

Chapter 6: IP Routing Essentials

Instructor Materials

CCNP Enterprise: Core Networking



Chapter 6 Content

This chapter covers the following content:

- **Routing Protocol Overview** - This section explains how different routing protocols advertise and identify routes.
- **Path Selection** - This section explains the logic a router uses to identify the best route and install it in the routing table.
- **Static Routing** - This section provides a brief overview of fundamental static route concepts.
- **Virtual Routing and Forwarding** - This section explains the creation of logical routers on a physical router.

Routing Protocol Overview

- A router is necessary to transmit packets between network segments.
- A router's primary function is to move a packet from one network to a different network.
- This section will cover distance vector algorithms, enhanced distance vector algorithms, link state algorithms, and path vector algorithms.

Routing Protocol Overview

Routing Protocol Overview

- A router's primary function is to move an IP packet from one network to a different network.
- A router learns about nonattached networks through configuration of static routes or through dynamic IP routing protocols.
- Dynamic IP routing protocols distribute network topology information between routers and provide updates without intervention when a topology change in the network occurs.
- With dynamic routing protocols, routers try to select the best loop-free path on which to forward a packet to its destination IP address.
- A network of interconnected routers and related systems managed under a common network administration is known as an autonomous system (AS), or a routing domain.

Routing Protocol Overview

Routing Protocol Overview (Cont.)

The common dynamic routing protocols found on most routing platforms today are as follows:

- Routing Information Protocol Version 2 (RIPv2)
- Enhanced Interior Gateway Routing (EIGRP)
- Open Shortest Path First (OSPF)
- Intermediate System-to-Intermediate System (IS-IS)
- Border Gateway Protocol (BGP)

With the exception of BGP, the protocols in this list are designed and optimized for routing within an autonomous system and are known as Interior Gateway Protocols (IGPs).

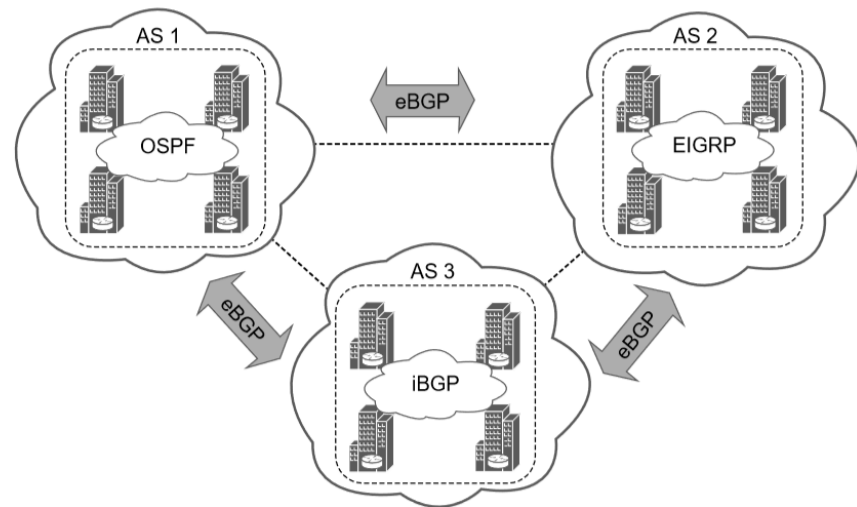


Figure 6-1 BGP Autonomous Systems and How They Interconnect

Routing Protocol Overview

Distance Vector Algorithms

Distance vector routing protocols, such as RIP, advertise routes as vectors, where distance is a metric (or cost) such as hop count, and vector is the next-hop router's IP used to reach the destination:

- **Distance** - The distance is the route metric to reach the network.
- **Vector** - The vector is the interface or direction to reach the network.

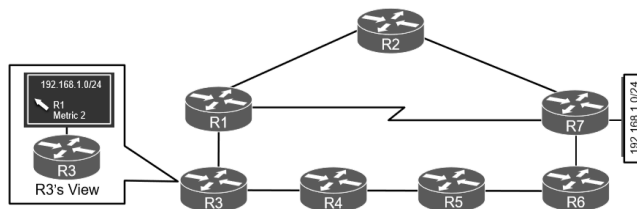


Figure 6-2 Distance Vector Protocol View of a Network

When a router receives routing information from a neighbor, it stores it in a local routing database, and the distance vector algorithm (such as the Bellman-Ford and Ford-Fulkerson algorithms) is used to determine which paths are the best loop-free paths to each reachable destination.

Enhanced Distance Vector Algorithms

The diffusing update algorithm (DUAL) is an enhanced distance vector algorithm that EIGRP uses to calculate the shortest path to a destination within a network. EIGRP advertises network information to its neighbors as other distance vector protocols do, but it has some enhancements:

- It offers rapid convergence time for changes in the network topology.
- It sends updates only when there is a change in the network. It does not send full routing table updates in a periodic fashion, as distance vector protocols do.
- It uses hellos and forms neighbor relationships just as link-state protocols do.
- It uses bandwidth, delay, reliability, load, and maximum transmission unit (MTU) size instead of hop count for path calculations.
- It has the option to load balance traffic across equal- or unequal-cost paths.

