

Module 3 – IPv6 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 (optionally) Module 2

Topology:

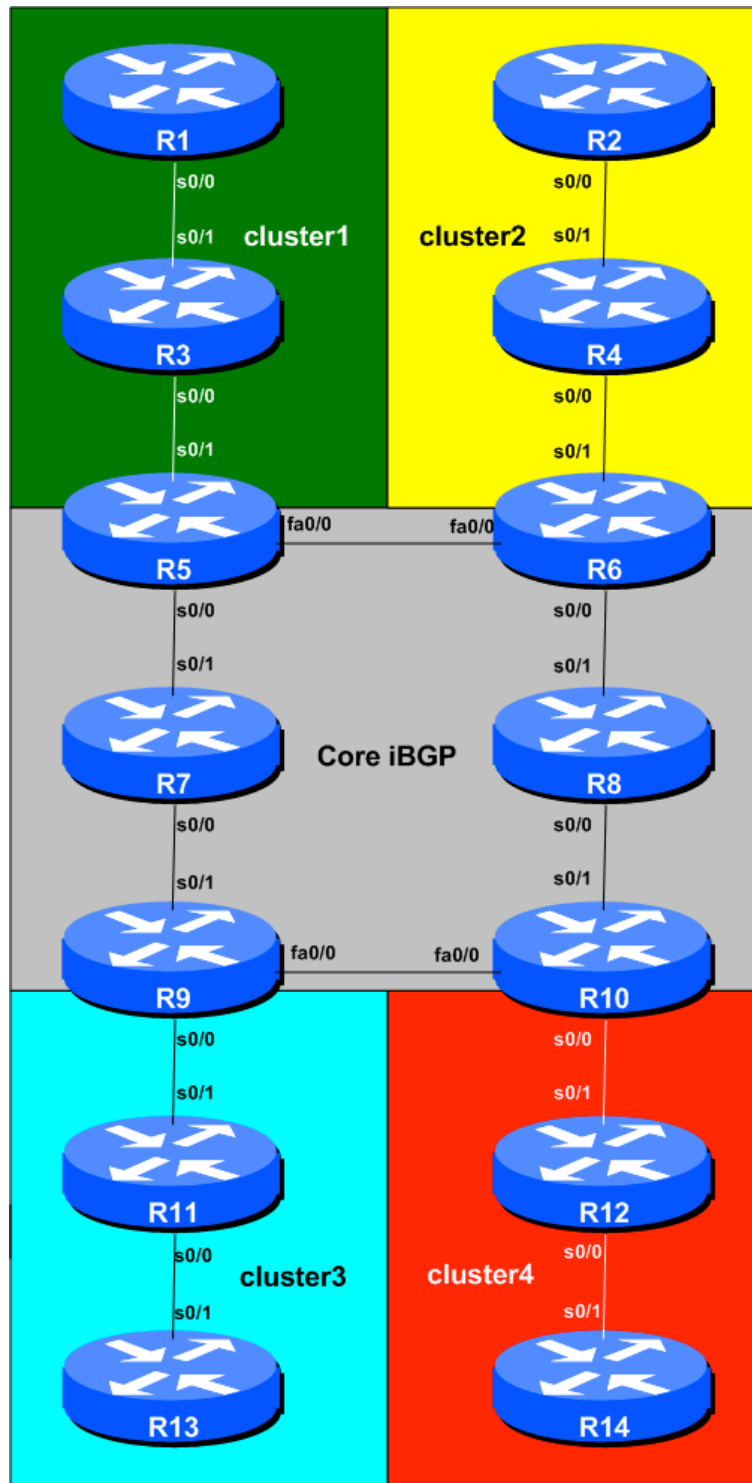


Figure 1 – Route Reflector Clusters

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, 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. **Retain the configuration** used in the previous Module. This means that all routers are using the same OSPF process ID (41). Before we start with this module, we need to remove the BGP configuration if we came from Module 1 – the options are to either completely delete BGP by doing

```
no router bgp 10
```

and starting from scratch, or using the more realistic alternative as it shows how to migrate from a full mesh iBGP to a route-reflector lay out. If the lab just completed was the IPv4 BGP route reflector module, then don't delete the BGP configuration – we are simply going to add IPv6 in as well.

2. **Configure full mesh iBGP in the core network.** Core routers should have a fully meshed iBGP among themselves. The teams operating the routers in the core of the network should now set up the full mesh iBGP within the core.

Example for Router 5:

```
router bgp 10
no bgp default ipv4-unicast
address-family ipv6
  bgp log-neighbor-changes
  no synchronization
  network 2001:db8::/32
  neighbor 2001:db8::6 remote-as 10
  neighbor 2001:db8::6 description iBGP with Router 6
  neighbor 2001:db8::6 update-source loopback0
  neighbor 2001:db8::7 remote-as 10
  neighbor 2001:db8::7 description iBGP with Router 7
  neighbor 2001:db8::7 update-source loopback0
```

```

neighbor 2001:db8::8 remote-as 10
neighbor 2001:db8::8 description iBGP with Router 8
neighbor 2001:db8::8 update-source loopback0
neighbor 2001:db8::9 remote-as 10
neighbor 2001:db8::9 description iBGP with Router 9
neighbor 2001:db8::9 update-source loopback0
neighbor 2001:db8::a remote-as 10
neighbor 2001:db8::a description iBGP with Router 10
neighbor 2001:db8::a update-source loopback0
. . .

```

- 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 still announce the /32 network which has been assigned to the AS.

Example for Router 5:

```

router bgp 10
no bgp default ipv4-unicast
address-family ipv6
no synchronization
bgp log-neighbor-changes
network 2001:db8::/32
neighbor 2001:db8::1 remote-as 10
neighbor 2001:db8::1 description iBGP with Router1
neighbor 2001:db8::1 update-source loopback0
neighbor 2001:db8::1 route-reflector-client
neighbor 2001:db8::3 remote-as 10
neighbor 2001:db8::3 description iBGP with Router3
neighbor 2001:db8::3 update-source loopback0
neighbor 2001:db8::3 route-reflector-client
. . .

```

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:

```

router bgp 10
no bgp default ipv4-unicast
address-family ipv6
no synchronization
bgp log-neighbor-changes
network 2001:db8::/32

```

Friday, July 17, 2009

```
neighbor 2001:db8::9 remote-as 10
neighbor 2001:db8::9 description iBGP with Router9
neighbor 2001:db8::9 update-source loopback0
. . .
```

5. Use “***show bgp ipv6 unicast*** <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.