

# Module 6(a) - iBGP and Basic eBGP

## Lab Topology

The following will be the topology used for the series of labs.

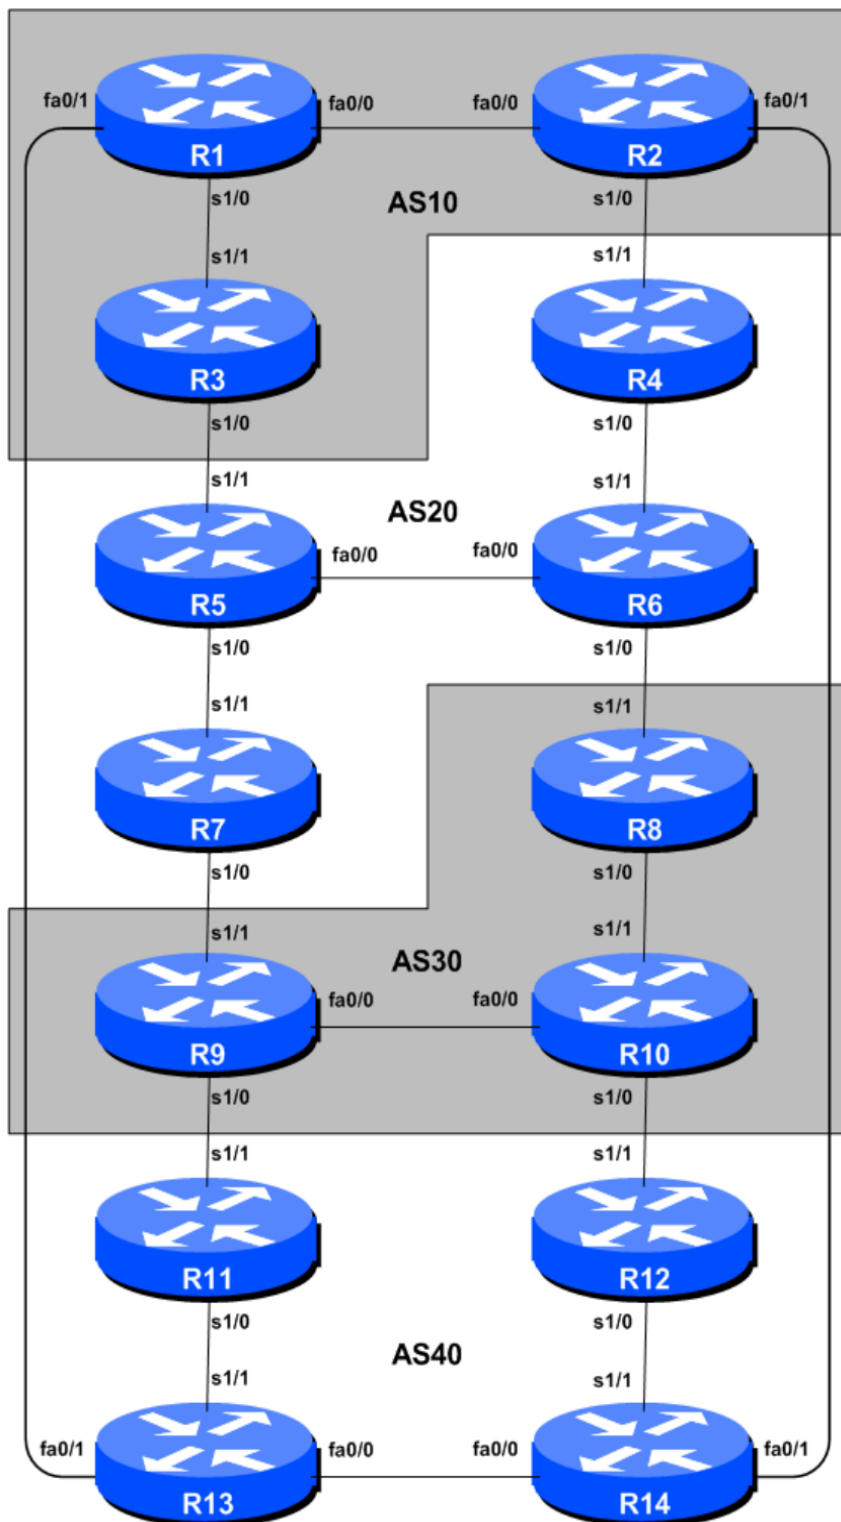



Figure 1 – BGP AS Numbers

## Lab Notes

The purpose of this module is to introduce external BGP (eBGP). This is the relationship between different autonomous systems on the "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.

