

Chapter 16: Route Redistribution

Instructor Materials

CCNP Enterprise: Advanced Routing



Chapter 16 Content

This chapter covers the following content:

- **Redistribution Overview** - This section provides an overview of redistribution fundamentals and rules of redistribution of routes between routing protocols.
- **Protocol-Specific Configuration** - This section explains protocol specific behaviors and configuration examples for redistribution of routes between routing protocols.

Redistribution Overview

- An organization might use multiple routing protocols, split up the routing domain between multiple instances (processes) of the same routing protocol, or need to merge networks with another organization that uses a different routing protocol.
- The routes from one routing protocol process need to be exchanged with a different routing protocol process to provide full connectivity.
- Redistribution is used to inject routes from one routing protocol into another routing protocol.

Multiple Routing Protocol Topology

Figure 16-1 illustrates a network that has multiple routing protocols that are not working together.

- R1, R2, and R3 exchange routes using Enhanced Interior Gateway Routing Protocol (EIGRP), and R3, R4, and R5 exchange routes with Open Shortest Path First (OSPF).
- R1 and R5 advertise their Loopback 0 interfaces (192.168.1.1/32 and 192.168.5.5/32) into their appropriate routing protocol, but they cannot establish connectivity to each other.
- Only R3 can connect to R1 and R5 because it is the only router that participates with both routing protocols and has a complete view of the network.

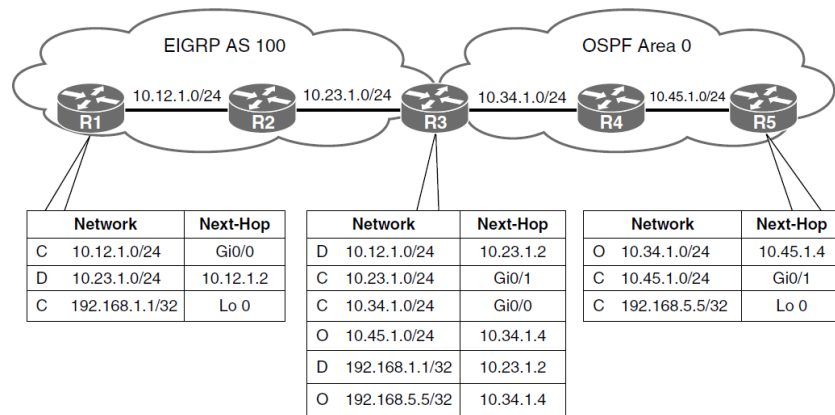


Figure 16-1 Topology with Multiple Routing Protocols

Mutual Redistribution Topology

The routes are not automatically redistributed between the routing protocols.

- Redistribution must be configured so that EIGRP routes are injected into OSPF and OSPF routes are injected into EIGRP.

In Figure 16-2, the R3 router is performing mutual redistribution. The OSPF routes are present in EIGRP as external routes, and the EIGRP routes are present in the OSPF routing domain as external routes (Type 5 link-state advertisements [LSAs]).

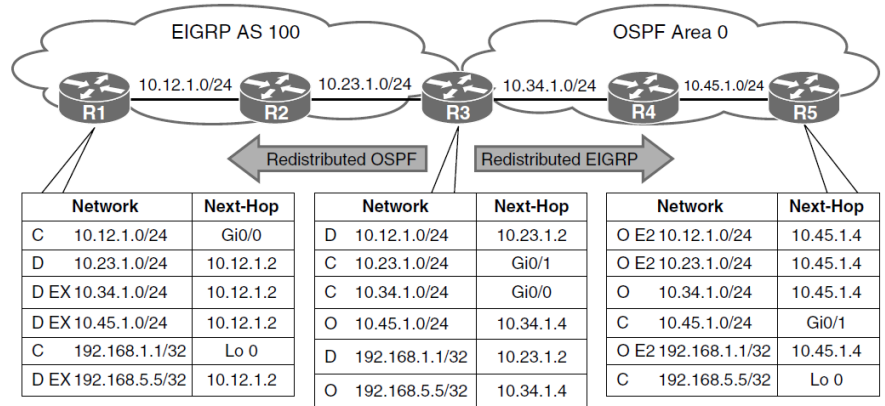


Figure 16-2 Mutual Redistribution Topology

Redistribution Source Protocols

Redistribution always encompasses two routing protocols:

- A source protocol and a destination protocol.
- The source protocol provides the network prefixes that are to be redistributed, and the destination protocol receives the injected network prefixes.
- The redistribution configuration exists under the destination protocol and identifies the source protocol.
- Using a route map allows for the filtering or modification of route attributes during the injection into the destination protocol.

Table 16-2 provides a list of source protocols for redistribution.

Table 16-2 Redistribution Source Protocols

Route Source	Description
Static	Any static route that is present in the Routing Information Base (RIB). A static route can only be a source protocol.
Connected	Any interface in an up state that is not associated with the destination protocol. A connected route can only be a source protocol.
EIGRP	Any routes in EIGRP, including EIGRP-enabled connected networks.
OSPF	Any routes in the OSPF link-state database (LSDB), including OSPF-enabled interfaces.
BGP	Any routes in the Border Gateway Protocol (BGP) Loc-RIB table learned externally. Internal BGP (iBGP) routes require the command bgp redistribute-internal for redistribution into Interior Gateway Protocol (IGP) routing protocols.

Redistribution Is Not Transitive

When redistributing between two or more routing protocols on a single router, redistribution is not transitive.

- When a router redistributes protocol 1 into protocol 2, and protocol 2 redistributes into protocol 3, the routes from protocol 1 are not redistributed into protocol 3.
- Example 16-1 provides sample logic for EIGRP mutually redistributing into OSPF and OSPF mutually redistributing into BGP.

Figure 16-3 illustrates redistribution on the router.

- The EIGRP route 172.16.1.0/24 redistributes into OSPF but does not redistribute into BGP.
- The BGP route 172.16.3.0/24 redistributes into OSPF but does not redistribute into EIGRP.
- The prefix 172.16.2.0/24 redistributes to both EIGRP and BGP.

