

## Module 7 – BGP Route Reflector Lab

**Objective:** To implement BGP route reflectors as an alternative to fully-meshed iBGP.

**Prerequisites:** Module 1, the Scaling BGP presentation and Module 6

Topology:

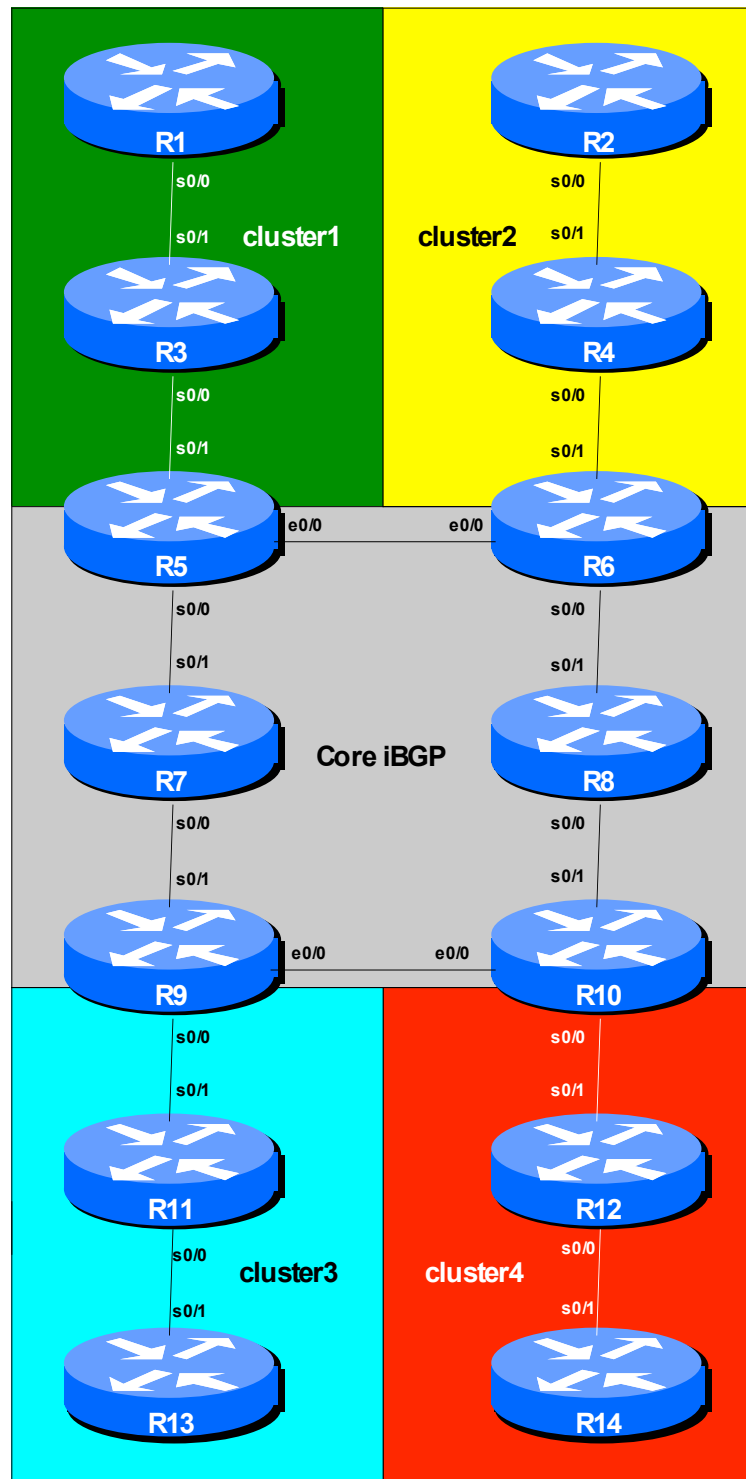


Figure 1 – Route Reflector Clusters

## Lab Notes

A pre-requisite for this Module is the Advanced OSPF Module which includes a description of OSPF areas. The design of this route-reflector network assumes the network layout and configuration used in the Advanced OSPF Module. While it will work without the OSPF area configuration, this is a good exercise in explaining by example how many ISPs overlay route reflectors on their IGP.

Before starting, decide which routers in the network will represent the core network, and which will represent other portions of the backbone. The example given in Figure 1 allows the student to study the best combination of route reflector clusters.

