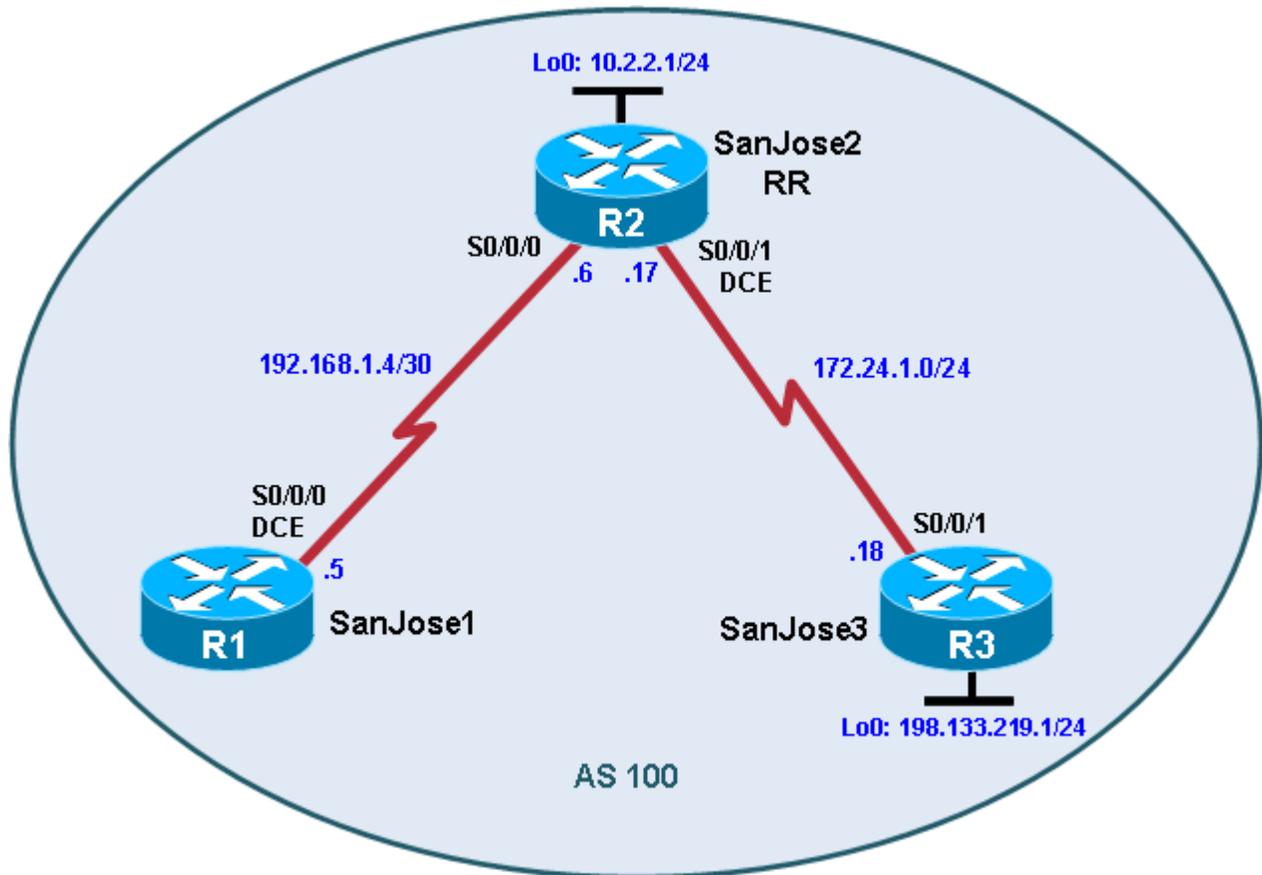


Chapter 6 Lab 6-4, BGP Route Reflectors and Route Filters

Topology



Objectives

- Configure IBGP routers to use a route reflector and a simple route filter.

Background

The International Travel Agency maintains a full-mesh IBGP network that has quickly scaled beyond 100 routers. The company wants to implement route reflectors to work around the full-mesh IBGP requirement. Configure a small cluster and observe how BGP operates in this configuration. Use IP prefix filters to control the updates between IBGP peers.

