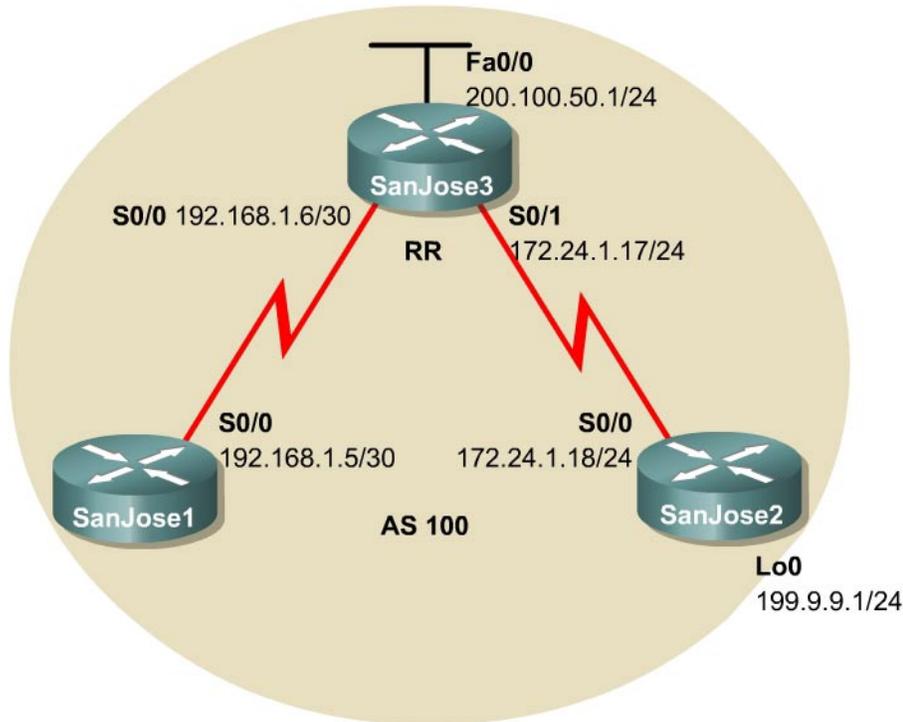


## Lab 9.11.4b BGP Route Reflectors and Route Filters



### Objective

In this lab, the student will configure IBGP routers to use a route reflector and a simple route filter.

### Scenario

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.

### Step 1

Build and configure the network according to the diagram, and use RIP as the IGP. Do not configure the 199.9.9.0 network. Use `ping` to test connectivity among all interfaces. Each router should have a complete routing table.

### Step 2

Configure the IBGP peers for BGP. SanJose3 will later be configured as the route reflector. However, first configure it to peer with both of the other routers, as shown in the following:

```
SanJose3(config)#router bgp 100
SanJose3(config-router)#neighbor 192.168.1.5 remote-as 100
SanJose3(config-router)#neighbor 172.24.1.18 remote-as 100
SanJose3(config-router)#no auto-summary
```

```
SanJose3(config-router)#no synchronization
SanJose3(config-router)#network 200.100.50.0
SanJose3(config-router)#network 172.24.1.0 mask 255.255.255.0
SanJose3(config-router)#network 192.168.1.4 mask 255.255.255.252
```

After SanJose3 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 to a full mesh.

Issue the following commands on SanJose1:

```
SanJose1(config)#router bgp 100
SanJose1(config-router)#neighbor 192.168.1.6 remote-as 100
SanJose1(config-router)#no auto-summary
SanJose1(config-router)#no synchronization
```

Issue the following commands on SanJose2:

```
SanJose2(config)#router bgp 100
SanJose2(config-router)#neighbor 172.24.1.17 remote-as 100
SanJose2(config-router)#no auto-summary
SanJose2(config-router)#no synchronization
```

Verify that SanJose3 has established a peering relationship with both SanJose1 and SanJose2. Troubleshoot, as necessary.

