Module 22 – Basic eBGP with ISIS Configuration

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

Prerequisites: Module 21

Topology :

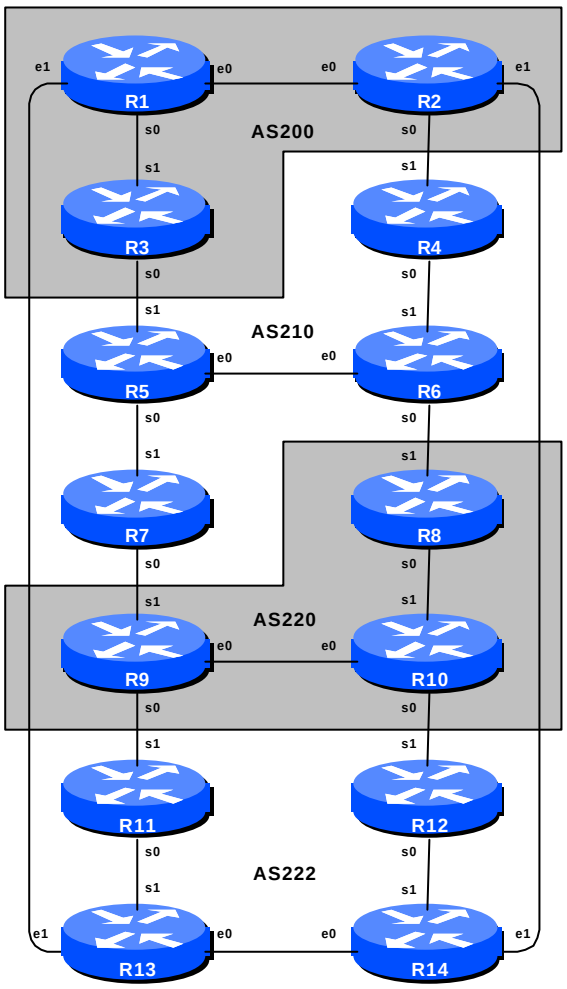

Figure 1 – BGP AS Numbers

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 **Error! Reference source not found.**.

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".

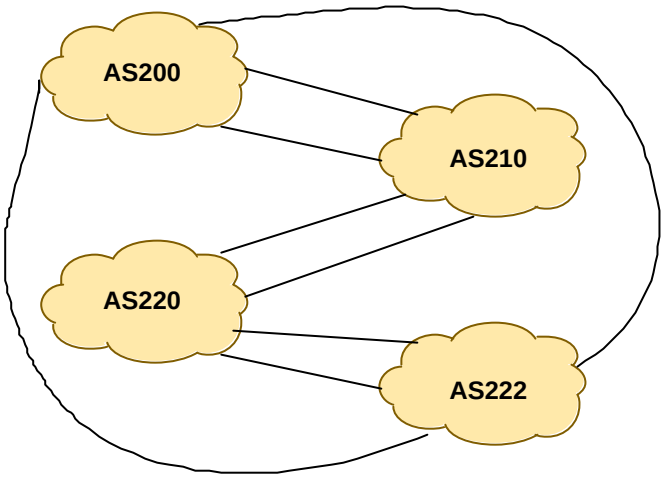

Figure 2 – AS relationship

2. The 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 from Module 1 by using the following three commands:

```
Router1 (config)# no router bgp 100
Router1 (config)# no router isis isp-workshop
```


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

```
Router1 (config)# no clns routing
```

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

4. **Add an Autonomous System Number pointer.** Even through it is not needed in BGP, the *autonomous-system* command helps people who are troubleshooting the router keep track of the AS number assigned to the router. Each Router Team will use this command in the router with their new AS number (see **Error! Reference source not found.**). Example:

```
Router1 (config)# autonomous-system 200
```

5. **Configure ISIS within the AS Set of routers.** Configure ISIS within each AS Set of routers. 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 *39.00z.z.x.x.x.x.00*, where *z.z* is two middle octets of the loopback IP address (this will set each router in different areas), and *x.x.x.x* is the loopback IP address. For example, Router Team 1, with two interfaces in AS 200 would have the following:

```
Router1 (config)# clns routing
Router1 (config)# router isis isp-as200
Router1 (config-router)# net 39.0020.0007.2002.0000.7224.00
Router1 (config-router)# is-type level-2-only
Router1 (config-router)# passive-interface Loopback0
Router1 (config-router)# passive-interface ethernet 0/1
Router1 (config-router)# log-adjacency-changes
!
Router1 (config)# interface ethernet 0/0
Router1 (config-if)# ip router isis isp-as200
Router1 (config-if)# isis metric 2 level-2
Router1 (config)# interface serial 0/0
Router1 (config-if)# ip router isis isp-as200
Router1 (config-if)# isis metric 20 level-2
```

Notes:

- *`passive-interface`* makes ISIS carry the DMZ networks (e1 subnet in this example) that will allow recursive routing to happen for BGP prefixes. It will also carry the loopback interfaces within the AS allowing iBGP peering via the loopback interface.
- 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 (*39.0000.0100*).

6. **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.

