Chapter 17: Troubleshoot Redistribution

Instructor Materials

CCNP Enterprise: Advanced Routing

Chapter 17 Content

This chapter covers the following content:

- **Troubleshooting Advanced Redistribution Issues -** This section explains how suboptimal routing and routing loops may occur when redistributing at multiple points in the network. In addition, you will discover how to recognize these redistribution issues and solve them.
- **Troubleshooting IPv4 and IPv6 Redistribution -** This section examines the issues that you should look out for when troubleshooting redistribution for IPv4 and IPv6 routing protocols such as EIGRP, OSPF, and BGP.
- **Redistribution Trouble Tickets -** This section provides trouble tickets that demonstrate how to use a structured troubleshooting process to solve a reported problem.

Troubleshoot Advanced Redistribution Issues

- Highly available network designs remove single points of failure through redundancy. When redistributing routes between protocols, there must be at least two points of redistribution in the network to ensure that connectivity is maintained during a failure.
- When performing multipoint redistribution between two protocols, the following issues may arise:
  - Suboptimal routing
  - Routing loops
- These issues can lead to loss of connectivity or slow connectivity for the end users.

Troubleshooting Suboptimal Routing Caused by Redistribution

If there are multiple points of redistribution between two sources, as shown in Figure 17-1, the suboptimal path may be chosen to reach networks.

- In Figure 17-1, focus on R1 and R2. The optimal path to reach 192.168.2.0/24 is from R2 because the 1 Gbps link is much faster than the 10 Mbps link.
  - When you perform redistribution on R1 and R2 into EIGRP, EIGRP does not know that the 10 Mbps link or the 1 Gbps link exists in the OSPF domain. Therefore, if an inappropriate seed metric is used during redistribution on R1 and R2, the traffic from 10.1.1.0/24 destined for 192.168.2.0/24 may take the suboptimal path through R1.

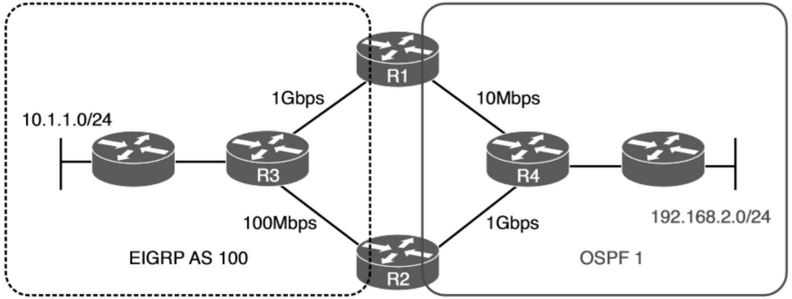

Figure 17-1 *Suboptimal Routing Topology*

- According to the EIGRP AS, it is the best path because all it sees is the seed metric and the 1 Gbps and 100 Mbps link in the EIGRP autonomous system. Therefore, if the seed metrics you define are the same on R1 and R2 when you redistribute into EIGRP, the 1 Gbps link in the EIGRP autonomous system is preferred, and traffic goes to R1. Then R1 sends it across the 10 Mbps link to 192.168.2.0/24, which is suboptimal. It works, but it is suboptimal.

Troubleshooting Suboptimal Routing Caused by Redistribution (Cont.)

In Figure 17-1, if the result of the traceroute from 10.1.1.0/24 to 192.168.2.0/24 goes through R1, you know that suboptimal routing is occurring because of redistribution.

You can solve this issue by providing different seed metrics on the boundary routers (R1 and R2 in this case) to ensure that a certain path is preferred because it has a lower overall metric. So, R2's EIGRP seed metric must be significantly better than R1's EIGRP seed metric to ensure that R3 chooses the path through R2, even though it is a slower link between R3 and R2 than between R3 and R1. The key is to make sure the traffic avoids the 10 Mbps link.

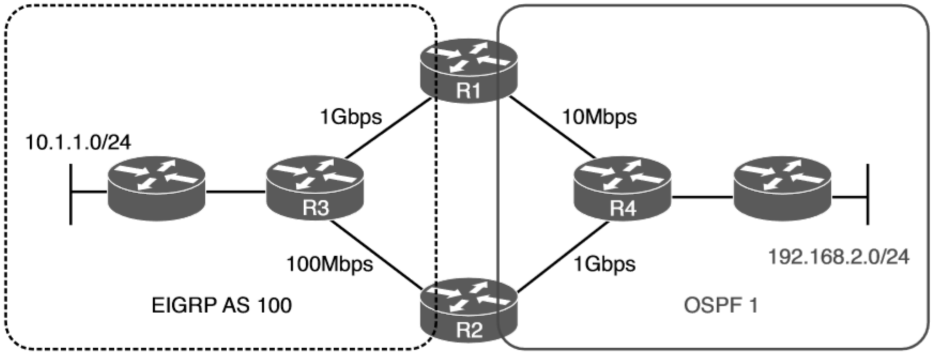

Figure 17-1 *Suboptimal Routing Topology*

Troubleshooting Suboptimal Routing Caused by Redistribution (Cont.)

When troubleshooting suboptimal routing caused by redistribution, keep in mind the following:

- Based on the topology, you need to be able to recognize that mutual redistribution is occurring at multiple points in the network.
- Based on the connections, you need to be able to recognize the different speeds of the links.
- Based on the routing protocols in use, you need to be able to identify how the seed metric is determined and how it behaves for the different protocols.
- Based on the business requirements, you need to know how to fix the suboptimal routing by manipulating the metrics on the boundary routers with the **default-metric** command, the **metric** parameter in the **redistribute** command, or within a route map.

Troubleshooting Routing Loops Caused by Redistribution

Examine Figure 17-2. The 10.1.1.0/24 network is redistributed into the EIGRP autonomous system, and then it is redistributed into the OSPF domain on R1 and R2. This does not appear to be an issue; however, it is an issue because of administrative distance (AD).

When the 10.1.1.0/24 network is redistributed from RIPv2 into EIGRP autonomous system 100, it is classified as an external route in the EIGRP autonomous system. R1 and R2 place the route in the routing table with the code D EX and an AD of 170, as shown in Figure 17-3.

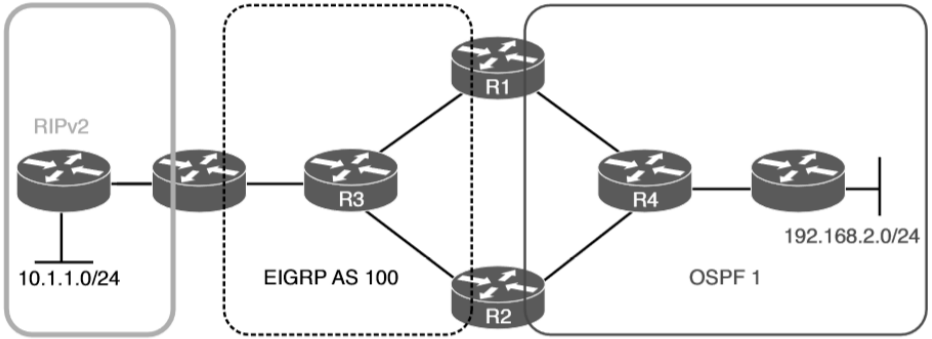

Figure 17-2 *Routing Loop Routing Topology*

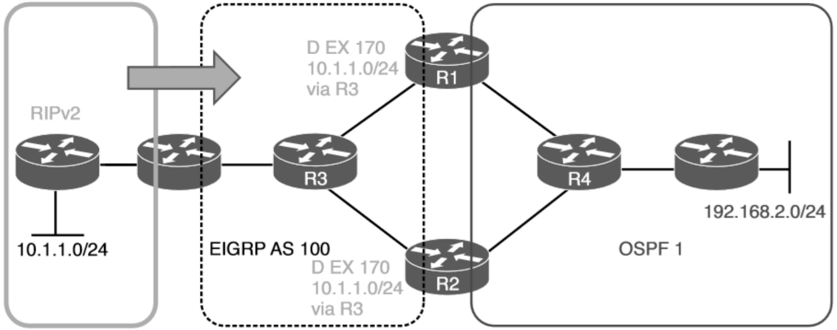

Figure 17-3 *Redistributing 10.1.1.0/24 into the EIGRP Autonomous System*

Troubleshooting Routing Loops Caused by Redistribution (Cont.)

- When R1 and R2 redistribute the 10.1.1.0/24 network in the OSPF domain, by default, the Type 5 link-state advertisement (LSA) is advertising 10.1.1.0/24 as an O E2 route, with an AD of 110, as shown in Figure 17-4.
- R1 will receive R2's LSA, and R2 will receive R1's LSA, which creates the problem.

