


Chapter 4 Lab 4-2, Redistribution Between EIGRP and OSPF

Topology



Objectives

- Review EIGRP and OSPF configuration.
- Redistribute into EIGRP.
- Redistribute into OSPF.
- Summarize routes in EIGRP.
- Filter routes using route maps.
- Modify EIGRP distances.
- Modify OSPF distances.
- Create passive interfaces in EIGRP.
- Summarize in OSPF at an ABR and an ASBR.

Background

R1 is running EIGRP, and R3 is running multi-area OSPF. In this lab, you configure redistribution on R2 to enable these two routing protocols to interact, allowing full connectivity between all networks. In Appendix A of this lab, you explore black hole operation.

Note: This lab uses Cisco 1841 routers with Cisco IOS Release 12.4(24)T1 and the Advanced IP Services image c1841-advipservicesk9-mz.124-24.T1.bin. You can use other routers (such as 2801 or 2811) and Cisco IOS Software versions if they have comparable capabilities and features. Depending on the router model and Cisco IOS Software version, the commands available and output produced might vary from what is shown in this lab.

Required Resources

- 3 routers (Cisco 1841 with Cisco IOS Release 12.4(24)T1 Advanced IP Services or comparable)
- Serial and console cables

Step 1: Configure loopbacks and additional addressing.

- a. Start with the final configurations of Lab 4.1, "Redistribution Between RIP and OSPF." On R1 and R2, remove the RIPv2 configuration and the static route with the following commands.

```
R1(config)# no router rip
R1(config)# no ip route 192.168.48.0 255.255.252.0 null0
R1(config)# no ip prefix-list RIP-OUT
```

```
R2(config)# no router rip
R2(config)# router ospf 1
R2(config-router)# no default-information originate
R2(config-router)# no redistribute rip
R2(config-router)# no default-metric 10000
```

- b. Configure the additional loopback interfaces on R2 and R3, as shown in the diagram.

```
R2(config)# interface loopback 100
R2(config-if)# ip address 172.16.100.1 255.255.255.0
```

```
R3(config)# interface loopback 8
R3(config-if)# ip address 192.168.8.1 255.255.255.0
R3(config-if)# interface loopback 9
R3(config-if)# ip address 192.168.9.1 255.255.255.0
R3(config-if)# interface loopback 10
R3(config-if)# ip address 192.168.10.1 255.255.255.0
R3(config-if)# interface loopback 11
R3(config-if)# ip address 192.168.11.1 255.255.255.0
```

Step 2: Configure EIGRP.

- a. Configure R1 and R2 to run EIGRP in autonomous system 1. On R1, add in all connected interfaces either with classful **network** commands or with wildcard masks. Use a classful **network** statement on R2 and disable automatic summarization.

```
R1(config)# router eigrp 1
R1(config-router)# no auto-summary
R1(config-router)# network 172.16.0.0
R1(config-router)# network 192.168.48.0
R1(config-router)# network 192.168.49.0
R1(config-router)# network 192.168.50.0
```

CCNPv6 ROUTE

```
R1(config-router)# network 192.168.51.0
R1(config-router)# network 192.168.70.0
```

or

```
R1(config)# router eigrp 1
R1(config-router)# no auto-summary
R1(config-router)# network 172.16.0.0
R1(config-router)# network 192.168.0.0 0.0.255.255
```

```
R2(config)# router eigrp 1
R2(config-router)# no auto-summary
R2(config-router)# network 172.16.0.0
```

- b. Verify the configuration with the **show ip eigrp neighbors** and **show ip route eigrp** commands on both routers.

```
R1# show ip eigrp neighbors
```

```
IP-EIGRP neighbors for process 1
```

H	Address	Interface	Hold (sec)	Uptime	SRTT (ms)	RTO	Q Cnt	Seq Num
0	172.16.12.2	Se0/0/0	11	00:00:30	36	216	0 3	

```
R1# show ip route eigrp
```

```
172.16.0.0/24 is subnetted, 5 subnets
D 172.16.23.0 [90/41024000] via 172.16.12.2, 00:01:38, Serial0/0/0
D 172.16.2.0 [90/40640000] via 172.16.12.2, 00:01:16, Serial0/0/0
D 172.16.100.0 [90/40640000] via 172.16.12.2, 00:02:13, Serial0/0/0
```

```
R2# show ip eigrp neighbors
```

```
IP-EIGRP neighbors for process 1
```

H	Address	Interface	Hold (sec)	Uptime	SRTT (ms)	RTO	Q Cnt	Seq Num
0	172.16.12.1	Se0/0/0	11	00:01:53	1604	5000	0 2	

```
R2# show ip route eigrp
```

```
172.16.0.0/24 is subnetted, 6 subnets
D 172.16.1.0 [90/40640000] via 172.16.12.1, 00:01:08, Serial0/0/0
D 192.168.70.0/24 [90/40640000] via 172.16.12.1, 00:01:08, Serial0/0/0
D 192.168.51.0/24 [90/40640000] via 172.16.12.1, 00:01:08, Serial0/0/0
D 192.168.50.0/24 [90/40640000] via 172.16.12.1, 00:01:08, Serial0/0/0
D 192.168.49.0/24 [90/40640000] via 172.16.12.1, 00:01:08, Serial0/0/0
D 192.168.48.0/24 [90/40640000] via 172.16.12.1, 00:01:08, Serial0/0/0
```

Step 3: Create passive interfaces in EIGRP.

- a. Issue the **show ip eigrp interfaces** command on R2.

```
R2# show ip eigrp interfaces
```

```
IP-EIGRP interfaces for process 1
```

		Xmit Queue	Mean	Pacing Time	Multicast	
Pending						
Interface	Peers	Un/Reliable	SRTT	Un/Reliable	Flow Timer	Routes
Se0/0/0	1	0/0	32	10/380	496	0
Se0/0/1	0	0/0	0	0/1	0	0
Lo0	0	0/0	0	0/1	0	0
Lo100	0	0/0	0	0/1	0	0

Because you used the classful **network** command, both serial interfaces are involved with EIGRP.