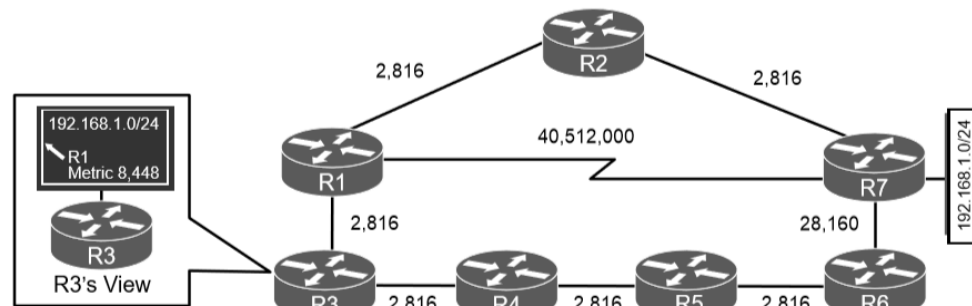


Figure 6-3 Distance Vector Protocol Versus Enhanced Distance Vector

EIGRP is sometimes referred to as a hybrid routing protocol.

Routing Protocol Overview

Link-State Algorithms

A link-state dynamic IP routing protocol advertises the link state and link metric for each of its connected links and directly connected routers to every router in the network. OSPF and IS-IS are two link-state routing protocols commonly used in enterprise and service provider networks. OSPF advertisements are called link-state advertisements (LSAs), and IS-IS uses link-state packets (LSPs) for its advertisements.

- A router receives an advertisement from a neighbor and stores it in the link-state database (LSDB).
- The link state information is advertised to each of its neighbor routers exactly as it was received. This enables all the routers in the network to have an identical map of the network.
- The Dijkstra shortest path first (SPF) algorithm is run by every router in the network in order to calculate the best, shortest loop-free paths.

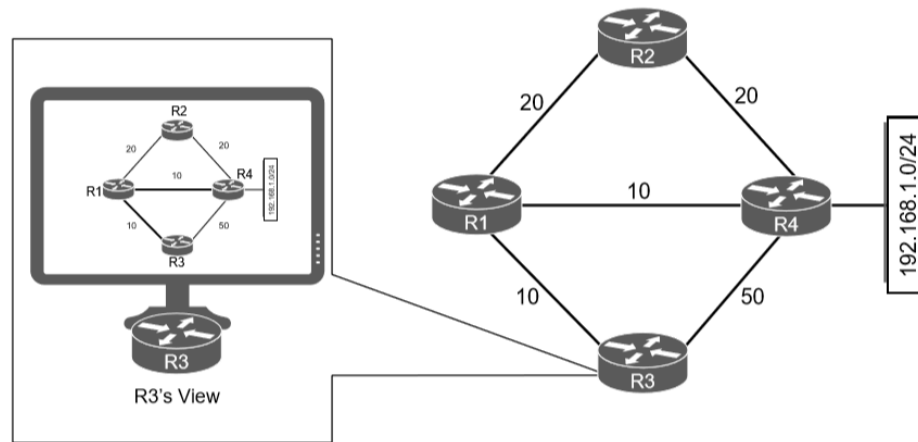


Figure 6-4 Link-State Protocol View of a Network

Routing Protocol Overview

Path Vector Algorithm

A path vector protocol, such as BGP, is similar to a distance vector protocol. The difference is that instead of looking at the distance to determine the best loop-free path, it looks at various BGP path attributes. BGP path attributes include autonomous system path (AS_Path), multi-exit discriminator (MED), origin, next hop, local preference, atomic aggregate, and aggregator.

Figure 6-5 illustrates the loop prevention concept over the following steps:

1. R1 (AS 1) advertises the 10.1.1.0/24 network to R2 (AS 2). R1 adds the AS 1 to the AS_Path during the network advertisement to R2.
2. R2 advertises the 10.1.1.0/24 network to R4 and adds AS 2 to the AS_Path during the network advertisement to R4.
3. R4 advertises the 10.1.1.0/24 network to R3 and adds AS 4 to the AS_Path during the network advertisement to R3.
4. R3 advertises the 10.1.1.0/24 network back to R1 and R2 after adding AS 3 to the AS_Path during the network advertisement.
5. As R1 receives the 10.1.1.0/24 network advertisement from R3, it discards the route advertisement because R1 detects its AS (AS 1) in the AS_Path “3 4 2 1” and considers the advertisement as a loop. R2 discards the 10.1.1.0/24 network advertisement from R3 as it detects its AS (AS 2) in the AS_Path “3 4 2 1” and considers it a loop, too.