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 a 2801 or 2811) and Cisco IOS Software versions if they have comparable capabilities and features. Depending on the router or switch 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: Prepare the routers for the lab.

Cable the network as shown in the topology diagram. Erase the startup configuration and reload each router to clear previous configurations. Do not configure Loopback 0 on SanJose3 at this time.

Step 2: Configure the hostname and interface addresses.

You can copy and paste the following configurations into your routers to begin.

Router R1 (hostname SanJose1)

```
hostname SanJose1
!
interface Serial0/0/0
 ip address 192.168.1.5 255.255.255.252
 clock rate 128000
 no shutdown
```

Router R2 (hostname SanJose2)

```
hostname SanJose2
!
interface Loopback0
 ip address 10.2.2.1 255.255.255.0
!
interface Serial0/0/0
 ip address 192.168.1.6 255.255.255.252
 no shutdown
!
interface Serial0/0/1
 ip address 172.24.1.17 255.255.255.0
 clock rate 128000
 no shutdown
```

Router R3 (hostname SanJose3)

```
hostname SanJose3
!
interface Serial0/0/1
 ip address 172.24.1.18 255.255.255.0
 no shutdown
```

Note: Do not configure R3 (SanJose3) with loopback 0 at this time. That will be done in a later step.

Step 3: Configure RIPv2.

- a. Build and configure the network according to the diagram. Use RIPv2 as the IGP. Do not configure the 198.133.219.0 network under the RIP process.

```
SanJose1(config)# router rip
SanJose1(config-router)# version 2
SanJose1(config-router)# no auto-summary
SanJose1(config-router)# network 192.168.1.0
```

```
SanJose2(config)# router rip
SanJose2(config-router)# version 2
```

```
SanJose2(config-router)# no auto-summary
SanJose2(config-router)# network 172.24.0.0
SanJose2(config-router)# network 192.168.1.0
SanJose2(config-router)# network 10.0.0.0
```

```
SanJose3(config)# router rip
SanJose3(config-router)# version 2
SanJose3(config-router)# no auto-summary
SanJose3(config-router)# network 172.24.0.0
```

- b. Issue the **show ip route** command on the routers to verify that each router has a complete routing table.

```
SanJose1# show ip route
```

```
Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static
route
       o - ODR, P - periodic downloaded static route
```

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

```
       172.24.0.0/24 is subnetted, 1 subnets
R       172.24.1.0 [120/1] via 192.168.1.6, 00:00:21, Serial0/0/0
       10.0.0.0/24 is subnetted, 1 subnets
R       10.2.2.0 [120/1] via 192.168.1.6, 00:00:21, Serial0/0/0
       192.168.1.0/30 is subnetted, 1 subnets
C       192.168.1.4 is directly connected, Serial0/0/0
```

- c. Run the following Tcl script on all routers to verify connectivity.

```
SanJose1# tclsh

foreach address {
10.2.2.1
192.168.1.5
192.168.1.6
172.24.1.17
172.24.1.18
} {
ping $address }

```

Step 4: Configure IBGP peers and route reflectors.

In this lab, you will configure a route reflector. By default, a router that receives an EBGP route advertises it to its EBGP and IBGP peers. However, if it receives it through IBGP, it does not advertise it to its IBGP peers, as a loop prevention mechanism. To maintain loop prevention, a route reflector adds two optional, nontransitive BGP attributes to each reflected route, the `ORIGINATOR_ID` and `CLUSTER_LIST`. It uses these attributes in a similar way to `AS_PATH` list to prevent routing loops from occurring. See <http://tools.ietf.org/html/rfc4456> for more information.

However, because of this behavior, the only way for all IBGP routers to receive a route after it is originated into the AS is to have a full mesh of IBGP peers. This can get complex with a large number of peers. A route reflector allows a topology to get around the IBGP limitation of having to have a full mesh. To do this, a route reflector specifies some of its neighbors as route reflector clients. When a route reflector receives an update from a route reflector client, it can pass it on to its other clients. The route reflector would also pass that client-

learned route on to its other non-client peers (both IBGP and EBGP peers). Similarly, a route learned from a non-client peer (again, from either an IBGP or EBGP peer) would be passed on to its client peers. This greatly simplifies configuration because only the route reflector needs to know all the other peers. The clients do not even know that they are clients. To them, it is just a normal IBGP peering relationship. You can even set up multiple route reflectors in a more advanced configuration for redundancy.