- b. To stop EIGRP from sending or processing received EIGRP packets on the serial interface going to R3, use the **passive-interface** *interface_type interface_number* command.

```
R2(config)# router eigrp 1
R2(config-router)# passive-interface serial 0/0/1
```

- c. Verify the change with the **show ip eigrp interfaces** and **show ip protocols** commands.

```
R2# show ip eigrp interfaces
IP-EIGRP interfaces for process 1
```

		Xmit Queue	Mean	Pacing Time	Multicast	
Pending						
Interface	Peers	Un/Reliable	SRTT	Un/Reliable	Flow Timer	Routes
Se0/0/0	1	0/0	32	10/380	496	0
Lo0	0	0/0	0	0/1	0	0
Lo100	0	0/0	0	0/1	0	0

```
R2# show ip protocols
Routing Protocol is "ospf 1"
<output omitted>
```

```
Routing Protocol is "eigrp 1"
  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
  EIGRP metric weight K1=1, K2=0, K3=1, K4=0, K5=0
  EIGRP maximum hopcount 100
  EIGRP maximum metric variance 1
  Redistributing: eigrp 1
  EIGRP NSF-aware route hold timer is 240s
  Automatic network summarization is not in effect
  Maximum path: 4
  Routing for Networks:
    172.16.0.0
  Passive Interface(s):
    Serial0/0/1
  Routing Information Sources:
    Gateway         Distance      Last Update
    172.16.12.1     90           00:27:57
  Distance: internal 90 external 170
```

How does preventing hello packets out of an interface affect the update capabilities of EIGRP out that interface?

Is this behavior more like RIP or like OSPF in regard to the **passive-interface** command?

Step 4: Manually summarize with EIGRP.

You can have EIGRP summarize routes sent out an interface to make routing updates more efficient by using the **ip summary-address eigrp as network mask** command.

- a. Have R1 advertise one supernet for loopbacks 48 and 49 to R2. Do not summarize loopbacks 50 and 51 in this statement, because these will be summarized in Step 9.

```
R1(config)# interface serial 0/0/0
R1(config-if)# ip summary-address eigrp 1 192.168.48.0 255.255.254.0

*Feb 23 18:20:21.219: %DUAL-5-NBRCHANGE: IP-EIGRP(0) 1: Neighbor 172.16.12.2
(Serial0/0/0) is resync: summary configured
```

```
R1# show ip route eigrp
 172.16.0.0/24 is subnetted, 6 subnets
D    172.16.23.0 [90/41024000] via 172.16.12.2, 00:45:21, Serial0/0/0
D    172.16.2.0 [90/40640000] via 172.16.12.2, 00:45:21, Serial0/0/0
D    172.16.100.0 [90/40640000] via 172.16.12.2, 00:08:12, Serial0/0/0
D    192.168.48.0/23 is a summary, 04:27:07, Null0
```

- b. Verify the configuration with the **show ip route eigrp** and **show ip route 192.168.48.0 255.255.254.0** commands on R1. Notice the administrative distance for this route.

```
R1# show ip route 192.168.48.0 255.255.254.0
Routing entry for 192.168.48.0/23, supernet
  Known via "eigrp 1", distance 5, metric 128256, type internal
  Redistributing via eigrp 1
  Routing Descriptor Blocks:
  * directly connected, via Null0
    Route metric is 128256, traffic share count is 1
    Total delay is 5000 microseconds, minimum bandwidth is 10000000 Kbit
    Reliability 255/255, minimum MTU 1514 bytes
    Loading 1/255, Hops 0
```

Why does EIGRP make the administrative distance different for summary routes?

Step 5: Additional OSPF configuration.

OSPF is already partially configured on R2 and R3.

- a. You need to add the area 10 configuration to R2 and the area 20 configuration to R3 to complete the configuration.

```
R2(config)# router ospf 1
R2(config-router)# network 172.16.100.0 0.0.0.255 area 10
```

```
R3(config)# router ospf 1
R3(config-router)# network 192.168.8.0 0.0.3.255 area 20
```

- b. Verify that your adjacencies come up with the **show ip ospf neighbor** command, and make sure that you have routes from OSPF populating the R2 routing table using the **show ip route ospf** command.

```
R2# show ip ospf neighbor
```

Neighbor ID	Pri	State	Dead Time	Address	Interface
192.168.40.1	0	FULL/ -	00:00:35	172.16.23.3	Serial0/0/1

```
R3# show ip ospf neighbor
```

Neighbor ID	Pri	State	Dead Time	Address	Interface
172.16.2.1	0	FULL/ -	00:00:35	172.16.23.2	Serial0/0/1

```
R2# show ip route ospf
```

```
O 192.168.30.0/24 [110/1563] via 172.16.23.3, 00:12:10, Serial0/0/1
  192.168.8.0/32 is subnetted, 1 subnets
O IA 192.168.8.1 [110/1563] via 172.16.23.3, 00:20:48, Serial0/0/1
O 192.168.25.0/24 [110/1563] via 172.16.23.3, 00:12:10, Serial0/0/1
  192.168.9.0/32 is subnetted, 1 subnets
O IA 192.168.9.1 [110/1563] via 172.16.23.3, 00:20:48, Serial0/0/1
  192.168.10.0/32 is subnetted, 1 subnets
O IA 192.168.10.1 [110/1563] via 172.16.23.3, 00:20:48, Serial0/0/1
O 192.168.40.0/24 [110/1563] via 172.16.23.3, 00:12:10, Serial0/0/1
  172.16.0.0/24 is subnetted, 6 subnets
O 172.16.3.0 [110/1563] via 172.16.23.3, 00:21:26, Serial0/0/1
  192.168.11.0/32 is subnetted, 1 subnets
O IA 192.168.11.1 [110/1563] via 172.16.23.3, 00:20:48, Serial0/0/1
O 192.168.20.0/24 [110/1563] via 172.16.23.3, 00:12:10, Serial0/0/1
O 192.168.35.0/24 [110/1563] via 172.16.23.3, 00:12:10, Serial0/0/1
```

```
R3# show ip route ospf
```

```
172.16.0.0/16 is variably subnetted, 3 subnets, 2 masks
O IA 172.16.100.1/32 [110/1563] via 172.16.23.2, 00:00:15, Serial0/0/1
```

Notice that for the newly added loopback interfaces, OSPF advertised /32 destination prefixes (for example, R2 has a route to 192.168.8.1/32 in its routing table).

