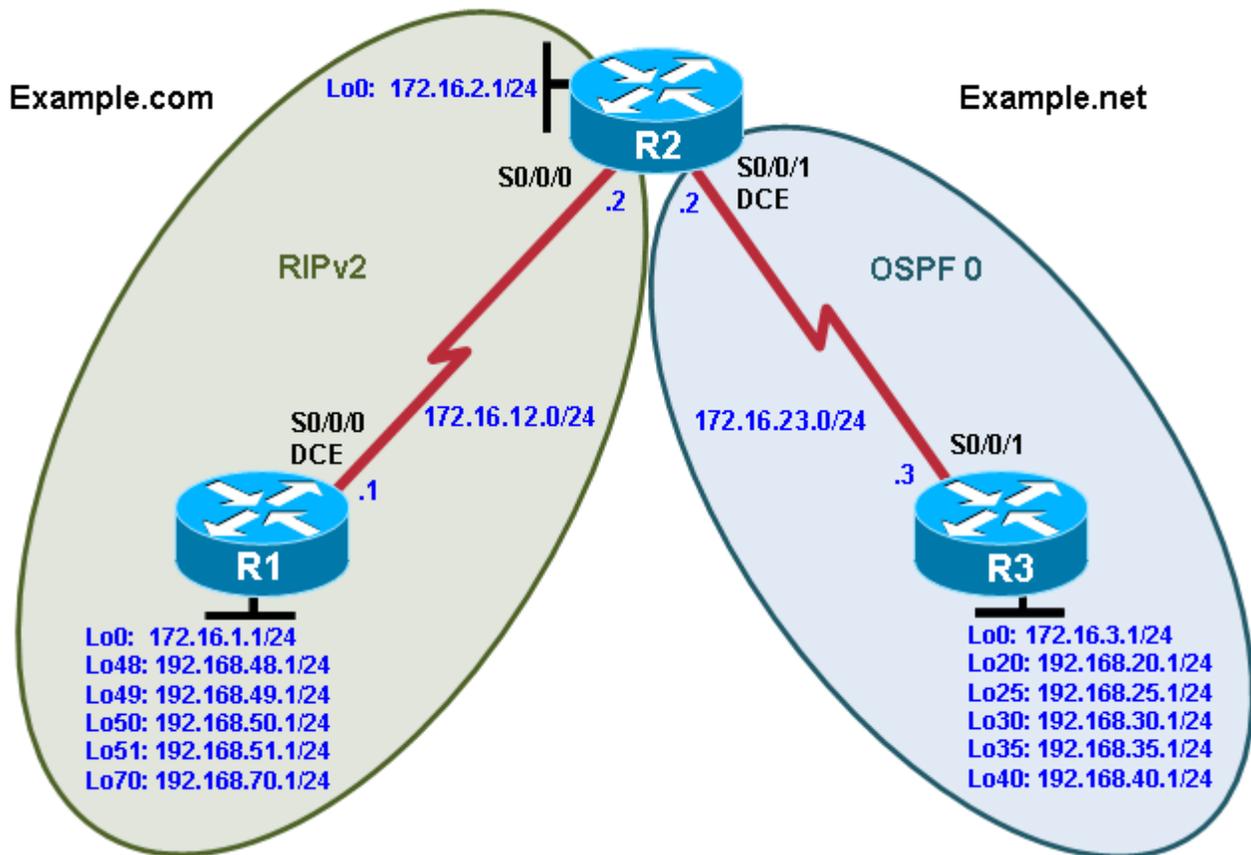


Chapter 4 Lab 4-1, Redistribution Between RIP and OSPF

Topology



Objectives

- Review configuration and verification of RIP and OSPF.
- Configure passive interfaces in both RIP and OSPF.
- Filter routing updates using distribute lists.
- Redistribute static routes into RIP.
- Redistribute RIP routes into OSPF.
- Redistribute OSPF routes into RIP.
- Originate a default route into OSPF.
- Set a default seed metric.
- Modify OSPF external network types.
- Configure summary addresses.

Background

Two online booksellers, Example.com and Example.net, have merged and now need a short-term solution to inter-domain routing. Since these companies provide client services to Internet users, it is essential to have minimal downtime during the transition.

Example.com is a small firm running RIP, while Example.net has a somewhat larger network running OSPF. The diagram identifies R2 as the router that will bridge the two networks. Because it is imperative that the two booksellers continuously deliver Internet services, you should bridge these two routing domains without interfering with each router's path through its own routing domain to the Internet.

The CIO determines that it is preferable to keep the two protocol domains shown in the diagram during the transition period, because the network engineers on each side need to understand the other's network before deploying a long-term solution. Redistribution will be a short-term solution.

In this scenario, R1 and R2 are running RIPv2, but the 172.16.23.0/24 network between R2 and R3 is running OSPF. You need to configure R2 to enable these two routing protocols to interact to allow full connectivity between all networks.

