

Module 2 – Basic eBGP Configuration Lab

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

Prerequisites: Module 1

Topology :

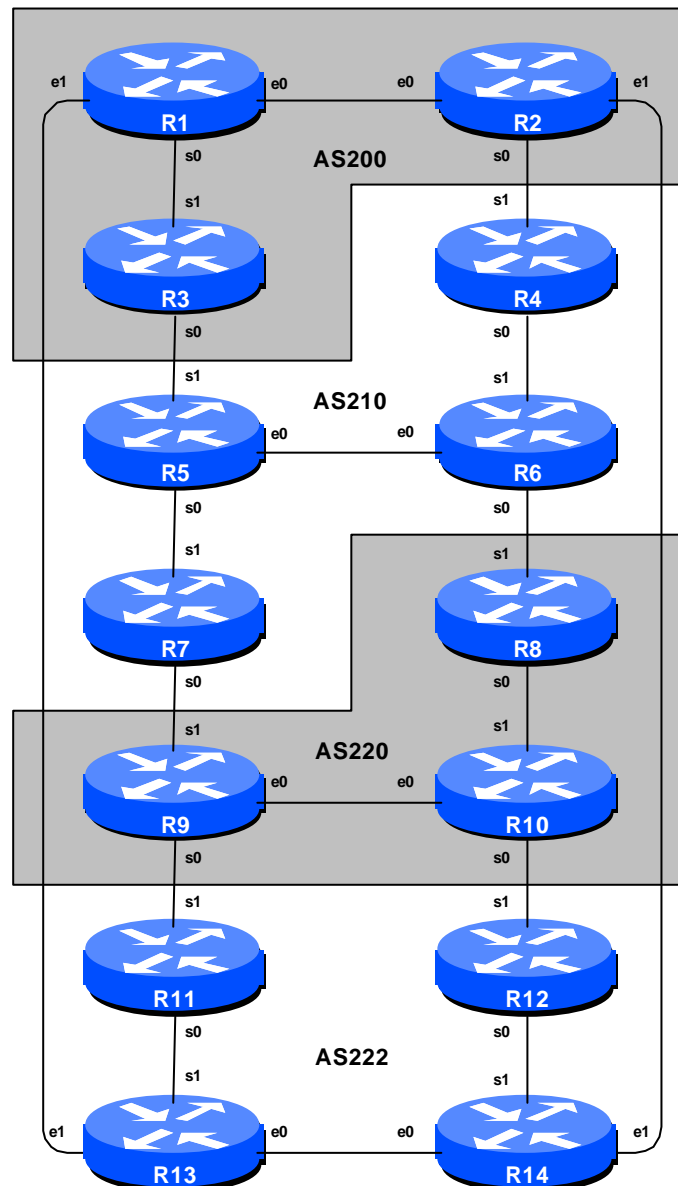


Figure 1 – BGP AS Numbers

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.

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

```
Router1 (config)# no router bgp 100
Router1 (config)# no router ospf 100
```

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

4. **Add an Autonomous System Number pointer.** Even though 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 Figure 1). Example:

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

5. **Configure OSPF on the routers within each AS.** In each AS configure OSPF routing. This means that networks for the links to each member in the AS must be entered in OSPF via the *network* statement. All the routers in an AS will be in the same OSPF *area 0* and use the same OSPF process ID. For example, Router Team one, with two interfaces in an AS would have the following:

```
Router1 (config)# router ospf 200
Router1 (config-router)# network 200.200.4.0 0.0.0.3 area 0      ! eth0/0
Router1 (config-router)# network 200.200.5.0 0.0.0.3 area 0      ! ser0/0
Router1 (config-router)# network 200.200.6.0 0.0.0.3 area 0      ! eth0/1
Router1 (config-router)# network 200.200.7.224 0.0.0.0 area 0    ! loop0
Router1 (config-router)# passive-interface loopback 0
Router1 (config-router)# passive-interface ethernet 0/1
Router1 (config-router)# log-adjacency-changes
```

