

This case study addresses the issue of integrating Routing Information Protocol (RIP) networks with Open Shortest Path First (OSPF) networks. Most OSPF networks also use RIP to communicate with hosts or to communicate with portions of the internetwork that do not use OSPF. Cisco supports both the RIP and OSPF protocols and provides a way to exchange routing information between RIP and OSPF networks. This case study provides examples of how to complete the following phases in redistributing information between RIP and OSPF networks, including the following topics:

- [Configuring a RIP Network](#)
- [Adding OSPF to the Center of a RIP Network](#)
- [Adding OSPF Areas](#)
- [Setting Up Mutual Redistribution](#)

Guide Contents
Internetworking Design Basics
Designing various internetworks
Network Enhancements
IP Routing Concepts
UDP Broadcast Flooding
Large-Scale H.323 Network Design for Service Providers
LAN Switching
Subnetting an IP Address Space
IBM Serial Link Implementation Notes
References and Recommended Reading

Contents

- [1 Configuring a RIP Network](#)
 - ◆ [1.1 Figure: A RIP network](#)
 - ◆ [1.2 Table: RIP Network Address Assignments](#)
 - ◆ [1.3 Configuration File Examples](#)
- [2 Adding OSPF to the Center of a RIP Network](#)
 - ◆ [2.1 Figure: RIP network with OSPF at the center](#)
 - ◆ [2.2 Configuration File Examples](#)
- [3 Adding OSPF Areas](#)
 - ◆ [3.1 Figure: Configuring route summarization between OSPF areas](#)
 - ◆ [3.2 Table: OSPF Address Assignments](#)
 - ◆ [3.3 Configuration File Examples](#)
- [4 Setting Up Mutual Redistribution](#)
 - ◆ [4.1 Figure: Mutual redistribution between RIP and OSPF networks](#)
 - ◆ [4.2 Router A](#)
- [5 Summary](#)

Configuring a RIP Network

[Figure: A RIP network](#) illustrates a RIP network. Three sites are connected with serial lines. The RIP network uses a Class B address and an 8-bit subnet mask. Each site has a contiguous set of network numbers.

Figure: A RIP network




Table: RIP Network Address Assignments lists the network address assignments for the RIP network, including the network number, subnet range, and subnet masks. All interfaces indicate network 130.10.0.0; however, the specific address includes the subnet and subnet mask. For example, serial interface 0 on Router C has an IP address of 130.10.63.3 with a subnet mask of 255.255.255.0.

Table: RIP Network Address Assignments

Network Number	Subnets	Subnet Masks
130.10.0.0	Site A: 8 through 15	255.255.255.0
130.10.0.0	Site B: 16 through 23	255.255.255.0
130.10.0.0	Site C: 24 through 31	255.255.255.0
130.10.0.0	Serial Backbone: 62 through 64	255.255.255.0

Configuration File Examples

The following commands in the configuration file for Router A determine the IP address for each interface and enable RIP on those interfaces:

```
interface serial 0
ip address 130.10.62.1 255.255.255.0
interface serial 1
ip address 130.10.63.1 255.255.255.0
interface ethernet 0
ip address 130.10.8.1 255.255.255.0
interface tokenring 0
ip address 130.10.9.1 255.255.255.0
router rip
network 130.10.0.0
```

The following commands in the configuration file for Router B determine the IP address for each interface and enable RIP on those interfaces:

```
interface serial 0
ip address 130.10.62.2 255.255.255.0
interface serial 1
```

Figure: A RIP network

Internetwork_Design_Guide_--_RIP_and_OSPF_Redistribution

```
ip address 130.10.64.2 255.255.255.0
interface ethernet 0
ip address 130.10.17.2 255.255.255.0
interface tokenring 0
ip address 130.10.16.2 255.255.255.0
router rip
network 130.10.0.0
```


The following commands in the configuration file for Router C determine the IP address for each interface and enable RIP on those interfaces:

```
interface serial 0
ip address 130.10.63.3 255.255.255.0
interface serial 1
ip address 130.10.64.3 255.255.255.0
interface ethernet 0
ip address 130.10.24.3 255.255.255.0
router rip
network 130.10.0.0
```

Adding OSPF to the Center of a RIP Network

A common first step in converting a RIP network to OSPF is to add backbone routers that run both RIP and OSPF, while the remaining network devices run RIP. These backbone routers are OSPF autonomous system boundary routers. Each autonomous system boundary router controls the flow of routing information between OSPF and RIP. In [Figure: RIP network with OSPF at the center](#), Router A is configured as the autonomous system boundary router.