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 assign addresses.

- a. Configure all loopback interfaces on the three routers in the diagram. Configure the serial interfaces with the IP addresses, bring them up, and set a DCE clock rate where appropriate.

```
R1(config)# interface Loopback0
R1(config-if)# ip address 172.16.1.1 255.255.255.0
R1(config-if)# interface Loopback48
R1(config-if)# ip address 192.168.48.1 255.255.255.0
R1(config-if)# interface Loopback49
R1(config-if)# ip address 192.168.49.1 255.255.255.0
R1(config-if)# interface Loopback50
R1(config-if)# ip address 192.168.50.1 255.255.255.0
R1(config-if)# interface Loopback51
R1(config-if)# ip address 192.168.51.1 255.255.255.0
R1(config-if)# interface Loopback70
R1(config-if)# ip address 192.168.70.1 255.255.255.0
R1(config-if)# interface Serial0/0/0
R1(config-if)# ip address 172.16.12.1 255.255.255.0
R1(config-if)# clock rate 64000
R1(config-if)# bandwidth 64
R1(config-if)# no shutdown

R2(config)# interface Loopback0
R2(config-if)# ip address 172.16.2.1 255.255.255.0
R2(config-if)# interface Serial0/0/0
R2(config-if)# ip address 172.16.12.2 255.255.255.0
```

```
R2(config-if)# bandwidth 64
R2(config-if)# no shutdown
R2(config-if)# interface Serial0/0/1
R2(config-if)# ip address 172.16.23.2 255.255.255.0
R2(config-if)# clock rate 64000
R2(config-if)# bandwidth 64
R2(config-if)# no shutdown
```

```
R3(config)# interface Loopback0
R3(config-if)# ip address 172.16.3.1 255.255.255.0
R3(config-if)# interface Loopback20
R3(config-if)# ip address 192.168.20.1 255.255.255.0
R3(config-if)# interface Loopback25
R3(config-if)# ip address 192.168.25.1 255.255.255.0
R3(config-if)# interface Loopback30
R3(config-if)# ip address 192.168.30.1 255.255.255.0
R3(config-if)# interface Loopback35
R3(config-if)# ip address 192.168.35.1 255.255.255.0
R3(config-if)# interface Loopback40
R3(config-if)# ip address 192.168.40.1 255.255.255.0
R3(config-if)# interface Serial0/0/1
R3(config-if)# ip address 172.16.23.3 255.255.255.0
R3(config-if)# bandwidth 64
R3(config-if)# no shutdown
```

- b. (Optional) On each router, create an enable secret password. Configure the console line for synchronous logging and no timeout. Configure the vty lines to allow Telnet to and remote configuration of network devices.

R1 example:

```
R1(config)# enable secret cisco

R1(config)# line con 0
R1(config-line)# logging synchronous
R1(config-line)# exec-timeout 0 0

R1(config)# line vty 0 4
R1(config-line)# password cisco
R1(config-line)# login
```

- c. Verify that you can ping across the serial links when you are finished. Use the following Tcl script to check full and partial connectivity throughout this lab.

```
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
172.16.12.2
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
172.16.23.3
} { ping $address }
```

At this point, the only pings that you should receive back are those of the connected networks of the router from which you are pinging.

Step 2: Configure RIPv2.

Configuring RIPv2 on a router is fairly simple:

- Type the global configuration command **router rip** to enter RIP configuration mode.
- Enable RIPv2 with the **version 2** command.
- Enter the **no auto-summary** command to disable automatic summarization at classful network boundaries.
- Add the networks you want using the **network network** command.

Unlike EIGRP and OSPF, the RIP **network** command only requires the classful network address to be entered and does not support a wildcard mask. This behavior is inherited from the classful RIPv1 protocol configuration and is kept for backward compatibility with older Cisco IOS versions that would not otherwise be able to process **network** commands with wildcard masks. Classful protocols do not support subnets; therefore, subnet or wildcard masks are unnecessary.

Based on the topology diagram, which major networks need to be advertised into RIP for R1?

Which major networks need to be advertised into RIP for R2?

- a. Apply the following commands to R1 and R2.

```
R1(config)# router rip
R1(config-router)# version 2
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
R1(config-router)# network 192.168.51.0
R1(config-router)# network 192.168.70.0
```

```
R2(config)# router rip
R2(config-router)# version 2
R2(config-router)# no auto-summary
```

```
R2(config-router)# network 172.16.0.0
```

