Module 2 – More iBGP, and Basic eBGP Configuration

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

Prerequisites: Module 1 (ISIS)

Topology :

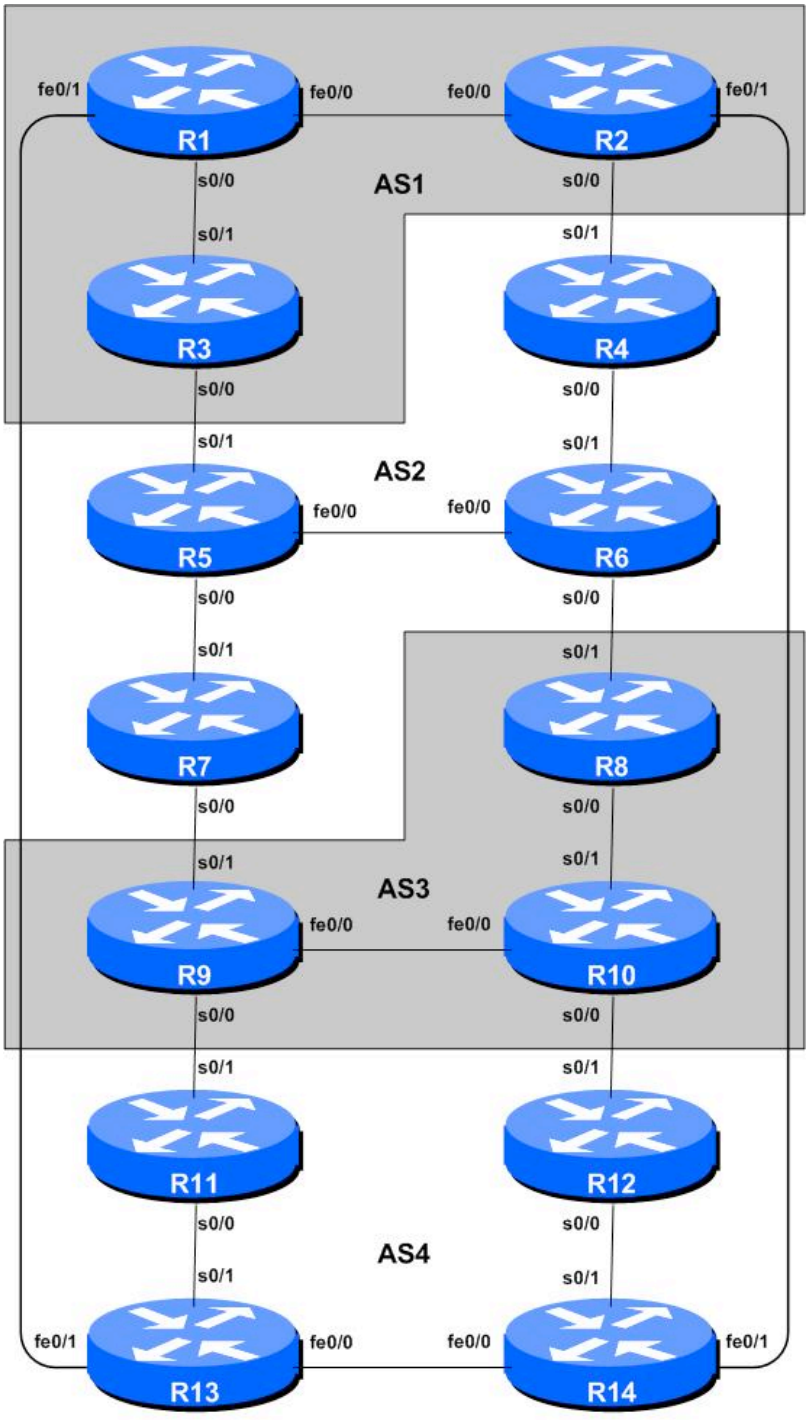

Figure 1 – BGP AS Numbers

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.

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".