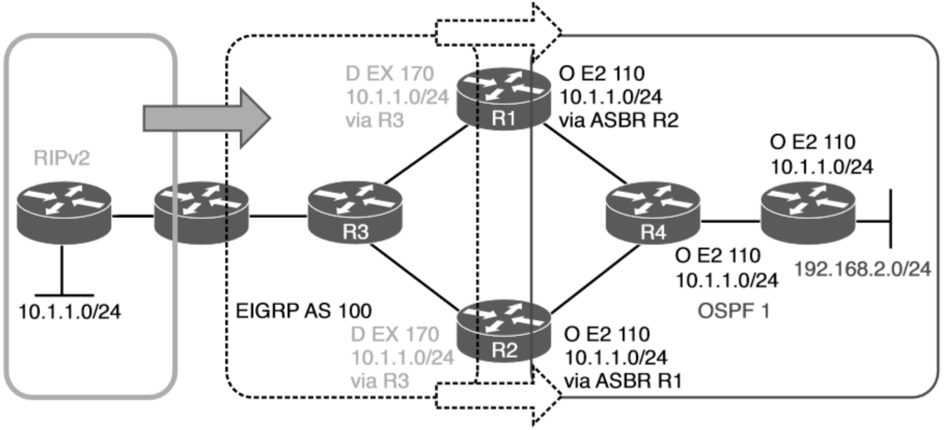

Figure 17-4 *Redistributing 10.1.1.0/24 into the OSPF Domain on R1 and R2*

- The OSPF route is preferred because it has a lower AD. Therefore, R1 points to R2 through the OSPF domain to reach 10.1.1.0/24
- Now when traffic is sent from 192.168.2.0/24 to 10.1.1.0/24, it bounces back and forth between R1 and R2, and this is classified as a routing loop.

Troubleshooting Routing Loops Caused by Redistribution (Cont.)

This scenario gets worse because of how redistribution works:

- To redistribute a route from one routing source to another (EIGRP into OSPF, for example), that route must be in the routing table as an entry for the routing source that you are redistributing the route from.
- In Figure 17-4 when R1 and R2 originally learned about the network 10.1.1.0/24 from R3, it was an EIGRP external route.
- There was no other source of information in the routing table at the time for 10.1.1.0/24; therefore, it was considered the best source and installed in the routing table as an EIGRP route.

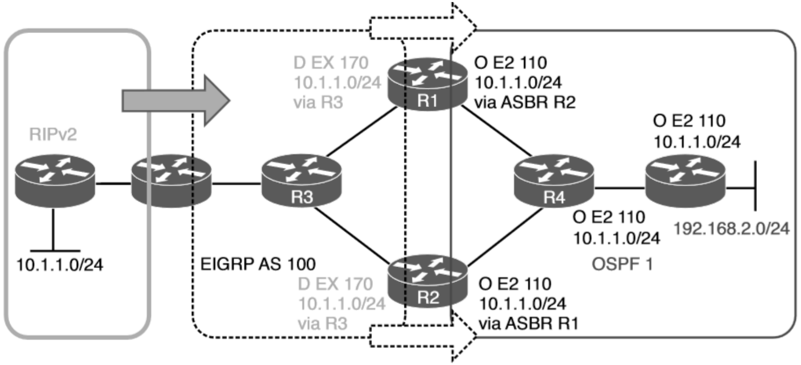

Figure 17-4 *Redistributing 10.1.1.0/24 into the OSPF Domain on R1 and R2*

- Because redistribution is occurring from EIGRP into OSPF, the 10.1.1.0/24 network is redistributed from the routing table into the OSPF process and advertised.

Troubleshooting Routing Loops Caused by Redistribution (Cont.)

When R1 and R2 learn about the OSPF 10.1.1.0/24 route from each other, they notice that it is a better source of information because the AD is lower (110) than the one for EIGRP (170) that is currently in the routing table. Therefore, the OSPF route replaces the EIGRP route.

Because the EIGRP route is still in the topology table but not in the routing table, it is no longer available for redistribution into OSPF, and therefore, there are no more Type 5 LSAs to advertise in the OSPF domain. As a result, R1 and R2 have to notify the routers in the OSPF domain that 10.1.1.0/24 no longer exists. When this happens, R1 and R2 no longer have the 10.1.1.0/24 network that they learned through OSPF from each other in the routing table.

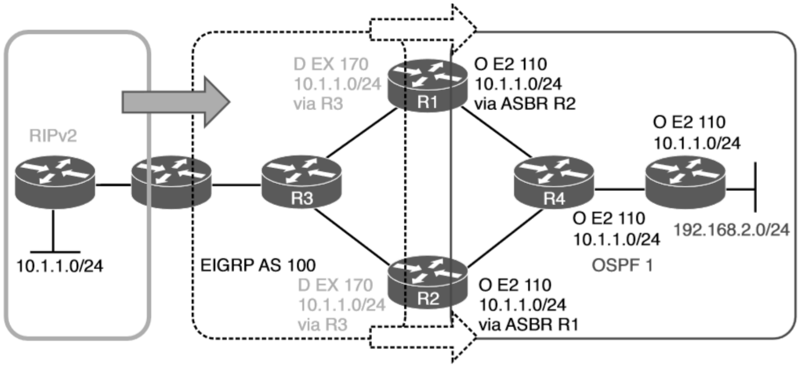

Figure 17-4 *Redistributing 10.1.1.0/24 into the OSPF Domain on R1 and R2*

The EIGRP external route 10.1.1.0/24 is reinstalled in the routing table, and because redistribution from EIGRP into OSPF is occurring, the issue repeats all over again. The routing table is not stable.

Troubleshooting Routing Loops Caused by Redistribution (Cont.)

Examine Figure 17-5, which shows the 10.1.1.0/24 network being redistributed back into the EIGRP autonomous system on R1 and R2 when the OSPF route is in the routing table on R1 and R2. This is known as route feedback.

Now R3 thinks that 10.1.1.0/24 is reachable through the boundary router (R5) between the RIPv2 domain and the EIGRP autonomous system, as well as through R1 and R2 between EIGRP and OSPF. Depending on the metric for each of the learned paths to 10.1.1.0/24, R3 may choose the correct path to the RIP domain or the path through R1 or R2, which would eventually blackhole the traffic.

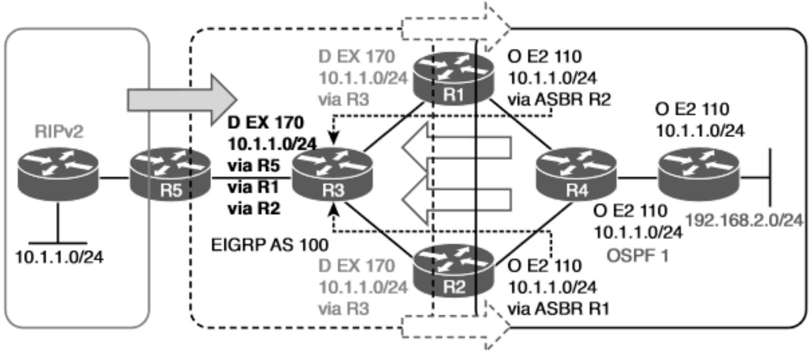

Figure 17-5 *Redistributing 10.1.1.0/24 Back into the EIGRP Autonomous System from OSPF*

Remember that this issue was caused by AD. Either lower the AD of the EIGRP external routes on R1 and R2 for 10.1.1.0/24 or increase the AD of the OSPF Type 5 learned routes on R1 and R2 for 10.1.1.0/24. Make sure the EIGRP learned route is the preferred route. Regardless, use the distance command on R1 and R2 and specify what the AD will be for the 10.1.1.0/24 network.

Troubleshooting Routing Loops Caused by Redistribution (Cont.)

EIGRP already differentiates between routes learned from within the autonomous system and routes learned from outside the autonomous system by assigning different administrative distance:

- Internal EIGRP: 90
- External EIGRP: 170

To modify the default administrative distance on IOS routers, use the **distance eigrp** *ad-internal ad-external* EIGRP configuration command . Valid values for the AD are between 1 and 255; a value of 255 stops the installation of the route into the Routing Information Base (RIB).

Cisco IOS routers also allow selective AD modification for specific internal networks with this command:

distance *ad source-ip source-ip-wildcard [acl-number* | *acl-name]*

The *source-ip* option restricts the modification of routes in the EIGRP table that were learned from a specific router, and the optional *acl* restricts to a specific network prefix. Note that EIGRP does not allow the selective AD modification based on prefixes for external EIGRP routes.

Troubleshooting Routing Loops Caused by Redistribution (Cont.)