- b. Verify that the RIP routes were learned from the other routers using the **show ip route rip** command on each router.

```
R1# show ip route rip
    172.16.0.0/24 is subnetted, 4 subnets
R       172.16.23.0 [120/1] via 172.16.12.2, 00:00:03, Serial0/0/0
R       172.16.2.0 [120/1] via 172.16.12.2, 00:00:03, Serial0/0/0
```

```
R2# show ip route rip
    172.16.0.0/24 is subnetted, 4 subnets
R       172.16.1.0 [120/1] via 172.16.12.1, 00:00:29, Serial0/0/0
R       192.168.51.0/24 [120/1] via 172.16.12.1, 00:00:29, Serial0/0/0
R       192.168.50.0/24 [120/1] via 172.16.12.1, 00:00:29, Serial0/0/0
R       192.168.49.0/24 [120/1] via 172.16.12.1, 00:00:29, Serial0/0/0
R       192.168.70.0/24 [120/1] via 172.16.12.1, 00:00:29, Serial0/0/0
R       192.168.48.0/24 [120/1] via 172.16.12.1, 00:00:29, Serial0/0/0
```

- c. You can also verify which routes are coming in from RIP advertisements with the **show ip rip database** command.

```
R1# show ip rip database
172.16.0.0/16      auto-summary
172.16.1.0/24     directly connected, Loopback0
172.16.2.0/24     [1] via 172.16.12.2, 00:00:06, Serial0/0/0
172.16.12.0/24    directly connected, Serial0/0/0
172.16.23.0/24    [1] via 172.16.12.2, 00:00:06, Serial0/0/0
192.168.48.0/24   auto-summary
192.168.48.0/24   directly connected, Loopback48
192.168.49.0/24   auto-summary
192.168.49.0/24   directly connected, Loopback49
192.168.50.0/24   auto-summary
192.168.50.0/24   directly connected, Loopback50
192.168.51.0/24   auto-summary
192.168.51.0/24   directly connected, Loopback51
192.168.70.0/24   auto-summary
192.168.70.0/24   directly connected, Loopback70
```

```
R2# show ip rip database
172.16.0.0/16      auto-summary
172.16.1.0/24     [1] via 172.16.12.1, 00:00:10, Serial0/0/0
172.16.2.0/24     directly connected, Loopback0
172.16.12.0/24    directly connected, Serial0/0/0
172.16.23.0/24    directly connected, Serial0/0/1
192.168.48.0/24   auto-summary
192.168.48.0/24   [1] via 172.16.12.1, 00:00:10, Serial0/0/0
192.168.49.0/24   auto-summary
192.168.49.0/24   [1] via 172.16.12.1, 00:00:10, Serial0/0/0
192.168.50.0/24   auto-summary
192.168.50.0/24   [1] via 172.16.12.1, 00:00:10, Serial0/0/0
192.168.51.0/24   auto-summary
192.168.51.0/24
```

```
[1] via 172.16.12.1, 00:00:10, Serial0/0/0
192.168.70.0/24 auto-summary
192.168.70.0/24
[1] via 172.16.12.1, 00:00:10, Serial0/0/0
```

Step 3: Configure passive interfaces in RIP.

- a. On R1, use the **show ip route rip** command to view the RIP routes in the routing table. Notice that the network for the serial interface of R2 that connects to R3 is present, even though you do not have a RIP neighbor on that interface. This is because the entire class B network 172.16.0.0 /16 was added to RIP on R2.

```
R1# show ip route rip
      172.16.0.0/24 is subnetted, 4 subnets
R       172.16.23.0 [120/1] via 172.16.12.2, 00:00:03, Serial0/0/0
R       172.16.2.0 [120/1] via 172.16.12.2, 00:00:03, Serial0/0/0
```

- b. Issue the **show ip protocols** command to verify that RIPv2 updates are being sent out both serial interfaces.

```
R2# show ip protocols
Routing Protocol is "rip"
  Outgoing update filter list for all interfaces is not set
  Incoming update filter list for all interfaces is not set
  Sending updates every 30 seconds, next due in 13 seconds
  Invalid after 180 seconds, hold down 180, flushed after 240
  Redistributing: rip
  Default version control: send version 2, receive version 2
  Interface          Send  Recv  Triggered RIP  Key-chain
  Serial0/0/0         2    2
  Serial0/0/1         2    2
  Loopback0           2    2
  Automatic network summarization is not in effect
  Maximum path: 4
  Routing for Networks:
    172.16.0.0
  Routing Information Sources:
    Gateway           Distance      Last Update
    172.16.12.1       120           00:00:26
  Distance: (default is 120)
```

For security reasons and to reduce unnecessary traffic, RIP updates should not be propagated into the OSPF domain. You can disable sending updates with the **passive-interface interface_type interface_number** router configuration command

- c. On R2, configure the serial interface connecting to R3 as passive. Notice that the interface is no longer listed in the output of the **show ip protocols** command.