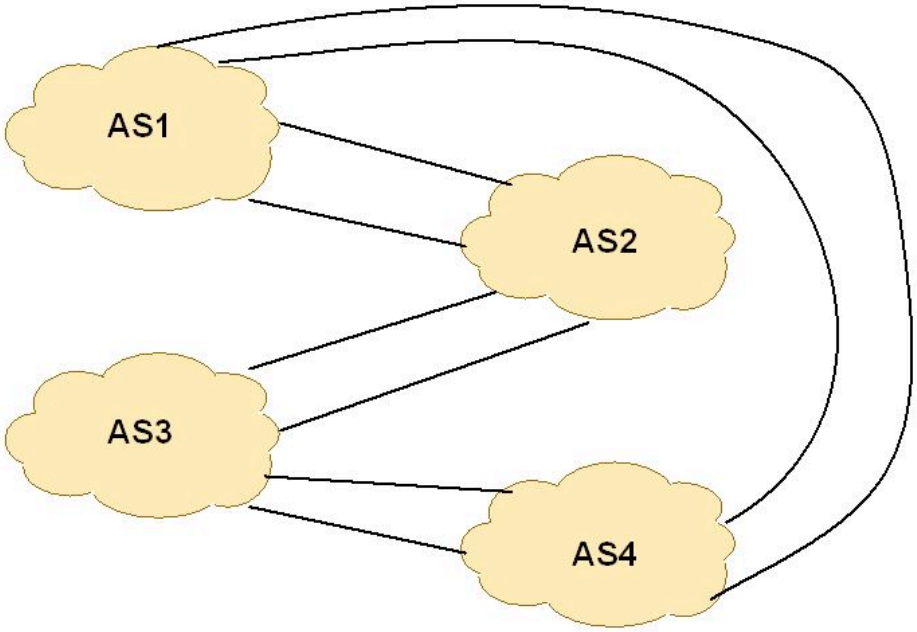

Figure 2 – AS relationship

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.** On each router, remove the BGP and ISIS processes from Module 1 by using the following two commands:

```
Router1 (config)# no router bgp 10
Router1 (config)# no router isis isp-workshop
```

This will clear the BGP and ISIS configuration for the current module.

(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

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

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 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 *49.00zz.x.x.x.x.00*, where *zz* is the router number (this will set each router in different areas), and *x.x.x.x* is the loopback IP 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 *ip router isis* commands on external interfaces. 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.

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

```
Router1 (config)# clns routing
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)# passive-interface Loopback0
Router1 (config-router)# passive-interface Fastethernet 0/1
Router1 (config-router)# log-adjacency-changes
!
Router1 (config)# interface fastethernet 0/0
Router1 (config-if)# ip router isis isp-as1
Router1 (config-if)# isis metric 2 level-2
!
Router1 (config)# interface serial 0/0
Router1 (config-if)# ip router isis isp-as1
Router1 (config-if)# isis metric 20 level-2
```

   Notes:
   - ISIS by default will only set up adjacencies and announce the prefixes of the interfaces which are activated by the "`ip 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. In this module, we assign all routers in different areas, and all are level-2 routers. However, in module 1, all routers are level-2 in one area (*49.0001*).

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 clns 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.

```
Router1 (config)# router bgp 1
Router1 (config-router)# no synchronization
Router1 (config-router)# network 100.1.0.0 mask 255.255.240.0
Router1 (config-router)# neighbor 100.1.31.224 remote-as 1
Router1 (config-router)# neighbor 100.1.31.224 update-source loopback 0
Router1 (config-router)# neighbor 100.1.31.224 description iBGP Link to R2
Router1 (config-router)# neighbor 100.1.63.224 remote-as 1
Router1 (config-router)# neighbor 100.1.63.224 update-source loopback 0
Router1 (config-router)# neighbor 100.1.63.224 description iBGP Link to R3
Router1 (config-router)# no auto-summary
Router1 (config-router)# exit
Router1 (config)# ip route 100.1.0.0 255.255.240.0 Null0
```

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:

```
Router2 (config)# router bgp 1
Router2 (config-router)# neighbor 100.1.63.224 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. A missing password on one side of the BGP session will result in the neighbouring router producing these errors:

```
%TCP-6-BADAUTH: No MD5 digest from 3.3.3.3:179 to 2.2.2.2:11272
%TCP-6-BADAUTH: No MD5 digest from 3.3.3.3:179 to 2.2.2.2:11272
%TCP-6-BADAUTH: No MD5 digest from 3.3.3.3:179 to 2.2.2.2:11272
```

   whereas a mismatch in the configured passwords will result in these messages:

```
%TCP-6-BADAUTH: Invalid MD5 digest from 3.3.3.3:11024 to 2.2.2.2:179
%TCP-6-BADAUTH: Invalid MD5 digest from 3.3.3.3:11024 to 2.2.2.2:179
%TCP-6-BADAUTH: Invalid MD5 digest from 3.3.3.3:11024 to 2.2.2.2:179
```

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).