Figure: RIP network with OSPF at the center



RIP does not need to run between the backbone routers; therefore, RIP is suppressed on Router A with the following commands:

```
router rip
passive-interface serial 0
passive-interface serial 1
```

The RIP routes are redistributed into OSPF by all three routers with the following commands:

Internetwork_Design_Guide_--_RIP_and_OSPF_Redistribution

```
router ospf 109
redistribute rip subnets
```

The `subnets` keyword tells OSPF to redistribute all subnet routes. Without the `subnets` keyword, only networks that are not subnetted will be redistributed by OSPF. Redistributed routes appear as external type 2 routes in OSPF. Each RIP domain receives information about networks in other RIP domains and in the OSPF backbone area from the following commands that redistribute OSPF routes into RIP:

```
router rip
redistribute ospf 109 match internal external 1 external 2
default-metric 10
```

The **redistribute** command uses the `ospf` keyword to specify that OSPF routes are to be redistributed into RIP. The keyword **internal** indicates the OSPF intra-area and interarea routes: External 1 is the external route type 1, and external 2 is the external route type 2. Because the command in the example uses the default behavior, these keywords may not appear when you use the **write terminal** or **show configuration** commands.

Because metrics for different protocols cannot be directly compared, you must specify the default metric in order to designate the cost of the redistributed route used in RIP updates. All routes that are redistributed will use the default metric.

In [Figure: RIP network with OSPF at the center](#), there are no paths directly connecting the RIP clouds. However, in typical networks, these paths, or "back doors," frequently exist, allowing the potential for feedback loops. You can use access lists to determine the routes that are advertised and accepted by each router. For example, access list 11 in the configuration file for Router A allows OSPF to redistribute information learned from RIP only for networks 130.10.8.0 through 130.10.15.0:

```
router ospf 109
redistribute rip subnet
distribute-list 11 out rip
access-list 11 permit 130.10.8.0 0.0.7.255
access-list 11 deny 0.0.0.0 255.255.255.255
```

These commands prevent Router A from advertising networks in other RIP domains onto the OSPF backbone, thereby preventing other boundary routers from using false information and forming a loop.

Configuration File Examples

The full configuration for Router A follows:

```
interface serial 0
ip address 130.10.62.1 255.255.255.0
interface serial 1
ip address 130.10.63.1 255.255.255.0
interface ethernet 0
ip address 130.10.8.1 255.255.255.0
interface tokenring 0
ip address 130.10.9.1 255.255.255.0
!
router rip
default-metric 10
network 130.10.0.0
passive-interface serial 0
passive-interface serial 1
redistribute ospf 109 match internal external 1 external 2
!
router ospf 109
```

Figure: RIP network with OSPF at the center

Internetwork_Design_Guide_--_RIP_and_OSPF_Redistribution

```
network 130.10.62.0 0.0.0.255 area 0
network 130.10.63.0 0.0.0.255 area 0
redistribute rip subnets
distribute-list 11 out rip
!
access-list 11 permit 130.10.8.0 0.0.7.255
access-list 11 deny 0.0.0.0 255.255.255.255
```

The full configuration for Router B follows:

```
interface serial 0
ip address 130.10.62.2 255.255.255.0
interface serial 1
ip address 130.10.64.2 255.255.255.0
interface ethernet 0
ip address 130.10.17.2 255.255.255.0
interface tokenring 0
ip address 130.10.16.2 255.255.255.0
!
router rip
default-metric 10
network 130.10.0.0
passive-interface serial 0
passive-interface serial 1
redistribute ospf 109 match internal external 1 external 2
!
router ospf 109
network 130.10.62.0 0.0.0.255 area 0
network 130.10.64.0 0.0.0.255 area 0
redistribute rip subnets
distribute-list 11 out rip
access-list 11 permit 130.10.16.0 0.0.7.255
access-list 11 deny 0.0.0.0 255.255.255.255
```

