

## 2.0 Static Routing

### 2.0.1.1 Introduction

Routing is at the core of every data network, moving information across an internetwork 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, or 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

Refer to  
Lab Activity  
for this chapter

### 2.0.1.2 Class Activity - Which Way Should We Go

#### Activity - 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 it 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.

Refer to  
Online Course  
for Illustration

## 2.1 Static Routing Implementation

### 2.1.1 Static Routing

#### 2.1.1.1 Reach Remote Networks

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

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.

Refer to  
Online Course  
for Illustration

#### 2.1.1.2 Why Use Static Routing?

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, 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 is 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 the figure, dynamic and static routing features are compared. Notice that the advantages of one method are the disadvantages of the other.

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.

Refer to  
**Online Course**  
for Illustration

### 2.1.1.3 When to Use Static Routes

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.

The figure 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.

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.

Refer to  
**Interactive Graphic**  
in online course.

### 2.1.1.4 Activity - Identify the Advantages and Disadvantages of Static Routing

Refer to  
**Online Course**  
for Illustration

## 2.1.2 Types of Static Routes

### 2.1.2.1 Static Route Applications

As shown in the figure, 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

Refer to  
Online Course  
for Illustration

### 2.1.2.2 Standard Static Route

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

The figure 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.

Refer to  
Online Course  
for Illustration

### 2.1.2.3 Default Static Route

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. 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 the figure for a sample scenario of implementing default static routing.

Refer to  
Online Course  
for Illustration

### 2.1.2.4 Summary Static Route

To reduce the number of routing table entries, 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.

In the figure, 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.

Refer to  
Online Course  
for Illustration

### 2.1.2.5 Floating Static Route

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 the figure, 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.

Refer to  
Interactive Graphic  
in online course.

### 2.1.2.6 Activity - Identify the Type of Static Route

Refer to  
Online Course  
for Illustration

## 2.2 Configure Static and Default Routes

### 2.2.1 Configure IPv4 Static Routes

#### 2.2.1.1 ip route Command

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.
- *exit-intf* - The outgoing interface to use to forward the packet to the next hop.

As shown in the figure, 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.

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

Refer to  
Online Course  
for Illustration

### 2.2.1.2 Next-Hop Options

In this example, Figures 1 to 3 display 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.

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 4 displays a successful ping from R1 to R2. Figure 5 displays an unsuccessful ping to the R3 LAN. This is because R1 does not have an entry in its routing table for the R3 LAN network.

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.

Refer to  
Online Course  
for Illustration

### 2.2.1.3 Configure a Next-Hop Static Route

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 1, three next-hop static routes are configured on R1 using the IP address of the next hop, R2.

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 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 the figure. 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 the figure. 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.

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: a FIB (Forwarding Information Base), which is a copy of the routing table and an adjacency table that includes Layer 2 addressing information. The information combined in both of these tables work together so there is no recursive lookup needed for next-hop IP address lookups. In other words, a static route using a next-hop IP requires only a single lookup when CEF is enabled on the router.

Use the Syntax Checker in Figures 3 and 4 to configure and verify next-hop static routes on R2 and R3.

Refer to  
Online Course  
for Illustration

### 2.2.1.4 Configure a Directly Connected Static Route

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 1, three directly connected static routes are configured on R1 using the exit interface. The routing table for R1 in Figure 2 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.

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.

Use the Syntax Checker in Figures 3 and 4 to configure and verify directly connected static routes on R2 and R3.

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.