1. SanJose1 and SanJose2 should not have established a connection. Why?

---

---

---

SanJose2 was not configured with the appropriate BGP `neighbor` command. As a route reflector client, SanJose1 will not need to reach an Established state with SanJose2.

### Step 3

To observe the full effect of using a route reflector, configure SanJose2 to inject external routing information into BGP as follows:

```
SanJose2(config)#int lo0
SanJose2(config-if)#ip address 199.9.9.1 255.255.255.0
SanJose2(config)#router bgp 100
SanJose2(config-router)#network 199.9.9.0
```

This configuration forces SanJose2 to inject the external route 199.9.9.0 into BGP. Use the `show ip route` command to check if SanJose3 has picked up this route through BGP. SanJose3 should have a route to 199.9.9.0.

2. Is the next hop for this route 172.24.1.18?

---

199.9.9.1 should `ping` from SanJose3. If not, troubleshoot.

Now check the routing table of SanJose1.

3. There should not be a route to 199.9.9.0. Why?

---

---

Remember that SanJose1 is not configured to peer with SanJose2. To eliminate the need for a full IBGP mesh, SanJose3 must be configured as a route reflector server. Issue the following commands on SanJose3:

```
SanJose3(config)#router bgp 100
SanJose3(config-router)#neighbor 192.168.1.5 route-reflector-client
SanJose3(config-router)#neighbor 172.24.1.18 route-reflector-client
```

Verify that an IBGP cluster was successfully created by issuing the `show ip protocols` command on SanJose3. The output of this command should indicate that SanJose3 is a route reflector.

4. How many clients does SanJose3 have?

---

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.

Finally, verify that route reflection is working by checking the routing table on SanJose1. SanJose1 should, at least, have a route to 199.9.9.0.

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

---

6. 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?

---

---

`Ping` 199.9.9.1 from SanJose1. This ping should be successful.

Notice that SanJose1 pings 199.9.9.1 even though the next hop address is not on a directly connected network. For example, the next hop address could be 172.24.1.18.

## Step 4

For the purposes of this lab, configure SanJose2 to inject a summary address into BGP, as shown in the following:

```
SanJose2(config)#router bgp 100
SanJose2(config-router)#aggregate-address 199.0.0.0 255.0.0.0
```

BGP should now send the supernet route, 199.0.0.0/8, to SanJose3 with the `ATOMIC_AGGREGATE` attribute set.

On SanJose3, issue the following command:

```
SanJose3#show ip bgp 199.0.0.0
BGP routing table entry for 199.0.0.0/8, version 6
Paths: (1 available, best #1)
Bestpath transition flag: 0x208
      Advertised to non peer-group peers:
        192.168.1.5
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, localpref 100, valid, internal, atomic-aggregate, best,
  ref 2
```

7. According to the output of this command, what address aggregated this route?

---

8. What indicates that route reflection is involved in this process?

---

9. Is there an indication that the ATOMIC\_AGGREGATE attribute has been set?

---

SanJose3 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 199.9.9.0 and the supernet route, 199.0.0.0, should be installed in both the SanJose1 routing table and the BGP table.

The International Travel Agency has decided to filter specific routes to the 199.0.0.0/8 address space. Configure a route filter to prevent SanJose3 from sending the 199.9.9.0/24 route to its other clients, in this case to SanJose1.

Issue the following commands on SanJose3:

```
SanJose3(config)#ip prefix-list supernetonly permit 199.0.0.0/8
SanJose3(config)#ip prefix-list supernetonly permit 172.24.1.0/24
SanJose3(config)#ip prefix-list supernetonly permit 200.100.50.0/24
SanJose3(config)#router bgp 100
SanJose3(config-router)#neighbor 192.168.1.5 prefix-list supernetonly out
```

Return to SanJose1, issue the `clear ip bgp *` 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 199.9.9.0 and 199.0.0.0 were present, now only one route to 199.0.0.0 in the routing and BGP tables should be seen. Troubleshoot as necessary.