Example 16-1 *Problematic Multiprotocol Redistribution Logic*

```
router eigrp
 redistribute ospf
router ospf
 redistribute eigrp
 redistribute bgp
router bgp
 redistribute ospf
```

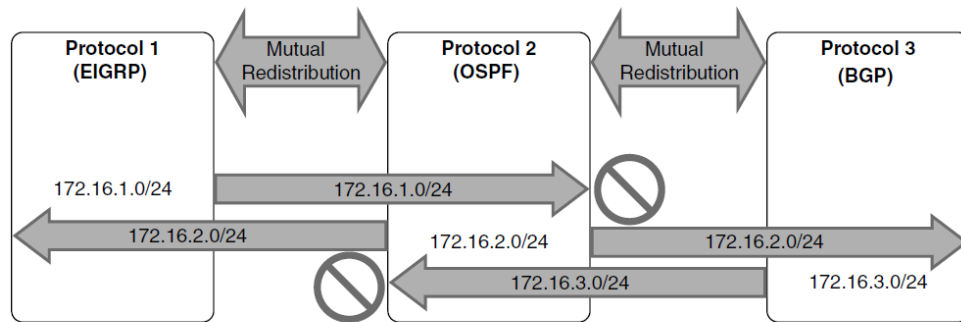


Figure 16-3 *Nontransitive Redistribution Logic*

Redistribution Is Not Transitive (Cont.)

For routes to be exchanged between all three routing protocols, mutual redistribution must be configured between all three protocols, as shown in Example 16-2.

Now that all three routing protocols are mutually redistributed, EIGRP's 172.16.1.0/24 network exists in OSPF and BGP, OSPF's 172.16.2.0/24 network exists in EIGRP and BGP, and BGP's 172.16.3.0/24 network exists in OSPF and EIGRP. Figure 16-4 illustrates the router processing for all three network prefixes.

Example 16-2 Multiprotocol Redistribution Logic

```
router eigrp
 redistribute ospf
 redistribute bgp
router ospf
 redistribute eigrp
 redistribute bgp
router bgp
 redistribute ospf
 redistribute eigrp
```

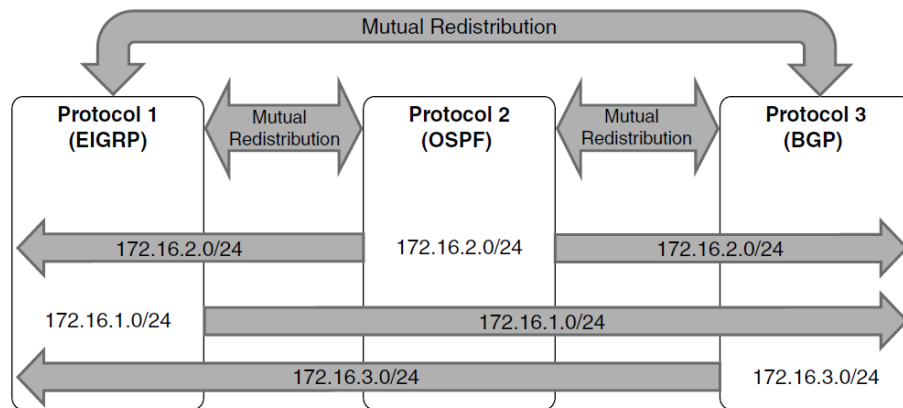


Figure 16-4 Multiprotocol Mutual Redistribution

Sequential Protocol Redistribution

Sequential protocol redistribution is redistribution between multiple protocols over a series of routers, as shown in Figure 16-5.

- R2 redistributes the EIGRP 192.168.1.1/32 prefix into BGP, and R4 redistributes the BGP 192.168.1.1/32 prefix into OSPF.
- All three routing protocols contain the 192.168.1.1/32 prefix.

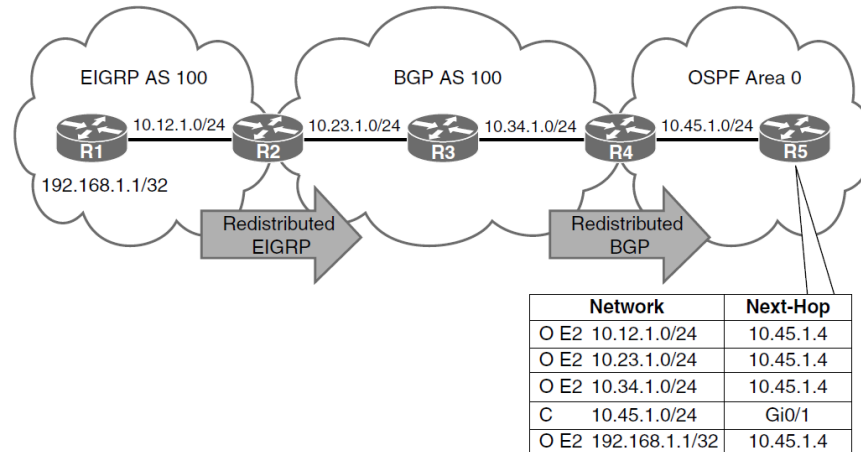


Figure 16-5 Sequential Protocol Redistribution on Different Routers

Routes Must Exist in the RIB

A route must exist in the RIB in order for it to be redistributed into the destination protocol. This ensures that the route is deemed reachable by the redistributing router.

- The source protocol that redistributes into the destination protocol must be the source for the route in the RIB. This ensures that the router redistributes only the route it deems the best for the destination protocol.
- The only exception for this logic is for directly connected interfaces participating in the source protocol because they have an administrative distance (AD) of 0.

In Figure 16-6, R1 and R3 both advertise the 10.13.1.0/24 network, and R5 is redistributing RIP and OSPF routes into EIGRP.




Figure 16-6 *Identification of Source Protocol Topology*

Routes Must Exist in the RIB (Cont.)

R5 receives route information for 10.13.1.0/24 from both RIP and OSPF routing protocols.

- Example 16-3 provides verification that R5 contains an entry for the 10.13.1.0/24 network in RIP's database and OSPF's LSDB.
- R5 determines that the most desirable path to reach 10.13.1.0/24 is the OSPF route because it has a lower AD than RIP. Example 16-4 displays R5's routing table.
- The 10.13.1.0/24 OSPF route is inserted into the RIB.
- Example 16-5 shows the EIGRP topology table for the 10.13.1.0/24 network.
- R5 checks the Routing Information Base (RIB) for the 10.13.1.0/24 network and verifies its existence in the RIB and confirms that the source protocol is the protocol that installed the route.
- EIGRP identifies the source protocol as OSPF, with a path metric of 3.