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

```
R2# show ip protocols
Routing Protocol is "rip"
  Outgoing update filter list for all interfaces is not set
  Incoming update filter list for all interfaces is not set
  Sending updates every 30 seconds, next due in 23 seconds
  Invalid after 180 seconds, hold down 180, flushed after 240
  Redistributing: rip
  Default version control: send version 2, receive version 2
```

Interface	Send	Recv	Triggered	RIP	Key-chain
Serial0/0/0	2	2			
Loopback0	2	2			

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	120	00:00:17

Distance: (default is 120)

- d. On R1, issue the **show ip route rip** command. Notice that the 172.16.23.0 network is still in the routing table and being sourced from RIP.

```
R1# show ip route rip
    172.16.0.0/24 is subnetted, 4 subnets
R       172.16.23.0 [120/1] via 172.16.12.2, 00:00:19, Serial0/0/0
R       172.16.2.0 [120/1] via 172.16.12.2, 00:00:19, Serial0/0/0
```

Making an interface in RIP passive only disables updates from being sent through RIP. It does not affect routes being received through it.

What are some reasons to prevent RIP from sending updates out a particular interface?

Putting a RIPv2 interface in passive mode saves the router from sending multicast RIP packets out an interface that has no neighbors.

Does RIPv2 send advertisements out loopback interfaces?

- e. If you are unsure, monitor the output of the **debug ip rip** command to verify your answer. On R1 and R2, configure all loopbacks from which RIPv2 is sending advertisements in passive state with the **passive-interface** command.

```
R1(config)# router rip
R1(config-router)# passive-interface loopback 0
R1(config-router)# passive-interface loopback 48
R1(config-router)# passive-interface loopback 49
```

```
R1(config-router)# passive-interface loopback 50
R1(config-router)# passive-interface loopback 51
R1(config-router)# passive-interface loopback 70
```

```
R2(config)# router rip
R2(config-router)# passive-interface loopback 0
```

When running RIPv2, implement passive interfaces as a common practice to save CPU processor cycles and bandwidth on interfaces that do not have multicast RIPv2 neighbors.

Note: An alternative to making each loopback interface on R1 passive is to make all interfaces passive with the **passive-interface default** command in router configuration mode. Then make any interfaces that need to send updates, such as S0/0/0, nonpassive.

```
R1(config)# router rip
R1(config-router)# passive-interface default
R1(config-router)# no passive-interface Serial0/0/0
```

Step 4: Summarize a supernet with RIP.

- On R2, issue the **show ip route rip** command. Notice that you can see all prefixes from R1 in the R2 routing table.

```
R2# show ip route rip
    172.16.0.0/24 is subnetted, 4 subnets
R       172.16.1.0 [120/1] via 172.16.12.1, 00:00:29, Serial0/0/0
R       192.168.51.0/24 [120/1] via 172.16.12.1, 00:00:29, Serial0/0/0
R       192.168.50.0/24 [120/1] via 172.16.12.1, 00:00:29, Serial0/0/0
R       192.168.49.0/24 [120/1] via 172.16.12.1, 00:00:29, Serial0/0/0
R       192.168.70.0/24 [120/1] via 172.16.12.1, 00:00:29, Serial0/0/0
R       192.168.48.0/24 [120/1] via 172.16.12.1, 00:00:29, Serial0/0/0
```

In preparing for redistribution, you want to redistribute the minimum number of destination prefixes into each of the routing protocols. Which RIP routes should you summarize because they are contiguous and which mask should you use?

Under normal circumstances, you could simply summarize the four consecutive class-C networks with the **ip summary address rip** command on the R1 serial 0/0/0 interface. However, the RIP implementation in the Cisco IOS Software does not allow summarizing to a mask length that is less than the classful network prefix (in this case, 24 bits). This limitation does not affect other routing protocols. If you do try, you receive the following error message:

```
R1(config)# interface serial 0/0/0
R1(config-if)# ip summary-address rip 192.168.48.0 255.255.252.0
Summary mask must be greater or equal to major net
```

Recall from the EIGRP labs that summary routes display in the summarizing device's routing table as having the next hop being the Null0 interface. You can create an entry manually using the **ip route** command and redistribute it into RIP, thereby emulating the approach of EIGRP to a certain extent.

- To get around the **ip summary-address rip** message error, create a static route on R1 to summarize the networks of loopbacks 48 through 51. Then redistribute the route on R1.

```
R1(config)# ip route 192.168.48.0 255.255.252.0 null0
R1(config)# router rip
R1(config-router)# redistribute static
```

This solution might seem unusual, but for RIPv2, it resembles many effects of summarization as performed in other routing protocols like EIGRP or OSPF. Again, this is not a limitation of RIPv2, but rather a Cisco IOS implementation issue.

- c. On R1 and R2, verify that the RIP supernet has been added to the routing table with the **show ip route** command.