## Lab Exercise

1. The relationship between the ASes is shown in Figure 2 (to relate to the "real" world).



**Figure 2 – AS relationship**

2. **Clean configuration.** The instructors will have erased the entire configuration of each router before starting this lab.
3. **IP Addressing.** Each AS gets their own /20 address space:

AS10 10.10.0.0/20	AS30 10.30.0.0/20
AS20 10.20.0.0/20	AS40 10.40.0.0/20

- 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`.
- Please **configure the addresses on each interface** which will be used for this module, and **check basic IP connectivity** with your immediately adjacent neighbours.

4. **Router Loopback Interface Addressing.** We have set aside a /27 for loopbacks even though each AS has either 3 or 4 routers in it – this leaves enough room for future expansion. The loopback address assignments which will be used for this module are below:

Router	Loopback	Router	Loopback
Router1	10.10.15.224	Router8	10.30.15.224
Router2	10.10.15.225	Router9	10.30.15.225
Router3	10.10.15.226	Router10	10.30.15.226
Router4	10.20.15.224	Router11	10.40.15.224
Router5	10.20.15.225	Router12	10.40.15.225
Router6	10.20.15.226	Router13	10.40.15.226
Router7	10.20.15.227	Router14	10.40.15.227

- **Configure the loopback** addresses for each router.
5. **Configure OSPF on the routers within each AS.** In each AS configure OSPF routing. This means setting up the OSPF process, marking internal interfaces as non-passive, and then configuring each internal interface and the loopback with the OSPF process. All the routers in an AS will be in the same OSPF area 0 and use the same OSPF process ID. OSPF 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 ospf` commands on external interfaces. A side effect of this is that external link addresses will not appear in the IGP (see the next section discussion iBGP deployment).

For example, Router 1, with two interfaces connecting to other routers in their AS, would have the following:

```
Router1(config)# router ospf 10
Router1(config-router)# passive-interface default
Router1(config-router)# no passive-interface fastethernet 0/0
Router1(config-router)# no passive-interface serial 1/0
Router1(config-router)# log-adjacency-changes
!
Router1(config)# interface fastethernet 0/0
Router1(config-interface)# ip ospf 10 area 0
!
Router1(config-interface)# interface serial 1/0
Router1(config-interface)# ip ospf 10 area 0
!
Router1(config-interface)# interface loopback 0
Router1(config-interface)# ip ospf 10 area 0
!
```

### Notes:

- `passive-interface default` makes sure that OSPF does not attempt to set up adjacencies on any interfaces apart from those specified by the `no passive-interface` commands.
- While the OSPF process ID can be any number, for this lab we recommend the OSPF process ID be the same as the AS number (which is the convention used by a number of ISPs).

6. **OSPF on Point-to-Point Ethernet Links.** We can modify OSPF'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 OSPF does not try and determine a designated or backup designated router.

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

```
Router1(config)# interface fastethernet 0/0
Router1(config-interface)# ip ospf network point-to-point
```

7. **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.
8. **Save the Configuration.**

```
Router1(config)#^Z
Router1# write memory
Building configuration...
[OK]
Router1#
```

