

Module 2 – IPv6 iBGP and Basic eBGP

Objective: Using IPv6, simulate four different interconnected ISP backbones using a combination of ISIS, internal BGP, and external BGP.

Prerequisites: Module 1

Topology :

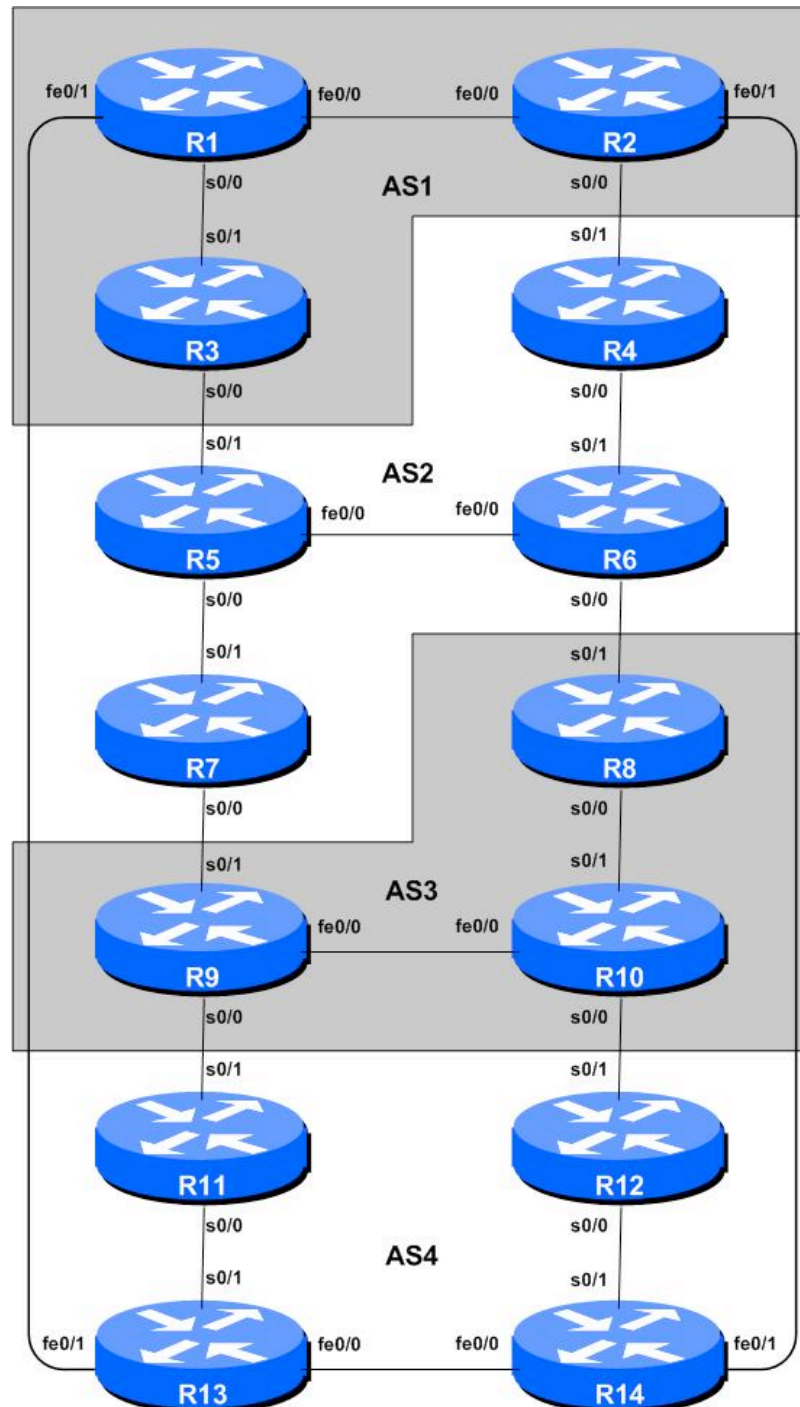


Figure 1 – BGP AS Numbers

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”. The classroom is split into four distinct networks, and the teams belonging to each network work together as a typical ISP. Each AS has two links to its neighbouring ASes, and this feature will be 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.

Note: this IPv6 module can be completed independently of the IPv4 version of this module. It only requires the basic topology and connectivity provided by Module 2. If IPv4 is not configured at all, remember to manually set the BGP router-id.

Lab Exercises

1. Connect routers as shown in Figure 1. All routers within an AS must be physically connected and reachable. The relationship between the ASes is as drawn in Figure 2 and gives a view which can be related to the “real world”.




Figure 2 – AS relationship

2. The address assignments and addresses used for links between routers should be left the same as those chosen for Module 1.
3. **Re-configure BGP and ISIS (if coming from Module 1).** On each router, remove the BGP process from Module 1 by using the following command:

```
Router1(config)# no router bgp 10
```

This will clear the BGP configuration for the current module. You will need to carefully remove the ISIS IPv6 configuration by hand though, by going to each interface and removing the ISIS IPv6 configuration directly.