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

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

```
      172.16.0.0/24 is subnetted, 4 subnets
R       172.16.23.0 [120/1] via 172.16.12.2, 00:00:27, Serial0/0/0
C       172.16.12.0 is directly connected, Serial0/0/0
C       172.16.1.0 is directly connected, Loopback0
R       172.16.2.0 [120/1] via 172.16.12.2, 00:00:27, Serial0/0/0
C      192.168.51.0/24 is directly connected, Loopback51
C      192.168.50.0/24 is directly connected, Loopback50
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
S      192.168.48.0/22 is directly connected, Null0
```

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

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

```
      172.16.0.0/24 is subnetted, 4 subnets
C       172.16.23.0 is directly connected, Serial0/0/1
C       172.16.12.0 is directly connected, Serial0/0/0
R       172.16.1.0 [120/1] via 172.16.12.1, 00:00:05, Serial0/0/0
C       172.16.2.0 is directly connected, Loopback0
R      192.168.51.0/24 [120/1] via 172.16.12.1, 00:00:05, Serial0/0/0
R      192.168.50.0/24 [120/1] via 172.16.12.1, 00:00:05, Serial0/0/0
R      192.168.49.0/24 [120/1] via 172.16.12.1, 00:00:05, Serial0/0/0
R      192.168.70.0/24 [120/1] via 172.16.12.1, 00:00:07, Serial0/0/0
R      192.168.48.0/24 [120/1] via 172.16.12.1, 00:00:07, Serial0/0/0
R      192.168.48.0/22 [120/1] via 172.16.12.1, 00:00:07, Serial0/0/0
```

Will this route to Null0 affect routing to prefixes with longer addresses on R1? Explain.

Step 5: Suppress routes using prefix lists.

Sometimes you might not want to advertise certain networks out a particular interface, or you might want to filter updates as they come in. This is possible with distance-vector routing protocols, such as RIP or EIGRP. However, link-state protocols are less flexible, because every router in an area is required to have a synchronized database as a condition for full adjacency.

Distribute lists can be used with either access lists or prefix lists to filter routes by network address. With prefix lists, they can also be configured to filter routes by subnet masks.

In this scenario, you want to filter updates from R1 to R2, allowing only the networks of Loopback 0 and Loopback 70 and the summary route to be advertised. You want to suppress the more specific prefixes so that routing tables are kept small, and CPU processor cycles on the routers are not wasted.

The 22-bit summary and the 24-bit major network address both have the same address, so access lists will not accomplish the filtering correctly. Therefore, it is necessary to use prefix lists.

To create a prefix list or add a prefix list entry, use the **ip prefix-list** command in global configuration mode.

```
ip prefix-list {list-name | list-number} {deny network/length | permit
network/length} [ge ge-length] [le le-length]
```

The **ge** keyword represents the “greater than or equal to” operator. The **le** keyword represents the “less than or equal to” operator. If both the **ge** and **le** keywords are omitted, the prefix list is processed using an exact match.

- a. On R1, use a prefix list as a distribution filter to prevent the more specific routes to loopbacks 48 through 51 from being advertised. Allow all other destination networks, including the summary route.

```
R1(config)# ip prefix-list RIP-OUT permit 192.168.48.0/22
R1(config)# ip prefix-list RIP-OUT deny 192.168.48.0/22 le 24
R1(config)# ip prefix-list RIP-OUT permit 0.0.0.0/0 le 32
```

Line 1 of the prefix list permits the summary route and nothing else, because no other route can match that network address with a mask of exactly 22 bits.

Line 2 denies all prefixes with a network address in the 192.168.48.0/22 block of addresses that have subnet masks from 22 bits to 24 bits. This removes exactly four network addresses matching the 22, 23, and 24 bits in length of the subnet mask. Line 2 would deny the 192.168.48.0/22 summary route you created if Line 1 did not explicitly permit the summary route.

Line 3 allows all IPv4 prefixes that are not explicitly denied in previous statements of the prefix list.

- b. From the RIP configuration prompt on R1, apply this access list with the **distribute-list** command.

```
R1(config)# router rip
R1(config-router)# distribute-list prefix RIP-OUT out serial0/0/0
```

- c. On R2, verify that the filtering has taken place using the **show ip route rip** and **show ip rip database** commands.

```
R2# show ip route rip
    172.16.0.0/24 is subnetted, 4 subnets
R       172.16.1.0 [120/1] via 172.16.12.1, 00:00:12, Serial0/0/0
R       192.168.70.0/24 [120/1] via 172.16.12.1, 00:00:12, Serial0/0/0
R       192.168.48.0/22 [120/1] via 172.16.12.1, 00:00:12, Serial0/0/0
```

Note: You might need to issue the **clear ip route *** command on R2 to see the removal of the more specific R1 prefixes. Also, if the network 192.168.48.0/22 does not appear on R2, this is incorrect behavior and might be corrected in recent versions of Cisco IOS Software. A workaround is to remove the **network 192.168.48.0** command from RIP and issue the **clear ip route *** command on R1.

```
R2# show ip rip database
172.16.0.0/16      auto-summary
172.16.1.0/24
    [1] via 172.16.12.1, 00:00:11, Serial0/0/0
172.16.2.0/24    directly connected, Loopback0
172.16.12.0/24   directly connected, Serial0/0/0
```

```
172.16.23.0/24    directly connected, Serial0/0/1
192.168.48.0/22
    [1] via 172.16.12.1, 00:00:11, Serial0/0/0
192.168.70.0/24    auto-summary
192.168.70.0/24
    [1] via 172.16.12.1, 00:00:11, Serial0/0/0
```

Why would you want to filter updates being sent out or coming in?

