More iBGP and Basic eBGP - IPv6

- Objective: Connect your ISP using IPv6 to a Transit provider and the Internet Exchange Point using a combination of ISIS, internal BGP, and external BGP.
- Prerequisites: Basic Topology and ISIS, iBGP and iBGP - IPv6 modules

Lab Notes

The purpose of this module is to introduce the student to external BGP (eBGP). This is the relationship between different autonomous systems in an “Internet”. Each virtual IXP is split into seven distinct networks, and the teams belonging to each network work together as a typical ISP. Each ISP has a link to its Transit provider and to the IXP, and this feature will used throughout a significant portion of this workshop.

The connectivity shown in the diagrams represents links between AS’s. It is assumed that all the routers within an AS are physically connected to each other as per Figure 1.

Lab Exercises

1. Final ISIS check

Use the various *“show isis”* commands to see the ISIS status of the lab network now. Check the IPv6 routing and the routing table.

4. Test internal BGP connectivity.

Remember that the loopback addresses are used as the end points of the BGP connections. If IS-IS is not working then the BGP sessions can't be established.

Use the BGP Show commands to ensure you are still receiving IPv6 routes from within your AS.

2. eBGP peering information.

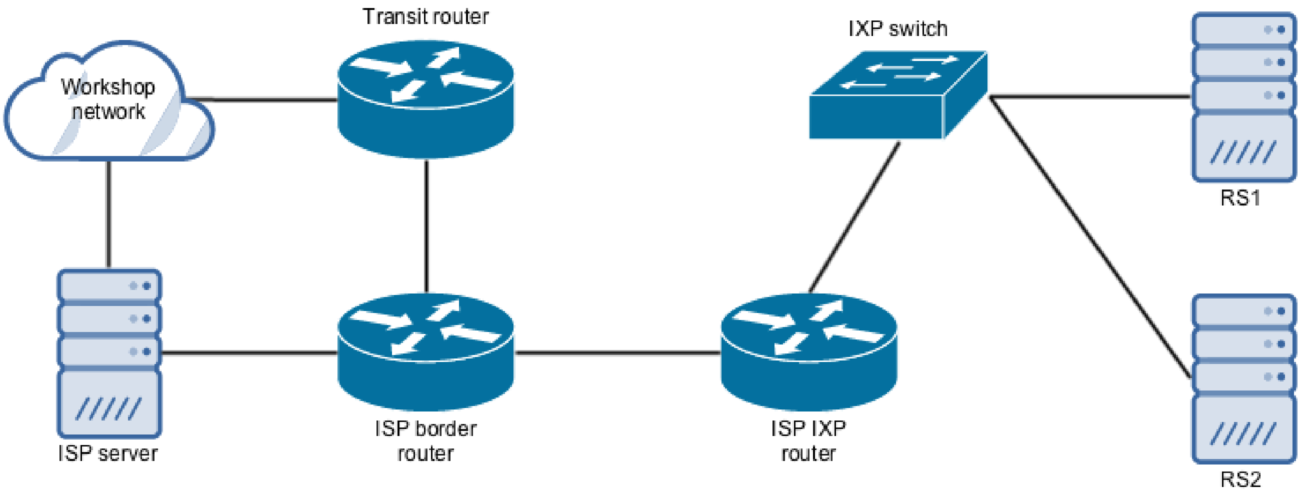

The diagram above shows the connections from your ISP routers to your Transit provider and to your IXP.

IXP1	IPv4	IPv6	AS number
Transit	100.66.101.254	FD90:100:66:101::254	64110
rs1	192.168.1.8	FD90:192:168:1::8	64010
rs2	192.168.1.9	FD90:192:168:1::9	64010
IXP2	**IPv4**	**IPv6**	**AS number**
Transit	100.66.102.254	FD90:100:66:102::254	64120
rs1	192.168.2.8	FD90:192:168:2::8	64020
rs2	192.168.2.9	FD90:192:168:2::9	64020

These tables show the relevant BGP parameters for the Transit and IXP BGP sessions. Now we'll configure the IPv6 sessions.

Q. When you make a BGP connection to the Transit provider which IPv6 address will you use as the **neighbor**? Which AS number will you use?

A.

```
IPv6 address:
AS number:
```

Q. When you make a BGP connection to the IXP route servers which IPv6 addresses will you use as the **neighbor**? Which AS number will you use?

A.

```
IPv6 addresses:
AS number:
```

Q. Why can’t the loopback interfaces be used for the eBGP peerings?

A. The IP address of a router’s loopback interface is not known to external BGP peers, so the external peers will have no way of knowing how to contact each other to establish the peering.

Q. Which BGP show command allows you to see the state of the BGP connection to your peer?

A. Try show bgp ipv6 unicast neighbor X:X:X::X – this will give detailed information about the state of the peer. There are subcommands of this one, giving more information about the peering.

Q. Which BGP Show command will allow you to see exactly which networks you are advertising and receiving from your eBGP peers?