The full configuration for Router C follows:

```
interface serial 0
ip address 130.10.63.3 255.255.255.0
interface serial 1
ip address 130.10.64.3 255.255.255.0
interface ethernet 0
ip address 130.10.24.3 255.255.255.0
!
router rip
default-metric 10
!
network 130.10.0.0
passive-interface serial 0
passive-interface serial 1
redistribute ospf 109 match internal external 1 external 2
!
router ospf 109
network 130.10.63.0 0.0.0.255 area 0
network 130.10.64.0 0.0.0.255 area 0
redistribute rip subnets
distribute-list 11 out rip
access-list 11 permit 130.10.24.0 0.0.7.255
access-list 11 deny 0.0.0.0 255.255.255.255
```

Adding OSPF Areas

Figure: Configuring route summarization between OSPF areas illustrates how each of the RIP clouds can be converted into an OSPF area. All three routers are area border routers. Area border routers control network information distribution between OSPF areas and the OSPF backbone. Each router keeps a detailed record of the topology of its area and receives summarized information from the other area border routers on their respective areas.

Figure: Configuring route summarization between OSPF areas




Figure: Configuring route summarization between OSPF areas also illustrates variable-length subnet masks (VLSMs). VLSMs use different size network masks in different parts of the network for the same network number. VLSM conserves address space by using a longer mask in portions of the network that have fewer hosts. Table: OSPF Address Assignments lists the network address assignments for the network, including the network number, subnet range, and subnet masks. All interfaces indicate network 130.10.0.0.

Table: OSPF Address Assignments

Network Number	Subnets	Subnet Masks
130.10.0.0	Area 0: 62 through 64	255.255.255.248
130.10.0.0	Area 1: 8 through 15	255.255.255.0
130.10.0.0	Area 2: 16 through 23	255.255.255.0
130.10.0.0	Area 3: 24 through 31	255.255.255.0

To conserve address space, a mask of 255.255.255.248 is used for all the serial lines in area 0. If an area contains a contiguous range of network numbers, an area border router uses the **range** keyword with the **area** command to summarize the routes that are injected into the backbone:

```
router ospf 109
network 130.10.8.0 0.0.7.255 area 1
area 1 range 130.10.8.0 255.255.248.0
```

Internetwork_Design_Guide_--_RIP_and_OSPF_Redistribution

These commands allow Router A to advertise one route, 130.10.8.0 255.255.248.0, which covers all subnets in Area 1 into Area 0. Without the **range** keyword in the **area** command, Router A would advertise each subnet individually; for example, one route for 130.10.8.0 255.255.255.0, one route for 130.10.9.0 255.255.255.0, and so forth.

Because Router A no longer needs to redistribute RIP routes, the **router rip** command can now be removed from the configuration file; however, it is common in some environments for hosts to use RIP to discover routers. When RIP is removed from the routers, the hosts must use an alternative technique to find the routers. Cisco routers support the following alternatives to RIP:

- **ICMP Router Discovery Protocol (IRDP)**-This technique is illustrated in the example at the end of this section. IRDP is the recommended method for discovering routers. The **ip irdp** command enables IRDP on the router. Hosts must also run IRDP.
- **Proxy Address Resolution Protocol (ARP)**-If the router receives an ARP request for a host that is not on the same network as the ARP request sender, and if the router has the best route to that host, the router sends an ARP reply packet giving the router's own local data link address. The host that sent the ARP request then sends its packets to the router, which forwards them to the intended host. Proxy ARP is enabled on routers by default. Proxy ARP is transparent to hosts.