- a. Configure the IBGP peers for BGP. Later, you will configure SanJose2 as the route reflector. However, first configure it to peer with both of the other routers.

```
SanJose2(config)# router bgp 100
SanJose2(config-router)# neighbor 192.168.1.5 remote-as 100
SanJose2(config-router)# neighbor 172.24.1.18 remote-as 100
```

After SanJose2 is configured, configure the other two routers as route reflector clients. Remember that to set up clients simply, configure peering between the client and the server. IBGP does not need to be configured in a full mesh.

- b. Issue the following commands on SanJose1:

```
SanJose1(config)# router bgp 100
SanJose1(config-router)# neighbor 192.168.1.6 remote-as 100
```

- c. Issue the following commands on SanJose3:

```
SanJose3(config)# router bgp 100
SanJose3(config-router)# neighbor 172.24.1.17 remote-as 100
```

- d. Use the **show ip bgp neighbors** command to verify that SanJose2 has established a peering relationship with both SanJose1 and SanJose3. Troubleshoot as necessary.

```
SanJose2# show ip bgp neighbors
BGP neighbor is 172.24.1.18, remote AS 100, internal link
  BGP version 4, remote router ID 172.24.1.18
  BGP state = Established, up for 00:02:10
```

<output omitted>

```
BGP neighbor is 192.168.1.5, remote AS 100, internal link
  BGP version 4, remote router ID 192.168.1.5
  BGP state = Established, up for 00:04:15
```

SanJose1 and SanJose3 should not have established a connection. Why?

SanJose1 and SanJose3 were not configured with the appropriate BGP **neighbor** command. As route reflector clients, SanJose1 and SanJose3 do not need to reach an established state.

Step 5: Inject a network into BGP.

- a. To observe the full effect of using a route reflector, configure SanJose3 to inject external routing information into BGP.

```
SanJose3(config)# interface loopback 0
SanJose3(config-if)# ip address 198.133.219.1 255.255.255.0
```

```
SanJose3(config-if)# router bgp 100
SanJose3(config-router)# network 198.133.219.0
```

This configuration forces SanJose3 to inject the external route 198.133.219.0 into BGP. Use the **show ip route** command to check if SanJose2 has picked up this route through BGP. SanJose2 should have a route to 198.133.219.0.

```
SanJose2# show ip route
```

```
Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static
route
       o - ODR, P - periodic downloaded static route
```

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

```
       172.24.0.0/24 is subnetted, 1 subnets
C       172.24.1.0 is directly connected, Serial0/0/1
       10.0.0.0/24 is subnetted, 1 subnets
C       10.2.2.0 is directly connected, Loopback0
B       198.133.219.0/24 [200/0] via 172.24.1.18, 00:01:48
C       10.2.2.0 is directly connected, Loopback0
       192.168.1.0/30 is subnetted, 1 subnets
C       192.168.1.4 is directly connected, Serial0/0/0
```

What is the next hop for this route? Explain.

- b. Verify that you can ping 198.133.219.1 from SanJose2. If not, troubleshoot.
- c. Check the routing table of SanJose1. There should not be a route to 198.133.219.0. Why?
- d. Remember that SanJose1 is not configured to peer with SanJose3. To eliminate the need for a full IBGP mesh, SanJose2 must be configured as a route reflector. Issue the following commands on SanJose2:

```
SanJose2(config)# router bgp 100
```

```
SanJose2(config-router)# neighbor 192.168.1.5 route-reflector-client
```

```
SanJose2(config-router)# neighbor 172.24.1.18 route-reflector-client
```

```
*Mar  9 19:02:27.831: %BGP-5-ADJCHANGE: neighbor 192.168.1.5 Down RR client
conf
```

```
ig change
```

```
*Mar  9 19:02:27.931: %BGP-5-ADJCHANGE: neighbor 172.24.1.18 Down RR client
conf
```

```
ig change
```

```
*Mar  9 19:02:32.387: %BGP-5-ADJCHANGE: neighbor 172.24.1.18 Up
```

```
*Mar  9 19:02:37.507: %BGP-5-ADJCHANGE: neighbor 192.168.1.5 Up
```

