iBGP - IPv6

- Objective: Add IPv6 to the iBGP setup
- Ensure that all routers, interfaces, cables and connections are working properly.
- Prerequisites: IS-IS, iBGP - IPv4.

1. Create Loopback Interfaces.

Our network setup is very simple. There are only two routers in our **ISP** network but in a real ISP network we'd expect to see more. In an IPv6 network we'd allocate a /64 to provide a loopback address for each router. We each have 2 routers in our lab so this allows us to grow our network without any restrictions. The assigned loopback addresses are:

Router	Loopback address	Router	Loopback address
r1-isp1-ixp1	FD90:100:66:22::251/128	r1-isp1-ixp2	FD90:100:66:182::251/128
r2-isp1-ixp1	FD90:100:66:22::252/128	r2-isp1-ixp2	FD90:100:66:182::252/128
r1-isp2-ixp1	FD90:100:66:42::251/128	r1-isp2-ixp2	FD90:100:66:202::251/128
r2-isp2-ixp1	FD90:100:66:42::252/128	r2-isp2-ixp2	FD90:100:66:202::252/128
r1-isp3-ixp1	FD90:100:66:62::251/128	r1-isp3-ixp2	FD90:100:66:222::251/128
r2-isp3-ixp1	FD90:100:66:62::252/128	r2-isp3-ixp2	FD90:100:66:222::252/128
r1-isp4-ixp1	FD90:100:66:82::251/128	r1-isp4-ixp2	FD90:100:66:242::251/128
r2-isp4-ixp1	FD90:100:66:82::252/128	r2-isp4-ixp2	FD90:100:66:242::252/128
r1-isp5-ixp1	FD90:100:66:102::251/128	r1-isp5-ixp2	FD90:100:66:262::251/128
r2-isp5-ixp1	FD90:100:66:102::252/128	r1-isp5-ixp2	FD90:100:66:262::251/128
r1-isp6-ixp1	FD90:100:66:122::251/128	r1-isp5-ixp2	FD90:100:66:262::251/128
r2-isp6-ixp1	FD90:100:66:122::252/128	r2-isp6-ixp2	FD90:100:66:282::252/128
r1-isp7-ixp1	FD90:100:66:142::251/128	r1-isp7-ixp2	FD90:100:66:1282::251/128
r2-isp7-ixp1	FD90:100:66:142::252/128	r2-isp7-ixp2	FD90:100:66:1282::252/128

For example, ISP Team 1 on IXP 1 would assign the following address and mask to the loopback on r1-isp1-ixp1:

```
r1-isp1-ixp1(config)#interface loopback 0
r1-isp1-ixp1(config-if)#ipv6 address FD90:100:66:22::251/128
```

Q: Why do we use /128 masks for the loopback interface address?

A: There is no physical network attached to the loopback so there can only be one device there. So we only need to assign a /128 mask – it is a waste of address space to use anything else.

Activating Multi-Topology ISIS

We also need to activate multi-topology ISIS to roll out ISIS support for IPv6 if the existing network is already using IPv4 ISIS. This allows the IPv6 topology to be incrementally rolled out, very useful during

deployment of IPv6. This means that each team can add ISIS IPv6 support without having to coordinate with their neighbouring teams.

```
Router1(config)# router isis workshop
Router1(config-router)# address-family ipv6
Router1(config-router-af)# multi-topology
```

NB. If we do not enable multi-topology, then each team will have to coordinate the enabling of IPv6 ISIS on each interface with their respective neighbouring teams. Failure to do so will result in the ISIS session going down, as there will be a topology mismatch on that interface.

NB. Multi-topology is configurable in IOS 12.3 and 12.4 but it does not work due to a bug which Cisco refuses to fix. The workaround is to use single topology, noting the caveat above, or use 12.2SRE, 12.2SXH, 12.4T, 15.0 or later IOS images.