- c. Override this default behavior by using the **ip ospf network point-to-point** command on the OSPF loopback interfaces on R2 and R3. You can copy and paste the following configurations to save time.

Router R2:

```
interface loopback 100
 ip ospf network point-to-point
```

Router R3:

(Only configure the point-to-point network type for the newly added loopbacks in area 20. The area 0 loopbacks were configured in Lab 4-1.)

```
interface loopback 8
 ip ospf network point-to-point
```

```
interface loopback 9
 ip ospf network point-to-point
```

```
interface loopback 10
```

```
ip ospf network point-to-point
```

```
interface loopback 11
ip ospf network point-to-point
```

Note: You can also use the **interface range** command to configure multiple interfaces simultaneously, as shown below.

```
interface range lo 8 - 11
ip ospf network point-to-point
```

- d. Verify the configuration with the **show ip route** command on R2. Notice that the routes now each show on one line with the /24 major network mask.

```
R2# show ip route
<output omitted>
```

Gateway of last resort is not set

```
O   192.168.30.0/24 [110/1563] via 172.16.23.3, 00:27:11, Serial0/0/1
O IA 192.168.8.0/24 [110/1563] via 172.16.23.3, 00:08:39, Serial0/0/1
O   192.168.25.0/24 [110/1563] via 172.16.23.3, 00:27:11, Serial0/0/1
O IA 192.168.9.0/24 [110/1563] via 172.16.23.3, 00:08:39, Serial0/0/1
O IA 192.168.10.0/24 [110/1563] via 172.16.23.3, 00:08:39, Serial0/0/1
O   192.168.40.0/24 [110/1563] via 172.16.23.3, 00:27:11, Serial0/0/1
    172.16.0.0/24 is subnetted, 6 subnets
C     172.16.23.0 is directly connected, Serial0/0/1
C     172.16.12.0 is directly connected, Serial0/0/0
D     172.16.1.0 [90/40640000] via 172.16.12.1, 00:47:33, Serial0/0/0
C     172.16.2.0 is directly connected, Loopback0
O     172.16.3.0 [110/1563] via 172.16.23.3, 00:36:27, Serial0/0/1
C     172.16.100.0 is directly connected, Loopback100
O IA 192.168.11.0/24 [110/1563] via 172.16.23.3, 00:08:41, Serial0/0/1
O   192.168.20.0/24 [110/1563] via 172.16.23.3, 00:27:13, Serial0/0/1
D   192.168.51.0/24 [90/40640000] via 172.16.12.1, 00:47:36, Serial0/0/0
D   192.168.50.0/24 [90/40640000] via 172.16.12.1, 00:47:36, Serial0/0/0
O   192.168.35.0/24 [110/1563] via 172.16.23.3, 00:27:13, Serial0/0/1
D   192.168.70.0/24 [90/40640000] via 172.16.12.1, 00:47:36, Serial0/0/0
D   192.168.48.0/23 [90/40640000] via 172.16.12.1, 00:40:01, Serial0/0/0
```

Notice that R2 is the only router with knowledge of all routes in the topology at this point, because it is involved with both routing protocols.

Step 6: Summarize OSPF areas at the ABR.

Review the R2 routing table. Notice the inter-area routes for the R3 loopbacks in area 20.

- a. Summarize the areas into a single inter-area route using the **area area range network mask** command on R3.

```
R3(config)# router ospf 1
R3(config-router)# area 20 range 192.168.8.0 255.255.252.0
```

- b. On R2, verify the summarization with the **show ip route ospf** command on R2.

```
R2# show ip route ospf
O   192.168.30.0/24 [110/1563] via 172.16.23.3, 02:38:46, Serial0/0/1
O   192.168.25.0/24 [110/1563] via 172.16.23.3, 02:38:46, Serial0/0/1
O   192.168.40.0/24 [110/1563] via 172.16.23.3, 02:38:46, Serial0/0/1
    172.16.0.0/24 is subnetted, 6 subnets
O     172.16.3.0 [110/1563] via 172.16.23.3, 02:38:46, Serial0/0/1
```

```
O 192.168.20.0/24 [110/1563] via 172.16.23.3, 02:38:46, Serial0/0/1
O 192.168.35.0/24 [110/1563] via 172.16.23.3, 02:38:46, Serial0/0/1
O IA 192.168.8.0/22 [110/1563] via 172.16.23.3, 00:00:07, Serial0/0/1
```

Where can you summarize in OSPF?

Compare and contrast OSPF and EIGRP in terms of where summarization takes place.

Explain the synchronization requirement in OSPF that eliminates other routers as points of summarization.

Why or why not does EIGRP have this requirement?

Step 7: Configure mutual redistribution between OSPF and EIGRP.

- a. Under the OSPF process on R2, issue the **redistribute eigrp 1 subnets** command. The **subnets** command is necessary because, by default, OSPF only redistributes classful networks and supernets. A default seed metric is not required for OSPF. Under the EIGRP process, issue the **redistribute ospf 1 metric 10000 100 255 1 1500** command, which tells EIGRP to redistribute OSPF process 1 with these metrics: bandwidth of 10000, delay of 100, reliability of 255/255, load of 1/255, and a MTU of 1500. Like RIP, EIGRP requires a seed metric. You can also set a default seed metric with the **default-metric** command.

```
R2(config)# router ospf 1
R2(config-router)# redistribute eigrp 1 subnets
R2(config-router)# exit
```

```
R2(config)# router eigrp 1
R2(config-router)# redistribute ospf 1 metric 10000 100 255 1 1500
```


or

```
R2(config-router)# default-metric 10000 100 255 1 1500  
R2(config-router)# redistribute ospf 1
```