```
Router1 (config)# router bgp 1
Router1 (config-router)# neighbor 100.1.2.2 remote-as 4
Router1 (config-router)# neighbor 100.1.2.2 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 ip bgp 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 ip bgp 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:

```
Router2 (config)# router bgp 1
Router2 (config-router)# neighbor 100.1.17.2 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 /20 address block or a /19 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 /18 address block.

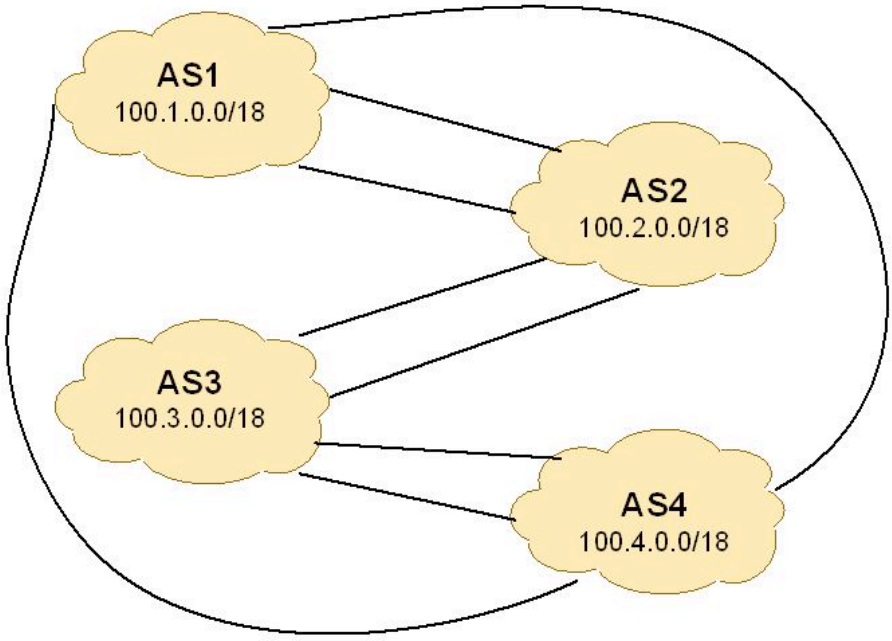

Figure 3 – Aggregates for each ASN

For example, AS1 has three routers in it. Router 1 was allocated 100.1.0.0/20, Router 2 was allocated 100.1.16.0/20 and Router 3 was allocated 100.1.32.0/19. These three address blocks can be aggregated into the 100.1.0.0/18 network. And this /18 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:

```
Router2(config)# router bgp 1
Router2(config)# aggregate-address 100.1.0.0 255.255.192.0
```

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

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

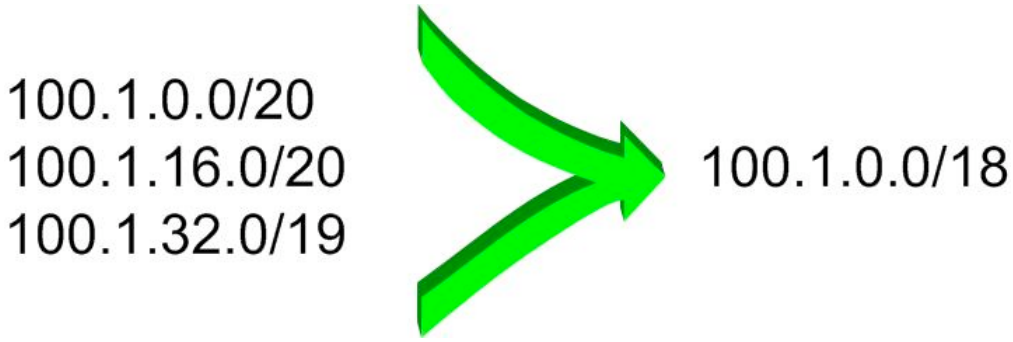

Figure 4 – Aggregating AS1 address space into a /18

13. 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 ip bgp sum", "show ip bgp", "show ip route", and "show ip bgp neigh x.x.x.x route | advertise". There should be 13 specific prefixes and 4 aggregate prefixes (one for each ISP).*

14. BGP Update Activity (Optional). Use *debug ip bgp update* to see BGP update activity after clearing a BGP session. To stop the debug running, do *undebug ip bgp update.*

Warning: it might not be such a good idea to run this debug command on a router receiving the full Internet routing table; using this command in a lab network such as this might show you why!

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.

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