Refer to  
Online Course  
for Illustration

### 2.2.1.5 Configure a Fully Specified Static Route

#### Fully Specified Static Route

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 IOSs, 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 1. 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
```

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, that a fully specified static route is used including both the exit interface and the next-hop address.

As shown in Figure 2, 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.

Use the Syntax Checker in Figure 3 and 4 to configure and verify fully specified static routes on R2 and R3.

Refer to  
Online Course  
for Illustration

### 2.2.1.6 Verify a Static Route

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

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



Figure 1 displays sample output of the `show ip route static` command. In the example, the output is filtered using the pipe and `begin` parameter. The output reflects the use of static routes using the next-hop address.

Figure 2 displays sample output of the `show ip route 192.168.2.1` command.

Figure 3 verifies the `ip route` configuration in the running configuration.

Use the Syntax Checker in Figure 4 to verify the routing settings of R2.

Use the Syntax Checker in Figure 5 to verify the routing settings of R3.

Refer to  
Online Course  
for Illustration

## 2.2.2 Configure IPv4 Default Routes

### 2.2.2.1 Default Static Route

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 the figure, 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.

Refer to  
Online Course  
for Illustration

### 2.2.2.2 Configure a Default Static Route

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 the figure 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.

Refer to  
Online Course  
for Illustration

### 2.2.2.3 Verify a Default Static Route

In the figure, 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 the figure, the asterisk indicates that this static route is a candidate default route, which is why it is selected as the Gateway of Last Resort.

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.

Refer to **Packet Tracer Activity** for this chapter

### 2.2.2.4 Packet Tracer - Configuring IPv4 Static and Default Routes

#### Background/Scenario

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.

Refer to **Lab Activity** for this chapter

### 2.2.2.5 Lab - 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

Refer to **Online Course** for Illustration

## 2.2.3 Configure IPv6 Static Routes

### 2.2.3.1 The ipv6 route Command

Static routes for IPv6 are configured using the `ipv6 route` global configuration command. Figure 1 shows the simplified version of the command syntax:

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

Most of 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 displays the enabling of IPv6 unicast routing.

Use the Syntax Checker in Figures 3 and 4 to enable IPv6 unicast routing on R2 and R3.

Refer to  
Online Course  
for Illustration

### 2.2.3.2 Next-Hop Options

In this example, Figures 1 to 3 display 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.

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 4 displays a successful ping from R1 to R2. Figure 5 displays an unsuccessful ping to the R3 LAN. This is because R1 does not have an entry in its routing table for that network.

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.

Refer to  
Online Course  
for Illustration

### 2.2.3.3 Configure a Next-Hop Static IPv6 Route

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 1, three next-hop static routes are configured on R1.

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 (Cisco Express Forwarding) is the default behavior on most platforms running IOS 12.0 or later.

Figure 2 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.
2. 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.

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.

Use the Syntax Checker in Figure 3 and Figure 4 to configure next-hop static IPv6 routes.

Refer to  
Online Course  
for Illustration

### 2.2.3.4 Configure a Directly Connected Static IPv6 Route

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 IOSs or whenever CEF is disabled, to avoid the recursive lookup problem.

For instance, in Figure 1, three directly connected static routes are configured on R1 using the exit interface.

The IPv6 routing table for R1 in Figure 2 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.

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.

Use the Syntax Checker in Figure 3 and Figure 4 to configure directly connected static IPv6 routes.

Refer to  
Online Course  
for Illustration

### 2.2.3.5 Configure a Fully Specified Static IPv6 Route

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 1 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 1, 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 shows the IPv6 routing table entry for this route. Notice that both the next-hop link-local address and the exit interface are included.

Use the Syntax Checker in Figure 3 to configure fully specified static IPv6 routes on R2 to reach R1's LAN using a link-local address.

Refer to  
Online Course  
for Illustration

### 2.2.3.6 Verify IPv6 Static Routes

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

- `show ipv6 route`
- `show ipv6 route static`
- `show ipv6 route network`

Figure 1 displays sample output of the `show ipv6 route static` command. The output reflects the use of static routes using next-hop global unicast addresses.

Figure 2 displays sample output from the `show ip route 2001:DB8:ACAD:3::` command.

Refer to  
Online Course  
for Illustration

Figure 3 verifies the `ipv6 route` configuration in the running configuration.

## 2.2.4 Configure IPv6 Default Routes

### 2.2.4.1 Default Static IPv6 Route

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

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 the figure, 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 }`

Refer to  
**Online Course**  
for Illustration

### 2.2.4.2 Configure a Default Static IPv6 Route

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 the figure displays a configuration for a default static IPv6 route on R1.

Refer to  
**Online Course**  
for Illustration

### 2.2.4.3 Verify a Default Static Route

In Figure 1, the `show ipv6 route static` command output displays the contents of the routing table.

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.

Figure 2 displays a successful ping to the R3 LAN interface.

Refer to **Packet Tracer Activity** for this chapter

### 2.2.4.4 Packet Tracer - 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.

Refer to  
**Lab Activity**  
for this chapter

### 2.2.4.5 Lab - 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

Refer to  
