Module 3 – BGP Route Reflector Lab

Objective: To migrate the network to BGP route reflectors as an alternative to using a fully-meshed iBGP.

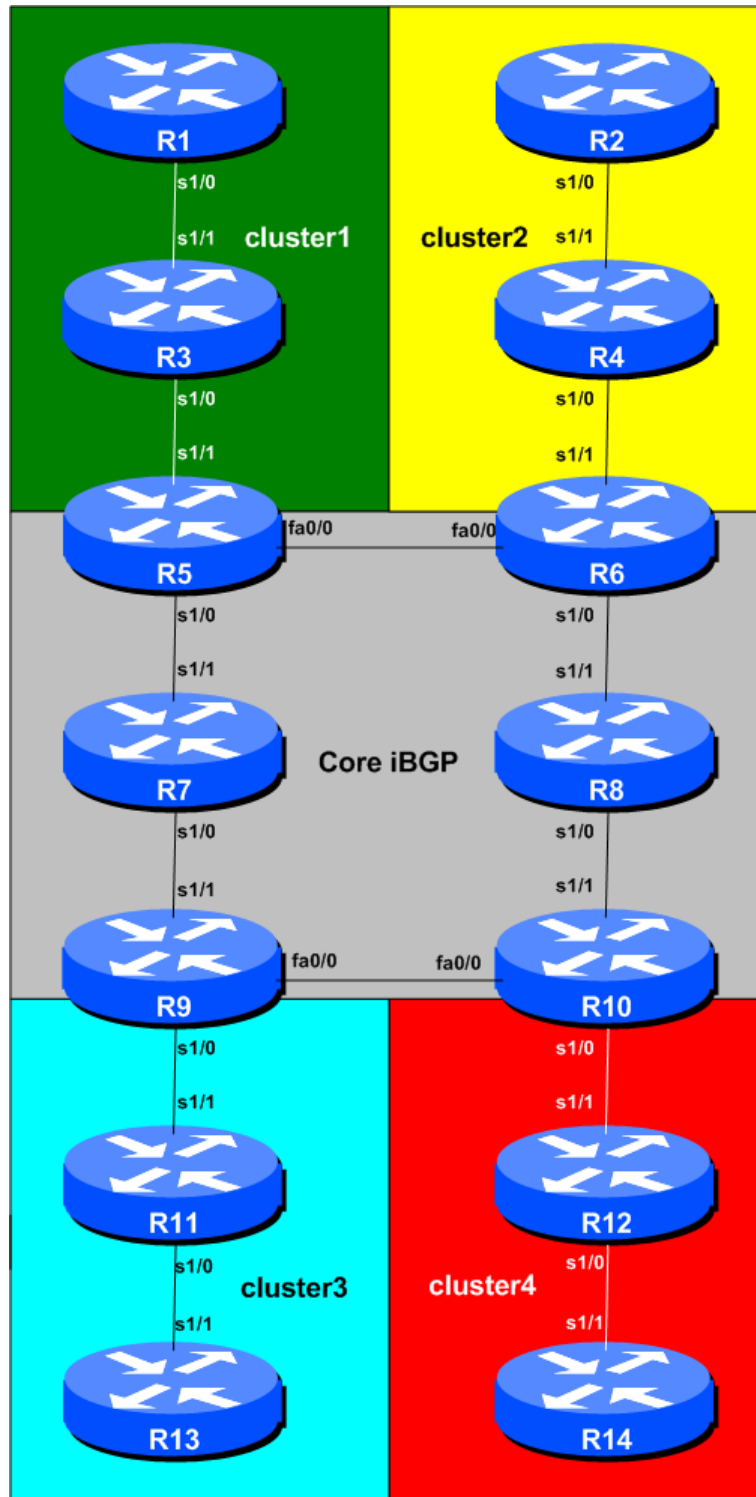**Prerequisites: Module 1, the Scaling BGP presentation and (optionally) Module 2**

Topology:

Figure 1 – Route Reflector Clusters

Lab Notes

While not a pre-requisite, it is helpful if the OSPF Areas Module is completed before this Module. The design of this route-reflector network assumes the network layout and configuration used in the OSPF Area Module. While it will work without the OSPF area configuration (for example, if ISIS is chosen as the IGP instead), this is a good exercise in explaining by example how many ISPs overlay route reflectors on their IGP.

The topology shown 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. **Retain the configuration** used in the previous Module. This means that all routers are using the same IGP (OSPF or ISIS). Before we start with this module, we need to ensure that the BGP configuration is exactly as it was coming from Module 1 – ensure that all the iBGP sessions are up and running.