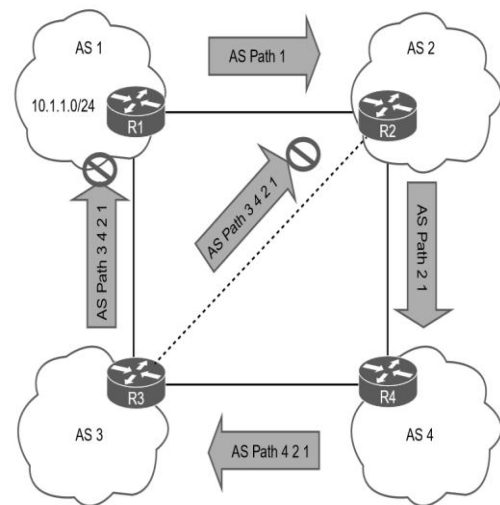


Figure 6-5 Path Vector Loop Avoidance

Path Selection

- A router identifies the path a packet should take by evaluating the prefix length that is programmed in the Forwarding Information Base (FIB).
- The FIB is programmed through the routing table, which is also known as the Routing Information Base (RIB).
- This section will cover prefix length, administrative distance, metrics, equal cost multipathing, and unequal-cost load balancing.

Path Selection

Path Selection

Path selection has three main components:

- **Prefix length** - The prefix length represents the number of leading binary bits in the subnet mask that are in the on position.
- **Administrative distance** - Administrative distance (AD) is a rating of the trustworthiness of a routing information source. If a router learns about a route to a destination from more than one routing protocol, and all the routes have the same prefix length, then the AD is compared.
- **Metrics** - A metric is a unit of measure used by a routing protocol in the best-path calculation. The metrics vary from one routing protocol to another.

Path Selection

Prefix Length

Assume that a router has the following routes with various prefix lengths in the routing table:
10.0.3.0/28, 10.0.3.0/26, and 10.0.3.0/24.

Each of these routes, also known as prefix routes or simply prefixes, has a different prefix length (subnet mask). The routes are considered to be different destinations, and they will all be installed into the RIB, also known as the routing table. The routing table also includes the outgoing interface and the next-hop IP address (unless the prefix is a connected network). Table 6-2 shows this routing table.

Table 6-2
Representation of
Routing Table

| Prefix | IP Address Range | Next Hop | Outgoing Interface |
|-------------|---------------------|----------|----------------------|
| 10.0.3.0/28 | 10.0.3.0–10.0.3.15 | 10.1.1.1 | Gigabit Ethernet 1/1 |
| 10.0.3.0/26 | 10.0.3.0–10.0.3.63 | 10.2.2.2 | Gigabit Ethernet 2/2 |
| 10.0.3.0/24 | 10.0.3.0–10.0.3.255 | 10.3.3.3 | Gigabit Ethernet 3/3 |

If a packet needs to be forwarded, the route chosen depends on the prefix length, where the longest prefix length is always preferred. The forwarding decision is a function of the FIB and results from the calculations performed in the RIB. The RIB is calculated through the combination of routing protocol metrics and administrative distance.

Administrative Distance

As each routing protocol receives routing updates and other routing information, it chooses the best path to any given destination and attempts to install this path into the routing table.

The RIB is programmed from the various routing protocol processes. Every routing protocol presents the same information to the RIB for insertion: the destination network, the next-hop IP address, the AD, and metric values.

| Routing Protocol | Default Administrative Distance |
|---------------------|---------------------------------|
| Command | 0 |
| Static | 1 |
| EIGRP summary route | 5 |
| External BGP (eBGP) | 20 |
| EIGRP (internal) | 90 |
| OSPF | 110 |
| IS-IS | 115 |
| RIP | 120 |
| EIGRP (external) | 170 |
| Internal BGP (iBGP) | 200 |

The RIB accepts or rejects a route based on the following logic:

- If the route does not exist in the RIB, the route is accepted.
- If the route exists in the RIB, the AD must be compared. Routes with the lower AD values are installed in the routing table, routes with higher AD values are rejected and the submitting routing process is notified.

Administrative Distance (Cont.)

Each of these three protocols attempts to install the route to 10.3.3.0/24 into the routing table. Because the prefix length is the same, the next decision point is the AD, where the routing protocol with the lowest AD installs the route into the routing table.

Because the EIGRP internal route has the best AD, it is the one installed into the routing table, as demonstrated in Table 6-4.