Step 6: Configure OSPF.

- Configure single-area OSPF between R2 and R3. On R2, include just the serial link connecting to R3. On R3, include the serial link and all loopback interfaces.

```
R2(config)# router ospf 1
R2(config-router)# network 172.16.23.0 0.0.0.255 area 0
```

```
R3(config)# router ospf 1
R3(config-router)# network 172.16.0.0 0.0.255.255 area 0
R3(config-router)# network 192.168.0.0 0.0.255.255 area 0
```

```
15:01:37.047: %OSPF-5-ADJCHG: Process 1, Nbr 172.16.2.1 on Serial0/0/1 from
LOADING to FULL, Loading Done
```

- On R3, change the network type for the loopback interfaces to point-to-point so that they are advertised with the correct subnet mask (/24 instead of /32).

```
R3(config)# interface Loopback0
R3(config-if)# ip ospf network point-to-point
R3(config-if)# interface Loopback20
R3(config-if)# ip ospf network point-to-point
R3(config-if)# interface Loopback25
R3(config-if)# ip ospf network point-to-point
R3(config-if)# interface Loopback30
R3(config-if)# ip ospf network point-to-point
R3(config-if)# interface Loopback35
R3(config-if)# ip ospf network point-to-point
R3(config-if)# interface Loopback40
R3(config-if)# ip ospf network point-to-point
```

- Verify the OSPF adjacencies on R2 and R3 with the **show ip ospf neighbors** command. Also make sure that you have routes from OSPF populating the routing tables with 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/ -	00:00:37	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/ -	00:00:39	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, 00:01:23, Serial0/0/1
O   192.168.25.0/24 [110/1563] via 172.16.23.3, 00:01:23, Serial0/0/1
O   192.168.40.0/24 [110/1563] via 172.16.23.3, 00:01:23, Serial0/0/1
    172.16.0.0/24 is subnetted, 5 subnets
O       172.16.3.0 [110/1563] via 172.16.23.3, 00:01:23, Serial0/0/1
O   192.168.20.0/24 [110/1563] via 172.16.23.3, 00:01:23, Serial0/0/1
O   192.168.35.0/24 [110/1563] via 172.16.23.3, 00:01:23, Serial0/0/1
```

R3# **show ip route ospf**

R3#

Note that output of the **show ip route ospf** command on R3 is blank.

The **network 192.168.0.0 0.0.255.255 area 0** command allows OSPF to involve interfaces that have IP addresses in that range.

A common misconception is that OSPF advertises the entire range of the network given in the router's **network** statement; it does not. However, it does advertise any connected subnets in that range of addresses to adjacent routers. You can verify this by viewing the output of the **show ip route** command on R2. Do you see a 192.168.0.0/16 supernet?

R2 is the only router with all routes in the topology (except for those that were filtered out), because it is involved with both routing protocols.

Step 7: Configure passive interfaces in OSPF.

Passive interfaces save CPU cycles, router memory, and link bandwidth by preventing broadcast and multicast routing updates on interfaces that have no neighbors. In link-state protocols, adjacencies must be formed before routers exchange routing information. The **passive-interface** command in OSPF configuration mode prevents an interface from sending or processing OSPF packets on that interface.

OSPF included the R3 loopback interfaces in its **network** statements shown in Step 6.

- a. On R3, configure Loopback0 as a passive interface in OSPF. At the OSPF router configuration prompt, use the **passive-interface interface_type interface_number** command.

```
R3(config-router)# passive-interface loopback 0
```

How is this different from the RIP version of this command?

- b. Cisco IOS Software provides a quick way of selecting interfaces for passive mode. Use the **passive-interface default** command to make all interfaces passive. Then use the **no passive-interface interface interface_number** command to bring the Serial0/0/1 interface out of passive mode.

```
R3(config)# router ospf 1
R3(config-router)# passive-interface default
R3(config-router)#
*Oct 15 01:49:44.174: %OSPF-5-ADJCHG: Process 1, Nbr 172.16.2.1 on
Serial0/0/1 from FULL to DOWN, Neighbor Down: Interface down or detached
R3(config-router)# no passive-interface serial 0/0/1
R3(config-router)#
*Oct 15 01:49:55.438: %OSPF-5-ADJCHG: Process 1, Nbr 172.16.2.1 on
Serial0/0/1 from LOADING to FULL, Loading Done
```

- c. You can verify the application of this command by issuing the **show ip protocols** command.