OSPF uses the same default AD 110 value for routes learned within the OSPF routing domain and routes learned outside the OSPF routing domain. On IOS routers, you modify the default AD with this OSPF configuration command:

distance ospf {external | inter-area | intra-area} *ad*

The command allows you to set a different AD for each OSPF network type.

IOS routers allow selective AD modification for specific networks with this command:

distance *ad source-ip source-ip-wildcard [acl-number* | *acl-name]*

The *source-ip* option restricts the modification to routes in the OSPF link-state database (LSDB) learned from the advertising router of the LSA. The *source-ip-wildcard* address fields match the router ID (RID) for the advertising route. The optional *acl* is used to restrict to a specific network prefix.

Troubleshooting Routing Loops Caused by Redistribution (Cont.)

Example 17-2 demonstrates how to modify R1 and R2 so that OSPF external routes are set with an AD of 171, which is higher than the AD of external EIGRP routes (170).

Example 17-2 *OSPF Customized AD Configuration*

```
R1(config)# router ospf 1
R1(config-rtr)# distance ospf external 171
R1(config-rtr)#

R2(config)# router ospf 1
R2(config-rtr)# distance ospf external 171
R2(config-rtr)#
```

Troubleshooting Routing Loops Caused by Redistribution (Cont.)

BGP differentiates between routes learned from internal BGP (iBGP) peers, routes learned from external BGP (eBGP) peers, and routes learned locally. On IOS routers, you use the BGP configuration command **distance bgp** *external-ad internal-ad local-routes* to set the AD for each BGP network type and the address family, command **distance** *ad source-ip source-wildcard [acl-number* | *acl-name]* to modify AD for routes received from a specific neighbor.

For example, the BGP command **distance 44 55 66** sets the AD for eBGP routes to 44, the AD for iBGP routes to 55, and the AD for locally learned routes to 66.

Troubleshooting Routing Loops Caused by Redistribution (Cont.)

There is another way to solve the issue presented in Figure 17-5. You could attach a distribute list to the OSPF process on R1 and R2. When a distribute list is used with OSPF, it can control what routes are installed in the routing table from the OSPF database. Therefore, if you deny the 10.1.1.0/24 route in the OSPF database with an AD of 110 from being installed in the routing table with a distribute list on R1 and R2, the EIGRP route with an AD of 170 is installed in the routing table instead.

Example 17-3 shows how you can configure R1 and R2 to accomplish this.

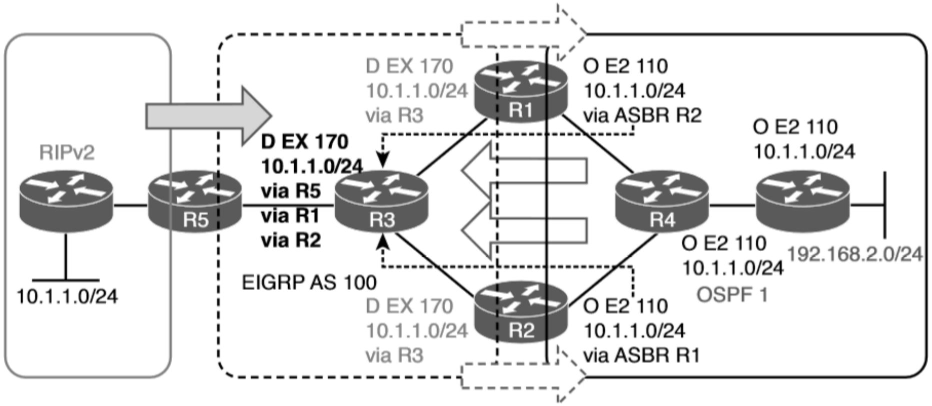

Figure 17-5 *Redistributing 10.1.1.0/24 Back into the EIGRP Autonomous System from OSPF*

Example 17-3 *Using a Distribute List to Control OSPF Routes Installed in the Routing Table*

```
R1(config)# ip prefix-list PREFER_EIGRP seq 10 deny 10.1.1.0/24
R1(config)# ip prefix-list PREFER_EIGRP seq 20 permit 0.0.0.0/0 le 32
R1(config)# router ospf 1
R1(config-rtr)# distribute-list prefix PREFER_EIGRP in
R1(config-rtr)#


R2(config)# ip prefix-list PREFER_EIGRP seq 10 deny 10.1.1.0/24
R2(config)# ip prefix-list PREFER_EIGRP seq 20 permit 0.0.0.0/0 le 32
R2(config)# router ospf 1
R2(config-rtr)# distribute-list prefix PREFER_EIGRP in
R2(config-rtr)#
```

Troubleshooting Routing Loops Caused by Redistribution (Cont.)

You do not want the routes that are redistributed from EIGRP into OSPF to be redistributed back into the EIGRP autonomous system or vice versa. Such redistribution can cause routing issues such as loops, which prevent packets from being correctly delivered to their destination (in addition to wasting CPU and memory resources on various devices in the network).

The most robust way to deal with this is to use route tags.

Figure 17-6 shows how R1 and R2 can add a tag when the route is redistributed. This is accomplished with route maps. In this example, when R1 redistributes the 10.1.1.0/24 route into the OSPF domain, it adds the tag 10. When R2 redistributes the 10.1.1.0/24 route into the OSPF domain, it adds the tag 20.

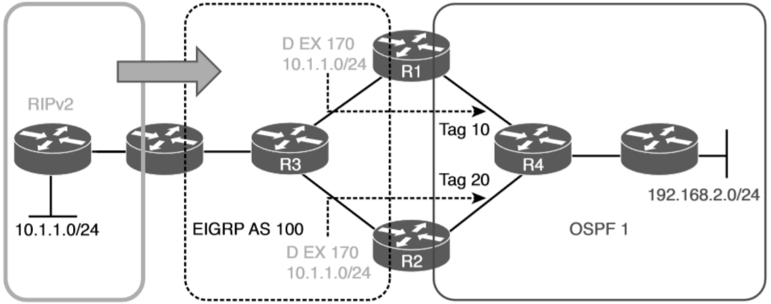

Figure 17-6 *Adding Tags to Routes During Redistribution*

Troubleshooting Routing Loops Caused by Redistribution (Cont.)

Example 17-4 shows the commands that you could use to tag the 10.1.1.0/24 routes as they are redistributed on R1 and R2.

- Define the routes you want to tag with an ACL or a prefix list.
- Create a route map that has a sequence that matches the ACL or prefix list created, and will set the desired tag when there is a match.
- Then attach the route map to the redistribution command.

Example 17-4 *Tagging Routes as They Are Being Redistributed*

```
R1#
ip prefix-list TAG_10.1.1.0/24 seq 5 permit 10.1.1.0/24
!
route-map REDIS_EIGRP_TO_OSPF permit 10
 match ip address prefix-list TAG_10.1.1.0/24
 set tag 10
route-map REDIS_EIGRP_TO_OSPF permit 20
!
router ospf 1
 redistribute eigrp 100 subnets route-map REDIS_EIGRP_TO_OSPF

R2#
ip prefix-list TAG_10.1.1.0/24 seq 5 permit 10.1.1.0/24
!
route-map REDIS_EIGRP_TO_OSPF permit 10
 match ip address prefix-list TAG_10.1.1.0/24
 set tag 20
route-map REDIS_EIGRP_TO_OSPF permit 20
!
router ospf 1
 redistribute eigrp 100 subnets route-map REDIS_EIGRP_TO_OSPF
```

Troubleshooting Routing Loops Caused by Redistribution (Cont.)

To prevent R1 and R2 from redistributing the OSPF-learned 10.1.1.0/24 routes with their tags back into EIGRP, you deny the routes based on their tags.

As shown in Figure 17-7, on R1 you deny the routes with the tag 20 from being redistributed into the EIGRP autonomous system, and on R2 you deny the routes with the tag 10 from being redistributed into the EIGRP autonomous system.

