route` command in Figure 1-31 can be viewed in the online course on page 1.3.1.2.

Remote Network Routing Entries (1.3.1.3)

As a network administrator, it is imperative to know how to interpret the content of an IPv4 and IPv6 routing table. Figure 1-32 displays an IPv4 routing table entry on R1 for the route to remote network 10.1.1.0.

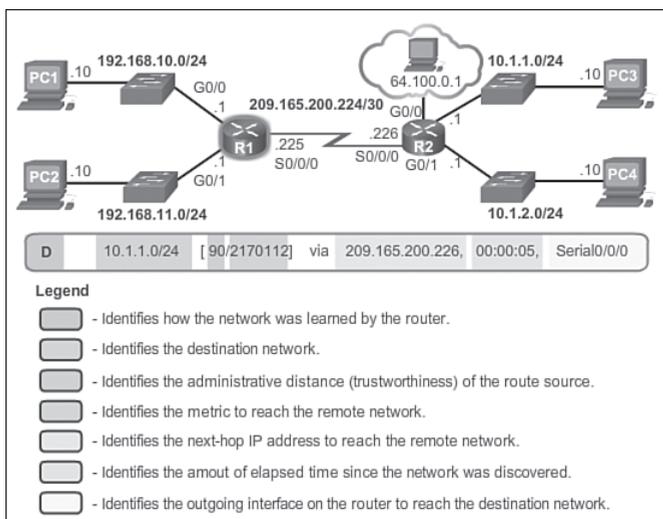


Figure 1-32 Remote Network Entry Identifiers

The entry identifies the following information:

- **Route source:** Identifies how the route was learned.
- **Destination network:** Identifies the address of the remote network.
- **Administrative distance:** Identifies the trustworthiness of the route source. Lower values indicate preferred route source.
- **Metric:** Identifies the value assigned to reach the remote network. Lower values indicate preferred routes.
- **Next-hop:** Identifies the IPv4 address of the next router to forward the packet to.
- **Route timestamp:** Identifies how much time has passed since the route was learned.
- **Outgoing interface:** Identifies the exit interface to use to forward a packet toward the final destination.

Interactive
Graphic

Activity 1.3.1.4: Interpret the Content of a Routing Table Entry

Go to the online course to perform this practice activity.

Directly Connected Routes (1.3.2)

How does a router add its interfaces to the routing table? Well, whenever an interface is configured with an IP address and enabled, it is automatically added as a directly connected network.

Directly Connected Interfaces (1.3.2.1)

A newly deployed router, without any configured interfaces, has an empty routing table, as shown in Figure 1-33.

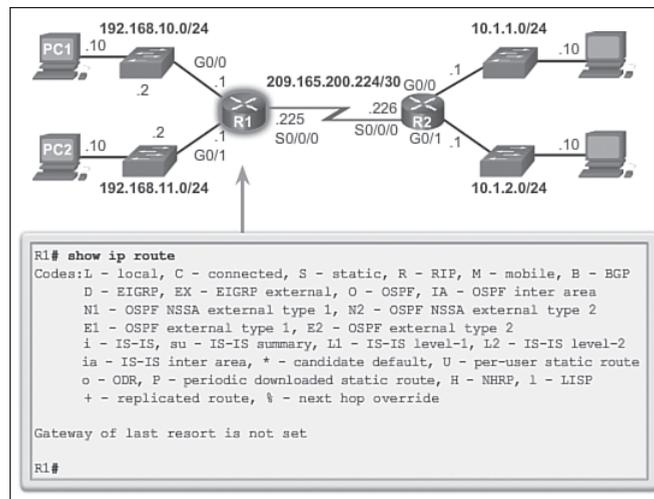


Figure 1-33 Empty Routing Table

Before the interface state is considered “up/up” and added to the IPv4 routing table, the interface must:

- Be assigned a valid IPv4 or IPv6 address
- Be activated with the **no shutdown** command
- Receive a carrier signal from another device (router, switch, host, etc.)

When the interface is up, the network of that interface is added to the routing table as a directly connected network.

Directly Connected Route Table Entries (1.3.2.2)

An active, properly configured, directly connected interface actually creates two routing table entries. Figure 1-34 displays the IPv4 routing table entries on R1 for the directly connected network 192.168.10.0.

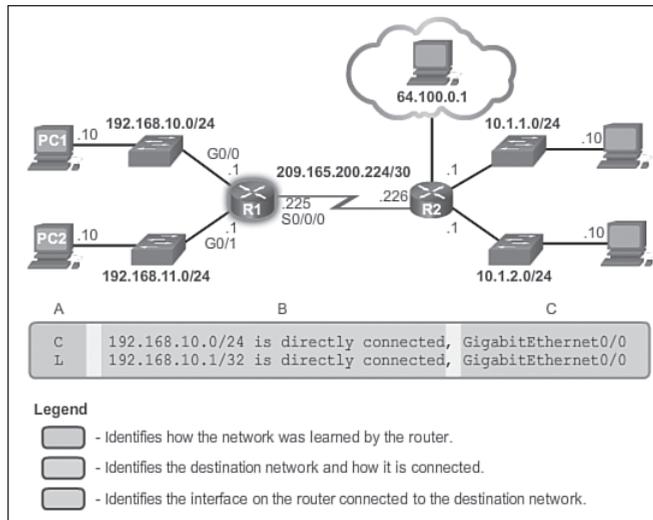


Figure 1-34 Directly Connected Network Entry Identifiers

The routing table entry for directly connected interfaces is simpler than the entries for remote networks. The entries contain the following information:

- **Route source:** Identifies how the route was learned. Directly connected interfaces have two route source codes. 'c' identifies a directly connected network. 'L' identifies the IPv4 address assigned to the router's interface.
- **Destination network:** The address of the remote network.
- **Outgoing interface:** Identifies the exit interface to use when forwarding packets to the destination network.

Note

Prior to IOS 15, local route routing table entries (L) were not displayed in the IPv4 routing table. Local route (L) entries have always been a part of the IPv6 routing table.

Directly Connected Examples (1.3.2.3)

When directly connected interfaces are enabled, Layer 1 and 2 informational messages are automatically generated.

For example, configuring the following commands on R1 would enable the directly connected Gigabit Ethernet 0/0 interface and generate the following messages:

```
R1(config)# interface gigabitethernet 0/0
R1(config-if)# description Link to LAN 1
R1(config-if)# ip address 192.168.10.1 255.255.255.0
```

```

R1(config-if)# no shutdown
R1(config-if)# exit
R1(config)#
*Jan 30 22:04:47.551: %LINK-3-UPDOWN: Interface GigabitEthernet0/0, changed state to
down
R1(config)#
*Jan 30 22:04:50.899: %LINK-3-UPDOWN: Interface GigabitEthernet0/0, changed state to
up
*Jan 30 22:04:51.899: %LINEPROTO-5-UPDOWN: Line protocol on Interface
GigabitEthernet0/0, changed state to up
R1(config)#

```

As each interface is enabled, the routing table automatically adds the connected ('c') and local ('L') entries.

The following provides an example of the routing table with the directly connected interfaces of R1 configured and activated:

```

R1# show ip route | begin Gateway
Gateway of last resort is not set

    192.168.10.0/24 is variably subnetted, 2 subnets, 2 masks
C       192.168.10.0/24 is directly connected, GigabitEthernet0/0
L       192.168.10.1/32 is directly connected, GigabitEthernet0/0
    192.168.11.0/24 is variably subnetted, 2 subnets, 2 masks
C       192.168.11.0/24 is directly connected, GigabitEthernet0/1
L       192.168.11.1/32 is directly connected, GigabitEthernet0/1
    209.165.200.0/24 is variably subnetted, 2 subnets, 2 masks
C       209.165.200.224/30 is directly connected, Serial0/0/0
L       209.165.200.225/32 is directly connected, Serial0/0/0
R1#

```

Interactive Graphic

Activity 1.3.2.3: Configure and Activate the Interfaces on R2

Go to the online course to use the Syntax Checker in the fifth graphic to configure the interfaces of the R2 router.

Directly Connected IPv6 Example (1.3.2.4)

Enabling directly connected IPv6 interfaces also generates Layer 1 and Layer 2 informational messages.

For example, configuring the following commands on R1 would enable the directly connected IPv6 Gigabit Ethernet 0/0 interface and generate the following messages:

```

R1(config)# ipv6 unicast-routing
R1(config)# interface gigabitethernet 0/0
R1(config-if)# description Link to LAN 1

```

```
R1(config-if)# ipv6 address 2001:db8:acad:1::1/64
R1(config-if)# no shutdown
R1(config-if)# exit
*Feb  3 21:38:37.279: %LINK-3-UPDOWN: Interface GigabitEthernet0/0, changed state to
down
*Feb  3 21:38:40.967: %LINK-3-UPDOWN: Interface GigabitEthernet0/0, changed state to
up
*Feb  3 21:38:41.967: %LINEPROTO-5-UPDOWN: Line protocol on Interface
GigabitEthernet0/0, changed state to up
R1(config)#
```

The following provides an example of the routing table with the directly connected interfaces of R1 configured and activated:

```
R1# show ipv6 route
IPv6 Routing Table - default - 5 entries
Codes: C - Connected, L - Local, S - Static, U - Per-user Static route
        B - BGP, R - RIP, H - NHRP, I1 - ISIS L1
        I2 - ISIS L2, IA - ISIS interarea, IS - ISIS summary, D - EIGRP
        EX - EIGRP external, ND - ND Default, NDP - ND Prefix, DCE - Destination
        NDR - Redirect, 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:ACAD:1::/64 [0/0]
    via GigabitEthernet0/0, directly connected
L   2001:DB8:ACAD:1::1/128 [0/0]
    via GigabitEthernet0/0, receive
C   2001:DB8:ACAD:2::/64 [0/0]
    via GigabitEthernet0/1, directly connected
L   2001:DB8:ACAD:2::1/128 [0/0]
    via GigabitEthernet0/1, receive
L   FF00::/8 [0/0]
    via Null0, receive
R1#
```

When the **show ipv6 route** command reveals a ‘c’ next to a route, that indicates that this is a directly connected network. An ‘l’ indicates the local route. In an IPv6 network, the local route has a /128 prefix. Local routes are used by the routing table to efficiently process packets with a destination address of the interface of the router.

Notice that there is also a route installed to the FF00::/8 network. This route is required for multicast routing.

Figure 1-35 displays how the **show ipv6 route** command can be combined with a specific network destination to display the details of how that route was learned by the router.

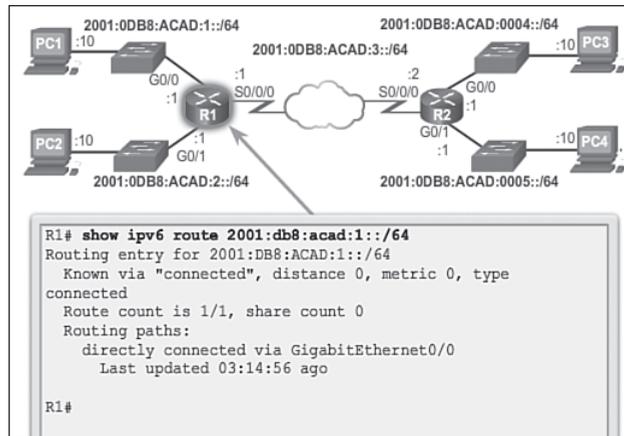


Figure 1-35 Show the IPv6 Route Entry

The following displays how connectivity to R2 can be verified using the **ping** command:

```

R1# ping 2001:db8:acad:3::2
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2001:DB8:ACAD:3::2, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 12/13/16 ms
R1#

```

Notice what happens when the G0/0 LAN interface of R2 is the target of the **ping** command:

```

R1# ping 2001:db8:acad:4::1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2001:DB8:ACAD:4::1, timeout is 2 seconds:

% No valid route for destination
Success rate is 0 percent (0/1)
R1#

```

The pings are unsuccessful. This is because R1 does not have an entry in the routing table to reach the 2001:DB8:ACAD:4::/64 network.

R1 requires additional information to reach a remote network. Remote network route entries can be added to the routing table using either:

- Static routing
- Dynamic routing protocols

Packet Tracer
Activity**Packet Tracer Activity 1.3.2.5: Investigating Directly Connected Routes**

The network in the activity is already configured. You will log in to the routers and use **show** commands to discover and answer the questions below about the directly connected routes.

Statically Learned Routes (1.3.3)

Routers can dynamically learn about remote networks using a routing protocol or statically. Statically learned routes must be manually configured by an administrator. This topic introduces how to enter remote routes manually.

Static Routes (1.3.3.1)

After directly connected interfaces are configured and added to the routing table, then static or dynamic routing can be implemented.

Static routes are manually configured. They define an explicit path between two networking devices. Unlike a dynamic routing protocol, static routes are not automatically updated and must be manually reconfigured if the network topology changes. The benefits of using static routes include improved security and resource efficiency. Static routes use less bandwidth than dynamic routing protocols, and no CPU cycles are used to calculate and communicate routes. The main disadvantage to using static routes is the lack of automatic reconfiguration if the network topology changes.

There are two common types of static routes in the routing table:

- Static route to a specific network
- Default static route

A static route can be configured to reach a specific remote network. IPv4 static routes are configured using the **ip route *network mask {next-hop-ip | exit-intf}*** global configuration command. A static route is identified in the routing table with the code 's'.

A default static route is similar to a default gateway on a host. The default static route specifies the exit point to use when the routing table does not contain a path for the destination network.

A default static route is useful when a router has only one exit point to another router, such as when the router connects to a central router or service provider.

To configure an IPv4 default static route, use the **ip route 0.0.0.0 0.0.0.0 {exit-intf | next-hop-ip}** global configuration command.

Figure 1-36 provides a simple scenario of how default and static routes can be applied.

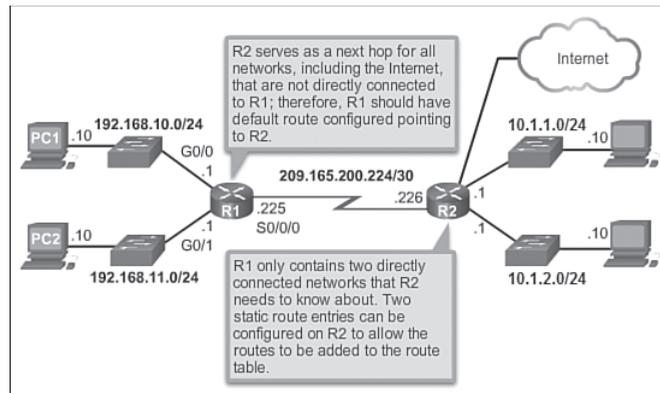


Figure 1-36 Static and Default Route Scenario

Static Route Examples (1.3.3.2)

Figure 1-37 shows the configuration of an IPv4 default static route on R1 to the Serial 0/0/0 interface. Notice that the configuration of the route generated an ‘s*’ entry in the routing table. The ‘s’ signifies that the route source is a static route, while the asterisk (*) identifies this route as a possible candidate to be the default route. In fact, it has been chosen as the default route as evidenced by the line that reads, “Gateway of last resort is 0.0.0.0 to network 0.0.0.0.”

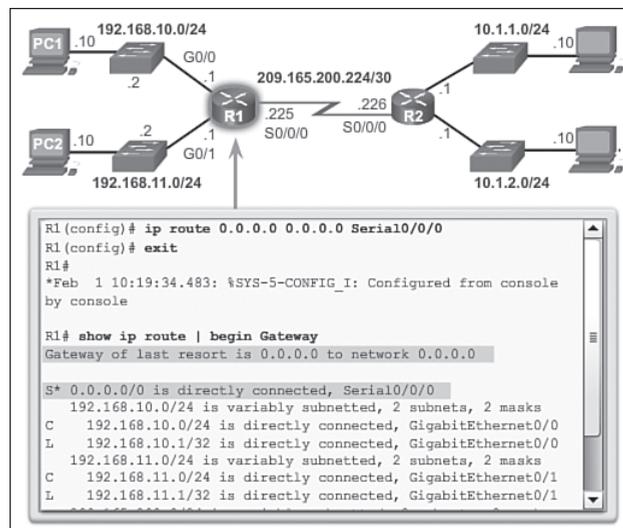


Figure 1-37 Entering and Verifying a Static Default Route

Note

The entire output of the **show ip route** command in Figure 1-37 can be viewed in the online course on page 1.3.3.2 graphic number 1.

Figure 1-38 shows the configuration of two static routes from R2 to reach the two LANs on R1. The route to 192.168.10.0/24 has been configured using the exit interface while the route to 192.168.11.0/24 has been configured using the next-hop IPv4 address. Although both are acceptable, there are some differences in how they operate. For instance, notice how different they look in the routing table. Also notice that because these static routes were to specific networks, the output indicates that the Gateway of Last Resort is not set.

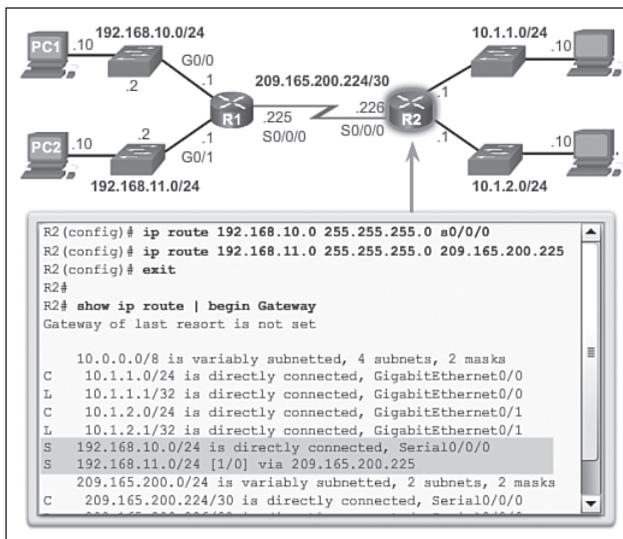


Figure 1-38 Entering and Verifying Static Routes

Note

The entire output of the **show ip route** command in Figure 1-38 can be viewed in the online course on page 1.3.3.2 graphic number 2.

Note

Static and default static routes are discussed in detail in the next chapter.

**Interactive
Graphic**
Activity 1.3.3.2: Entering and Verifying the Static and Default Routes on R1 and R2

Go to the online course to use the Syntax Checker in the third and fourth graphics to configure a default static route on router R1 and a static route on router R2.

Static IPv6 Route Examples (1.3.3.3)

Like IPv4, IPv6 supports static and default static routes. They are used and configured like IPv4 static routes.

To configure a default static IPv6 route, use the **ipv6 route ::/0** {*ipv6-address* | *interface-type interface-number*} global configuration command.

The following configures a default static route on R1 exiting out of the Serial 0/0/0 interface:

```
R1(config)# ipv6 route ::/0 s0/0/0
R1(config)# exit
R1#
```

Notice in the following output that the default static route configuration generated an 's' entry in the routing table. The 's' signifies that the route source is a static route. Unlike the IPv4 static route, there is no asterisk (*) or Gateway of Last Resort explicitly identified.

```
R1# show ipv6 route
IPv6 Routing Table - default - 8 entries
Codes: C - Connected, L - Local, S - Static, U - Per-user Static route
       B - BGP, R - RIP, H - NHRP, I1 - ISIS L1
       I2 - ISIS L2, IA - ISIS interarea, IS - ISIS summary, D - EIGRP
       EX - EIGRP external, ND - ND Default, NDp - ND Prefix, DCE - Destination
       NDr - Redirect, O - OSPF Intra, OI - OSPF Inter, OE1 - OSPF ext 1
       OE2 - OSPF ext 2, ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2
S    ::/0 [1/0]
     via Serial0/0/0, directly connected
C    2001:DB8:ACAD:1::/64 [0/0]
     via GigabitEthernet0/0, directly connected
L    2001:DB8:ACAD:1::1/128 [0/0]
     via GigabitEthernet0/0, receive
C    2001:DB8:ACAD:2::/64 [0/0]
     via GigabitEthernet0/1, directly connected
L    2001:DB8:ACAD:2::1/128 [0/0]
     via GigabitEthernet0/1, receive
C    2001:DB8:ACAD:3::/64 [0/0]
     via Serial0/0/0, directly connected
L    2001:DB8:ACAD:3::1/128 [0/0]
     via Serial0/0/0, receive
L    FF00::/8 [0/0]
     via Null0, receive
```

Like IPv4, static routes are routes explicitly configured to reach a specific remote network. Static IPv6 routes are configured using the **ipv6 route** *ipv6-prefix/prefix-length {ipv6-address | interface-type interface-number}* global configuration command.

The following example configures two static routes from R2 to reach the two LANs on R1:

```
R2(config)# ipv6 route 2001:DB8:ACAD:1::/64 2001:DB8:ACAD:3::1
R2(config)# ipv6 route 2001:DB8:ACAD:2::/64 s0/0/0
R2(config)# ^Z
R2#
```

The route to the 2001:0DB8:ACAD:2::/64 LAN is configured with an exit interface, while the route to the 2001:0DB8:ACAD:1::/64 LAN is configured with the next-hop IPv6 address. The next-hop IPv6 address can be either an IPv6 global unicast or link-local address.

The following output displays the routing table with the new static routes installed:

```
R2# show ipv6 route
IPv6 Routing Table - default - 9 entries
Codes: C - Connected, L - Local, S - Static, U - Per-user Static route
       B - BGP, R - RIP, H - NHRP, I1 - ISIS L1
       I2 - ISIS L2, IA - ISIS interarea, IS - ISIS summary, D - EIGRP
       EX - EIGRP external, ND - ND Default, NDp - ND Prefix, DCE - Destination
       Ndr - Redirect, O - OSPF Intra, OI - OSPF Inter, OE1 - OSPF ext 1
       OE2 - OSPF ext 2, ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2
S 2001:DB8:ACAD:1::/64 [1/0]
   via 2001:DB8:ACAD:3::1
S 2001:DB8:ACAD:2::/64 [1/0]
   via Serial0/0/0, directly connected
C 2001:DB8:ACAD:3::/64 [0/0]
   via Serial0/0/0, directly connected
L 2001:DB8:ACAD:3::2/128 [0/0]
   via Serial0/0/0, receive
C 2001:DB8:ACAD:4::/64 [0/0]
   via GigabitEthernet0/0, directly connected
L 2001:DB8:ACAD:4::1/128 [0/0]
   via GigabitEthernet0/0, receive
C 2001:DB8:ACAD:5::/64 [0/0]
   via GigabitEthernet0/1, directly connected
L 2001:DB8:ACAD:5::1/128 [0/0]
   via GigabitEthernet0/1, receive
```

```
L   FF00::/8 [0/0]
    via Null0, receive
R2#
```

The following confirms remote network connectivity to the 2001:0DB8:ACAD:4::/64 LAN on R2 from R1:

```
R1# ping 2001:db8:acad:4::1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2001:DB8:ACAD:4::1, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 12/13/16 ms
R1#
```

Dynamic Routing Protocols (1.3.4)

You just saw how a router can be manually configured with static routes to reach remote networks. In this section you will see how a dynamic routing protocol can be used to achieve the same result.

Dynamic Routing (1.3.4.1)

Dynamic routing protocols are used by routers to share information about the reachability and status of remote networks. Dynamic routing protocols perform several activities, including network discovery and maintaining routing tables.

Network discovery is the ability of a routing protocol to share information about the networks that it knows about with other routers that are also using the same routing protocol. Instead of depending on manually configured static routes to remote networks on every router, a dynamic routing protocol allows the routers to automatically learn about these networks from other routers. These networks, and the best path to each, are added to the routing table of the router, and identified as a network learned by a specific dynamic routing protocol.

During network discovery, routers exchange routes and update their routing tables. Routers have converged after they have finished exchanging and updating their routing tables. Routers then maintain the networks in their routing tables.

Figure 1-39 provides a simple scenario of how two neighboring routers would initially exchange routing information. In this simplified message exchange, R1 introduces itself and the networks it can reach. R2 responds and provides R1 with its networks.

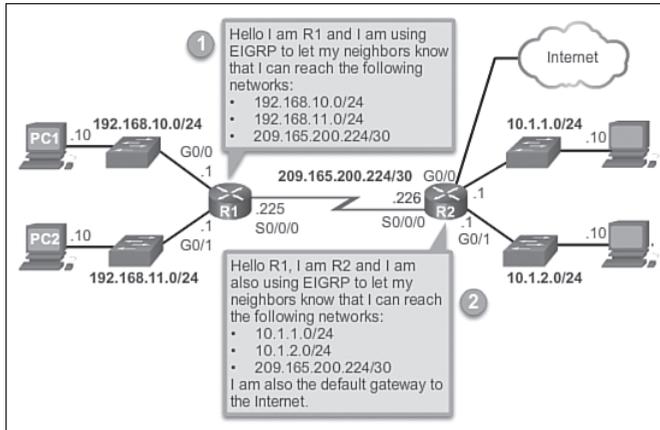


Figure 1-39 Dynamic Routing Scenario

IPv4 Routing Protocols (1.3.4.2)

A router running a dynamic routing protocol does not only make a best path determination to a network, it also determines a new best path if the initial path becomes unusable (or if the topology changes). For these reasons, dynamic routing protocols have an advantage over static routes. Routers that use dynamic routing protocols automatically share routing information with other routers and compensate for any topology changes without involving the network administrator.

Cisco ISR routers can support a variety of dynamic IPv4 routing protocols, including:

- **EIGRP:** Enhanced Interior Gateway Routing Protocol
- **OSPF:** Open Shortest Path First
- **IS-IS:** Intermediate System-to-Intermediate System
- **RIP:** Routing Information Protocol

To determine which routing protocols are supported by the IOS, use the **router ?** command in global configuration mode as shown in Figure 1-40.

Note

The focus of this course is on EIGRP and OSPF. RIP will be discussed only for legacy reasons; the other routing protocols supported by the IOS are beyond the scope of the CCNA certification.

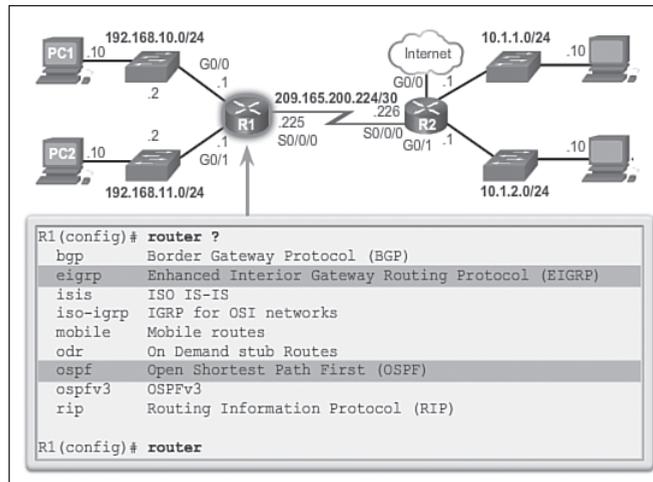


Figure 1-40 Supported IPv4 Routing Protocols

IPv4 Dynamic Routing Examples (1.3.4.3)

In this dynamic routing example, assume that R1 and R2 have been configured to support the dynamic routing protocol EIGRP. The routers also advertise directly connected networks. R2 advertises that it is the default gateway to other networks.

The output in Figure 1-41 displays the routing table of R1 after the routers have exchanged updates and converged.

```

R1# show ip route | begin Gateway
Gateway of last resort is 209.165.200.226 to network 0.0.0.0

D*EX 0.0.0.0/0 [170/2297856] via 209.165.200.226, 00:07:29, Serial0/0/0
    10.0.0.0/24 is subnetted, 2 subnets
    D    10.1.1.0 [90/2172416] via 209.165.200.226, 00:07:29, Serial0/0/0
    D    10.1.2.0 [90/2172416] via 209.165.200.226, 00:07:29, Serial0/0/0
    192.168.10.0/24 is variably subnetted, 2 subnets, 2 masks
    C    192.168.10.0/24 is directly connected, GigabitEthernet0/0
    L    192.168.10.1/32 is directly connected, GigabitEthernet0/0
    192.168.11.0/24 is variably subnetted, 2 subnets, 2 masks
    C    192.168.11.0/24 is directly connected, GigabitEthernet0/1
    L    192.168.11.1/32 is directly connected, GigabitEthernet0/1
    209.165.200.0/24 is variably subnetted, 2 subnets, 2 masks
    C    209.165.200.224/30 is directly connected, Serial0/0/0
    L    209.165.200.225/32 is directly connected, Serial0/0/0
R1#
  
```

Figure 1-41 Verifying Dynamic Routes

Along with the connected and link-local interfaces, there are three 'd' entries in the routing table:

- The entry beginning with 'D*EX' identifies that the source of this entry was EIGRP ('D'). The route is a candidate to be a default route ('*'), and the route is an external route ('EX') forwarded by EIGRP.

- The other two 'D' entries are routes installed in the routing table based on the update from R2 advertising its LANs.

IPv6 Routing Protocols (1.3.4.4)

As shown in Figure 1-42, ISR routers can support dynamic IPv6 routing protocols, including:

- RIPng (RIP next generation)
- OSPFv3
- EIGRP for IPv6

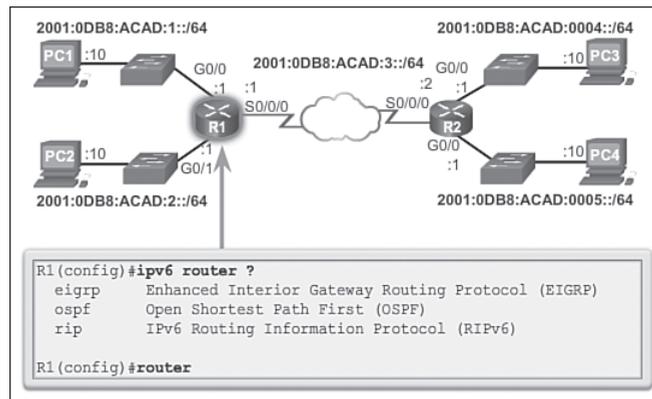


Figure 1-42 Supported IPv6 Routing Protocols

Support for dynamic IPv6 routing protocols is dependent on hardware and IOS version. Most of the modifications in the routing protocols are to support the longer IPv6 addresses and different header structures.

To enable IPv6 routers to forward traffic, you must configure the **ipv6 unicast-routing** global configuration command.

IPv6 Dynamic Routing Examples (1.3.4.5)

Routers R1 and R2 have been configured with the dynamic routing protocol EIGRP for IPv6. (This is the IPv6 equivalent of EIGRP for IPv4.)

To view the routing table on R1, enter the **show ipv6 route** command, as shown in the following output. The output displays the routing table of R1 after the routers have exchanged updates and converged. Along with the connected and local routes, there are two 'D' entries (EIGRP routes) in the routing table.

```
R1# show ipv6 route
IPv6 Routing Table - default - 9 entries
Codes: C - Connected, L - Local, S - Static, U - Per-user Static route
       B - BGP, R - RIP, H - NHRP, I1 - ISIS L1
       I2 - ISIS L2, IA - ISIS interarea, IS - ISIS summary, D - EIGRP
       EX - EIGRP external, ND - ND Default, NDp - ND Prefix, DCE - Destination
       NDr - Redirect, 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:ACAD:1::/64 [0/0]
  via GigabitEthernet0/0, directly connected
L 2001:DB8:ACAD:1::1/128 [0/0]
  via GigabitEthernet0/0, receive
C 2001:DB8:ACAD:2::/64 [0/0]
  via GigabitEthernet0/1, directly connected
L 2001:DB8:ACAD:2::1/128 [0/0]
  via GigabitEthernet0/1, receive
C 2001:DB8:ACAD:3::/64 [0/0]
  via Serial0/0/0, directly connected
L 2001:DB8:ACAD:3::1/128 [0/0]
  via Serial0/0/0, receive
D 2001:DB8:ACAD:4::/64 [90/2172416]
  via FE80::D68C:B5FF:FECE:A120, Serial0/0/0
D 2001:DB8:ACAD:5::/64 [90/2172416]
  via FE80::D68C:B5FF:FECE:A120, Serial0/0/0
L FF00::/8 [0/0]
  via Null0, receive
R1#
```

Summary (1.4)



Class Activity 1.4.1.1: We Really Could Use a Map!

Scenario

Use the Ashland and Richmond routing tables shown in the file provided with this activity.

With the help of a classmate, draw a network topology using the information from the tables.

To assist you with this activity, follow these guidelines:

- Start with the Ashland router—use its routing table to identify ports and IP addresses/networks.
- Add the Richmond router—use its routing table to identify ports and IP addresses/networks.
- Add any other intermediary and end devices as specified by the tables.

In addition, record answers from your group to the reflection questions provided with this activity.

Be prepared to share your work with another group and/or the class.

This chapter introduced the router. The main purpose of a router is to connect multiple networks and forward packets from one network to the next. This means that a router typically has multiple interfaces. Each interface is a member or host on a different IP network.

Cisco IOS uses what is known as the administrative distance (AD) to determine the route to install into the IP routing table. The routing table is a list of networks known by the router. The routing table includes network addresses for its own interfaces, which are the directly connected networks, as well as network addresses for remote networks. A remote network is a network that can only be reached by forwarding the packet to another router.

Remote networks are added to the routing table in one of two ways: either by the network administrator manually configuring static routes or by implementing a dynamic routing protocol. Static routes do not have as much overhead as dynamic routing protocols; however, static routes can require more maintenance if the topology is constantly changing or is unstable.

Dynamic routing protocols automatically adjust to changes without any intervention from the network administrator. Dynamic routing protocols require more CPU processing and also use a certain amount of link capacity for routing updates and messages. In many cases, a routing table will contain both static and dynamic routes.

Routers make their primary forwarding decision at Layer 3, the network layer. However, router interfaces participate in Layers 1, 2, and 3. Layer 3 IP packets are encapsulated into a Layer 2 data link frame and encoded into bits at Layer 1. Router interfaces participate in Layer 2 processes associated with their encapsulation. For example, an Ethernet interface on a router participates in the ARP process like other hosts on that LAN.

The Cisco IP routing table is not a flat database. The routing table is actually a hierarchical structure that is used to speed up the lookup process when locating routes and forwarding packets.

Components of the IPv6 routing table are very similar to the IPv4 routing table. For instance, it is populated using directly connected interfaces, static routes, and dynamically learned routes.

Practice

The following activities provide practice with the topics introduced in this chapter. The Labs and Class Activities are available in the companion *Routing Protocols Lab Manual* (978-1-58713-322-0). The Packet Tracer Activities PKA files are found in the online course.



Class Activities

Class Activity 1.0.1.2: Do We Really Need a Map?

Class Activity 1.4.1.1: We Really Could Use a Map!



Labs

Lab 1.1.1.9: Mapping the Internet

Lab 1.1.4.6: Configuring Basic Router Settings with IOS CLI

Lab 1.1.4.7: Configuring Basic Router Settings with CCP



Packet Tracer Activities

Packet Tracer Activity 1.1.1.8: Using Traceroute to Discover the Network

Packet Tracer Activity 1.1.2.9: Documenting the Network

Packet Tracer Activity 1.1.3.5: Configuring IPv4 and IPv6 Interfaces

Packet Tracer Activity 1.1.4.5: Configuring and Verifying a Small Network

Packet Tracer Activity 1.3.2.5: Investigating Directly Connected Routes

Check Your Understanding Questions

Complete all the review questions listed here to test your understanding of the topics and concepts in this chapter. The appendix, “Answers to the ‘Check Your Understanding’ Questions,” lists the answers.

1. Which of the following matches a router component with its function?
 - A. Flash: Permanently stores the bootstrap program
 - B. ROM: Permanently stores the startup configuration file
 - C. NVRAM: Permanently stores the operating system image
 - D. RAM: Stores the routing tables and ARP cache

2. Which command can a technician use to determine whether router serial ports have IP addresses that are assigned to them?
 - A. `show interfaces ip brief`
 - B. `show controllers all`
 - C. `show ip config`
 - D. `show ip interface brief`

3. Which of the following commands will set and automatically encrypt the privileged EXEC mode password to “quiz”?
 - A. `R1(config)# enable secret quiz`
 - B. `R1(config)# password secret quiz`
 - C. `R1(config)# enable password secret quiz`
 - D. `R1(config)# enable secret password quiz`

4. Which routing principle is correct?
 - A. If one router has certain information in its routing table, all adjacent routers have the same information.
 - B. Routing information about a path from one network to another implies routing information about the reverse, or return, path.
 - C. Every router makes its routing decisions alone, based on the information it has in its own routing table.
 - D. Every router makes its routing decisions based on the information it has in its own routing table and the information in its neighbor routing tables.

5. What two tasks do dynamic routing protocols perform? (Choose two.)
 - A. Discover hosts
 - B. Update and maintain routing tables
 - C. Propagate host default gateways
 - D. Network discovery
 - E. Assign IP addressing

6. A network engineer is configuring a new router. The interfaces have been configured with IP addresses and activated, but no routing protocols or static routes have been configured yet. What routes are present in the routing table?
 - A. Default routes
 - B. Remote network routes
 - C. Directly connected routes
 - D. No route as the routing table is empty

7. Which statements are correct regarding how a router forwards packets? (Choose two.)
 - A. If the packet is destined for a remote network, the router forwards the packet out all interfaces that might be a next hop to that network.
 - B. If the packet is destined for a directly connected network, the router forwards the packet out the exit interface indicated by the routing table.
 - C. If the packet is destined for a remote network, the router forwards the packet based on the information in the router host table.
 - D. If the packet is destined for a remote network, the router sends the packet to the next-hop IP address in the routing table.
 - E. If the packet is destined for a directly connected network, the router forwards the packet based on the destination MAC address.
 - F. If the packet is destined for a directly connected network, the router forwards the packet to the switch on the next-hop VLAN.

8. Which command is used to explicitly configure a local IPv6 address on a router interface?
 - A. `ipv6 enable`
 - B. `ipv6 address ipv6-address/prefix-length`
 - C. `ipv6 address ipv6-address/prefix-length eui-64`
 - D. `ipv6 address ipv6-address/prefix-length link-local`

9. Which statement is true regarding metrics used by routing protocols?
- A. A metric is the quantitative value that a routing protocol uses to measure a given route.
 - B. A metric is a Cisco-proprietary means to convert distances to a standard unit.
 - C. Metrics represent a composite value of the amount of packet loss occurring for all routing protocols.
 - D. Metrics are used by the router to determine whether a packet has an error and should be dropped.
10. The network administrator configured the `ip route 0.0.0.0 0.0.0.0 serial 0/0/0` command on the router. How will this command appear in the routing table, assuming that the Serial 0/0/0 interface is up?
- A. D 0.0.0.0/0 is directly connected, Serial0/0/0
 - B. S* 0.0.0.0/0 is directly connected, Serial0/0/0
 - C. S* 0.0.0.0/0 [1/0] via 192.168.2.2
 - D. C 0.0.0.0/0 [1/0] via 192.168.2.2
11. How many equal-cost paths can a dynamic routing protocol use for load balancing by default?
- A. 2
 - B. 3
 - C. 4
 - D. 6
12. What two statements correctly describe the concepts of administrative distance and metric? (Choose two.)
- A. Administrative distance refers to the trustworthiness of a particular route.
 - B. A router first installs routes with higher administrative distances in its routing table.
 - C. Routes with the smallest metric to the destination indicate the best path.
 - D. Metrics are used by the router to determine whether a packet has an error and should be dropped.
 - E. The metric is always determined based on hop count.

13. When a packet travels from router to router to its destination, what address continually changes from hop to hop?
 - A. Source and destination Layer 2 address
 - B. Source Layer 3 address
 - C. Destination Layer 3 address
 - D. Destination port
14. Describe the internal router hardware components, and outline the purpose of each.
15. Describe the router bootup process from power on to final configuration.
16. What are two important functions that a router performs?
17. Describe the steps necessary to configure basic settings on a router.
18. Describe the importance of the routing table. What purposes does it serve?
19. What are the three basic ways a router learns about networks?
20. What three pieces of information must be configured on a host to forward packets to remote networks? (Choose three.)
 - A. Clock rate
 - B. Default gateway
 - C. DHCP server address hostname
 - D. DNS server address
 - E. IP address
 - F. Subnet mask
21. A serial interface has been configured with an IP address and the clock rate. However, the **show ip interface brief** command indicates that the interface is administratively down. What must be done to correct the problem?
22. What type of IPv6 address must be configured on an IPv6-enabled interface?
23. When a computer is pinging another computer for the first time, which type of message does it send first to determine the MAC address of the other device?

Static Routing

Objectives

Upon completion of this chapter, you will be able to answer the following questions:

- What are the advantages and disadvantages of static routes?
- Can you explain the purpose of different types of static routes?
- Can you configure IPv4 and IPv6 static routes by specifying a next-hop address?
- How is legacy classful addressing used in network implementation?
- What is the purpose of CIDR in replacing classful addressing?
- How do you design and implement a hierarchical addressing scheme?
- How do you configure an IPv4 and IPv6 summary network address to reduce the number of routing table entries?
- Can you configure a floating static route to provide a backup connection?
- How does a router process packets when a static route is configured?
- How do you troubleshoot common static and default route configuration issues?

Key Terms

This chapter uses the following key terms. You can find the definitions in the Glossary.

stub network page 77

summary static route page 80

floating static route page 81

recursive lookup page 86

fully specified static route page 89

quad-zero route page 93

stub router page 106

summary route page 128

route summarization page 128

Introduction (2.0.1.1)

Routing is at the core of every data network, moving information across an internet-network from source to destination. Routers are the devices responsible for the transfer of packets from one network to the next.

Routers learn about remote networks either dynamically, using routing protocols, or manually, using static routes. In many cases, routers use a combination of both dynamic routing protocols and static routes. This chapter focuses on static routing.

Static routes are very common and do not require the same amount of processing and overhead as dynamic routing protocols.

In this chapter, sample topologies will be used to configure IPv4 and IPv6 static routes and to present troubleshooting techniques. In the process, several important IOS commands and the resulting output will be examined. An introduction to the routing table using both directly connected networks and static routes will be included.

This chapter will also contrast classful routing and the widely implemented classless routing methods. It will cover Classless Inter-Domain Routing (CIDR) and the variable-length subnet mask (VLSM) methods. CIDR and VLSM have helped conserve the IPv4 address space using subnetting and summarization techniques.



Class Activity 2.0.1.2: Which Way Should We Go?

A huge sporting event is about to take place in your city. To attend the event, you make concise plans to arrive at the sports arena on time to see the entire game.

There are two routes you can take to drive to the event:

- **Highway route:** It is easy to follow and fast driving speeds are allowed.
- **Alternative, direct route:** You found this route using a city map. Depending on conditions, such as the amount of traffic or congestion, this just may be the way to get to the arena on time!

With a partner, discuss these options. Choose a preferred route to arrive at the arena in time to see every second of the huge sporting event.

Compare your optional preferences to network traffic, which route would you choose to deliver data communications for your small- to medium-sized business? Would your network route be the fastest, easiest route or the alternative, direct route? Justify your choice.

Complete the modeling activity .pdf and be prepared to justify your answers to the class or with another group.

Static Routing Implementation (2.1)

As previously stated, static routes are widely used in networks today. Static routes are used in networks of all sizes, and are used along with a dynamic routing protocol. For this reason, a good understanding of static routes is a requirement for implementing routing on a network.

Reach Remote Networks (2.1.1.1)

A router can learn about remote networks in one of two ways:

- **Manually:** Remote networks are manually entered into the route table using static routes.
- **Dynamically:** Remote routes are automatically learned using a dynamic routing protocol.

Figure 2-1 provides a sample scenario of static routing. Figure 2-2 provides a sample scenario of dynamic routing using EIGRP.

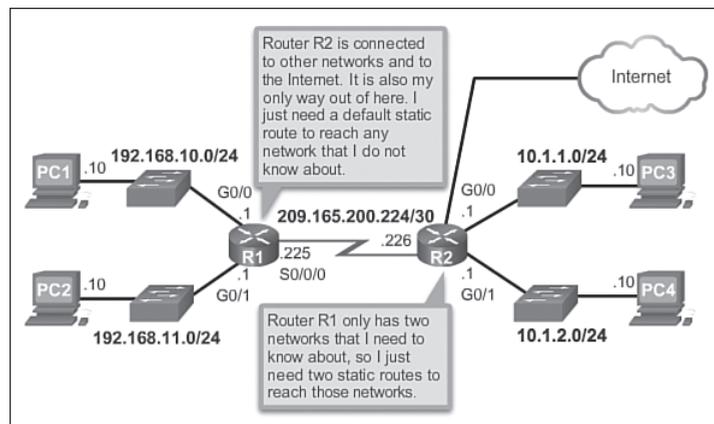


Figure 2-1 Static and Default Route Scenario

A network administrator can manually configure a static route to reach a specific network. Unlike a dynamic routing protocol, static routes are not automatically updated and must be manually reconfigured any time the network topology changes. A static route does not change until the administrator manually reconfigures it.