```
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
  Number of areas in this router is 1. 1 normal 0 stub 0 nssa
  Maximum path: 4
  Routing for Networks:
    172.16.0.0 0.0.255.255 area 0
    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
    Loopback20
    Loopback25
    Loopback30
    Loopback35
    Loopback40
  Routing Information Sources:
    Gateway         Distance      Last Update
    172.16.2.1      110          00:03:04
  Distance: (default is 110)
```

Step 8: Allow one-way redistribution.

- a. On R2, configure OSPF to redistribute into RIP under the RIP configuration prompt with the **redistribute ospf process metric metric** command, where *process* is the OSPF process number, and *metric* is the default metric with which you want to originate the routes into RIP. If you do not specify a default metric in RIP, it gives routes an infinite metric and they are not advertised.

```
R2(config)# router rip
R2(config-router)# redistribute ospf 1 metric 4
```

- b. Verify the redistribution with the **show ip protocols** command.

```
R2# show ip protocols
Routing Protocol is "rip"
  Outgoing update filter list for all interfaces is not set
  Incoming update filter list for all interfaces is not set
```

```
Sending updates every 30 seconds, next due in 24 seconds
Invalid after 180 seconds, hold down 180, flushed after 240
Redistributing: rip, ospf 1
Default version control: send version 2, receive version 2
  Interface          Send Recv Triggered RIP Key-chain
  Serial0/0/0        2     2
Automatic network summarization is not in effect
Maximum path: 4
Routing for Networks:
  172.16.0.0
Passive Interface(s):
  Serial0/0/1
  Loopback0
Routing Information Sources:
  Gateway           Distance      Last Update
  172.16.12.1       120          00:00:19
Distance: (default is 120)
```

<output omitted>

- c. On R1, look at the routing table with the **show ip route rip** command. It has all the routes in the topology.

```
R1# show ip route rip
R       192.168.30.0 [120/4] via 172.16.12.2, 00:00:11, Serial0/0/0
R       192.168.25.0 [120/4] via 172.16.12.2, 00:00:11, Serial0/0/0
R       192.168.40.0 [120/4] via 172.16.12.2, 00:00:11, Serial0/0/0
       172.16.0.0/24 is subnetted, 5 subnets
R       172.16.23.0 [120/1] via 172.16.12.2, 00:00:11, Serial0/0/0
R       172.16.2.0 [120/1] via 172.16.12.2, 00:00:11, Serial0/0/0
R       172.16.3.0 [120/4] via 172.16.12.2, 00:00:11, Serial0/0/0
R       192.168.20.0 [120/4] via 172.16.12.2, 00:00:11, Serial0/0/0
R       192.168.35.0 [120/4] via 172.16.12.2, 00:00:11, Serial0/0/0
```

- d. On R1, ping a loopback on R3. Notice that it shows that R1 has a route to R3, but R3 does not have a route back to R1.

```
R1# ping 192.168.30.1
```

```
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.168.30.1, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
```

- e. On R1, verify that R3 does not have a route back with the **traceroute** command.

```
R1# traceroute 192.168.30.1
```

```
Type escape sequence to abort.
Tracing the route to 192.168.30.1

 0 172.16.12.2 12 msec 12 msec 16 msec
 1 * * *
 2 * * *
 3 * * *
 4 * * *
```

<output omitted>

To address this problem, you can originate a default route into OSPF that points toward R2 so that the pings are routed back toward R2. R2 uses its information from RIPv2 to send pings back to R1.

- f. From the OSPF configuration prompt, issue the **default-information originate always** command to force R2 to advertise a default route in OSPF.

```
R2(config)# router ospf 1
R2(config-router)# default-information originate always
```

- g. Verify that this route is present in the R3 routing table.

```
R3# show ip route ospf
O*E2 0.0.0.0/0 [110/1] via 172.16.23.2, 00:05:13, Serial0/0/1
```

You should now have full connectivity between all networks in the diagram.

- h. Use the Tcl script from Step 1 to verify full connectivity.

Step 9: Redistribute between two routing protocols.

You can substitute this default route with actual, more specific routes.

- a. On R2, under the OSPF router configuration prompt, remove the default route advertisement with the **no default-information originate always** command. Next, use the **redistribute rip** command. You do not need to specify a default metric in OSPF. Notice the warning.

```
R2(config)# router ospf 1
R2(config-router)# no default-information originate always
R2(config-router)# redistribute rip
% Only classful networks will be redistributed
```

- b. If you display the routing table on R3, the external OSPF routes that were added are the 192.168.70.0/24 and 192.168.48.0/22 networks.

```
R3# show ip route ospf
O E2 192.168.70.0/24 [110/20] via 172.16.23.2, 00:00:51, Serial0/0/1
O E2 192.168.48.0/22 [110/20] via 172.16.23.2, 00:00:51, Serial0/0/1
```

This is because, by default, OSPF only accepts classful networks and supernets when redistributing into it. The only classful network coming into R2 from RIP is the class C network 192.168.70.0, and the only supernet is the 192.168.48.0/22.

- c. You can modify this behavior by adding the **subnets** keyword to the **redistribute** command.

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

- d. On R3, verify the configuration with the **show ip route ospf** command.