- b. Issue the **show ip protocols** command on the redistributing router, R2. Compare your output with the following output.

```
R2# show ip protocols  
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 172.16.2.1  
  It is an area border and autonomous system boundary router  
  Redistributing External Routes from,  
  eigrp 1, includes subnets in redistribution  
  Number of areas in this router is 2. 2 normal 0 stub 0 nssa  
  Maximum path: 4  
  Routing for Networks:  
    172.16.23.0 0.0.0.255 area 0  
    172.16.100.0 0.0.0.255 area 10  
  Reference bandwidth unit is 100 mbps  
  Routing Information Sources:  
    Gateway          Distance      Last Update  
    192.168.40.1      110          00:00:33  
  Distance: (default is 110)
```

```
Routing Protocol is "eigrp 1"  
  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  
  EIGRP metric weight K1=1, K2=0, K3=1, K4=0, K5=0  
  EIGRP maximum hopcount 100  
  EIGRP maximum metric variance 1  
  Redistributing: ospf 1, eigrp 1  
  EIGRP NSF-aware route hold timer is 240s  
  Automatic network summarization is not in effect  
  Maximum path: 4  
  Routing for Networks:  
    172.16.0.0  
  Passive Interface(s):  
    Serial0/0/1  
  Routing Information Sources:  
    Gateway          Distance      Last Update  
    172.16.12.1      90           00:00:49  
  Distance: internal 90 external 170
```

- c. Display the routing tables on R1 and R3 so that you can see the redistributed routes. Redistributed OSPF routes display on R1 as D EX, which means that they are external EIGRP routes. Redistributed EIGRP routes are tagged in the R3 routing table as O E2, which means that they are OSPF external type 2. Type 2 is the default OSPF external type.

```
R1# show ip route  
<output omitted>
```

```
Gateway of last resort is not set
```

CCNPv6 ROUTE

```
D EX 192.168.30.0/24 [170/40537600] via 172.16.12.2, 00:00:05, Serial0/0/0
D EX 192.168.25.0/24 [170/40537600] via 172.16.12.2, 00:00:05, Serial0/0/0
D EX 192.168.40.0/24 [170/40537600] via 172.16.12.2, 00:00:05, Serial0/0/0
    172.16.0.0/24 is subnetted, 6 subnets
D      172.16.100.0 [90/40640000] via 172.16.12.2, 00:38:02, Serial0/0/0
D      172.16.23.0 [90/41024000] via 172.16.12.2, 00:38:02, Serial0/0/0
C      172.16.12.0 is directly connected, Serial0/0/0
C      172.16.1.0 is directly connected, Loopback0
D      172.16.2.0 [90/40640000] via 172.16.12.2, 00:38:02, Serial0/0/0
D EX 172.16.3.0 [170/40537600] via 172.16.12.2, 00:00:06, Serial0/0/0
D EX 192.168.20.0/24 [170/40537600] via 172.16.12.2, 00:00:06, Serial0/0/0
C      192.168.51.0/24 is directly connected, Loopback51
C      192.168.50.0/24 is directly connected, Loopback50
D EX 192.168.35.0/24 [170/40537600] via 172.16.12.2, 00:00:06, Serial0/0/0
C      192.168.49.0/24 is directly connected, Loopback49
C      192.168.70.0/24 is directly connected, Loopback70
C      192.168.48.0/24 is directly connected, Loopback48
D EX 192.168.8.0/22 [170/40537600] via 172.16.12.2, 00:00:07, Serial0/0/0
D      192.168.48.0/23 is a summary, 04:19:50, Null0
```

R3# **show ip route**

<output omitted>

Gateway of last resort is not set

```
C      192.168.30.0/24 is directly connected, Loopback30
C      192.168.8.0/24 is directly connected, Loopback8
C      192.168.25.0/24 is directly connected, Loopback25
C      192.168.9.0/24 is directly connected, Loopback9
C      192.168.10.0/24 is directly connected, Loopback10
C      192.168.40.0/24 is directly connected, Loopback40
    172.16.0.0/24 is subnetted, 6 subnets
C      172.16.23.0 is directly connected, Serial0/0/1
O E2 172.16.12.0 [110/20] via 172.16.23.2, 00:41:48, Serial0/0/1
O E2 172.16.1.0 [110/20] via 172.16.23.2, 00:41:48, Serial0/0/1
O E2 172.16.2.0 [110/20] via 172.16.23.2, 00:41:48, Serial0/0/1
O IA 172.16.100.0 [110/1563] via 172.16.23.2, 00:41:48, Serial0/0/1
C      172.16.3.0 is directly connected, Loopback0
C      192.168.11.0/24 is directly connected, Loopback11
C      192.168.20.0/24 is directly connected, Loopback20
O E2 192.168.51.0/24 [110/20] via 172.16.23.2, 00:41:48, Serial0/0/1
O E2 192.168.50.0/24 [110/20] via 172.16.23.2, 00:41:48, Serial0/0/1
C      192.168.35.0/24 is directly connected, Loopback35
O E2 192.168.70.0/24 [110/20] via 172.16.23.2, 00:41:48, Serial0/0/1
O      192.168.8.0/22 is a summary, 01:34:48, Null0
O E2 192.168.48.0/23 [110/20] via 172.16.23.2, 00:41:48, Serial0/0/1
```

- d. Verify full connectivity with the following Tcl script:

R1# **tclsh**

```
foreach address {
172.16.1.1
192.168.48.1
192.168.49.1
192.168.50.1
192.168.51.1
192.168.70.1
172.16.12.1
```

```

172.16.2.1
172.16.100.1
172.16.12.2
172.16.23.2
172.16.3.1
192.168.20.1
192.168.25.1
192.168.30.1
192.168.35.1
192.168.40.1
192.168.8.1
192.168.9.1
192.168.10.1
192.168.11.1
172.16.23.3
} { ping $address }

```

Step 8: Filter redistribution with route maps.