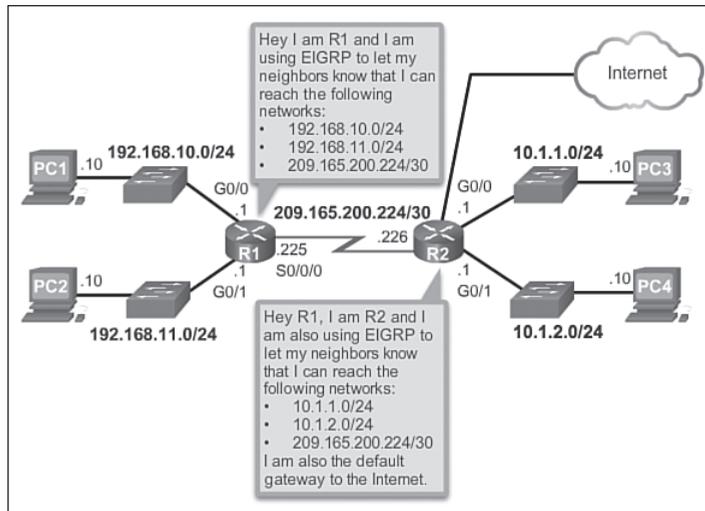


Figure 2-2 Dynamic Routing Scenario

Why Use Static Routing? (2.1.1.2)

Static routing provides some advantages over dynamic routing, including:

- Static routes are not advertised over the network, resulting in better security.
- Static routes use less bandwidth than dynamic routing protocols, as routers do not exchange routes.
- No CPU cycles are used to calculate and communicate routes.
- The path a static route uses to send data is known.

Static routing has the following disadvantages:

- Initial configuration and maintenance is time-consuming.
- Configuration can be error-prone, especially in large networks.
- Administrator intervention is required to maintain changing route information.
- Does not scale well with growing networks; maintenance becomes cumbersome.
- Requires complete knowledge of the whole network for proper implementation.

In Table 2-1, dynamic and static routing features are compared. Notice that the advantages of one method are the disadvantages of the other.

Table 2-1 Dynamic Routing Versus Static Routing

	Dynamic Routing	Static Routing
Configuration Complexity	Generally independent of the network size	Increases with the network size
Topology Changes	Automatically adapts to topology changes	Administrator intervention required
Scaling	Suitable for simple and complex topologies	Suitable for simple topologies
Security	Less secure	More secure
Resource Usage	Uses CPU, memory, link bandwidth	No extra resources needed
Predictability	Route depends on the current topology	Route to destination is always the same

Static routes are useful for smaller networks with only one path to an outside network. They also provide security in a larger network for certain types of traffic or links to other networks that need more control. It is important to understand that static and dynamic routing are not mutually exclusive. Rather, most networks use a combination of dynamic routing protocols and static routes. This may result in the router having multiple paths to a destination network via static routes and dynamically learned routes. However, the administrative distance (AD) of a static route is 1. Therefore, a static route will take precedence over all dynamically learned routes.

When to Use Static Routes (2.1.1.3)

Static routing has three primary uses:

- Providing ease of routing table maintenance in smaller networks that are not expected to grow significantly.
- Routing to and from stub networks. A *stub network* is a network accessed by a single route, and the router has only one neighbor.
- Using a single default route to represent a path to any network that does not have a more specific match with another route in the routing table. Default routes are used to send traffic to any destination beyond the next upstream router.

Figure 2-3 shows an example of a stub network connection and a default route connection. Notice in the figure that any network attached to R1 would only have one

way to reach other destinations, whether to networks attached to R2, or to destinations beyond R2. This means that network 172.16.3.0 is a stub network and R1 is a stub router. Running a routing protocol between R2 and R1 is a waste of resources.

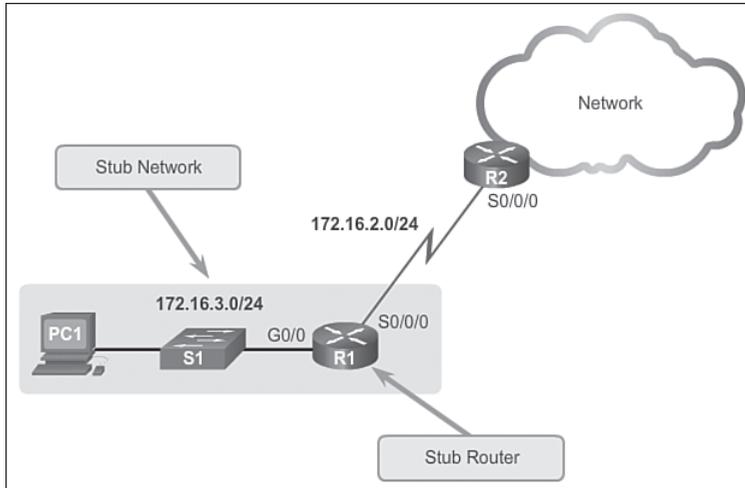


Figure 2-3 Stub Networks and Stub Routers

In this example, a static route can be configured on R2 to reach the R1 LAN. Additionally, because R1 has only one way to send out non-local traffic, a default static route can be configured on R1 to point to R2 as the next hop for all other networks.

**Interactive
Graphic**

Activity 2.1.1.4: Identify the Advantages and Disadvantages of Static Routing

Go to the online course to perform this practice activity.

Static Route Applications (2.1.2.1)

Static routes are most often used to connect to a specific network or to provide a Gateway of Last Resort for a stub network. They can also be used to:

- Reduce the number of routes advertised by summarizing several contiguous networks as one static route
- Create a backup route in case a primary route link fails

The following types of IPv4 and IPv6 static routes will be discussed:

- Standard static route
- Default static route

- Summary static route
- Floating static route

Standard Static Route (2.1.2.2)

Both IPv4 and IPv6 support the configuration of static routes. Static routes are useful when connecting to a specific remote network.

Figure 2-4 shows that R2 can be configured with a static route to reach the stub network 172.16.3.0/24.

Note

The example is highlighting a stub network, but in fact, a static route can be used to connect to any network.

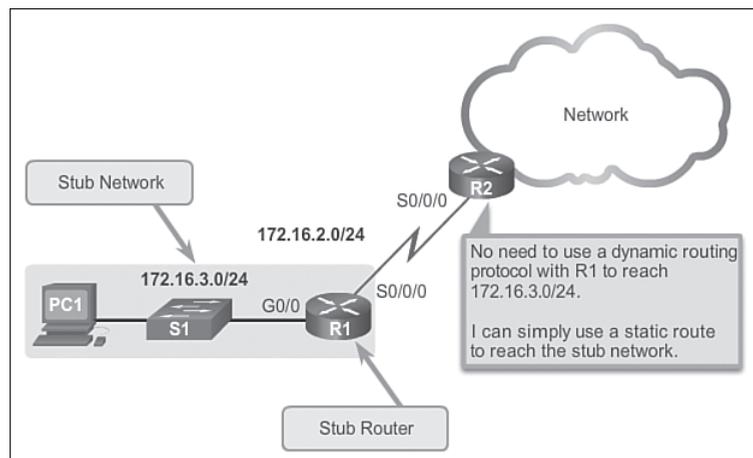


Figure 2-4 Connecting to a Stub Network

Default Static Route (2.1.2.3)

A default static route is a route that matches all packets. A default route identifies the gateway IP address to which the router sends all IP packets that it does not have a learned or static route for. A default static route is simply a static route with 0.0.0.0/0 as the destination IPv4 address. Configuring a default static route creates a Gateway of Last Resort.

Note

All routes that identify a specific destination with a larger subnet mask take precedence over the default route.

Default static routes are used:

- When no other routes in the routing table match the packet destination IP address. In other words, when a more specific match does not exist. A common use is when connecting a company's edge router to the ISP network.
- When a router has only one other router to which it is connected. This condition is known as a stub router.

Refer to Figure 2-5 for a sample scenario of implementing default static routing.

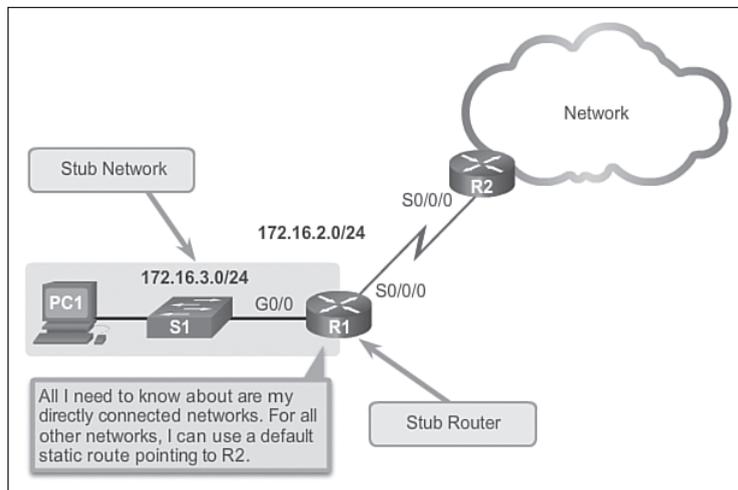


Figure 2-5 Connecting to a Stub Router

Summary Static Route (2.1.2.4)

To reduce the number of routing table entries, multiple static routes can be summarized into a single *summary static route* if:

- The destination networks are contiguous and can be summarized into a single network address.
- The multiple static routes all use the same exit interface or next-hop IP address.

In Figure 2-6, R1 would require four separate static routes to reach the 172.20.0.0/16 to 172.23.0.0/16 networks. Instead, one *summary static route* can be configured and still provide connectivity to those networks.

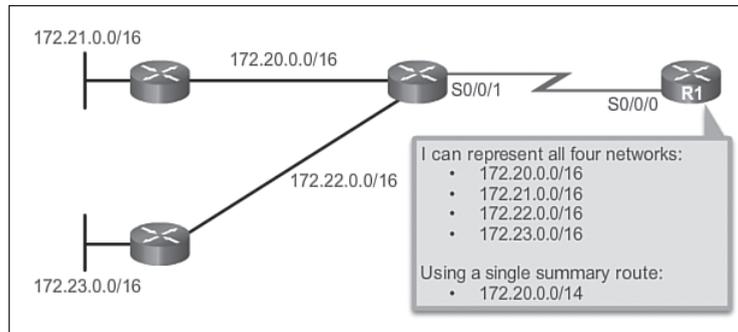


Figure 2-6 Using One Summary Static Route

Floating Static Route (2.1.2.5)

Another type of static route is a *floating static route*. Floating static routes are static routes that are used to provide a backup path to a primary static or dynamic route, in the event of a link failure. The floating static route is only used when the primary route is not available.

To accomplish this, the floating static route is configured with a higher administrative distance than the primary route. Recall that the administrative distance represents the trustworthiness of a route. If multiple paths to the destination exist, the router will choose the path with the lowest administrative distance.

For example, assume that an administrator wants to create a floating static route as a backup to an EIGRP-learned route. The floating static route must be configured with a higher administrative distance than EIGRP. EIGRP has an administrative distance of 90. If the floating static route is configured with an administrative distance of 95, the dynamic route learned through EIGRP is preferred to the floating static route. If the EIGRP-learned route is lost, the floating static route is used in its place.

In Figure 2-7, the Branch router typically forwards all traffic to the HQ router over the private WAN link. In this example, the routers exchange route information using EIGRP. A floating static route, with an administrative distance of 91 or higher, could be configured to serve as a backup route. If the private WAN link fails and the EIGRP route disappears from the routing table, the router selects the floating static route as the best path to reach the HQ LAN.

Interactive
Graphic

Activity 2.1.2.6: Identify the Type of Static Route

Go to the online course to perform this practice activity.

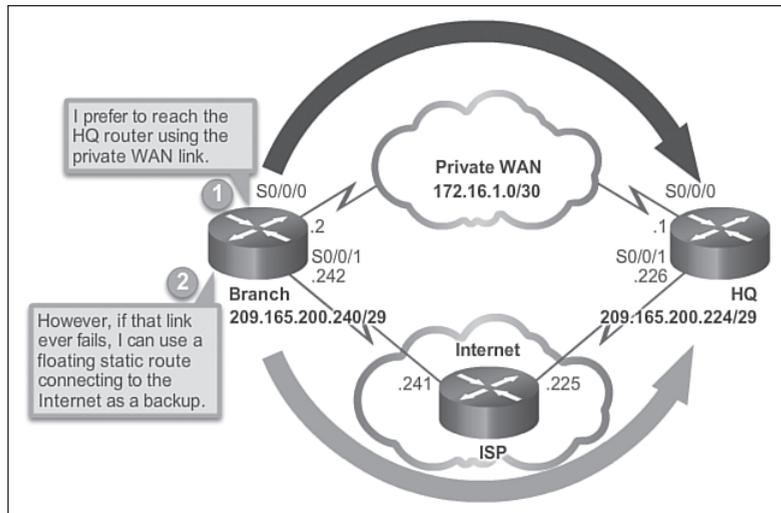


Figure 2-7 Configuring a Backup Route

Configure Static and Default Routes (2.2)

Recall that a router can learn about remote networks in one of two ways:

- Manually, from configured static routes
- Automatically, from a dynamic routing protocol

This chapter focuses on configuring static routes. Dynamic routing protocols are introduced in the next chapter. Even with the implementation of a dynamic routing protocol, static routes are also commonly used.

Configure IPv4 Static Routes (2.2.1)

Routers can be configured for both IPv4 and IPv6 static routes. This section will focus on the configuration of IPv4 static routes. IPv4 static routes are manually configured routing entries for reaching IPv4 networks. The configuration of IPv6 static routes is covered later in this chapter.

ip route Command (2.2.1.1)

Static routes are configured using the **ip route** global configuration command. The syntax of the command is:

```
Router(config)# ip route network-address subnet-mask { ip-address | interface-type
interface-number [ ip-address ] } [ distance ] [ name name ] [ permanent ] [ tag tag ]
```

The following parameters are required to configure static routing:

- *network-address*: Destination network address of the remote network to be added to the routing table; often this is referred to as the prefix.
- *subnet-mask*: Subnet mask, or just mask, of the remote network to be added to the routing table. The subnet mask can be modified to summarize a group of networks.

One or both of the following parameters must also be used:

- *ip-address*: The IP address of the connecting router to use to forward the packet to the remote destination network. Commonly referred to as the next hop.
- *interface-type interface-number or exit-intf*: The outgoing interface to use to forward the packet to the next hop.

As shown in Table 2-2, the command syntax commonly used is **ip route *network-address subnet-mask* {*ip-address* | *exit-intf*}**.

The *distance* parameter is used to create a floating static route by setting an administrative distance that is higher than a dynamically learned route.

Table 2-2 iproute Command Syntax

Parameter	Description
<i>network-address</i>	Destination network address of the remote network to be added to the routing table.
<i>subnet-mask</i>	<ul style="list-style-type: none"> ■ Subnet mask of the remote network to be added to the routing table. ■ The subnet mask can be modified to summarize a group of networks.
<i>ip-address</i>	<ul style="list-style-type: none"> ■ Commonly referred to as the next-hop router's IP address. ■ Typically used when connecting to a broadcast media (i.e., Ethernet). ■ Commonly creates a recursive lookup.
<i>exit-intf</i>	<ul style="list-style-type: none"> ■ Use the outgoing interface to forward packets to the destination network. ■ Also referred to as a directly attached static route. ■ Typically used when connecting in a point-to-point configuration.

Note

The remaining parameters are not relevant for this chapter or for CCNA studies.

Next-Hop Options (2.2.1.2)

Figure 2-8 and the subsequent outputs display the topology and the routing tables of R1, R2, and R3. Notice that each router has entries only for directly connected networks and their associated local addresses. None of the routers have any knowledge of any networks beyond their directly connected interfaces.

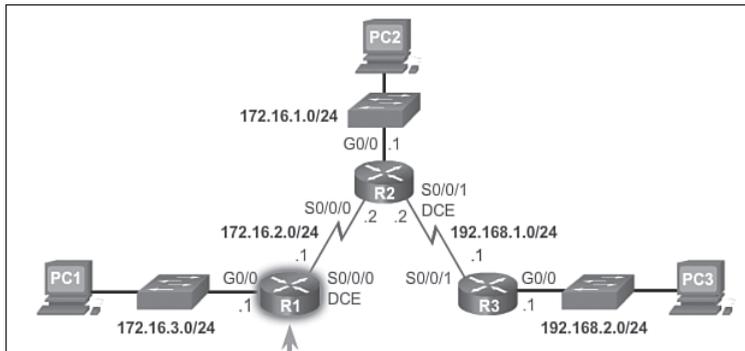


Figure 2-8 Verify the Routing Table of R1, R2, R3

```
R1# show ip route | begin Gateway
Gateway of last resort is not set
    172.16.0.0/16 is variably subnetted, 4 subnets, 2 masks
C       172.16.2.0/24 is directly connected, Serial0/0/0
L       172.16.2.1/32 is directly connected, Serial0/0/0
C       172.16.3.0/24 is directly connected, GigabitEthernet0/0
L       172.16.3.1/32 is directly connected, GigabitEthernet0/0
R1#

R2# show ip route | begin Gateway
Gateway of last resort is not set
    172.16.0.0/16 is variably subnetted, 4 subnets, 2 masks
C       172.16.1.0/24 is directly connected, GigabitEthernet0/0
L       172.16.1.1/32 is directly connected, GigabitEthernet0/0
C       172.16.2.0/24 is directly connected, Serial0/0/0
L       172.16.2.2/32 is directly connected, Serial0/0/0
    192.168.1.0/24 is variably subnetted, 2 subnets, 2 masks
C       192.168.1.0/24 is directly connected, Serial0/0/1
L       192.168.1.2/32 is directly connected, Serial0/0/1
R2#

R3# show ip route | include C
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
C       192.168.1.0/24 is directly connected, Serial0/0/1
C       192.168.2.0/24 is directly connected, GigabitEthernet0/0
R3#
```

For example, R1 has no knowledge of networks:

- 172.16.1.0/24: LAN on R2
- 192.168.1.0/24: Serial network between R2 and R3
- 192.168.2.0/24: LAN on R3

Figure 2-9 displays a successful ping from R1 to R2. Figure 2-10 displays an unsuccessful ping to the R3 LAN. This is unsuccessful because R1 does not have an entry in its routing table for the R3 LAN network.

```
R1# ping 172.16.2.2
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.16.2.2, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 12/13/16
ms
R1#
```

Figure 2-9 Verify Connectivity from R1 to R2

```
R1# ping 192.168.2.1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.168.2.1, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
R1#
```

Figure 2-10 Verify Connectivity from R1 to R3 LAN

The next hop can be identified by an IP address, exit interface, or both. How the destination is specified creates one of the three following route types:

- **Next-hop route:** Only the next-hop IP address is specified.
- **Directly connected static route:** Only the router exit interface is specified.
- **Fully specified static route:** The next-hop IP address and exit interface are specified.

Configure a Next-Hop Static Route (2.2.1.3)

In a next-hop static route, only the next-hop IP address is specified. The output interface is derived from the next hop. For example, in Figure 2-11, three next-hop static routes are configured on R1 using the IP address of the next hop, R2.

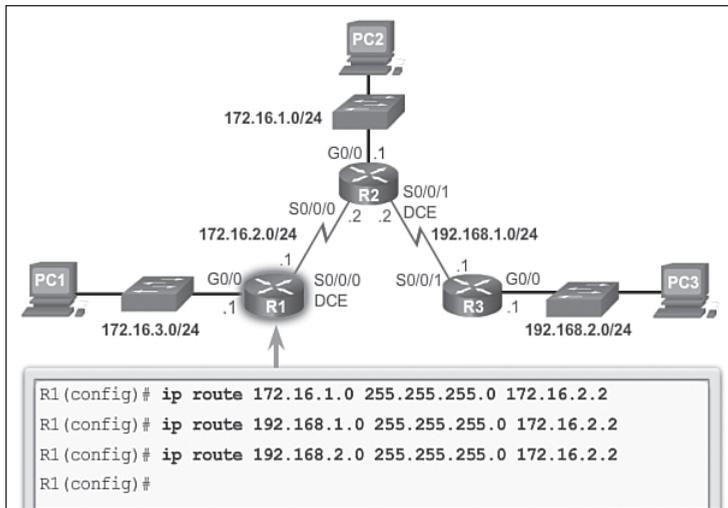


Figure 2-11 Configuring Next-Hop Static Routes on R1

Before any packet is forwarded by a router, the routing table process must determine the exit interface to use to forward the packet. This is known as route resolvability. The route resolvability process will vary depending upon the type of forwarding mechanism being used by the router. CEF (Cisco Express Forwarding) is the default behavior on most platforms running IOS 12.0 or later.

Figure 2-12 details the basic packet forwarding process in the routing table for R1 without the use of CEF. When a packet is destined for the 192.168.2.0/24 network, R1:

1. Looks for a match in the routing table and finds that it has to forward the packets to the next-hop IPv4 address 172.16.2.2, as indicated by the label 1 in Figure 2-12. Every route that references only a next-hop IPv4 address and does not reference an exit interface must have the next-hop IPv4 address resolved using another route in the routing table with an exit interface.
2. R1 must now determine how to reach 172.16.2.2; therefore, it searches a second time for a 172.16.2.2 match. In this case, the IPv4 address matches the route for the directly connected network 172.16.2.0/24 with the exit interface Serial 0/0/0, as indicated by the label 2 in Figure 2-12. This lookup tells the routing table process that this packet is forwarded out of that interface.

It actually takes two routing table lookup processes to forward any packet to the 192.168.2.0/24 network. When the router performs multiple lookups in the routing table before forwarding a packet, it is performing a process known as a *recursive lookup*. Because recursive lookups consume router resources, they should be avoided when possible.

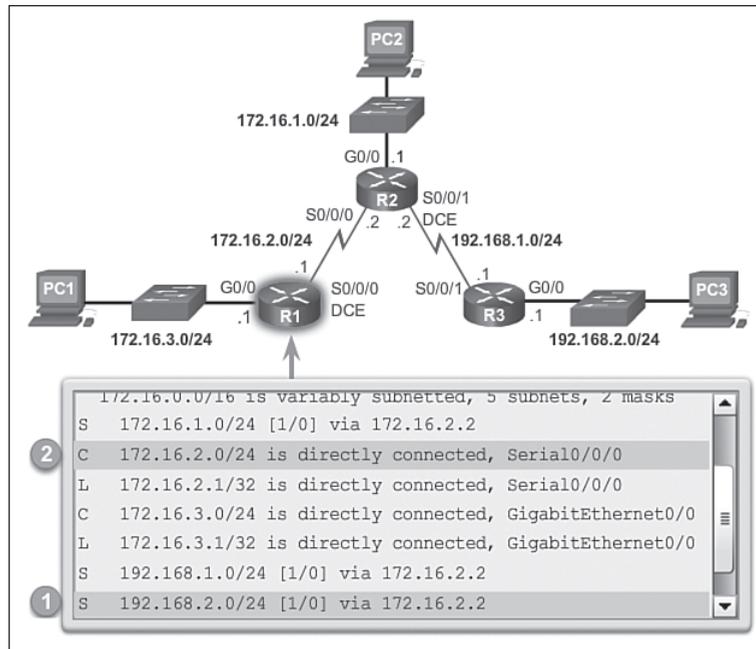


Figure 2-12 Verify the Routing Table of R1

A recursive static route is valid (that is, it is a candidate for insertion in the routing table) only when the specified next hop resolves, either directly or indirectly, to a valid exit interface.

Note

CEF provides optimized lookup for efficient packet forwarding by using two main data structures stored in the data plane: an FIB (Forwarding Information Base), which is a copy of the routing table, and an adjacency table that includes Layer 2 addressing information. By combining the information from these two tables, CEF eliminates the need for recursive lookup for next-hop IP address lookups. In other words, a static route using a next-hop IP address requires only a single lookup when CEF is enabled on the router.

Configure a Directly Connected Static Route (2.2.1.4)

When configuring a static route, another option is to use the exit interface to specify the next-hop address. In older IOS versions, prior to CEF, this method is used to avoid the recursive lookup problem.

In Figure 2-13, three directly connected static routes are configured on R1 using the exit interface. The routing table for R1 in Figure 2-14 shows that when a packet is destined for the 192.168.2.0/24 network, R1 looks for a match in the routing table,

and finds that it can forward the packet out of its Serial 0/0/0 interface. No other lookups are required.

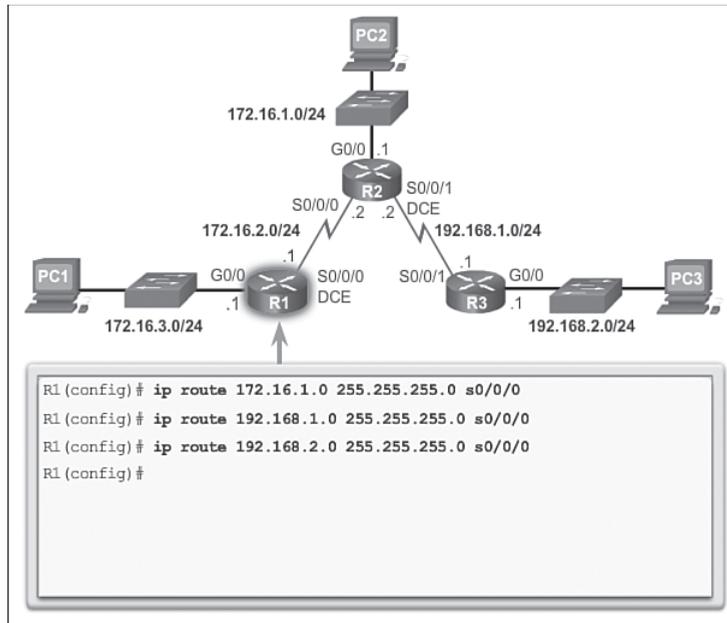


Figure 2-13 Configure Directly Connected Static Routes on R1

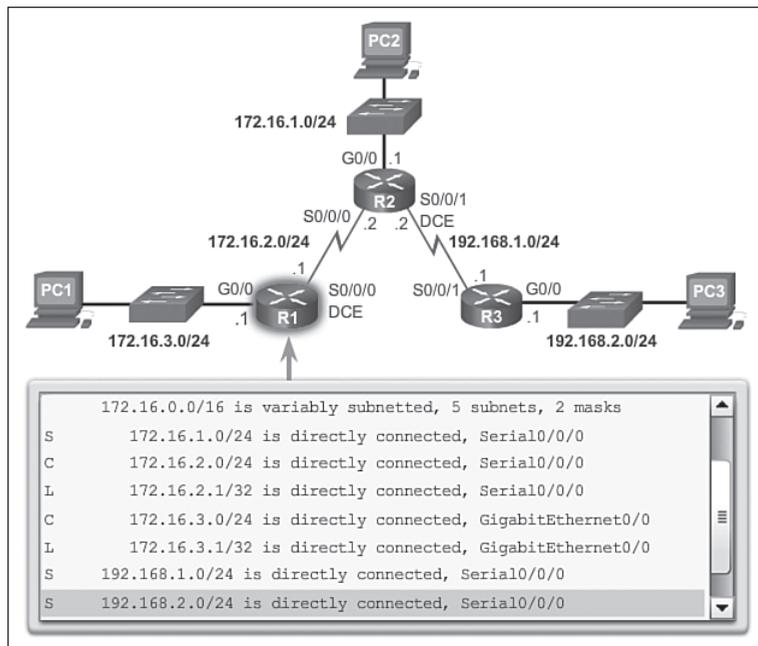


Figure 2-14 Verify the Routing Table of R1

Notice how the routing table looks different for the route configured with an exit interface than the route configured with a recursive entry.

Configuring a directly connected static route with an exit interface allows the routing table to resolve the exit interface in a single search, instead of two searches. Although the routing table entry indicates “directly connected,” the administrative distance of the static route is still 1. Only a directly connected interface can have an administrative distance of 0.

Note

For point-to-point interfaces, you can use static routes that point to the exit interface or to the next-hop address. For multipoint/broadcast interfaces, it is more suitable to use static routes that point to a next-hop address.

Interactive Graphic

Activity 2.2.1.4 Part 1: Configure Directly Connected Static Routes on R2

Go to the online course to use the Syntax Checker in the third graphic to configure a static route to the 172.16.3.0/24 network using exit interface S0/0/0.

Interactive Graphic

Activity 2.2.1.4 Part 2: Configure Directly Connected Static Routes on R3

Go to the online course to use the Syntax Checker in the fourth graphic to configure static routes to the 172.16.1.0/24, 172.16.2.0/24, and 172.16.3.0/24 networks using the exit interface S0/0/1.

Although static routes that use only an exit interface on point-to-point networks are common, the use of the default CEF forwarding mechanism makes this practice unnecessary.

Configure a Fully Specified Static Route (2.2.1.5)

In a *fully specified static route*, both the output interface and the next-hop IP address are specified. This is another type of static route that is used in older IOS versions, prior to CEF. This form of static route is used when the output interface is a multi-access interface and it is necessary to explicitly identify the next hop. The next hop must be directly connected to the specified exit interface.

Suppose that the network link between R1 and R2 is an Ethernet link and that the GigabitEthernet 0/1 interface of R1 is connected to that network, as shown in Figure 2-15. CEF is not enabled. To eliminate the recursive lookup, a directly connected static route can be implemented using the following command:

```
R1(config)# ip route 192.168.2.0 255.255.255.0 GigabitEthernet 0/1
```

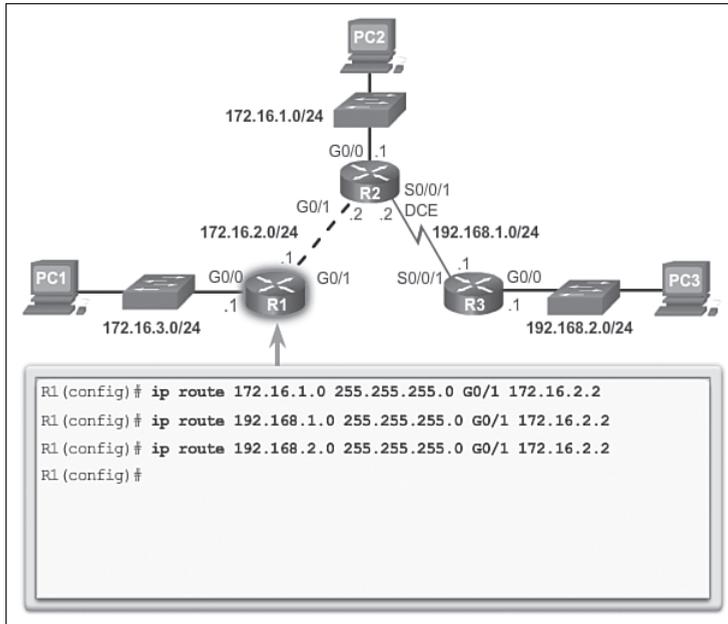


Figure 2-15 Configure Fully Specified Static Routes on R1

However, this may cause unexpected or inconsistent results. The difference between an Ethernet multi-access network and a point-to-point serial network is that a point-to-point network has only one other device on that network, the router at the other end of the link. With Ethernet networks, there may be many different devices sharing the same multi-access network, including hosts and even multiple routers. By only designating the Ethernet exit interface in the static route, the router will not have sufficient information to determine which device is the next-hop device.

R1 knows that the packet needs to be encapsulated in an Ethernet frame and sent out the GigabitEthernet 0/1 interface. However, R1 does not know the next-hop IPv4 address and therefore it cannot determine the destination MAC address for the Ethernet frame.

Depending upon the topology and the configurations on other routers, this static route may or may not work. It is recommended that when the exit interface is an Ethernet network, a fully specified static route is used including both the exit interface and the next-hop address.

As shown in Figure 2-16, when forwarding packets to R2, the exit interface is GigabitEthernet 0/1 and the next-hop IPv4 address is 172.16.2.2.

Note

With the use of CEF, a fully specified static route is no longer necessary. A static route using a next-hop address should be used.

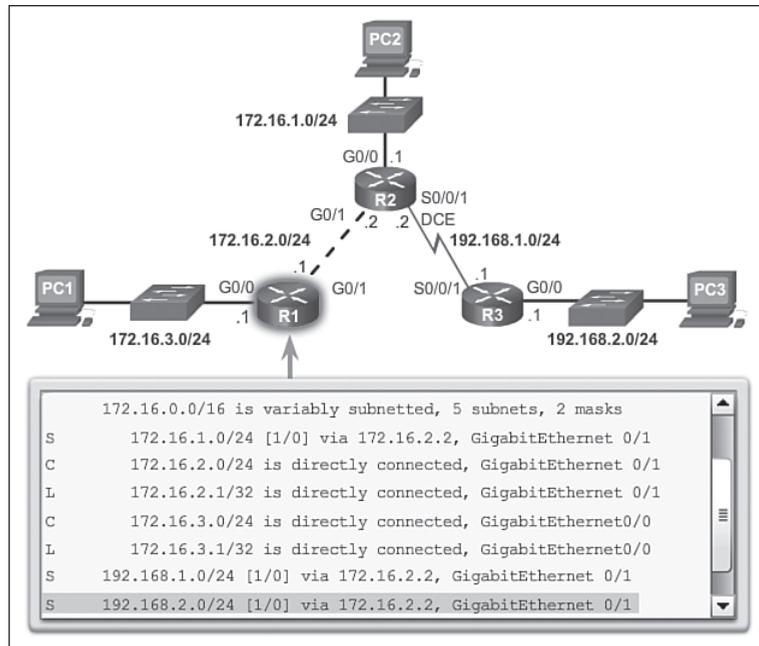


Figure 2-16 Verify the Routing Table of R1

Interactive
Graphic

Activity 2.2.1.5 Part 1: Configure Fully Specified Static Routes on R2

Go to the online course to use the Syntax Checker in the third graphic to configure a static route to the 172.16.3.0/24 network using the exit interface/next-hop pair: G0/1 172.16.2.1.

Interactive
Graphic

Activity 2.2.1.5 Part 2: Configure Fully Specified Static Routes on R3

Go to the online course to use the Syntax Checker in the fourth graphic to configure a static route to the 172.16.1.0/24, 172.16.2.0/24, and 172.16.3.0/24 networks using the exit interface S0/0/1 and next-hop address 192.168.1.2.

Verify a Static Route (2.2.1.6)

Along with `ping` and `tracert`, useful commands to verify static routes include:

- `show ip route`
- `show ip route static`
- `show ip route network`

**Interactive
Graphic****Activity 2.2.1.6 Part 1: Verify the Static Routing Settings on R2**

Go to the online course to use the Syntax Checker in the third graphic to display only the static routes in the routing table of R2.

**Interactive
Graphic****Activity 2.2.1.6 Part 2: Verify the Static Routing Settings on R3**

Go to the online course to use the Syntax Checker in the third graphic to display only the static routes in the routing table of R3.

Configure IPv4 Default Routes (2.2.2)

If a router does not have a route entry in its routing table for a destination network, a default route entry can be used to forward packets to another router. The use of a static default route is common with dynamic routing protocols and will be discussed in later chapters.

Default Static Route (2.2.2.1)

A default route is a static route that matches all packets. Rather than storing all routes to all networks in the routing table, a router can store a single default route to represent any network that is not in the routing table.

Routers commonly use default routes that are either configured locally or learned from another router, using a dynamic routing protocol. A default route is used when no other routes in the routing table match the destination IP address of the packet. In other words, if a more specific match does not exist, then the default route is used as the Gateway of Last Resort.

Default static routes are commonly used when connecting:

- An edge router to a service provider network
- A stub router (a router with only one upstream neighbor router)

As shown in Table 2-3, the command syntax for a default static route is similar to any other static route, except that the network address is **0.0.0.0** and the subnet mask is **0.0.0.0**. The basic command syntax of a default static route is:

- `ip route 0.0.0.0 0.0.0.0 { ip-address | exit-intf }`

Note

An IPv4 default static route is commonly referred to as a *quad-zero route*.

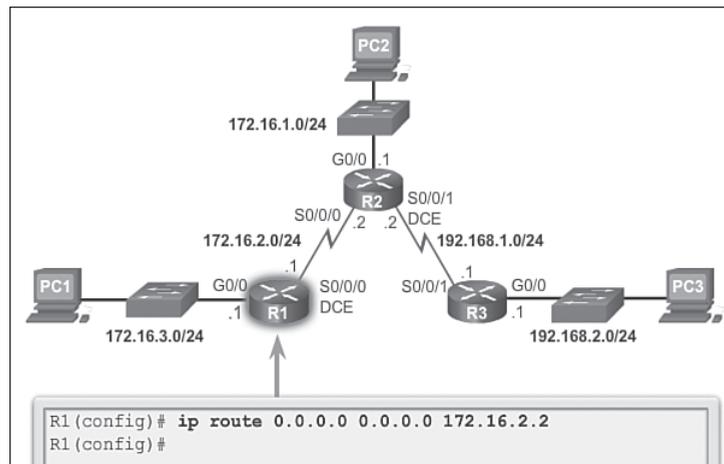
Table 2-3 Default Static Route Syntax

Parameter	Description
0.0.0.0	Matches any network address.
0.0.0.0	Matches any subnet mask.
<i>ip-address</i>	<ul style="list-style-type: none"> Commonly referred to as the next-hop router's IP address. Typically used when connecting to a broadcast media (i.e., Ethernet).
<i>exit-intf</i>	<ul style="list-style-type: none"> Use the outgoing interface to forward packets to the destination network. Also referred to as a directly attached static route.

Configure a Default Static Route (2.2.2.2)

R1 can be configured with three static routes to reach all of the remote networks in the example topology. However, R1 is a stub router because it is only connected to R2. Therefore, it would be more efficient to configure a default static route.

The example in Figure 2-18 configures a default static route on R1. With the configuration shown in the example, any packets not matching more specific route entries are forwarded to 172.16.2.2.

**Figure 2-18** Configuring a Default Static Route

Verify a Default Static Route (2.2.2.3)

In Figure 2-19, the `show ip route static` command output displays the contents of the routing table. Note the asterisk (*) next to the route with code 's'. As displayed in the

Codes table in Figure 2-19, the asterisk indicates that this static route is a candidate default route, which is why it is selected as the Gateway of Last Resort.

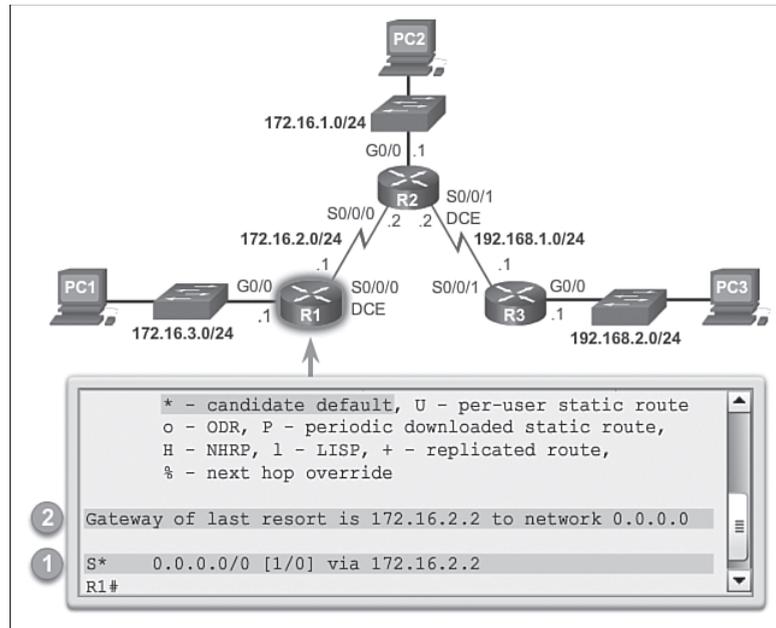


Figure 2-19 Verifying the Routing Table of R1

The key to this configuration is the /0 mask. Recall that the subnet mask in a routing table determines how many bits must match between the destination IP address of the packet and the route in the routing table. A binary 1 indicates that the bits must match. A binary 0 indicates that the bits do not have to match. A /0 mask in this route entry indicates that none of the bits are required to match. The default static route matches all packets for which a more specific match does not exist.

Packet Tracer
Activity

Packet Tracer Activity 2.2.2.4: Configuring IPv4 Static and Default Routes

In this activity, you will configure static and default routes. A static route is a route that is entered manually by the network administrator to create a route that is reliable and safe. There are four different static routes that are used in this activity: a recursive static route, a directly connected static route, a fully specified static route, and a default route.



Lab 2.2.2.5: Configuring IPv4 Static and Default Routes

In this lab, you will complete the following objectives:

- Part 1: Set Up the Topology and Initialize Devices
- Part 2: Configure Basic Device Settings and Verify Connectivity
- Part 3: Configure Static Routes
- Part 4: Configure and Verify a Default Route

Configure IPv6 Static Routes (2.2.3)

This section focuses on the configuration of IPv6 static routes. IPv6 static routes are similar to IPv4 static routes. IPv4 static routes are manually configured routes for reaching IPv4 networks, whereas IPv6 static routes are configured for reaching IPv6 networks.

The `ipv6 route` Command (2.2.3.1)

Static routes for IPv6 are configured using the `ipv6 route` global configuration command. Table 2-4 shows the simplified version of the command syntax:

```
Router(config)# ipv6 route ipv6-prefix/prefix-length { ipv6-address | exit-intf }
```

Table 2-4 IPv6 Command Syntax

Parameter	Description
<i>ipv6-prefix</i>	Destination network address of the remote network to be added to the routing table.
<i>prefix-length</i>	Prefix length of the remote network to be added to the routing table.
<i>ipv6-address</i>	<ul style="list-style-type: none"> ■ Commonly referred to as the next-hop router's IP address. ■ Typically used when connecting to a broadcast media (i.e., Ethernet).
<i>exit-intf</i>	<ul style="list-style-type: none"> ■ Use the outgoing interface to forward packets to the destination network. ■ Also referred to as a directly attached static route. ■ Typically used when connecting in a point-to-point configuration.

Most of the parameters are identical to the IPv4 version of the command. IPv6 static routes can also be implemented as:

- Standard IPv6 static route
- Default IPv6 static route

- Summary IPv6 static route
- Floating IPv6 static route

As with IPv4, these routes can be configured as recursive, directly connected, or fully specified.

The **ipv6 unicast-routing** global configuration command must be configured to enable the router to forward IPv6 packets. Figure 2-20 displays the enabling of IPv6 unicast routing.

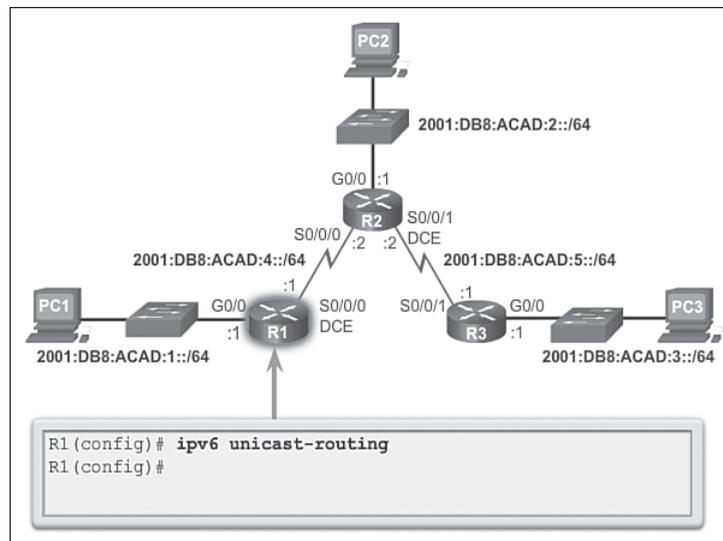


Figure 2-20 Enabling IPv6 Unicast Routing

Interactive Graphic

Activity 2.2.3.1 Part 1: Enabling IPv6 Unicast Routing on R2

Go to the online course to use the Syntax Checker in the third graphic to enable IPv6 unicast routing on R2.

Interactive Graphic

Activity 2.2.3.1 Part 2: Enabling IPv6 Unicast Routing on R3

Go to the online course to use the Syntax Checker in the third graphic to enable IPv6 unicast routing on R3.

Next-Hop Options (2.2.3.2)

The following example displays the routing tables of R1, R2, and R3. Each router has entries only for directly connected networks and their associated local addresses.