Example 16-3 *Verification of the Network in the Link-State Databases*

```
R5# show ip rip database 10.13.1.0 255.255.255.0 10.13.1.0/24
[2] via 10.25.1.2, 00:00:30, GigabitEthernet0/0

R5# show ip ospf database router 192.168.3.3
! Output omitted for brevity
      OSPF Router with ID (192.168.5.5) (Process ID 1)
      Link State ID: 192.168.3.3
      Advertising Router: 192.168.3.3
      ..
      Link connected to: a Stub Network
      (Link ID) Network/subnet number: 10.13.1.0
      (Link Data) Network Mask: 255.255.255.0
```

Example 16-4 *R5's Routing Table*

```
R5# show ip route
! Output omitted for brevity

R      10.12.1.0/24 [120/1] via 10.25.1.2, 22:46:01, GigabitEthernet0/1
O      10.13.1.0/24 [110/3] via 10.45.1.4, 00:04:27, GigabitEthernet0/0
C      10.25.1.0/24 is directly connected, GigabitEthernet0/1
O      10.34.1.0/24 [110/2] via 10.45.1.4, 22:48:24, GigabitEthernet0/0
C      10.45.1.0/24 is directly connected, GigabitEthernet0/0
C      10.56.1.0/24 is directly connected, GigabitEthernet0/2
```

Example 16-5 *EIGRP Topology Table for the 10.13.1.0/24 Network*

```
R5# show ip eigrp topology 10.13.1.0/24
! Output omitted for brevity
EIGRP-IPv4 Topology Entry for AS(100)/ID(10.56.1.5) for 10.13.1.0/24
  State is Passive, Query origin flag is 1, 1 Successor(s), FD is 2560000256
  Descriptor Blocks:
    10.45.1.4, from Redistributed, Send flag is 0x0
    External data:
      AS number of route is 1
      External protocol is OSPF, external metric is 3
```

Redistribution Overview

Seed Metrics

Every routing protocol uses a different methodology for calculating the best path for a route.

- For example, EIGRP can use bandwidth, delay, load, and reliability, whereas OSPF primarily uses the path metric for calculating the shortest path first (SPF) tree (SPT). OSPF cannot calculate the SPT using EIGRP path attributes, and EIGRP cannot run diffusing update algorithm (DUAL) using only the total path metric.
- The destination protocol must provide relevant metrics to the destination protocols so that the destination protocol can calculate the best path for the redistributed routes.
- Every protocol provides a seed metric at the time of redistribution that allows the destination protocol to calculate a best path. The seed metric is a baseline and may reflect a loss of information during redistribution when redistribution occurs between two different protocol types.
- A route map modifies the seed metric for a route during redistribution. Table 16-3 provides a list of seed metrics for the destination routing protocol.

Table 16-3 Redistribution Source Protocol Chart

Protocol	Default Seed Metric
EIGRP	Infinity. Routes set with infinity are not installed into the EIGRP topology table.
OSPF	All routes are Type 2 external. Routes sourced from BGP use a seed metric of 1, and all other protocols uses a seed metric of 20.
BGP	Origin is set to incomplete, the multi-exit discriminator (MED) is set to the IGP metric, and the weight is set to 32,768.

Protocol-Specific Configuration

Every routing protocol has a unique redistribution behavior.

Route Map Match Command

Every routing protocol has a unique redistribution behavior.

- IOS and IOS XE routers use the following command syntax in the destination protocol to identify the source routing protocol:

redistribute {**connected** | **static** | **eigrp** *as-number* | **ospf** *process-id* [**match** {**internal** | **external** [1|2]]} | **bgp** *as-number*} [*destination-protocoloptions*]
[**route-map** *route-map-name*].

- Redistribution commonly uses route maps to manipulate or filter routes on the redistributing router. Table 16-4 lists additional conditional matching commands for route selection during redistribution.

Table 16-4 Route Map *match* Command Options

match Command	Description
match interface <i>interface-type</i> <i>interface-number</i>	Selects prefixes based on the outbound interface for the selection of routes
match route-type { external [type-1 type-2] internal local nssa-external [type-1 type-2]}	Selects prefixes based on routing protocol characteristics: <ul style="list-style-type: none"> ■ external: External BGP, EIGRP, or OSPF ■ internal: Internal EIGRP or intra-area/interarea OSPF routes ■ local: Locally generated BGP routes ■ nssa-external: NSSA external (Type 7 LSAs)

Route Map Match Set Command

Table 16-5 lists the route map **set** actions that modify the route as it is redistributed into the destination protocol.

Table 16-5 Route Map set Actions

set Action	Description
set as-path prepend { <i>as-number-pattern</i> last-as 1-10}	Prepends that AS Path for the network prefix with the pattern specified or from multiple iterations from the neighboring autonomous system.
set ip next-hop { <i>ip-address</i> peer-address self}	Sets the next-hop IP address for any matching prefix. BGP dynamic manipulation requires the peer-address or self keywords.
set local-preference 0-4294967295	Sets the BGP PA local preference.

set Action	Description
set metric {+ <i>value</i> - <i>value</i> <i>value</i> } * <i>value</i> parameters are 0–4294967295	Modifies the existing metric or sets the metric for a route.
set origin {igp incomplete}	Sets the BGP PA origin.
set weight 0-65535	Sets the BGP PA weight.

Source Specific Behaviors: Connected Networks

Connected networks are networks associated with primary and secondary IP addresses for any up interfaces that are not participating with the destination protocol.

At times during redistribution, only select interfaces need to be redistributed.

- This is accomplished using a route map that selects only the desired interfaces.
- Example 16-6 provides a reference route map for selecting the specific connected network 192.168.1.1/32 on the Loopback 0 interface.

Example 16-6 *Selective Connected Network Redistribution*

```
router bgp 65100
  address-family ipv4
    redistribute connected route-map RM-LOOPBACK0
  !
route-map RM-LOOPBACK0 permit 10
  match interface Loopback0
```


Source Specific Behaviors: BGP

By default, BGP redistributes only eBGP routes into IGP protocols.