Configuration File Examples

The full configuration for Router A follows:

```
interface serial 0
ip address 130.10.62.1 255.255.255.248
interface serial 1
ip address 130.10.63.1 255.255.255.248
interface ethernet 0
ip address 130.10.8.1 255.255.255.0
ip irdp
interface tokenring 0
ip address 130.10.9.1 255.255.255.0
ip irdp
router ospf 109
network 130.10.62.0 0.0.0.255 area 0
network 130.10.63.0 0.0.0.255 area 0
network 130.10.8.0 0.0.7.255 area 1
area 1 range 130.10.8.0 255.255.248.0
```

The full configuration for Router B follows:

```
interface serial 0
ip address 130.10.62.2 255.255.255.248
interface serial 1
ip address 130.10.64.2 255.255.255.248
interface ethernet 0
ip address 130.10.17.2 255.255.255.0
ip irdp
interface tokenring 0
ip address 130.10.16.2 255.255.255.0
ip irdp
router ospf 109
network 130.10.62.0 0.0.0.255 area 0
network 130.10.64.0 0.0.0.255 area 0
network 130.10.16.0 0.0.7.255 area 2
area 2 range 130.10.16.0 255.255.248.0
```

The full configuration for Router C follows:


Table: OSPF Address Assignments

```
interface serial 0
ip address 130.10.63.2 255.255.255.248
interface serial 1
ip address 130.10.64.2 255.255.255.248
interface ethernet 0
ip address 130.10.24.3 255.255.255.0
ip irdp
router ospf 109
network 130.10.63.0 0.0.0.255 area 0
network 130.10.64.0 0.0.0.255 area 0
network 130.10.24.0 0.0.0.255 area 3
area 3 range 130.10.24.0 255.255.248.0
```

Setting Up Mutual Redistribution

It is sometimes necessary to accommodate more complex network topologies such as independent RIP and OSPF clouds that must perform mutual redistribution. In this scenario, it is critically important to prevent potential routing loops by filtering routes. The router in [Figure: Mutual redistribution between RIP and OSPF networks](#) is running both OSPF and RIP.

Figure: Mutual redistribution between RIP and OSPF networks



With the following commands, OSPF routes will be redistributed into RIP. You must specify the default metric to designate the cost of the redistributed route in RIP updates. All routes redistributed into RIP will have this default metric.

```
! passive interface subcommand from previous example is left out for clarity!
router rip
default-metric 10
network 130.10.0.0
redistribute ospf 109
```

It is a good practice to strictly control which routes are advertised when redistribution is configured. In the following example, a **distribute-list out** command causes RIP to ignore routes coming from the OSPF that originated from the RIP domain.

```
router rip
distribute-list 10 out ospf 109
!
access-list 10 deny 130.10.8.0 0.0.7.255
access-list 10 permit 0.0.0.0 255.255.255.255
```

Router A

The full configuration for the router follows:

```
interface serial 0
ip add 130.10.62.1 255.255.255.0
!
interface serial 1
ip add 130.10.63.1 255.255.255.0
!
interface ethernet 0
```


Internetwork_Design_Guide_-_RIP_and_OSPF_Redistribution

```
ip add 130.10.8.1 255.255.255.0
!
interface tokenring 0
ip add 130.10.9.1 255.255.255.0
!
router rip
default-metric 10
network 130.10.0.0
passive-interface serial 0
passive-interface serial 1
redistribute ospf 109
distribute-list 10 out ospf 109
!
router ospf 109
network 130.10.62.0 0.0.0.255 area 0
network 130.10.63.0 0.0.0.255 area 0
redistribute rip subnets
distribute-list 11 out rip
!
access-list 10 deny 130.10.8.0 0.0.7.255
access-list 10 permit 0.0.0.0 255.255.255.255
access-list 11 permit 130.10.8.0 0.0.7.255
access-list 11 deny 0.0.0.0 255.255.255.255
```

Summary

Because it is common for OSPF and RIP to be used together, it is important to use the practices described here in order to provide functionality for both protocols on an internetwork. You can configure autonomous system boundary routers that run both RIP and OSPF and redistribute RIP routes into the OSPF and vice versa. You can also create OSPF areas using area border routers that provide route summarizations. Use VLSM to conserve address space.