None of the routers have any knowledge of any networks beyond their directly connected interfaces.

```
R1# show ipv6 route
<output omitted>
C   2001:DB8:ACAD:1::/64 [0/0]
    via GigabitEthernet0/0, directly connected
L   2001:DB8:ACAD:1::1/128 [0/0]
    via GigabitEthernet0/0, receive
C   2001:DB8:ACAD:4::/64 [0/0]
    via Serial0/0/0, directly connected
L   2001:DB8:ACAD:4::1/128 [0/0]
    via Serial0/0/0, receive
L   FF00::/8 [0/0]
    via Null0, receive
R1#
```

```
R2# show ipv6 route
<output omitted>
C   2001:DB8:ACAD:2::/64 [0/0]
    via GigabitEthernet0/0, directly connected
L   2001:DB8:ACAD:2::1/128 [0/0]
    via GigabitEthernet0/0, receive
C   2001:DB8:ACAD:4::/64 [0/0]
    via Serial0/0/0, directly connected
L   2001:DB8:ACAD:4::2/128 [0/0]
    via Serial0/0/0, receive
C   2001:DB8:ACAD:5::/64 [0/0]
    via Serial0/0/1, directly connected
L   2001:DB8:ACAD:5::2/128 [0/0]
    via Serial0/0/1, receive
L   FF00::/8 [0/0]
    via Null0, receive
R2#
```

```
R3# show ipv6 route
<output omitted>
C   2001:DB8:ACAD:3::/64 [0/0]
    via GigabitEthernet0/0, directly connected
L   2001:DB8:ACAD:3::1/128 [0/0]
    via GigabitEthernet0/0, receive
C   2001:DB8:ACAD:5::/64 [0/0]
    via Serial0/0/1, directly connected
L   2001:DB8:ACAD:5::1/128 [0/0]
    via Serial0/0/1, receive
L   FF00::/8 [0/0]
    via Null0, receive
R3#
```

For example, R1 has no knowledge of networks:

- 2001:DB8:ACAD:2::/64: LAN on R2
- 2001:DB8:ACAD:5::/64: Serial network between R2 and R3
- 2001:DB8:ACAD:3::/64: LAN on R3

Figure 2-21 displays a successful ping from R1 to R2. And the subsequent output shows an unsuccessful ping to the R3 LAN. This is unsuccessful because R1 does not have an entry in its routing table for that network.

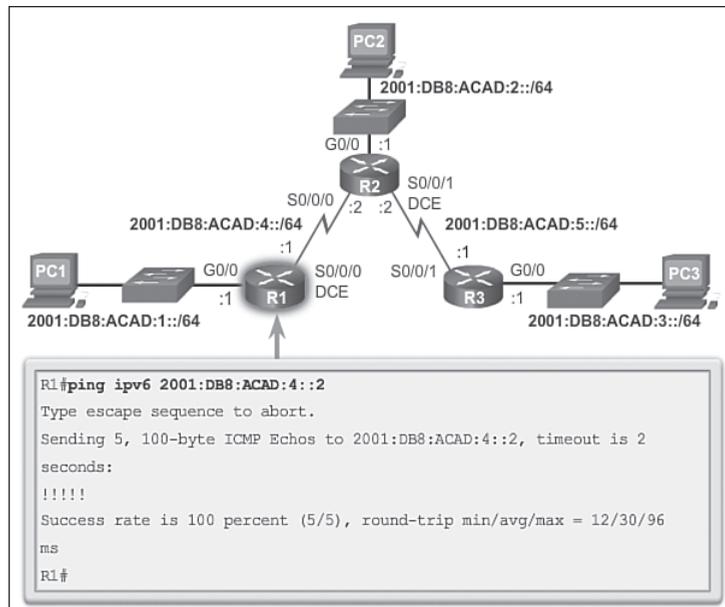


Figure 2-21 Verify Connectivity from R1 to R2

```

R1# ping ipv6 2001:DB8:ACAD:3::1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2001:DB8:ACAD:3::1, timeout is 2 seconds:

% No valid route for destination
Success rate is 0 percent (0/1)
R1#
  
```

The next hop can be identified by an IPv6 address, exit interface, or both. How the destination is specified creates one of three route types:

- **Next-hop static IPv6 route:** Only the next-hop IPv6 address is specified.
- **Directly connected static IPv6 route:** Only the router exit interface is specified.
- **Fully specified static IPv6 route:** The next-hop IPv6 address and exit interface are specified.

Configure a Next-Hop Static IPv6 Route (2.2.3.3)

In a next-hop static route, only the next-hop IPv6 address is specified. The output interface is derived from the next hop. For instance, in Figure 2-22, three next-hop static routes are configured on R1.

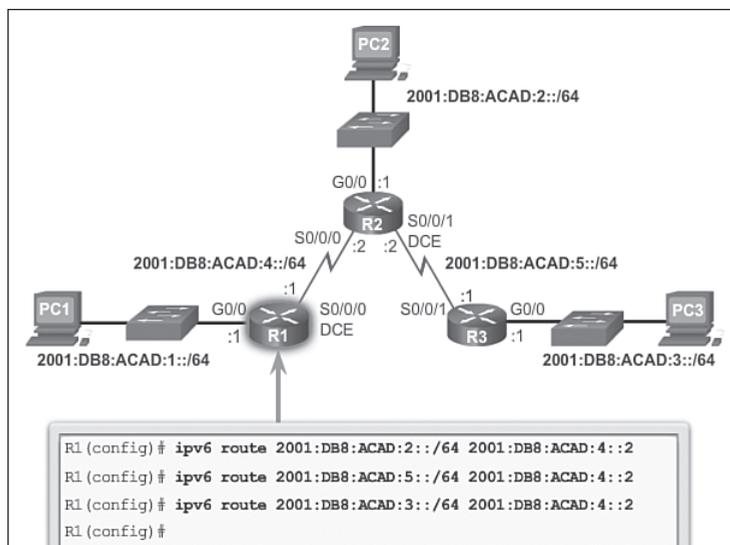


Figure 2-22 Configure Next-Hop Static IPv6 Routes

As with IPv4, before any packet is forwarded by the router, the routing table process must resolve the route to determine the exit interface to use to forward the packet. The route resolvability process will vary depending upon the type of forwarding mechanism being used by the router. CEF is the default behavior on most platforms running IOS 12.0 or later.

Figure 2-23 details the basic packet forwarding route resolvability process in the routing table for R1 without the use of CEF. When a packet is destined for the 2001:DB8:ACAD:3::/64 network, R1:

1. Looks for a match in the routing table and finds that it has to forward the packets to the next-hop IPv6 address 2001:DB8:ACAD:4::2. Every route that references

only a next-hop IPv6 address and does not reference an exit interface must have the next-hop IPv6 address resolved using another route in the routing table with an exit interface.

- R1 must now determine how to reach 2001:DB8:ACAD:4::2; therefore, it searches a second time looking for a match. In this case, the IPv6 address matches the route for the directly connected network 2001:DB8:ACAD:4::/64 with the exit interface Serial 0/0/0. This lookup tells the routing table process that this packet is forwarded out of that interface.

```

OE2 - OSPF ext 2, ON1 - OSPF NSSA ext 1,
ON2 - OSPF NSSA ext 2
C 2001:DB8:ACAD:1::/64 [0/0]
   via GigabitEthernet0/0, directly connected
L 2001:DB8:ACAD:1::1/128 [0/0]
   via GigabitEthernet0/0, receive
S 2001:DB8:ACAD:2::/64 [1/0]
   via 2001:DB8:ACAD:4::2
1 S 2001:DB8:ACAD:3::/64 [1/0]
   via 2001:DB8:ACAD:4::2
2 C 2001:DB8:ACAD:4::/64 [0/0]
   via Serial0/0/0, directly connected
L 2001:DB8:ACAD:4::1/128 [0/0]
   via Serial0/0/0, receive
S 2001:DB8:ACAD:5::/64 [1/0]
   via 2001:DB8:ACAD:4::2
L FF00::/8 [0/0]
   via Null0, receive
R1#

```

Figure 2-23 Verifying an IPv6 Next-Hop Lookup

Therefore, it actually takes two routing table lookup processes to forward any packet to the 2001:DB8:ACAD:3::/64 network. When the router has to perform multiple lookups in the routing table before forwarding a packet, it is performing a process known as a recursive lookup.

A recursive static IPv6 route is valid (that is, it is a candidate for insertion in the routing table) only when the specified next hop resolves, either directly or indirectly, to a valid exit interface.

**Interactive
Graphic**

Activity 2.2.3.3 Part 1: Configure Next-Hop Static IPv6 Routing on R2

Go to the online course to use the Syntax Checker in the third graphic to configure an IPv6 route to network 2001:DB8:ACAD:1::/64 using the next-hop address 2001:DB8:ACAD:4::1.

Interactive
Graphic**Activity 2.2.3.3 Part 2: Configure Next-Hop Static IPv6 Routing on R3**

Go to the online course to use the Syntax Checker in the fourth graphic to configure an IPv6 route to networks 2001:DB8:ACAD:1::/64, 2001:DB8:ACAD:2::/64, and 2001:DB8:ACAD:4::/64 using the next-hop address 2001:DB8:ACAD:5::2.

Configure a Directly Connected Static IPv6 Route (2.2.3.4)

When configuring a static route on point-to-point networks, an alternative to using the next-hop IPv6 address is to specify the exit interface. This is an alternative used in older IOS versions or whenever CEF is disabled, to avoid the recursive lookup problem.

For instance, in Figure 2-24, three directly connected static routes are configured on R1 using the exit interface.

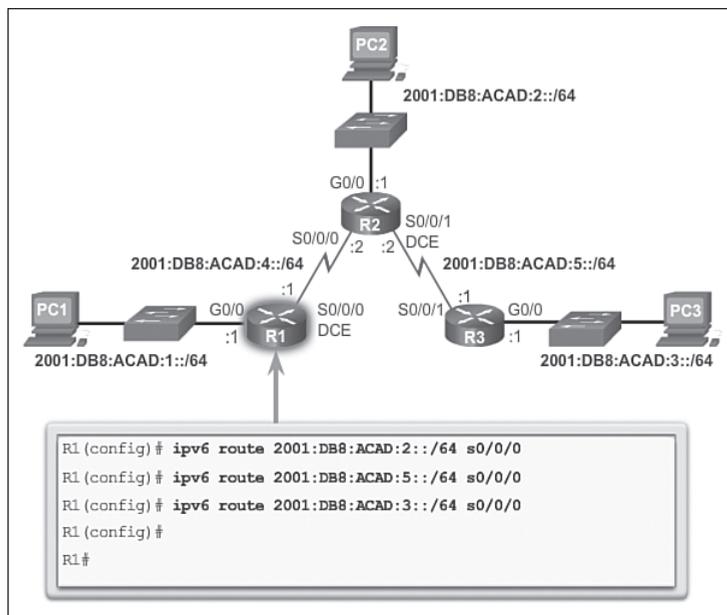


Figure 2-24 Configure Directly Connected Static IPv6 Routes on R1

The IPv6 routing table for R1 in the following output shows that when a packet is destined for the 2001:DB8:ACAD:3::/64 network, R1 looks for a match in the routing table and finds that it can forward the packet out of its Serial 0/0/0 interface. No other lookups are required.

```

R1# show ipv6 route
IPv6 Routing Table - default - 8 entries
Codes: C - Connected, L - Local, S - Static, U - Per-user Static route
       B - BGP, R - RIP, I1 - ISIS L1, I2 - ISIS L2
       IA - ISIS interarea, IS - ISIS summary, D - EIGRP, EX - EIGRP external
       ND - ND Default, NDP - ND Prefix, DCE - Destination, NDR - Redirect
       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:ACAD:1::/64 [0/0]
    via GigabitEthernet0/0, directly connected
L   2001:DB8:ACAD:1::1/128 [0/0]
    via GigabitEthernet0/0, receive
S   2001:DB8:ACAD:2::/64 [1/0]
    via Serial0/0/0, directly connected
S   2001:DB8:ACAD:3::/64 [1/0]
    via Serial0/0/0, directly connected
C   2001:DB8:ACAD:4::/64 [0/0]
    via Serial0/0/0, directly connected
L   2001:DB8:ACAD:4::1/128 [0/0]
    via Serial0/0/0, receive
S   2001:DB8:ACAD:5::/64 [1/0]
    via Serial0/0/0, directly connected
L   FF00::/8 [0/0]
    via Null0, receive
R1#

```

Notice how the routing table looks different for the route configured with an exit interface than the route configured with a recursive entry.

Configuring a directly connected static route with an exit interface allows the routing table to resolve the exit interface in a single search instead of two searches. Recall that with the use of the CEF forwarding mechanism, static routes with an exit interface are considered unnecessary. A single lookup is performed using a combination of the FIB and adjacency table stored in the data plane.

Interactive Graphic

Activity 2.2.3.4 Part 1: Configure Directly Connected Static IPv6 Routes on R2

Go to the online course to use the Syntax Checker in the third graphic to configure an IPv6 route to network 2001:DB8:ACAD:1::/64 using exit interface S0/0/0.

Interactive Graphic

Activity 2.2.3.4 Part 2: Configure Directly Connected Static IPv6 Routes on R3

Go to the online course to use the Syntax Checker in the fourth graphic to configure an IPv6 route to the 2001:DB8:ACAD:1::/64, 2001:DB8:ACAD:2::/64, and 2001:DB8:ACAD:4::/64 networks using exit interface S0/0/1.

Configure a Fully Specified Static IPv6 Route (2.2.3.5)

In a fully specified static route, both the output interface and the next-hop IPv6 address are specified. Similar to fully specified static routes used with IPv4, this would be used if CEF were not enabled on the router and the exit interface was on a multi-access network. With CEF, a static route using only a next-hop IPv6 address would be the preferred method even when the exit interface is a multi-access network.

Unlike IPv4, there is a situation in IPv6 when a fully specified static route must be used. If the IPv6 static route uses an IPv6 link-local address as the next-hop address, a fully specified static route including the exit interface must be used. Figure 2-25 shows an example of a fully qualified IPv6 static route using an IPv6 link-local address as the next-hop address.

The reason a fully specified static route must be used is because IPv6 link-local addresses are not contained in the IPv6 routing table. Link-local addresses are only unique on a given link or network. The next-hop link-local address may be a valid address on multiple networks connected to the router. Therefore, it is necessary that the exit interface be included.

In Figure 2-25, a fully specified static route is configured using R2's link-local address as the next-hop address. Notice that IOS requires that an exit interface be specified.

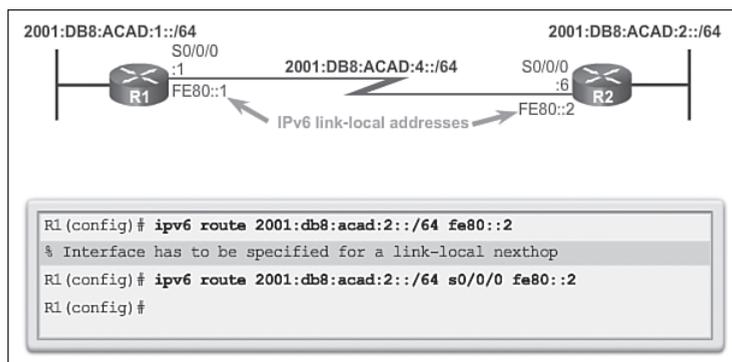


Figure 2-25 Configure a Fully Specified Static IPv6 Route on R1

The following output shows the IPv6 routing table entry for this route. Notice that both the next-hop link-local address and the exit interface are included.

```
R1# show ipv6 route static 2001:DB8:ACAD:2::/64
S 2001:DB8:ACAD:2::/64 (1/0)
  via FE80::2, Serial0/0/0
```

Interactive
Graphic**Activity 2.2.3.5: Configure a Fully Specified IPv6 Route on R2**

Go to the online course to use the Syntax Checker in the third graphic to configure a fully specified static IPv6 route to the R1 LAN using the R1 link-local address as the next-hop address.

Verify IPv6 Static Routes (2.2.3.6)

Along with **ping** and **tracert**, useful commands to verify static routes include:

- **show ipv6 route**
- **show ipv6 route static**
- **ipv6 route *network***

The following displays sample output of the **show ipv6 route static** command. The output reflects the use of static routes using next-hop global unicast addresses.

```
R1# show ipv6 route static
IPv6 Routing Table - default - 8 entries
Codes: C - Connected, L - Local, S - Static, U - Per-user Static route
        B - BGP, R - RIP, I1 - ISIS L1, I2 - ISIS L2
        IA - ISIS interarea, IS - ISIS summary, D - EIGRP, EX - EIGRP external
        ND - ND Default, NDp - ND Prefix, DCE - Destination, NDr - Redirect
        O - OSPF Intra, OI - OSPF Inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2
        ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2
S   2001:DB8:ACAD:2::/64 [1/0]
    via 2001:DB8:ACAD:4::2, Serial0/0/0
S   2001:DB8:ACAD:3::/64 [1/0]
    via 2001:DB8:ACAD:4::2, Serial0/0/0
S   2001:DB8:ACAD:5::/64 [1/0]
    via 2001:DB8:ACAD:4::2, Serial0/0/0
R1#
```

The following output displays sample output from the **show ip route 2001:DB8:ACAD:3::1** command:

```
R1# show ipv6 route 2001:0DB8:ACAD:3::1
Routing entry for 2001:DB8:ACAD:3::/64
  Known via "static", distance 1, metric 0
  Route count is 1/1, share count 0
  Routing paths:
    2001:DB8:ACAD:4::2, Serial0/0/0
    Last updated 00:19:11 ago
R1#
```

The following output verifies the **ipv6 route** configuration in the running configuration with the output filtered using the pipe and **section** parameter:

```
R1# show running-config | section ipv6 route
ipv6 route 2001:DB8:ACAD:2::/64 Serial0/0/0 2001:DB8:ACAD:4::2
ipv6 route 2001:DB8:ACAD:3::/64 Serial0/0/0 2001:DB8:ACAD:4::2
ipv6 route 2001:DB8:ACAD:5::/64 Serial0/0/0 2001:DB8:ACAD:4::2
R1#
```

Configure IPv6 Default Routes (2.2.4)

Similar to IPv4, an IPv6 default route entry can be used to forward packets to another router when there is not a specific IPv6 route in the IPv6 routing table.

Default Static IPv6 Route (2.2.4.1)

A default route is a static route that matches all packets. Instead of routers storing routes for all of the networks in the Internet, they can store a single default route to represent any network that is not in the routing table.

Routers commonly use default routes that are either configured locally or learned from another router, using a dynamic routing protocol. They are used when no other routes match the packet's destination IP address in the routing table. In other words, if a more specific match does not exist, then use the default route as the Gateway of Last Resort.

Default static routes are commonly used when connecting:

- A company's edge router to a service provider network.
- A router with only an upstream neighbor router. The router has no other neighbors and is, therefore, referred to as a *stub router*.

As shown in Table 2-5, the command syntax for a default static route is similar to any other static route, except that the ipv6-prefix/prefix-length is **::/0**, which matches all routes.

The basic command syntax of a default static route is:

- **ipv6 route ::/0 { ipv6-address | exit-intf }**

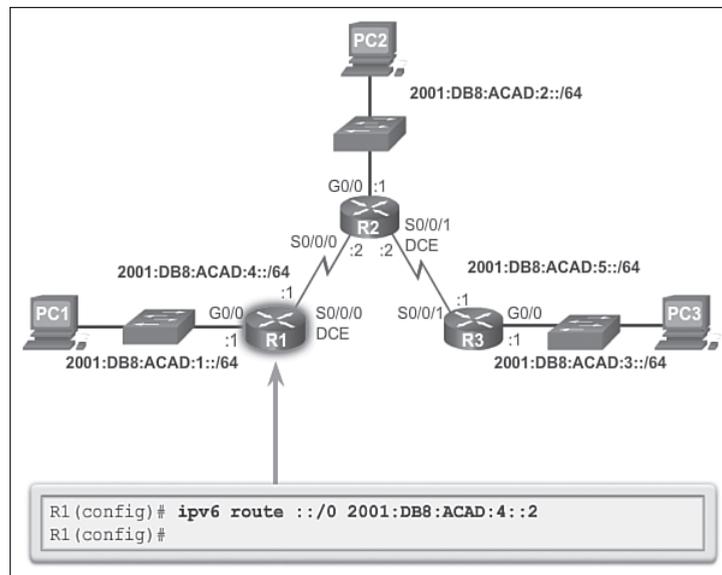
Table 2-5 Default Static IPv6 Route Syntax

Parameter	Description
<code>::/0</code>	Matches any IPv6 prefix regardless of prefix length.
<code>ipv6-address</code>	<ul style="list-style-type: none"> Commonly referred to as the next-hop router's IP address. Typically used when connecting to a broadcast media (i.e., Ethernet).
<code>exit-intf</code>	<ul style="list-style-type: none"> Use the outgoing interface to forward packets to the destination network. Also referred to as a directly attached static route.

Configure a Default Static IPv6 Route (2.2.4.2)

R1 can be configured with three static routes to reach all of the remote networks in our topology. However, R1 is a stub router because it is only connected to R2. Therefore, it would be more efficient to configure a default static IPv6 route.

The example in Figure 2-26 displays a configuration for a default static IPv6 route on R1.

**Figure 2-26** Default Static IPv6 Route Syntax

Verify a Default Static Route (2.2.4.3)

In the following output, the `show ipv6 route static` command output displays the contents of the routing table:

```
R1# show ipv6 route static
IPv6 Routing Table - default - 6 entries
Codes: C - Connected, L - Local, S - Static, U - Per-user Static route
       B - BGP, R - RIP, I1 - ISIS L1, I2 - ISIS L2
       IA - ISIS interarea, IS - ISIS summary, D - EIGRP, EX - EIGRP external
       ND - ND Default, NDP - ND Prefix, DCE - Destination, NDr - Redirect
       O - OSPF Intra, OI - OSPF Inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2
       ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2
S ::/0 [1/0]
  via 2001:DB8:ACAD:4::2, Serial0/0/0
R1#
```

Unlike IPv4, IPv6 does not explicitly state that the default IPv6 is the Gateway of Last Resort.

The key to this configuration is the `::/0` mask. Recall that the `ipv6` prefix-length in a routing table determines how many bits must match between the destination IP address of the packet and the route in the routing table. The `::/0` mask indicates that none of the bits are required to match. As long as a more specific match does not exist, the default static IPv6 route matches all packets.

The following output displays a successful ping to the R3 LAN interface:

```
R1# ping 2001:0DB8:ACAD:3::1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2001:DB8:ACAD:3::1, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 28/28/28 ms
R1#
```

Packet Tracer
Activity

Packet Tracer Activity 2.2.4.4: Configuring IPv6 Static and Default Routes

In this activity, you will configure IPv6 static and default routes. A static route is a route that is entered manually by the network administrator to create a route that is reliable and safe. There are four different static routes used in this activity: a recursive static route; a directly connected static route; a fully specified static route; and a default route.



Lab 2.2.4.5: Configuring IPv6 Static and Default Routes

In this lab, you will complete the following objectives:

- Part 1: Build the Network and Configure Basic Device Settings
 - Part 2: Configure IPv6 Static and Default Routes
-

Review of CIDR and VLSM (2.3)

Recall that VLSM (variable-length subnet mask) subnetting is similar to the traditional subnetting. The difference is that with VLSM, the network is first subnetted, and then subnets are subnetted again. This can occur in several iterations. CIDR (Classless Inter-Domain Routing) was introduced by IETF in 1993 to replace class network assignments. VLSM and CIDR helped make the allocation of the limited IPv4 address space more efficient.

Classful Addressing (2.3.1)

Although CIDR and classless addressing obsoleted classful addressing, an understanding of classful addressing is still important. Routing protocols such as RIP and EIGRP can both be configured to summarize on classful network boundaries. The IPv4 routing table is also structured in a classful hierarchy.

Classful Network Addressing (2.3.1.1)

Released in 1981, RFC 790 and RFC 791 describe how IPv4 network addresses were initially allocated based on a classification system. In the original specification of IPv4, the authors established the classes to provide three different sizes of networks for large, medium, and small organizations. As a result, class A, B, and C addresses were defined with a specific format for the high order bits. High order bits are the far left bits in a 32-bit address.

As shown in Table 2-6:

- **Class A addresses begin with 0:** Intended for large organizations; includes all addresses from 0.0.0.0 (00000000) to 127.255.255.255 (01111111). The 0.0.0.0 address is reserved for default routing and the 127.0.0.0 address is reserved for loopback testing.
- **Class B addresses begin with 10:** Intended for medium-to-large organizations; includes all addresses from 128.0.0.0 (10000000) to 191.255.255.255 (10111111).
- **Class C addresses begin with 110:** Intended for small-to-medium organizations; includes all addresses from 192.0.0.0 (11000000) to 223.255.255.255 (11011111).

The remaining addresses were reserved for multicasting and future uses:

- **Class D Multicast addresses begin with 1110:** Multicast addresses are used to identify a group of hosts that are part of a multicast group. This helps reduce the amount of packet processing that is done by hosts, particularly on broadcast media (i.e., Ethernet LANs). Routing protocols such as RIPv2, EIGRP, and OSPF use designated multicast addresses (RIP = 224.0.0.9, EIGRP = 224.0.0.10, OSPF = 224.0.0.5, and 224.0.0.6).
- **Class E Reserved IP addresses begin with 1111:** These addresses were reserved for experimental and future use.

Table 2-6 High Order Bits

Class	High Order Bits	Start	End
Class A	0xxxxxxx	0.0.0.0	127.255.255.255
Class B	10xxxxxx	128.0.0.0	191.255.255.255
Class C	110xxxxx	192.0.0.0	223.255.255.255
Class D (Multicast)	111xxxx	224.0.0.0	239.255.255.255
Class E (Reserved)	1111xxxx	240.0.0.0	255.255.255.255

Links:

“Internet Protocol,” <http://www.ietf.org/rfc/rfc791.txt>

“IPv4 Multicast Address Space Registry,” <http://www.iana.org/assignments/multicast-addresses>

Classful Subnet Masks (2.3.1.2)

As specified in RFC 790, each network class has a default subnet mask associated with it.

As shown in Figure 2-27, class A networks used the first octet to identify the network portion of the address. This is translated to a 255.0.0.0 classful subnet mask. Because only 7 bits were left in the first octet (remember, the first bit is always 0), this made 2 to the 7th power, or 128 networks. The actual number is 126 networks, because there are two reserved class A addresses (i.e., 0.0.0.0/8 and 127.0.0.0/8). With 24 bits in the host portion, each class A address had the potential for over 16 million individual host addresses.

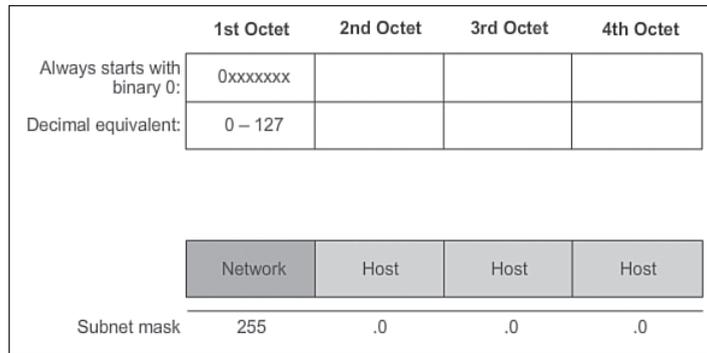


Figure 2-27 Class A Networks

As shown in Figure 2-28, class B networks used the first two octets to identify the network portion of the network address. With the first two bits already established as 1 and 0, 14 bits remained in the first two octets for assigning networks, which resulted in 16,384 class B network addresses. Because each class B network address contained 16 bits in the host portion, it controlled 65,534 addresses. (Recall that two addresses were reserved for the network and broadcast addresses.)

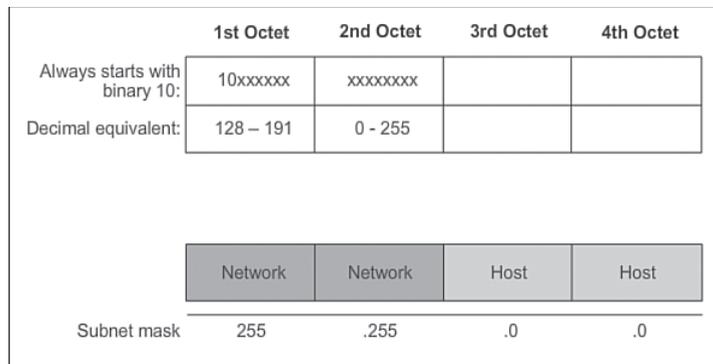


Figure 2-28 Class B Networks

As shown in Figure 2-29, class C networks used the first three octets to identify the network portion of the network address. With the first three bits established as 1 and 1 and 0, 21 bits remained for assigning networks for over 2 million class C networks. But, each class C network only had 8 bits in the host portion, or 254 possible host addresses.

An advantage of assigning specific default subnet masks to each class is that it made routing update messages smaller. Classful routing protocols do not include the subnet mask information in their updates. The receiving router applies the default mask based on the value of the first octet, which identifies the class.

	1st Octet	2nd Octet	3rd Octet	4th Octet
Always starts with binary 110:	110xxxxx	xxxxxxx	xxxxxxx	
Decimal equivalent:	192 – 223	0 - 255	0 - 255	
	Network	Network	Network	Host
Subnet mask	255	.255	.255	.0

Figure 2-29 Class C Networks

Classful Routing Protocol Example (2.3.1.3)

Using classful IP addresses meant that the subnet mask of a network address could be determined by the value of the first octet, or more accurately, the first three bits of the address. Routing protocols, such as RIPv1, only need to propagate the network address of known routes and do not need to include the subnet mask in the routing update. This is due to the router receiving the routing update determining the subnet mask simply by examining the value of the first octet in the network address, or by applying its ingress interface mask for subnetted routes. The subnet mask was directly related to the network address.

In Figure 2-30, R1 sends an update to R2. In the example, R1 knows that subnet 172.16.1.0 belongs to the same major classful network as the outgoing interface. Therefore, it sends an RIP update to R2 containing subnet 172.16.1.0. When R2 receives the update, it applies the receiving interface subnet mask (/24) to the update and adds 172.16.1.0 to the routing table.

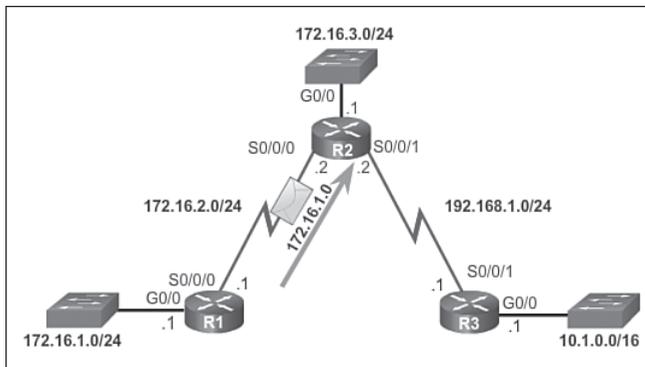


Figure 2-30 Classful Routing Update: R1 to R2

In Figure 2-31, R2 sends an update to R3. When sending updates to R3, R2 summarizes subnets 172.16.1.0/24, 172.16.2.0/24, and 172.16.3.0/24 into the major classful network 172.16.0.0. Because R3 does not have any subnets that belong to 172.16.0.0, it applies the classful mask for a class B network, which is /16.

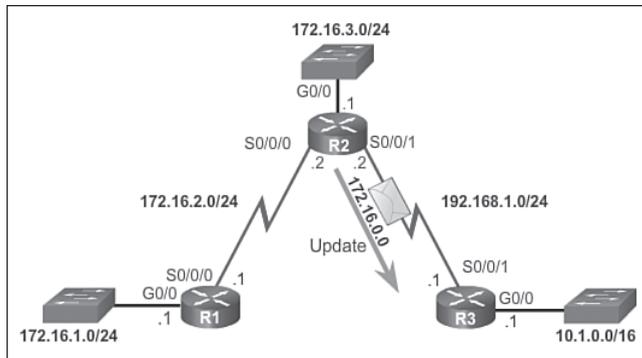


Figure 2-31 Classful Routing Update: R2 to R3

Classful Addressing Waste (2.3.1.4)

The classful addressing specified in RFCs 790 and 791 resulted in a tremendous waste of address space. In the early days of the Internet, organizations were assigned an entire classful network address from the A, B, or C class.

As illustrated in Figure 2-32:

- Class A had 50% of the total address space. However, only 126 organizations could be assigned a class A network address. Ridiculously, each of these organizations could provide addresses for up to 16 million hosts. Very large organizations were allocated entire class A address blocks. Some companies and governmental organizations still have class A addresses. For example, General Electric owns 3.0.0.0/8, Apple Computer owns 170.0.0/8, and the U.S. Postal Service owns 56.0.0.0/8.
- Class B had 25% of the total address space. Up to 16,384 organizations could be assigned a class B network address, and each of these networks could support up to 65,534 hosts. Only the largest organizations and governments could ever hope to use all 65,000 addresses. Like class A networks, many IP addresses in the class B address space were wasted.
- Class C had 12.5% of the total address space. Many more organizations were able to get class C networks, but were limited in the total number of hosts that they could connect. In fact, in many cases, class C addresses were often too small for most midsize organizations.
- Classes D and E are used for multicasting and reserved addresses.

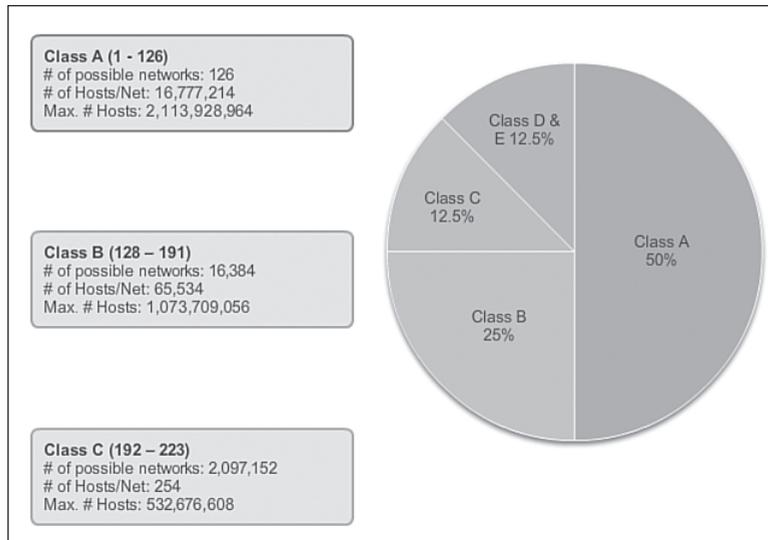


Figure 2-32 Classful IP Address Allocation = Inefficient

The overall result was that the classful addressing was a very wasteful addressing scheme. A better network addressing solution had to be developed. For this reason, Classless Inter-Domain Routing (CIDR) was introduced in 1993.

CIDR (2.3.2)

To help solve the limitation of the IPv4 address space, at least for a relatively short term, CIDR replaced classful addressing to make the distribution of IPv4 addresses more efficient.

Classless Inter-Domain Routing (2.3.2.1)

Just as the Internet was growing at an exponential rate in the early 1990s, so was the size of the routing tables that were maintained by Internet routers under classful IP addressing. For this reason, the IETF introduced CIDR in RFC 1517 in 1993.

CIDR replaced the classful network assignments, and address classes (A, B, and C) became obsolete. Using CIDR, the network address is no longer determined by the value of the first octet. Instead, the network portion of the address is determined by the subnet mask, also known as the network prefix, or prefix length (i.e., /8, /19, etc.).

ISPs are no longer limited to a /8, /16, or /24 subnet mask. They can now more efficiently allocate address space using any prefix length, starting with /8 and larger (i.e., /8, /9, /10, etc.). Figure 2-33 shows how blocks of IP addresses can be assigned to a network based on the requirements of the customer, ranging from a few hosts to hundreds or thousands of hosts.

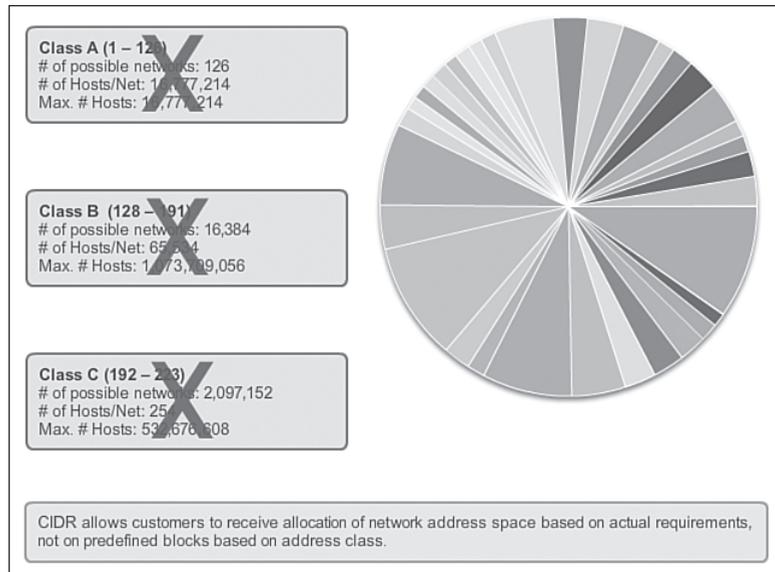


Figure 2-33 CIDR = Efficient

CIDR also reduces the size of routing tables and manages the IPv4 address space more efficiently using:

- **Route summarization:** Also known as prefix aggregation, routes are summarized into a single route to help reduce the size of routing tables. For instance, one summary static route can replace several specific static route statements.
- **Supernetting:** Occurs when the route summarization mask is a smaller value than the default traditional classful mask.

Note

A supernet is always a route summary, but a route summary is not always a supernet.

Classless Inter-Domain Routing (2.3.2.2)

In Figure 2-34, notice that ISP1 has four customers, and that each customer has a variable amount of IP address space. The address space of the four customers can be summarized into one advertisement to ISP2. The 192.168.0.0/20 summarized or aggregated route includes all the networks belonging to Customers A, B, C, and D. This type of route is known as a supernet route. A supernet summarizes multiple network addresses with a mask that is smaller than the classful mask.

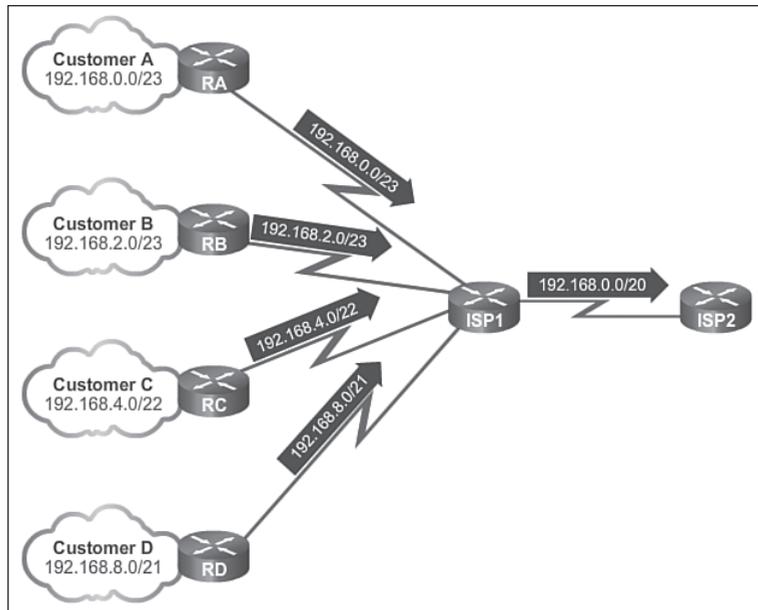


Figure 2-34 Summarizing Supernet Routes

Determining the summary route and subnet mask for a group of networks can be done in the following three steps:



- Step 1.** List the networks in binary format.
- Step 2.** Count the number of far left matching bits. This identifies the prefix length or subnet mask for the summarized route.
- Step 3.** Copy the matching bits and then add zero bits to the rest of the address to determine the summarized network address.

The summarized network address and subnet mask can now be used as the summary route for this group of networks.

Summary routes can be configured by both static routes and classless routing protocols.

Note

If a routing table contains both a summarized route and a more specific route, a route with a longer subnet mask (prefix length), the more specific route is always preferred.

Static Routing CIDR Example (2.3.2.3)

Creating smaller routing tables makes the routing table lookup process more efficient, because there are fewer routes to search. If one static route can be used instead of multiple static routes, the size of the routing table is reduced. In many cases, a single static route can be used to represent dozens, hundreds, or even thousands of routes.

Summary CIDR routes can be configured using static routes. This helps to reduce the size of routing tables.

In Figure 2-35, R1 has been configured to reach the identified networks in the topology. Although acceptable, it would be more efficient to configure a summary static route.

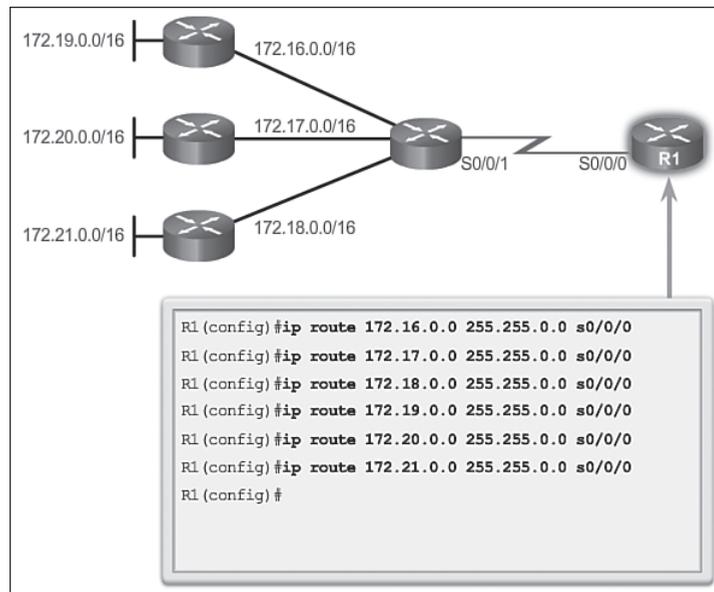


Figure 2-35 Six Static Routes

Figure 2-36 provides a solution using CIDR summarization. The six static route entries could be reduced to a 172.16.0.0/13 entry. The example removes the six static route entries and replaces them with a summary static route.

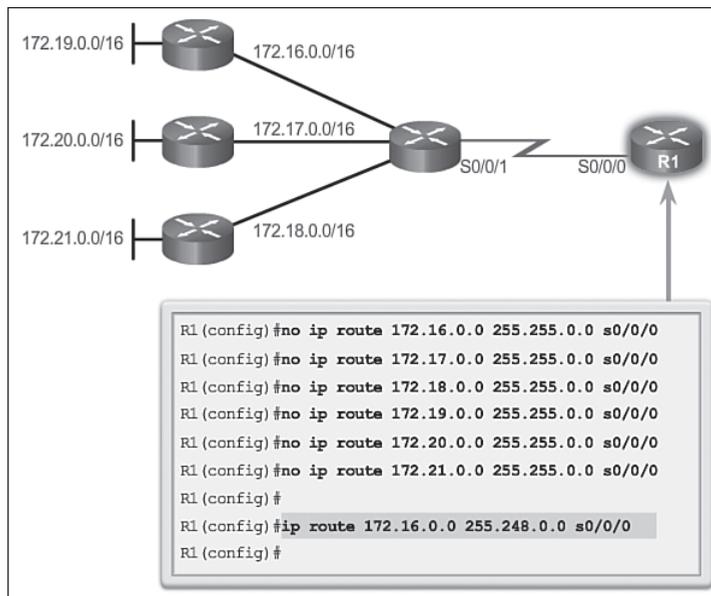


Figure 2-36 One Summary Static Route

Classless Routing Protocol Example (2.3.2.4)

Classful routing protocols cannot send supernet routes. This is because the receiving router automatically applies the default classful subnet mask to the network address in the routing update. If the topology in Figure 2-37 contained a classful routing protocol, then R3 would only install 172.16.0.0/16 in the routing table.

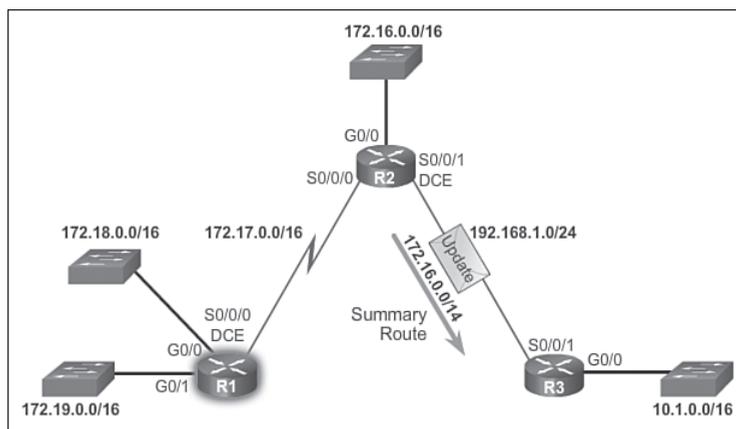


Figure 2-37 Classless Routing Update

Propagating VLSM and supernet routes requires a classless routing protocol such as RIPv2, OSPF, or EIGRP. Classless routing protocols advertise network addresses with their associated subnet masks. With a classless routing protocol, R2 can summarize networks 172.16.0.0/16, 172.17.0.0/16, 172.18.0.0/16, and 172.19.0.0/16, and advertise a supernet summary static route 172.16.0.0/14 to R3. R3 then installs the supernet route 172.16.0.0/14 in its routing table.