- In Figure 16-7, R3 advertises the 192.168.3.3/32 network, and R4 advertises the 192.168.4.4/32 network into BGP.
- R2 is redistributing BGP into OSPF, but only the 192.168.4.4/32 address is redistributed because it is an eBGP route. The iBGP route from R3 was not included because of BGP loop-prevention rules. It is assumed that the IGP routing topology already has a path to reach the network.
- BGP's default behavior requires that a route have an AS_Path to redistribute into an IGP routing protocol.

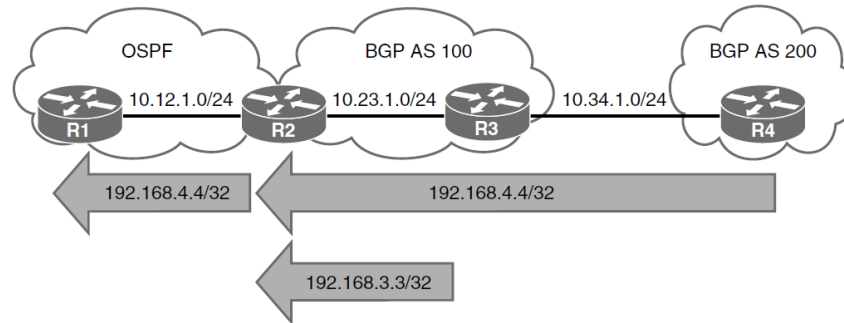


Figure 16-7 Visualizing the Redistribution of BGP Routes

Source Specific Behaviors: BGP (Cont.)

BGP behavior can be changed so that all BGP routes are redistributed with the BGP configuration command **bgp redistribute-internal**. To enable the iBGP route 192.168.3.3/32 to redistribute into OSPF, the **bgp redistribute-internal** command is required on R2.

BGP is designed to handle a large routing table, whereas IGP's are not. Redistributing BGP into an IGP on a router with a larger BGP table (for example, the Internet table with 800,000+ routes) should use selective route redistribution. Otherwise, the IGP can become unstable in the routing domain, which can lead to packet loss.

Redistributing iBGP routes into an IGP could result in routing loops. A more logical solution is to advertise the network into the IGP.

Destination Specific Behaviors: EIGRP

External EIGRP routes are given an AD of 170 and use a default seed metric of infinity, which prevents the installation of the routes into the EIGRP topology table.

- If an EIGRP autonomous system redistributes into another EIGRP autonomous system, all the path metrics are included during redistribution.
- The default path metric can be changed from infinity to specific values for bandwidth, load, delay, reliability, and maximum transmission unit (MTU), thereby allowing for the installation into the EIGRP topology table.
- Routers can set the default metric with the address family configuration command **default-metric** *bandwidth delay reliability load mtu*.
- Delay is entered in tens of microseconds (μ s).

The metric can also be set within a route map or at the time of redistribution with this command:

redistribute *source-protocol* [*metric bandwidth delay reliability load mtu*] [**route-map** *route-map-name*]

Named mode EIGRP configuration occurs in the topology base configuration.

Destination Specific Behaviors: EIGRP Topology

Figure 16-8 provides a topology example in which R2 mutually redistributes OSPF into EIGRP and R3 mutually redistributes BGP into EIGRP.

- R1 is advertising the Loopback 0 address 192.168.1.1/32, and R4 is advertising the Loopback 0 address 192.168.4.4/32.
- Example 16-7 shows the relevant EIGRP configuration. R2 uses the **default-metric** configuration command and is displayed with classic and named mode configurations.
- Notice where the default metrics are placed with EIGRP named mode configuration. R3's configuration specifies the seed metrics with the redistribution command.




Figure 16-8 EIGRP Redistribution Topology

Example 16-7 EIGRP Redistribution Configuration

R2 (AS Classic Configuration)

```
router eigrp 100
 default-metric 1000000 1 255 1 1500
 network 10.23.1.0 0.0.0.255
 redistribute ospf 1
```

R2 (Named Mode Configuration)

```
router eigrp EIGRP-NAMED
 address-family ipv4 unicast autonomous-system 100
 topology base
 default-metric 1000000 1 255 1 1500
 redistribute ospf 1
 exit-af-topology
 network 10.23.1.0 0.0.0.255
```

R3 (Named Mode Configuration)

```
router eigrp EIGRP-NAMED
 address-family ipv4 unicast autonomous-system 100
 topology base
 redistribute bgp 65100 metric 1000000 1 255 1 1500
 exit-af-topology
 network 10.23.1.0 0.0.0.255
 exit-address-family
```

Destination Specific Behaviors: EIGRP Metrics

You can overwrite EIGRP seed metrics by setting K values with the route map command **set metric bandwidth delay reliability load mtu**.

- Example 16-8 shows the configuration without the use of the **default-metric** command but by setting the EIGRP metric using a route map.
- Example 16-9 shows the EIGRP topology table with the locally redistributed routes highlighted.
- The redistributed routes are shown in the routing table with D EX and an AD of 170, as shown in Example 16-10.

Example 16-8 EIGRP Redistribution with Route Map Configuration

```
R2
router eigrp 100
 network 10.23.1.0 0.0.0.255
 redistribute ospf 1 route-map OSPF-2-EIGRP
!
route-map OSPF-2-EIGRP permit 10
 set metric 1000000 1 255 1 1500
```

Example 16-9 EIGRP Topology Table of Redistributed Routes

```
R2# show ip eigrp topology
EIGRP-IPv4 Topology Table for AS(100)/ID(192.168.2.2)
Codes: P - Passive, A - Active, U - Update, Q - Query, R - Reply,
       r - reply Status, s - sia Status

P 10.34.1.0/24, 1 successors, FD is 3072
   via 10.23.1.3 (3072/2816), GigabitEthernet0/1
P 192.168.4.4/32, 1 successors, FD is 3072, tag is 65200
   via 10.23.1.3 (3072/2816), GigabitEthernet0/1
P 10.12.1.0/24, 1 successors, FD is 2816
   via Redistributed (2816/0)
P 192.168.1.1/32, 1 successors, FD is 2816
   via Redistributed (2816/0)
P 10.23.1.0/24, 1 successors, FD is 2816
   via Connected, GigabitEthernet0/1
```