**Online Course**  
for Illustration

## 2.3 Review of CIDR and VLSM

### 2.3.1 Classful Addressing

#### 2.3.1.1 Classful Network Addressing

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 the figure:

- **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.

Links:

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

“Internet Multicast Addresses,” <http://www.iana.org/assignments/multicast-addresses>

Refer to  
Online Course  
for Illustration

### 2.3.1.2 Classful Subnet Masks

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

As shown in Figure 1, 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.

As shown in Figure 2, 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.)

As shown in Figure 3, 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.

Refer to  
Online Course  
for Illustration

### 2.3.13 Classful Routing Protocol Example

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 1, 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 a 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.

In Figure 2, 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.

Refer to  
Online Course  
for Illustration

### 2.3.14 Classful Addressing Waste

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 the figure:

- 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 170.0.0.0/8, 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 mid-size organizations.
- Classes D and E are used for multicasting and reserved addresses.

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.



Refer to  
Online Course  
for Illustration

## 2.3.2 CIDR

### 2.3.2.1 Classless Inter-Domain Routing

Just as the Internet was growing at an exponential rate in the early 1990s, so were 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.). The figure 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.

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.

Refer to  
Online Course  
for Illustration

### 2.3.2.2 CIDR and Route Summarization

In the figure, 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.

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.

Refer to  
Online Course  
for Illustration

### 2.3.2.3 Static Routing CIDR Example

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

Refer to  
Online Course  
for Illustration

### 2.3.2.4 Classless Routing Protocol Example

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 the figure contained a classful routing protocol, then R3 would only install 172.16.0.0/16 in the routing table.

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.

Refer to  
Online Course  
for Illustration

## 2.3.3 VLSM

### 2.3.3.1 Fixed-Length Subnet Masking

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 1 requires that network address 192.168.20.0/24 be subnetted into seven subnets: one subnet for each of the four LANs (Building A to D), and one for each of the three WAN connections between routers.

Figure 2 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 hosts IP addresses per subnet. This scheme creates the needed subnets and meets the host requirement of the largest LAN.

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 3, this results in 84 unused addresses (28 x 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.

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 mask (VLSM), was designed to avoid wasting addresses.

Refer to  
Online Course  
for Illustration

### 2.3.3.2 Variable-Length Subnet Masking

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 1, traditional subnetting creates subnets of equal size. Each subnet in a traditional scheme uses the same subnet mask.

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, 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.

Refer to  
Online Course  
for Illustration

### 2.3.3.3 VLSM in Action

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.

The figure 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, ..., 10.255.0.0/16. Four of these /16 subnets are displayed in the figure. Any of these /16 subnets can be subnetted further.

Click the Play button in the figure to view the animation. In the animation:

- 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, the figure 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.

Refer to  
Online Course  
for Illustration

### 2.3.3.4 Subnetting Subnets

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

In Figure 1, 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.

In Figure 2, 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.

In Figure 3, 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.

In Figure 4, 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.

In Figure 5, 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.

Refer to  
Online Course  
for Illustration

### 2.3.3.5 VLSM Example

Careful consideration must be given to the design of a network addressing scheme. For example, the sample topology in Figure 1 requires seven subnets.

Using traditional subnetting, the first seven address blocks are allocated for LANs and WANs, as shown in Figure 2. 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.

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 3, 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.

Figures 4 to 7 display sample configurations on all four routers to implement the VLSM addressing scheme.

Refer to Packet Tracer Activity for this chapter

### 2.3.3.6 Packet Tracer - Designing and Implementing a VLSM Addressing Scheme

#### Background/Scenario

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

Refer to Lab Activity for this chapter

### 2.3.3.7 Lab - 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

Refer to Online Course for Illustration

## 2.4 Configure Summary and Floating Static Routes

### 2.4.1 Configure IPv4 Summary Routes

#### 2.4.1.1 Route Summarization

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 the figure, R1 requires a summary static route to reach networks in the range of 172.20.0.0/16 to 172.23.0.0/16.

Refer to Online Course for Illustration

#### 2.4.1.2 Calculate a Summary Route

Summarizing networks into a single address and mask can be done in three steps:

- Step 1.** List the networks in binary format. Figure 1 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 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 3 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 4 displays R1 configured with a summary static route to reach networks 172.20.0.0/16 to 172.23.0.0/16.

Refer to  
Online Course  
for Illustration

### 2.4.1.3 Summary Static Route Example

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 1. All routers have connectivity using static routes.

Figure 2 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.

