Module 6 – IPv6 iBGP and Basic eBGP

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

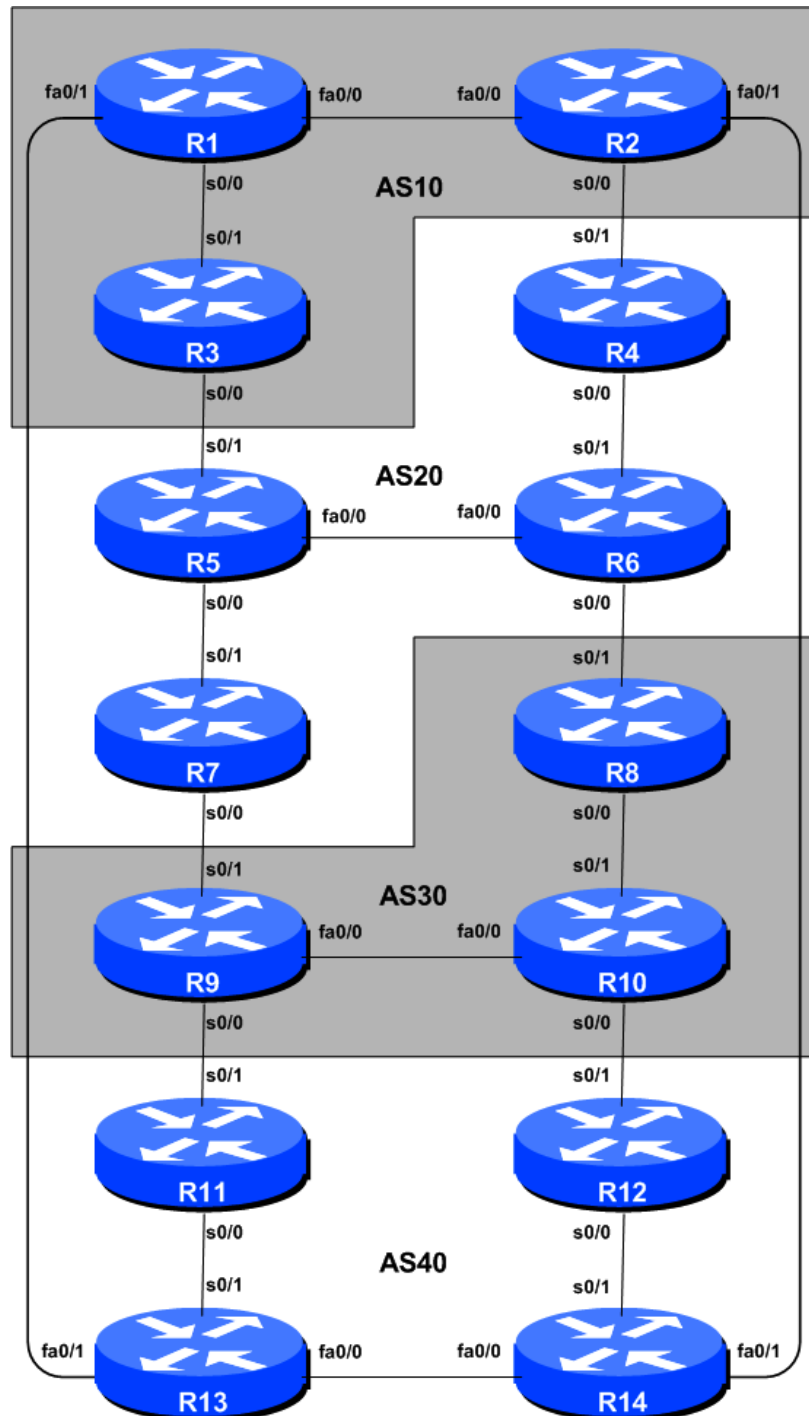**Prerequisites: Module 1**

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 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. **Remove the IP addressing.** Before we proceed to even think about configuring the routing protocols to match the above AS layout, we need to tidy up the addressing from previous modules. In this step we remove any remaining IPv6 addresses from all physical interfaces, and the loopback interface. This effectively takes the lab back to the configuration it had prior to step 10 in Module 1. Don't forget to remove **all** the IPv6 addressing.

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.

4. **IPv6 Addressing.** As we did in step 5 of Module 1, we need to come up with a sensible and scalable addressing plan for each AS in this network. Each AS gets their own address block, again a /20 (typical minimum allocation for a starter ISP). This address block should be assigned to links and loopbacks on the routers making up each ASN. The allocations are as follows:

AS10	**2001:db8::/32**	**AS30**	**2001:dba::/32**
AS20	**2001:db9::/32**	**AS40**	**2001:dbb::/32**

Again we need to divide up each address block so that we have customer address space, network infrastructure address space, and some space for loopbacks. Figure 3 below reminds how this could be done:


Figure 3 – Dividing allocated block of /20 into Customer, Infrastructure and Loopbacks

Please refer to the accompanying hand out for the address plan which should be used for this module onwards – it is entitled "Addressing Plan – Modules 6 to 9". As with Module 1, configure the addresses on each interface which will be used for this module, and check basic IP connectivity with your immediately adjacent neighbours.