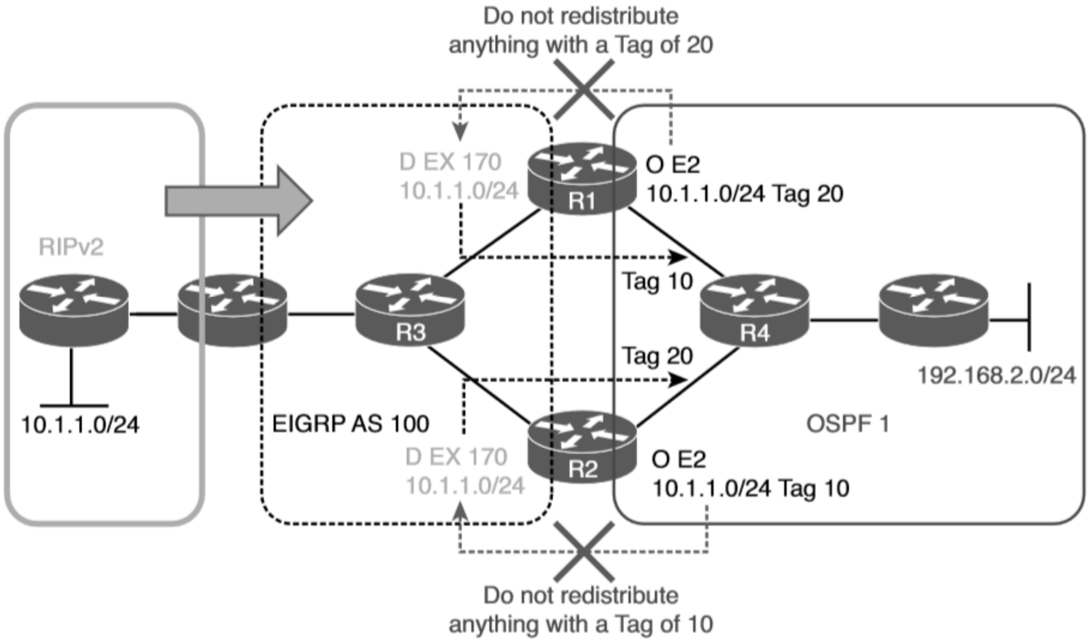

Figure 17-7 *Denying Routes with Certain Tags During Redistribution*

Troubleshooting Routing Loops Caused by Redistribution (Cont.)

Example 17-5 shows the commands that can be used to ensure that R1 and R2 do not redistribute the 10.1.1.0/24 networks back into the EIGRP autonomous system.

Example 17-5 *Using Route Tags to Prevent Routes from Being Reinjected*

```
R1#
route-map REDIS_OSPF_INTO_EIGRP deny 10
 match tag 20
route-map REDIS_OSPF_INTO_EIGRP permit 20
!
router eigrp 100
 redistribute ospf 1 metric 100000 100 255 1 1500 route-map REDIS_OSPF_INTO_EIGRP

R2#
route-map REDIS_OSPF_INTO_EIGRP deny 10
 match tag 10
route-map REDIS_OSPF_INTO_EIGRP permit 20
!
router eigrp 100
 redistribute ospf 1 metric 100000 100 255 1 1500 route-map REDIS_OSPF_INTO_EIGRP
```

Troubleshooting IPV4 and IPV6 Redistribution

- Route redistribution allows routes learned through one source (for example, statically configured routes, locally connected routes, or routes learned through a routing protocol) to be injected into a routing protocol.
- If two routing protocols are mutually redistributed, the routes learned through each routing protocol are injected into the other routing protocol.
- This section reviews route redistribution and explains how to troubleshoot redistribution issues.

Route Redistribution Review

A router that connects two or more routing domains and that will be the point of redistribution is known as a boundary router (see Figure 17-8). A boundary router can redistribute static routes, connected routes, and routes learned through one routing protocol into another routing protocol.

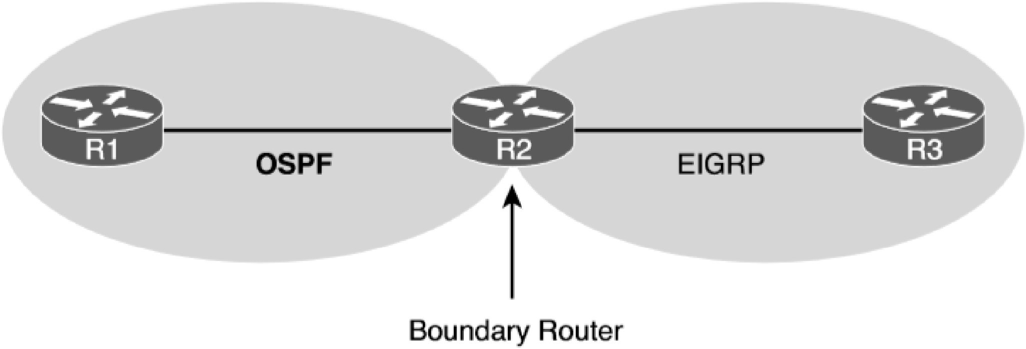

Figure 17-8 *Boundary Router*

Route Redistribution Review (Cont.)

Redistribution occurs from the routing table into a routing protocol's data structure (such as the EIGRP topology table or the OSPF LSDB), as shown in Figure 17-9. This is a key concept for troubleshooting purposes because if the route is not in the routing table, it cannot be redistributed.

Different routing protocols use different types of metrics, as illustrated in Figure 17-10. Therefore, when a route is redistributed into a routing protocol, a metric used by the destination routing protocol needs to be associated with the route being redistributed.

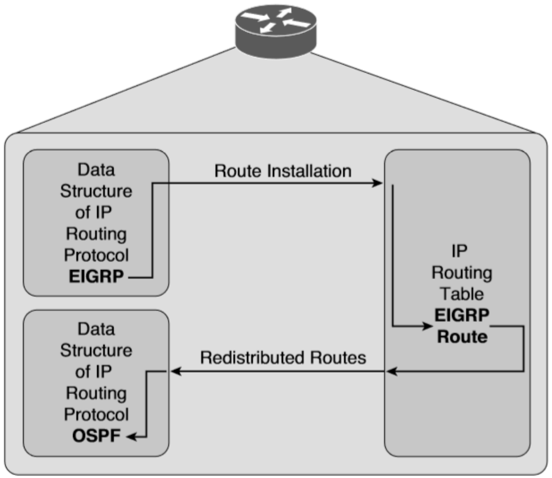

Figure 17-9 *Redistribution Occurs from the Routing Table into a Routing Protocol's Data Structure*

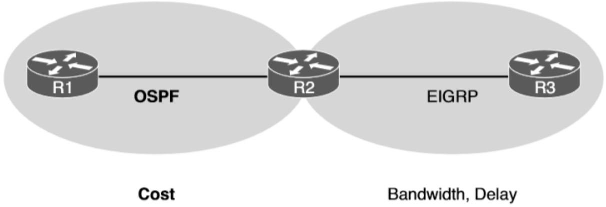

Figure 17-10 *Differing Metrics Between Routing Protocols*

Seed Metric

The metric assigned to a route being redistributed into another routing process is called a seed metric. The seed metric is needed to communicate relative levels of reachability between dissimilar routing protocols. A seed metric can be defined in one of three ways:

- Using the **default-metric** command
- Using the **metric** parameter with the **redistribute** command
- Applying a route map configuration to the **redistribute** command

If multiple seed metrics are defined with the commands, the order of preference is (1) metric defined in the route map that was applied to the **redistribute** command; (2) metric parameter defined with the **redistribute** command; (3) metric defined with the **default-metric** command.

If a seed metric is not specified, a default seed metric is used.

Redistribution Troubleshooting Targets

Two prerequisites must be met for the routes of one IP routing protocol to be redistributed into another IP routing protocol:

- The route needs to be installed in the IP routing table of the border router (the router performing redistribution) by the protocol being redistributed.
- The destination IP routing protocol needs a reachable metric to assign to the redistributed routes.

Based on these two prerequisites, Table 17-2 lists various redistribution troubleshooting targets and recommendations for dealing with them.

Table 17-2 Troubleshooting Targets for Route Redistribution

Troubleshooting Target	Troubleshooting Recommendation
Source routing protocol	Verify that a route to be redistributed from a routing protocol has been learned by that routing protocol. Issue appropriate **show** commands for the data structures of the source routing protocol to ensure that the source routing protocol has learned the route in question.
Route selection	Because a route must be in a router's IP routing table to be redistributed, ensure that the routes of the source routing protocol are indeed being injected into the router's IP routing table.
Redistribution configuration	If a route has been injected into a router's IP routing table from a source routing protocol but not redistributed into the destination routing protocol, check the redistribution configuration. This involves checking the metric applied to routes as they are redistributed into the destination routing protocol, checking for any route filtering that might be preventing redistribution, and checking the redistribution syntax to confirm that the correct routing process ID or autonomous system number is specified.
Destination routing protocol	If a route has been successfully redistributed into a destination routing protocol but the route has not been successfully learned by neighboring routers, you should investigate the destination routing protocol. You could use traditional methods of troubleshooting a destination routing protocol; however, keep in mind that the redistributed route might be marked as an external route. Therefore, check the characteristics of the destination routing protocol to determine whether it treats external routes differently from internal ones.

Troubleshooting Redistribution into EIGRP

When redistributing into EIGRP for IPv4, you can apply a metric with the **metric** keyword or a route map with the **route-map** keyword. If you are redistributing OSPF into EIGRP, as shown in Example 17-6, you also have the option to specify the **match** option, which allows you to match just **internal** routes, just **external** routes, just **nssa-external** routes, or a combination of them.