Note

When a supernet route is in a routing table, for example, as a static route, a classful routing protocol does not include that route in its updates.

VLSM (2.3.3)

Along with CIDR, VLSM helped more efficiently allocate the IPv4 address space. VLSM permits network administrators to subnet one or more specific subnets, allowing for different sized subnets. With VLSM, network administrators are no longer required to create a one-size-fits-all subnet.

Fixed-Length Subnet Masking (2.3.3.1)

With fixed-length subnet masking (FLSM), the same number of addresses is allocated for each subnet. If all the subnets have the same requirements for the number of hosts, these fixed-size address blocks would be sufficient. However, most often that is not the case.

Note

FLSM is also referred to as traditional subnetting.

The topology shown in Figure 2-38 requires that network address 192.168.20.0/24 be subnetted into seven subnets: one subnet for each of the four LANs (Buildings A to D), and one for each of the three WAN connections between routers.

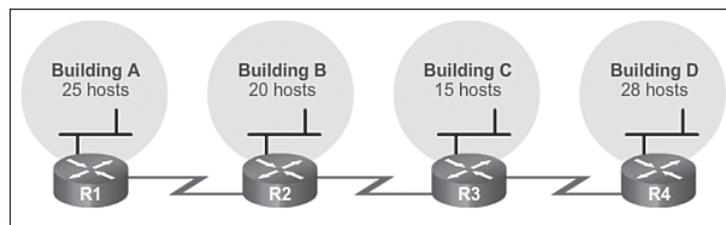


Figure 2-38 Network Topology: Basic Subnets

Figure 2-39 highlights how traditional subnetting can borrow 3 bits from the Host portion in the last octet to meet the subnet requirement of seven subnets. For example, under the Host portion, the Subnet portion highlights how borrowing 3 bits creates 8 subnets, while the Host portion highlights 5 host bits providing 30 usable host IP addresses per subnet. This scheme creates the needed subnets and meets the host requirement of the largest LAN.

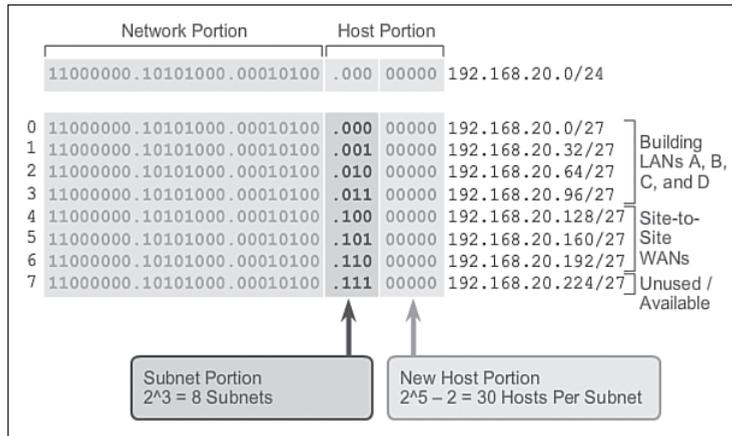


Figure 2-39 Basic Subnet Scheme

Although this traditional subnetting meets the needs of the largest LAN and divides the address space into an adequate number of subnets, it results in significant waste of unused addresses.

For example, only two addresses are needed in each subnet for the three WAN links. Because each subnet has 30 usable addresses, there are 28 unused addresses in each of these subnets. As shown in Figure 2-40, this results in 84 unused addresses (28×3). Further, this limits future growth by reducing the total number of subnets available. This inefficient use of addresses is characteristic of traditional subnetting of classful networks.

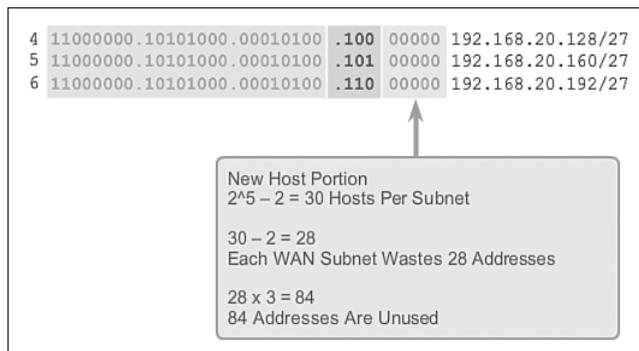


Figure 2-40 Unused Addresses on WAN Subnets

Applying a traditional subnetting scheme to this scenario is not very efficient and is wasteful. In fact, this example is a good model for showing how subnetting a subnet can be used to maximize address utilization. Subnetting a subnet, or using variable-length subnet masking (VLSM), was designed to avoid wasting addresses.

Variable-Length Subnet Masking (2.3.3.2)

In traditional subnetting the same subnet mask is applied for all the subnets. This means that each subnet has the same number of available host addresses.

As illustrated in Figure 2-41, traditional subnetting creates subnets of equal size. Each subnet in a traditional scheme uses the same subnet mask.

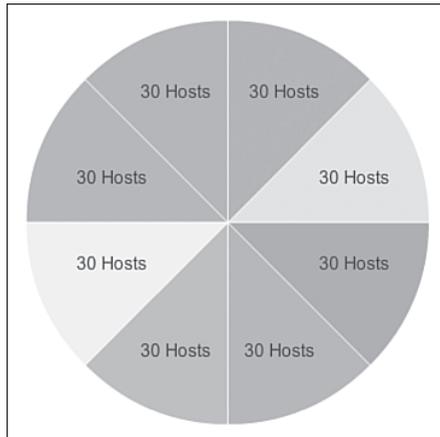


Figure 2-41 Traditional Subnetting Creates Equal Sized Subnets

With VLSM the subnet mask length varies depending on how many bits have been borrowed for a particular subnet, thus the “variable” part of variable-length subnet mask. As shown in Figure 2-42, VLSM allows a network space to be divided into unequal parts.

VLSM subnetting is similar to traditional subnetting in that bits are borrowed to create subnets. The formulas to calculate the number of hosts per subnet and the number of subnets created still apply. The difference is that subnetting is not a single-pass activity. With VLSM, the network is first subnetted, and then the subnets are subnetted again. This process can be repeated multiple times to create subnets of various sizes.

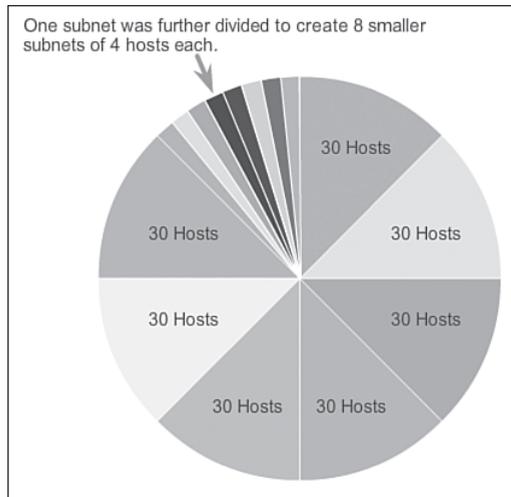


Figure 2-42 Subnets of Varying Sizes

VLSM in Action (2.3.3.3)

VLSM allows the use of different masks for each subnet. After a network address is subnetted, those subnets can be further subnetted. VLSM is simply subnetting a subnet. VLSM can be thought of as sub-subnetting.

Figure 2-43 shows the network 10.0.0.0/8 that has been subnetted using the subnet mask of /16, which makes 256 subnets; that is, 10.0.0.0/16, 10.1.0.0/16, 10.2.0.0/16, and so forth through 10.255.0.0/16. Four of these /16 subnets are displayed in Figure 2-43. Any of these /16 subnets can be subnetted further.

Figure 2-43 shows four subnets, subnetted a second time using three different subnet masks:

- The 10.1.0.0/16 subnet is subnetted again with the /24 mask.
- The 10.2.0.0/16 subnet is subnetted again with the /24 mask.
- The 10.3.0.0/16 subnet is subnetted again with the /28 mask
- The 10.4.0.0/16 subnet is subnetted again with the /20 mask.

Individual host addresses are assigned from the addresses of sub-subnets. For example, Figure 2-43 shows the 10.1.0.0/16 subnet divided into /24 subnets. The 10.1.4.10 address would now be a member of the more specific subnet 10.1.4.0/24.

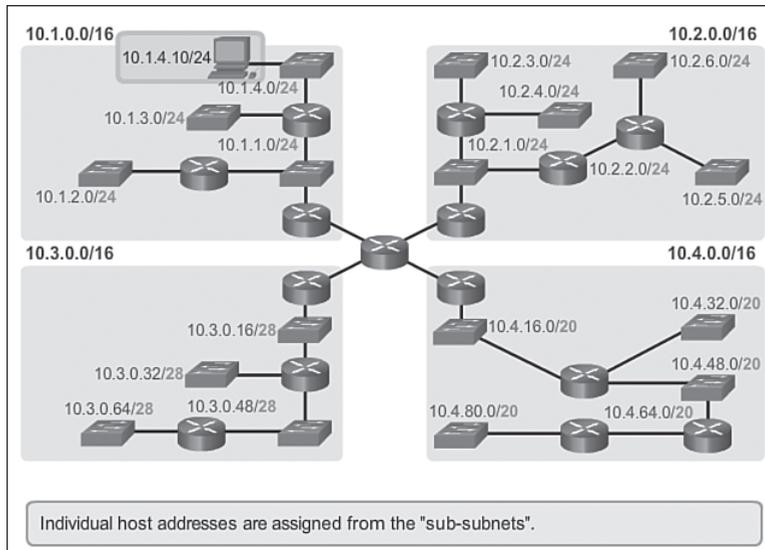


Figure 2-43 VLSM Subnets

Subnetting Subnets (2.3.3.4)

Another way to view the VLSM subnets is to list each subnet and its sub-subnets.

In Figure 2-44, the 10.0.0.0/8 network is the starting address space and is subnetted with a /16 mask. Borrowing 8 bits (going from /8 to /16) creates 256 subnets that range from 10.0.0.0/16 to 10.255.0.0/16.

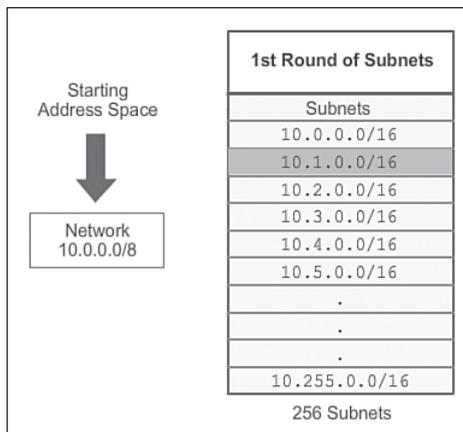


Figure 2-44 Subnetting 10.0.0.0/8 to 10.0.0.0/16

In Figure 2-45, the 10.1.0.0/16 subnet is further subnetted by borrowing 8 more bits. This creates 256 subnets with a /24 mask. This mask allows 254 host addresses per subnet. The subnets ranging from 10.1.0.0/24 to 10.1.255.0/24 are subnets of the subnet 10.1.0.0/16.

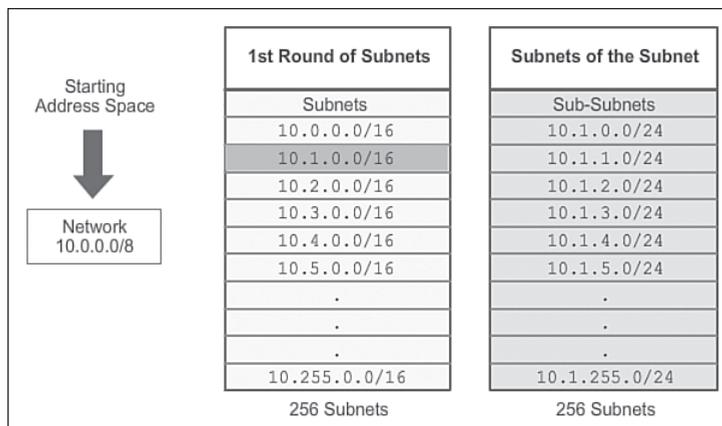


Figure 2-45 Subnetting 10.1.0.0/16 to 10.1.0.0/24

In Figure 2-46, the 10.2.0.0/16 subnet is also further subnetted with a /24 mask allowing 254 host addresses per subnet. The subnets ranging from 10.2.0.0/24 to 10.2.255.0/24 are subnets of the subnet 10.2.0.0/16.

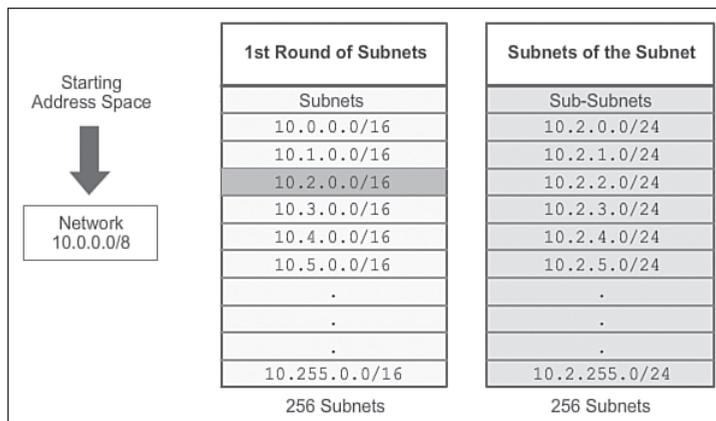


Figure 2-46 Subnetting 10.2.0.0/16 to 10.2.0.0/24

In Figure 2-47, the 10.3.0.0/16 subnet is further subnetted with a /28 mask, thus creating 4,096 subnets and allowing 14 host addresses per subnet. The subnets ranging from 10.3.0.0/28 to 10.3.255.240/28 are subnets of the subnet 10.3.0.0/16.

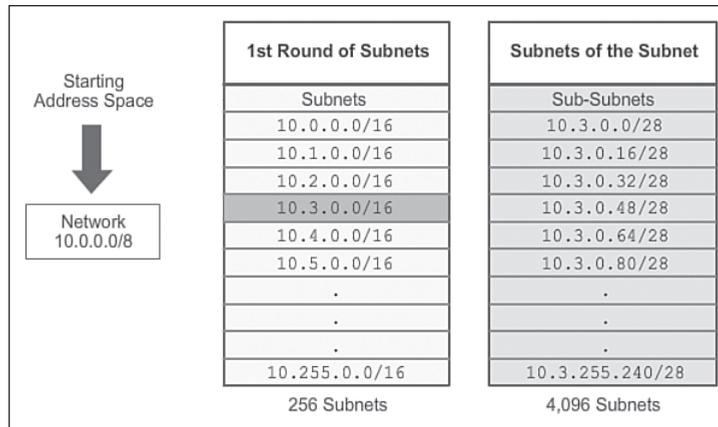


Figure 2-47 Subnetting 10.3.0.0/16 to 10.3.0.0/28

In Figure 2-48, the 10.4.0.0/16 subnet is further subnetted with a /20 mask, thus creating 16 subnets and allowing 4,094 host addresses per subnet. The subnets ranging from 10.4.0.0/20 to 10.4.240.0/20 are subnets of the subnet 10.4.0.0/16. These /20 subnets are big enough to subnet even further, allowing more networks.

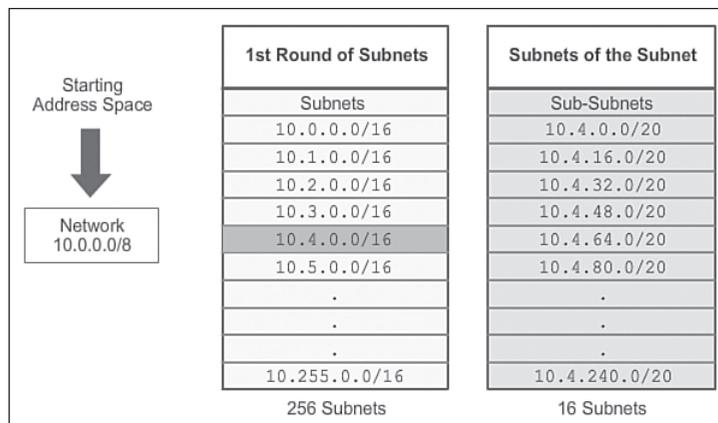


Figure 2-48 Subnetting 10.4.0.0/16 to 10.4.0.0/20

VLSM Example (2.3.3.5)

Careful consideration must be given to the design of a network addressing scheme. For example, the sample topology in Figure 2-49 requires seven subnets.

Using traditional subnetting, the first seven address blocks are allocated for LANs and WANs, as shown in Figure 2-50. This scheme results in 8 subnets with 30 usable addresses each (/27). While this scheme works for the LAN segments, there are many wasted addresses in the WAN segments.

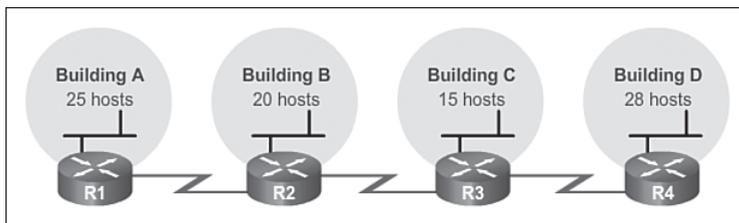


Figure 2-49 Basic Topology

	/27 Network	Hosts
Building A	.0	.1 - .30
Building B	.32	.33 - .62
Building C	.64	.65 - .94
Building D	.96	.97 - .126
WAN R1 – R2	.128	.129 - .158
WAN R2 – R3	.160	.161 - .190
WAN R3 – R4	.192	.193 - .222
Unused	.224	.225 - .254

Figure 2-50 Subnetting 192.168.20.0/24 to 192.168.20.0/27

If an addressing scheme is designed for a new network, the address blocks can be assigned in a way that minimizes waste and keeps unused blocks of addresses contiguous. It can be more difficult to do this when adding to an existing network.

As shown in Figure 2-51, to use the address space more efficiently, /30 subnets are created for WAN links. To keep the unused blocks of addresses together, the last /27 subnet is further subnetted to create the /30 subnets. The first 3 subnets were assigned to WAN links, creating subnets 192.168.20.224/30, 192.168.20.228/30, and 192.168.20.232/30. Designing the addressing scheme in this way leaves three unused /27 subnets and five unused /30 subnets.

	/27 Network	Hosts
Building A	.0	.1 - .30
Building B	.32	.33 - .62
Building C	.64	.65 - .94
Building D	.96	.97 - .126
Unused	.128	.129 - .158
Unused	.160	.161 - .190
Unused	.192	.193 - .222
	.224	.225 - .254

	/30 Network	Hosts
WAN R1–R2	.224	.225 - .226
WAN R2–R3	.228	.229 - .230
WAN R3–R4	.232	.233 - .234

Figure 2-51 Subnetting 192.168.20.224/27 to 192.168.20.224/30

The next four CLI outputs display sample configurations on all four routers to implement the VLSM addressing scheme.

Configuring VLSM on R1:

```
R1(config)# interface gigabitethernet 0/0
R1(config-if)# ip address 192.168.20.1 255.255.255.224
R1(config-if)# exit
R1(config)# interface serial 0/0/0
R1(config-if)# ip address 192.168.20.225 255.255.255.252
R1(config-if)# end
R1#
```

Configuring VLSM on R2:

```
R2(config)# interface gigabitethernet 0/0
R2(config-if)# ip address 192.168.20.33 255.255.255.224
R2(config-if)# exit
R2(config)# interface serial 0/0/0
R2(config-if)# ip address 192.168.20.226 255.255.255.252
R2(config-if)# exit
R2(config)# interface serial 0/0/1
R2(config-if)# ip address 192.168.20.229 255.255.255.252
R2(config-if)# end
R2#
```

Configuring VLSM on R3:

```
R3(config)# interface gigabitethernet 0/0
R3(config-if)# ip address 192.168.20.65 255.255.255.224
R3(config-if)# exit
R3(config)# interface serial 0/0/0
R3(config-if)# ip address 192.168.20.230 255.255.255.252
R3(config-if)# exit
R3(config)# interface serial 0/0/1
R3(config-if)# ip address 192.168.20.233 255.255.255.252
R3(config-if)# end
R3#
```

Configuring VLSM on R4:

```
R4(config)# interface gigabitethernet 0/0
R4(config-if)# ip address 192.168.20.97 255.255.255.224
R4(config-if)# exit
R4(config)# interface serial 0/0/0
R4(config-if)# ip address 192.168.20.234 255.255.255.252
R4(config-if)# end
R4#
```

Packet Tracer
Activity**Packet Tracer Activity 2.3.3.6: Designing and Implementing a VLSM Addressing Scheme**

In this activity, you are given a network address to develop a VLSM addressing scheme for the network shown in the included topology.

**Lab 2.3.3.7: Designing and Implementing Addressing with VLSM**

In this lab, you will complete the following objectives:

- Part 1: Examine the Network Requirements
 - Part 2: Design the VLSM Address Scheme
 - Part 3: Cable and Configure the IPv4 Network
-

Configure Summary and Floating Static Routes (2.4)

Summary static routes can be used to help minimize the number of static routes in the routing table. Using summary static routes can also make management of a large number of static routes easier and less prone to errors. Floating static routes can be used as a backup route for another static route or a dynamic routing protocol.

Configure IPv4 Summary Routes (2.4.1)

A single IPv4 static *summary route* can be used to replace multiple static routes when those routes can be summarized with a common prefix length. The configuration of a summary static route is similar to the configuration of other IPv4 static routes.

Route Summarization (2.4.1.1)

Route summarization, also known as route aggregation, is the process of advertising a contiguous set of addresses as a single address with a less-specific, shorter subnet mask. CIDR is a form of route summarization and is synonymous with the term supernetting.

CIDR ignores the limitation of classful boundaries, and allows summarization with masks that are smaller than that of the default classful mask. This type of summarization helps reduce the number of entries in routing updates and lowers the number of entries in local routing tables. It also helps reduce bandwidth utilization for routing updates and results in faster routing table lookups.

In Figure 2-52, R1 requires a summary static route to reach networks in the range of 172.20.0.0/16 to 172.23.0.0/16.

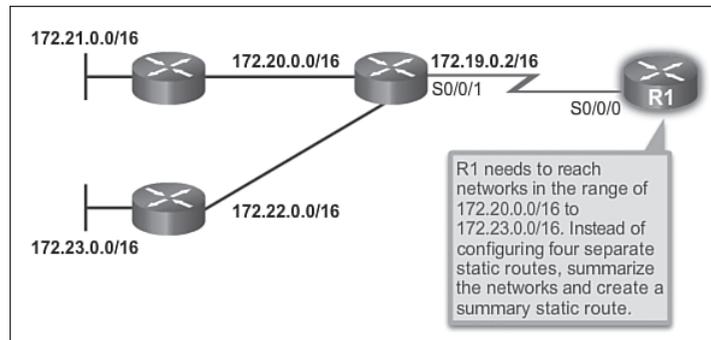


Figure 2-52 Basic Topology

Calculate a Summary Route (2.4.1.2)

Summarizing networks into a single address and mask can be done in three steps, as shown in Figure 2-53:



- Step 1.** List the networks in binary format. Figure 2-53 lists networks 172.20.0.0/16 to 172.23.0.0/16 in binary format.
- Step 2.** Count the number of far left matching bits to determine the mask for the summary route. Figure 2-53 highlights the 14 far left matching bits. This is the prefix, or subnet mask, for the summarized route: /14 or 255.252.0.0.
- Step 3.** Copy the matching bits and then add zero bits to determine the summarized network address. Figure 2-53 shows that the matching bits with zeros at the end results in the network address 172.20.0.0. The four networks—172.20.0.0/16, 172.21.0.0/16, 172.22.0.0/16, and 172.23.0.0/16—can be summarized into the single network address and prefix 172.20.0.0/14.

Figure 2-54 displays R1 configured with a summary static route to reach networks 172.20.0.0/16 to 172.23.0.0/16.

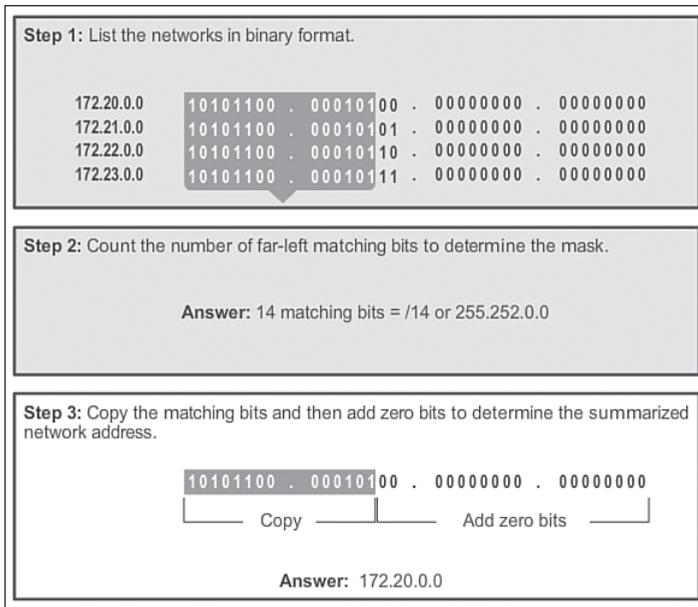


Figure 2-53 Calculating a Route Summary

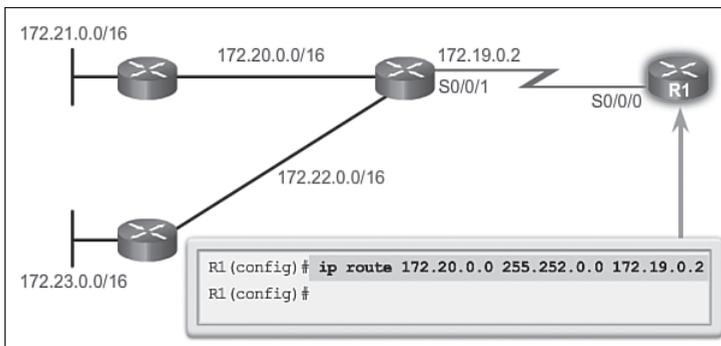


Figure 2-54 One Summary Static Route

Summary Static Route Example (2.4.1.3)

Multiple static routes can be summarized into a single static route if:

- The destination networks are contiguous and can be summarized into a single network address.
- The multiple static routes all use the same exit interface or next-hop IP address.

Consider the example in Figure 2-55. All routers have connectivity using static routes.

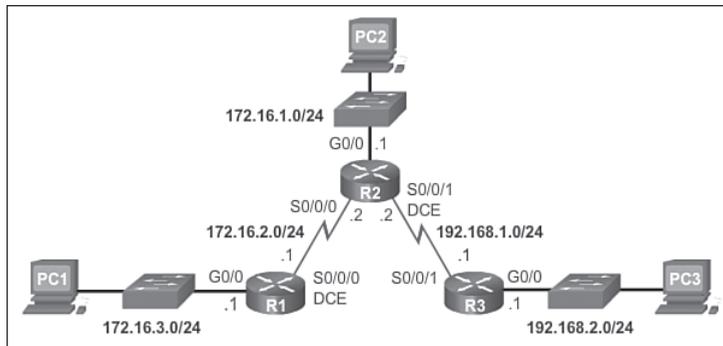


Figure 2-55 Basic Topology

The following output displays the static routing table entries for R3. Notice that it has three static routes that can be summarized because they share the same two first octets.

```
R3# show ip route static | begin Gateway
Gateway of last resort is not set
    172.16.0.0/24 is subnetted, 3 subnets
S       172.16.1.0 is directly connected, Serial0/0/1
S       172.16.2.0 is directly connected, Serial0/0/1
S       172.16.3.0 is directly connected, Serial0/0/1
R3#
```

Figure 2-56 displays the steps to summarize those three networks:



- Step 1.** Write out the networks to summarize in binary.
- Step 2.** To find the subnet mask for summarization, start with the far left bit, work to the right, finding all the bits that match consecutively until a column of bits that do not match is found, identifying the summary boundary.
- Step 3.** Count the number of far left matching bits; in our example, it is 22. This number identifies the subnet mask for the summarized route as /22 or 255.255.252.0.
- Step 4.** To find the network address for summarization, copy the matching 22 bits and add all 0 bits to the end to make 32 bits.

After the summary route is identified, replace the existing routes with the one summary route.

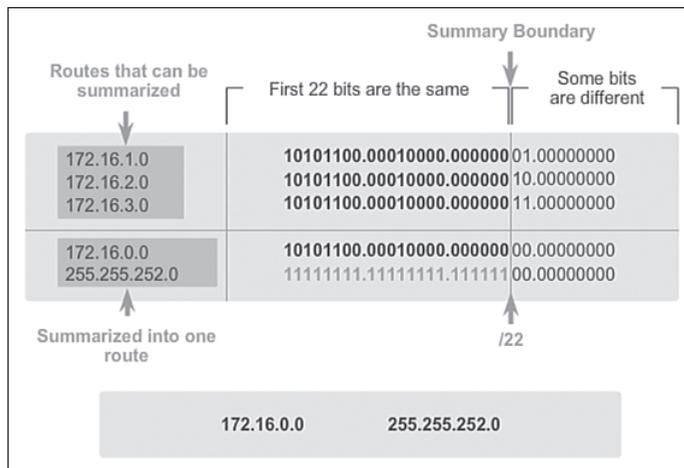


Figure 2-56 Summarize the Networks

The following output shows how the three existing routes are removed and then the new summary static route is configured:

```
R3(config)# no ip route 172.16.1.0 255.255.255.0 s0/0/1
R3(config)# no ip route 172.16.2.0 255.255.255.0 s0/0/1
R3(config)# no ip route 172.16.3.0 255.255.255.0 s0/0/1
R3(config)# ip route 172.16.0.0 255.255.252.0 s0/0/1
R3(config)#
```

The following output confirms that the summary static route is in the routing table of R3:

```
R3# show ip route static | begin Gateway
Gateway of last resort is not set
    172.16.0.0/22 is subnetted, 1 subnets
S       172.16.0.0 is directly connected, Serial0/0/1
R3#
```

Interactive
Graphic

Activity 2.4.1.4: Determine the Summary Network Address and Prefix

Go to the online course to perform this practice activity.

Packet Tracer
Activity**Packet Tracer Activity 2.4.1.5: Configuring IPv4 Route Summarization – Scenario 1**

In this activity, you will calculate and configure summary routes. Route summarization, also known as route aggregation, is the process of advertising a contiguous set of addresses as a single address.

Packet Tracer
Activity**Packet Tracer Activity 2.4.1.6: Configuring IPv4 Route Summarization – Scenario 2**

In this activity, you will calculate and configure summary routes. Route summarization, also known as route aggregation, is the process of advertising a contiguous set of addresses as a single address. After calculating summary routes for each LAN, you must summarize a route which includes all networks in the topology in order for the ISP to reach each LAN.

Configure IPv6 Summary Routes (2.4.1)

Similar to IPv4, a single IPv6 static summary route can be used to replace multiple IPv6 static routes with a common prefix length. The calculation and configuration of an IPv6 summary static route is similar to the configuration of an IPv4 static summary route.

Summarize IPv6 Network Addresses (2.4.2.1)

Aside from the fact that IPv6 addresses are 128 bits long and written in hexadecimal, summarizing IPv6 addresses is actually similar to the summarization of IPv4 addresses. It just requires a few extra steps due to the abbreviated IPv6 addresses and hex conversion.

Multiple static IPv6 routes can be summarized into a single static IPv6 route if:

- The destination networks are contiguous and can be summarized into a single network address.
- The multiple static routes all use the same exit interface or next-hop IPv6 address.

Refer to the network in Figure 2-57. R1 currently has four static IPv6 routes to reach networks 2001:DB8:ACAD:1::/64 to 2001:DB8:ACAD:4::/64.

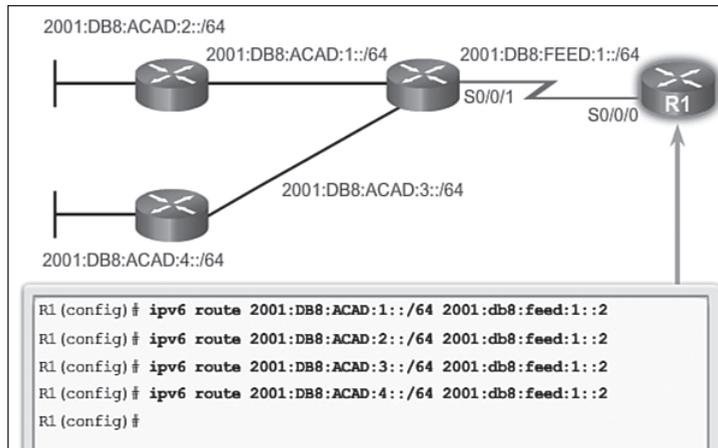


Figure 2-57 Basic Topology

The following output displays the IPv6 static routes installed in the IPv6 routing table:

```

R1# show ipv6 route static
IPv6 Routing Table - default - 7 entries
Codes: C - Connected, L - Local, S - Static, U - Per-user Static route
       B - BGP, R - RIP, I1 - ISIS L1, I2 - ISIS L2
       IA - ISIS interarea, IS - ISIS summary, D - EIGRP, EX - EIGRP external
       ND - ND Default, NDP - ND Prefix, DCE - Destination, NDr - Redirect
       O - OSPF Intra, OI - OSPF Inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2
       ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2
S    2001:DB8:ACAD:1::/64 [1/0]
    via 2001:DB8:FEED:1::2
S    2001:DB8:ACAD:2::/64 [1/0]
    via 2001:DB8:FEED:1::2
S    2001:DB8:ACAD:3::/64 [1/0]
    via 2001:DB8:FEED:1::2
S    2001:DB8:ACAD:4::/64 [1/0]
    via 2001:DB8:FEED:1::2
R1#

```

Calculate IPv6 Network Addresses (2.4.2.2)

Summarizing IPv6 networks into a single IPv6 prefix and prefix length can be done in seven steps as shown in Figures 2-58 to 2-64:



- Step 1.** List the network addresses (prefixes) and identify the part where the addresses differ.
- Step 2.** Expand the IPv6 if it is abbreviated.
- Step 3.** Convert the differing section from hex to binary.
- Step 4.** Count the number of far left matching bits to determine the prefix length for the summary route.
- Step 5.** Copy the matching bits and then add zero bits to determine the summarized network address (prefix).
- Step 6.** Convert the binary section back to hex.
- Step 7.** Append the prefix of the summary route (result of Step 4).

2001:0DB8:ACAD:1::/64
2001:0DB8:ACAD:2::/64
2001:0DB8:ACAD:3::/64
2001:0DB8:ACAD:4::/64

Figure 2-58 Identify the Part Where the Addresses Differ

2001:0DB8:ACAD:0001::/64
2001:0DB8:ACAD:0002::/64
2001:0DB8:ACAD:0003::/64
2001:0DB8:ACAD:0004::/64

Figure 2-59 Identify the Part Where the Addresses Differ – Expanded View

2001:0DB8:ACAD:0000000000000001::/64
2001:0DB8:ACAD:0000000000000010::/64
2001:0DB8:ACAD:0000000000000011::/64
2001:0DB8:ACAD:0000000000000100::/64

Figure 2-60 Convert the Section from Hex to Binary

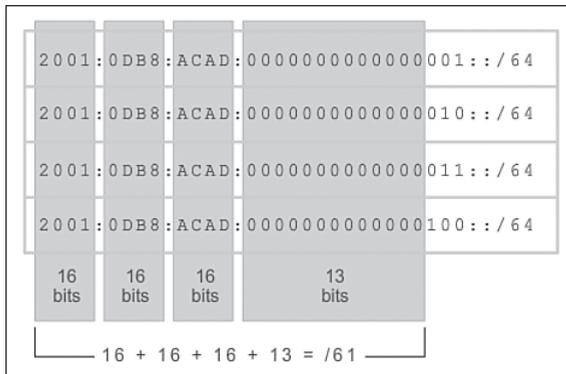


Figure 2-61 Count the Number of Far Left Matching Bits

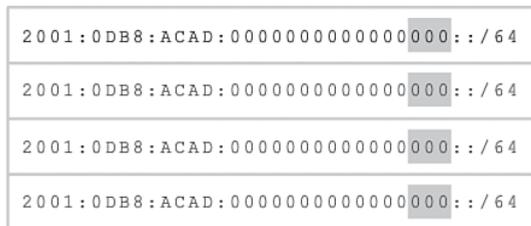


Figure 2-62 Add Zero Bits to Determine the Summarized Network Address

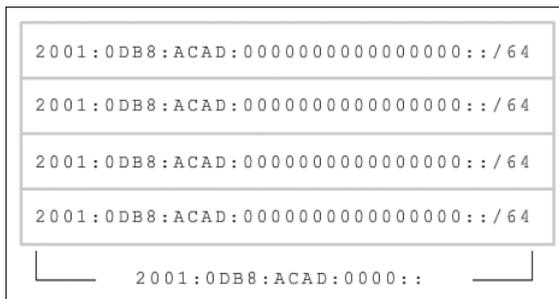


Figure 2-63 Convert the Binary Section Back to Hex

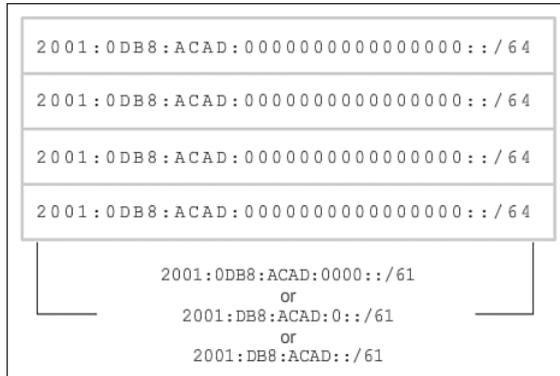


Figure 2-64 Count the Number of Far Left Matching Bits

Configure an IPv6 Summary Address (2.4.2.3)

After the summary route is identified, replace the existing routes with the single summary route.

Figure 2-65 displays how the four existing routes are removed and then the new summary static IPv6 route is configured.

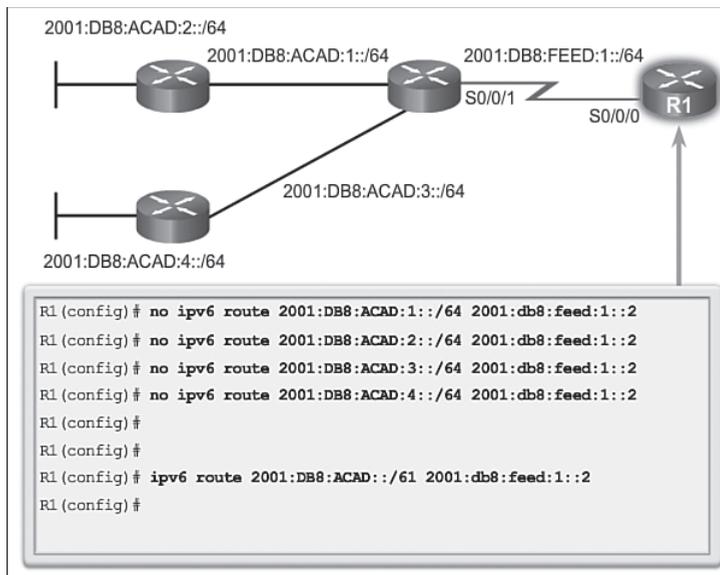


Figure 2-65 Remove Static Routes and Configure Summary IPv6 Route

The following output confirms that the summary static route is in the routing table of R1:

```
R1# show ipv6 route static
IPv6 Routing Table - default - 4 entries
Codes: C - Connected, L - Local, S - Static, U - Per-user Static route
        B - BGP, R - RIP, I1 - ISIS L1, I2 - ISIS L2
        IA - ISIS interarea, IS - ISIS summary, D - EIGRP, EX - EIGRP external
        ND - ND Default, NDP - ND Prefix, DCE - Destination, NDR - Redirect
        O - OSPF Intra, OI - OSPF Inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2
        ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2
S      2001:DB8:ACA8::/45 [1/0]
      via 2001:DB8:FEED:1::2
R1#
```

Packet Tracer
Activity

Packet Tracer Activity 2.4.2.4: Configuring IPv6 Route Summarization

In this activity, you will calculate, configure, and verify a summary route for all the networks R1 can access through R2. R1 is configured with a loopback interface. Instead of adding a LAN or another network to R1, we can use a loopback interface to simplify testing when verifying routing.



Lab 2.4.2.5: Calculating Summary Routes with IPv4 and IPv6

In this lab, you will complete the following objectives:

- Part 1: Calculate IPv4 Summary Routes
 - Part 2: Calculate IPv6 Summary Routes
-

Configure Floating Static Routes (2.4.3)

There may be times when a primary route fails due to physical layer problems, hardware issues, a misconfiguration, or many other reasons. A floating static route can be used as a backup route when there is a secondary path available.

Floating Static Routes (2.4.3.1)

Floating static routes are static routes that have an administrative distance greater than the administrative distance of another static route or dynamic routes. They are very useful when providing a backup to a primary link, as shown in Figure 2-66.

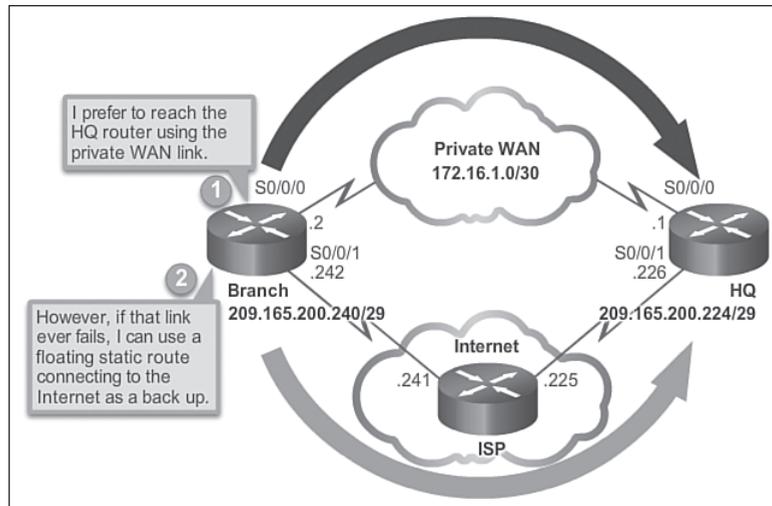


Figure 2-66 Why Configure a Floating Static Route?

By default, static routes have an administrative distance of 1, making them preferable to routes learned from dynamic routing protocols. For example, the administrative distances of some common dynamic routing protocols are:

- EIGRP = 90
- IGRP = 100
- OSPF = 110
- IS-IS = 115
- RIP = 120

The administrative distance of a static route can be increased to make the route less desirable than that of another static route or a route learned through a dynamic routing protocol. In this way, the static route “floats” and is not used when the route with the better administrative distance is active. However, if the preferred route is lost, the floating static route can take over, and traffic can be sent through this alternate route.

A floating static route can be used to provide a backup route to multiple interfaces or networks on a router. It is also encapsulation independent, meaning it can be used to forward packets out any interface, regardless of encapsulation type.

An important consideration of a floating static route is that it is affected by convergence time. A route that is continuously dropping and re-establishing a connection can cause the backup interface to be activated unnecessarily.

Configure a Floating Static Route (2.4.3.2)

IPv4 static routes are configured using the `ip route` global configuration command and specifying an administrative distance. If no administrative distance is configured, the default value (1) is used.

Refer to the topology in Figure 2-67. In this scenario, the preferred route from R1 is to R2. The connection to R3 should be used for backup only.

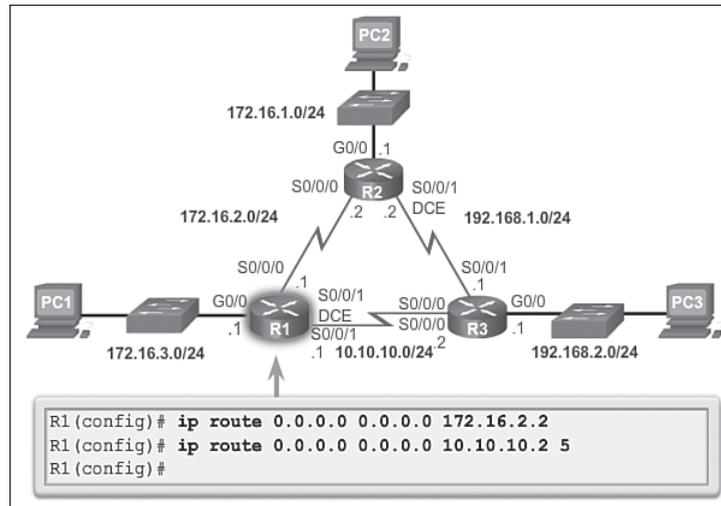


Figure 2-67 Configure a Floating Static Route to R3

R1 is configured with a default static route pointing to R2. Because no administrative distance is configured, the default value (1) is used for this static route. R1 is also configured with a floating static default pointing to R3 with an administrative distance of 5. This value is greater than the default value of 1 and, therefore, this route floats and is not present in the routing table, unless the preferred route fails.