2. **Migrating from full mesh iBGP to Route Reflector based iBGP.** The process for migration is relatively straightforward, as long as simple rules are adhered to. The key is to configure a client to be a route reflector client before removing the redundant iBGP sessions. This ensures that the client iBGP sessions are up and running and passing on prefixes, meaning that removing the surplus sessions should not cause any break in iBGP routing information. The next steps in the module work through the process, and each Router Team affected should follow them carefully. Teamwork is important – with good teamwork it is more than likely that the network can be migrated without any major connectivity impact.

3. **Configure the iBGP peerings with the Route Reflector Clients.** Routers 1 to 4 and Routers 11 to 14 will become route reflector clients. So the teams operating Routers 5, 6, 9 and 10 should reconfigure their peerings with Routers 1 and 3, 2 and 4, 11 and 13, and 12 and 14 respectively to be route-reflector clients. To configure a neighbour as a client, we simply add in the **route-reflector-client** BGP neighbour command as per the following example:

   Example for Router 5:

   ```
   router bgp 10
    no synchronization
    bgp log-neighbor-changes
    network 10.0.0.0 mask 255.255.240.0
    neighbor 10.0.15.241 remote-as 10
    neighbor 10.0.15.241 description iBGP with Router1
    neighbor 10.0.15.241 update-source loopback0
    neighbor 10.0.15.241 route-reflector-client
   ```

```
    neighbor 10.0.15.243 remote-as 10
    neighbor 10.0.15.243 description iBGP with Router3
    neighbor 10.0.15.243 update-source loopback0
    neighbor 10.0.15.243 route-reflector-client
```

Remember that if you have multiple address families configured for iBGP, you will need to enter the IPv4 address family to add these extra lines (if this workshop is done in conjunction with the IPv6 materials). The configuration to do this would be:

```
    router bgp 10
     address-family ipv4
      neighbor 10.0.15.241 ...
```

and then the rest of the configuration above would apply.

Checkpoint #1: *Once the teams operating Router 5, 6, 9 and 10 have completed adding the two route reflector client lines, they should call the lab instructors and show the function of their router. The lab instructors will then tell the class to proceed with remaining steps in this module.*

STOP AND WAIT HERE

4. **Removing redundant iBGP sessions – Step 1.** Once the teams operating Routers 5, 6, 9 and 10 have **ALL** completed step 3, they should remove the redundant iBGP sessions with route reflector clients which are not in their cluster.

    This means that:
    - Router5 removes iBGP sessions with Routers 2, 4, 11, 12, 13 and 14
    - Router6 removes iBGP sessions with Routers 1, 3, 11, 12, 13 and 14
    - Router9 removes iBGP sessions with Routers 1, 2, 3, 4, 12 and 14
    - Router10 removes iBGP sessions with Routers 1, 2, 3, 4, 11 and 13

    Example for Router 9:

    ```
    router bgp 10
     no neighbor 10.0.15.241
     no neighbor 10.0.15.242
     no neighbor 10.0.15.243
     no neighbor 10.0.15.244
     no neighbor 10.0.15.252
     no neighbor 10.0.15.254
    ```

5. **Remove redundant iBGP sessions – Step 2.** The clients of the route reflectors can also remove their redundant iBGP sessions. Each client should have only ONE iBGP session – the other 12 set up in Module 1 can now be removed.

    Example for Router 1:

    ```
    router bgp 10
     no neighbor 10.0.15.242
     no neighbor 10.0.15.243
     no neighbor 10.0.15.244
    ```

```
    no neighbor 10.0.15.246
    no neighbor 10.0.15.247
    no neighbor 10.0.15.248
    no neighbor 10.0.15.249
    no neighbor 10.0.15.250
    no neighbor 10.0.15.251
    no neighbor 10.0.15.252
    no neighbbor 10.0.15.253
    no neighbor 10.0.15.254
```

Router 1 only retains its iBGP session with Router5 – every other one should be removed as in the example above. The other teams operating route reflector clients (Routers 2, 3, 4, 11, 12, 13 and 14) should do the same thing.

6. **Remove redundant iBGP sessions – Step 3.** Finally Router7 and Router8, the two core routers which are not route reflectors, can remove their redundant iBGP sessions – they no longer need to run iBGP with any of the 8 clients (Routers 1 to 4, 11 to 14). They still need to run iBGP with other core routers though, as they are part of the core iBGP mesh.

   Example for Router 8:

```
router bgp 10
 no neighbor 10.0.15.241
 no neighbor 10.0.15.242
 no neighbor 10.0.15.243
 no neighbor 10.0.15.244
 no neighbor 10.0.15.251
 no neighbor 10.0.15.252
 no neighbor 10.0.15.253
 no neighbor 10.0.15.254
```

7. **Final Check.** The network has now been migrated from full mesh IBGP with each router having 13 neighbours, to a Route Reflector set up, with core routers being route reflectors, and edge routers being clients. 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.*