The most common issue you will run into when redistributing into EIGRP for IPv4 is related to the metric. Remember that the seed metric is, by default, set to infinity (unreachable). Therefore, if you fail to manually set the metric using any of the options listed earlier in the chapter, routes will not be advertised to the other routers in the EIGRP autonomous system.

Also, if the wrong route map is applied, or if there is an error in the route map, routes will not be redistributed properly.

Example 17-6 *EIGRP for IPv4 Redistribution Options*

```
R1(config)# router eigrp 1
R1(config-router)# redistribute ospf 1 ?
 match Redistribution of OSPF routes
 metric Metric for redistributed routes
 route-map Route map reference
 <cr>
```


Troubleshooting Redistribution into EIGRP (Cont.)

- With EIGRP for IPv6, you have the same **match**, **metric**, and **route-map** keywords, as well as the **include-connected** keyword.
- By default, with EIGRP for IPv4, the networks associated with the local interfaces participating in the redistributed routing process are also redistributed.
- With EIGRP for IPv6, the networks associated with the local interfaces participating in the redistributed routing process are not redistributed. Therefore, if you want to include the networks associated with the local interfaces participating in the routing process that is being redistributed, you need to use the **include-connected** keyword, as shown in Example 17-7.

Example 17-7 *EIGRP for IPv6 Redistribution Options*

```
R1(config)# ipv6 router eigrp 1
R1(config-rtr)# redistribute ospf 1 ?
 include-connected Include connected
 match Redistribution of OSPF routes
 metric Metric for redistributed routes
 route-map Route map reference
 <cr>
```

Troubleshooting Redistribution into EIGRP (Cont.)

- On the boundary router, you can verify which protocols are being redistributed into EIGRP for IPv4 with the **show ip protocols** command. Example 17-8 shows that OSPF routes are being redistributed into EIGRP for IPv4.
- When reviewing the EIGRP for IPv4 topology table with the **show ip eigrp topology** command, you can identify the routes that have been injected into the EIGRP process via redistribution because the output states via Redistributed, as shown in Example 17-9.

Example 17-8 *Verifying Protocols That Are Being Redistributed into EIGRP for IPv4*

```
R2# show ip protocols
*** IP Routing is NSF aware ***

Routing Protocol is "eigrp 100"
 Outgoing update filter list for all interfaces is not set
 Incoming update filter list for all interfaces is not set
 Default networks flagged in outgoing updates
 Default networks accepted from incoming updates
 Redistributing: ospf
 EIGRP-IPv4 Protocol for AS(100)
 Metric weight K1=1, K2=0, K3=1, K4=0, K5=0
...output omitted...
```

Example 17-9 *Verifying Routes Redistributed into EIGRP for IPv4 (Topology Table)*

```
R2# show ip eigrp topology
EIGRP-IPv4 Topology Table for AS(100)/ID(203.0.113.1)
     Codes: P - Passive, A - Active, U - Update, Q - Query, R - Reply,
            r - reply Status, s - sia Status

P 10.1.12.0/24, 1 successors, FD is 2560000256
     via Redistributed (2560000256/0)
P 10.1.14.0/24, 1 successors, FD is 2560000256
     via Redistributed (2560000256/0)
P 10.1.3.0/24, 1 successors, FD is 3072
     via 10.1.23.3 (3072/2816), GigabitEthernet1/0
P 10.1.23.0/24, 1 successors, FD is 2816
     via Connected, GigabitEthernet1/0
P 10.1.1.0/24, 1 successors, FD is 2560000256
     via Redistributed (2560000256/0)
```

Troubleshooting Redistribution into EIGRP (Cont.)

- When examining a redistributed route in the routing table on the boundary router, as shown in Example 17-10, with the **show ip route** *ip-address* command, the output indicates how the route is known, how it is being redistributed, and the EIGRP metric values that are being used at the redistribution point.

- When examining the routing tables on other routers (not the boundary router) in the EIGRP for IPv4 autonomous system, the redistributed routes have an AD of 170 by default and the code D EX, as shown in Example 17-11.

Example 17-10 *Verifying Routes Redistributed into EIGRP for IPv4 (Routing Table)*

```
R2# show ip route 10.1.1.0
Routing entry for 10.1.1.0/24
 Known via "ospf", distance 110, metric 1
 Redistributing via eigrp 100, ospf
 Advertised by eigrp 100 metric 1 1 1 1 1
 Last update from 10.1.12.1 on GigabitEthernet0/0, 00:00:19 ago
 Routing Descriptor Blocks:
 * 10.1.12.1, from 10.1.12.1, 00:00:19 ago, via GigabitEthernet0/0
 Route metric is 1, traffic share count is 1
```

Example 17-11 *Examining EIGRP for IPv4 Redistributed Routes in a Routing Table*

```
R3# show ip route
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
      D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
      N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
      E1 - OSPF external type 1, E2 - OSPF external type 2
      i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
      ia - IS-IS inter area, * - candidate default, U - per-user static route
      o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP
      + - replicated route, % - next hop override

Gateway of last resort is not set

 10.0.0.0/8 is variably subnetted, 7 subnets, 2 masks
D EX 10.1.1.0/24
        [170/2560000512] via 10.1.23.2, 00:04:38, GigabitEthernet1/0
C 10.1.3.0/24 is directly connected, GigabitEthernet0/0
L 10.1.3.3/32 is directly connected, GigabitEthernet0/0
D EX 10.1.12.0/24
        [170/2560000512] via 10.1.23.2, 00:04:38, GigabitEthernet1/0
D EX 10.1.14.0/24
        [170/2560000512] via 10.1.23.2, 00:04:38, GigabitEthernet1/0
C 10.1.23.0/24 is directly connected, GigabitEthernet1/0
L 10.1.23.3/32 is directly connected, GigabitEthernet1/0
```

Troubleshooting Redistribution into EIGRP (Cont.)

- For EIGRP for IPv6, **the show ipv6 protocols** output is more detailed for redistribution, as shown in Example 17-12. Notice that it states the protocol, the seed metric, and whether connected networks are included.

- The output of **show ipv6 eigrp topology** on the boundary router also indicates which routes are redistributed, as shown in Example 17-13.

- When examining the routing tables on other routers besides the boundary router in the EIGRP for IPv6 autonomous system, the redistributed routes have an administrative distance of 170 by default and the code EX.

Example 17-12 *Verifying EIGRP for IPv6 Redistribution with show ipv6 protocols*

```
R2# show ipv6 protocols
...output omitted...
IPv6 Routing Protocol is "eigrp 100"
EIGRP-IPv6 Protocol for AS(100)
 Metric weight K1=1, K2=0, K3=1, K4=0, K5=0
 NSF-aware route hold timer is 240
 Router-ID: 203.0.113.1
 Topology : 0 (base)
 Active Timer: 3 min
 Distance: internal 90 external 170
 Maximum path: 16
Maximum hopcount 100
 Maximum metric variance 1

 Interfaces:
   GigabitEthernet1/0
 Redistribution:
   Redistributing protocol OSPF 1 with metric 1 1 1 1 1 include-connected
```

Example 17-13 *Verifying EIGRP for IPv6 Redistribution with show ipv6 eigrp topology*

```
R2# show ipv6 eigrp topology
EIGRP-IPv6 Topology Table for AS(100)/ID(203.0.113.1)
Codes: P - Passive, A - Active, U - Update, Q - Query, R - Reply,
 r - reply Status, s - sia Status

P 2001:DB8:0:1::/64, 1 successors, FD is 2560000256
        via Redistributed (2560000256/0)
P 2001:DB8:0:3::/64, 1 successors, FD is 3072
        via FE80::C804:10FF:FE2C:1C (3072/2816), GigabitEthernet1/0
P 2001:DB8:0:12::/64, 1 successors, FD is 2560000256
        via Redistributed (2560000256/0)
P 2001:DB8:0:23::/64, 1 successors, FD is 2816
        via Connected, GigabitEthernet1/0
```

Troubleshooting Redistribution into OSPF

- When redistributing into OSPF, you have more options than you have with other routing protocols.
- The **metric** option allows you to provide a seed metric at the redistribution point. The default seed metric is 20 with OSPF; therefore, providing a metric is not mandatory.
- With the **subnets** keyword, all classless and classful networks are redistributed.
- The **tag** keyword can be used to add a numeric ID (tag) to the route so the route can be referenced by the tag at a later point for filtering or manipulation purposes.