The following output verifies that the default route to R2 is installed in the routing table. Note that the backup route to R3 is not present in the routing table.

```

R1# show ip route static | begin Gateway
Gateway of last resort is 0.0.0.0 to network 0.0.0.0
S*    0.0.0.0/0 [1/0] via 172.16.2.2
R1#
  
```

Interactive
Graphic

Activity 2.4.3.2: Configure a Default Static Route on R3

Go to the online course to use the Syntax Checker in the third graphic to configure a default route using the next-hop address 192.168.1.2.

Test the Floating Static Route (2.4.3.3)

Because the default static route on R1 to R2 has an administrative distance of 1, traffic from R1 to R3 should go through R2. The output in Figure 2-68 confirms that traffic between R1 and R3 flows through R2.

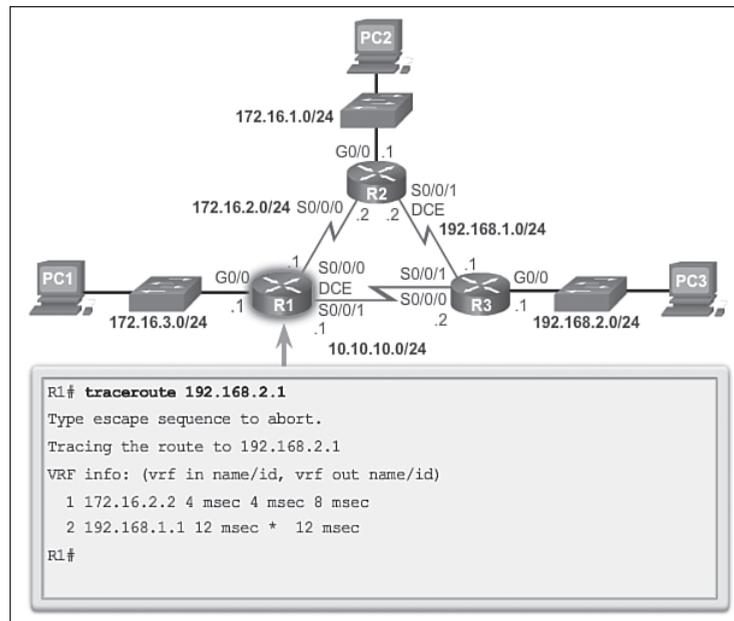


Figure 2-68 Verify the Path to the R3 LAN

What would happen if R2 failed? To simulate this failure, both serial interfaces of R2 are shut down, as shown in the following output:

```

R2(config)# int s0/0/0
R2(config-if)# shut
*Feb 21 16:33:35.939: %LINK-5-CHANGED: Interface Serial0/0/0, changed state to administratively down
*Feb 21 16:33:36.939: %LINEPROTO-5-UPDOWN: Line protocol on Interface Serial0/0/0, changed state to down
R2(config-if)# int s0/0/1
R2(config-if)# shut
R2(config-if)#
*Feb 21 16:33:42.543: %LINK-5-CHANGED: Interface Serial0/0/1, changed state to administratively down
*Feb 21 16:33:43.543: %LINEPROTO-5-UPDOWN: Line protocol on Interface Serial0/0/1, changed state to down

```

Notice in the following output that R1 automatically generates messages indicating that the serial interface to R2 is down. A look at the routing table verifies that the

default route is now pointing to R3 using the floating static default route configured for next-hop 10.10.10.2.

```
*Feb 21 16:35:58.435: %LINK-3-UPDOWN: Interface Serial0/0/0, changed state to down
*Feb 21 16:35:59.435: %LINEPROTO-5-UPDOWN: Line protocol on Interface Serial0/0/0,
  changed state to down
R1#
R1# show ip route static | begin Gateway
Gateway of last resort is 0.0.0.0 to network 0.0.0.0
S*    0.0.0.0/0 [5/0] via 10.10.10.2
R1#
```

The output confirms that traffic now flows directly between R1 and R3:

```
R1# traceroute 192.168.2.1
Type escape sequence to abort.
Tracing the route to 192.168.2.1
VRF info: (vrf in name/id, vrf out name/id)
  1 10.10.10.2 4 msec 4 msec *
R1#
```

Note

Configuring IPv6 floating static routes is outside of the scope of this chapter.

Packet Tracer Activity

Packet Tracer Activity 2.4.3.4: Configuring a Floating Static Route

In this activity, you will configure a floating static route. A floating static route is used as a backup route. It has a manually configured administrative distance greater than that of the primary route and therefore would not be in the routing table until the primary route fails. You will test failover to the backup route, and then restore connectivity to the primary route.

Troubleshoot Static and Default Route Issues (2.5)

Now that you have learned how to configure different types of static routes, this section discusses how to troubleshoot some of the common problems you might encounter. Troubleshooting exercises are an excellent method to help better understand network protocols and configurations. When a static route is no longer needed, that static route should be deleted from the running and startup configuration files.

Packet Processing with Static Routes (2.5.1)

Now that you have configured static routes, you need to learn about the process that a packet goes through as it is forwarded by a router.

Static Routes and Packet Forwarding (2.5.1.1)

The following example describes the packet forwarding process with static routes.

Examine Figure 2-69, where PC1 is sending a packet to PC3:

1. The packet arrives on the FastEthernet 0/0 interface of R1.
2. R1 does not have a specific route to the destination network, 192.168.2.0/24; therefore, R1 uses the default static route.
3. R1 encapsulates the packet in a new frame. Because the link to R2 is a point-to-point link, R1 adds an “all 1s” address for the Layer 2 destination address.
4. The frame is forwarded out of the Serial 0/0/0 interface. The packet arrives on the Serial 0/0/0 interface on R2.
5. R2 de-encapsulates the frame and looks for a route to the destination. R2 has a static route to 192.168.2.0/24 out of the Serial 0/0/1 interface.
6. R2 encapsulates the packet in a new frame. Because the link to R3 is a point-to-point link, R2 adds an “all 1s” address for the Layer 2 destination address.
7. The frame is forwarded out of the Serial 0/0/1 interface. The packet arrives on the Serial 0/0/1 interface on R3.
8. R3 de-encapsulates the frame and looks for a route to the destination. R3 has a connected route to 192.168.2.0/24 out of the FastEthernet 0/0 interface.
9. R3 looks up the ARP table entry for 192.168.2.10 to find the Layer 2 Media Access Control (MAC) address for PC3. If no entry exists, R3 sends an Address Resolution Protocol (ARP) request out of the FastEthernet 0/0 interface, and PC3 responds with an ARP reply, which includes the PC3 MAC address.
10. R3 encapsulates the packet in a new frame with the MAC address of the FastEthernet 0/0 interface as the source Layer 2 address and the MAC address of PC3 as the destination MAC address.
11. The frame is forwarded out of the FastEthernet 0/0 interface. The packet arrives on the network interface card (NIC) interface of PC3.

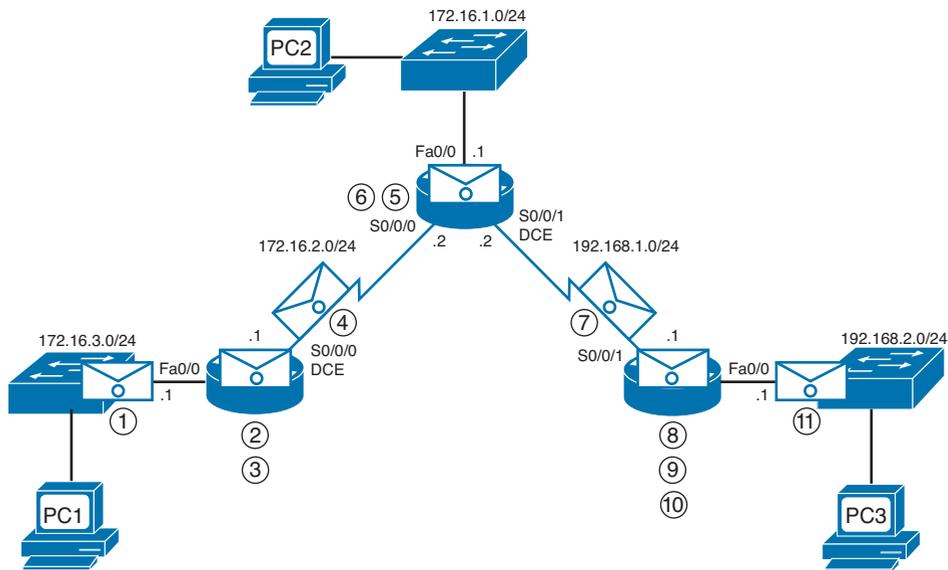


Figure 2-69 Static Routes and Packet Forwarding

Troubleshoot IPv4 Static and Default Route Configuration (2.5.2)

Troubleshooting is a skill that develops as you gain experience. It is always best to look for the most obvious and simplest issues first, such as an interface still in shutdown mode or an interface with the wrong IP address. After these items have been verified, begin looking for more complicated possibilities like an error in the static route configuration.

Troubleshooting a Missing Route (2.5.2.1)

When end-to-end connectivity is a problem, begin by making sure that you can ping your own interface and other devices on your own directly connected networks. When this has been verified, begin testing connectivity to remote networks from other devices.

Networks are subject to forces that can cause their status to change quite often:

- An interface fails.
- A service provider drops a connection.

- Links become oversaturated.
- An administrator enters a wrong configuration.

When there is a change in the network, connectivity may be lost. Network administrators are responsible for pinpointing and solving the problem. To find and solve these issues, a network administrator must be familiar with tools to help isolate routing problems quickly.

Common IOS troubleshooting commands include:

- ping
- traceroute
- show ip route
- show interface brief
- show cdp neighbors detail

Figure 2-70 displays the result of an extended ping from the source interface of R1 to the LAN interface of R3. An extended ping is when the source interface or source IP address is specified.

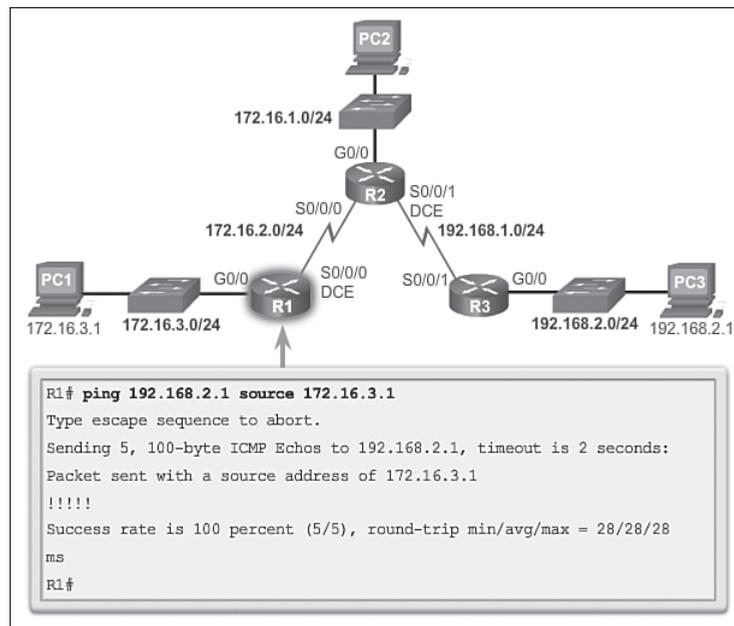


Figure 2-70 Extended Ping

The following output displays the result of a traceroute from R1 to the R3 LAN:

```
R1# traceroute 192.168.2.1
Type escape sequence to abort.
Tracing the route to 192.168.2.1
VRF info: (vrf in name/id, vrf out name/id)
  1 172.16.2.2 4 msec 4 msec 8 msec
  2 192.168.1.1 12 msec 12 msec *
```

The following output displays the routing table of R1:

```
R1# show ip route | begin Gateway
Gateway of last resort is not set
      172.16.0.0/16 is variably subnetted, 5 subnets, 2 masks
S       172.16.1.0/24 [1/0] via 172.16.2.2
C       172.16.2.0/24 is directly connected, Serial0/0/0
L       172.16.2.1/32 is directly connected, Serial0/0/0
C       172.16.3.0/24 is directly connected, GigabitEthernet0/0
L       172.16.3.1/32 is directly connected, GigabitEthernet0/0
S       192.168.1.0/24 [1/0] via 172.16.2.2
S       192.168.2.0/24 [1/0] via 172.16.2.2
R1#
```

The following output provides a quick status of all interfaces on the router:

```
R1# show ip interface brief
Interface                IP-Address      OK? Method Status              Protocol
Embedded-Service-Engine0/0 unassigned      YES unset  administratively down down
GigabitEthernet0/0       172.16.3.1     YES manual up                    up
GigabitEthernet0/1       unassigned      YES unset  administratively down down
Serial0/0/0               172.16.2.1     YES manual up                    up
Serial0/0/1               unassigned      YES unset  administratively down down
R1#
```

The **show cdp neighbors** command in the following output provides a list of directly connected Cisco devices. This command validates Layer 2 (and therefore Layer 1) connectivity. For example, if a neighbor device is listed in the command output, but it cannot be pinged, then Layer 3 addressing should be investigated.

```
R1# show cdp neighbors
Capability Codes: R - Router, T - Trans Bridge, B - Source Route Bridge
                  S - Switch, H - Host, I - IGMP, r - Repeater, P - Phone,
                  D - Remote, C - CVTA, M - Two-port Mac Relay

Device ID    Local Intrfce    Holdtme    Capability  Platform  Port ID
netlab-cs5   Gig 0/0          156        S I         WS-C2960- Fas 0/1
R2           Ser 0/0/0        153        R S I       CISCO1941 Ser 0/0/0
R1#
```

Solve a Connectivity Problem (2.5.2.2)

Finding a missing (or misconfigured) route is a relatively straightforward process, if the right tools are used in a methodical manner.

For instance, in this example, the user at PC1 reports that he cannot access resources on the R3 LAN. This can be confirmed by pinging the LAN interface of R3 using the LAN interface of R1 as the source (see Figure 2-71). The results show that there is no connectivity between these LANs.

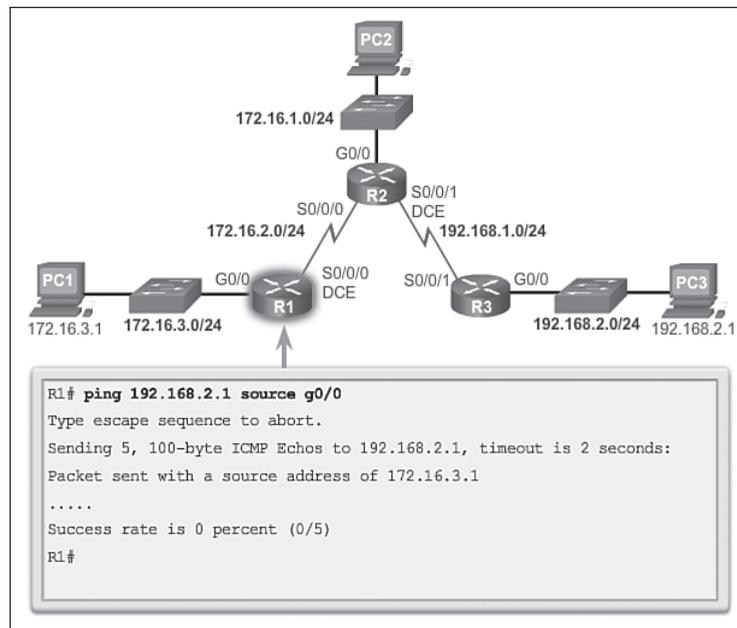


Figure 2-71 Verify Connectivity to the R3 LAN

A traceroute in the following output reveals that R2 is not responding as expected. For some reason, R2 forwards the traceroute back to R1. R1 returns it to R2. This loop would continue until the time to live (TTL) value decrements to zero, in which case, the router would then send an Internet Control Message Protocol (ICMP) Destination Unreachable message to R1.

```

R1# traceroute 192.168.2.1
Type escape sequence to abort.
Tracing the route to 192.168.2.1
VRF info: (vrf in name/id, vrf out name/id)
 0 172.16.2.2 4 msec 4 msec 8 msec
 1 172.16.2.1 12 msec 12 msec 12 msec
 2 172.16.2.2 12 msec 8 msec 8 msec
 3 172.16.2.1 20 msec 16 msec 20 msec
 4 172.16.2.2 16 msec 16 msec 16 msec

```

```
6 172.16.2.1 20 msec 20 msec 24 msec
7 172.16.2.2 20 msec
R1#
```

The next step is to investigate the routing table of R2, because it is the router displaying a strange forwarding pattern. Using the **show ip route | begin Gateway** command, the routing table in the following output reveals that the 192.168.2.0/24 network is configured incorrectly. A static route to the 192.168.2.0/24 network has been configured using the next-hop address 172.16.2.1. Using the configured next-hop address, packets destined for the 192.168.2.0/24 network are sent back to R1. It is clear from the topology that the 192.168.2.0/24 network is connected to R3, not R1. Therefore, the static route to the 192.168.2.0/24 network on R2 must use next-hop 192.168.1.1, not 172.16.2.1.

```
R2# show ip route | begin Gateway
Gateway of last resort is not set
    172.16.0.0/16 is variably subnetted, 5 subnets, 2 masks
C       172.16.1.0/24 is directly connected, GigabitEthernet0/0
L       172.16.1.1/32 is directly connected, GigabitEthernet0/0
C       172.16.2.0/24 is directly connected, Serial0/0/0
L       172.16.2.2/32 is directly connected, Serial0/0/0
S       172.16.3.0/24 [1/0] via 172.16.2.1
    192.168.1.0/24 is variably subnetted, 2 subnets, 2 masks
C       192.168.1.0/24 is directly connected, Serial0/0/1
L       192.168.1.2/32 is directly connected, Serial0/0/1
S       192.168.2.0/24 [1/0] via 172.16.2.1
R2#
```

The following shows output from the running configuration that reveals the incorrect **ip route** statement. The incorrect route is removed and the correct route is then entered.

```
R2# show running-config | section ip route
ip route 172.16.3.0 255.255.255.0 172.16.2.1
ip route 192.168.2.0 255.255.255.0 172.16.2.1
R2#
R2# conf t
Enter configuration commands, one per line. End with CNTL/Z.
R2(config)# no ip route 192.168.2.0 255.255.255.0 172.16.2.1
R2(config)# ip route 192.168.2.0 255.255.255.0 192.168.1.1
R2(config)#
```

The following output verifies that R1 can now reach the LAN interface of R3. As a last step in confirmation, the user on PC1 should also test connectivity to the 192.168.2.0/24 LAN.

```
R1# ping 192.168.2.1 source g0/0
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.168.2.1, timeout is 2 seconds:
Packet sent with a source address of 172.16.3.1
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 28/28/28 ms
R1#
```

Packet Tracer
Activity

Packet Tracer Activity 2.5.2.3: Solving the Missing Route

In this activity, PC1 reports that it cannot access resources at Server. Locate the problem, decide on an appropriate solution, and resolve the issue.

Packet Tracer
Activity

Packet Tracer Activity 2.5.2.4: Troubleshooting VLSM and Route Summarization

In this activity, the network is already addressed using VLSM and configured with static routes. But there is a problem. Locate the issue or issues, determine the best solution, implement the solution, and verify connectivity.



Lab 2.5.2.5: Troubleshooting Static Routes

In this lab, you will complete the following objectives:

- Part 1: Build the Network and Configure Basic Device Settings
 - Part 2: Troubleshoot Static Routes in an IPv4 Network
 - Part 3: Troubleshoot Static Routes in an IPv6 Network
-

Summary (2.6)



Class Activity 2.6.1.1: Make it Static

Go to the online course to perform this practice activity.

As the use of IPv6 addressing becomes more prevalent, it is important for network administrators to be able to direct network traffic between routers.

To prove that you are able to direct IPv6 traffic correctly and review the IPv6 default static route curriculum concepts, use the topology as shown in the .pdf file provided specifically for this activity.

Work with a partner to write an IPv6 statement for each of the three scenarios. Try to write the route statements without the assistance of completed labs, Packet Tracer files, etc.

Scenario 1

IPv6 default static route from R2 directing all data through your S0/0/0 interface to the next-hop address on R1.

Scenario 2

IPv6 default static route from R3 directing all data through your S0/0/1 interface to the next-hop address on R2.

Scenario 3

IPv6 default static route from R2 directing all data through your S0/0/1 interface to the next-hop address on R3.

When complete, get together with another group and compare your written answers. Discuss any differences found in your comparisons.



Packet Tracer Activity 2.6.1.2: Packet Tracer Skills Integration Challenge

The network administrator asked you to implement IPv4 and IPv6 static and default routing in the test environment shown in the topology. Configure each static and default route as directly connected.

In this chapter, you learned how IPv4 and IPv6 static routes can be used to reach remote networks. Remote networks are networks that can only be reached by forwarding the packet to another router. Static routes are easily configured. However, in large networks, this manual operation can become quite cumbersome. Static routes are still used, even when a dynamic routing protocol is implemented.

Static routes can be configured with a next-hop IP address, which is commonly the IP address of the next-hop router. When a next-hop IP address is used, the routing table process must resolve this address to an exit interface. On point-to-point serial links, it is usually more efficient to configure the static route with an exit interface. On multi-access networks, such as Ethernet, both a next-hop IP address and an exit interface can be configured on the static route.

Static routes have a default administrative distance of 1. This administrative distance also applies to static routes configured with a next-hop address, as well as an exit interface.

A static route is only entered in the routing table if the next-hop IP address can be resolved through an exit interface. Whether the static route is configured with a next-hop IP address or exit interface, if the exit interface that is used to forward that packet is not in the routing table, the static route is not included in the routing table.

Using CIDR, several static routes can be configured as a single summary route. This means fewer entries in the routing table and results in a faster routing table lookup process. CIDR also manages the IPv4 address space more efficiently.

VLSM subnetting is similar to traditional subnetting in that bits are borrowed to create subnets. With VLSM, the network is first subnetted, and then the subnets are subnetted again. This process can be repeated multiple times to create subnets of various sizes.

The ultimate summary route is a default route, configured with a 0.0.0.0 network address and a 0.0.0.0 subnet mask for IPv4, and the prefix/prefix-length `::/0` for IPv6. If there is not a more specific match in the routing table, the routing table uses the default route to forward the packet to another router.

A floating static route can be configured to back up a main link by manipulating its administrative value.

Practice

The following activities provide practice with the topics introduced in this chapter. The Labs and Class Activities are available in the companion *Routing Protocols Lab Manual* (978-1-58713-322-0). The Packet Tracer Activities PKA files are found in the online course.



Class Activities

Class Activity 2.0.1.2: Which Way Should We Go?

Class Activity 2.6.1.1: Make it Static



Labs

Lab 2.2.2.5: Configuring IPv4 Static and Default Routes

Lab 2.2.4.5: Configuring IPv6 Static and Default Routes

Lab 2.3.3.7: Designing and Implementing IPv4 Addressing with VLSM

Lab 2.4.2.5: Calculating Summary Routes with IPv4 and IPv6

Lab 2.5.2.5: Troubleshooting Static Routes

Packet Tracer
Activity

Packet Tracer Activities

Packet Tracer Activity 2.2.2.4: Configuring IPv4 Static and Default Routes

Packet Tracer Activity 2.2.4.4: Configuring IPv6 Static and Default Routes

Packet Tracer Activity 2.3.3.6: Designing and Implementing a VLSM Addressing Scheme

Packet Tracer Activity 2.4.1.5: Configuring IPv4 Route Summarization – Scenario 1

Packet Tracer Activity 2.4.1.6: Configuring IPv4 Route Summarization – Scenario 2

Packet Tracer Activity 2.4.2.4: Configuring IPv6 Route Summarization

Packet Tracer Activity 2.4.3.4: Configuring a Floating Static Route

Packet Tracer Activity 2.5.2.3: Solving the Missing Route

Packet Tracer Activity 2.5.2.4: Troubleshooting VLSM and Route Summarization

Packet Tracer Activity 2.6.1.2: Packet Tracer Skills Integration Challenge

Check Your Understanding Questions

Complete all the review questions listed here to test your understanding of the topics and concepts in this chapter. The appendix, “Answers to the ‘Check Your Understanding’ Questions,” lists the answers.

1. Refer to Figure 2-72. Which two commands must be configured to allow communications between the 192.168.10/24 and 10.0.0.0/8 networks? (Choose two.)



Figure 2-72 Topology for Quiz Question 1

- A. A(config)# ip route 10.0.0.0 255.0.0.0 172.16.40.2
 - B. A(config)# ip route 10.0.0.0 255.0.0.0 s0/0/0
 - C. A(config)# ip route 10.0.0.0 255.0.0.0 10.0.0.1
 - D. B(config)# ip route 192.168.1.0 255.255.255.0 172.16.40.1
 - E. B(config)# ip route 192.168.1.0 255.255.255.0 172.16.40.2
 - F. B# ip route 192.168.1.0 255.255.255.0 192.168.1.1
2. Which two statements are true concerning configuring static routes using next-hop addresses? (Choose two.)
- A. Next-hop addresses can only be used with IPv4 static routes. They cannot be used for IPv6 static routes.
 - B. When configuring a static route with a next-hop address, the exit interface must also be included in the configuration.
 - C. Routers configured with static routes using a next-hop address must either have the exit interface listed in the route or have another route with the network of the next hop and an associated exit interface.
 - D. With CEF enabled, there is no need for a recursive lookup when using static routes with next-hop addresses.
3. Which of the following are three characteristics of a static route? (Choose three.)
- A. Less memory and processing requirements than a dynamic routing protocol
 - B. Ensures that there is always a path available
 - C. Used to dynamically find the best path to a destination network
 - D. Used for routers that connect to stub networks
 - E. Used to indicate a default route or a Gateway of Last Resort
 - F. Reduces configuration time on large networks
4. Which global configuration command configures an IPv4 static default route using the next-hop address 10.0.0.1?
- A. Router(config)# ip route 0.0.0.0/0 10.0.0.1
 - B. Router(config)# ip route 0.0.0.0 10.0.0.1
 - C. Router(config)# ipv4 route 0.0.0.0 0.0.0.0 10.0.0.1
 - D. Router(config)# ip route 0.0.0.0 0.0.0.0 10.0.0.1

5. Which global configuration command configures an IPv6 static default route using the next-hop address 2001:DB8:ACAD:1::1?
 - A. Router(config)# **ipv6 route 0.0.0.0 0.0.0.0 2001:DB8:ACAD:1::1**
 - B. Router(config)# **ip route 0/0 2001:DB8:ACAD:1::1**
 - C. Router(config)# **ipv6 route ::/0 2001:DB8:ACAD:1::1**
 - D. Router(config)# **ip route ::/0 2001:DB8:ACAD:1::1**
6. True/False: A static route configured with an exit interface has an administrative distance of 0, the same as a directly connected network.
7. Summarize the following addresses using the shortest valid subnet mask:
10.0.12.0
10.0.13.0
10.0.14.0
10.0.15.0
8. What type of static route can be configured to be a backup route in case the primary route fails?
 - A. Floating static route
 - B. Default route
 - C. Backup static route
 - D. Summary route
9. True/False: Static routes are commonly configured along with a dynamic routing protocol.
10. Which command will only display the IPv6 static routes in the IPv6 routing table?
 - A. Router# **show ip route static**
 - B. Router# **show ip static route**
 - C. Router# **show ipv6 route static**
 - D. Router# **show static ipv6 route**

Routing Dynamically

Objectives

Upon completion of this chapter, you will be able to answer the following questions:

- What is the purpose of dynamic routing protocols?
- How does dynamic routing compare with static routing?
- How do dynamic routing protocols share route information and achieve convergence?
- What are the differences between the categories of dynamic routing protocols?
- How does the algorithm used by distance vector routing protocols determine the best path?
- What are the different types of distance vector routing protocols?
- How do you configure the RIP routing protocol?
- How do you configure the RIPng routing protocol?
- How does the algorithm used by link-state routing protocols determine the best path?
- How do link-state routing protocols use information sent in link-state updates?
- What are the advantages and disadvantages of using link-state routing protocols?
- How do you determine the source route, administrative distance, and metric for a given route?
- How do you explain the concept of a parent/child relationship in a dynamically built routing table?
- How do you describe the differences between the IPv4 route lookup process and the IPv6 route lookup process?
- Can you determine which route will be used to forward a packet upon analyzing a routing table?

Key Terms

This chapter uses the following key terms. You can find the definitions in the Glossary.

dynamic routing protocols page 158

Routing Information Protocol (RIP) page 158

Advanced Research Projects Agency Network (ARPANET) page 158

Open Shortest Path First (OSPF) page 158

Intermediate System-to-Intermediate System (IS-IS) page 159

Interior Gateway Routing Protocol (IGRP) page 159

Enhanced IGRP (EIGRP) page 159

Border Gateway Protocol (BGP) page 159

data structures page 159

routing protocol messages page 160

- algorithm* page 160
- convergence* page 170
- classful routing protocols* page 171
- classless routing protocols* page 171
- autonomous system (AS)* page 172
- Interior Gateway Protocols (IGP)* page 172
- Exterior Gateway Protocols (EGP)* page 172
- distance vector routing protocols* page 174
- link-state routing protocols* page 174
- discontiguous network* page 177
- variable-length subnet mask (VLSM)* page 179
- metrics* page 180
- bounded triggered updates* page 185
- Hello keepalive mechanism* page 185
- automatic summarization* page 193
- passive-interface* page 194
- default static route* page 195
- RIPng* page 196
- Dijkstra's algorithm* page 201
- shortest path first (SPF)* page 201
- link-state packet (LSP)* page 204
- SPF tree* page 211
- event-driven updates* page 213
- OSPFv3* page 214
- ultimate route* page 220
- level 1 route* page 220
- supernet route* page 221
- level 1 parent route* page 221
- level 2 child route* page 222

Introduction (3.0.1.1)

The data networks that we use in our everyday lives to learn, play, and work range from small, local networks to large, global internetworks. At home, a user may have a router and two or more computers. At work, an organization may have multiple routers and switches servicing the data communication needs of hundreds or even thousands of PCs.

Routers forward packets by using information in the routing table. Routes to remote networks can be learned by the router in two ways: static routes and dynamic routes.

In a large network with numerous networks and subnets, configuring and maintaining static routes between these networks requires a great deal of administrative and operational overhead. This operational overhead is especially cumbersome when changes to the network occur, such as a down link or implementing a new subnet. Implementing dynamic routing protocols can ease the burden of configuration and maintenance tasks and give the network scalability.

This chapter introduces dynamic routing protocols. It explores the benefits of using dynamic routing protocols, how different routing protocols are classified, and the metrics routing protocols use to determine the best path for network traffic. Other topics covered in this chapter include the characteristics of dynamic routing protocols and how the various routing protocols differ. Network professionals must understand the different routing protocols available in order to make informed decisions about when to use static or dynamic routing. They also need to know which dynamic routing protocol is most appropriate in a particular network environment.



Class Activity 3.0.1.2: How Much Does This Cost?

This modeling activity illustrates the network concept of routing cost.

You will be a member of a team of five students who travel routes to complete the activity scenarios. One digital camera or bring your own device (BYOD) with camera, a stopwatch, and the student file for this activity will be required per group. One person will function as the photographer and event recorder, as selected by each group. The remaining four team members will actively participate in the following scenarios.

A school or university classroom, hallway, outdoor track area, school parking lot, or any other location can serve as the venue for these activities.

Activity 1

The tallest person in the group establishes a start and finish line by marking 15 steps from start to finish, indicating the distance of the team route. Each student will take 15 steps from the start line toward the finish line and then stop on the 15th step—no further steps are allowed.

Note

Not all of the students may reach the same distance from the start line due to their height and stride differences. The photographer will take a group picture of the entire team's final location after taking the 15 steps required.

Activity 2

A new start and finish line will be established; however, this time, a longer distance for the route will be established than the distance specified in Activity 1. No maximum steps are to be used as a basis for creating this particular route. One at a time, students will walk the new route from beginning to end twice.

Each team member will count the steps taken to complete the route. The recorder will time each student and, at the end of each team member's route, record the time that it took to complete the full route and how many steps were taken, as recounted by each team member and recorded on the team's student file.

After both activities have been completed, teams will use the digital picture taken for Activity 1 and their recorded data from Activity 2 file to answer the reflection questions.

Group answers can be discussed as a class, time permitting.

Dynamic Routing Protocols (3.1)

Dynamic routing protocols play an important role in today's networks. The following sections describe several important benefits that dynamic routing protocols provide. In many networks, dynamic routing protocols are typically used with static routes.

The Evolution of Dynamic Routing Protocols (3.1.1.1)

Dynamic routing protocols have been used in networks since the late 1980s. One of the first routing protocols was *Routing Information Protocol (RIP)*. RIP version 1 (RIPv1) was released in 1988, but some of the basic algorithms within the protocol were used on the *Advanced Research Projects Agency Network (ARPANET)* as early as 1969.

As networks evolved and became more complex, new routing protocols emerged. The RIP routing protocol was updated to accommodate growth in the network environment, into RIPv2. However, the newer version of RIP still does not scale to the larger network implementations of today. To address the needs of larger networks, two advanced routing protocols were developed: *Open Shortest Path First (OSPF)*

and *Intermediate System-to-Intermediate System (IS-IS)*. Cisco developed the *Interior Gateway Routing Protocol (IGRP)* and *Enhanced IGRP (EIGRP)*, which also scales well in larger network implementations.

Additionally, there was the need to connect different internetworks and provide routing between them. The *Border Gateway Protocol (BGP)* is now used between Internet service providers (ISPs). BGP is also used between ISPs and their larger private clients to exchange routing information.

Table 3-1 classifies the protocols.

Table 3-1 Routing Protocol Classification

Interior Gateway Protocols				Exterior Gateway Protocols	
	Distance Vector		Link-State		Path Vector
IPv4	RIPv2	EIGRP	OSPFv2	IS-IS	BGP-4
IPv6	RIPng	EIGRP for IPv6	OSPFv3	IS-IS for IPv6	MBGP

With the advent of numerous consumer devices using IP, the IPv4 addressing space is nearly exhausted; thus, IPv6 has emerged. To support the communication based on IPv6, newer versions of the IP routing protocols have been developed, as shown by the IPv6 row in Table 3-1.

RIP is the simplest of dynamic routing protocols and is used in this section to provide a basic level of routing protocol understanding.

Purpose of Dynamic Routing Protocols (3.1.1.2)

Routing protocols are used to facilitate the exchange of routing information between routers. A routing protocol is a set of processes, algorithms, and messages that are used to exchange routing information and populate the routing table with the routing protocol's choice of best paths. The purpose of dynamic routing protocols includes:

- Discovery of remote networks
- Maintaining up-to-date routing information
- Choosing the best path to destination networks
- Ability to find a new best path if the current path is no longer available

The main components of dynamic routing protocols include:

- **Data structures:** Routing protocols typically use tables or databases for their operations. This information is kept in RAM.

- **Routing protocol messages:** Routing protocols use various types of messages to discover neighboring routers, exchange routing information, and perform other tasks to learn and maintain accurate information about the network.
- **Algorithm:** An algorithm is a finite list of steps used to accomplish a task. Routing protocols use algorithms for facilitating routing information and for best path determination.

Figure 3-1 highlights the data structures, routing protocol messages, and routing algorithm used by EIGRP.

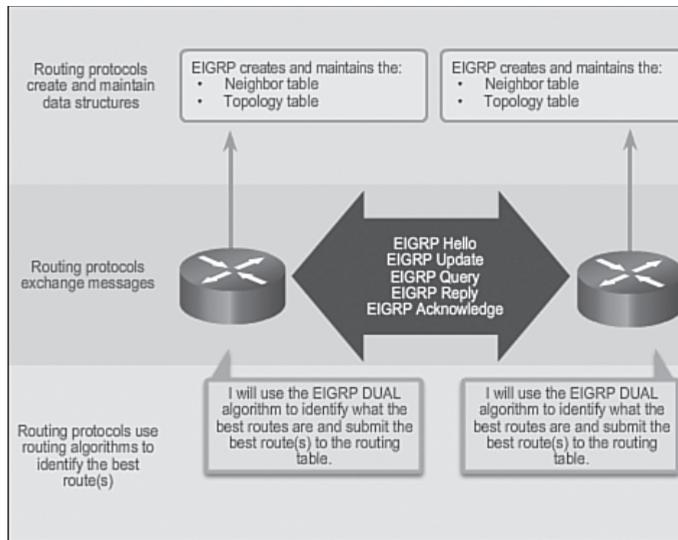


Figure 3-1 Components of Routing Protocols

The Role of Dynamic Routing Protocols (3.1.1.3)

Routing protocols allow routers to dynamically share information about remote networks and automatically add this information to their own routing tables.

Video

Video 3.1.1.3: Routers Dynamically Share Updates

Go to the online course and play the animation of three routers sharing updates dynamically.

Routing protocols determine the best path, or route, to each network. That route is then added to the routing table. A primary benefit of dynamic routing protocols is that routers exchange routing information when there is a topology change. This exchange allows routers to automatically learn about new networks and also to find alternate paths when there is a link failure to a current network.

Compared to static routing, dynamic routing protocols require less administrative overhead. However, the expense of using dynamic routing protocols is dedicating part of a router's resources for protocol operation, including CPU time and network link bandwidth. Despite the benefits of dynamic routing, static routing still has its place. There are times when static routing is more appropriate and other times when dynamic routing is the better choice. Networks with moderate levels of complexity may have both static and dynamic routing configured.

**Interactive
Graphic****Activity 3.1.1.4: Identify Components of a Routing Protocol (EIGRP)**

Go to the online course to perform these three practice activities.

Dynamic versus Static Routing (3.1.2)

Routing tables can contain directly connected, manually configured static routes and routes learned dynamically using a routing protocol. Network professionals must understand when to use static or dynamic routing. This section compares static routing and dynamic routing.

Using Static Routing (3.1.2.1)

Before identifying the benefits of dynamic routing protocols, consider the reasons why network professionals use static routing. Dynamic routing certainly has several advantages over static routing; however, static routing is still used in networks today. In fact, networks typically use a combination of both static and dynamic routing.

Static routing has several primary uses, including:

- Providing ease of routing table maintenance in smaller networks that are not expected to grow significantly.
- Routing to and from a stub network, which is a network with only one default route out and no knowledge of any remote networks.
- Accessing a single default route (which is used to represent a path to any network that does not have a more specific match with another route in the routing table).

Figure 3-2 provides a sample static routing scenario.

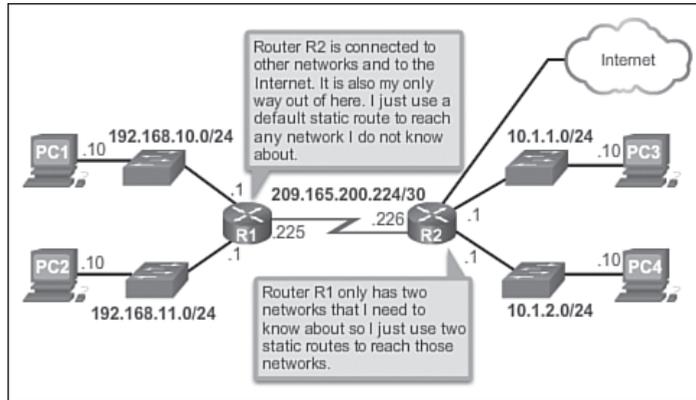


Figure 3-2 Static Routing Scenario

Static Routing Scorecard (3.1.2.2)

Static routing is easy to implement in a small network. Static routes stay the same, which makes them fairly easy to troubleshoot. Static routes do not send update messages and, therefore, require very little overhead.

The disadvantages of static routing include:

- They are not easy to implement in a large network.
- Managing the static configurations can become time consuming.
- If a link fails, a static route cannot reroute traffic.

Table 3-2 highlights the advantages and disadvantages of static routing.

Table 3-2 Static Routing Advantages and Disadvantages

Advantages	Disadvantages
Easy to implement in a small network.	Suitable for simple topologies or for special purposes such as a default static route.
Very secure. No advertisements are sent, unlike with dynamic routing protocols.	Configuration complexity increases dramatically as the network grows. Managing the static configurations in large networks can become time consuming.
It is very predictable, as the route to the destination is always the same.	If a link fails, a static route cannot reroute traffic. Therefore, manual intervention is required to re-route traffic.
No routing algorithm or update mechanisms are required. Therefore, extra resources (CPU and memory) are not required.	

Using Dynamic Routing Protocols (3.1.2.3)

Dynamic routing protocols help the network administrator manage the time-consuming and exacting process of configuring and maintaining static routes.

Imagine maintaining the static routing configurations for the seven routers in Figure 3-3.

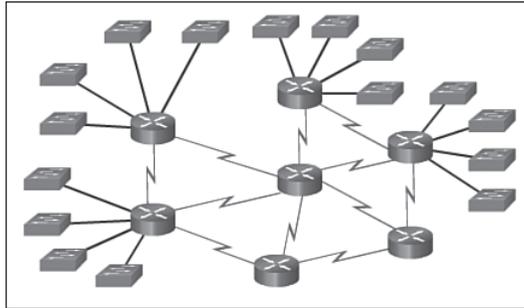


Figure 3-3 Small Dynamic Routing Scenario

What if the company grew and now has four regions and 28 routers to manage, as shown in Figure 3-4? What happens when a link goes down? How do you ensure that redundant paths are available?

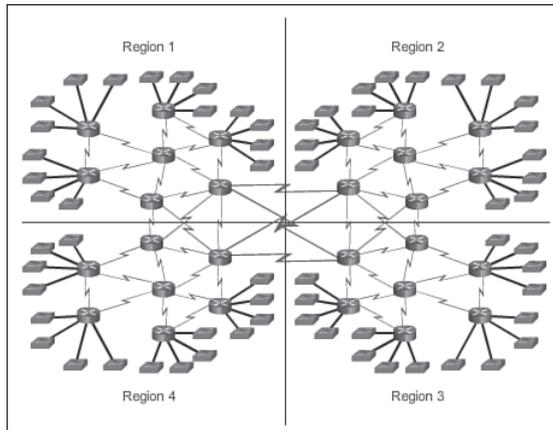


Figure 3-4 Large Dynamic Routing Scenario

Dynamic routing is the best choice for large networks like the one shown in Figure 3-4.

Dynamic Routing Scorecard (3.1.2.4)

Dynamic routing protocols work well in any type of network consisting of several routers. They are scalable and automatically determine better routes if there is a

change in the topology. Although there is more to the configuration of dynamic routing protocols, they are simpler to configure in a large network.

There are disadvantages to dynamic routing. Dynamic routing requires knowledge of additional commands. It is also less secure than static routing because the interfaces identified by the routing protocol send routing updates out. Routes taken may differ between packets. The routing algorithm uses additional CPU, RAM, and link bandwidth.

Table 3-3 highlights the advantages and disadvantages of dynamic routing.

Table 3-3 Dynamic Routing Advantages and Disadvantages

Advantages	Disadvantages
Suitable in all topologies where multiple routers are required.	Can be more complex to initially implement.
Generally independent of the network size.	Less secure due to the broadcast and multicast routing updates. Additional configuration settings such as passive interfaces and routing protocol authentication are required to increase security.
Automatically adapts topology to reroute traffic if possible.	Route depends on the current topology.
	Requires additional resources such as CPU, memory, and link bandwidth.

Notice how dynamic routing addresses the disadvantages of static routing.

**Interactive
Graphic**

Activity 3.1.2.5: Compare Static and Dynamic Routing

Go to the online course to perform this practice activity.

Routing Protocol Operating Fundamentals (3.1.3)

All routing protocols basically perform the same tasks. They all exchange routing updates and converge to build routing tables that are used by the router to make packet forwarding decisions. This section provides an overview of routing protocol fundamentals.

Dynamic Routing Protocol Operation (3.1.3.1)

All routing protocols are designed 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 depends upon the algorithm it uses and the operational characteristics of that protocol.

In general, the operations of a dynamic routing protocol can be described as follows:

1. The router sends and receives routing messages on its interfaces.
2. The router shares routing messages and routing information with other routers that are using the same routing protocol.
3. Routers exchange routing information to learn about remote networks.
4. When a router detects a topology change, the routing protocol can advertise this change to other routers.

Video

Video 3.1.3.1: Routing Protocol Operation

Go to the online course and play the animation of two routers sharing routing updates.

Cold Start (3.1.3.2)

All routing protocols follow the same patterns of operation. When a router powers up, it knows nothing about the network topology. It does not even know that there are devices on the other end of its links. The only information that a router has is from its own saved configuration file stored in NVRAM.

After a router boots successfully, it applies the saved configuration. If the IP addressing is configured correctly, then the router initially discovers its own directly connected networks.

Video

Video 3.1.3.2: Directly Connected Networks Detected

Go to the online course to view an animation of the initial discovery of connected networks for each router.