Example 16-10 Verification of External EIGRP Routes

```
R2# show ip route | begin Gateway
! Output omitted for brevity
Gateway of last resort is not set
  10.0.0.0/8 is variably subnetted, 5 subnets, 2 masks
C       10.12.1.0/24 is directly connected, GigabitEthernet0/0
C       10.23.1.0/24 is directly connected, GigabitEthernet0/1
D EX    10.34.1.0/24 [170/3072] via 10.23.1.3, 00:07:43, GigabitEthernet0/1
O       192.168.1.1 [110/2] via 10.12.1.1, 00:29:22, GigabitEthernet0/0
D EX    192.168.4.4 [170/3072] via 10.23.1.3, 00:08:49, GigabitEthernet0/1
```

```
R3# show ip route | begin Gateway
! Output omitted for brevity

D EX    10.12.1.0/24 [170/15360] via 10.23.1.2, 00:22:27, GigabitEthernet0/1
C       10.23.1.0/24 is directly connected, GigabitEthernet0/1
C       10.34.1.0/24 is directly connected, GigabitEthernet0/0
D EX    192.168.1.1 [170/15360] via 10.23.1.2, 00:22:27, GigabitEthernet0/1
B       192.168.4.4 [20/0] via 10.34.1.4, 00:13:21
```

EIGRP-to-EIGRP Redistribution

Redistributing routes between EIGRP autonomous systems preserves the path metrics during redistribution.

- Figure 16-9 shows a topology that includes multiple EIGRP autonomous systems.
- R2 mutually redistributes routes between AS 10 and AS 20, and R3 mutually redistributes routes between AS 20 and AS 30.
- R1 advertises the Loopback 0 interface (192.168.1.1/32) into EIGRP AS 10.
- R4 advertises the Loopback 0 interface (192.168.4.4/32) into EIGRP AS 30.




Figure 16-9 *Mutual EIGRP Redistribution Topology*

EIGRP-to-EIGRP Redistribution Configuration

Example 16-11 shows the configurations for R2 and R3.

- The default seed metrics do not need to be set because they are maintained between EIGRP ASs.
- R2 is using classic configuration mode, and R3 is using EIGRP named configuration mode.

Example 16-11 *EIGRP Mutual Redistribution Configuration*

```
R2
router eigrp 10
 network 10.12.1.0 0.0.0.255
 redistribute eigrp 20
router eigrp 20
 network 10.23.1.0 0.0.0.255
 redistribute eigrp 10

R3
router eigrp EIGRP-NAMED-20
 address-family ipv4 unicast autonomous-system 20
  topology base
   redistribute eigrp 30
 exit-af-topology
 network 10.23.1.0 0.0.0.255
!
router eigrp EIGRP-NAMED-30
 address-family ipv4 unicast autonomous-system 30
  topology base
   redistribute eigrp 20
 exit-af-topology
 network 10.34.1.0 0.0.0.255
exit-address-family
```

Protocol-Specific Configuration

EIGRP-to-EIGRP Redistribution Verification

Example 16-12 provides verification that R1 has routes learned from AS 20 and AS 30, and R4 has learned routes from AS 10 and AS 20.

Example 16-12 *Verification of Routes*

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

      10.0.0.0/8 is variably subnetted, 4 subnets, 2 masks
D EX    10.23.1.0/24 [170/3072] via 10.12.1.2, 00:09:07, GigabitEthernet0/0
D EX    10.34.1.0/24 [170/3328] via 10.12.1.2, 00:05:48, GigabitEthernet0/0
      192.168.4.0/32 is subnetted, 1 subnets
D EX    192.168.4.4 [170/131328] via 10.12.1.2, 00:05:48, GigabitEthernet0/0

R4# show ip route eigrp | begin Gateway
Gateway of last resort is not set

      10.0.0.0/8 is variably subnetted, 4 subnets, 2 masks
D EX    10.12.1.0/24 [170/3328] via 10.34.1.3, 00:07:31, GigabitEthernet0/0
D EX    10.23.1.0/24 [170/3072] via 10.34.1.3, 00:07:31, GigabitEthernet0/0
      192.168.1.0/32 is subnetted, 1 subnets
D EX    192.168.1.1 [170/131328] via 10.34.1.3, 00:07:31, GigabitEthernet0/0
```


Protocol-Specific Configuration

EIGRP-to-EIGRP Redistribution Verification (Cont.)

Example 16-13 shows the EIGRP topology table for the route 192.168.4.4/32 in AS 10 and AS 20.

- The EIGRP path metrics for bandwidth, reliability, load, and delay are the same between the autonomous systems.
- Notice that the feasible distance (131,072) is the same for both autonomous systems, but the reported distance (RD) is 0 for AS 10 and 130,816 for AS 20.
- The RD was reset when it was redistributed into AS 10.

Example 16-13 Topology Table for 192.168.4.4/32

```
R2# show ip eigrp topology 192.168.4.4/32
! Output omitted for brevity
EIGRP-IPv4 Topology Entry for AS(10)/ID(192.168.2.2) for 192.168.4.4/32
  State is Passive, Query origin flag is 1, 1 Successor(s), FD is 131072
  Descriptor Blocks:
    10.23.1.3, from Redistributed, Send flag is 0x0
      Composite metric is (131072/0), route is External
      Vector metric:
        Minimum bandwidth is 1000000 Kbit
        Total delay is 5020 microseconds
        Reliability is 255/255
        Load is 1/255
        Minimum MTU is 1500
        Hop count is 2
        Originating router is 192.168.2.2
      External data:
        AS number of route is 20
        External protocol is EIGRP, external metric is 131072
        Administrator tag is 0 (0x00000000)
EIGRP-IPv4 Topology Entry for AS(20)/ID(192.168.2.2) for 192.168.4.4/32
  State is Passive, Query origin flag is 1, 1 Successor(s), FD is 131072
  Descriptor Blocks:
    10.23.1.3 (GigabitEthernet0/1), from 10.23.1.3, Send flag is 0x0
      Composite metric is (131072/130816), route is External
      Vector metric:
        Minimum bandwidth is 1000000 Kbit
        Total delay is 5020 microseconds
        Reliability is 255/255
        Load is 1/255
        Minimum MTU is 1500
        Hop count is 2
        Originating router is 192.168.3.3
      External data:
        AS number of route is 30
        External protocol is EIGRP, external metric is 2570240
```

Destination Specific Behaviors: OSPF

The AD is set to 110 for intra-area, interarea, and external OSPF routes. External OSPF routes are classified as Type 1 or Type 2, with Type 2 as the default setting.

- The seed metric is 1 for BGP-sourced routes and 20 for all other protocols.
- The exception is that if OSPF redistributes from another OSPF process, the path metric is transferred.

The main differences between Type 1 and Type 2 External OSPF routes follow:

- Type 1 routes are preferred over Type 2.
 - The Type 1 metric equals the redistribution metric plus the total path metric to the autonomous system boundary router (ASBR). In other words, as the LSA propagates away from the originating ASBR, the metric increases.
 - The Type 2 metric equals only the redistribution metric. The metric is the same for the router next to the ASBR as the router 30 hops away from the originating ASBR.
 - If two Type 2 paths have exactly the same metric, the lower forwarding cost is preferred.
- This is the default external metric type used by OSPF.

Protocol-Specific Configuration

Redistribution into OSPF

For redistribution into OSPF, you use this command:

```
redistribute source-protocol [subnets] [metric metric] [metric-type {1 | 2}]  
[tag 0-4294967295] [route-map route-map-name]
```

- If the optional **subnets** keyword is not included, only the classful networks are redistributed.
- The optional **tag** keyword allows for a 32-bit route tag to be included on all redistributed routes.
- The **metric** and **metric-type** keywords can be set during redistribution.

Figure 16-10 provides a topology example in which R2 mutually redistributes EIGRP into OSPF and R3 mutually redistributes RIP into OSPF. R1 is advertising the Loopback 0 interface 192.168.1.1/32, and R4 is advertising the Loopback 0 interface 192.168.4.4/32.




Figure 16-10 *OSPF Redistribution Topology*

Redistribution into OSPF Configuration

Example 16-14 shows the relevant OSPF configuration. Notice that R2 and R3 use different OSPF process numbers but are still able to form an adjacency. The OSPF process numbers are locally significant in linking the OSPF-enabled interfaces to a process, as shown later in this section.

Example 16-14 *OSPF Redistribution Configuration*

```
R2
router ospf 2
  router-id 192.168.2.2
  network 10.23.1.0 0.0.0.255 area 0
  redistribute eigrp 100 subnets
```

```
R3
router ospf 3
  router-id 192.168.3.3
  redistribute rip subnets
  network 10.23.1.3 0.0.0.0 area 0
```

Protocol-Specific Configuration

Redistribution into OSPF LSDB

Example 16-15 shows the Type 5 LSAs for the external networks in the OSPF domain.

- This example shows that the routes 10.12.1.0/24, 10.34.1.0/24, 192.168.1.1/32, and 192.168.4.4/32 successfully redistributed into OSPF.
- The redistributed networks are Type 2 with a metric of 20.
- The redistributed routes appear in the routing table with O E2 for Type 2 and O E1 for Type 1 external routes.

Example 16-15 *OSPF LSDB from R3*

```
R3# show ip ospf database external
! Output omitted for brevity

OSPF Router with ID (192.168.3.3) (Process ID 2)

Type-5 AS External Link States

Link State ID: 10.12.1.0 (External Network Number )
Advertising Router: 192.168.2.2
Network Mask: /24
Metric Type: 2 (Larger than any link state path)
Metric: 20

Link State ID: 10.34.1.0 (External Network Number )
Advertising Router: 192.168.3.3
Network Mask: /24
Metric Type: 2 (Larger than any link state path)
Metric: 20

Link State ID: 192.168.1.1 (External Network Number )
Advertising Router: 10.23.1.2
Network Mask: /32
Metric Type: 2 (Larger than any link state path)
Metric: 20

Link State ID: 192.168.4.4 (External Network Number )
Advertising Router: 192.168.3.3
Network Mask: /32
Metric Type: 2 (Larger than any link state path)
Metric: 20
```

Protocol-Specific Configuration

Redistribution into OSPF Verification

In Example 16-16, the routers do not explicitly set a metric type, so all the redistributed routes from the topology are type E2.

Example 16-16 OSPF Route Redistribution Verification

```
R2# show ip route | begin Gateway
Gateway of last resort is not set

    10.0.0.0/8 is variably subnetted, 5 subnets, 2 masks
C       10.12.1.0/24 is directly connected, GigabitEthernet0/0
C       10.23.1.0/24 is directly connected, GigabitEthernet0/1
O E2    10.34.1.0/24 [110/20] via 10.23.1.3, 00:04:44, GigabitEthernet0/1
        192.168.1.0/32 is subnetted, 1 subnets
D       192.168.1.1 [90/130816] via 10.12.1.1, 00:03:56, GigabitEthernet0/0
        192.168.2.0/32 is subnetted, 1 subnets
C       192.168.2.2 is directly connected, Loopback0
O E2    192.168.4.0/24 [110/20] via 10.23.1.3, 00:04:42, GigabitEthernet0/1