Table 6-4 Route Selection for the RIB

| Routing Protocol | AD | Network | Installs in the RIB |
|------------------|-----|-------------|---------------------|
| EIGRP | 90 | 10.3.3.0/24 | ✓ |
| OSPF | 110 | 10.3.3.0/24 | X |
| IS-IS | 115 | 10.3.3.0/24 | X |

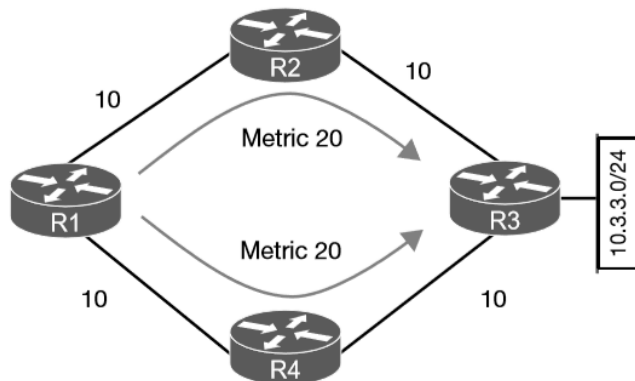
The routing protocol or protocols that failed to install their route into the table (in this example, OSPF and IS-IS) hang on to the route and tell the routing table process to report to them if the best path fails so that they can try to reinstall this route. Understanding the order of processing from a router is critical because in some scenarios the path with the lowest AD may not always be installed in the RIB.

Metrics – Equal-Cost Multipathing

The logic for selecting the best path for a routing protocol can vary. Most IGPs prefer internally learned routes over external routes and further prioritize the path with the lowest metric.

Equal-Cost Multipathing

If a routing protocol identifies multiple paths as a best path and supports multiple path entries, the router installs the maximum number of paths allowed per destination. This is known as equal-cost multipathing (ECMP) and provides load sharing across all links. RIP, EIGRP, OSPF, and IS-IS all support ECMP. ECMP provides a mechanism to increase bandwidth across multiple paths by splitting traffic equally across the links. This figure shows an example and the routing table.



Example 6-1 R1's Routing Table, Showing the ECMP Paths to 10.3.3.0/24

```
R1# show ip route
! Output omitted for brevity
O    10.3.3.0/24 [110/30] via 10.12.1.2, 00:49:12, GigabitEthernet0/2
                        [110/30] via 10.14.1.4, 00:49:51, GigabitEthernet0/4
```

Metrics- Unequal-Cost Load Balancing

By default, routing protocols install only routes with the lowest path metric. However, EIGRP can be configured to install multiple routes with different path metrics. This allows for unequal-cost load balancing across multiple paths. Traffic is transmitted out the router's interfaces based on that path's metrics in ratio to other the interface's metrics.

Figure 6-7 shows a topology with four routers running EIGRP. The delay has been incremented on R1's Gi0/2 interface from 1 μ to 10 μ . R1 sees the two paths with different metrics. The path from R1 to R3 via R1–R2–R3 has been assigned a path metric of 5632, and the path via R1–R4–R3 has been assigned a path metric of 3328.

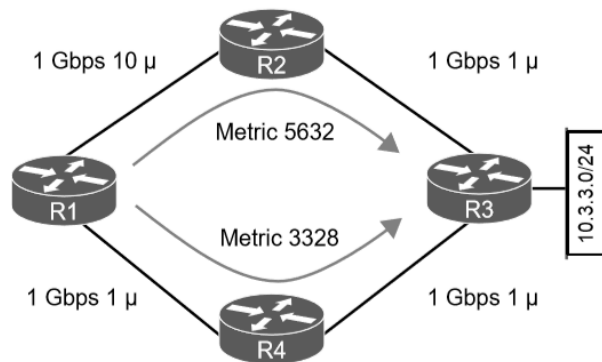


Figure 6-7 EIGRP Unequal-Cost Load Balancing

Example 6-2 R1's Routing Table, Showing the Unequal-Cost Load Balancing

```
R1# show ip route eigrp
! Output omitted for brevity
Gateway of last resort is not set

10.0.0.0/8 is variably subnetted, 7 subnets, 2 masks
D    10.3.3.0/24 [90/3328] via 10.14.1.4, 00:00:02, GigabitEthernet0/4
     [90/5632] via 10.12.1.2, 00:00:02, GigabitEthernet0/2
```

Note: The explicit path must be viewed to see the traffic ratios with unequal-cost load balancing.

Static Routing

- Using static routing requires zero network bandwidth because implementing manual route entries does not require communication with other routers.
- Because the routers are not communicating, there is no network intelligence. If a link goes down, other routers will not be aware that the network path is no longer valid.
- This section covers directly attached static routes, recursive static routes, and fully specified static routes.

Static Routing and Static Route Types

Static routes are useful when:

- Dynamic routing protocols cannot be used on a router because of limited router CPU or memory.
- Routes learned from dynamic routing protocols need to be superseded.

Static routes can be classified as one of the following:

- Directly attached static routes
- Recursive static route
- Fully specified static route

Static Routing

Directly Attached Static Routes

A static route that uses only the outbound next-hop interface is known as a directly attached static route. The outbound interface specified must be in an up state for the route to be installed into the RIB.