9. **Final check.** Use the various `show ip ospf` commands to see the OSPF status of the lab network now. Check the routing and the routing table. If you are missing any adjacencies, work with your neighbouring routers in your AS to work out why, and what might have gone wrong with the neighbour authentication.
10. **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 AS (aggregate) for advertisement in BGP. Each router team should announce this /20 address block from their routers.

```
Router1(config)# router bgp 10
Router1(config-router)# distance bgp 200 200 200
Router1(config-router)# no synchronization
Router1(config-router)# network 10.10.0.0 mask 255.255.240.0
Router1(config-router)# neighbor 10.10.15.225 remote-as 10
Router1(config-router)# neighbor 10.10.15.225 update-source loopback 0
Router1(config-router)# neighbor 10.10.15.225 next-hop-self
Router1(config-router)# neighbor 10.10.15.225 description iBGP Link to R2
Router1(config-router)# neighbor 10.10.15.226 remote-as 10
Router1(config-router)# neighbor 10.10.15.226 update-source loopback 0
Router1(config-router)# neighbor 10.10.15.226 next-hop-self
Router1(config-router)# neighbor 10.10.15.226 description iBGP Link to R3
Router1(config-router)# no auto-summary
Router1(config-router)# exit
Router1(config)# ip route 10.10.0.0 255.255.240.0 Null0
```

#### Notes:

- **Next-hop-self configuration.** The next-hop-self configuration makes the iBGP speaking router use the iBGP source address (in this case the loopback) rather than the external next-hop address (as per the BGP specification). This is industry best practice and means that ISPs do not need to carry external next-hops in their IGP.
- `distance bgp <external-routes> <internal-routes> <local-routes>`. The IOS defaults are not optimised for Service Provider networks, so before we bring up BGP sessions, we should set the defaults that we require. The default distance for eBGP is 20, the default distance for iBGP is 200, and the default distance for OSPF is 120. This means that there is a potential for a prefix learned by eBGP to override the identical prefix carried by ISIS. Prefixes present in OSPF will never be found in BGP, and vice-versa.

- Discuss within your group why `update-source loopback 0` is necessary on iBGP?

11. **Test iBGP connectivity.** Use the following BGP commands to ensure you are receiving everyone's routes from within your AS.

```
show ip bgp summary : see a list of BGP peers that the router sees
```

```
show ip bgp/sh bgp ipv4 unicast : see a list of BGP paths that the router sees
```

```
show ip route : see all the routes that the router has installed
```

```
show ip bgp neighbor <neighbour-address> route : will show which routes you are receiving from your peer
```

```
show ip bgp neighbor <neighbour-address> advertised-routes : to see what you are exporting to the other router
```

12. **Configure Deterministic MED.** Another industry best practice is to configure deterministic MED for BGP. This means that IOS will order by AS Number the same prefix heard from multiple paths, and do the best path selection per ASN group. The IOS default is to compare the paths for the same prefix from most recent to the oldest, which can result in non-deterministic (ie different) path selection each time the path selection process is run.

For example, for Router5:

```
Router5(config)# router bgp 20
Router5(config-router)# bgp deterministic-med
```

13. **Configure eBGP peering.** Use Figure 1 to determine the links between the ASes. Addressing for eBGP session between 2 ASes 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 10
Router1(config-router)# neighbor 10.10.15.14 remote-as 40
Router1(config-router)# neighbor 10.10.15.14 description eBGP to Router13
```

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

14. **Adding a “customer” route into BGP.** 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 /26, for simplicity's sake.

Router	Customer	Router	Customer
Router1	10.10.0.0/26	Router8	10.30.0.0/26
Router2	10.10.0.64/26	Router9	10.30.0.64/26
Router3	10.10.0.128/26	Router10	10.30.0.128/26
Router4	10.20.0.0/26	Router11	10.40.0.0/26
Router5	10.20.0.64/26	Router12	10.40.0.64/26
Router6	10.20.0.128/26	Router13	10.40.0.128/26
Router7	10.20.0.192/26	Router14	10.40.0.192/26

Each team should now set up a static route pointing to the NULL0 interface for the /26 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:

```
Router8(config)# ip route 10.30.0.0 255.255.255.192 Null0
Router8(config)# router bgp 30
Router8(config-router)# network 10.30.0.0 mask 255.255.255.192
```

15. **Check the BGP table.** Are there routes seen via `show ip bgp` ? If not, why not? Once every team in the class has done their configuration, each team should see the aggregate (/20) from each AS (4 in total) as well as the 14 customer /26s introduced in the previous step. If this is not happening, work with your neighbours to fix the problem.
16. **Save the configuration.**

```
Router1(config)# ^Z
Router1# wr mem
```