R3# show ip route | begin Gateway
Gateway of last resort is not set

    10.0.0.0/8 is variably subnetted, 5 subnets, 2 masks
O E2    10.12.1.0/24 [110/20] via 10.23.1.2, 00:05:41, GigabitEthernet0/1
C       10.23.1.0/24 is directly connected, GigabitEthernet0/1
C       10.34.1.0/24 is directly connected, GigabitEthernet0/0
        192.168.1.0/32 is subnetted, 1 subnets
O E2    192.168.1.1 [110/20] via 10.23.1.2, 00:05:41, GigabitEthernet0/1
        192.168.3.0/32 is subnetted, 1 subnets
C       192.168.3.3 is directly connected, Loopback0
R       192.168.4.0/24 [120/1] via 10.34.1.4, 00:00:00, GigabitEthernet0/0
```

OSPF-to-OSPF Redistribution

Redistributing routes between OSPF processes preserves the path metric during redistribution, independent of the metric type.

Figure 16-11 shows a topology with multiple OSPF processes and areas.

Although the topology in Figure 16-11 looks discontinuous, OSPF is redistributing the routes between processes. This technique can be used to advertise routes over discontinuous OSPF networks, but it results in the loss of path information as the Type 1, Type 2, and Type 3 LSAs are not propagated through route redistribution.

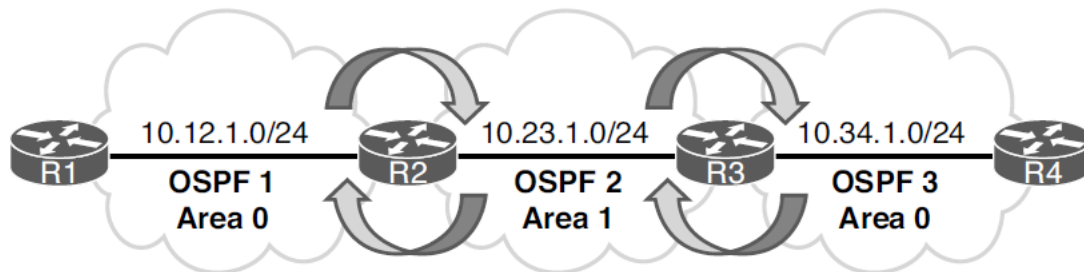


Figure 16-11 *OSPF Multiprocess Redistribution*

Protocol-Specific Configuration

OSPF-to-OSPF Redistribution (Cont.)

Example 16-17 shows the relevant configurations for R2 and R3. Notice that the metric type is set at the time of redistribution into the destination protocol so that you can see the metric increase as the route travels between the OSPF processes.

Example 16-18 provides verification that R1 has routes learned from OSPF process 3 (R3 and R4) and that R4 has routes learned from OSPF process 1 (R1 and R2). Notice that the metrics have carried over through the redistribution.

Example 16-17 *OSPF Multiprocess Redistribution*

```
R2# show running-config | section router ospf
router ospf 1
 redistribute ospf 2 subnets metric-type 1
 network 10.12.1.0 0.0.0.255 area 0
router ospf 2
 redistribute ospf 1 subnets metric-type 1
```

```
R3# show running-config | section router ospf
router ospf 2
 redistribute ospf 3 subnets metric-type 1
 network 10.23.1.0 0.0.0.255 area 1
router ospf 3
 redistribute ospf 2 subnets metric-type 1
 network 10.34.1.0 0.0.0.255 area 0
```

Example 16-18 *Verification of OSPF Redistribution*

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

    10.0.0.0/8 is variably subnetted, 4 subnets, 2 masks
O E1   10.23.1.0/24 [110/2] via 10.12.1.2, 00:00:21, GigabitEthernet0/0
O E1   10.34.1.0/24 [110/3] via 10.12.1.2, 00:00:21, GigabitEthernet0/0
    192.168.4.0/32 is subnetted, 1 subnets
O E1   192.168.4.4 [110/4] via 10.12.1.2, 00:00:21, GigabitEthernet0/0
```

```
R4# show ip route ospf | begin Gateway
Gateway of last resort is not set

    10.0.0.0/8 is variably subnetted, 4 subnets, 2 masks
O E1   10.12.1.0/24 [110/3] via 10.34.1.3, 00:01:36, GigabitEthernet0/0
O E1   10.23.1.0/24 [110/2] via 10.34.1.3, 00:01:46, GigabitEthernet0/0
    192.168.1.0/32 is subnetted, 1 subnets
O E1   192.168.1.1 [110/4] via 10.34.1.3, 02:38:49, GigabitEthernet0/0
```


Protocol-Specific Configuration

OSPF Forwarding Address

OSPF Type 5 LSAs include a field known as the forwarding address that optimizes forwarding traffic when the source uses a shared network segment.

- The scenario defined in RFC 2328 is not common but is shown in Figure 16-12.
- Example 16-19 shows the Type 5 LSA for the AS 100 route for 192.168.1.1/32.
- Network traffic from R3 (and R5) takes the suboptimal route (R3→R5→R4→R2→R1), as shown in Example 16-20. The optimal route would use the directly connected 10.123.1.0/24 network.




Figure 16-12 OSPF Forwarding Address Set to the Default

Example 16-19 OSPF External LSA with Forwarding Address 0.0.0.0

```
R3# show ip ospf database external
! Output omitted for brevity
Type-5 AS External Link States

Routing Bit Set on this LSA in topology Base with MTID 0
LS Type: AS External Link
Link State ID: 192.168.1.1 (External Network Number)
Advertising Router: 10.123.1.2
Network Mask: /32
Metric Type: 2 (Larger than any link state path)
Metric: 1
Forward Address: 0.0.0.0
```

Example 16-20 Verification of Suboptimal Routing

```
R3# trace 192.168.1.1
Tracing the route to 192.168.1.1
 1 10.35.1.5 0 msec 0 msec 1 msec
 2 10.45.1.4 0 msec 0 msec 0 msec
 3 10.24.1.2 1 msec 0 msec 0 msec
 4 10.123.1.1 1 msec * 0 msec

R5# trace 192.168.1.1
Tracing the route to 192.168.1.1
 1 10.45.1.4 0 msec 0 msec 0 msec
 2 10.24.1.2 1 msec 0 msec 0 msec
 3 10.123.1.1 1 msec * 0 msec
```

Protocol-Specific Configuration

OSPF Forwarding Nondefault

The OSPF forwarding address changes from 0.0.0.0 to the next-hop IP address in the source routing protocol when:

- OSPF is enabled on the ASBR's interface that points to the next-hop IP address.
- That interface is not set to passive.
- That interface is a broadcast or nonbroadcast OSPF network type.

When the forwarding address is set to a value besides 0.0.0.0, the OSPF routers forward traffic only to the forwarding address.

Example 16-21 provides the Type 5 LSA for the 192.168.1.1/32 network. Now that OSPF has been enabled on R2's 10.123.1.2 interface and the interface is a broadcast network type, the forwarding address has changed from 0.0.0.0 to 10.123.1.1.