- e. Verify that an IBGP cluster was successfully created by issuing the **show ip protocols** command on SanJose2. The output of this command should indicate that SanJose2 is a route reflector.

```
SanJose2# show ip protocols
```

```
Routing Protocol is "rip"
```

```
Outgoing update filter list for all interfaces is not set
```

CCNPv6 ROUTE

```
Incoming update filter list for all interfaces is not set
Sending updates every 30 seconds, next due in 26 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:
  10.0.0.0
  172.24.0.0
  192.168.1.0
Routing Information Sources:
  Gateway            Distance      Last Update
Distance: (default is 120)
```

Routing Protocol is "bgp 100"

```
Outgoing update filter list for all interfaces is not set
Incoming update filter list for all interfaces is not set
Route Reflector for address family IPv4 Unicast, 2 clients
Route Reflector for address family IPv6 Unicast, 2 clients
Route Reflector for address family IPv4 MDT, 2 clients
Route Reflector for address family VPNv4 Unicast, 2 clients
Route Reflector for address family VPNv6 Unicast, 2 clients
Route Reflector for address family IPv4 Multicast, 2 clients
Route Reflector for address family IPv6 Multicast, 2 clients
Route Reflector for address family NSAP Unicast, 2 clients
IGP synchronization is disabled
Automatic route summarization is disabled
Neighbor(s):
  Address            FiltIn FiltOut DistIn DistOut Weight RouteMap
  172.24.1.18
  192.168.1.5
Maximum path: 1
Routing Information Sources:
  Gateway            Distance      Last Update
  172.24.1.18        200           00:01:43
Distance: external 20 internal 200 local 200
```

How many clients does SanJose2 have?

- f. Issue the **show ip protocols** command on SanJose1. The output of this command does not include information about route reflectors. Remember that SanJose1 is a client and not a route reflector server, so it is unaware of route reflection.
- g. Finally, verify that route reflection is working by checking the routing table on SanJose1. SanJose1 will have a route to network 198.133.219.0.

```
SanJose1# show ip route
```

```
Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
```

CCNPv6 ROUTE

i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
ia - IS-IS inter area, * - candidate default, U - per-user static
route
o - ODR, P - periodic downloaded static route

Gateway of last resort is not set

```
172.24.0.0/24 is subnetted, 1 subnets
R    172.24.1.0 [120/1] via 192.168.1.6, 00:00:08, Serial0/0/0
10.0.0.0/24 is subnetted, 1 subnets
R    10.2.2.0 [120/1] via 192.168.1.6, 00:00:08, Serial0/0/0
B   198.133.219.0/24 [200/0] via 172.24.1.18, 00:01:25
192.168.1.0/30 is subnetted, 1 subnets
C    192.168.1.4 is directly connected, Serial0/0/0
```

Is 172.24.1.18 the IP address of the next hop of this route on the SanJose1 table? Explain.

Notice that SanJose1 is not directly connected to the IP network for the next hop. Why?

Hint: From which router did SanJose1 learn the route?

- h. Ping 198.133.219.1 from SanJose1. This ping should be successful.

Notice that SanJose1 pings to R3 198.133.219.1 are successful even though the next-hop address is not on a directly-connected network. For example, the next-hop address could be 192.168.1.6 on R2 if it were not for the behavior of IBGP.

Step 6: Inject a summary address into BGP.

- a. For the purpose of this lab, configure SanJose3 to inject a summary address into BGP.

```
SanJose3(config)# router bgp 100
SanJose3(config-router)# aggregate-address 198.0.0.0 255.0.0.0
```

BGP should now send the supernet route 198.0.0.0/8 to SanJose2 with the attribute ATOMIC_AGGREGATE set.