Notes:

- *Passive-interface* makes sure that OSPF does not attempt to set up adjacencies on the DMZ networks. It **MUST** be used for every interface which is not running OSPF.
 - The number following “*router ospf*” is a process ID and is used inside the router only (so it can be any number). But for this lab we recommend the OSPF process ID be the same as the AS number (which is convention for a large number of ISPs).
6. **Ping Test.** Check the routes via OSPF. Make sure you can see all the networks within your AS. Ping all loopback interfaces within your AS. Use the “*show ip ospf neighbor*” and “*show ip route*” commands. If you cannot see the other routers in your AS, you will not be able to bring up BGP in the next steps.
 7. **Save the configuration.** Don’t forget to save the configuration to NVRAM!

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

Sunday, January 12, 2003

8. **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-router)# no synchronization
Router1 (config-router)# network 200.200.4.0 mask 255.255.252.0
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 200.200.11.224 description iBGP Link to R2
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)# neighbor 200.200.19.224 description iBGP Link to R3
Router1 (config-router)# no auto-summary
Router1 (config-router)# exit
Router1 (config)# ip route 200.200.4.0 255.255.252.0 Null0 250
```

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.

Q. Why do you need the BGP command *no auto-summary*?

A. By default BGP automatically summarises prefixes injected into BGP to the nearest **classful** boundary. As the Internet is classless, auto summarisation needs to be disabled. For example, if an AS was using addresses from the former class A address space, BGP would announce the /8 prefix if auto summarisation was still enabled. While this lab only uses address space from the former class C block and therefore the *no auto-summary* is not required, it is good practice to ensure that this command is part of all ISP BGP configurations.

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

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

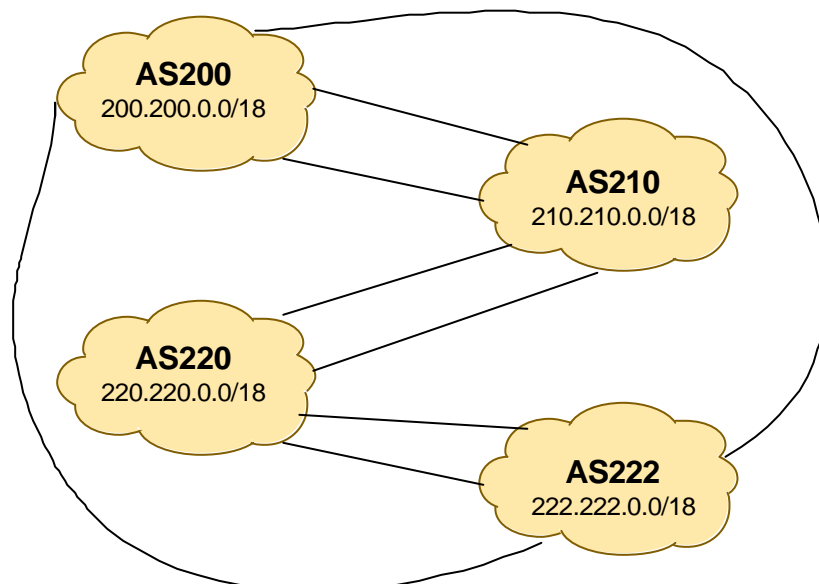


Figure 3 – Aggregates for each ASN

Sunday, January 12, 2003

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.)

11. Aggregate each AS'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.

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

Checkpoint #2 : 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).

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

14. Use the BGP aggregate-address command to summarise all advertisements between eBGP peers. There should only be the iBGP prefix generated by the local router (as the other routers in the AS will also have implemented the summary-only option), the summary aggregate from the other routers in the AS, and three 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 one local specific prefix and 4 aggregate prefixes (one for each ISP).

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

16. 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
```

17. 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?

Sunday, January 12, 2003

CONFIGURATION NOTES

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