One way to filter prefixes is with a route map. When used for filtering prefixes, a route map works like an access list. It has multiple statements that are read in a sequential order. Each statement can be a deny or permit and can have a match clause for a variety of attributes, such as the route or a route tag. You can also include route attributes in each statement that will be set if the match clause is met.

- a. Before filtering the R3 loopback 25 and 30 networks from being redistributed into EIGRP on R2, display the R1 routing table and verify that those two routes currently appear there.

```

R1# show ip route eigrp
D EX 192.168.30.0/24 [170/40537600] via 172.16.12.2, 00:04:28, Serial0/0/0
D EX 192.168.25.0/24 [170/40537600] via 172.16.12.2, 00:04:28, Serial0/0/0
D EX 192.168.40.0/24 [170/40537600] via 172.16.12.2, 00:04:28, Serial0/0/0
    172.16.0.0/24 is subnetted, 6 subnets
D      172.16.23.0 [90/41024000] via 172.16.12.2, 00:42:25, Serial0/0/0
D      172.16.2.0 [90/40640000] via 172.16.12.2, 00:42:25, Serial0/0/0
D EX   172.16.3.0 [170/40537600] via 172.16.12.2, 00:04:28, Serial0/0/0
D      172.16.100.0 [90/40640000] via 172.16.12.2, 01:34:26, Serial0/0/0
D EX 192.168.20.0/24 [170/40537600] via 172.16.12.2, 00:04:28, Serial0/0/0
D EX 192.168.35.0/24 [170/40537600] via 172.16.12.2, 00:04:28, Serial0/0/0
D EX 192.168.8.0/22 [170/40537600] via 172.16.12.2, 00:04:28, Serial0/0/0
D     192.168.48.0/23 is a summary, 04:24:12, Null0

```

There are multiple ways to configure this filtering. For this exercise, configure an access list that matches these two network addresses and a route map that denies based on a match for that access list.

- b. Configure the access list as follows:

```

R2(config)# access-list 1 permit 192.168.25.0
R2(config)# access-list 1 permit 192.168.30.0

```

- c. Configure a route map with a statement that denies based on a match with this access list. Then add a permit statement without a match statement, which acts as an explicit permit all.

```

R2(config)# route-map SELECTED-DENY deny 10
R2(config-route-map)# match ip address 1
R2(config-route-map)# route-map SELECTED-DENY permit 20

```

- d. Apply this route map by redoing the **redistribute** command with the route map under the EIGRP process.

```

R2(config)# router eigrp 1
R2(config-router)# redistribute ospf 1 route-map SELECTED-DENY metric 64 100
255 1 1500

```

- e. As an alternative, if you previously configured a default metric under EIGRP, you can simply use the following command.

```
R2(config-router)# redistribute ospf 1 route-map SELECTED-DENY
```

- f. Verify that these routes are filtered out in the R1 routing table.

```
R1# show ip route eigrp
D EX 192.168.40.0/24 [170/40537600] via 172.16.12.2, 00:07:24, Serial0/0/0
    172.16.0.0/24 is subnetted, 6 subnets
D     172.16.23.0 [90/41024000] via 172.16.12.2, 00:45:21, Serial0/0/0
D     172.16.2.0 [90/40640000] via 172.16.12.2, 00:45:21, Serial0/0/0
D EX  172.16.3.0 [170/40537600] via 172.16.12.2, 00:07:24, Serial0/0/0
D     172.16.100.0 [90/40640000] via 172.16.12.2, 00:45:21, Serial0/0/0
D EX 192.168.20.0/24 [170/40537600] via 172.16.12.2, 00:07:24, Serial0/0/0
D EX 192.168.35.0/24 [170/40537600] via 172.16.12.2, 00:07:24, Serial0/0/0
D EX 192.168.8.0/22 [170/40537600] via 172.16.12.2, 00:07:24, Serial0/0/0
D    192.168.48.0/23 is a summary, 04:27:07, Null0
```

Step 9: Summarize external routes into OSPF at the ASBR.

You cannot summarize routes redistributed into OSPF using the **area range** command. This command is effective only on routes internal to the specified area. Instead, use the OSPF **summary-address network mask** command.

- a. Before you make any changes, display the R3 routing table.

```
R3# show ip route ospf
    172.16.0.0/24 is subnetted, 6 subnets
O E2   172.16.12.0 [110/20] via 172.16.23.2, 00:00:07, Serial0/0/1
O E2   172.16.1.0 [110/20] via 172.16.23.2, 00:00:07, Serial0/0/1
O E2   172.16.2.0 [110/20] via 172.16.23.2, 00:00:07, Serial0/0/1
O IA   172.16.100.0 [110/1563] via 172.16.23.2, 00:00:07, Serial0/0/1
O E2 192.168.70.0/24 [110/20] via 172.16.23.2, 00:00:07, Serial0/0/1
O    192.168.8.0/22 is a summary, 00:00:07, Null0
O E2 192.168.51.0/24 [110/20] via 172.16.23.2, 00:00:07, Serial0/0/1
O E2 192.168.50.0/24 [110/20] via 172.16.23.2, 00:00:07, Serial0/0/1
O E2 192.168.48.0/23 [110/20] via 172.16.23.2, 00:00:07, Serial0/0/1
```

Notice the three external routes for the R1 loopback interfaces 48 through 51. Two of the loopbacks are already summarized to one /23.

Which mask should you use to summarize all four of the loopbacks to one prefix?

- b. You can summarize this all into one supernet on R2 using the following commands.

```
R2(config)# router ospf 1
R2(config-router)# summary-address 192.168.48.0 255.255.252.0
```

- c. Verify this action in the R3 routing table.

```
R3# show ip route ospf
    172.16.0.0/24 is subnetted, 6 subnets
O E2   172.16.12.0 [110/20] via 172.16.23.2, 01:40:45, Serial0/0/1
O E2   172.16.1.0 [110/20] via 172.16.23.2, 00:48:54, Serial0/0/1
O E2   172.16.2.0 [110/20] via 172.16.23.2, 01:40:45, Serial0/0/1
O IA   172.16.100.0 [110/1563] via 172.16.23.2, 01:40:45, Serial0/0/1
O E2 192.168.70.0/24 [110/20] via 172.16.23.2, 00:48:54, Serial0/0/1
```

```
O 192.168.8.0/22 is a summary, 01:41:55, Null0
O E2 192.168.48.0/22 [110/20] via 172.16.23.2, 00:00:08, Serial0/0/1
```

What would happen if loopback 50 on R1 were to become unreachable by R2?