(The alternative is to simply erase the entire router configuration using *write erase*, and then reload the router and start again, completing all steps of Module 1 up to Step 11. Or copying the configuration you saved at the end of Step 11, just prior to starting the ISIS configuration. Which you probably forgot to do, even with both the lab instructor and these written notes telling you to do so.)

4. **Configure ISIS for IPv6 on the routers within each AS.** In each AS configure ISIS routing. This means that each router team should configure *router ISIS* with ISIS ID *isp-asy* on the router, where *y* is the AS number. And the links to each member in the AS must be configured with *ip router ISIS isp-asy*. The NET should be *49.0001.x.x.x.x.00*, where *x.x.x.x* is built from the loopback IPv4 address.

ISIS should be configured on internal interfaces **only**. You do not want to set up adjacencies with devices outside your AS. Make sure that there are no *ipv6 router isis* commands on external interfaces. We also need to configure the IPv6 address family and enable multi-topology IS.

As an example, Router Team 1, with two interfaces in AS 1 would have the following:

```
Router1 (config)# router isis isp-as1
Router1 (config-router)# net 49.0001.1000.0101.5224.00
Router1 (config-router)# is-type level-2-only
Router1 (config-router)# metric-style wide
Router1 (config-router)# log-adjacency-changes
Router1 (config-router)# address-family ipv6
Router1 (config-router-af)# multi-topology
!
Router1 (config)# interface fastethernet 0/0
Router1 (config-if)# ipv6 router isis isp-as1
Router1 (config-if)# isis ipv6 metric 2 level-2
!
Router1 (config)# interface serial 0/0
Router1 (config-if)# ipv6 router isis isp-as1
Router1 (config-if)# isis ipv6 metric 20 level-2
```

Now mark the interfaces on which you do not want to run ISIS as *passive*. For ISIS, marking an interface as *passive* means that CLNS adjacencies are not solicited **and** the IP subnet used on the interface is inserted into ISIS. Note that you cannot mark interfaces as passive until you have ISIS assigned to at least one physical interface on the router.

```
Router1 (config)# router isis isp-as1
Router1 (config-router)# passive-interface Loopback0
Router1 (config-router)# passive-interface Fastethernet 0/1
```

Notes:

- ISIS by default will only set up adjacencies and announce the prefixes of the interfaces which are activated by the “*ipv6 router isis*” command. This is different behaviour from OSPF which will attempt to set up adjacencies on interfaces covered by the *network* statement (and hence require the use of *passive* and *no passive* to control its behaviour).

- Different ISPs use different NET addressing scheme. But it is common using router loopback IP address as the system ID in either hex or decimal format.

5. **Ping Test.** Check the routes via ISIS. Make sure you can see all the networks within your AS, and see no networks from other ASs. Ping all loopback interfaces within your AS Set. Use the “*show cns neighbor*” and “*show ip route*” commands.
6. **Save the configuration.** Don’t forget to save the configuration to NVRAM!

Checkpoint #1: *call the lab assistant to verify the connectivity.*

7. **Configure iBGP peering between routers within an AS.** Use the loopback address for the iBGP peerings. Also, configure the *network* command to add the address block assigned to each Router Team for advertisement in BGP.

```
router bgp 1
no bgp default ipv4-unicast
bgp log-neighbor-changes
address-family ipv6
no synchronization
network 1001:10::/32
neighbor 1001:11::1 remote-as 1
neighbor 1001:11::1 update-source loopback 0
neighbor 1001:11::1 description iBGP Link to R2
neighbor 1001:13::1 remote-as 1
neighbor 1001:13::1 update-source loopback 0
neighbor 1001:13::1 description iBGP Link to R3
!
ipv6 route 1001:10::/32 Null0
```

Q: Do you need the BGP command *no synchronization*? Why?

A: An ISP network is a **transit** network, meaning it accepts packets from one peering AS, carries it across the backbone, then hands it over to the next AS toward the destination. To ensure that routers internal to the AS are able to forward transit packets (from the ingress border router to the egress border router), all BGP border routers will wait for a network prefix to arrive in the IGP (as they all participate in the same IGP) before advertising them to external BGP peers. This is referred to as **synchronization**. In other words, internal routers must be aware of those prefixes learned via IGP that border routers learn via iBGP.

As you can see, this applies to an environment where BGP routes are redistributed into the IGP. A typical ISP usually doesn’t do that as the Internet routing table is somewhat large. Instead, it runs a fully meshed iBGP (or uses route-reflectors) among all routers in the backbone. Therefore synchronisation should be turned off in this kind of environment.