Note: Configuring a directly attached static route to an interface that uses Address Resolution Protocol (ARP), such as an Ethernet network, causes problems and is not recommended. The router must repeat the ARP process for every destination that matches the static route, which consumes CPU and memory. In larger networks, this can cause router instability issues.

Point-to-point (P2P) serial interfaces do not use ARP, so static routes can directly reference the outbound interface of a router. Directly attached static routes are configured with the command **ip route network subnetmask next-hop-interface-id..**

Figure 6-8 shows a P2P serial topology with R1 and R2 directly attached with a serial connection.

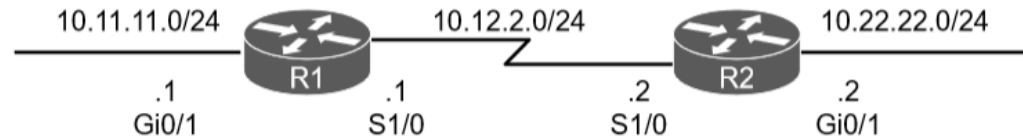


Figure 6-8 R1 and R2 Connected with a Serial Connection

Static Routing

Directly Attached Static Routes (Cont.)

Example 6-4 shows the configuration of R1 and R2 using static routes with serial 1/0 interfaces. R1 indicates that the 10.22.22.0/24 network is reachable via the S1/0 interface, and R2 indicates that the 10.11.11.0/24 network is reachable via the S1/0 interface.

Example 6-5 shows the routing table with the static route configured. A directly attached static route does not display [AD/Metric] information when looking at the routing table. Notice that the static route displays directly connected with the outbound interface.

Example 6-4 *Configuring Directly Attached Static Routes*

```
R1# configure term
Enter configuration commands, one per line. End with CNTL/Z.
R1(config)# ip route 10.22.22.0 255.255.255.0 Serial 1/0

R2# configure term
Enter configuration commands, one per line. End with CNTL/Z.
R2(config)# ip route 10.11.11.0 255.255.255.0 Serial 1/0
```

Example 6-5 *R1 and R2 Routing Table*

```
R1# show ip route
! Output omitted for brevity
Gateway of last resort is not set

    10.0.0.0/8 is variably subnetted, 5 subnets, 2 masks
C       10.11.11.0/24 is directly connected, GigabitEthernet0/1
C       10.12.2.0/24 is directly connected, Serial1/0
S       10.22.22.0/24 is directly connected, Serial1/0

R2# show ip route
! Output omitted for brevity
Gateway of last resort is not set

    10.0.0.0/8 is variably subnetted, 5 subnets, 2 masks
S       10.11.11.0/24 is directly connected, Serial1/0
C       10.12.2.0/24 is directly connected, Serial1/0
C       10.22.22.0/24 is directly connected, GigabitEthernet0/1
```

Static Routing

Recursive Static Routes

The forwarding engine on Cisco devices needs to know which interface an outbound packet should use.

- A recursive static route specifies the IP address of the next-hop address.
- The recursive lookup occurs when the router queries the RIB to locate the route toward the next-hop IP address (connected, static, or dynamic) and then cross-references the adjacency table.
- Recursive static routes are configured with the command **ip route** *network subnet-mask next-hop-ip*.
- Recursive static routes require the route's next-hop address to exist in the routing table to install the static route into the RIB. A recursive static route may not resolve the next-hop forwarding address using the default route (0.0.0.0/0) entry. The static route will fail next-hop reachability requirements and will not be inserted into the RIB.

Static Routing

Recursive Static Routes (Cont.)

Figure 6-9 shows a topology with R1 and R2 connected using the Gi0/0 port. R1 uses a recursive static route to the 10.22.22.0/24 network, and R2 uses a recursive static route to the 10.11.11.0/24 network to allow connectivity between these networks.

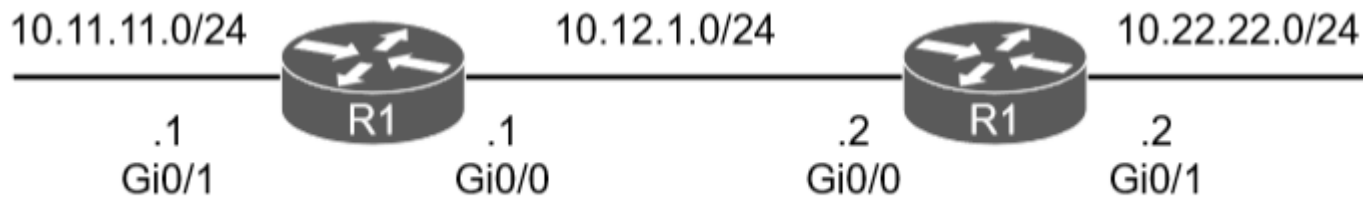


Figure 6-9 *R1 and R2 Connected by Ethernet*

In Example 6-6, R1's configuration states that the 10.22.22.0/24 network is reachable via the 10.12.1.2 IP address, and R2's configuration states that the 10.11.11.0/24 network is reachable via the 10.12.1.1 IP address.