```
R3# show ip route ospf
172.16.0.0/24 is subnetted, 5 subnets
O E2 172.16.12.0 [110/20] via 172.16.23.2, 00:00:01, Serial0/0/1
O E2 172.16.1.0 [110/20] via 172.16.23.2, 00:00:01, Serial0/0/1
O E2 172.16.2.0 [110/20] via 172.16.23.2, 00:00:01, Serial0/0/1
O E2 192.168.70.0/24 [110/20] via 172.16.23.2, 00:04:19, Serial0/0/1
O E2 192.168.48.0/22 [110/20] via 172.16.23.2, 00:04:19, Serial0/0/1
```

You should again have full connectivity between all networks in the diagram.

- e. Run the Tcl script on each router to verify full connectivity.

Step 10: Set a default seed metric.

Under any routing protocol, you can specify a default seed metric to be used for redistribution instead of, or in addition to, setting metrics on a per-protocol basis. A seed metric is a protocol-independent feature of the Cisco IOS Software that is usually configured when redistributing into distance-vector protocols.

Notice that the metric listed in the R3 routing table is 20.

```
R3# show ip route ospf
      172.16.0.0/24 is subnetted, 5 subnets
O E2   172.16.12.0 [110/20] via 172.16.23.2, 00:00:01, Serial0/0/1
O E2   172.16.1.0 [110/20] via 172.16.23.2, 00:00:01, Serial0/0/1
O E2   172.16.2.0 [110/20] via 172.16.23.2, 00:00:01, Serial0/0/1
O E2  192.168.70.0/24 [110/20] via 172.16.23.2, 00:04:19, Serial0/0/1
O E2  192.168.48.0/22 [110/20] via 172.16.23.2, 00:04:19, Serial0/0/1
```

You can override the global creation of a default seed metric on a per-protocol basis by using the *metric* argument in a redistribution command. You can also use the **metric** command under other routing protocols.

- a. On R2, in OSPF configuration mode, issue the **default-metric metric** command to configure a default metric for redistributed routes. The default metric for all OSPF redistributed routes is 20, except for BGP, which is 1. Setting the metric for RIP to a higher number makes it less preferable to routes redistributed from other routing protocols.

```
R2(config)# router ospf 1
R2(config-router)# default-metric 10000
```

- b. Verify the new metric in the R3 routing table. It might take some time for the new metric to propagate.

```
R3# show ip route ospf
      172.16.0.0/24 is subnetted, 5 subnets
O E2   172.16.12.0 [110/10000] via 172.16.23.2, 00:02:56, Serial0/0/1
O E2   172.16.1.0 [110/10000] via 172.16.23.2, 00:02:56, Serial0/0/1
O E2   172.16.2.0 [110/10000] via 172.16.23.2, 00:02:56, Serial0/0/1
O E2  192.168.70.0/24 [110/10000] via 172.16.23.2, 00:02:56, Serial0/0/1
O E2  192.168.48.0/22 [110/10000] via 172.16.23.2, 00:02:56, Serial0/0/1
```

Step 11: Change the OSPF external network type.

Look at the R3 routing table. Notice that the external (redistributed) routes have O E2 as their type. In the output, O means OSPF, and E2 means external, type 2. OSPF has two external metric types, and E2 is the default. External type 1 metrics increase like a usual route, whereas external type 2 metrics do not increase as they get advertised through the OSPF domain. Also notice that the metric is exactly the same as the seed metric in the previous step.

Where would an external type 1 metric be useful?

Where would an external type 2 metric be useful?

- a. You can change the external type using the **metric-type** argument with the **redistribute** command. Change the type to E1 for RIP redistributed routes.

```
R2(config)# router ospf 1
R2(config-router)# redistribute rip subnets metric-type 1
```

CCNPv6 ROUTE

- b. Display the R3 routing table again.

```
R3# show ip route ospf
      172.16.0.0/24 is subnetted, 5 subnets
O E1   172.16.12.0 [110/11562] via 172.16.23.2, 00:03:05, Serial0/0/1
O E1   172.16.1.0 [110/11562] via 172.16.23.2, 00:03:05, Serial0/0/1
O E1   172.16.2.0 [110/11562] via 172.16.23.2, 00:03:05, Serial0/0/1
O E1 192.168.70.0/24 [110/11562] via 172.16.23.2, 00:03:05, Serial0/0/1
O E1 192.168.48.0/22 [110/11562] via 172.16.23.2, 00:03:05, Serial0/0/1
```

Which attributes of the routes changed?

Note: Be sure to save your final configurations through Step 11 for use in Lab 4-2, "Redistribution Between EIGRP and OSPF."

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.