Note: By default, BGP on Cisco routers advertises both aggregate routes and the individual component routes. If only the aggregate route is to be advertised, use the **aggregate-address network mask summary-only** command.

- b. On SanJose2, issue the following command:

```
SanJose2# show ip bgp 198.0.0.0
BGP routing table entry for 198.0.0.0/8, version 8
Paths: (1 available, best #1, table Default-IP-Routing-Table)
Flag: 0x820
  Advertised to update-groups:
    1
  Local, (aggregated by 100 172.24.1.18), (Received from a RR-client)
    172.24.1.18 from 172.24.1.18 (172.24.1.18)
    Origin IGP, metric 0, localpref 100, valid, internal, atomic-aggregate,
  best
```

According to the output of this command, which address aggregated this route?

What indicates that route reflection is involved in this process?

Is there an indication that the ATOMIC_AGGREGATE attribute has been set?

- c. SanJose2 should, in turn, reflect this route to SanJose1. Check both the routing table and BGP table on SanJose1 to be sure. Both the route to 198.133.219.0 and the supernet route 198.0.0.0 should be installed in the SanJose1 routing table and the BGP table.

```
SanJose1# show ip route
```

```
Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static
route
       o - ODR, P - periodic downloaded static route
```

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

```
       172.24.0.0/24 is subnetted, 1 subnets
R       172.24.1.0 [120/1] via 192.168.1.6, 00:00:20, Serial0/0/0
       10.0.0.0/24 is subnetted, 1 subnets
R       10.2.2.0 [120/1] via 192.168.1.6, 00:00:20, Serial0/0/0
B       198.133.219.0/24 [200/0] via 172.24.1.18, 00:08:34
       192.168.1.0/30 is subnetted, 1 subnets
C       192.168.1.4 is directly connected, Serial0/0/0
B       198.0.0.0/8 [200/0] via 172.24.1.18, 00:04:19
```

The International Travel Agency has decided to filter specific routes to the 198.0.0.0/8 address space. Configure a route filter to prevent SanJose2 from sending the 198.133.219.0/24 route to its other clients, in this case to SanJose1.

- d. Issue the following commands on SanJose2:

```
SanJose2(config)# ip prefix-list SUPERNETONLY permit 198.0.0.0/8
SanJose2(config)# router bgp 100
SanJose2(config-router)# neighbor 192.168.1.5 prefix-list SUPERNETONLY out
```

- e. Return to SanJose1, issue the **clear ip bgp * soft** command, and verify that the prefix list has done its job by issuing a **show ip bgp** command. Troubleshoot as necessary.

Unlike before, where routes to 198.133.219.0 and 198.0.0.0 were present, now only one route to 198.0.0.0 in the routing and BGP tables should be seen. Troubleshoot as necessary.

```
SanJose1# show ip route
```

```
Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
```

CCNPv6 ROUTE

N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
E1 - OSPF external type 1, E2 - OSPF external type 2
i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
ia - IS-IS inter area, * - candidate default, U - per-user static

route

o - ODR, P - periodic downloaded static route

Gateway of last resort is not set

```
172.24.0.0/24 is subnetted, 1 subnets
R   172.24.1.0 [120/1] via 192.168.1.6, 00:00:20, Serial0/0/0
10.0.0.0/24 is subnetted, 1 subnets
R   10.2.2.0 [120/1] via 192.168.1.6, 00:00:20, Serial0/0/0
192.168.1.0/30 is subnetted, 1 subnets
C   192.168.1.4 is directly connected, Serial0/0/0
B   198.0.0.0/8 [200/0] via 172.24.1.18, 00:04:19
```

- f. Run the following Tcl script on all routers to verify full connectivity. All pings should be successful.

```
SanJose1# tclsh
```

```
foreach address {
10.2.2.1
198.133.219.1
192.168.1.5
192.168.1.6
172.24.1.17
172.24.1.18
} {
ping $address }
```

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 try to list all the 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. An example of this is an ISDN BRI interface. The string in parenthesis is the legal abbreviation that can be used in Cisco IOS commands to represent the interface.				