Routers 5 to 10 represent the “core network” running a fully meshed iBGP, with Routers 5, 6, 9 and 10 being route reflectors. In a typical ISP backbone, these routers would carry all the routes known in the ISPs network, and possibly all the Internet routes too. Routers 1, 3 and 5 represent a cluster – say the “distribution network”, carrying out the function of aggregating customer connections. Routers 2, 4 and 6 form another cluster. Routers 9, 11 and 13 form another cluster. And Routers 10, 12 and 14 form a cluster.

## Lab Exercise

1. **Reset the configuration** so that all routers are in the same AS (AS 10 is used in this example). Also all routers should use the same OSPF process ID (41) and should be as configured in the Advanced OSPF Module.
2. **Configure full mesh iBGP in the core network.** The core router should be a route reflector for the distribution router in the same cluster. Core routers should have a fully meshed iBGP among themselves. The routers in a cluster do not require fully meshed iBGP among themselves – they should only peer with the route reflector.

### Example for Router 5 using peer groups (recommended):

```
router bgp 10
  bgp log-neighbor-changes
  no synchronization
  network 100.2.16.0 mask 255.255.240.0
  neighbor CORE-IBGP peer-group
  neighbor CORE-IBGP description iBGP for the CORE routers
  neighbor CORE-IBGP remote-as 10
  neighbor CORE-IBGP password cisco
  neighbor CORE-IBGP update-source loopback0
  neighbor 100.2.47.224 peer-group CORE-IBGP      ! ibgp with Router 6
  neighbor 100.2.63.224 peer-group CORE-IBGP      ! ibgp with Router 7
  neighbor 100.3.15.224 peer-group CORE-IBGP      ! ibgp with Router 8
  neighbor 100.3.31.224 peer-group CORE-IBGP      ! ibgp with Router 9
  neighbor 100.3.63.224 peer-group CORE-IBGP      ! ibgp with Router 10
  . . .
```

3. **Configuring route-reflector-client peers:** On the routers that will be router-reflectors (Routers 5, 6, 9, and 10), configure peers inside the cluster as **route-reflector-clients**. Each router team should announce the /20 network (or /19 in the case of Router 3 and Router 10) which has been assigned to them.

**Example for Router 5 using peer groups (recommended):**

```

router bgp 10
  no synchronization
  bgp log-neighbor-changes
  network 100.2.16.0 mask 255.255.240.0
  neighbor RR-CLIENT peer-group
  neighbor RR-CLIENT remote-as 10
  neighbor RR-CLIENT description iBGP route reflector client configuration
  neighbor RR-CLIENT password cisco
  neighbor RR-CLIENT update-source loopback0
  neighbor RR-CLIENT route-reflector-client
  neighbor 100.1.15.224 peer-group RR-CLIENT      ! ibgp with Router 1
  neighbor 100.1.63.224 peer-group RR-CLIENT      ! ibgp with Router 3
  . . .

```

Note that the clients don't require an iBGP peering between each other – the route reflector announces one client's network to all other clients.

**Aside:** If a full-mesh amongst clients is required (usually due to operational conditions), then the router requires the *no bgp-client-to-client reflection* bgp command to be configured.

**Q.** What does the *no bgp-client-to-client reflection* bgp command do?

**A.** By default, a router reflector will reflect all routes it hears to its clients, even the routes it hears from its clients. This command turns off the functionality so that the route reflectors don't announce client routes back to the members of the cluster.

**4. Route-reflector clients should configure iBGP peering to the router reflector inside the cluster.**

**Example for Router 11 using peer groups (recommended):**

```

router bgp 10
  no synchronization
  bgp log-neighbor-changes
  network 100.4.0.0 mask 255.255.240.0
  neighbor RR peer-group
  neighbor RR remote-as 10
  neighbor RR description iBGP configuration for RR clients
  neighbor RR password cisco
  neighbor RR update-source loopback0
  neighbor 100.3.31.224 peer-group RR
  . . .

```

**5. Use “*show ip bgp <address>*” to see how reflected prefixes show up on the clients. How do you explain the path choices which you see?**

**Checkpoint #2:** Call the lab instructors and show the function of your router. You should have peerings with all the routers in your cluster, and any peers/clients. You should also demonstrate the output from “*show ip route*” so that you can see which routes you are hearing from which routers.

## ***CONFIGURATION NOTES***

Documentation is critical! You should record the configuration at each ***Checkpoint***, as well as the configuration at the end of the module.