Would data destined for 192.168.50.0/24 from R3 still be sent to R2?

Would data destined for 192.168.50.0/24 from R2 continue to be sent to R1?

- d. If you are unsure of the outcome, shut down the interface on R1. Issue the ICMP **tracert** command to 192.168.50.1 from R3 and then from R2. Check your output against the output and analysis in Appendix A. Remember to issue the **no shutdown** command when you are finished checking.

Is this a desirable outcome? Explain.

Step 10: Modify EIGRP distances.

- a. By default, EIGRP uses an administrative distance of 90 for internal routes and 170 for external routes. You can see this in the R1 routing table and in the output of the **show ip protocols** command.

```
R1# show ip route eigrp
D EX 192.168.40.0/24 [170/40537600] via 172.16.12.2, 00:04:03, Serial0/0/0
    172.16.0.0/24 is subnetted, 6 subnets
D     172.16.23.0 [90/41024000] via 172.16.12.2, 00:04:03, Serial0/0/0
D     172.16.2.0 [90/40640000] via 172.16.12.2, 00:04:03, Serial0/0/0
D EX   172.16.3.0 [170/40537600] via 172.16.12.2, 00:04:03, Serial0/0/0
D     172.16.100.0 [90/40640000] via 172.16.12.2, 00:04:03, Serial0/0/0
D EX 192.168.20.0/24 [170/40537600] via 172.16.12.2, 00:04:03, Serial0/0/0
D EX 192.168.35.0/24 [170/40537600] via 172.16.12.2, 00:04:03, Serial0/0/0
D EX 192.168.8.0/22 [170/40537600] via 172.16.12.2, 00:04:03, Serial0/0/0
D   192.168.48.0/23 is a summary, 3d17h, Null0
D EX 192.168.48.0/22 [170/40537600] via 172.16.12.2, 00:04:03, Serial0/0/0
```

```
R1# show ip protocols
<output omitted>
```

```
Address Summarization:
  192.168.48.0/23 for Serial0/0/0
    Summarizing with metric 128256
Maximum path: 4
Routing for Networks:
  172.16.0.0
  192.168.0.0/16
```

```

Routing Information Sources:
  Gateway           Distance      Last Update
  (this router)      90           00:34:33
  172.16.12.2        90           00:16:35
Distance: internal 90 external 170

```

- b. You can change the administrative distance with the **distance eigrp internal external** command. This command is only applicable locally. Change the distance to 95 for internal routes and 165 for external routes.

```

R1(config)# router eigrp 1
R1(config-router)# distance eigrp 95 165

```

Note: The EIGRP neighbor adjacency will be re-negotiated:

```

R1#
*May  3 00:28:38.379: %DUAL-5-NBRCHANGE: IP-EIGRP(0) 1: Neighbor 172.16.12.2
(Serial0/0/0) is down: route configuration changed
*May  3 00:28:41.503: %DUAL-5-NBRCHANGE: IP-EIGRP(0) 1: Neighbor 172.16.12.2
(Serial0/0/0) is up: new adjacency

```

- c. Verify the change in the routing table with the **show ip route eigrp** and **show ip protocols** commands.

```

R1# show ip route eigrp
D EX 192.168.40.0/24 [165/40537600] via 172.16.12.2, 00:04:03, Serial0/0/0
    172.16.0.0/24 is subnetted, 6 subnets
D     172.16.23.0 [95/41024000] via 172.16.12.2, 00:04:03, Serial0/0/0
D     172.16.2.0 [95/40640000] via 172.16.12.2, 00:04:03, Serial0/0/0
D EX   172.16.3.0 [165/40537600] via 172.16.12.2, 00:04:03, Serial0/0/0
D     172.16.100.0 [95/40640000] via 172.16.12.2, 00:04:03, Serial0/0/0
D EX 192.168.20.0/24 [165/40537600] via 172.16.12.2, 00:04:03, Serial0/0/0
D EX 192.168.35.0/24 [165/40537600] via 172.16.12.2, 00:04:03, Serial0/0/0
D EX 192.168.8.0/22 [165/40537600] via 172.16.12.2, 00:04:03, Serial0/0/0
D   192.168.48.0/23 is a summary, 3d17h, Null0
D EX 192.168.48.0/22 [165/40537600] via 172.16.12.2, 00:04:03, Serial0/0/0

```

```

R1# show ip protocols
Routing Protocol is "eigrp 1"
<output omitted>

```

```

Routing Information Sources:
  Gateway           Distance      Last Update
  172.16.12.2        95           00:00:00
Distance: internal 95 external 165

```

Step 11: Modify OSPF distances.

You can also modify individual OSPF distances. By default, all OSPF distances are 110, but you can change the intra-area, inter-area, and external route distances using the **distance ospf intra-area distance inter-area distance external distance** command. All the command arguments are optional, so you can change only what you need to.