8. **Test internal BGP connectivity.** Use the BGP Show commands to ensure you are receiving everyone's routes from within your AS.
9. **Configure passwords on the iBGP sessions.** Passwords should now be configured on the iBGP sessions. Review the presentation why this is necessary. Agree amongst all your team members in your AS what the password should be on the iBGP session, and then apply it to all the iBGP

peerings on your router. For example, on Router2's peering with Router3, with "cisco" used as the password:

```
router bgp 1
address-family ipv6
neighbor 1001:13::1 password cisco
```

IOS currently resets the iBGP session between you and your neighbouring router whenever an MD5 password is added. So when passwords are added to BGP sessions on live operational networks, this work should be done during a maintenance period when customers know to expect disruptions to service. In the workshop lab, it doesn't matter so much. (Future IOS releases will avoid having this rather serious service disruption.)

Watch the router logs – with the BGP session neighbour changes being logged, any mismatch in the password should be easy to spot.

Checkpoint #2: *Call the lab assistant and demonstrate the password as set on the iBGP session. Once confirmed by the lab assistant, move on to the next steps.*

10. Configure eBGP peering. Use Figure 1 to determine the links between the AS's. Addressing for eBGP links between 2 AS's will use the point-to-point interface addresses, **NOT** the loopback addresses (review the BGP presentation if you don't understand why). So, for Router1's peering with Router13, the configuration might look like:

```
router bgp 1
address-family ipv6
neighbor 1001:10:0:2::2 remote-as 4
neighbor 1001:10:0:2::2 description eBGP to Router13
```

Use the BGP Show commands to ensure you are sending and receiving the BGP advertisements from your eBGP neighbours.

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.)

- 11. 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. For example, on Router2's peering with Router4, with "cisco" used as the password:

```
router bgp 1
  address-family ipv6
    neighbor 1001:11:0:1::2 password cisco
```

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.

Checkpoint #3: Call the lab assistant and demonstrate the password as set on the eBGP session. Once confirmed by the lab assistant, move on to the next steps.

- 12. Aggregate each AS's CIDR Blocks.** Each router team was allocated either a /32 address block or a /31 address block in the first Module. However, each AS has three or four routers in it, so we need to take the address space from each router team in the AS and aggregate it before we make any announcement to any external AS. 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 – but leaking this subdivided address space out to the Internet at large is considered antisocial and unfriendly by many ISPs. In this case the address blocks belonging to each AS can be aggregated into a larger /30 address block.

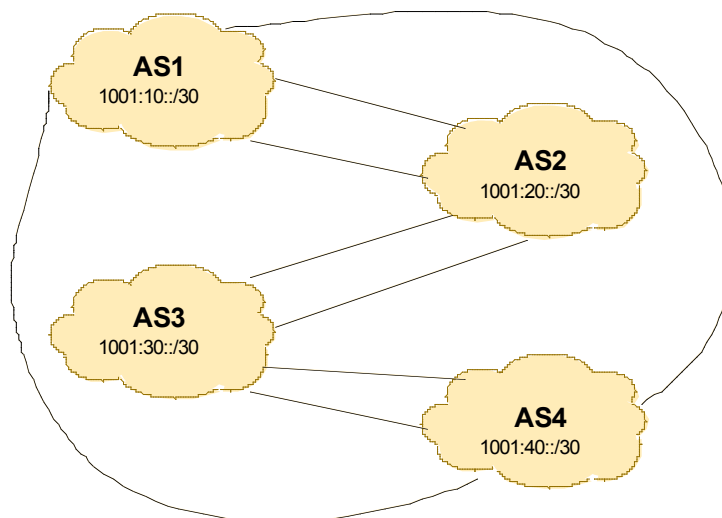


Figure 3 – Aggregates for each ASN




Figure 4 – Aggregating AS1 IPv6 address space into a /30

For example, AS1 has three routers in it. Router 1 was allocated 1001:10::/32, Router 2 was allocated 1001:11::/32 and Router 3 was allocated 1001:12::/31. These three address blocks can be aggregated into the 1001:10::/30 network. And this /30 is what should be announced to eBGP neighbours.

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. Configure:

```
router bgp 1
  address-family ipv6
    aggregate-address 1001:10::/30
```

Type ? after the command to see what options this command has.

13. Examine the *origin* of the network prefixes. What is the origin type for the aggregated prefixes? Write your answer here:

14. Check the network paths. Do traceroutes to hosts nominated on the network by the lab instructor.

Checkpoint #4: Call the lab assistant to verify the connectivity. Use commands such as “show ip route sum”, “show bgp ipv6 unicast sum”, “show bgp ipv6 unicast”, “show ipv6 route”, and “show bgp ipv6 unicast neigh x.x.x.x route | advertise”. There should be 13 specific prefixes and 4 aggregate prefixes (one for each ISP).

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?

CONFIGURATION NOTES

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