Notice how the routers proceed through the boot process and then discover any directly connected networks and subnet masks. This information is added to their routing tables as follows:

- R1 adds the 10.1.0.0 network available through interface FastEthernet 0/0 and adds 10.2.0.0 available through interface Serial 0/0/0.

- R2 adds the 10.2.0.0 network available through interface Serial 0/0/0 and adds 10.3.0.0 available through interface Serial 0/0/1.
- R3 adds the 10.3.0.0 network available through interface Serial 0/0/1 and adds 10.4.0.0 available through interface FastEthernet 0/0.

With this initial information, the routers then proceed to find additional route sources for their routing tables.

Network Discovery (3.1.3.3)

After initial boot up and discovery, the routing table is updated with all directly connected networks and the interfaces those networks reside on.

If a routing protocol is configured, the next step is for the router to begin exchanging routing updates to learn about any remote routes.

The router sends an update packet out all interfaces that are enabled on the router. The update contains the information in the routing table, which currently comprises all directly connected networks.

At the same time, the router also receives and processes similar updates from other connected routers. Upon receiving an update, the router checks it for new network information. Any networks that are not currently listed in the routing table are added.

Figure 3-5 depicts an example topology setup between three routers, R1, R2, and R3. Notice that only the directly connected networks are listed in each router's respective routing table.

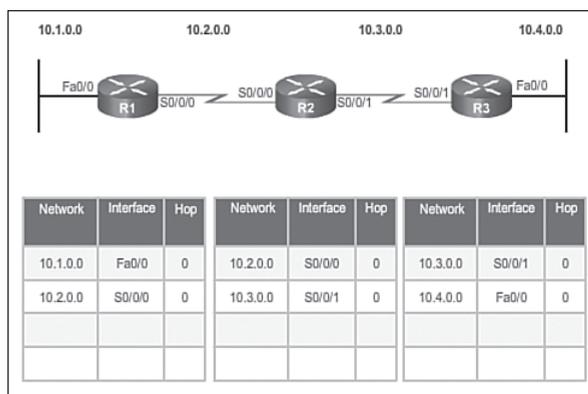


Figure 3-5 Initial Routing Table Before Exchange

Based on this topology, a listing of the different updates that R1, R2, and R3 send and receive during initial convergence is provided:

R1:

- Sends an update about network 10.1.0.0 out the Serial 0/0/0 interface
- Sends an update about network 10.2.0.0 out the FastEthernet 0/0 interface
- Receives update from R2 about network 10.3.0.0 and increments the hop count by 1
- Stores network 10.3.0.0 in the routing table via Serial 0/0/0 with a metric of 1

R2:

- Sends an update about network 10.3.0.0 out the Serial 0/0/0 interface
- Sends an update about network 10.2.0.0 out the Serial 0/0/1 interface
- Receives an update from R1 about network 10.1.0.0 and increments the hop count by 1
- Stores network 10.1.0.0 in the routing table via Serial 0/0/0 with a metric of 1
- Receives an update from R3 about network 10.4.0.0 and increments the hop count by 1
- Stores network 10.4.0.0 in the routing table via Serial 0/0/1 with a metric of 1

R3:

- Sends an update about network 10.4.0.0 out the Serial 0/0/1 interface
- Sends an update about network 10.3.0.0 out the FastEthernet 0/0 interface
- Receives an update from R2 about network 10.2.0.0 and increments the hop count by 1
- Stores network 10.2.0.0 in the routing table via Serial 0/0/1 with a metric of 1

Figure 3-6 displays the routing tables after the initial exchange.

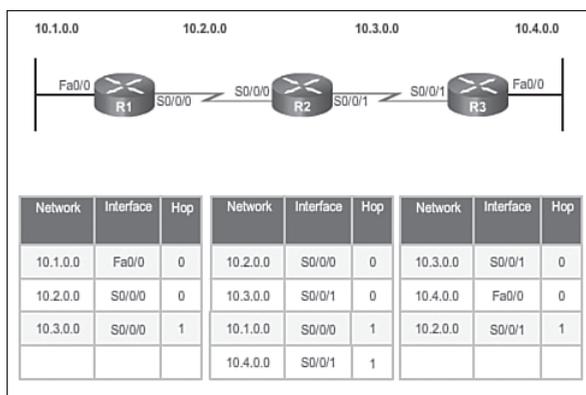


Figure 3-6 Routing Table After Initial Exchange

Video**Video 3.1.3.3: Initial Exchange**

Go to the online course and play the animation of R1, R2, and R3 starting the initial exchange.

After this first round of update exchanges, each router knows about the connected networks of its directly connected neighbors. However, did you notice that R1 does not yet know about 10.4.0.0 and that R3 does not yet know about 10.1.0.0? Full knowledge and a converged network do not take place until there is another exchange of routing information.

Exchanging the Routing Information (3.1.3.4)

At this point the routers have knowledge about their own directly connected networks and about the connected networks of their immediate neighbors. Continuing the journey toward convergence, the routers exchange the next round of periodic updates. Each router again checks the updates for new information.

After initial discovery is complete, each router continues the convergence process by sending and receiving the following updates.

R1:

- Sends an update about network 10.1.0.0 out the Serial 0/0/0 interface
- Sends an update about networks 10.2.0.0 and 10.3.0.0 out the FastEthernet 0/0 interface
- Receives an update from R2 about network 10.4.0.0 and increments the hop count by 1
- Stores network 10.4.0.0 in the routing table via Serial 0/0/0 with a metric of 2
- Same update from R2 contains information about network 10.3.0.0 with a metric of 1. There is no change; therefore, the routing information remains the same.

R2:

- Sends an update about networks 10.3.0.0 and 10.4.0.0 out of Serial 0/0/0 interface
- Sends an update about networks 10.1.0.0 and 10.2.0.0 out of Serial 0/0/1 interface
- Receives an update from R1 about network 10.1.0.0. There is no change; therefore, the routing information remains the same.
- Receives an update from R3 about network 10.4.0.0. There is no change; therefore, the routing information remains the same.

R3:

- Sends an update about network 10.4.0.0 out the Serial 0/0/1 interface
- Sends an update about networks 10.2.0.0 and 10.3.0.0 out the FastEthernet 0/0 interface
- Receives an update from R2 about network 10.1.0.0 and increments the hop count by 1
- Stores network 10.1.0.0 in the routing table via Serial 0/0/1 with a metric of 2
- Same update from R2 contains information about network 10.2.0.0 with a metric of 1. There is no change; therefore, the routing information remains the same.

Figure 3-7 displays the routing tables after the routers have converged.

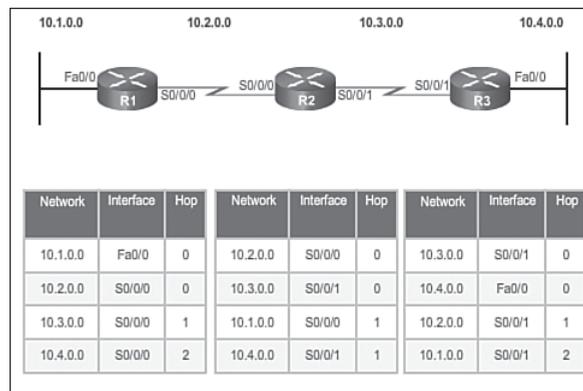


Figure 3-7 Routing Table After Convergence

Video

Video 3.1.3.4: Next Update

Go to the online course and play an animation of R1, R2, and R3 sending the latest routing table to their neighbors.

Distance vector routing protocols typically implement a routing loop prevention technique known as split horizon. Split horizon prevents information from being sent out the same interface from which it was received. For example, R2 does not send an update containing the network 10.1.0.0 out of Serial 0/0/0, because R2 learned about network 10.1.0.0 through Serial 0/0/0.

After routers within a network have converged, the router can then use the information within the route table to determine the best path to reach a destination. Different routing protocols have different ways of calculating the best path.

Achieving Convergence (3.1.3.5)

The network has converged when all routers have complete and accurate information about the entire network, as shown in Figure 3-7. *Convergence* is the time it takes routers to share information, calculate best paths, and update their routing tables. A network is not completely operable until the network has converged; therefore, most networks require short convergence times.

Convergence is both collaborative and independent. The routers share information with each other, but must independently calculate the impacts of the topology change on their own routes. Because they develop an agreement with the new topology independently, they are said to converge on this consensus.

Convergence properties include the speed of propagation of routing information and the calculation of optimal paths. The speed of propagation refers to the amount of time it takes for routers within the network to forward routing information.

As shown in Figure 3-8, routing protocols can be rated based on the speed to convergence; the faster the convergence, the better the routing protocol. Generally, older protocols, such as RIP, are slow to converge, whereas modern protocols, such as EIGRP and OSPF, converge more quickly.

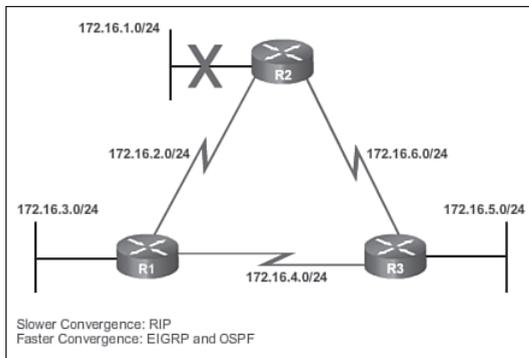


Figure 3-8 Converging

Packet Tracer
Activity

Packet Tracer Activity 3.1.3.6: Investigating Convergence

This activity will help you identify important information in routing tables and witness the process of network convergence.

Types of Routing Protocols (3.1.4)

Table 3-1 showed how routing protocols can be classified according to various characteristics. This section gives an overview of the most common IP routing protocols. Most of these routing protocols will be examined in detail in other chapters. For now, this section gives a very brief overview of each protocol.

Classifying Routing Protocols (3.1.4.1)

Routing protocols can be classified into different groups according to their characteristics. Specifically, routing protocols can be classified by their:

- **Purpose:** Interior Gateway Protocol (IGP) or Exterior Gateway Protocol (EGP)
- **Operation:** Distance vector protocol, link-state protocol, or path-vector protocol
- **Behavior:** Classful (legacy) or classless protocol

For example, IPv4 routing protocols are classified as follows:

- **RIPv1 (legacy):** IGP, distance vector, classful protocol
- **IGRP (legacy):** IGP, distance vector, classful protocol developed by Cisco (deprecated from 12.2 IOS and later)
- **RIPv2:** IGP, distance vector, classless protocol
- **EIGRP:** IGP, distance vector, classless protocol developed by Cisco
- **OSPF:** IGP, link-state, classless protocol
- **IS-IS:** IGP, link-state, classless protocol
- **BGP:** EGP, path-vector, classless protocol

The *classful routing protocols*, RIPv1 and IGRP, are legacy protocols and are only used in older networks. These routing protocols have evolved into the *classless routing protocols*, RIPv2 and EIGRP, respectively. Link-state routing protocols are classless by nature.

Figure 3-9 displays a hierarchical view of dynamic routing protocol classification.

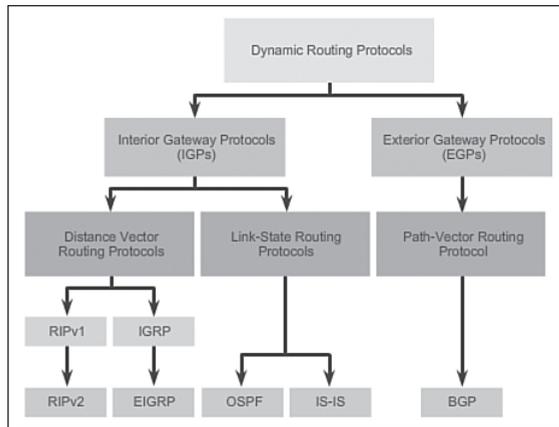


Figure 3-9 Routing Protocol Classification

IGP and EGP Routing Protocols (3.1.4.2)

An *autonomous system (AS)* is a collection of routers under a common administration such as a company or an organization. An AS is also known as a routing domain. Typical examples of an AS are a company's internal network and an ISP's network.

The Internet is based on the AS concept; therefore, two types of routing protocols are required:

- **Interior Gateway Protocols (IGP):** Used for routing within an AS. It is also referred to as intra-AS routing. Companies, organizations, and even service providers use an IGP on their internal networks. IGPs include RIP, EIGRP, OSPF, and IS-IS.
- **Exterior Gateway Protocols (EGP):** Used for routing between autonomous systems. It is also referred to as inter-AS routing. Service providers and large companies may interconnect using an EGP. The Border Gateway Protocol (BGP) is the only currently viable EGP and is the official routing protocol used by the Internet.

Note

Because BGP is the only EGP available, the term EGP is rarely used; instead, most engineers simply refer to BGP.

The example in Figure 3-10 provides simple scenarios highlighting the deployment of IGPs, BGP, and static routing.

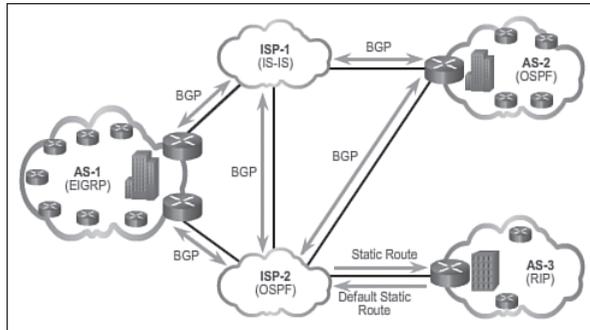


Figure 3-10 IGP versus EGP Routing Protocols

There are five individual autonomous systems in the scenario:

- **ISP-1:** This is an AS and it uses IS-IS as the IGP. It interconnects with other autonomous systems and service providers using BGP to explicitly control how traffic is routed.
- **ISP-2:** This is an AS and it uses OSPF as the IGP. It interconnects with other autonomous systems and service providers using BGP to explicitly control how traffic is routed.
- **AS-1:** This is a large organization and it uses EIGRP as the IGP. Because it is multihomed (i.e., connects to two different service providers), it uses BGP to explicitly control how traffic enters and leaves the AS.
- **AS-2:** This is a medium-sized organization and it uses OSPF as the IGP. It is also multihomed; therefore, it uses BGP to explicitly control how traffic enters and leaves the AS.
- **AS-3:** This is a small organization with older routers within the AS; it uses RIP as the IGP. BGP is not required because it is single-homed (i.e., connects to one service provider). Instead, static routing is implemented between the AS and the service provider.

Note

BGP is beyond the scope of this course and is not discussed in detail.

Distance Vector Routing Protocols (3.1.4.3)

Distance vector means that routes are advertised by providing two characteristics:

- **Distance:** Identifies how far it is to the destination network and is based on a metric such as the hop count, cost, bandwidth, delay, and more

- **Vector:** Specifies the direction of the next-hop router or exit interface to reach the destination

For example, in Figure 3-11, R1 knows that the distance to reach network 172.16.3.0/24 is one hop and that the direction is out of the interface Serial 0/0/0 toward R2.

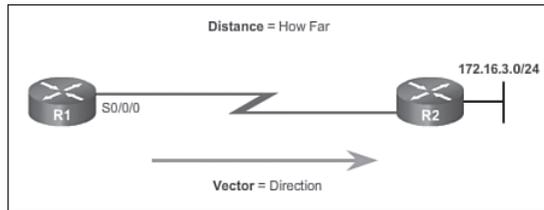


Figure 3-11 The Meaning of Distance Vector

A router using a *distance vector routing protocol* does not have the knowledge of the entire path to a destination network. Distance vector protocols use routers as sign posts along the path to the final destination. The only information a router knows about a remote network is the distance or metric to reach that 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.

There are four distance vector IPv4 IGPs:

- **RIPv1:** First generation legacy protocol
- **RIPv2:** Simple distance vector routing protocol
- **IGRP:** First generation Cisco proprietary protocol (obsolete and replaced by EIGRP)
- **EIGRP:** Advanced version of distance vector routing

Link-State Routing Protocols (3.1.4.4)

In contrast to distance vector routing protocol operation, a router configured with a *link-state routing protocol* can create a complete view or topology of the network by gathering information from all of the other routers.

To continue our analogy of sign posts, using a link-state routing protocol is like having a complete map of the network topology. The sign posts along the way from source to destination are not necessary, because all link-state routers are using an identical map of the network. 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.

RIP-enabled routers send periodic updates of their routing information to their neighbors. Link-state routing protocols do not use periodic updates. After the network has

converged, a link-state update is only sent when there is a change in the topology. For example, in Figure 3-12, the link-state update is sent when the 172.16.3.0 network goes down.

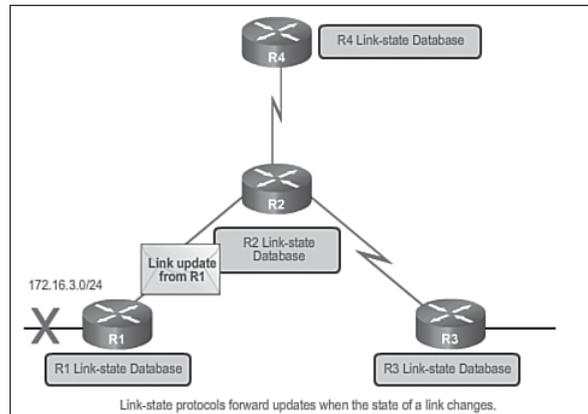


Figure 3-12 Link-State Protocol Operation

Video

Video 3.1.4.4: Link-State Protocol Operation

Go to the online course and play the animation to see how a link-state update is only sent when the 172.16.3.0 network goes down.

Link-state protocols work best in situations where:

- The network design is hierarchical, usually occurring in large networks
- Fast convergence of the network is crucial
- The administrators have good knowledge of the implemented link-state routing protocol

There are two link-state IPv4 IGPs:

- OSPF: Popular standards-based routing protocol
- IS-IS: Popular in provider networks

Classful Routing Protocols (3.1.4.5)

The biggest distinction between classful and classless routing protocols is that classful routing protocols do not send subnet mask information in their routing updates. Classless routing protocols include subnet mask information in the routing updates.

The two original IPv4 routing protocols developed were RIPv1 and IGRP. They were created when network addresses were allocated based on classes (i.e., class A, B, or

C). At that time, a routing protocol did not need to include the subnet mask in the routing update, because the network mask could be determined based on the first octet of the network address.

Note

Only RIPv1 and IGRP are classful. All other IPv4 and IPv6 routing protocols are classless. Classful addressing has never been a part of IPv6.

The fact that RIPv1 and IGRP do not include subnet mask information in their updates means that they cannot provide variable-length subnet masks (VLSMs) and Classless Inter-Domain Routing (CIDR).

Classful routing protocols also create problems in discontinuous networks. A discontinuous network is when subnets from the same classful major network address are separated by a different classful network address.

To illustrate the shortcoming of classful routing, refer to the topology in Figure 3-13.

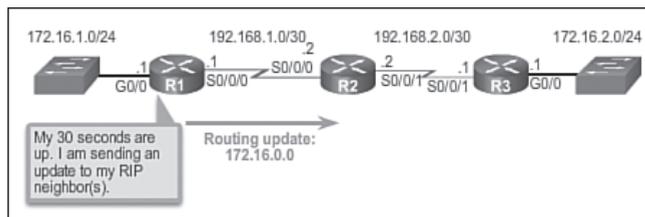


Figure 3-13 R1 Forwards a Classful Update to R2

Notice that the LANs of R1 (172.16.1.0/24) and R3 (172.16.2.0/24) are both subnets of the same class B network (172.16.0.0/16). They are separated by different classful network addresses (192.168.1.0/30 and 192.168.2.0/30).

When R1 forwards an update to R2, RIPv1 does not include the subnet mask information with the update; it only forwards the class B network address 172.16.0.0.

R2 receives and processes the update. It then creates and adds an entry for the class B 172.16.0.0/16 network in the routing table, as shown in Figure 3-14.

```
R2# show ip route | begin Gateway
Gateway of last resort is not set

R    172.16.0.0/16 [120/1] via 192.168.1.1, 00:00:11,
     Serial0/0/0
     192.168.1.0/24 is variably subnetted, 2 subnets,
     2 masks
C    192.168.1.0/30 is directly connected, Serial0/0/0
L    192.168.1.2/32 is directly connected, Serial0/0/0
     192.168.2.0/24 is variably subnetted, 2 subnets, .
     2 masks
C    192.168.2.0/30 is directly connected, Serial0/0/1
L    192.168.2.2/32 is directly connected, Serial0/0/1
R2#
```

Figure 3-14 R2 Adds the Entry for 172.16.0.0 via R1

When R3 forwards an update to R2, it also does not include the subnet mask information and therefore only forwards the classful network address 172.16.0.0.

R2 receives and processes the update and adds another entry for the classful network address 172.16.0.0/16 to its routing table, as shown in Figure 3-15. When there are two entries with identical metrics in the routing table, the router shares the load of the traffic equally among the two links. This is known as load balancing.

```

R2# show ip route | begin Gateway
Gateway of last resort is not set

R    172.16.0.0/16 [120/1] via 192.168.2.1, 00:00:14,
      Serial0/0/1
      [120/1] via 192.168.1.1, 00:00:16,
      Serial0/0/0
      192.168.1.0/24 is variably subnetted, 2 subnets,
      2 masks
C    192.168.1.0/30 is directly connected, Serial0/0/0
L    192.168.1.2/32 is directly connected, Serial0/0/0
L    192.168.2.0/24 is variably subnetted, 2 subnets,
      2 masks
C    192.168.2.0/30 is directly connected, Serial0/0/1
L    192.168.2.2/32 is directly connected, Serial0/0/1
R2#

```

Figure 3-15 R2 Adds the Entry for 172.16.0.0 via R3

Discontiguous networks have a negative impact on a network. For example, a ping to 172.16.1.1 would return “U.U.U” because R2 would forward the first ping out its Serial 0/0/1 interface toward R3, and R3 would return a Destination Unreachable (U) error code to R2. The second ping would exit out of R2’s Serial 0/0/0 interface toward R1, and R1 would return a successful code (.). This pattern would continue until the ping command is done.

Classless Routing Protocols (3.1.4.6)

Modern networks no longer use classful IP addressing and the subnet mask cannot be determined by the value of the first octet. The classless IPv4 routing protocols (RIPv2, EIGRP, OSPF, and IS-IS) all include the subnet mask information with the network address in routing updates. Classless routing protocols support VLSM and CIDR.

IPv6 routing protocols are classless. The distinction whether a routing protocol is classful or classless typically only applies to IPv4 routing protocols. All IPv6 routing protocols are considered classless because they include the prefix-length with the IPv6 address.

Figures 3-16 through 3-18 illustrate how classless routing solves the issues created with classful routing.

In the *discontiguous network* design of Figure 3-16, the classless protocol RIPv2 has been implemented on all three routers. When R1 forwards an update to R2, RIPv2 includes the subnet mask information with the update 172.16.1.0/24.

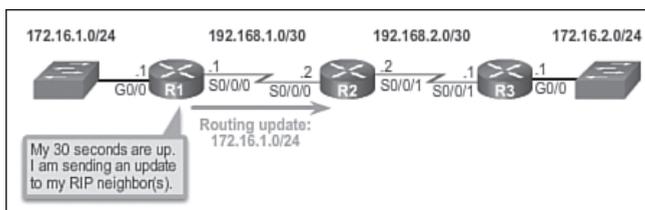


Figure 3-16 R1 Forwards a Classless Update to R2

In Figure 3-17, R2 receives, processes, and adds two entries in the routing table. The first line displays the classful network address 172.16.0.0 with the /24 subnet mask of the update. This is known as the parent route. The second entry displays the VLSM network address 172.16.1.0 with the exit and next-hop address. This is referred to as the child route. Parent routes never include an exit interface or next-hop IP address.

```
R2# show ip route | begin Gateway
Gateway of last resort is not set

172.16.0.0/24 is subnetted, 1 subnets
R    172.16.1.0 [120/1] via 192.168.1.1, 00:00:06,
    Serial0/0/0
192.168.1.0/24 is variably subnetted, 2 subnets,
2 masks
C    192.168.1.0/30 is directly connected, Serial0/0/0
L    192.168.1.2/32 is directly connected, Serial0/0/0
R2#
```

Figure 3-17 R2 Adds the Entry for the 172.16.1.0/24 Network via R1

When R3 forwards an update to R2, RIPv2 includes the subnet mask information with the update 172.16.2.0/24.

R2 receives, processes, and adds another child route entry 172.16.2.0/24 under the parent route entry 172.16.0.0, as shown in Figure 3-18.

```
R2# show ip route | begin Gateway
Gateway of last resort is not set

172.16.0.0/24 is subnetted, 2 subnets
R    172.16.1.0 [120/1] via 192.168.1.1, 00:00:03,
    Serial0/0/0
R    172.16.2.0 [120/1] via 192.168.2.1, 00:00:03,
    Serial0/0/1
192.168.1.0/24 is variably subnetted, 2 subnets,
2 masks
C    192.168.1.0/30 is directly connected, Serial0/0/0
L    192.168.1.2/32 is directly connected, Serial0/0/0
192.168.2.0/24 is variably subnetted, 2 subnets,
2 masks
C    192.168.2.0/30 is directly connected, Serial0/0/1
L    192.168.2.2/32 is directly connected, Serial0/0/1
R2#
```

Figure 3-18 Entry for the 172.16.2.0/24 Network via R3

A ping from R2 to 172.16.1.1 would now be successful.

Routing Protocol Characteristics (3.1.4.7)

Routing protocols can be compared based on the following characteristics:

- **Speed of convergence:** Speed of convergence defines how quickly the routers in the network topology share routing information and reach a state of consistent knowledge. The faster the convergence, the more preferable the protocol. Routing loops can occur when inconsistent routing tables are not updated due to slow convergence in a changing network.
- **Scalability:** Scalability defines how large a network can become, based on the routing protocol that is deployed. The larger the network is, the more scalable the routing protocol needs to be.
- **Classful or classless (use of VLSM):** Classful routing protocols do not include the subnet mask and cannot support *variable-length subnet mask (VLSM)*. Classless routing protocols include the subnet mask in the updates. Classless routing protocols support VLSM and better route summarization.
- **Resource usage:** Resource usage includes the requirements of a routing protocol such as memory space (RAM), CPU utilization, and link bandwidth utilization. Higher resource requirements necessitate more powerful hardware to support the routing protocol operation, in addition to the packet forwarding processes.
- **Implementation and maintenance:** Implementation and maintenance describes the level of knowledge that is required for a network administrator to implement and maintain the network based on the routing protocol deployed.

Table 3-4 summarizes the characteristics of each routing protocol.

Table 3-4 Comparing Routing Protocols

	Distance Vector			Link-State		
	RIPv1	RIPv2	IGRP	EIGRP	OSPF	IS-IS
Speed of Convergence	Slow	Slow	Slow	Fast	Fast	Fast
Scalability – Size of Network	Small	Small	Small	Large	Large	Large
Use of VLSM	No	Yes	No	Yes	Yes	Yes
Resource Usage	Low	Low	Low	Medium	High	High
Implementation and Maintenance	Simple	Simple	Simple	Complex	Complex	Complex

Routing Protocol Metrics (3.1.4.8)

There are cases when a routing protocol learns of more than one route to the same destination. To select the best path, the routing protocol must be able to evaluate and differentiate between the available paths. This is accomplished through the use of routing *metrics*.

A metric is a measurable value that is assigned by the routing protocol to different routes based on the usefulness of that route. In situations where there are multiple paths to the same remote network, the routing metrics are used to determine the overall “cost” of a path from source to destination. Routing protocols determine the best path based on the route with the lowest cost.

Different routing protocols use different metrics. The metric used by one routing protocol is not comparable to the metric used by another routing protocol. Two different routing protocols might choose different paths to the same destination.

For example, assume that PC1 wants to send a packet to PC2. In Figure 3-19, the RIP routing protocol has been enabled on all routers and the network has converged. RIP makes a routing protocol decision based on the least number of hops. Therefore, when the packet arrives on R1, the best route to reach the PC2 network would be to send it directly to R2 even though the link is much slower than all other links.

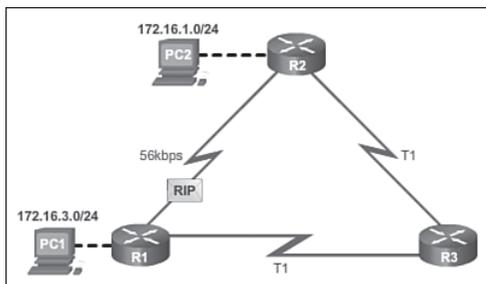


Figure 3-19 RIP Uses Shortest Hop Count Path

In Figure 3-20, the OSPF routing protocol has been enabled on all routers and the network has converged. OSPF makes a routing protocol decision based on the best bandwidth. Therefore, when the packet arrives on R1, the best route to reach the PC2 network would be to send it to R3, which would then forward it to R2.

Video

Video 3.1.4.8: Routing Protocols and Their Metrics

Go to the online course and play the animation showing that RIP would choose the path with the least number of hops, whereas OSPF would choose the path with the highest bandwidth.

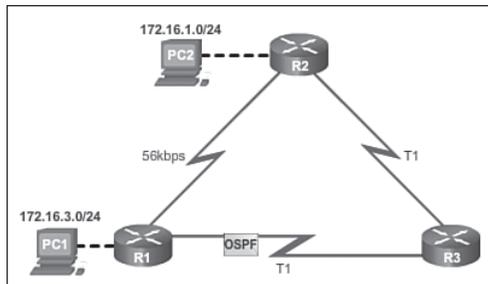


Figure 3-20 OSPF Uses Faster Links

**Interactive
Graphic**

Activity 3.1.4.9: Classify Dynamic Routing Protocols

Go to the online course to perform this practice activity.

**Interactive
Graphic**

Activity 3.1.4.10: Compare Routing Protocols

Go to the online course to perform this practice activity.

**Interactive
Graphic**

Activity 3.1.4.11: Match the Metric to the Protocol

Go to the online course to perform this practice activity.

Distance Vector Dynamic Routing (3.2)

This section describes the characteristics, operations, and functionality of distance vector routing protocols. Understanding the operation of distance vector routing is critical to enabling, verifying, and troubleshooting these protocols.

Distance Vector Technologies (3.2.1.1)

Distance vector routing protocols share updates between neighbors. Neighbors are routers that share a link and are configured to use the same routing protocol. The router is only aware of the network addresses of its own interfaces and the remote network addresses it can reach through its neighbors. Routers using distance vector routing are not aware of the network topology.

Some distance vector routing protocols send periodic updates. For example, RIP sends a periodic update to all of its neighbors every 30 seconds. RIP does this even if the topology has not changed; it continues to send updates. RIPv1 reaches all of

its neighbors by sending updates to the all-hosts IPv4 address of 255.255.255.255, a broadcast.

The broadcasting of periodic updates is inefficient because the updates consume bandwidth and consume network device CPU resources. Every network device has to process a broadcast message. RIPv2 and EIGRP, instead, use multicast addresses so that only neighbors that need updates will receive them. EIGRP can also send a unicast message to only the affected neighbor. Additionally, EIGRP only sends an update when needed, instead of periodically.

As shown in Figure 3-21, the two modern IPv4 distance vector routing protocols are RIPv2 and EIGRP. RIPv1 and IGRP are listed only for historical accuracy.

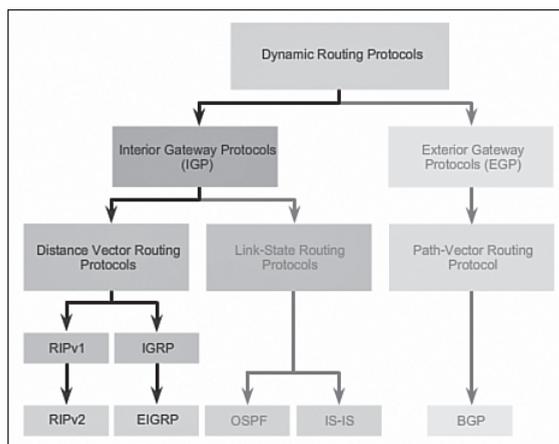


Figure 3-21 Distance Vector Routing Protocols

Distance Vector Algorithm (3.2.1.2)

At the core of the distance vector protocol is the routing algorithm. The algorithm is used to calculate the best paths and then send that information to the neighbors.

The algorithm used for the routing protocols defines the following processes:

- Mechanism for sending and receiving routing information
- Mechanism for calculating the best paths and installing routes in the routing table
- Mechanism for detecting and reacting to topology changes

Video

Video 3.2.1.2: Routers Route Packets

Go to the online course and play the animation to see how the RIP routing protocol adds and deletes routes from a routing table.

In the animation in the online course, R1 and R2 are configured with the RIP routing protocol. The algorithm sends and receives updates. Both R1 and R2 then glean new information from the update. In this case, each router learns about a new network. The algorithm on each router makes its calculations independently and updates the routing table with the new information. When the LAN on R2 goes down, the algorithm constructs a triggered update and sends it to R1. R1 then removes the network from the routing table.

Different routing protocols use different algorithms to install routes in the routing table, send updates to neighbors, and make path determination decisions. For example:

- RIP uses the Bellman-Ford algorithm as its routing algorithm. It is based on two algorithms developed in 1958 and 1956 by Richard Bellman and Lester Ford, Jr.
- IGRP and EIGRP use the Diffusing Update Algorithm (DUAL) routing algorithm developed by Dr. J.J. Garcia-Luna-Aceves at SRI International.

**Interactive
Graphic**

Activity 3.2.1.3: Identify Distance Vector Terminology

Go to the online course to perform this practice activity.

Types of Distance Vector Routing Protocols (3.2.2)

There are two main distance vector routing protocols. This section highlights similarities and differences between RIP and EIGRP.

Routing Information Protocol (3.2.2.1)

The Routing Information Protocol (RIP) was a first generation routing protocol for IPv4 originally specified in RFC 1058. It is easy to configure, making it a good choice for small networks.

RIPv1 has the following key characteristics:

- Routing updates are broadcasted (255.255.255.255) every 30 seconds.
- The hop count is used as the metric for path selection.
- A hop count greater than 15 hops is deemed infinite (too far). That 15th hop router would not propagate the routing update to the next router.

In 1993, RIPv1 evolved to a classless routing protocol known as RIP version 2 (RIPv2). RIPv2 introduced the following improvements:

- **Classless routing protocol:** It supports VLSM and CIDR, because it includes the subnet mask in the routing updates.
- **Increased efficiency:** It forwards updates to multicast address 224.0.0.9, instead of the broadcast address 255.255.255.255.
- **Reduced routing entries:** It supports manual route summarization on any interface.
- **Secure:** It supports an authentication mechanism to secure routing table updates between neighbors.

Table 3-5 summarizes the differences between RIPv1 and RIPv2.

Table 3-5 RIPv1 versus RIPv2

Characteristics and Features	RIPv1	RIPv2
Metric	Both use hop count as a simple metric. The maximum number of hops is 15.	
Updates Forwarded to Address	255.255.255.255	224.0.0.9
Supports VLSM	No	Yes
Supports CIDR	No	Yes
Supports Summarization	No	Yes
Supports Authentication	No	Yes

RIP updates are encapsulated into a UDP segment, with both source and destination port numbers set to UDP port 520.

In 1997, the IPv6-enabled version of RIP was released. RIPng is based on RIPv2. It still has a 15-hop limitation and the administrative distance is 120.

Enhanced Interior Gateway Routing Protocol (3.2.2.2)

The Interior Gateway Routing Protocol (IGRP) was the first proprietary IPv4 routing protocol, developed by Cisco in 1984. It used the following design characteristics:

- Bandwidth, delay, load, and reliability are used to create a composite metric.
- Routing updates are broadcast every 90 seconds, by default.

In 1992, IGRP was replaced by Enhanced IGRP (EIGRP). Like RIPv2, EIGRP also introduced support for VLSM and CIDR. EIGRP increases efficiency, reduces routing updates, and supports secure message exchange.

Table 3-6 summarizes the differences between IGRP and EIGRP.

Table 3-6 IGRP versus EIGRP

Characteristics and Features	IGRP	EIGRP
Metric	Both use a composite metric based on bandwidth and delay. Reliability and load can also be included in the metric calculation if configured.	
Updates Forwarded to Address	255.255.255.255	224.0.0.10
Supports VLSM	No	Yes
Supports CIDR	No	Yes
Supports Summarization	No	Yes
Supports Authentication	No	Yes

EIGRP also introduced:

- **Bounded triggered updates:** It does not send periodic updates. Only routing table changes are propagated, whenever a change occurs. This reduces the amount of load the routing protocol places on the network. Bounded triggered updates means that EIGRP only sends to the neighbors that need it. It uses less bandwidth, especially in large networks with many routes.
- **Hello keepalive mechanism:** A small Hello message is periodically exchanged to maintain adjacencies with neighboring routers. This means a very low usage of network resources during normal operation, instead of the periodic updates.
- **Maintains a topology table:** Maintains all the routes received from neighbors (not only the best paths) in a topology table. DUAL can insert backup routes into the EIGRP topology table.
- **Rapid convergence:** In most cases, it is the fastest IGP to converge because it maintains alternate routes, enabling almost instantaneous convergence. If a primary route fails, the router can use the alternate route identified. The switchover to the alternate route is immediate and does not involve interaction with other routers.

- **Multiple network layer protocol support:** EIGRP uses Protocol Dependent Modules (PDM), which means that it is the only protocol to include support for protocols other than IPv4 and IPv6, such as legacy IPX and AppleTalk.

**Interactive
Graphic****Activity 3.2.2.3: Compare RIP and EIGRP**

Go to the online course to perform this practice activity.

**Packet Tracer
Activity****Packet Tracer Activity 3.2.2.4: Comparing RIP and EIGRP Path Selection**

PCA and PCB need to communicate. The path that the data takes between these end devices can travel through R1, R2, and R3, or it can travel through R4 and R5. The process by which routers select the best path depends on the routing protocol. We will examine the behavior of two distance vector routing protocols, Enhanced Interior Gateway Routing Protocol (EIGRP) and Routing Information Protocol version 2 (RIPv2).

RIP and RIPv2 Routing (3.3)

Although the use of RIP has decreased in the past decade, it is still important to your networking studies because it might be encountered in a network implementation. As well, understanding how RIP operates and knowing its implementation will make learning other routing protocols easier.

Configuring the RIP Protocol (3.3.1)

In this section, you will learn how to configure, verify, and troubleshoot RIPv2.

Router RIP Configuration Mode (3.3.1.1)

Although RIP is rarely used in modern networks, it is useful as a foundation for understanding basic network routing. For this reason, this section provides a brief overview of how to configure basic RIP settings and to verify RIPv2.

Refer to the reference topology in Figure 3-22 and the addressing table in Table 3-7.

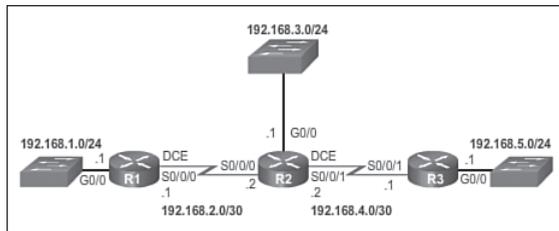


Figure 3-22 RIP Reference Topology

Table 3-7 Addressing Table

Device	Interface	IP Address	Subnet Mask
R1	G0/0	192.168.1.1	255.255.255.0
	S0/0/0	192.168.2.1	255.255.255.0
R2	G0/0	192.168.3.1	255.255.255.0
	S0/0/0	192.168.2.2	255.255.255.0
	S0/0/1	192.168.4.2	255.255.255.0
R3	G0/0	192.168.5.1	255.255.255.0
	S0/0/1	192.168.4.1	255.255.255.0

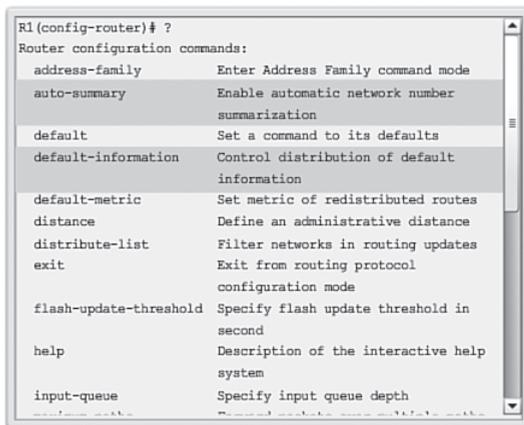
In this scenario, all routers have been configured with basic management features and all interfaces identified in the reference topology are configured and enabled. There are no static routes configured and no routing protocols enabled; therefore, remote network access is currently impossible. RIPv2 is used as the dynamic routing protocol.

To enable RIP, use the **router rip** command to enter router configuration mode, as shown in the following output. This command does not directly start the RIP process. Instead, it provides access to the router configuration mode where the RIP routing settings are configured.

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

To disable and eliminate RIP, use the **no router rip** global configuration command. This command stops the RIP process and erases all existing RIP configurations.

Figure 3-23 displays a partial list of the various RIP commands that can be configured. This section covers the two highlighted commands as well as **network**, **passive-interface**, and **version**.



```
R1(config-router)# ?
Router configuration commands:
  address-family  Enter Address Family command mode
  auto-summary    Enable automatic network number
                  summarization
  default         Set a command to its defaults
  default-information  Control distribution of default
                  information
  default-metric  Set metric of redistributed routes
  distance        Define an administrative distance
  distribute-list  Filter networks in routing updates
  exit           Exit from routing protocol
                  configuration mode
  flash-update-threshold  Specify flash update threshold in
                  second
  help           Description of the interactive help
                  system
  input-queue    Specify input queue depth
```

Figure 3-23 RIP Configuration Options

Note

The entire output in Figure 3-23 can be viewed in the online course on page 3.3.1.1 graphic number 4.

Advertising Networks (3.3.1.2)

By entering the RIP router configuration mode, the router is instructed to run RIP. But the router still needs to know which local interfaces it should use for communication with other routers, as well as which locally connected networks it should advertise to those routers.

To enable RIP routing for a network, use the **network** *network-address* router configuration mode command. Enter the classful network address for each directly connected network. This command:

- Enables RIP on all interfaces that belong to a specific network. Associated interfaces now both send and receive RIP updates.
- Advertises the specified network in RIP routing updates sent to other routers every 30 seconds.

Note

If a subnet address is entered, the IOS automatically converts it to the classful network address. Remember RIPv1 is a classful routing protocol for IPv4. For example, entering the **network 192.168.1.32** command would automatically be converted to **network 192.168.1.0** in the running configuration file. The IOS does not give an error message, but instead corrects the input and enters the classful network address.

In the following command sequence, the **network** command is used to advertise the R1 directly connected networks.

```
R1(config)# router rip
R1(config-router)# network 192.168.1.0
R1(config-router)# network 192.168.2.0
R1(config-router)#
```

Interactive Graphic

Activity 3.3.1.2: Advertising the R2 and R3 Networks

Go to the online course to use the Syntax Checker in the second graphic to configure a similar configuration on R2 and R3.

Examining Default RIP Settings (3.3.1.3)

The output of the **show ip protocols** command in Figure 3-24 displays the IPv4 routing protocol settings currently configured on the router.

```
R1# show ip protocols
*** IP Routing is NSF aware ***

1 Routing Protocol is "rip"
   Outgoing update filter list for all interfaces is not set
   Incoming update filter list for all interfaces is not set
2   Sending updates every 30 seconds, next due in 16 seconds
   Invalid after 180 seconds, hold down 180, flushed after 240
   Redistributing: rip

3   Default version control: send version 1, receive any version
   Interface          Send Recv  Triggered RIP  Key-chain
   GigabitEthernet0/0  1       1 2
   Serial0/0/0        1       1 2

4   Automatic network summarization is in effect
   Maximum path: 4
5   Routing for Networks:
     192.168.1.0
     192.168.2.0

6   Routing Information Sources:
   Gateway         Distance   Last Update
   192.168.2.2     120       00:00:15
   Distance: (default is 120)

R1#
```

Figure 3-24 Verifying RIP Settings on R1

This output confirms that:

1. RIP routing is configured and running on router R1.
2. The values of various timers; for example, the next routing update is sent by R1 in 16 seconds.
3. The version of RIP configured is currently RIPv1.
4. R1 is currently summarizing at the classful network boundary.
5. The classful networks are advertised by R1. These are the networks that R1 includes in its RIP updates.

6. The RIP neighbors are listed, including their next-hop IP address, the associated AD that R2 uses for updates sent by this neighbor, and when the last update was received from this neighbor.

Note

This command is also very useful when verifying the operations of other routing protocols (i.e., EIGRP and OSPF).