- a. Before changing anything, display the R3 routing table.

```

R3# show ip route ospf
    172.16.0.0/24 is subnetted, 6 subnets
O E2   172.16.12.0 [110/20] via 172.16.23.2, 01:40:45, Serial0/0/1
O E2   172.16.1.0 [110/20] via 172.16.23.2, 00:48:54, Serial0/0/1
O E2   172.16.2.0 [110/20] via 172.16.23.2, 01:40:45, Serial0/0/1
O IA   172.16.100.0 [110/1563] via 172.16.23.2, 01:40:45, Serial0/0/1

```

CCNPv6 ROUTE

```
O E2 192.168.70.0/24 [110/20] via 172.16.23.2, 00:48:54, Serial0/0/1
O   192.168.8.0/22 is a summary, 01:41:55, Null0
O E2 192.168.48.0/22 [110/20] via 172.16.23.2, 00:00:08, Serial0/0/1
```

- b. Change the intra-area distance to 105, inter-area distance to 115, and external routes to 175 on R3.

```
R3(config)# router ospf 1
R3(config-router)# distance ospf intra-area 105 inter-area 115 external 175
```

- c. Verify the change in the routing table. Unfortunately, the only information that you can get from the output of the **show ip protocols** command is the default distance, which is the intra-area distance.

```
R3# show ip route ospf
 172.16.0.0/24 is subnetted, 6 subnets
O E2   172.16.12.0 [175/20] via 172.16.23.2, 00:00:05, Serial0/0/1
O E2   172.16.1.0 [175/20] via 172.16.23.2, 00:00:05, Serial0/0/1
O E2   172.16.2.0 [175/20] via 172.16.23.2, 00:00:05, Serial0/0/1
O IA   172.16.100.0 [115/1563] via 172.16.23.2, 00:00:05, Serial0/0/1
O E2  192.168.70.0/24 [175/20] via 172.16.23.2, 00:00:05, Serial0/0/1
O   192.168.8.0/22 is a summary, 00:00:05, Null0
O E2  192.168.48.0/22 [175/20] via 172.16.23.2, 00:00:05, Serial0/0/1
```

```
R3# show ip protocols
```

```
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 192.168.40.1
  It is an area border router
  Number of areas in this router is 2. 2 normal 0 stub 0 nssa
  Maximum path: 4
  Routing for Networks:
    172.16.0.0 0.0.255.255 area 0
    192.168.8.0 0.0.3.255 area 20
    192.168.0.0 0.0.255.255 area 0
  Reference bandwidth unit is 100 mbps
  Passive Interface(s):
    FastEthernet0/0
    FastEthernet0/1
    Serial0/0/0
    Serial0/1/0
    Serial0/1/1
    Loopback0
    Loopback8
    Loopback9
    Loopback10
    Loopback11
  Passive Interface(s):
    Loopback20
    Loopback25
    Loopback30
    Loopback35
    Loopback40
    VoIP-Null0
  Routing Information Sources:
    Gateway         Distance      Last Update
    (this router)   110          00:03:04
    172.16.2.1     110          00:03:04
  Distance: (default is 105)
```

Challenge: Change the Administrative Distance on R2

The previous two steps demonstrated using the **distance** command in a fairly inconsequential environment. In which types of scenarios would the **distance** command be more valuable?

On R2, you are running both EIGRP and OSPF. Imagine a fourth router, R4, connected to both R1 and R3. R4 is redistributing between the two routing protocols.

Using the default administrative distances for EIGRP and OSPF, which protocol would be preferred in the routing table for destination prefixes in native OSPF networks and why?

Which protocol would be preferred in the routing table for destination prefixes for native EIGRP networks?

Instead of adding the 172.16.1.0/24 networks natively to EIGRP using a **network** statement, add the networks using the **redistribute connected** command in EIGRP configuration mode on R1.

With the default administrative distances set, what would the administrative distance be for that prefix on R2 in EIGRP and in OSPF? Explain why.

How could you make the EIGRP path prefer this route? Is there more than one way?

Could using the **distance** command in this situation cause asymmetric routing? Explain.

Router Interface Summary Table

Router Interface Summary				
Router Model	Ethernet Interface #1	Ethernet Interface #2	Serial Interface #1	Serial Interface #2
1700	Fast Ethernet 0 (FA0)	Fast Ethernet 1 (FA1)	Serial 0 (S0)	Serial 1 (S1)
1800	Fast Ethernet 0/0 (FA0/0)	Fast Ethernet 0/1 (FA0/1)	Serial 0/0/0 (S0/0/0)	Serial 0/0/1 (S0/0/1)
2600	Fast Ethernet 0/0 (FA0/0)	Fast Ethernet 0/1 (FA0/1)	Serial 0/0 (S0/0)	Serial 0/1 (S0/1)
2800	Fast Ethernet 0/0 (FA0/0)	Fast Ethernet 0/1 (FA0/1)	Serial 0/0/0 (S0/0/0)	Serial 0/0/1 (S0/0/1)

Note: To find out how the router is configured, look at the interfaces to identify the type of router and how many interfaces the router has. Rather than list all combinations of configurations for each router class, this table includes identifiers for the possible combinations of Ethernet and serial interfaces in the device. The table does not include any other type of interface, even though a specific router might contain one. For example, for an ISDN BRI interface, the string in parenthesis is the legal abbreviation that can be used in Cisco IOS commands to represent the interface.

Appendix A: Exploring Black Hole Operations

- a. Configure R1 and shut down the loopback 50 interface:

```
R1(config)# interface loopback 50
R1(config-if)# shutdown
```

- b. On R2, you **should** see the following output.

```
R2# show ip route
```

```
Gateway of last resort is not set
```

```
O   192.168.30.0/24 [110/1563] via 172.16.23.3, 18:53:52, Serial0/0/1
O   192.168.25.0/24 [110/1563] via 172.16.23.3, 18:53:52, Serial0/0/1
O   192.168.40.0/24 [110/1563] via 172.16.23.3, 18:53:52, Serial0/0/1
    172.16.0.0/24 is subnetted, 6 subnets
C     172.16.100.0 is directly connected, Loopback100
C     172.16.23.0 is directly connected, Serial0/0/1
C     172.16.12.0 is directly connected, Serial0/0/0
D     172.16.1.0 [90/40640000] via 172.16.12.1, 18:54:06, Serial0/0/0
C     172.16.2.0 is directly connected, Loopback0
O     172.16.3.0 [110/1563] via 172.16.23.3, 18:53:53, Serial0/0/1
O   192.168.20.0/24 [110/1563] via 172.16.23.3, 18:53:53, Serial0/0/1
D   192.168.51.0/24 [90/40640000] via 172.16.12.1, 18:54:07, Serial0/0/0
O   192.168.35.0/24 [110/1563] via 172.16.23.3, 18:53:53, Serial0/0/1
D   192.168.70.0/24 [90/40640000] via 172.16.12.1, 18:54:07, Serial0/0/0
O IA 192.168.8.0/22 [110/1563] via 172.16.23.3, 18:53:54, Serial0/0/1
D   192.168.48.0/23 [90/40640000] via 172.16.12.1, 18:54:08, Serial0/0/0
O   192.168.48.0/22 is a summary, 17:16:44, Null0
```

Notice the absence of 192.168.50.0/24 in a specific route in the R2 routing table.