Figure 3 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 4 displays how the three existing routes are removed and then the new summary static route is configured.

Figure 5 confirms that the summary static route is in the routing table of R3.

Refer to  
Interactive Graphic  
in online course.

### 2.4.1.4 Activity - Determine the Summary Network Address and Prefix

Refer to Packet  
Tracer Activity  
for this chapter

### 2.4.1.5 Packet Tracer - Configuring IPv4 Route Summarization - Scenario 1

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

Refer to **Packet Tracer Activity** for this chapter

### 2.4.1.6 Packet Tracer - Configuring IPv4 Route Summarization - Scenario 2

In this activity, you will calculate and configure summary routes. Router 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.

Refer to **Online Course** for Illustration

## 2.4.2 Configure IPv6 Summary Routes

### 2.4.2.1 Summarize IPv6 Network Addresses

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 the Figure 1. R1 currently has four static IPv6 routes to reach networks 2001:DB8:ACAD:1::/64 to 2001:DB8:ACAD:4::/64.

Figure 2 displays the IPv6 static routes installed in the IPv6 routing table.

Refer to **Online Course** for Illustration

### 2.4.2.2 Calculate IPv6 Network Addresses

Summarizing IPv6 networks into a single IPv6 prefix and prefix-length can be done in seven steps as shown in Figures 1 to 7:

- 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).

Refer to **Online Course** for Illustration

### 2.4.2.3 Configure an IPv6 Summary Address

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

Figure 1 displays how the four existing routes are removed and then the new summary static IPv6 route is configured.

Figure 2 confirms that the summary static route is in the routing table of R1.

Refer to **Packet Tracer Activity** for this chapter

#### 2.4.2.4 Packet Tracer - 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.

Refer to **Lab Activity** for this chapter

#### 2.4.2.5 Lab - 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

Refer to **Online Course** for illustration

### 2.4.3 Configure Floating Static Routes

#### 2.4.3.1 Floating Static Routes

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 the figure.

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.

Refer to  
Online Course  
for Illustration

### 2.4.3.2 Configure a Floating Static Route

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 1. In this scenario, the preferred route from R1 is to R2. The connection to R3 should be used for backup only.

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.

Figure 2 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.

Use the Syntax Checker in Figure 3 to configure R3 similarly to R1.

Refer to  
Online Course  
for Illustration

### 2.4.3.3 Test the Floating Static Route

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 1 confirms that traffic between R1 and R3 flows through R2.

What would happen if R2 failed? To simulate this failure both serial interfaces of R2 are shut down, as shown in Figure 2.

Notice in Figure 3 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.

The output in Figure 4 confirms that traffic now flows directly between R1 and R3.

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

Refer to Packet  
Tracer Activity  
for this chapter

### 2.4.3.4 Packet Tracer - 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.

Refer to  
Online Course  
for Illustration

## 2.5 Troubleshoot Static and Default Route Issues

### 2.5.1 Packet Processing with Static Routes

#### 2.5.1.1 Static Routes and Packet Forwarding

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

In the figure, click the Play button to see the animation, where PC1 is sending a packet to PC3:

1. The packet arrives on the GigabitEthernet 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 GigabitEthernet 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 GigabitEthernet 0/0 interface, and PC3 responds with an ARP reply, which includes the PC3 MAC address.
10. R3 encapsulates the packet in a new frame with the MAC address of the GigabitEthernet 0/0 interface as the source Layer 2 address and the MAC address of PC3 as the destination MAC address.
11. The frame is forwarded out of GigabitEthernet 0/0 interface. The packet arrives on the network interface card (NIC) interface of PC3.

Refer to  
Online Course  
for Illustration

### 2.5.2 Troubleshoot IPv4 Static and Default Route Configuration

#### 2.5.2.1 Troubleshoot a Missing Route

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`
- `tracert`
- `show ip route`
- `show ip interface brief`
- `show cdp neighbors detail`

Figure 1 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 displays the result of a traceroute from R1 to the R3 LAN.

Figure 3 displays the routing table of R1.

Figure 4 provides a quick status of all interfaces on the router.

Figure 5 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.

Refer to  
Online Course  
for Illustration

### 2.5.2.2 Solve a Connectivity Problem

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 1). The results show that there is no connectivity between these LANs.

A traceroute in Figure 2 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.

The next step is to investigate the routing table of R2, because it is the router displaying a strange forwarding pattern. The routing table in Figure 3 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.

Figure 4 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.

Figure 5 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.

Refer to **Packet Tracer Activity** for this chapter

### 2.5.2.3 Packet Tracer - Solving the Missing Route

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

Refer to **Packet Tracer Activity** for this chapter

### 2.5.2.4 Packet Tracer - 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.

Refer to **Lab Activity** for this chapter

### 2.5.2.5 Lab - 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

Refer to **Online Course** for Illustration

## 2.6 Summary

Refer to **Lab Activity** for this chapter

### 2.6.1.1 Class Activity - Make It Static

**Activity - Make It Static!**

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.

Refer to **Packet Tracer Activity** for this chapter

## 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.

Refer to **Online Course** for illustration

## 2.6.1.3 Summary

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.

Go to the online course to take the quiz and exam.

## Chapter 2 Quiz

This quiz is designed to provide an additional opportunity to practice the skills and knowledge presented in the chapter and to prepare for the chapter exam. You will be allowed multiple attempts and the grade does not appear in the gradebook.

## Chapter 2 Exam

The chapter exam assesses your knowledge of the chapter content.

## Your Chapter Notes