The `show ip route` command displays the RIP routes installed in the routing table. In Figure 3-25, R1 now knows about the highlighted networks.

```
R1# show ip route | begin Gateway
Gateway of last resort is not set

  192.168.1.0/24 is variably subnetted, 2 subnets, 2 masks
C   192.168.1.0/24 is directly connected, GigabitEthernet0/0
L   192.168.1.1/32 is directly connected, GigabitEthernet0/0
  192.168.2.0/24 is variably subnetted, 2 subnets, 2 masks
C   192.168.2.0/24 is directly connected, Serial0/0/0
L   192.168.2.1/32 is directly connected, Serial0/0/0
R   192.168.3.0/24 [120/1] via 192.168.2.2, 00:00:24, Serial0/0/0
R   192.168.4.0/24 [120/1] via 192.168.2.2, 00:00:24, Serial0/0/0
R   192.168.5.0/24 [120/2] via 192.168.2.2, 00:00:24, Serial0/0/0
R1#
```

Figure 3-25 Verifying RIP Routes on R1

Interactive Graphic

Activity 3.3.1.3: Advertising the R2 and R3 Networks

Go to the online course to use the Syntax Checker in the third graphic to verify the R2 and R3 RIP settings and routes.

Enabling RIPv2 (3.3.1.4)

By default, when a RIP process is configured on a Cisco router, it is running RIPv1, as shown in the following output:

```
R1# show ip protocols
*** IP Routing is NSF aware ***

Routing Protocol is "rip"
  Outgoing update filter list for all interfaces is not set
  Incoming update filter list for all interfaces is not set
  Sending updates every 30 seconds, next due in 16 seconds
  Invalid after 180 seconds, hold down 180, flushed after 240
  Redistributing: rip
  Default version control: send version 1, receive any version
  Interface          Send  Recv  Triggered RIP  Key-chain
  GigabitEthernet0/0  1    1 2
  Serial0/0/0        1    1 2
```

```

Automatic network summarization is in effect
Maximum path: 4
Routing for Networks:
  192.168.1.0
  192.168.2.0
Routing Information Sources:
  Gateway         Distance      Last Update
  192.168.2.2         120          00:00:15
Distance: (default is 120)

R1#

```

However, even though the router only sends RIPv1 messages, it can interpret both RIPv1 and RIPv2 messages. A RIPv1 router ignores the RIPv2 fields in the route entry.

Use the **version 2** router configuration mode command to enable RIPv2, as shown in Figure 3-26.

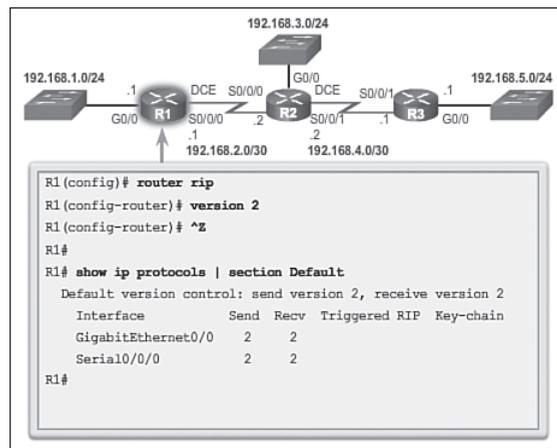


Figure 3-26 Enable and Verify RIPv2 on R1

Notice how the **show ip protocols** command verifies that R2 is now configured to send and receive version 2 messages only. The RIP process now includes the subnet mask in all updates, making RIPv2 a classless routing protocol.

Note

Configuring **version 1** enables RIPv1 only, while configuring **no version** returns the router to the default setting of sending version 1 updates but listening for version 1 or version 2 updates.

The following output verifies that there are no RIP routes still in the routing table:

```
R1# show ip route | begin Gateway
Gateway of last resort is not set

      192.168.1.0/24 is variably subnetted, 2 subnets, 2 masks
C       192.168.1.0/24 is directly connected, GigabitEthernet0/0
L       192.168.1.1/32 is directly connected, GigabitEthernet0/0
      192.168.2.0/24 is variably subnetted, 2 subnets, 2 masks
C       192.168.2.0/24 is directly connected, Serial0/0/0
L       192.168.2.1/32 is directly connected, Serial0/0/0
R1#
```

There are no RIP routes because R1 is now only listening for RIPv2 updates. R2 and R3 are still sending RIPv1 updates. Therefore, the **version 2** command must be configured on all routers in the routing domain.

Interactive Graphic

Activity 3.3.1.4: Enable and Verify RIPv2 on R2 and R3

Go to the online course to use the Syntax Checker in the fourth graphic to enable RIPv2 on R2 and R3.

Disabling Auto Summarization (3.3.1.5)

As shown in Figure 3-27, RIPv2 automatically summarizes networks at major network boundaries by default, just like RIPv1.

```
R1# show ip protocols
*** IP Routing is NSF aware ***

Routing Protocol is "rip"
  Outgoing update filter list for all interfaces is not set
  Incoming update filter list for all interfaces is not set
  Sending updates every 30 seconds, next due in 16 seconds
  Invalid after 180 seconds, hold down 180, flushed after
  240
  Redistributing: rip
  Default version control: send version 2, receive version 2
  Interface          Send Recv Triggered RIP Key-chain
  GigabitEthernet0/0  1     1 2
  Serial0/0/0        1     1 2
  Automatic network summarization is in effect
  Maximum path: 4
  Routing for Networks:
    192.168.1.0
    192.168.2.0
  Routing Information Sources:
    Gateway         Distance    Last Update
    192.168.2.2     120        00:00:15
  Distance: (default is 120)
R1#
```

Figure 3-27 Verify RIPv2 Route Summarization

To modify the default RIPv2 behavior of *automatic summarization*, use the **no auto-summary** router configuration mode command as shown in the following command sequence:

```
R1(config)# router rip
R1(config-router)# no auto-summary
R1(config-router)# end
R1#
*Mar 10 14:11:49.659: %SYS-5-CONFIG_I: Configured from console by console
R1# show ip protocols | section Automatic
    Automatic network summarization is not in effect
R1#
```

This command has no effect when using RIPv1. When automatic summarization has been disabled, RIPv2 no longer summarizes networks to their classful address at boundary routers. RIPv2 now includes all subnets and their appropriate masks in its routing updates. The **show ip protocols** output now states that automatic network summarization is not in effect.

Note

RIPv2 must be enabled before automatic summarization is disabled.

Interactive Graphic

Activity 3.3.1.5: Disable Automatic Summarization on R2 and R3

Go to the online course to use the Syntax Checker in the third graphic to disable automatic summarization on R2 and R3.

Configuring Passive Interfaces (3.3.1.6)

By default, RIP updates are forwarded out all RIP-enabled interfaces. However, RIP updates really only need to be sent out interfaces connecting to other RIP-enabled routers.

For instance, refer to the topology in Figure 3-22. RIP sends updates out of its Gigabit Ethernet 0/0 interface even though no RIP device exists on that LAN. R1 has no way of knowing this and, as a result, sends an update every 30 seconds. Sending out unneeded updates on a LAN impacts the network in three ways:

- **Wasted bandwidth:** Bandwidth is used to transport unnecessary updates. Because RIP updates are either broadcasted or multicasted, switches also forward the updates out all ports.
- **Wasted resources:** All devices on the LAN must process the update up to the transport layers, at which point the devices will discard the update.

- **Security risk:** Advertising updates on a broadcast network is a security risk. RIP updates can be intercepted with packet sniffing software. Routing updates can be modified and sent back to the router, corrupting the routing table with false metrics that misdirect traffic.

To address these problems, an interface can be configured to stop sending routing updates. This is referred to as configuring a *passive interface*. Use the `passive-interface` router configuration command to prevent the transmission of routing updates through a router interface but still allow that network to be advertised to other routers. The command stops routing updates out the specified interface. However, the network that the specified interface belongs to is still advertised in routing updates that are sent out other interfaces.

There is no need for R1, R2, and R3 to forward RIP updates out of their LAN interfaces. The configuration in Figure 3-28 identifies the R1 Gigabit Ethernet 0/0 interface as passive.

```

R1(config)# router rip
R1(config-router)# passive-interface g0/0
R1(config-router)# end
R1#
R1# show ip protocols | begin Default
Default version control: send version 2, receive version 2
Interface      Send Recv Triggered RIP Key-chain
Serial0/0/0    2      2
Automatic network summarization is not in effect
Maximum path: 4
Routing for Networks:
 192.168.1.0
 192.168.2.0
Passive Interface(s):
GigabitEthernet0/0
Routing Information Sources:
Gateway      Distance    Last Update
192.168.2.2  120         00:00:06
Distance: (default is 120)
R1#

```

Figure 3-28 Configuring and Verifying a Passive Interface on R1

The `show ip protocols` command is then used to verify that the Gigabit Ethernet interface was passive. Notice that the Gigabit Ethernet 0/0 interface is no longer listed as sending or receiving version 2 updates, but instead is now listed under the Passive Interface(s) section. Also notice that the network 192.168.1.0 is still listed under Routing for Networks, which means that this network is still included as a route entry in RIP updates that are sent to R2.

Note

All routing protocols support the `passive-interface` command.

Interactive
Graphic**Activity 3.3.1.6: Configuring and Verifying a Passive Interface on R2 and R3**

Go to the online course to use the Syntax Checker in the third graphic to configure a passive interface on R2 and R3.

As an alternative, all interfaces can be made passive using the **passive-interface default** command. Interfaces that should not be passive can be re-enabled using the **no passive-interface** command.

Propagating a Default Route (3.3.1.7)

In the topology in Figure 3-29, R1 is single-homed to a service provider. Therefore, all that is required for R1 to reach the Internet is a *default static route* going out of the Serial 0/0/1 interface.

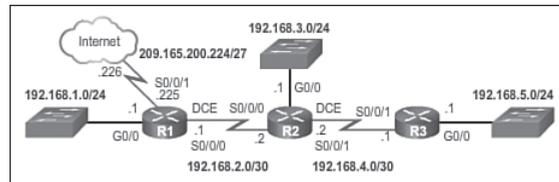


Figure 3-29 Propagating a Default Route on R1

Similar default static routes could be configured on R2 and R3, but it is much more scalable to enter it one time on the edge router R1 and then have R1 propagate it to all other routers using RIP. To provide Internet connectivity to all other networks in the RIP routing domain, the default static route needs to be advertised to all other routers that use the dynamic routing protocol.

To propagate a default route, the edge router must be configured with:

- A default static route using the **ip route 0.0.0.0 0.0.0.0 exit-intf next-hop-ip** command.
- The **default-information originate** router configuration command. This instructs R1 to originate default information, by propagating the static default route in RIP updates.

The example in Figure 3-30 configures a fully specified default static route to the service provider, and then the route is propagated by RIP. Notice that R1 now has a Gateway of Last Resort and default route installed in its routing table.

```

R1(config)# ip route 0.0.0.0 0.0.0.0 S0/0/1 209.165.200.226
R1(config)# router rip
R1(config-router)# default-information originate
R1(config-router)# ^Z
R1#
*Mar 10 23:33:51.801: %SYS-5-CONFIG_I: Configured from console by
console
R1# show ip route | begin Gateway
Gateway of last resort is 209.165.200.226 to network 0.0.0.0

S* 0.0.0.0/0 [1/0] via 209.165.200.226, Serial0/0/1
C 192.168.1.0/24 is variably subnetted, 2 subnets, 2 masks
C 192.168.1.0/24 is directly connected, GigabitEthernet0/0
L 192.168.1.1/32 is directly connected, GigabitEthernet0/0
C 192.168.2.0/24 is variably subnetted, 2 subnets, 2 masks
C 192.168.2.0/24 is directly connected, Serial0/0/0
L 192.168.2.1/32 is directly connected, Serial0/0/0
R 192.168.3.0/24 [120/1] via 192.168.2.2, 00:00:08,
Serial0/0/0
R 192.168.4.0/24 [120/1] via 192.168.2.2, 00:00:08,
Serial0/0/0
R 192.168.5.0/24 [120/2] via 192.168.2.2, 00:00:08,
Serial0/0/0
C 209.165.200.0/24 is variably subnetted, 2 subnets, 2 masks
C 209.165.200.0/24 is directly connected, Serial0/0/1

```

Figure 3-30 Configuring and Verifying a Default Route on R1

Interactive
Graphic

Activity 3.3.1.7: Verifying the Gateway of Last Resort on R2 and R3

Go to the online course to use the Syntax Checker in the third graphic to verify that the default route has been propagated to R2 and R3.

Packet Tracer
Activity

Packet Tracer Activity 3.3.1.8: Configuring RIPv2

Although RIP is rarely used in modern networks, it is useful as a foundation for understanding basic network routing. In this activity, you will configure a default route, configure RIP version 2 with appropriate network statements and passive interfaces, and verify full connectivity.

Configuring the RIPng Protocol (3.3.2)

In this section, you will learn how to configure, verify, and troubleshoot RIPng.

Advertising IPv6 Networks (3.3.2.1)

As with its IPv4 counterpart, RIPng is rarely used in modern networks. It is also useful as a foundation for understanding basic network routing. For this reason, this section provides a brief overview of how to configure basic *RIPng*.

Refer to the reference topology in Figure 3-31.

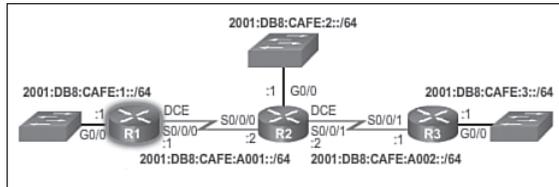


Figure 3-31 Enabling RIPng on the R1 Interfaces

In this scenario, all routers have been configured with basic management features and all interfaces identified in the reference topology are configured and enabled. There are no static routes configured and no routing protocols enabled; therefore, remote network access is currently impossible.

To enable an IPv6 router to forward IPv6 packets, **ipv6 unicast-routing** must be configured.

Unlike RIPv2, RIPng is enabled on an interface and not in router configuration mode. In fact, there is no **network network-address** command available in RIPng. Instead, use the **ipv6 rip domain-name enable** interface configuration command.

In the following output, IPv6 unicast routing is enabled and the Gigabit Ethernet 0/0 and Serial 0/0/0 interfaces are enabled for RIPng using the domain name RIP-AS:

```
R1(config)# ipv6 unicast-routing
R1(config)#
R1(config)# interface gigabitethernet 0/0
R1(config-if)# ipv6 rip RIP-AS enable
R1(config-if)# exit
R1(config)#
R1(config)# interface serial 0/0/0
R1(config-if)# ipv6 rip RIP-AS enable
R1(config-if)# no shutdown
R1(config-if)#
```

**Interactive
Graphic**

Activity 3.3.2.1: Enabling RIPng on the R2 and R3 Interfaces

Go to the online course to use the Syntax Checker in the second graphic to enable RIPng on the R2 and R3 interfaces.

The process to propagate a default route in RIPng is identical to RIPv2 except that an IPv6 default static route must be specified. For example, assume that R1 had an Internet connection from a Serial 0/0/1 interface to IP address 2001:DB8:FEED:1::1/64. To propagate a default route, R1 would have to be configured with:

- A default static route using the **ipv6 route 0::/0 2001:DB8:FEED:1::1** global configuration command.

- The `ipv6 rip domain-name default-information originate` interface configuration mode command. For example, the Serial 0/0/1 interface of R1 would have to be configured with the `ipv6 rip RIP-AS default-information originate` command. This would instruct R1 to be the source of the default route information and propagate the default static route in RIPng updates sent out of the RIPng-enabled interfaces.

Examining the RIPng Configuration (3.3.2.2)

In Figure 3-32, the `show ipv6 protocols` command does not provide the same amount of information as its IPv4 counterpart.

```

R1# show ipv6 protocols
IPv6 Routing Protocol is "connected"
IPv6 Routing Protocol is "ND"
IPv6 Routing Protocol is "rip RIP-AS"
  Interfaces:
    Serial0/0/0
    GigabitEthernet0/0
  Redistribution:
    None
R1#
  
```

Figure 3-32 Verifying RIPng Settings on R1

However, the command does confirm the following parameters:

1. That RIPng routing is configured and running on router R1.
2. The interfaces configured with RIPng.

The `show ipv6 route` command displays the routes installed in the routing table as shown in Figure 3-33. The output confirms that R1 now knows about the highlighted RIPng networks.

```

R1# show ipv6 route
IPv6 Routing Table - default - 8 entries
Codes: C - Connected, L - Local, S - Static, U - Per-user
Static route
  B - BGP, R - RIP, I1 - ISIS L1, I2 - ISIS L2
  IA - ISIS interarea, IS - ISIS summary, D - EIGRP,
  EX - EIGRP external, ND - ND Default,
  NDp - ND Prefix, DCE - Destination, NDr - Redirect,
  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:CAFE:1::/64 [0/0]
  via GigabitEthernet0/0, directly connected
L 2001:DB8:CAFE:1::1/128 [0/0]
  via GigabitEthernet0/0, receive
R 2001:DB8:CAFE:2::/64 [120/2]
  via FE80::FE99:47FF:FE71:78A0, Serial0/0/0
R 2001:DB8:CAFE:3::/64 [120/3]
  via FE80::FE99:47FF:FE71:78A0, Serial0/0/0
C 2001:DB8:CAFE:A001::/64 [0/0]
  via Serial0/0/0, directly connected
L 2001:DB8:CAFE:A001::1/128 [0/0]
  via Serial0/0/0, receive
R 2001:DB8:CAFE:A002::/64 [120/2]
  
```

Figure 3-33 Verifying Routes on R1

Notice that the R2 LAN is advertised as two hops away. This is because there is a difference in the way RIPv2 and RIPng calculate the hop counts. With RIPv2 (and RIPv1), the metric to the R2 LAN would be one hop. This is because the metric (hop count) that is displayed in the IPv4 routing table is the number of hops required to reach the remote network (counting the next-hop router as the first hop). In RIPng, the sending router already considers itself to be one hop away; therefore, R2 advertises its LAN with a metric of 1. When R1 receives the update, it adds another hop count of 1 to the metric. Therefore, R1 considers the R2 LAN to be two hops away. Similarly it considers the R3 LAN to be three hops away.

Appending the `rip` keyword to the command as shown in Figure 3-34 only lists RIPng networks.

```

R1# show ipv6 route rip
IPv6 Routing Table - default - 8 entries
Codes: C - Connected, L - Local, S - Static, U - Per-user
Static route
  B - BGP, R - RIP, I1 - ISIS L1, I2 - ISIS L2
  IA - ISIS interarea, IS - ISIS summary, D - EIGRP,
  EX - EIGRP external, ND - ND Default,
  NDp - ND Prefix, DCE - Destination, Ndr - Redirect,
  O - OSPF Intra, OI - OSPF Inter, OE1 - OSPF ext 1,
  OE2 - OSPF ext 2, ON1 - OSPF NSSA ext 1,
  ON2 - OSPF NSSA ext 2
R 2001:DB8:CAFE:2::/64 [120/2]
  via FE80::FE99:47FF:FE71:78A0, Serial0/0/0
R 2001:DB8:CAFE:3::/64 [120/3]
  via FE80::FE99:47FF:FE71:78A0, Serial0/0/0
R 2001:DB8:CAFE:A002::/64 [120/2]
  via FE80::FE99:47FF:FE71:78A0, Serial0/0/0
R1#

```

Figure 3-34 Verifying RIPng Routes on R1

**Interactive
Graphic**

Activity 3.3.2.2: Verifying RIPng Settings and Routes on R2 and R3

Go to the online course to use the Syntax Checker in the fourth graphic to verify RIPng settings and routes on R2 and R3.

**Packet Tracer
Activity**

Packet Tracer Activity 3.3.2.3: Configuring RIPng

RIPng (RIP Next Generation) is a distance vector routing protocol for routing IPv6 addresses. RIPng is based on RIPv2 and has the same administrative distance and 15-hop limitation. This activity will help you become more familiar with RIPng.



Lab 3.3.2.4: Configuring RIPv2

In this lab, you will complete the following objectives:

- Part 1: Build the Network and Configure Basic Device Settings
- Part 2: Configure and Verify RIPv2 Routing

- Part 3: Configure IPv6 on Devices
- Part 4: Configure and Verify RIPng Routing

Link-State Dynamic Routing (3.4)

Distance vector routing protocols are thought to be simple to understand, whereas link-state routing protocols have the reputation of being very complex, even intimidating. However, link-state routing protocols and concepts are not difficult to understand. In many ways, the link-state process is simpler to understand than distance vector concepts.

Link-State Routing Protocol Operation (3.4.1)

This section describes the characteristics, operations, and functionality of link-state routing protocols. Understanding the operation of link-state routing is critical to enabling, verifying, and troubleshooting these protocols.

Shortest Path First Protocols (3.4.1.1)

Link-state routing protocols are also known as shortest path first protocols and are built around Edsger Dijkstra's shortest path first (SPF) algorithm. The SPF algorithm is discussed in more detail in a later section.

The IPv4 link-state routing protocols are shown Figure 3-35:

- Open Shortest Path First (OSPF)
- Intermediate System-to-Intermediate System (IS-IS)

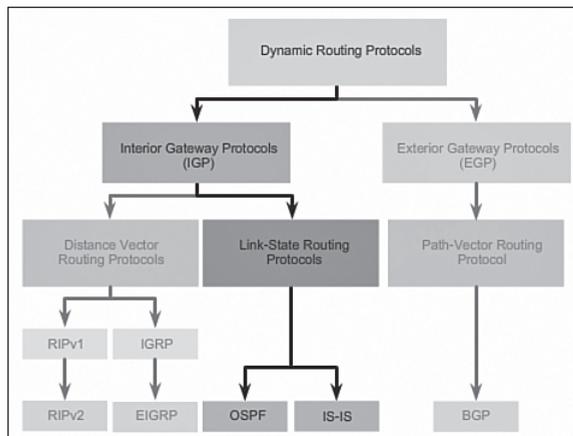


Figure 3-35 Link-State Routing Protocols

Link-state routing protocols have the reputation of being much more complex than their distance vector counterparts. However, the basic functionality and configuration of link-state routing protocols is equally straightforward.

Just like RIP and EIGRP, basic OSPF operations can be configured using the:

- `router ospf process-id` global configuration command
- `network` command to advertise networks

Dijkstra's Algorithm (3.4.1.2)

All link-state routing protocols apply *Dijkstra's algorithm* to calculate the best path route. The algorithm is commonly referred to as the *shortest path first (SPF)* algorithm. This algorithm uses accumulated costs along each path, from source to destination, to determine the total cost of a route.

In Figure 3-36, each path is labeled with an arbitrary value for cost.

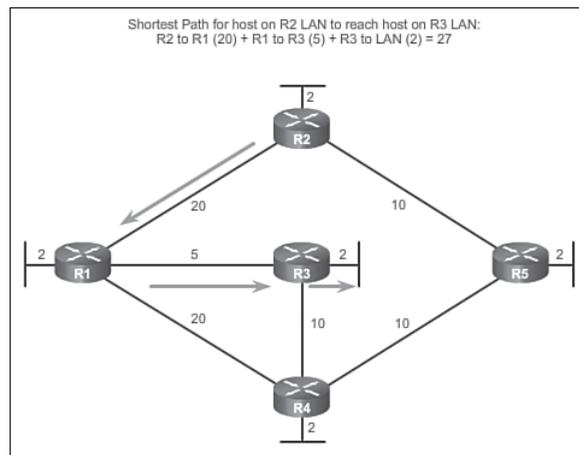


Figure 3-36 Dijkstra's Shortest Path First Algorithm

The cost of the shortest path for R2 to send packets to the LAN attached to R3 is 27. Specifically, the cost is R2 to R1 (20) plus R1 to R3 (5) plus R3 to LAN (2). Each router determines its own cost to each destination in the topology. In other words, each router calculates the SPF algorithm and determines the cost from its own perspective.

Note

The focus of this section is on cost, which is determined by the SPF tree. For this reason, the graphics throughout this section show the connections of the SPF tree, not the topology. All links are represented with a solid black line.

SPF Example (3.4.1.3)

The table in Figure 3-37 displays the shortest path and the accumulated cost to reach the identified destination networks from the perspective of R1.

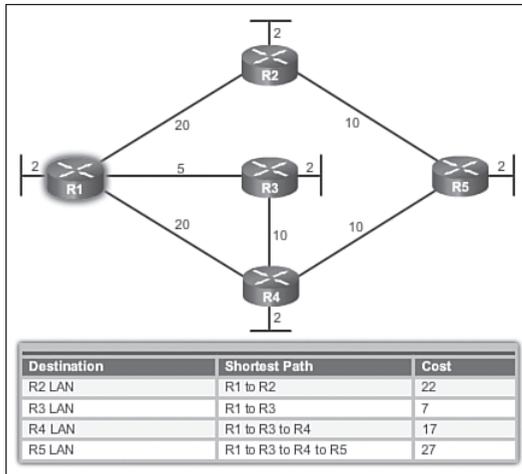


Figure 3-37 R1 SPF Tree

The shortest path is not necessarily the path with the least number of hops. For example, look at the path to the R5 LAN. It might be assumed that R1 would send directly to R4 instead of to R3. However, the cost to reach R4 directly (22) is higher than the cost to reach R4 through R3 (17).

Observe the shortest path for each router to reach each of the LANs, as shown in Tables 3-8 through 3-11.

Table 3-8 R2 SPF Tree

Destination	Shortest Path	Cost
R1 LAN	R2 to R1	22
R3 LAN	R2 to R1 to R3	27
R4 LAN	R2 to R5 to R4	22
R5 LAN	R2 to R5	12

Table 3-9 R3 SPF Tree

Destination	Shortest Path	Cost
R1 LAN	R3 to R1	7
R2 LAN	R3 to R1 to R2	27
R4 LAN	R3 to R4	12
R5 LAN	R3 to R4 to R5	22

Table 3-10 R4 SPF Tree

Destination	Shortest Path	Cost
R1 LAN	R4 to R3 to R1	17
R2 LAN	R4 to R5 to R2	22
R3 LAN	R4 to R3	12
R5 LAN	R4 to R5	12

Table 3-11 R5 SPF Tree

Destination	Shortest Path	Cost
R1 LAN	R5 to R4 to R3 to R1	27
R2 LAN	R5 to R2	12
R3 LAN	R5 to R4 to R3	22
R4 LAN	R5 to R4	12

Link-State Updates (3.4.2)

Link-state updates (LSUs) are the packets used for OSPF routing updates. This section discusses how OSPF exchanges LSUs to discover the best routes.

Link-State Routing Process (3.4.2.1)

So exactly how does a link-state routing protocol work? With link-state routing protocols, a link is an interface on a router. Information about the state of those links is known as link-states.

All routers in an OSPF area will complete the following generic link-state routing process to reach a state of convergence:

1. Each router learns about its own links and its own directly connected networks. This is done by detecting that an interface is in the up state.
2. Each router is responsible for meeting its neighbors on directly connected networks. Link-state routers do this by exchanging Hello packets with other link-state routers on directly connected networks.
3. Each router builds a *link-state packet (LSP)* containing the state of each directly connected link. This is done by recording all the pertinent information about each neighbor, including neighbor ID, link type, and bandwidth.
4. Each router floods the LSP to all neighbors. Those neighbors store all LSPs received in a database. They then flood the LSPs to their neighbors until all routers in the area have received the LSPs. Each router stores a copy of each LSP received from its neighbors in a local database.
5. Each router uses the database to construct a complete map of the topology and computes the best path to each destination network. Like having a road map, the router now has a complete map of all destinations in the topology and the routes to reach them. The SPF algorithm is used to construct the map of the topology and to determine the best path to each network.

Note

This process is the same for both OSPF for IPv4 and OSPF for IPv6. The examples in this section refer to OSPF for IPv4.

Link and Link-State (3.4.2.2)

The first step in the link-state routing process is that each router learns about its own links, its own directly connected networks. When a router interface is configured with an IP address and subnet mask, the interface becomes part of that network.

Refer to the topology in Figure 3-38. For purposes of this discussion, assume that R1 was previously configured and had full connectivity to all neighbors. However, R1 lost power briefly and had to restart.

During boot up R1 loads the saved startup configuration file. As the previously configured interfaces become active, R1 learns about its own directly connected networks. Regardless of the routing protocols used, these directly connected networks are now entries in the routing table.

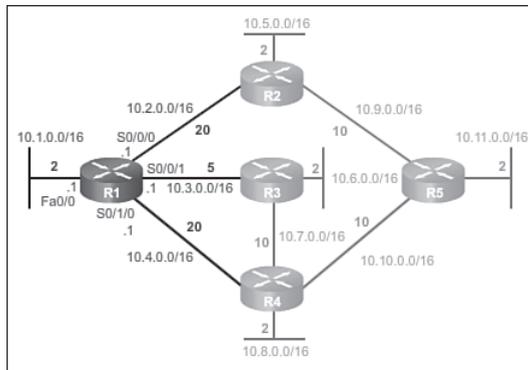


Figure 3-38 R1 Links

As with distance vector protocols and static routes, the interface must be properly configured with an IPv4 address and subnet mask, and the link must be in the up state before the link-state routing protocol can learn about a link. Also, like distance vector protocols, the interface must be included in one of the **network** router configuration statements before it can participate in the link-state routing process.

Figure 3-38 shows R1 linked to four directly connected networks:

- FastEthernet 0/0: 10.1.0.0/16
- Serial 0/0/0: 10.2.0.0/16
- Serial 0/0/1: 10.3.0.0/16
- Serial 0/1/0: 10.4.0.0/16

As shown in Figures 3-39 through 3-42, the link-state information includes:

- The interface's IPv4 address and subnet mask
- The type of network, such as Ethernet (broadcast) or Serial point-to-point link
- The cost of that link
- Any neighbor routers on that link

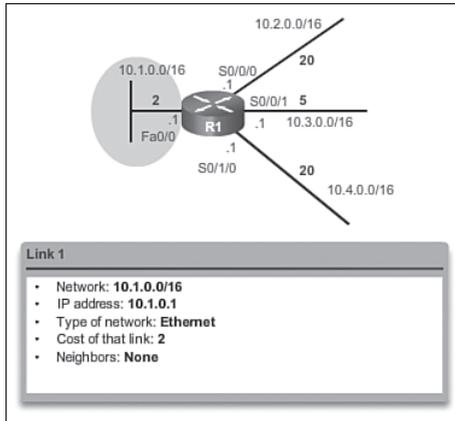


Figure 3-39 Link-State of Interface Fa0/0

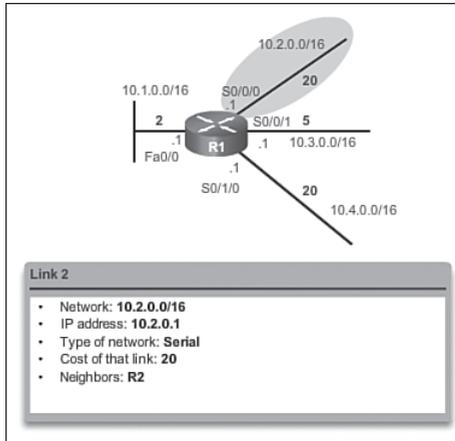


Figure 3-40 Link-State of Interface S0/0/0

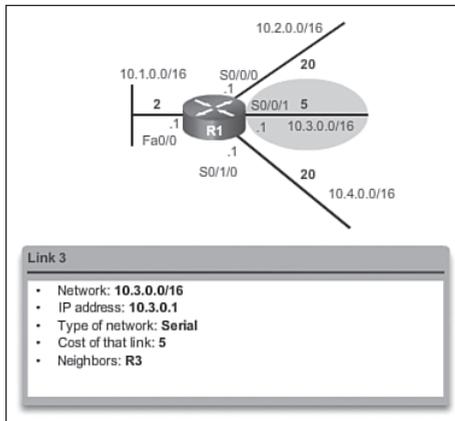


Figure 3-41 Link-State of Interface S0/0/1

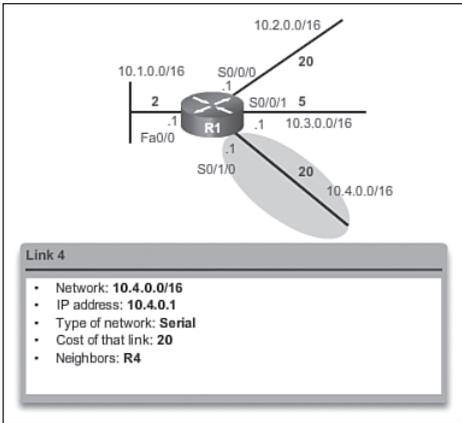


Figure 3-42 Link-State of Interface S0/1/0

Note

Cisco’s implementation of OSPF specifies the OSPF routing metric as the cost of the link based on the bandwidth of the outgoing interface. For the purposes of this chapter, we are using arbitrary cost values to simplify the demonstration.

Say Hello (3.4.2.3)

The second step in the link-state routing process is that each router is responsible for meeting its neighbors on directly connected networks.

Routers with link-state routing protocols use a Hello protocol to discover any neighbors on their links. A neighbor is any other router that is enabled with the same link-state routing protocol.

In Figure 3-43, R1 sends Hello packets out its links (interfaces) to discover if there are any neighbors.

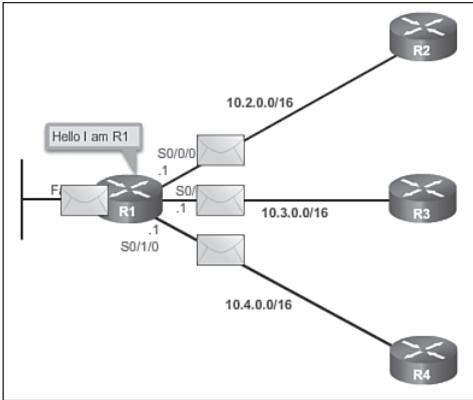


Figure 3-43 R1 Sends Hello Packets

In Figure 3-44, R2, R3, and R4 reply to the Hello packet with their own Hello packets because these routers are configured with the same link-state routing protocol. There are no neighbors out the FastEthernet 0/0 interface. Because R1 does not receive a Hello on this interface, it does not continue with the link-state routing process steps for the FastEthernet 0/0 link.

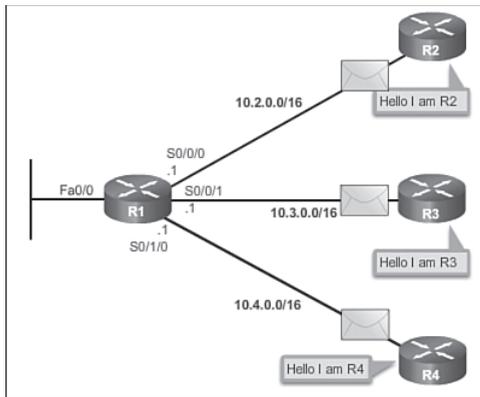


Figure 3-44 R2, R3, and R4 Reply with Hello Packets

Video

Video 3.4.2.3: Neighbor Discovery—Hello Packets

Go to the online course and play the animation to view the link-state neighbor discovery process with Hello packets.

When two link-state routers learn that they are neighbors, they form an adjacency. These small Hello packets continue to be exchanged between two adjacent neighbors and serve as a keepalive function to monitor the state of the neighbor. If a router stops receiving Hello packets from a neighbor, that neighbor is considered unreachable and the adjacency is broken.

Building the Link-State Packet (3.4.2.4)

The third step in the link-state routing process is that each router builds an LSP containing the state of each directly connected link.

After a router has established its adjacencies, it can build its LSPs that contain the link-state information about its links. A simplified version of the LSP from R1 displayed in Figure 3-45 would contain the following:

1. R1; Ethernet network 10.1.0.0/16; Cost 2
2. R1 -> R2; Serial point-to-point network; 10.2.0.0/16; Cost 20

- 3. R1 -> R3; Serial point-to-point network; 10.3.0.0/16; Cost 5
- 4. R1 -> R4; Serial point-to-point network; 10.4.0.0/16; Cost 20

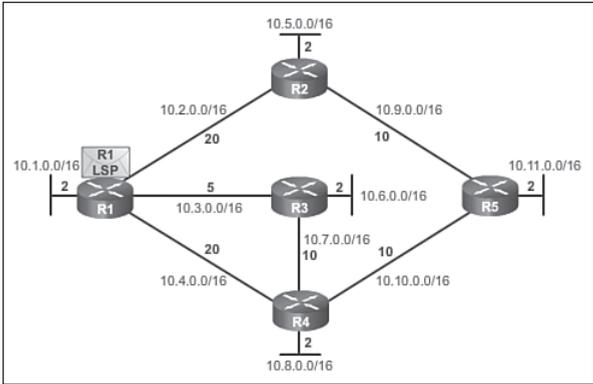


Figure 3-45 Building the LSP

Flooding the LSP (3.4.2.5)

The fourth step in the link-state routing process is that each router floods the LSP to all neighbors, who then store all LSPs received in a database.

Each router floods its link-state information to all other link-state routers in the routing area as shown in Figure 3-46.

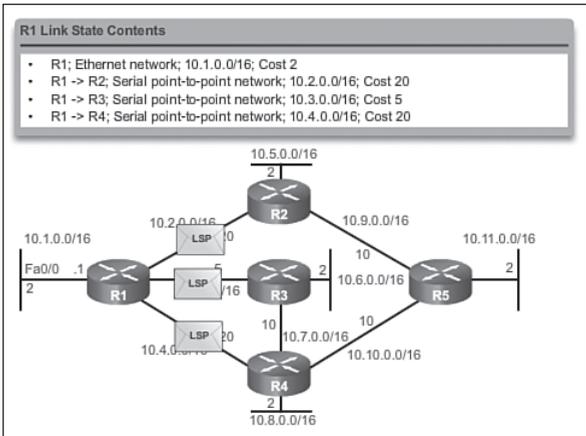


Figure 3-46 R1 Floods Its LSP

Whenever a router receives an LSP from a neighboring router, it immediately sends that LSP out all other interfaces except the interface that received the LSP. This process creates a flooding effect of LSPs from all routers throughout the routing area.

Video**Video 3.4.2.5: Routers Route Packets**

Go to the online course and play the animation to view the LSP flooding.

In the animation, the LSPs are flooded almost immediately after being received without any intermediate calculations. Link-state routing protocols calculate the SPF algorithm after the flooding is complete. As a result, link-state routing protocols reach convergence very quickly.

Remember that LSPs do not need to be sent periodically. An LSP only needs to be sent:

- During initial startup of the routing protocol process on that router (e.g., router restart)
- Whenever there is a change in the topology (e.g., a link going down or coming up, a neighbor adjacency being established or broken)

In addition to the link-state information, other information is included in the LSP, such as sequence numbers and aging information, to help manage the flooding process. This information is used by each router to determine if it has already received the LSP from another router or if the LSP has newer information than what is already contained in the link-state database. This process allows a router to keep only the most current information in its link-state database.

Building the Link-State Database (3.4.2.6)

The final step in the link-state routing process is that each router uses the database to construct a complete map of the topology and computes the best path to each destination network.

Eventually, all routers receive an LSP from every other link-state router in the routing area. These LSPs are stored in the link-state database.

Table 3-12 displays the link-state database content of R1.

Table 3-12 Link-State Database

R1 Link-states:

Connected to network 10.1.0.0/16, cost = 2
Connected to R2 on network 10.2.0.0/16, cost = 20
Connected to R3 on network 10.2.0.0/16, cost = 5
Connected to R4 on network 10.3.0.0/16, cost = 20

R2 Link-states:

Connected to network 10.5.0.0/16, cost = 2

Connected to R1 on network 10.2.0.0/16, cost = 20

Connected to R5 on network 10.9.0.0/16, cost = 10

R3 Link-states:

Connected to network 10.6.0.0/16, cost = 2

Connected to R1 on network 10.3.0.0/16, cost = 5

Connected to R4 on network 10.7.0.0/16, cost = 10

R4 Link-states:

Connected to network 10.8.0.0/16, cost = 2

Connected to R1 on network 10.4.0.0/16, cost = 20

Connected to R3 on network 10.7.0.0/16, cost = 10

Connected to R5 on network 10.10.0.0/16, cost = 10

R5 Link-states:

Connected to network 10.11.0.0/16, cost = 2

Connected to R2 on network 10.9.0.0/16, cost = 10

Connected to R4 on network 10.10.0.0/16, cost = 10

As a result of the flooding process, R1 has learned the link-state information for each router in its routing area. Notice that R1 also includes its own link-state information in the link-state database.

With a complete link-state database, R1 can now use the database and the shortest path first (SPF) algorithm to calculate the preferred path or shortest path to each network, resulting in the SPF tree.

Building the SPF Tree (3.4.2.7)

Each router in the routing area uses the link-state database and SPF algorithm to construct the *SPF tree*.

For example, using the link-state information from all other routers, R1 can now begin to construct an SPF tree of the network. To begin, the SPF algorithm interprets each router's LSP to identify networks and associated costs.

The SPF algorithm then calculates the shortest paths to reach each individual network, resulting in the SPF tree as shown in Figure 3-47. R1 now has a complete topology view of the link-state area.

Note

The entire process can be viewed in the online course on page 3.4.2.7 in Figures 1 through 6.

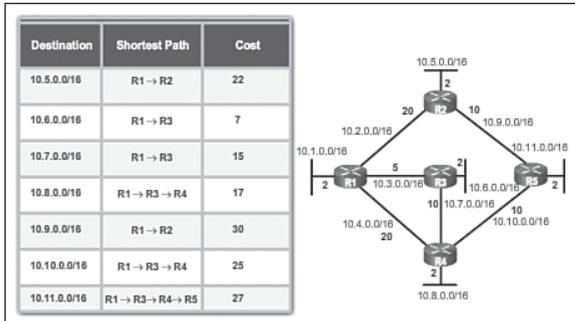


Figure 3-47 Resulting SPF Tree of R1

Each router constructs its own SPF tree independently from all other routers. To ensure proper routing, the link-state databases used to construct those trees must be identical on all routers.

Adding OSPF Routes to the Routing Table (3.4.2.8)

Using the shortest path information determined by the SPF algorithm, these paths can now be added to the routing table. Figure 3-48 shows the routes that have now been added to R1's IPv4 routing table.

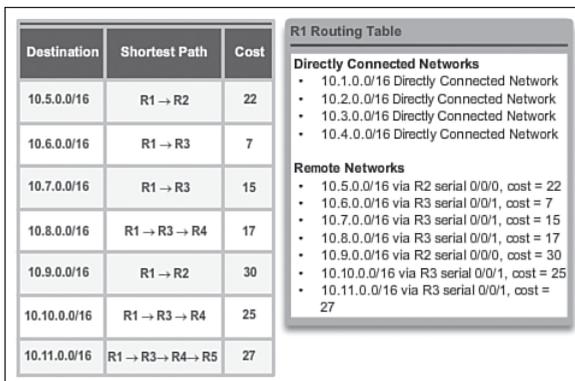


Figure 3-48 Populate the Routing Table

The routing table also includes all directly connected networks and routes from any other sources, such as static routes. Packets are now forwarded according to these entries in the routing table.

**Interactive
Graphic****Activity 3.4.2.9: Building the Link-State Database and SPF Tree**

Go to the online course to perform this practice activity.

Why Use Link-State Routing Protocols? (3.4.3)

This section discusses the advantages of using link-state routing protocols and compares the two types of link-state routing protocols.

Why Use Link-State Protocols? (3.4.3.1)

There are several advantages of link-state routing protocols compared to distance vector routing protocols.

- **Builds a topological map:** Link-state routing protocols create a topological map, or SPF tree of the network topology. Because link-state routing protocols exchange link-states, the SPF algorithm can build an SPF tree of the network. Using the SPF tree, each router can independently determine the shortest path to every network.
- **Fast convergence:** When receiving an LSP, link-state routing protocols immediately flood the LSP out all interfaces except for the interface from which the LSP was received. In contrast, RIP needs to process each routing update and update its routing table before flooding the routing update out other interfaces.
- **Event-driven updates:** After the initial flooding of LSPs, link-state routing protocols only send out an LSP when there is a change in the topology. The LSP contains only the information regarding the affected link. Unlike some distance vector routing protocols, link-state routing protocols do not send periodic updates.
- **Hierarchical design:** Link-state routing protocols use the concept of areas. Multiple areas create a hierarchical design to networks, allowing for better route aggregation (summarization) and the isolation of routing issues within an area.

Link-state protocols also have a few disadvantages compared to distance vector routing protocols:

- **Memory requirements:** Link-state protocols require additional memory to create and maintain the link-state database and SPF tree.
- **Processing requirements:** Link-state protocols can also require more CPU processing than distance vector routing protocols. The SPF algorithm requires more CPU time than distance vector algorithms such as Bellman-Ford, because link-state protocols build a complete map of the topology.
- **Bandwidth requirements:** The flooding of link-state packets can adversely affect the available bandwidth on a network. This should only occur during initial start-up of routers, but can also be an issue on unstable networks.

Link-State Protocols Support Multiple Areas (3.4.3.2)

Modern link-state routing protocols are designed to minimize the effects on memory, CPU, and bandwidth. The use and configuration of multiple areas can reduce the size of the link-state databases. Multiple areas can also limit the amount of link-state information flooding in a routing domain and send LSPs only to those routers that need them. When there is a change in the topology, only those routers in the affected area receive the LSP and run the SPF algorithm. This can help isolate an unstable link to a specific area in the routing domain.

For example, in Figure 3-49, there are three separate routing domains: area 1, area 0, and area 51.