Troubleshooting Redistribution into OSPF

- The **metric** option allows you to provide a seed metric at the redistribution point.
- The default seed metric is 20 with OSPF; therefore, providing a metric is not mandatory. If you forget to provide a metric, redistributed routes are still advertised to other routers in the OSPF domain.
- The **metric-type** option is used to define the type of OSPF external route the redistributed route will be. By default, it will be Type 2, which is represented as E2 in the routing table. With E2, each router preserves the seed metric for the external routes.
- Type 1, which is represented as E1 in the routing table, allows each router to add to the seed metric all the other link costs to reach the redistribution point in the domain. Therefore, each router has a metric that is a combination of the seed metric and the total cost to reach the redistribution router.
- With the **nssa-only** option, you can limit redistributed routes to the not-so-stubby area (NSSA) only, and with the **route-map** option, you can reference a route map that provides more granular control over the routes that are being redistributed.
- The **subnets** keyword is an extremely important option. Without the **subnets** keyword, only classful networks are redistributed (for example, a Class A address with a /8 mask, a Class B address with a /16 mask, and a Class C address with a /24 mask). With the **subnets** keyword, all classless and classful networks are redistributed.

Troubleshooting Redistribution into OSPF (Cont.)

- Example 17-16 displays the options available when redistributing into OSPFv3. The **include-connected** keyword has been added. By default, with OSPFv2, the networks associated with the local interfaces that are participating in the routing process that is being redistributed are also redistributed. However, with OSPFv3, they are not. To include the networks associated with the interfaces participating in the routing protocol that is being redistributed on the ASBR, use the **include-connected** keyword.
- The **subnets** keyword is not an option with OSPFv3 because the concepts classful and classless do not exist with IPv6.
- The **show ip protocols** command enables you to verify which routing protocols are being redistributed into the OSPFv2 process. In Example 17-17, you can see that EIGRP 100 routes, including subnets, are being redistributed into the OSPFv2 process.

Example 17-16 *OSPFv3 Redistribution Options*

```
R1(config)# ipv6 router ospf 1
R1(config-rtr)# redistribute eigrp 100 ?
 include-connected  Include connected
 metric             Metric for redistributed routes
 metric-type        OSPF/IS-IS exterior metric type for redistributed routes
 nssa-only          Limit redistributed routes to NSSA areas
 route-map          Route map reference
 tag                Set tag for routes redistributed into OSPF
 <cr>
```

Example 17-17 *Verifying Protocols Being Redistributed into OSPFv2*

```
R2# show ip protocols
...output omitted...
Routing Protocol is "ospf 1"
  Outgoing update filter list for all interfaces is not set
  Incoming update filter list for all interfaces is not set
  Router ID 203.0.113.1
  It is an autonomous system boundary router
 Redistributing External Routes from,
    eigrp 100, includes subnets in redistribution
  Number of areas in this router is 1. 1 normal 0 stub 0 nssa
  Maximum path: 4
  Routing for Networks:
    10.1.12.2 0.0.0.0 area 0
  Routing Information Sources:
    Gateway         Distance      Last Update
    10.1.14.1       110           00:19:48
  Distance: (default is 110)
```

Troubleshooting Redistribution into OSPF (Cont.)

- Routes that are redistributed into an OSPFv2 normal area are advertised within a Type 5 LSA. Routes that are redistributed into an OSPFv2 NSSA or totally NSSA area are advertised with a Type 7 LSA and then converted to a Type 5 LSA at an area border router (ABR).
- You can view the redistributed routes that are injected into the OSPFv2 LSDB with the **show ip ospf database** command, as shown in Example 17-18. In this example, the 10.1.3.0 and 10.1.23.0 networks have been redistributed into the OSPFv2 routing process.

Example 17-18 *Verifying Redistributed Routes in the OSPFv2 LSDB*

```
R2# show ip ospf database

            OSPF Router with ID (203.0.113.1) (Process ID 1)

                Router Link States (Area 0)

Link ID         ADV Router      Age         Seq#       Checksum    Link count
10.1.14.1       10.1.14.1       738         0x80000003 0x009AEA    3
203.0.113.1     203.0.113.1     596         0x80000003 0x005829    1

                Net Link States (Area 0)

Link ID         ADV Router      Age         Seq#       Checksum
10.1.12.1       10.1.14.1       738         0x80000002 0x001F8F

                Type-5 AS External Link States

Link ID         ADV Router      Age         Seq#       Checksum   Tag
10.1.3.0        203.0.113.1     596         0x80000002 0x00EB67   0
10.1.23.0       203.0.113.1     596         0x80000002 0x000F30   0
```

Troubleshooting Redistribution into OSPF (Cont.)

- When examining a redistributed route in the routing table on the boundary router (autonomous system boundary router [ASBR]), as shown in Example 17-19, with the **show ip route** *ip_address* command, the output indicates how the route is known, how it is being redistributed, and how it is being advertised. In this case, the route is known through EIGRP 100 and is being redistributed into the OSPF 1 process with the **subnets** keyword.

- For OSPFv3, the **show ipv6 protocols** output is shown in Example 17-21. Notice that it states the protocol, the seed metric, and whether connected networks are included.

- The output of **show ipv6 ospf database** on the ASBR identifies the external Type 5 routes just as OSPFv2 does.

Example 17-19 *Verifying Redistributed Routes in the ASBR's Routing Table*

```
R2# show ip route 10.1.3.0
Routing entry for 10.1.3.0/24
  Known via "eigrp 100", distance 90, metric 3072, type internal
  Redistributing via eigrp 100, ospf 1
  Advertised by ospf 1 subnets
  Last update from 10.1.23.3 on GigabitEthernet1/0, 00:50:19 ago
  Routing Descriptor Blocks:
  * 10.1.23.3, from 10.1.23.3, 00:50:19 ago, via GigabitEthernet1/0
      Route metric is 3072, traffic share count is 1
      Total delay is 20 microseconds, minimum bandwidth is 1000000 Kbit
      Reliability 255/255, minimum MTU 1500 bytes
      Loading 1/255, Hops 1
```

Example 17-21 *Verifying OSPFv3 Redistribution with show ipv6 protocols*

```
R2# show ipv6 protocols
...output omitted...
IPv6 Routing Protocol is "ospf 1"
  Router ID 2.2.2.2
  Autonomous system boundary router
  Number of areas: 1 normal, 0 stub, 0 nssa
  Interfaces (Area 0):
    GigabitEthernet0/0
  Redistribution:
    Redistributing protocol eigrp 100 with metric 10 include-connected
```

Troubleshooting Redistribution into OSPF (Cont.)

- When examining the routing table on other routers (not the ASBR) in the OSPFv3 domain, by default the redistributed routes have an administrative distance of 110 and the code OE2, as shown in Example 17-23. If the metric type is changed to Type 1, the code is OE1. In an NSSA or a totally NSSA, the redistributed routes are listed as ON1 or ON2.

- If you are redistributing from BGP into OSPF, EIGRP, or RIP, only eBGP routes are redistributed by default. If you want iBGP routes to be redistributed, you must issue the **bgp redistribute-internal** command in router BGP configuration mode.

Example 17-23 *Verifying OSPFv3 Redistributed Routes*

```
R1# show ipv6 route
IPv6 Routing Table - default - 9 entries
Codes: C - Connected, L - Local, S - Static, U - Per-user Static route
 B - BGP, R - RIP, H - NHRP, I1 - ISIS L1
 I2 - ISIS L2, IA - ISIS interarea, IS - ISIS summary, D - EIGRP
 EX - EIGRP external, ND - ND Default, NDp - ND Prefix, DCE - Destination
 NDr - Redirect, O - OSPF Intra, OI - OSPF Inter, OE1 - OSPF ext 1
 OE2 - OSPF ext 2, ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2, l - LISP
C    2001:DB8:0:1::/64 [0/0]
      via GigabitEthernet0/0, directly connected
L    2001:DB8:0:1::1/128 [0/0]
      via GigabitEthernet0/0, receive
OE2  2001:DB8:0:3::/64 [110/10]
      via FE80::C802:AFF:FE88:8, GigabitEthernet1/0
C    2001:DB8:0:12::/64 [0/0]
      via GigabitEthernet1/0, directly connected
L    2001:DB8:0:12::1/128 [0/0]
      via GigabitEthernet1/0, receive
C    2001:DB8:0:14::/64 [0/0]
      via FastEthernet3/0, directly connected
L    2001:DB8:0:14::1/128 [0/0]
      via FastEthernet3/0, receive
OE2  2001:DB8:0:23::/64 [110/10]
      via FE80::C802:AFF:FE88:8, GigabitEthernet1/0
L    FF00::/8 [0/0]
      via Null0, receive
```

Troubleshooting Redistribution into BGP

- When redistributing EIGRP into BGP for IPv4, you have the same options as when redistributing into EIGRP. You can apply a metric with the **metric** keyword or a route map with the **route-map** keyword.
- If you are redistributing OSPF into BGP, as shown in Example 17-24, you also have the option to specify the **match** option, which allows you to match just **internal** routes, just **external** routes, just **nssa-external** routes, or a combination of them. With BGP, only internal OSPF routes are redistributed by default.
- The **metric** keyword is not required because BGP uses the IGP metric by default. If the wrong route map is applied, or if there is an error in the route map, routes are not redistributed properly.