- c. Begin debugging all incoming IP packets on R2, and then issue the **ping 192.168.50.1** command.

```
R2# debug ip packet
```

CCNPv6 ROUTE

```
R2# ping 192.168.50.1
(debug output cleaned up so as to be readable)
```

```
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.168.50.1, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
```

```
*Oct 17 16:39:14.147: IP: s=172.16.2.1 (local), d=192.168.50.1 (Null0), len
100, sending
...
```

```
R2# undebug all
```

```
R2# traceroute 192.168.50.1
```

```
Type escape sequence to abort.
Tracing the route to 192.168.50.1
```

```
 1  *  *  *
 2  *  *  *
 3  *  *  *
 4  *  *  *
 5  *  *  *
 6  *  *  *
 7  *  *  *
 8  *  *  *
 9  *  *  *
```

<output omitted>

The summary route, pointing to the Null0 interface as the next hop, acts as a “catch all” for any traffic generated by R2 or forwarded to R2 with the destination network 192.168.48.0/24. R2 sends traffic to the Null0 virtual interface, as shown by the IP packet debugging output highlighted above.

R2 is not able to ping the R1 shutdown loopback interface because the 192.168.50.0/24 route no longer exists in the routing table.

Check to see if network 192.168.50.0/24, or a supernet of it, is in the routing table of R3.

```
R3# show ip route 192.168.50.1
```

```
Routing entry for 192.168.48.0/22, supernet
```

```
Known via "ospf 1", distance 110, metric 20, type extern 2, forward metric 1562
```

```
Last update from 172.16.23.2 on Serial0/0/1, 00:39:17 ago
```

```
Routing Descriptor Blocks:
```

```
* 172.16.23.2, from 172.16.2.1, 00:39:17 ago, via Serial0/0/1
```

```
Route metric is 20, traffic share count is 1
```

- d. Begin debugging all IP and ICMP packets on R3. Ping the address 192.168.50.1 from R3. Try to trace the route from R3 to 192.168.50.1.

```
R3# debug ip packet
R3# debug ip icmp
```

```
R3# ping 192.168.50.1
<output omitted>
```

Type escape sequence to abort.

Sending 5, 100-byte ICMP Echos to 192.168.50.1, timeout is 2 seconds:

U.U.U

Success rate is 0 percent (0/5)

*Oct 17 16:49:21.023: IP: tableid=0, s=172.16.23.3 (local), d=192.168.50.1 (Serial0/0/1), routed via FIB

*Oct 17 16:49:21.047: ICMP: dst (172.16.23.3) host unreachable rcv from 172.16.23.2

R3# **undebug all**

R3# **traceroute 192.168.50.1**

Type escape sequence to abort.

Tracing the route to 192.168.50.1

```
 1 172.16.23.2 12 msec 12 msec 16 msec
 2 172.16.23.2 !H !H *
```

Analyze the process indicated by the ICMP responses. You might also want to refer to debugging messages for ICMP and IP packets on R2.

1. R3 generates an ICMP echo request (ping) to 192.168.50.1.
2. R3 looks up the (next-hop address, outgoing interface) pair for the longest matching prefix containing 192.168.50.1 in the IP routing table. It finds (172.16.23.2, Serial0/0/1).
3. R3 routes the IP packet to (172.16.23.2, Serial0/0/1).
4. R2 receives the IP packet from R3 on interface Serial0/0/1.
5. R2 looks up the (next-hop address, outgoing interface) pair for the longest prefix matching containing 192.168.50.1 in the IP routing table. The longest matching prefix that the routing table returns is 192.168.48.0/22, for which the routing table responds with (null, Null0) because it has no next-hop address or physical outgoing interface.
6. R2 realizes that this packet was routed remotely to it but that it has no route, so it sends an ICMP Type 3, Code 1 (host unreachable) packet to the source address of the packet, 172.16.23.3.
7. R2 looks up the (next-hop address, outgoing interface) pair for 172.16.23.3 and resolves it to (172.16.23.3, Serial0/0/1).
8. R2 then routes the ICMP packet for destination 172.16.23.3, normally 172.16.23.3 through Serial0/0/1.
9. R3 receives a packet destined for its local address 172.16.23.3 and reads the packet, sending the ICMP "Host Unreachable" message to the ping output.

Note: For more information about how routers respond to unreachable hosts, see RFC 792 (ICMP) at <http://tools.ietf.org/html/rfc792> and RFC 4443 (ICMPv6) at <http://tools.ietf.org/html/rfc4443>.

Notice that R2 sends R3 an ICMP Type 3, Code 1 reply indicating that it does not have a route to the host 192.168.50.1. This ICMP "Host Unreachable" message is not only sent in response to pings or traceroutes (also a form of ICMP) but for all IP traffic.

- e. If you were to use Telnet to 192.168.50.1, you would receive the following message based on the ICMP response from R2:

```
R3#telnet 192.168.50.1
Trying 192.168.50.1 ...
```

```
% Destination unreachable; gateway or host down
```

```
R3#
```

This is not an example of Telnet timing out, but of intelligent network protocols responding to routing issues in the network.

This summarization problem is a classic example of a “black hole” in a domain, which simply means traffic passing through the network destined for that subnet is discarded at some point along the way. Thankfully, ICMP informs sources of when their traffic is being discarded.

- f. Do not forget to issue the **no shutdown** command on the R1 loopback 50 interface to re-enable routing to this network.

```
R1(config)# interface loopback 50  
R1(config-if)# no shutdown
```