2. Activating ISIS on each interface.

The ISIS process is already configured, so now we need to tell the system to pass IPv6 routing information as well as IPv4. Here is an example configuration as would be used on r1-isp6-ixp1:

```
r1-isp6-ixp1(config)# interface fastethernet 0/1
r1-isp6-ixp1(config-if)# ipv6 router isis workshop
```

and on r2-isp6-ixp1:

```
r2-isp6-ixp1(config)# interface fastethernet 0/0
r2-isp6-ixp1(config-if)# ipv6 router isis workshop
```

Note: The configuration is very simple as we've already done all the work for IPv4. This is one of the advantages of choosing ISIS instead of OSPF (which supports IPv4) and OSPFv3 (which supports IPv6).

Checkpoint #1: *call lab assistant to verify the connectivity. Demonstrate that you can ping and telnet to the adjacent routers.*

3. Configuring iBGP Neighbours.

Use your assigned AS number when you set up your iBGP sessions. When you have configured the sessions, use ***show bgp ipv6 unicast summary*** to check the peering. The BGP peering will be established using the loopback interfaces' IP address. For example, for r1-isp6-ixp1:

```
  router bgp 64016
  neighbor FD90:100:66:122::252 remote-as 64016
  neighbor FD90:100:66:122::252 description r2-isp6-ixp1 - IPv6
  neighbor FD90:100:66:122::252 update-source Loopback0

  address-family ipv4
    neighbor FD90:100:66:122::252 activate
    neighbor FD90:100:66:122::252 next-hop-self
```

```
    distance bgp 200 200 200
  exit-address-family
```

and for r2-isp6-ixp1:

```
router bgp 64016
  neighbor FD90:100:66:122::251 remote-as 64016
  neighbor FD90:100:66:122::251 description r1-isp6-ixp1 - IPv6
  neighbor FD90:100:66:122::251 update-source Loopback0

  address-family ipv4
    neighbor FD90:100:66:122::251 activate
    neighbor FD90:100:66:122::251 next-hop-self
    distance bgp 200 200 200
  exit-address-family
```

Q. Why is update-source loopback 0 necessary on iBGP?

Use **show bgp ipv6 summary** to check the status of the iBGP neighbour connections. If the iBGP session is not up and/or no updates are being sent, work with the Router Team for that neighbour connection to troubleshoot the problem.

4. Sanity Check.

Remember to use the following commands to ensure you are getting the information you are suppose to be getting:

```
show bgp ipv6 unicast summary : see a list of BGP peers that the router sees
show bgp ipv6 unicast : see a list of BGP paths that the router sees
show ipv6 route : see all the routes that the router has installed
```

Q. Are there routes seen via **show bgp ipv6 unicast**? If not, why not? Are there any routes tagged "B" when you do a **show ipv6 route**?

5. Add Networks via BGP.

Each Router Team will use BGP to advertise the address block used for the Module. For example, Router Team 6 would add the following configuration to their **core** router, r2-isp6-ixp1:

```
router bgp 64016
  address-family ipv6
    network FD90:100:66:120::/60
  exit-address-family
```

Use **show bgp ipv6 unicast** on neighbour's router to see if you are advertising your network via BGP.

Q. Does the network show up via BGP? If not, why?

Enter a static route for the CIDR block. For example, on r2-isp6-ixp1 we would use:

```
r2-isp6-ixp1 (config)# ipv6 route FD90:100:66:120::/60 Null0
```

Q. Does the network show up via the neighbour's BGP?

Use the command **show bgp ipv6 unicast neighbors <neighbour's IPv6 address> advertised-routes** to see what you are exporting to the other router. Physically go to one of your neighbour's routers and check their BGP Table.

Explain what you see.

Q. Does the network appear in the router's forwarding table?

Use the command **show ipv6 route** to check the local forwarding table. If not, why not?

Q. Why do we add the network statement and the static route on this router only?

Checkpoint #2: *call the lab assistant to verify the connectivity.*

6. Other Features in iBGP.

Review the documentation or use command line help by typing ? to see other show commands and other BGP configuration features.

Review Questions

1. What IOS show command(s) will display the router's BGP route table?
2. Why is the static pull up route necessary in relation to inserting a prefix into iBGP?