A. Try show bgp ipv6 unicast neighbor X:X:X::X route – this will show which routes you are receiving from your peer. Likewise, replacing route with advertised-routes will list the networks which are being announced to your peer. (Note that in general ISP operational practice, there are caveats here – if you apply route-maps and some BGP policies, these will not be processed by the advertised-routes command. Use the advertised-routes subcommand with due caution.)

8. Configure Transit peering.

The eBGP configuration on the Transit router is already configured. You need to configure your side of the connection to complete the setup.

For example on r2-isp4-ixp2:

```
router bgp 64024
  neighbor fd90:100:66:102::254 remote-as 64120
  neighbor fd90:100:66:102::254 description TRANSIT-V6

  address-family ipv6
    neighbor fd90:100:66:102::254 activate
  exit-address-family
```

Now use **show bgp ipv6 unicast summary** to check your BGP session with the transit provider.

Q. How many prefixes are you learning from the Transit provider?

A. It will depend on whether you are the first person to get things working! If other people have things setup then you will see their routes.

Q. How would you look at any routes you are receiving?

A. show bgp ipv6 unicast and **show ipv6 routes**

Checkpoint #1. *Call the lab assistant to verify the connectivity.*

9. Configure IXP peering.

The Internet Exchange Point offers a Route Server service. There are two Linux servers running the BIRD Routing Daemon software. The BIRD package talks a number of routing protocols but in this instance we will only be using BGP part of it. We run two servers to provide redundancy.

If each ISP peers with these servers then we will each be able to retrieve copies of each others routing announcements. Because we are all on the same Layer 2 switch we can get the routing information from the route servers but send the traffic directly from, for example, ISP1 to ISP4. Let's configure the connections and see how this works.

As we have two connections, that have features in common, to configure we will use a **peer-group** instead of building two completely separate connections. For example, on r1-isp5-ixp1:

```
router bgp 64015
  neighbor fd90:192:168:1::8 remote-as 64010
  neighbor fd90:192:168:1::8 description ROUTE-SERVERS-V6 - rs1

  address-family ipv6
    neighbor fd90:192:168:1::8 activate
  exit-address-family
```

This would add a peering session to **rs1** on **IXP1**.

Check that this is working using the commands you have already used for previous connections.

Now we can add the BGP session to **rs2**:

```
router bgp 64015
  neighbor fd90:192:168:1::9 remote-as 64010
  neighbor fd90:192:168:1::9 description ROUTE-SERVERS-V6 - rs2

  address-family ipv6
    neighbor fd90:192:168:1::9 activate
  exit-address-family
```

Check that this session is working as well. As before the number of prefixes you see will depend on who else is peering with the route servers.

10. Configure passwords on the eBGP session.

Passwords should now be configured on the eBGP sessions between your and your neighbouring ASes. Agree between you and your neighbouring AS what the password should be on the eBGP session, and then apply it to the eBGP peering. You should use the example we had in the previous section as a template.

There is a problem however. In the previous section, you controlled both routers and could make changes on both.

Checkpoint #2: *Call the lab assistant and work with them to set the password as set on the eBGP sessions with the Transit router and the IXP route servers. Once confirmed by the lab assistant, move on to the next steps.*

As previously for the iBGP session, watch the logs for password mismatches, or missing passwords. As with the iBGP sessions previously, you will find that the router will reset the eBGP session as soon as the password is applied.

Note: Wherever a BGP (either iBGP or eBGP) session is configured from now on in the workshop, all Router Teams MUST use passwords on these BGP sessions.

11. Check the BGP table.

Are there routes seen via **show ipv6 bgp**? If not, why not? Once every team in the class has done their configuration, each team should see the /60 aggregate from each AS. If this is not happening, work with your neighbours to fix the problem.

Checkpoint #3: *Call the lab assistant to verify the connectivity. Use commands such as "show ipv6 route sum", "show bgp ipv6 unicast sum", "show bgp ipv6 unicast", and "show bgp ipv6 unicast neigh X:X:X:X::X route | advertise". There should be 7 aggregate prefixes (one for each ISP) in the BGP table.*

12. The Importance of Aggregation.

Each AS was allocated a /60 address block. It is expected by all Internet operators that any address space an ISP is using is aggregated as much as possible before it is announced to the rest of the Internet. Subdividing the address space inside an AS is perfectly acceptable and obviously very common (as we have done here) – but most operators consider leaking this subdivided address space out to the Internet at large antisocial and unfriendly.

Q. How do you automatically aggregate via BGP smaller address blocks from within your network to a larger address block outside your network? Hint: Review the BGP documentation.

A. The "aggregate-address" command is quite often used to achieve this. We are not doing any filtering or limitation of the announcements of the "customer" address blocks we have introduced into each ASN. This will be one of the goals of the next modules in the workshop.

Review Questions

1. How many origin types exist in BGP?
2. List the origin types. Hint: Review the BGP presentations.
3. How are they used?
4. Why are passwords necessary on both iBGP and eBGP sessions? What do they protect against?
5. Why is aggregation important for the Internet?