Example 17-24 *BGP for IPv4 Redistribution Options*

```
R1(config)# router bgp 65001
R1(config-router)# address-family ipv4 unicast
R1(config-router-af)# redistribute ospf 1 ?
 match       Redistribution of OSPF routes
 metric      Metric for redistributed routes
 route-map   Route map reference
 vrf         VPN Routing/Forwarding Instance
 <cr>
```

Troubleshooting Redistribution into BGP (Cont.)

- With BGP for IPv6, you have the same **match**, **metric**, and **route-map** keywords, as well as the **include-connected** keyword.
- By default, with BGP for IPv4, the networks of the local interfaces participating in the routing protocol that is being redistributed on the border router are also redistributed. However, with BGP for IPv6, they are not.
- If you want to redistribute the networks associated with the local interfaces participating in the routing process being redistributed into BGP for IPv6, you need to use the **include-connected** keyword, as shown in Example 17-25.

Example 17-25 *BGP for IPv6 Redistribution Options*

```
R1(config)# router bgp 65001
R1(config-router) # address-family ipv6 unicast
R1(config-router-af) # redistribute ospf 1 ?
 include-connected   Include connected
 match               Redistribution of OSPF routes
 metric              Metric for redistributed routes
 route-map           Route map reference
 <cr>
```

- Using the commands **show ip protocols** and **show ipv6 protocols**, you can verify which protocols are being redistributed into the BGP routing process.

Troubleshooting Redistribution into BGP (Cont.)

- In the BGP table, redistributed routes appear with a question mark (?) under the Path column, as shown in Example 17-27.

Example 17-27 *Verifying Redistributed Routes in the BGP Table*

```
R2# show bgp all
For address family: IPv4 Unicast

BGP table version is 4, local router ID is 203.0.113.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
              x best-external, a additional-path, c RIB-compressed,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

 Network          Next Hop      Metric LocPrf Weight Path
 *> 10.1.1.0/24      10.1.12.1          2         32768 ?
 *> 10.1.12.0/24     0.0.0.0            0         32768 ?
 *> 10.1.14.0/24     10.1.12.1          2         32768 ?

For address family: IPv6 Unicast

BGP table version is 4, local router ID is 203.0.113.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
              x best-external, a additional-path, c RIB-compressed,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

 Network             Next Hop    Metric     LocPrf    Weight Path
 *> 2001:DB8:0:1::/64    ::               2              32768 ?
 *> 2001:DB8:0:12::/64   ::               0              32768 ?
 *> 2001:DB8:0:14::/64   ::               2              32768 ?

…output omitted…
```

Troubleshooting Redistribution with Route Maps

When applying a route map with the redistribution command, you have a few extra items to verify during the troubleshooting process:

- Is the correct route map applied?
- Is **permit** or **deny** specified for the sequence, and is it correct? A permit sequence indicates that what is matched is redistributed. A deny sequence indicates that what is matched is not redistributed.
- If there is an access list or a prefix list being used in the **match** statement, you need to verify that they are correct by using the **show** {**ip** | **ipv6**} **access-list** command or the **show {ip | ipv6} prefix-list** command.
- If there are **set** statements, you need to verify that the correct values have been specified to accomplish the desired goal.
- If a route does not match any of the **match** statements in any of the sequences, it falls into the implicit deny sequence at the end of the route map and is not redistributed.
- If a route map is attached to the **redistribution** command but that route map does not exist, none of the routes are redistributed.

Redistribution Trouble Tickets

- This section presents various trouble tickets related to the topics discussed in this chapter. The purpose of these trouble tickets is to show a process that you can follow when troubleshooting in the real world or in an exam environment.

Trouble Tickets

This section presents various trouble tickets related to the topics discussed in this chapter. The purpose of these trouble tickets is to show a process that you can follow when troubleshooting in the real world or in an exam environment.

All trouble tickets in this section are based on the topology shown in Figure 17-11.

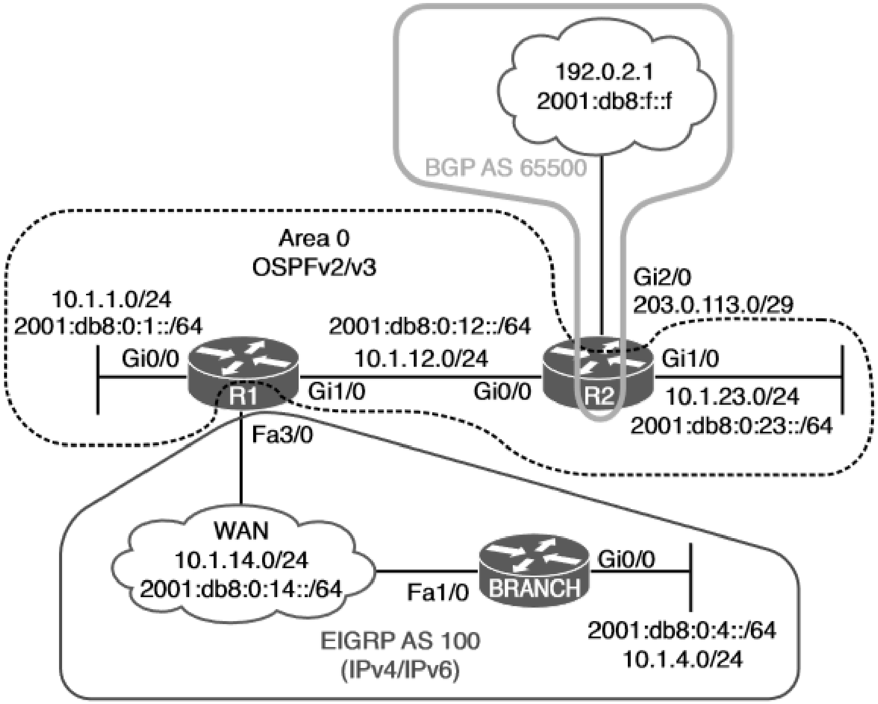

Figure 17-11 *Redistribution Trouble Tickets Topology*

Trouble Ticket 17-1

Problem: Users in the IPv4 Branch site indicate that they are not able to access any resources outside the Branch office.

On Branch, you first check the routing table to see which routes Branch knows; you do this by using the **show ip route** command, as shown in Example 17-28. The output indicates that Branch only knows about connected and local routes.

Refer to your text for next steps and examples to troubleshoot and resolve this trouble ticket.

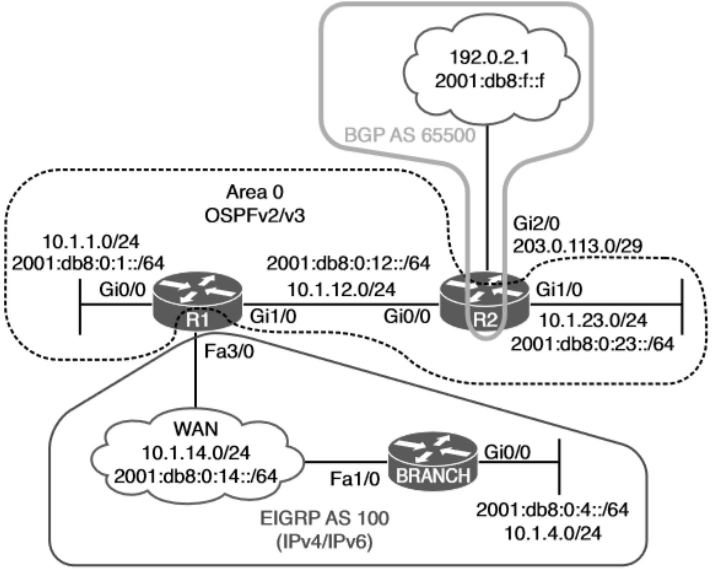

Figure 17-11 *Redistribution Trouble Tickets Topology*

Example 17-28 *Verifying the Routing Table on Branch*

```
Branch# show ip route
...output omitted...
      10.0.0.0/8 is variably subnetted, 4 subnets, 2 masks
C        10.1.4.0/24 is directly connected, GigabitEthernet0/0
L        10.1.4.4/32 is directly connected, GigabitEthernet0/0
C        10.1.14.0/24 is directly connected, FastEthernet1/0
L        10.1.14.4/32 is directly connected, FastEthernet1/0
```


Trouble Ticket 17-2

Problem: Users in the 10.1.23.0/24 network indicate that they are not able to access resources in the 10.1.4.0/24 network.