5. **Router Loopback Interface Addressing.** As the minimum subnet size possible for IPv6 is a /64, we will assign the first /64 out of our /48 infrastructure block to be used for loopbacks even though each AS has either 3 or 4 routers in it. The loopback address assignments which will be used for this module are below:

Router1	**2001:db8::1**	**Router8**	**2001:dba::1**
Router2	**2001:db8::2**	**Router9**	**2001:dba::2**
Router3	**2001:db8::3**	**Router10**	**2001:dba::3**
Router4	**2001:db9::1**	**Router11**	**2001:dbb::1**
Router5	**2001:db9::2**	**Router12**	**2001:dbb::2**
Router6	**2001:db9::3**	**Router13**	**2001:dbb::3**
Router7	**2001:db9::4**	**Router14**	**2001:dbb::4**

6. **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.00*, where *x.x.x* is built from the loopback IPv4 address (refer to Module 1 for details).

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.

As an example, Router Team 1, with two interfaces in AS 10 would have the following:

```
Router1 (config)# router isis as10
Router1 (config-router)# net 49.0001.0100.0001.5224.00
Router1 (config-router)# is-type level-2-only
Router1 (config-router)# metric-style wide
Router1 (config-router)# log-adjacency-changes
!
Router1 (config)# interface fastethernet 0/0
Router1 (config-if)# ipv6 router isis as10
Router1 (config-if)# isis ipv6 metric 2 level-2
!
Router1 (config)# interface serial 0/0
Router1 (config-if)# ipv6 router isis as10
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 as10
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 schemes. But it is common to use the router loopback IP address as the system ID in either hex or decimal format. In this module all routers in an AS are level-2 in and one area (*49.0001*) only.

7. **ISIS on Point-to-Point Ethernet Links.** One feature mentioned in the ISIS presentations was the option to modify ISIS's behaviour on point-to-point broadcast media links, such as Ethernet, when there are only two devices on that media. If we declare such a situation to be point to point, then ISIS does not try and determine a designated router; furthermore, there will be an improvement/simplification in SPF calculations and memory requirements on the router.

Those router teams which have ISIS configured over an Ethernet interface should now convert ISIS to point to point mode, for example:

```
Router1 (config)# interface fastethernet 0/0
```

```
Router1 (config-interface)# isis network point-to-point
```

The link is now treated like a point-to-point serial connection.

8. **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 clns neighbor*" and *"show ip route"* commands.

9. **Save the configuration.** Don't forget to save the configuration to NVRAM!

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

10. **Turning on neighbour authentication for ISIS – Part 1.** ISIS supports neighbour authentication; this is considered more and more important inside ISP networks as attacks on infrastructure increase and ISPs seek to use all available tools to secure their networks. (While an attack on ISIS is harder as it runs on the link layer alongside IP rather than on top of IP like OSPF, some ISPs are still prudent and implement neighbour authentication.)

    Each router team will now turn on neighbour authentication for ISIS. The first step is to set up the keychain to be used – we will use the key "cisco" for this lab:

    ```
    Router1(config)# key chain isis-level2
    Router1(config-keychain)# key 1
    Router1(config-keychain-key)# key-string cisco
    ```

11. **Turning on neighbour authentication for ISIS – Part 2.** Now that the keychain has been defined, we activate all the router interfaces to support authentication. The first step is to enable MD5 for level-2 IS's:

    ```
    Router1(config)# interface fastethernet 0/0
    Router1(config-if)# isis authentication mode md5 level-2
    ```

    And then associate the key-chain we defined earlier with the configured authentication:

    ```
    Router1(config)# interface fastethernet 0/0
    Router1(config-if)# isis authentication key-chain isis-level2 level-2
    ```

    Notice now that the ISIS adjacencies do not come up unless the neighbouring router has also entered the same configuration and key. Notice also how the ISIS adjacencies were reset as the configuration was entered – security is being introduced, so the adjacencies are reset.

12. **Final check.** Use the various "*show isis*" commands to see the ISIS status of the lab network now. Check the routing and the routing table. Make sure all the adjacencies have come back up again. If any adjacency has failed to come up, and you see several log messages saying:

    ```
    *Mar  1 00:05:17.825: %CLNS-4-AUTH_FAIL: ISIS: LAN IIH authentication failed
    ```

    you should reasonably expect that either you or your connected neighbour have forgotten to set up neighbour authentication.

Note: Wherever an ISIS session is configured from now on in the workshop, all Router Teams MUST use passwords on these ISIS sessions.

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

13. **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 10
 no bgp default ipv4-unicast
 bgp log-neighbor-changes
 address-family ipv6
  no synchronization
  network 2001:db8::/32
  neighbor 2001:db8::2 remote-as 10
  neighbor 2001:db8::2 update-source loopback 0
  neighbor 2001:db8::2 description iBGP Link to R2
  neighbor 2001:db8::3 remote-as 10
  neighbor 2001:db8::3 update-source loopback 0
  neighbor 2001:fb8::3 description iBGP Link to R3
 !
ipv6 route 2001:db8::/32 Null0
```

14. **Test internal BGP connectivity.** Use the BGP Show commands to ensure you are receiving everyone's routes from within your AS.