Example 6-6 *Configuring Recursive Static Routes*

```
R1# configure term
Enter configuration commands, one per line. End with CNTL/Z.
R1(config)# ip route 10.22.22.0 255.255.255.0 10.12.1.2

R2# configure term
Enter configuration commands, one per line. End with CNTL/Z.
R2(config)# ip route 10.11.11.0 255.255.255.0 10.12.1.1
```

Static Routing

Fully Specified Static Routes

Static route recursion can simplify topologies if a link fails because it may allow the static route to stay installed while it changes to a different outbound interface in the same direction as the destination. However, problems arise if the recursive lookup resolves to a different interface pointed in the opposite direction. The following will correct the issue:

- The static route configuration should use the outbound interface and the next-hop IP address (a fully specified static route).
- Specifying the next-hop address along with the physical interface removes the recursive lookup and does not involve the ARP processing problems that occur when using only the outbound interface.
- Fully specified static routes are configured with the command **ip route network subnet-mask interface-id next-hop-ip**.

Example 6-8 Configuring Fully Specified Static Routes

```
R1# configure term
Enter configuration commands, one per line. End with CNTL/Z.
R1(config)# ip route 10.22.22.0 255.255.255.0 GigabitEthernet0/0 10.12.1.2

R2# configure term
Enter configuration commands, one per line. End with CNTL/Z.
R2(config)# ip route 10.11.11.0 255.255.255.0 GigabitEthernet0/0 10.12.1.1
```

Static Routing

Floating Static Routing

Using a floating static route is a common technique for providing backup connectivity for prefixes learned via dynamic routing protocols. A floating static route is configured with an AD higher than that of the primary route (default AD on a static route is 1). Because the AD is higher than that of the primary route, it is installed in the RIB only when the primary route is withdrawn. In Figure 6-11, R1 and R2 are configured with two links. The 10.12.1.0/24 transit network is preferred to the 10.12.2.0/24 network.

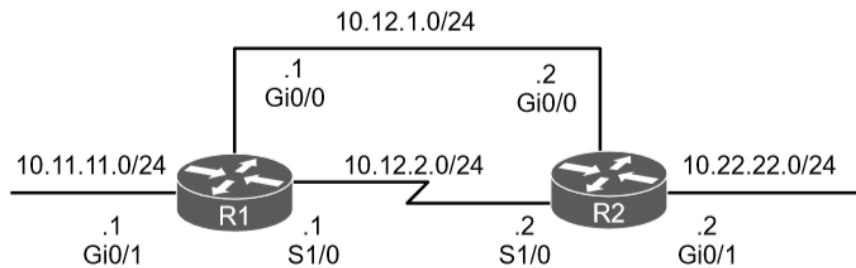


Figure 6-11 Floating Static Route Topology

Example 6-10 shows the configuration of the floating static route on R1, and R2 would be configured similarly. The static route using the Ethernet link (10.12.1.0/24) has an AD of 10, and the serial link (10.12.2.0/24) has an AD set to 210.

Example 6-10 Configuring the Floating Static Route for R1

```
R1# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
R1(config)# ip route 10.22.22.0 255.255.255.0 10.12.1.2 10
R1(config)# ip route 10.22.22.0 255.255.255.0 Serial 1/0 210
```

Static Routing

Static Null Routes

Configuring a static route to a null interface provides a method of dropping network traffic without requiring the configuration of an access list. Creating a static route to the Null0 interface is a common technique to prevent routing loops.

Figure 6-12 shows a common topology in which company ABC has acquired the 172.16.0.0/20 network range from its service provider. ABC uses only a portion of the given addresses but keeps the large network block in anticipation of future growth.

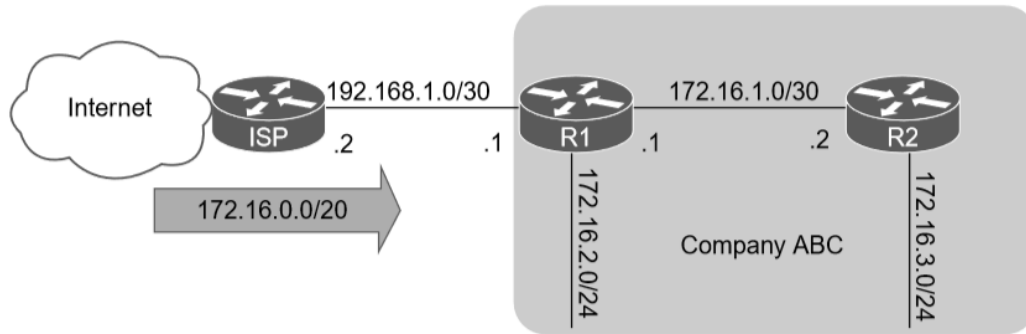


Figure 6-12 *Routing Loop Topology*

Static Routing

Static Null Routes (Cont.)