Checkpoint #1 : *call 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 *network* command to add the CIDR block assigned to each Router Team for advertisement in BGP.

```
Router1 (config)# router bgp 200
Router1 (config)# no synchronization
Router1 (config-router)# neighbor 200.200.11.224 remote-as 200
Router1 (config-router)# neighbor 200.200.11.224 update-source loopback 0
Router1 (config-router)# neighbor 200.200.19.224 remote-as 200
Router1 (config-router)# neighbor 200.200.19.224 update-source loopback 0
Router1 (config-router)# network 200.200.4.0 mask 255.255.252.0
Router1 (config-router)# exit
Router1 (config)# ip route 200.200.4.0 255.255.252.0 Null 0
```

   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 IGP. Typical ISP usually doesn't do that as the Internet routing table is somewhat large. Instead, it runs a fully meshed IBGP (or use route-reflector) among all routers in the backbone, so there is no need for BGP→IGP redistribution. Thus, synchronization 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 Set.

9. **Configure EBGP peering.** Use Figure 1 to determine the links between the AS Sets. Addressing for EBGP links between 2 AS's will use the point-to-point interface addresses, ***NOT*** the loopback addresses.

```
Router1 (config)# router bgp 200
Router1 (config-router)# neighbor 200.200.6.2 remote-as 222
```

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

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. What 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: there are caveats here – if you apply route-maps and some policy, these will not be processed by the *advertised-routes* command. Use with caution.)

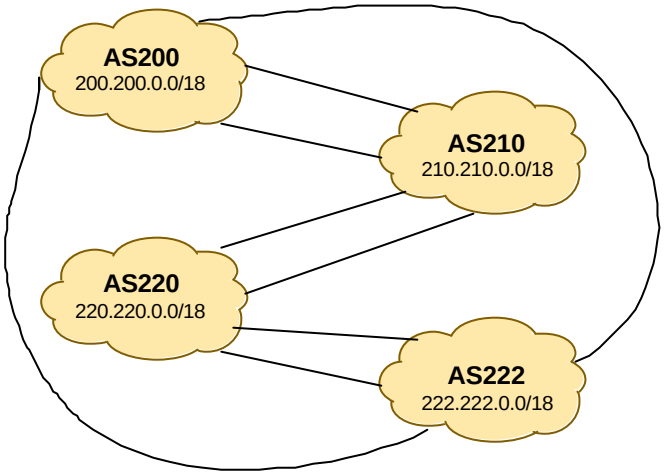

Figure 3 – Aggregates for each ASN

10. **Aggregate each AS Set's CIDR Blocks.** The /22 CIDR blocks belonging to each router in an AS can be aggregated into a larger /18 CIDR block. This will allow an ISP to have smaller prefixes inside their network while aggregating their network advertisement going outside their network.

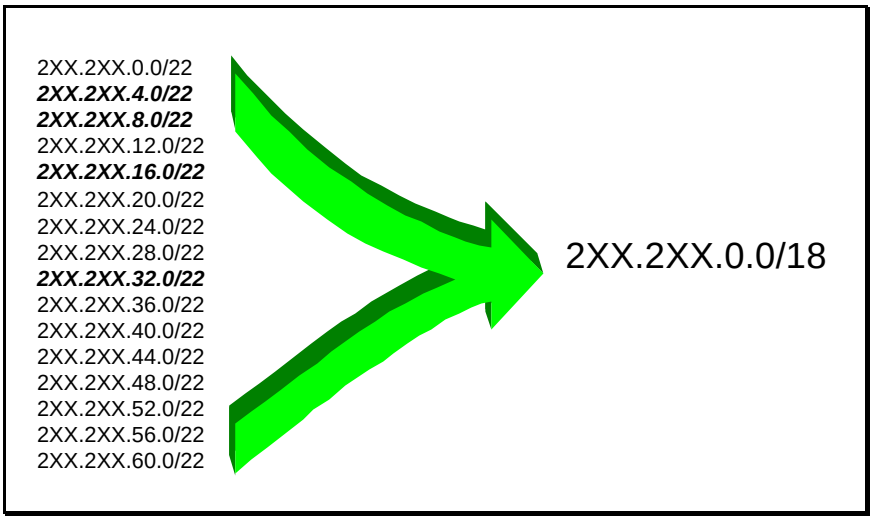

Figure 4 – Aggregating all potential /22s into a /18

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

A. Configure:

```
Router2(config)# router bgp 200
Router2(config)# aggregate-address 200.200.0.0 255.255.192.0
```

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

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

Checkpoint #2 : *call 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".*

12. **(Optional)** Use *debug ip bgp update* to see BGP update activity after clearing a BGP session.

13. Use the BGP aggregate-address command to summarise all advertisements between eBGP peers. There should only be the local iBGP prefixes and three other aggregated addresses from the other *ISPs* in the lab.

```
Router2(config)# router bgp 200
Router2(config)# aggregate-address 200.200.0.0 255.255.192.0 summary-only
```

Checkpoint #3 : *Call the lab assistant to verify the summarisation is working. 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 local specific prefixes (either 3 or 4 depending on the ISP) and 4 aggregate prefixes (one for each ISP).*

14. Examine the *origin* of the network prefixes. What is the origin type for the aggregated prefixes?

15. Remove the summarisation.

```
Router2(config)# router bgp 200
Router2(config)#no aggregate-address 200.200.0.0 255.255.192.0 summary-only
Router2(config)# aggregate-address 200.200.0.0 255.255.192.0
```

16. Examine the *origin* of the network prefixes. What is the origin type for the aggregated prefixes?

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?

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

CONFIGURATION NOTES

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