15. **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 10
 address-family ipv6
  neighbor 2001:db8::3 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 #3: *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.*

16. **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 10
 address-family ipv6
  neighbor 2001:db8:0:4::2 remote-as 40
  neighbor 2001:db8:0:4::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.)

17. **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 10
 address-family ipv6
  neighbor 2001:db8:0:3::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 #4: *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.*

18. **Adding a "customer" route into BGP.** As we did in Module 1, we are now going to add a "customer" route into BGP on each router. We don't have any "customers" as such connected to our routers in the lab, so we are going to simulate the connectivity by simply using a Null0 interface. The "customer" address space that each router team will introduce into the iBGP is listed below – again we will each use a /48, for simplicity's sake.

R1	`2001:db8:1::/48`		R8	`2001:dba:1::/48`
R2	`2001:db8:2::/48`		R9	`2001:dba:2::/48`
R3	`2001:db8:3::/48`		R10	`2001:dba:3::/48`
R4	`2001:db9:1::/48`		R11	`2001:dbb:1::/48`
R5	`2001:db9:2::/48`		R12	`2001:dbb:2::/48`
R6	`2001:db9:3::/48`		R13	`2001:dbb:3::/48`
R7	`2001:db9:4::/48`		R14	`2001:dbb:4::/48`

Each team should now set up a static route pointing to the **NULL0** interface for the /48 that they are to originate. Once the static is set up, the team should then add an entry into the BGP table. Here is an example for Router8:

```
ipv6 route 2001:dba:1::/48 Null0
!
router bgp 30
 address-family ipv6
  network 2001:dba:1::/48
!
```

19. **Check the BGP table.** Are there routes seen via *show bgp ipv6*? If not, why not? Once every team in the class has done their configuration, each team should see the aggregate from each AS as well as the fourteen /48s introduced in the previous step. If this is not happening, work with your neighbours to fix the problem.

Checkpoint #5: *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", "show ipv6 route", and "show bgp ipv6 unicast neigh x.x.x.x route | advertise". There should be 4 aggregate prefixes (one for each ISP) and the 14 customer /48's in the BGP table.*

20. **The Importance of Aggregation.** Each AS was allocated a /32 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.*

Cisco Systems Inc
170 West Tasman Drive.
San Jose, CA 95134-1706
Phone: +1 408 526-4000
Fax: +1 408 536-4100

8

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.
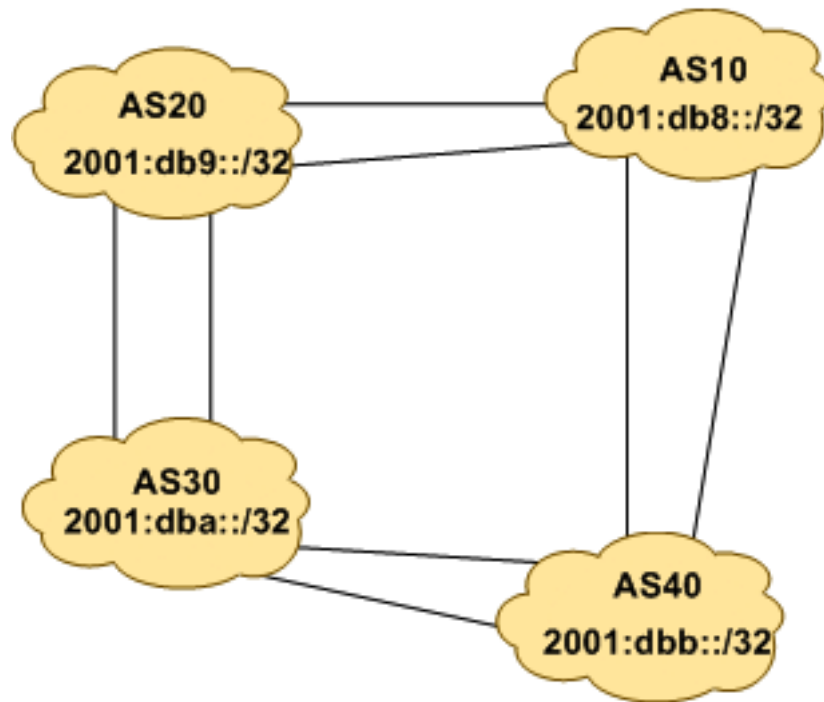
Figure 4 – Aggregates for each ASN

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.