Example 6-14 shows the routing loop that occurs when packets originate from R2 addressed to an unused address on the 172.16.0.0 network.

Packets that do not match any route in the table on R1, are sent using the default route to the ISP, who then resends them back to R1. This continues until TTL expires. Notice the IP address in the traceroute alternative between the ISP router (192.168.1.2) and R1 (192.168.1.1).

To prevent the routing loop, a static route is added for 172.16.0.0/20, pointed to the Null0 interface on R1. Any packets addressed to the 172.16.0.0/20 network that are not in the routing table of R1 will be dropped. Example 6-15 shows the static null route configuration for R1.

Example 6-14 *Packet Traces Demonstrating the Routing Loop*

```
R2# trace 172.16.5.5 source GigabitEthernet 0/2
```

```
Type escape sequence to abort.
```

```
Tracing the route to 172.16.5.5
```

```
 1 172.16.1.1 0 msec 0 msec 0 msec
 2 192.168.1.1 0 msec 0 msec 0 msec
 3 192.168.1.2 0 msec 4 msec 0 msec
 4 192.168.1.1 0 msec 0 msec 0 msec
 5 192.168.1.2 0 msec 0 msec 0 msec
```

```
! Output omitted for brevity
```

Example 6-15 *R1 Static Route for 172.16.0.0/20 to Null0*

```
R1
```

```
ip route 0.0.0.0 0.0.0.0 Gi0/0 192.168.1.2
```

```
ip route 172.16.3.0 255.255.255.0 Gi0/2 172.16.1.2
```

```
ip route 172.16.0.0 255.255.240.0 Null0
```

Static Routing

IPv6 Static Routes

The static routing principles for IPv4 routes are exactly the same for IPv6. IPv6 static routes are configured with the command **ipv6 route network/prefix-length { next-hop-interface-id | [next-hop-interface-id] next-ip-address}**.

Figure 6-13 shows R1 and R2 with IPv6 addressing to demonstrate static routing.

R1 needs a static route to R2's 2001:db8:22::/64 network, and R2 needs a static route to R1's 2001:db8:11::/64 network. Example 6-17 demonstrates the IPv6 static route configuration for R1 and R2.

The IPv6 routing table is displayed with the command **show ipv6 route**. Connectivity can be verified with the **traceroute** or **ping** command.

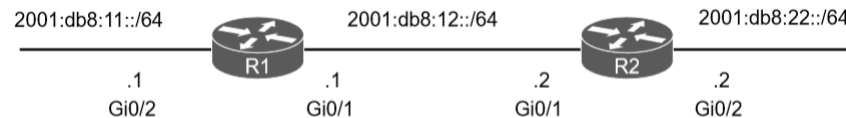


Figure 6-13 IPv6 Static Route Topology

Example 6-17 Configuring the IPv6 Static Route

```
R1# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
R1(config)# ipv6 unicast-routing
R1(config)# ipv6 route 2001:db8:22::/64 2001:db8:12::2

R2# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
R2(config)# ipv6 unicast-routing
R2(config)# ipv6 route 2001:db8:11::/64 2001:db8:12::1
```

Virtual Routing and Forwarding

- Virtual routing and forwarding (VRF) is a technology that creates separate virtual routers on a physical router. Router interfaces, routing tables, and forwarding tables are completely isolated between VRFs, preventing traffic from one VRF from forwarding into another VRF.
- All router interfaces belong to the global VRF until they are specifically assigned to a user-defined VRF.

Virtual Routing and Forwarding

Virtual Routing and Forwarding

The global VRF is identical to the regular routing table of non-VRF routers:

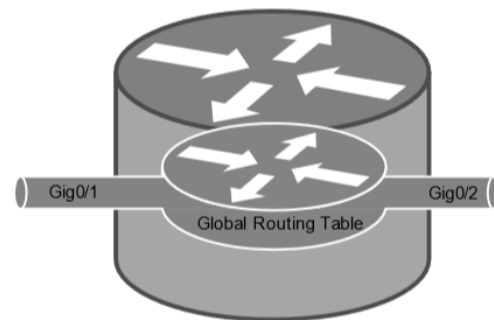
- Every router's VRF maintains a separate routing table; it is possible to allow for overlapping IP address ranges.
- VRF creates segmentation between network interfaces, network subinterfaces, IP addresses, and routing tables.
- Configuring VRF on a router ensures that the paths are isolated, network security is increased, and encrypting traffic on the network is not needed to maintain privacy between VRF instances.

Virtual Routing and Forwarding

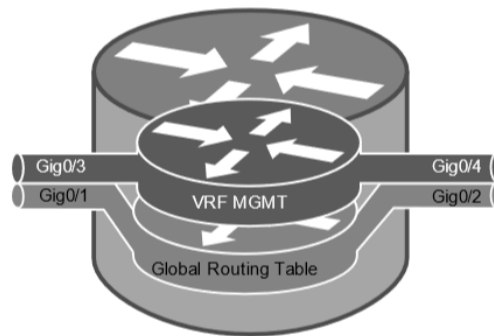
Virtual Routing and Forwarding (Cont.)