Figure 16-13 OSPF Forwarding Address Set to Nondefault

Example 16-21 OSPF External LSA with Forwarding Address 10.123.1.1

```
R3# show ip ospf database external
! Output omitted for brevity
Type-5 AS External Link States

Options: (No TOS-capability, DC)
LS Type: AS External Link
Link State ID: 192.168.1.1 (External Network Number )
Advertising Router: 10.123.1.2
Network Mask: /32

Metric Type: 2 (Larger than any link state path)
Metric: 1
Forward Address: 10.123.1.1
```

OSPF Forwarding Address Verification

Example 16-22 verifies that connectivity from R3 and R5 now takes the optimal path to R1 because the forwarding address has changed to 10.123.1.1.

- If the Type 5 LSA forwarding address is not a default value, the address must be an intra-area or interarea OSPF route.
- If the route does not exist, the LSA is ignored and is not installed into the RIB.

The OSPF forwarding address optimizes forwarding toward the destination network, but return traffic is unaffected. In Figure 16-13, outbound traffic from R3 or R5 still exits at R3's Gi0/0 interface, but return traffic is sent directly to R2.

Example 16-22 *Verification of Optimal Routing*

```
R3# trace 192.168.1.1
Tracing the route to 192.168.1.1
  0  10.123.1.1  0 msec *  1 msec

R5# trace 192.168.1.1
Tracing the route to 192.168.1.1
  0  10.35.1.3  0 msec  0 msec  1 msec
  1  10.123.1.1  0 msec *  1 msec
```




Figure 16-13 *OSPF Forwarding Address Set to Nondefault*

Protocol-Specific Configuration

Redistribution into BGP

Redistributing routes into BGP does not require a seed metric because BGP is a path vector protocol. Redistributed routes have the following BGP attributes set:

- The origin is set to incomplete.
- The next-hop address is set to the IP address of the source protocol.
- The weight is set to 32,768.
- The MED is set to the path metric of the source protocol.

Figure 16-14 shows a topology in which R2 mutually redistributes between OSPF and BGP, and R3 mutually redistributes between EIGRP AS 65100 and BGP.




Figure 16-14 *BGP Redistribution Topology*

Redistribution into BGP Configuration

Example 16-23 shows R2's BGP configuration for redistributing OSPF into BGP on R2 and R3's configuration for redistributing EIGRP into BGP.

- R3 has disabled the default IPv4 address family configuration.
- Notice that R2 and R3 have used the command **bgp redistribute-internal**, which allows for any iBGP learned prefixes to be redistributed into OSPF or EIGRP.

Example 16-24 shows the BGP table for AS 65100. Notice that the 192.168.1.1/32 and 192.168.4.4/32 networks have been installed into the BGP table. The metric is carried over from the IGP metric during redistribution.

Example 16-24 BGP Route Table

R2#	show bgp ipv4 unicast	begin	Network	Next Hop	Metric	LocPrf	Weight	Path
*>	10.12.1.0/24		0.0.0.0		0		32768	?
* i	10.23.1.0/24		10.23.1.3		0	100	0	i
*>			0.0.0.0		0		32768	i
*>i	10.34.1.0/24		10.23.1.3		0	100	0	?
*>	192.168.1.1/32		10.12.1.1		2		32768	?
*>i	192.168.4.4/32		10.34.1.4		130816	100	0	?

Example 16-23 BGP Redistribution Configuration

```
R2 (Default IPv4 Address Family Enabled)
router bgp 65100
  bgp redistribute-internal
  network 10.23.1.0 mask 255.255.255.0
  redistribute ospf 1
  neighbor 10.23.1.3 remote-as 65100
```

```
R3 (Default IPv4 Address Family Disabled)
router bgp 65100
  no bgp default ipv4-unicast
  neighbor 10.23.1.2 remote-as 65100
  !
  address-family ipv4
    bgp redistribute-internal
    network 10.23.1.0 mask 255.255.255.0
    redistribute eigrp 100
    neighbor 10.23.1.2 activate
  exit-address-family
```

Redistribution into BGP Verification

Example 16-25 shows detailed BGP path information for the redistributed routes. The origin is incomplete, and the BGP metric matches the IGP metric at the time of redistribution.

Redistributing routes from OSPF to BGP does not include OSPF external routes by default. **match external [1 | 2]** is required to redistribute OSPF external routes.

Example 16-25 *Verification of BGP Routes*

```
R2# show bgp ipv4 unicast 192.168.1.1
```

```
! Output omitted for brevity
```

```
BGP routing table entry for 192.168.1.1/32, version 3
```

```
Paths: (1 available, best #1, table default)
```

```
Local
```

```
10.12.1.1 from 0.0.0.0 (192.168.2.2)
```

```
Origin incomplete, metric 2, localpref 100, weight 32768, valid, sourced, best
```

```
R3# show bgp ipv4 unicast 192.168.4.4
```

```
BGP routing table entry for 192.168.4.4/32, version 3
```

```
Paths: (1 available, best #1, table default)
```

```
Local
```

```
10.34.1.4 from 0.0.0.0 (10.34.1.3)
```

```
Origin incomplete, metric 130816, localpref 100, weight 32768, valid, sourced,  
best
```

Prepare for the Exam

Prepare for the Exam

Key Topics for Chapter 16

Description	Description
Redistribution terminology	Destination protocol: EIGRP
Redistribution is not transitive	EIGRP seed metrics
Sequential protocol redistribution	EIGRP-to-EIGRP redistribution
Routes must exist in the RIB	Destination protocol: OSPF
Seed metrics	OSPF classful and classless redistribution
Redistribution source protocol chart	OSPF-to-OSPF redistribution
Source protocol: connected networks	OSPF forwarding address
Source protocol: BGP	Destination protocol: BGP

Prepare for the Exam

Key Terms for Chapter 16

Terms
destination protocol
mutual redistribution
source protocol
sequential protocol redistribution
seed metric

Prepare for the Exam

Command Reference for Chapter 16

Terms

Redistribute the source routing protocol into a destination routing protocol

redistribute {**connected** | **static** | **eigrp** *as-number* | **ospf** *process-id* [**match** {**internal** | **external** [1|2]}} | **bgp** *as-number* } [*destinationprotocol-options*] [**route-map** *route-map-name*]

Set the default EIGRP seed metric for redistributed prefixes

default-metric *bandwidth delay reliability load mtu*

Allow for the redistribution of iBGP learned prefixes into an IGP

bgp redistribute-internal