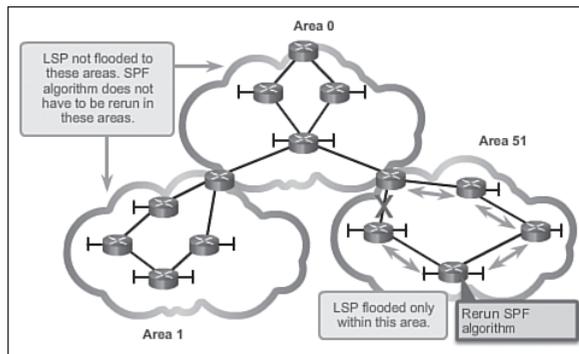


Figure 3-49 Create Areas to Minimize Router Resource Usage

If a network in area 51 goes down, the LSP with the information about this downed link is only flooded to other routers in that area. Only those routers in area 51 need to update their link-state databases, rerun the SPF algorithm, create a new SPF tree, and update their routing tables. Routers in other areas learn that this route is down, but this is done with a type of LSP that does not cause them to rerun their SPF algorithm. Routers in other areas can update their routing tables directly.

Protocols that Use Link-State (3.4.3.3)

There are only two link-state routing protocols, OSPF and IS-IS.

Open Shortest Path First (OSPF) is the most popular implementation. It was designed by the Internet Engineering Task Force (IETF) OSPF Working Group. The development of OSPF began in 1987 and there are two current versions in use:

- **OSPFv2:** OSPF for IPv4 networks (RFC 1247 and RFC 2328)
- **OSPFv3:** OSPF for IPv6 networks (RFC 2740)

Note

With the OSPFv3 Address Families feature, OSPFv3 includes support for both IPv4 and IPv6.

IS-IS was designed by International Organization for Standardization (ISO) and is described in ISO 10589. The first incarnation of this routing protocol was developed at Digital Equipment Corporation (DEC) and is known as DECnet Phase V. Radia Perlman was the chief designer of the IS-IS routing protocol.

IS-IS was originally designed for the OSI protocol suite and not the TCP/IP protocol suite. Later, Integrated IS-IS, or Dual IS-IS, included support for IP networks. Although IS-IS has been known as the routing protocol used mainly by ISPs and carriers, more enterprise networks are beginning to use IS-IS.

OSPF and IS-IS share many similarities and also have many differences. There are many pro-OSPF and pro-IS-IS factions who discuss and debate the advantages of one routing protocol over the other. Both routing protocols provide the necessary routing functionality.

The Routing Table (3.5)

As a network administrator, it is important to know the routing table in depth when troubleshooting network issues. Understanding the structure and lookup process of the routing table will help you diagnose any routing table issue, regardless of your level of familiarity with a particular routing protocol. For example, you might encounter a situation in which the routing table has all of the routes you would expect to see, but packet forwarding is not performing as expected. Knowing how to step through the lookup process of a destination IP address for a packet will enable you to determine whether the packet is being forwarded as expected, if and why the packet is being sent elsewhere, or whether the packet has been discarded.

Parts of an IPv4 Route Entry (3.5.1)

A routing table consists of directly connected networks and routes learned statically or dynamically. This section examines these two types of routing table entries.

Routing Table Entries (3.5.1.1)

The topology displayed in Figure 3-50 is used as the reference topology for this section.

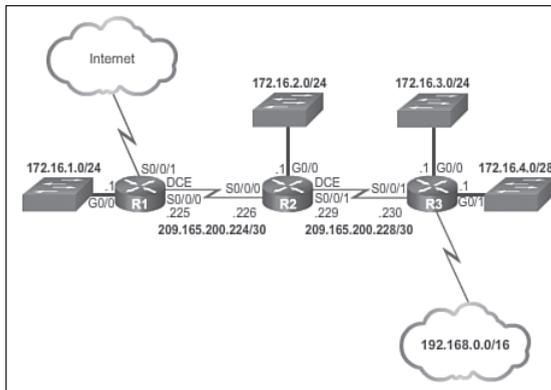


Figure 3-50 Reference Topology

Notice that in the topology:

- R1 is the edge router that connects to the Internet. Therefore, it is propagating a default static route to R2 and R3.
- R1, R2, and R3 contain discontinuous networks separated by another classful network.
- R3 is also introducing a 192.168.0.0/16 supernet route.

Figure 3-51 displays the IPv4 routing table of R1 with directly connected, static, and dynamic routes.

```

R1# show ip route | begin Gateway
Gateway of last resort is 209.165.200.234 to network 0.0.0.0

S* 0.0.0.0/0 [1/0] via 209.165.200.234, Serial0/0/1
    is directly connected, Serial0/0/1
172.16.0.0/16 is variably subnetted, 5 subnets, 3 masks
C   172.16.1.0/24 is directly connected, gigabitEthernet0/0
L   172.16.1.1/32 is directly connected, gigabitEthernet0/0
R   172.16.2.0/24 [120/1] via 209.165.200.226, 00:00:12, Serial0/0/0
R   172.16.3.0/24 [120/2] via 209.165.200.226, 00:00:12, Serial0/0/0
R   172.16.4.0/28 [120/2] via 209.165.200.226, 00:00:12, Serial0/0/0
R   192.168.0.0/16 [120/2] via 209.165.200.226, 00:00:03, Serial0/0/0
209.165.200.0/24 is variably subnetted, 5 subnets, 2 masks
C   209.165.200.224/30 is directly connected, Serial0/0/0
L   209.165.200.225/32 is directly connected, Serial0/0/0
R   209.165.200.228/30 [120/1] via 209.165.200.226, 00:00:12,
    Serial0/0/0
C   209.165.200.232/30 is directly connected, Serial0/0/1
L   209.165.200.233/30 is directly connected, Serial0/0/1
R1#

```

Figure 3-51 Routing Table of R1

Note

The routing table hierarchy in Cisco IOS was originally implemented with the classful routing scheme. Although the routing table incorporates both classful and classless addressing, the overall structure is still built around this classful scheme.

Directly Connected Entries (3.5.1.2)

As highlighted in Figure 3-52, the routing table of R1 contains three directly connected networks. Notice that two routing table entries are automatically created when an active router interface is configured with an IP address and subnet mask.

```
R1# show ip route | begin Gateway
Gateway of last resort is 209.165.200.234 to network 0.0.0.0

S* 0.0.0.0/0 [1/0] via 209.165.200.234, Serial0/0/1
    is directly connected, Serial0/0/1
    172.16.0.0/16 is variably subnetted, 5 subnets, 3 masks
C   172.16.1.0/24 is directly connected, GigabitEthernet0/0
L   172.16.1.1/32 is directly connected, GigabitEthernet0/0
R   172.16.2.0/24 [120/1] via 209.165.200.226,00:00:12, Serial0/0/0
R   172.16.3.0/24 [120/2] via 209.165.200.226, 00:00:12, Serial0/0/0
R   172.16.4.0/28 [120/2] via 209.165.200.226, 00:00:12, Serial0/0/0
R   192.168.0.0/16 [120/2] via 209.165.200.226, 00:00:03, Serial0/0/0
209.165.200.0/24 is variably subnetted, 5 subnets, 2 masks
C   209.165.200.224/30 is directly connected, Serial0/0/0
L   209.165.200.225/32 is directly connected, Serial0/0/0
R   209.165.200.228/30 [120/1] via 209.165.200.226, 00:00:12, Serial0/0/0
C   209.165.200.232/30 is directly connected, Serial0/0/1
L   209.165.200.233/32 is directly connected, Serial0/0/1
R1#
```

Figure 3-52 Directly Connected Interfaces of R1

Figure 3-53 displays one of the routing table entries on R1 for the directly connected network 172.16.1.0. These entries were automatically added to the routing table when the GigabitEthernet 0/0 interface was configured and activated.

Route Source	Destination Network	Outgoing Interface
C	172.16.1.0/24 is directly connected,	GigabitEthernet0/0
L	172.16.1.1/32 is directly connected,	GigabitEthernet0/0

Legend

- Identifies how the network was learned by the router.
- Identifies the destination network and how it is connected.
- Identifies the interface on the router connected to the destination network.

Figure 3-53 Directly Connected Routes of R1

The entries contain the following information:

- **Route source:** Identifies how the route was learned. Directly connected interfaces have two route source codes. C identifies a directly connected network. Directly connected networks are automatically created whenever an interface is configured with an IP address and activated. L identifies that this is a local route. Local routes are automatically created whenever an interface is configured with an IP address and activated.

- **Destination network:** The address of the remote network and how that network is connected.
- **Outgoing interface:** Identifies the exit interface to use when forwarding packets to the destination network.

Note

Local routing table entries did not appear in routing tables prior to IOS release 15.

A router typically has multiple interfaces configured. The routing table stores information about both directly connected and remote routes. As with directly connected networks, the route source identifies how the route was learned. For instance, common codes for remote networks include:

- **S:** Identifies that the route was manually created by an administrator to reach a specific network. This is known as a static route.
- **D:** Identifies that the route was learned dynamically from another router using the EIGRP routing protocol.
- **O:** Identifies that the route was learned dynamically from another router using the OSPF routing protocol.
- **R:** Identifies that the route was learned dynamically from another router using the RIP routing protocol.

Remote Network Entries (3.5.1.3)

Figure 3-54 displays an IPv4 routing table entry on R1 for the route to remote network 172.16.4.0 on R3.

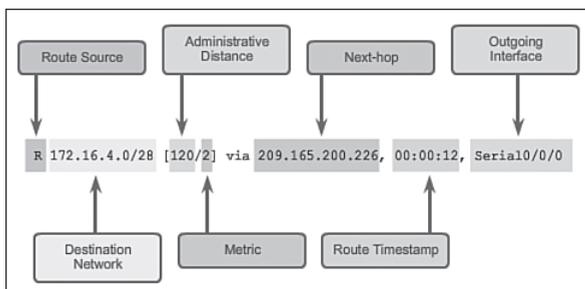


Figure 3-54 Remote Network Route Entry on R1

The entry identifies the following information:

- **Route source:** Identifies how the route was learned.
- **Destination network:** Identifies the address of the remote network.

- **Administrative distance:** Identifies the trustworthiness of the route source.
- **Metric:** Identifies the value assigned to reach the remote network. Lower values indicate preferred routes.
- **Next hop:** Identifies the IPv4 address of the next router to forward the packet to.
- **Route timestamp:** Identifies from when the route was last heard.
- **Outgoing interface:** Identifies the exit interface to use to forward a packet toward the final destination.

**Interactive
Graphic**

Activity 3.5.1.4: Identify Parts of an IPv4 Routing Table Entry

Go to the online course to perform this practice activity.

Dynamically Learned IPv4 Routes (3.5.2)

The structure or format of the routing table might seem obvious until you take a closer look. Understanding the structure of the routing table will help you verify and troubleshoot routing issues because you will understand the routing table lookup process.

Routing Table Terms (3.5.2.1)

A dynamically built routing table provides a great deal of information, as shown in Figure 3-55. Therefore, it is crucial to understand the output generated by the routing table. Special terms are applied when discussing the contents of a routing table.

```

R1#show ip route | begin Gateway
Gateway of last resort is 209.165.200.234 to network 0.0.0.0

S*   0.0.0.0/0 [1/0] via 209.165.200.234, Serial0/0/1
      is directly connected, Serial0/0/1
      172.16.0.0/16 is variably subnetted, 5 subnets, 3 masks
C     172.16.1.0/24 is directly connected, GigabitEthernet0/0
L     172.16.1.1/32 is directly connected, GigabitEthernet0/0
R     172.16.2.0/24 [120/1] via 209.165.200.226, 00:00:12,
      Serial0/0/0
R     172.16.3.0/24 [120/2] via 209.165.200.226, 00:00:12,
      Serial0/0/0
R     172.16.4.0/28 [120/2] via 209.165.200.226, 00:00:12,
      Serial0/0/0
R     192.168.0.0/16 [120/2] via 209.165.200.226, 00:00:03,
      Serial0/0/0
      209.165.200.0/24 is variably subnetted, 5 subnets, 2 masks
C     209.165.200.224/30 is directly connected, Serial0/0/0
L     209.165.200.225/32 is directly connected, Serial0/0/0
R     209.165.200.228/30 [120/1] via 209.165.200.226, 00:00:12,
      Serial0/0/0
C     209.165.200.232/30 is directly connected, Serial0/0/1
L     209.165.200.233/32 is directly connected, Serial0/0/1
R1#

```

Figure 3-55 Routing Table of R1

The Cisco IP routing table is not a flat database. The routing table is actually a hierarchical structure that is used to speed up the lookup process when locating routes and forwarding packets. Within this structure, the hierarchy includes several levels.

Routes are discussed in terms of:

- Ultimate route
- Level 1 route
- Level 1 parent route
- Level 2 child routes

Ultimate Route (3.5.2.2)

An *ultimate route* is a routing table entry that contains either a next-hop IPv4 address or an exit interface. Directly connected, dynamically learned, and local routes are ultimate routes.

In Figure 3-56, the highlighted areas are examples of ultimate routes. Notice that all of these routes specify either a next-hop IPv4 address or an exit interface.

```

R1# show ip route | begin Gateway
Gateway of last resort is 209.165.200.234 to network 0.0.0.0
S* 0.0.0.0/0 [1/0] via 209.165.200.234, Serial0/0/1
   is directly connected, Serial0/0/1
172.16.0.0/16 is variably subnetted, 5 subnets, 3 masks
C 172.16.1.0/24 is directly connected, GigabitEthernet0/0
L 172.16.1.1/32 is directly connected, GigabitEthernet0/0
R 172.16.2.0/24 [120/1] via 209.165.200.226, 00:00:12,
  Serial0/0/0
R 172.16.3.0/24 [120/2] via 209.165.200.226, 00:00:12,
  Serial0/0/0
R 172.16.4.0/28 [120/2] via 209.165.200.226, 00:00:12,
  Serial0/0/0
R 192.168.0.0/16 [120/2] via 209.165.200.226, 00:00:03,
  Serial0/0/0
209.165.200.0/24 is variably subnetted, 5 subnets, 2 masks
C 209.165.200.224/30 is directly connected, Serial0/0/0
L 209.165.200.225/32 is directly connected, Serial0/0/0
R 209.165.200.228/30 [120/1] via 209.165.200.226, 00:00:12,
  Serial0/0/0
C 209.165.200.232/30 is directly connected, Serial0/0/1
L 209.165.200.233/32 is directly connected, Serial0/0/1
R1#

```

Figure 3-56 Ultimate Routes of R1

Level 1 Route (3.5.2.3)

A *level 1 route* is a route with a subnet mask equal to or less than the classful mask of the network address. Therefore, a level 1 route can be a:

- **Network route:** A network route has a subnet mask equal to that of the classful mask.

- **Supernet route:** A supernet route is a network address with a subnet mask less than the classful mask, for example, a summary address.
- **Default route:** A default route is a static route with the address 0.0.0.0/0.

The source of the level 1 route can be a directly connected network, static route, or a dynamic routing protocol.

Figure 3-57 highlights how level 1 routes are also ultimate routes.

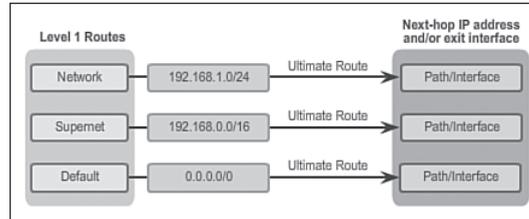


Figure 3-57 Sources of Level 1 Routes

Figure 3-58 highlights level 1 routes.

```
R1# show ip route | begin Gateway
Gateway of last resort is 209.165.200.234 to network
0.0.0.0
S* 0.0.0.0/0 [1/0] via 209.165.200.234, Serial0/0/1
    is directly connected, Serial0/0/1
    172.16.0.0/16 is variably subnetted, 5 subnets, 3
    masks
    C 172.16.1.0/24 is directly connected,
    GigabitEthernet0/0
    L 172.16.1.1/32 is directly connected,
    GigabitEthernet0/0
    R 172.16.2.0/24 [120/1] via 209.165.200.226,
    00:00:12, Serial0/0/0
    R 172.16.3.0/24 [120/2] via 209.165.200.226,
    00:00:12, Serial0/0/0
    R 172.16.4.0/28 [120/2] via 209.165.200.226,
    00:00:12, Serial0/0/0
    R 192.168.0.0/16 [120/2] via 209.165.200.226, 00:00:03,
    Serial0/0/0
    209.165.200.0/24 is variably subnetted, 5 subnets, 2
    masks
    C 209.165.200.224/30 is directly connected,
    Serial0/0/0
```

Figure 3-58 Example of Level 1 Routes

Level 1 Parent Route (3.5.2.4)

As illustrated in Figure 3-59, a *level 1 parent route* is a level 1 network route that is subnetted. A parent route can never be an ultimate route.

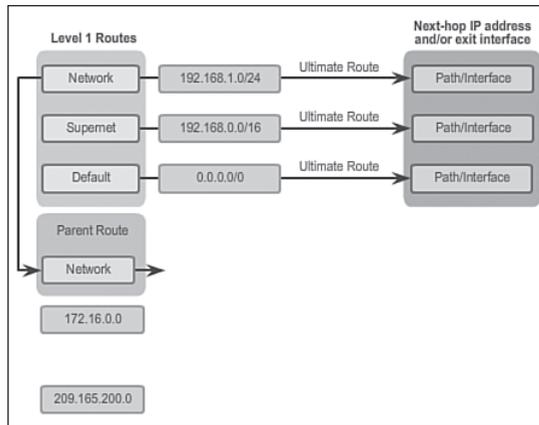


Figure 3-59 Level 1 Parent Route

Figure 3-60 highlights the level 1 parent routes in the routing table of R1. The routing table basically provides a heading for the specific subnets it contains. Each entry displays the classful network address, the number of subnets, and the number of different subnet masks that the classful address has been subdivided into.

```

R1# show ip route | begin Gateway
Gateway of last resort is 209.165.200.234 to network
0.0.0.0

S* 0.0.0.0/0 [1/0] via 209.165.200.234, Serial0/0/1
    is directly connected, Serial0/0/1
    172.16.0.0/16 is variably subnetted, 5 subnets, 3
    masks
    C 172.16.1.0/24 is directly connected,
    GigabitEthernet0/0
    L 172.16.1.1/32 is directly connected,
    GigabitEthernet0/0
    R 172.16.2.0/24 [120/1] via 209.165.200.226,
    00:00:12, Serial0/0/0
    R 172.16.3.0/24 [120/2] via 209.165.200.226,
    00:00:12, Serial0/0/0
    R 172.16.4.0/28 [120/2] via 209.165.200.226,
    00:00:12, Serial0/0/0
    R 192.168.0.0/16 [120/2] via 209.165.200.226, 00:00:03,
    Serial0/0/0
    209.165.200.0/24 is variably subnetted, 5 subnets, 2
    masks
    C 209.165.200.224/30 is directly connected,
    Serial0/0/0
  
```

Figure 3-60 Level 1 Parent Routes of R1

Level 2 Child Route (3.5.2.5)

A *level 2 child route* is a route that is a subnet of a classful network address. As illustrated in Figure 3-61, a level 1 parent route is a level 1 network route that is subnetted.

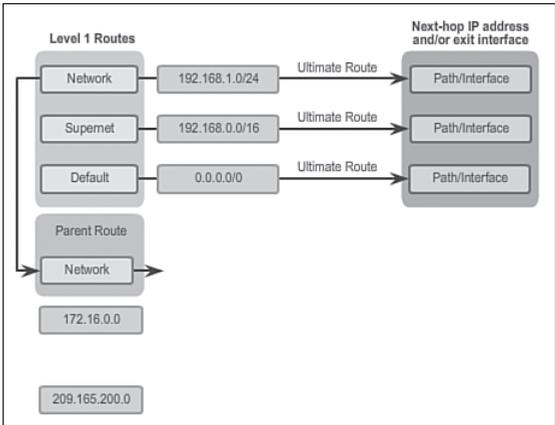


Figure 3-61 Level 2 Child Routes

A level 1 parent route contains level 2 child routes, as shown in Figure 3-62.

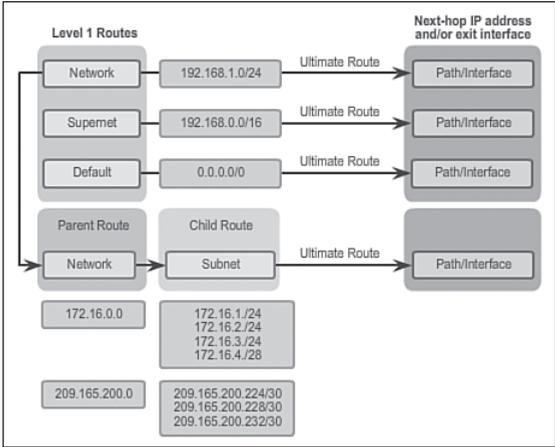


Figure 3-62 Child Routes Are Ultimate Routes

Like a level 1 route, the source of a level 2 route can be a directly connected network, a static route, or a dynamically learned route. Level 2 child routes are also ultimate routes.

Note
 The routing table hierarchy in Cisco IOS has a classful routing scheme. A level 1 parent route is the classful network address of the subnet route. This is the case even if a classless routing protocol is the source of the subnet route.

Figure 3-63 highlights the level 2 child routes in the routing table of R1.

```
R1#show ip route | begin Gateway
Gateway of last resort is 209.165.200.234 to network
0.0.0.0

S* 0.0.0.0/0 [1/0] via 209.165.200.234, Serial0/0/1
    is directly connected, Serial0/0/1
    172.16.0.0/16 is variably subnetted, 5 subnets, 3
    masks
    C 172.16.1.0/24 is directly connected,
    GigabitEthernet0/0
    L 172.16.1.1/32 is directly connected,
    GigabitEthernet0/0
    R 172.16.2.0/24 [120/1] via 209.165.200.226,
    00:00:12, Serial0/0/0
    R 172.16.3.0/24 [120/2] via 209.165.200.226,
    00:00:12, Serial0/0/0
    R 172.16.4.0/28 [120/2] via 209.165.200.226,
    00:00:12, Serial0/0/0
    R 192.168.0.0/16 [120/2] via 209.165.200.226, 00:00:03,
    Serial0/0/0
    209.165.200.0/24 is variably subnetted, 5 subnets, 2
    masks
    C 209.165.200.224/30 is directly connected,
    Serial0/0/0
```

Figure 3-63 Example of Level 2 Child Routes

Note

The entire output in Figure 3-63 can be viewed in the online course on page 3.5.2.5 graphic number 3.

Interactive Graphic

Activity 3.5.2.6: Identify Parent and Child IPv4 Routes

Go to the online course to perform this practice activity.

The IPv4 Route Lookup Process (3.5.3)

Now that you understand the structure of the routing table, this section will help you understand the routing table lookup process.

Route Lookup Process (3.5.3.1)

When a packet arrives on a router interface, the router examines the IPv4 header, identifies the destination IPv4 address, and proceeds through the router lookup process.

In Figure 3-64, the router examines level 1 network routes for the best match with the destination address of the IPv4 packet.

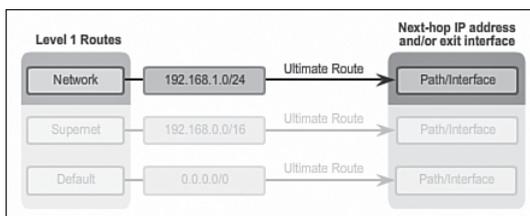


Figure 3-64 Match Level 1 Routes

Specifically, the router proceeds as follows:

1. If the best match is a level 1 ultimate route, then this route is used to forward the packet.
2. If the best match is a level 1 parent route, proceed to the next step.

In Figure 3-65, the router examines child routes (the subnet routes) of the parent route for a best match.

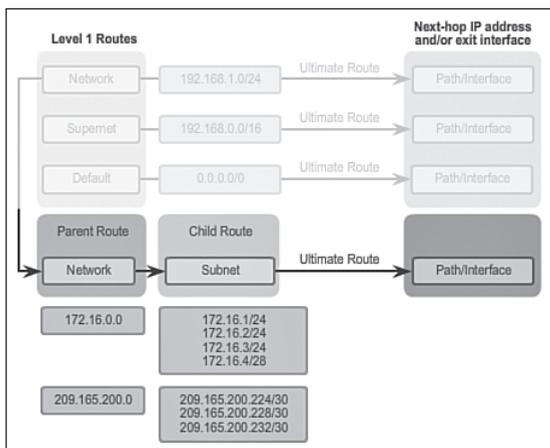


Figure 3-65 Match Level 2 Child Routes

3. If there is a match with a level 2 child route, that subnet is used to forward the packet.
4. If there is not a match with any of the level 2 child routes, proceed to the next step.

In Figure 3-66, the router continues searching level 1 supernet routes in the routing table for a match, including the default route, if there is one.

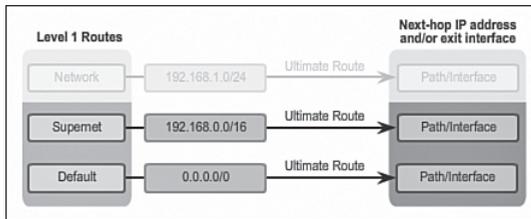


Figure 3-66 Match Supernet and Then Default Route

5. If there is now a lesser match with a level 1 supernet or default routes, the router uses that route to forward the packet.
6. If there is not a match with any route in the routing table, the router drops the packet.

Note

A route referencing only a next-hop IP address and not an exit interface must be resolved to a route with an exit interface. A recursive lookup is performed on the next-hop IP address until the route is resolved to an exit interface.

Best Route = Longest Match (3.5.3.2)

What is meant by the router must find the best match in the routing table? Best match is equal to the longest match.

For there to be a match between the destination IPv4 address of a packet and a route in the routing table, a minimum number of far left bits must match between the IPv4 address of the packet and the route in the routing table. The subnet mask of the route in the routing table is used to determine the minimum number of far left bits that must match. Remember that an IPv4 packet only contains the IPv4 address and not the subnet mask.

The best match is the route in the routing table that has the most number of far left matching bits with the destination IPv4 address of the packet. The route with the greatest number of equivalent far left bits, or the longest match, is always the preferred route.

In Figure 3-67, a packet is destined for 172.16.0.10.

The router has three possible routes that match this packet: 172.16.0.0/12, 172.16.0.0/18, and 172.16.0.0/26. Of the three routes, 172.16.0.0/26 has the longest match and is therefore chosen to forward the packet. Remember, for any of these routes to be considered a match there must be at least the number of matching bits indicated by the subnet mask of the route.

IP Packet Destination	172.16.0.10	10101100.00010000.00000000.00001010
Route 1	172.16.0.0/12	10101100.00010000.00000000.00000000
Route 2	172.16.0.0/18	10101100.00010000.00000000.00000000
Route 3	172.16.0.0/26	10101100.00010000.00000000.00000000

↑
Longest Match to IP Packet Destination

Figure 3-67 Matches for Packets Destined to 172.16.0.10

**Interactive
Graphic**

Activity 3.5.3.3: Determine the Longest Match Route

Go to the online course to perform this practice activity.

Analyze an IPv6 Routing Table (3.5.4)

The IPv6 routing table shares many similarities with the IPv4 routing table. It also consists of directly connected networks and routes learned statically or dynamically. However, the entries are displayed somewhat differently than IPv4 entries. This section examines the IPv6 routing table.

IPv6 Routing Table Entries (3.5.4.1)

Components of the IPv6 routing table are very similar to the IPv4 routing table. For instance, it is populated using directly connected interfaces, static routes, and dynamically learned routes.

Because IPv6 is classless by design, all routes are effectively level 1 ultimate routes. There is no level 1 parent of level 2 child routes.

The topology displayed in Figure 3-68 is used as the reference topology for this section.

Notice that in the topology:

- R1, R2, and R3 are configured in a full mesh topology. All routers have redundant paths to various networks.
- R2 is the edge router and connects to the ISP; however, a default static route is not being advertised.
- EIGRP for IPv6 has been configured on all three routers.

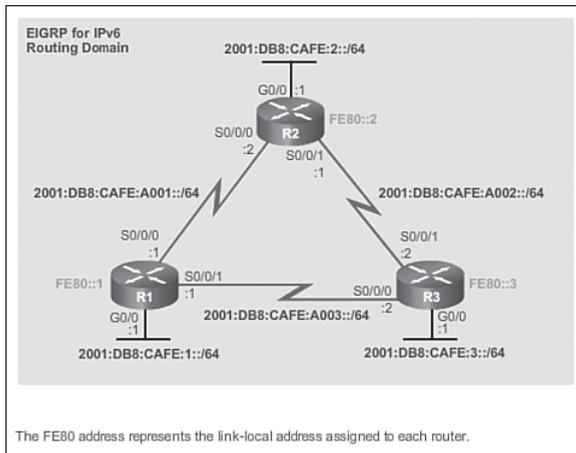


Figure 3-68 Reference IPv6 Topology

Directly Connected Entries (3.5.4.2)

The routing table of R1 is displayed in Figure 3-69 using the `show ipv6 route` command. Although the command output is displayed slightly differently than in the IPv4 version, it still contains the relevant route information.

```

R1# show ipv6 route
<output omitted>
C 2001:DB8:CAFE:1::/64 [0/0]
   via GigabitEthernet0/0, directly connected
L 2001:DB8:CAFE:1::1/128 [0/0]
   via GigabitEthernet0/0, receive
D 2001:DB8:CAFE:2::/64 [90/3524096]
   via FE80::3, Serial0/0/1
D 2001:DB8:CAFE:3::/64 [90/2170112]
   via FE80::3, Serial0/0/1
C 2001:DB8:CAFE:A001::/64 [0/0]
   via Serial0/0/0, directly connected
L 2001:DB8:CAFE:A001::1/128 [0/0]
   via Serial0/0/0, receive
D 2001:DB8:CAFE:A002::/64 [90/3523840]
   via FE80::3, Serial0/0/1
C 2001:DB8:CAFE:A003::/64 [0/0]
   via Serial0/0/1, directly connected
L 2001:DB8:CAFE:A003::1/128 [0/0]
   via Serial0/0/1, receive
L FF00::/8 [0/0]
   via Null0, receive
R1#

```

Figure 3-69 IPv6 Routing Table of R1

Figure 3-70 highlights the connected network and local routing table entries of the directly connected interfaces. The three entries were added when the interfaces were configured and activated.

```

R1# show ipv6 route
<output omitted>

C   2001:DB8:CAFE:1::/64 [0/0]
    via GigabitEthernet0/0, directly connected
L   2001:DB8:CAFE:1::1/128 [0/0]
    via GigabitEthernet0/0, receive
D   2001:DB8:CAFE:2::/64 [90/3524096]
    via FE80::3, Serial0/0/1
D   2001:DB8:CAFE:3::/64 [90/2170112]
    via FE80::3, Serial0/0/1
C   2001:DB8:CAFE:A001::/64 [0/0]
    via Serial0/0/0, directly connected
L   2001:DB8:CAFE:A001::1/128 [0/0]
    via Serial0/0/0, receive
D   2001:DB8:CAFE:A002::/64 [90/3523840]
    via FE80::3, Serial0/0/1
C   2001:DB8:CAFE:A003::/64 [0/0]
    via Serial0/0/1, directly connected
L   2001:DB8:CAFE:A003::1/128 [0/0]
    via Serial0/0/1, receive
L   FF00::/8 [0/0]
    via Null0, receive

R1#

```

Figure 3-70 Directly Connected Routes on R1

As shown in Figure 3-71, directly connected route entries display the following information:

- **Route source:** Identifies how the route was learned. Directly connected interfaces have two route source codes (C identifies a directly connected network while L identifies that this is a local route).
- **Directly connected network:** The IPv6 address of the directly connected network.
- **Administrative distance:** Identifies the trustworthiness of the route source. IPv6 uses the same distances as IPv4. A value of 0 indicates the best, most trustworthy source.
- **Metric:** Identifies the value assigned to reach the remote network. Lower values indicate preferred routes.
- **Outgoing interface:** Identifies the exit interface to use when forwarding packets to the destination network.

Note

The serial links have reference bandwidths configured to observe how EIGRP metrics select the best route. The reference bandwidth is not a realistic representation of modern networks. It is used only to provide a visual sense of link speed.

```

R1# show ipv6 route
<output omitted>
C 2001:DB8:CAFE:1::/64 [0/0]
  via GigabitEthernet0/0, directly connected
L 2001:DB8:CAFE:1:1/128 [0/0]
  via GigabitEthernet0/0, receive
C 2001:DB8:CAFE:2::/64 [90/3524096]
  via FE80::3, Serial0/0/1
  via Serial0/0/1, receive
C 2001:DB8:CAFE:3::/64 [90/2170112]
  via FE80::3, Serial0/0/1
  via Serial0/0/1, receive
C 2001:DB8:CAFE:A001::/64 [0/0]
  via Serial0/0/0, directly connected
L 2001:DB8:CAFE:A001:1/128 [0/0]
  via Serial0/0/0, receive
C 2001:DB8:CAFE:A002::/64 [90/3523840]
  via FE80::3, Serial0/0/1
  via Serial0/0/1, directly connected
L 2001:DB8:CAFE:A003:1/128 [0/0]
  via Serial0/0/1, receive
L FF00::/8 [0/0]
  via Null0, receive
R1#

```

Figure 3-71 Directly Connected Routes on R1

Remote IPv6 Network Entries (3.5.4.3)

Figure 3-72 highlights the routing table entries for the three remote networks (i.e., R2 LAN, R3 LAN, and the link between R2 and R3). The three entries were added by the EIGRP.

```

R1# show ipv6 route
<output omitted>
C 2001:DB8:CAFE:1::/64 [0/0]
  via GigabitEthernet0/0, directly connected
L 2001:DB8:CAFE:1:1/128 [0/0]
  via GigabitEthernet0/0, receive
D 2001:DB8:CAFE:2::/64 [90/3524096]
  via FE80::3, Serial0/0/1
D 2001:DB8:CAFE:3::/64 [90/2170112]
  via FE80::3, Serial0/0/1
C 2001:DB8:CAFE:A001::/64 [0/0]
  via Serial0/0/0, directly connected
L 2001:DB8:CAFE:A001:1/128 [0/0]
  via Serial0/0/0, receive
D 2001:DB8:CAFE:A002::/64 [90/3523840]
  via FE80::3, Serial0/0/1
C 2001:DB8:CAFE:A003::/64 [0/0]
  via Serial0/0/1, directly connected
L 2001:DB8:CAFE:A003:1/128 [0/0]
  via Serial0/0/1, receive
L FF00::/8 [0/0]
  via Null0, receive
R1#

```

Figure 3-72 Remote Networks Entries on R1

Figure 3-73 displays a routing table entry on R1 for the route to remote network 2001:DB8:CAFE:3::/64 on R3.

```

R1# show ipv6 route
<output omitted>
C 2001:DB8:CAFE:2::/64 [0/0]
  via GigabitEthernet0/0, receive
  2001:DB8:CAFE:2::/64 [0/0]
  via FE80::3, Serial0/0/1
D 2001:DB8:CAFE:3::/64 [90/2170112]
  via FE80::3, Serial0/0/1
C 2001:DB8:CAFE:A001::/64 [0/0]
  via Serial0/0/0, directly connected
C 2001:DB8:CAFE:A003::/64 [0/0]
  via Serial0/0/1, directly connected
L 2001:DB8:CAFE:A003::1/128 [0/0]
  via Serial0/0/1, receive
L FF00::/8 [0/0]
  via Null0, receive
R1#

```

The screenshot shows a Cisco router's IPv6 routing table. A callout box labeled 'Route Source' points to the 'D' in the entry 'D 2001:DB8:CAFE:3::/64'. Another callout labeled 'Destination Network' points to '2001:DB8:CAFE:3::/64'. A callout labeled 'Administrative Distance' points to '[90]'. A callout labeled 'Metric' points to '[2170112]'. A callout labeled 'Next Hop' points to 'via FE80::3, Serial0/0/1'. A callout labeled 'Outgoing Interface' points to 'Serial0/0/1'.

Figure 3-73 Remote Networks Entries on R1

The entry identifies the following information:

- **Route source:** Identifies how the route was learned. Common codes include O (OSPF), D (EIGRP), R (RIP), and S (Static route).
- **Destination network:** Identifies the address of the remote IPv6 network.
- **Administrative distance:** Identifies the trustworthiness of the route source. IPv6 uses the same distances as IPv4.
- **Metric:** Identifies the value assigned to reach the remote network. Lower values indicate preferred routes.
- **Next hop:** Identifies the IPv6 address of the next router to forward the packet to.
- **Outgoing interface:** Identifies the exit interface to use to forward a packet toward the final destination.

When an IPv6 packet arrives on a router interface, the router examines the IPv6 header and identifies the destination IPv6 address. The router then proceeds through the following router lookup process.

The router examines level 1 network routes for the best match with the destination address of the IPv6 packet. Just like IPv4, the longest match is the best match. For example, if there are multiple matches in the routing table, the router chooses the route with the longest match. A match is made by matching the far left bits of the packet's destination IPv6 address with the IPv6 prefix and prefix-length in the IPv6 routing table.

Activity 3.5.4.4: Identify Parts of an IPv6 Routing Table Entry

Go to the online course to perform this practice activity.

Summary (3.6)



Class Activity 3.6.1.1: IPv6 Details, Details...

After studying the concepts presented in this chapter concerning IPv6, you should be able to read a routing table easily and interpret the IPv6 routing information listed within it.

With a partner, use the IPv6 routing table diagram and the .pdf provided with this activity.

Record your answers to the Reflection questions.

Then compare your answers with, at least, one other group from the class.

Dynamic routing protocols are used by routers to facilitate the exchange of routing information between routers. The purpose of dynamic routing protocols includes: discovery of remote networks, maintaining up-to-date routing information, choosing the best path to destination networks, and ability to find a new best path if the current path is no longer available. While dynamic routing protocols require less administrative overhead than static routing, they do require dedicating part of a router's resources for protocol operation, including CPU time and network link bandwidth.

Networks typically use a combination of both static and dynamic routing. Dynamic routing is the best choice for large networks and static routing is better for stub networks.

Routing protocols are responsible for discovering remote networks, as well as maintaining accurate network information. When there is a change in the topology, routing protocols propagate that information throughout the routing domain. The process of bringing all routing tables to a state of consistency, where all of the routers in the same routing domain or area have complete and accurate information about the network, is called convergence. Some routing protocols converge faster than others.

Routing protocols can be classified as either classful or classless, as distance vector or link-state, and as an Interior Gateway Protocol or an Exterior Gateway Protocol.

Distance vector protocols use routers as “sign posts” along the path to the final destination. The only information a router knows about a remote network is the distance or metric to reach that 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.

A router configured with a link-state routing protocol can create a complete view or topology of the network by gathering information from all of the other routers.

Metrics are used by routing protocols to determine the best path or shortest path to reach a destination network. Different routing protocols may use different metrics. Typically, a lower metric means a better path. Metrics can be determined by hops, bandwidth, delay, reliability, and load.

Routers sometimes learn about multiple routes to the same network from both static routes and dynamic routing protocols. When a router learns about a destination network from more than one routing source, Cisco routers use the administrative distance value to determine which source to use. Each dynamic routing protocol has a unique administrative value, along with static routes and directly connected networks. The lower the administrative value, the more preferred the route source. A directly connected network is always the preferred source, followed by static routes and then various dynamic routing protocols.

The **show ip protocols** command displays the IPv4 routing protocol settings currently configured on the router. For IPv6, use **show ipv6 protocols**.

With link-state routing protocols such as OSPF, a link is an interface on a router. Information about the state of those links is known as link-states. All link-state routing protocols apply Dijkstra's algorithm to calculate the best path route. The algorithm is commonly referred to as the shortest path first (SPF) algorithm. This algorithm uses accumulated costs along each path, from source to destination, to determine the total cost of a route.

Practice

The following activities provide practice with the topics introduced in this chapter. The Labs and Class Activities are available in the companion *Routing Protocols Lab Manual* (978-1-58713-322-0). The Packet Tracer Activities PKA files are found in the online course.



Class Activities

Class Activity 3.0.1.2: How Much Does This Cost?

Class Activity 3.6.1.1: IPv6 Details, Details...



Lab

Lab 3.3.2.4: Configuring RIPv2



Packet Tracer Activities

Packet Tracer Activity 3.1.3.6: Investigating Convergence

Packet Tracer Activity 3.2.2.4: Comparing RIP and EIGRP Path Selection

Packet Tracer Activity 3.3.1.8: Configuring RIPv2

Packet Tracer Activity 3.3.2.3: Configuring RIPng

Check Your Understanding Questions

Complete all the review questions listed here to test your understanding of the topics and concepts in this chapter. The appendix, “Answers to the ‘Check Your Understanding’ Questions,” lists the answers.

1. What are two advantages of static routing over dynamic routing? (Choose two.)
 - A. The configuration is less error prone.
 - B. Static routing is more secure because routers do not advertise routes.
 - C. Growing the network usually does not present a problem.
 - D. No computing overhead is involved.
 - E. The administrator has less work maintaining the configuration.
2. Match the description to the proper routing protocol.

Routing protocols:

RIP
IGRP
OSPF
EIGRP
BGP

Description:

- A. Path vector exterior routing protocol:
- B. Cisco advanced interior routing protocol:
- C. Link-state interior routing protocol:
- D. Distance vector interior routing protocol:
- E. Cisco distance vector interior routing protocol:

3. Which statement best describes convergence on a network?
 - A. The amount of time required for routers to share administrative configuration changes, such as password changes, from one end of a network to the other end
 - B. The time required for the routers in the network to update their routing tables after a topology change has occurred
 - C. The time required for the routers in one autonomous system to learn routes to destinations in another autonomous system
 - D. The time required for routers running disparate routing protocols to update their routing tables

4. Dynamic routing protocols perform which two tasks? (Choose two.)
 - A. Assign IP addressing
 - B. Discover hosts
 - C. Network discovery
 - D. Propagate host default routes
 - E. Update and maintain routing tables

5. Which of the following parameters are used to calculate metrics? (Choose two.)
 - A. Hop count
 - B. Uptime
 - C. Bandwidth
 - D. Convergence time
 - E. Administrative distance

6. Which routing protocol has the most trustworthy administrative distance by default?
 - A. EIGRP internal routes
 - B. IS-IS
 - C. OSPF
 - D. RIPv1
 - E. RIPv2

7. Which command will show the administrative distance of routes?
 - A. **show interfaces**
 - B. **show ip route**
 - C. **show ip interfaces**
 - D. **debug ip routing**

8. When do directly connected networks appear in the routing table?
 - A. When they are included in a static route
 - B. When they are used as an exit interface
 - C. As soon as they are addressed and operational at Layer 2
 - D. As soon as they are addressed and operational at Layer 3
 - E. Always when a **no shutdown** command is issued

9. Router R1 is using the RIPv2 routing protocol and has discovered multiple unequal paths to reach a destination network. How will Router R1 determine which path is the best path to the destination network?
 - A. Lowest metric
 - B. Highest metric
 - C. Lowest administrative distance
 - D. Highest administrative distance
 - E. By load-balancing between up to four paths

10. Which of the following will trigger the sending of a link-state packet by OSPF? (Choose two.)
 - A. A change in the topology
 - B. A link to a neighbor router has become congested
 - C. The initial startup of the routing protocol process
 - D. The router update timer expiring

11. After examining its routing table for a best match with the destination address, which route will a router use to forward an IPv4 packet?
 - A. A level 1 child route
 - B. A level 1 parent route
 - C. A level 1 ultimate route
 - D. A level 2 supernet route
 - E. A level 2 parent route

-
- 12.** What is different between IPv6 routing table entries and IPv4 routing table entries?
- A. By design IPv6 is classless, so all routes are effectively level 1 ultimate routes.
 - B. Unlike IPv4, IPv6 does not use static routes to populate the routing table.
 - C. IPv6 routing tables include local route entries, which IPv4 routing tables do not.
 - D. The selection of IPv6 routes is based on the shortest matching prefix, unlike IPv4 route selection, which is based on the longest matching prefix.
- 13.** Enter the proper administrative distance for each routing protocol.
- A. eBGP:
 - B. EIGRP (Internal):
 - C. EIGRP (External):
 - D. IS-IS:
 - E. OSPF:
 - F. RIP:
- 14.** Designate the following characteristics as belonging to either a classful routing protocol or a classless routing protocol.
- A. Does not support discontinuous networks:
 - B. EIGRP, OSPF, and BGP:
 - C. Sends subnet mask information in routing updates:
 - D. Supports discontinuous networks:
 - E. RIP version 1 and IGRP:
 - F. Does not send subnet mask in its routing updates:
 - G. Allows for use of both 172.16.1.0/26 and 172.16.1.128/27 subnets in the same topology:
- 15.** Explain why static routing might be preferred over dynamic routing.
- 16.** What are four ways of classifying dynamic routing protocols?
- 17.** What are the most common metrics used in IP dynamic routing protocols?
- 18.** What is administrative distance, and why is it important?
- 19.** What is the purpose of a passive interface?