Figure 6-14 shows two routers to help visualize the VRF routing table concept. One of the routers has no VRFs configured, and the other one has a management VRF instance named MGMT.

The creation of multiprotocol VRF instances requires the global configuration command **vrf definition** *vrf-name*. Under the VRF definition submode, the command **address-family** {**ipv4** | **ipv6**} is required to specify the appropriate address family. The VRF instance is then associated to the interface with the command **vrf forwarding** *vrf-name under the interface configuration submode.*



Without VRF Configuration



With VRF Configuration

Figure 6-14 Comparison of a Router with no VRF Instances and a Router with a VRF Instance

Virtual Routing and Forwarding (Cont.)

The following steps are required to create a VRF and assign it to an interface:

Step 1. Create a multiprotocol VRF routing table by using the command **vrf definition** *vrf-name*.

Step 2. Initialize the appropriate address family by using the command **address-family** {**ipv4** | **ipv6**}. The address family can be IPv4, IPv6, or both.

Step 3. Enter interface configuration submode and specify the interface to be associated with the VRF instance by using the command **interface** *interface-id*.

Step 4. Associate the VRF instance to the interface or subinterface by entering the command **vrf forwarding** *vrf-name* under interface configuration submode.

Step 5. Configure an IP address (IPv4, IPv6, or both) on the interface or subinterface by entering either or both of the following commands:

IPv4 - **ip address** *ip-address subnet-mask* [**secondary**]

IPv6 - **ipv6 address** *ipv6-address/prefix-length*

Virtual Routing and Forwarding

Virtual Routing and Forwarding (Cont.)

Table 6-5 provides a set of interfaces and IP addresses that overlap between the global routing table and the VRF instance.

Table 6-5
Sample
Interfaces and IP
Address

| Interface | IP Address | VRF | Global |
|----------------------|-------------|------|--------|
| Gigabit Ethernet 0/1 | 10.0.3.1/24 | -- | ✓ |
| Gigabit Ethernet 0/2 | 10.0.4.1/24 | -- | ✓ |
| Gigabit Ethernet 0/3 | 10.0.3.1/24 | MGMT | -- |
| Gigabit Ethernet 0/4 | 10.0.4.1/24 | MGMT | -- |

Example 6-20 shows how the IP addresses are assigned to the interfaces in the global routing table, along with the creation of the VRF instance named MGMT and two interfaces associated with it (refer to Table 6-5). The IP addresses in the MGMT VRF instance overlap with the ones configured in the global table, but there is no conflict because they are in a different routing table.

Example 6-20 IP Address Configuration in the Global Routing Table

```
R1(config)# interface GigabitEthernet0/1
R1(config-if)# ip address 10.0.3.1 255.255.255.0
R1(config)# interface GigabitEthernet0/2
R1(config-if)# ip address 10.0.4.1 255.255.255.0
R1(config)# vrf definition MGMT
R1(config-vrf)# address-family ipv4
R1(config)# interface GigabitEthernet0/3
R1(config-if)# vrf forwarding MGMT
R1(config-if)# ip address 10.0.3.1 255.255.255.0
R1(config)# interface GigabitEthernet0/4
R1(config-if)# vrf forwarding MGMT
R1(config-if)# ip address 10.0.4.1 255.255.255.0
```

Prepare for the Exam

Prepare for the Exam

Key Topics for Chapter 6

| Description | |
|------------------------------------|-----------------------------------|
| Distance vector algorithms | Order of processing from a router |
| Distance vector perspective | Equal-cost multipathing |
| Enhanced distance vector algorithm | Unequal-cost load balancing |
| Hybrid routing protocol | Directly attached static routes |
| Link-state algorithms | Recursive static routes |
| Path vector algorithm | Fully specified static routes |
| Path selection | Floating static routing |
| Longest match | Static null routes |
| RIB route installation | IPv6 static routes |

Prepare for the Exam

Key Terms for Chapter 6

| Term | |
|---|------------------------------|
| administrative distance | link-state routing protocol |
| directly attached static route | path vector routing protocol |
| Distance vector routing protocol | prefix length |
| enhanced distance vector routing protocol | recursive static route |
| equal-cost multipathing | static null route |
| floating static route | unequal-cost load balancing |
| fully specified static route | |

Prepare for the Exam

Command Reference for Chapter 6

| Task | Command Syntax |
|--|--|
| Configure a directly attached static route | ip route network subnet-mask next-hop-interface-id |
| Configure a recursive static route | ip route subnet-mask next-hop-ip |
| Configure a fully specified static route | ip route network subnet-mask interface-id next-hop-ip |