You begin troubleshooting by verifying the problem on R2. You issue a ping to 10.1.4.4 from 10.1.23.2, but it fails, as shown in Example 17-38. Because R2 is not able to ping the destination network, you confirm that the clients in 10.1.23.0/24 are not able to connect with resources in 10.1.4.0/24.

Refer to your text for next steps and examples to troubleshoot and resolve this trouble ticket.

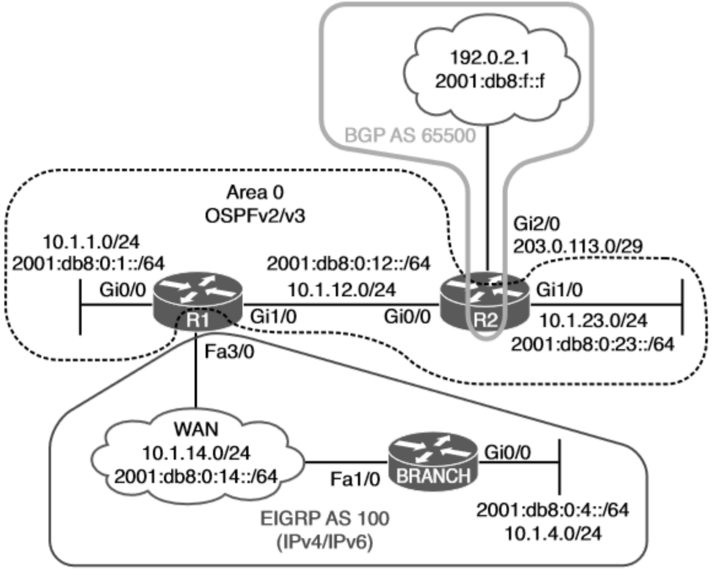

Figure 17-11 *Redistribution Trouble Tickets Topology*

Example 17-38 *Verifying the Problem from R2*

```
R2# ping 10.1.4.4 source 10.1.23.2
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.4.4, timeout is 2 seconds:
Packet sent with a source address of 10.1.23.2
.....
Success rate is 0 percent (0/5)
```

Trouble Ticket 17-3

Problem: IPv6 users in the 2001:db8:0:4::/64 network report that they are not able to access resources in the 2001:db8:0:1::/64 network.

You begin troubleshooting by confirming the problem on Branch. As shown in Example 17-50, the ping from 2001:db8:0:4::4 to 2001:db8:0:1::1 fails.

Refer to your text for next steps and examples to troubleshoot and resolve this trouble ticket.

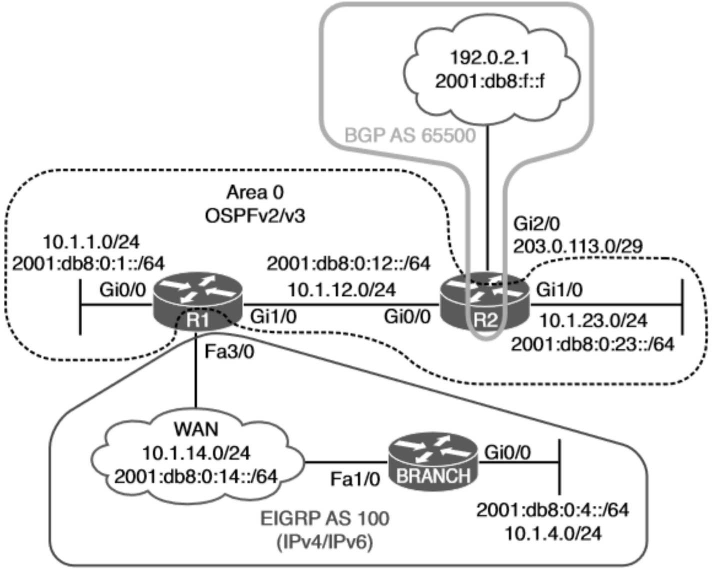

Figure 17-11 *Redistribution Trouble Tickets Topology*

Example 17-50 *Confirming the Problem with a Ping*

```
Branch# ping 2001:db8:0:1::1 source 2001:db8:0:4::4
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2001:DB8:0:1::1, timeout is 2 seconds:
Packet sent with a source address of 2001:DB8:0:4::4
.....
Success rate is 0 percent (0/5)
```

Trouble Ticket 17-4

Problem: A junior administrator has approached you, asking for help. He claims that users in BGP autonomous system 65500 are unable to access IPv4 resources in the IPv4 EIGRP autonomous system 100. However, they can access resources in the OSPFv2 domain. Because access to routers in BGP autonomous system 65500 is limited to only R2, the junior administrator has asked you for help.

You start by reviewing Figure 17-11 to confirm which local router is running BGP.

Refer to your text for next steps and examples to troubleshoot and resolve this trouble ticket.

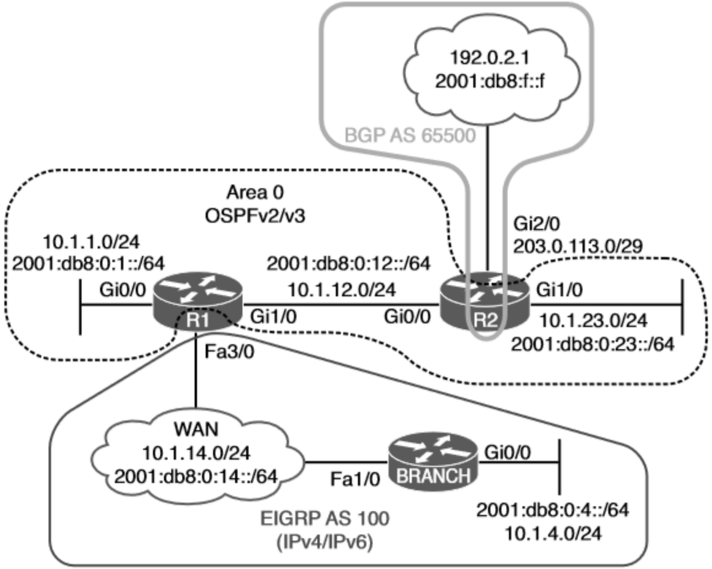

Figure 17-11 *Redistribution Trouble Tickets Topology*

Prepare for the Exam

Prepare for the Exam

Key Topics for Chapter 17

Description	
Solving and preventing suboptimal routing caused by redistribution	Using route tags to prevent routes from being reinjected
Troubleshooting suboptimal routing issues caused by redistribution	The redistribution process
Routing loops with redistribution at multiple points	Three methods for configuring a seed metric
Routing table instability with redistribution at multiple points	Prerequisites for redistributing a route
How traffic might get blackholed with redistribution at multiple points	Troubleshooting targets for route redistribution
Modifying the EIGRP administrative distance	Troubleshooting redistribution into EIGRP
Modifying the OSPF administrative distance	Troubleshooting redistribution into OSPF
Modifying the BGP administrative distance	Troubleshooting redistribution into BGP
Using a distribute list to control OSPF routes that are installed in the routing table	Troubleshooting redistribution that uses route maps
Tagging routes as they are being redistributed	

Prepare for the Exam
Key Terms for Chapter 17

Term	Term
redistribution	ASBR
boundary router	routing loop
metric	single-point redistribution
seed metric	multipoint redistribution
subnets keyword	route tag
Type 5 LSA	administrative distance

Prepare for the Exam

Command Reference for Chapter 17

Task	Command Syntax
Display the IPv4 sources of routing information that are being redistributed into the IPv4 routing protocols enabled on the device	**show ip protocols**
Display the IPv6 sources of routing information that are being redistributed into the IPv6 routing protocols enabled on the device	**show ipv6 protocols**
Show which IPv4 routes have been redistributed into the IPv4 EIGRP process on the boundary router	**show ip eigrp topology**
Show which IPv6 routes have been redistributed into the IPv6 EIGRP process on the boundary router	**show ipv6 eigrp topology**
Show which IPv4 routes have been redistributed into the OSPFv2 process; they are represented as Type 5 or Type 7 LSAs	**show ip ospf database**

Prepare for the Exam

Command Reference for Chapter 17 (Cont.)

Task	Command Syntax
Show which IPv6 routes have been redistributed into the OSPFv3 process; they are represented as Type 5 or Type 7 LSAs	**show ipv6 ospf database**
Display the IPv4 and IPv6 BGP learned routes; routes originally learned through redistribution have a question mark (?) in the Path column	**show bgp all**
Display a router's BGP router ID, autonomous system number, information about the BGP's memory usage, and summary information about IPv4 unicast BGP neighbors	**show bgp ipv4 unicast summary**
Display detailed information about all the IPv4 BGP neighbors of a router	**show bgp ipv4 unicast neighbors**