

Module 1c – iBGP

Objective: Create a basic physical lab interconnection with IGP and one BGP AS number. Ensure that all routers, interfaces, cables and connections are working properly.

Prerequisites: Module 1a (OSPF) or Module 1b (ISIS).

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

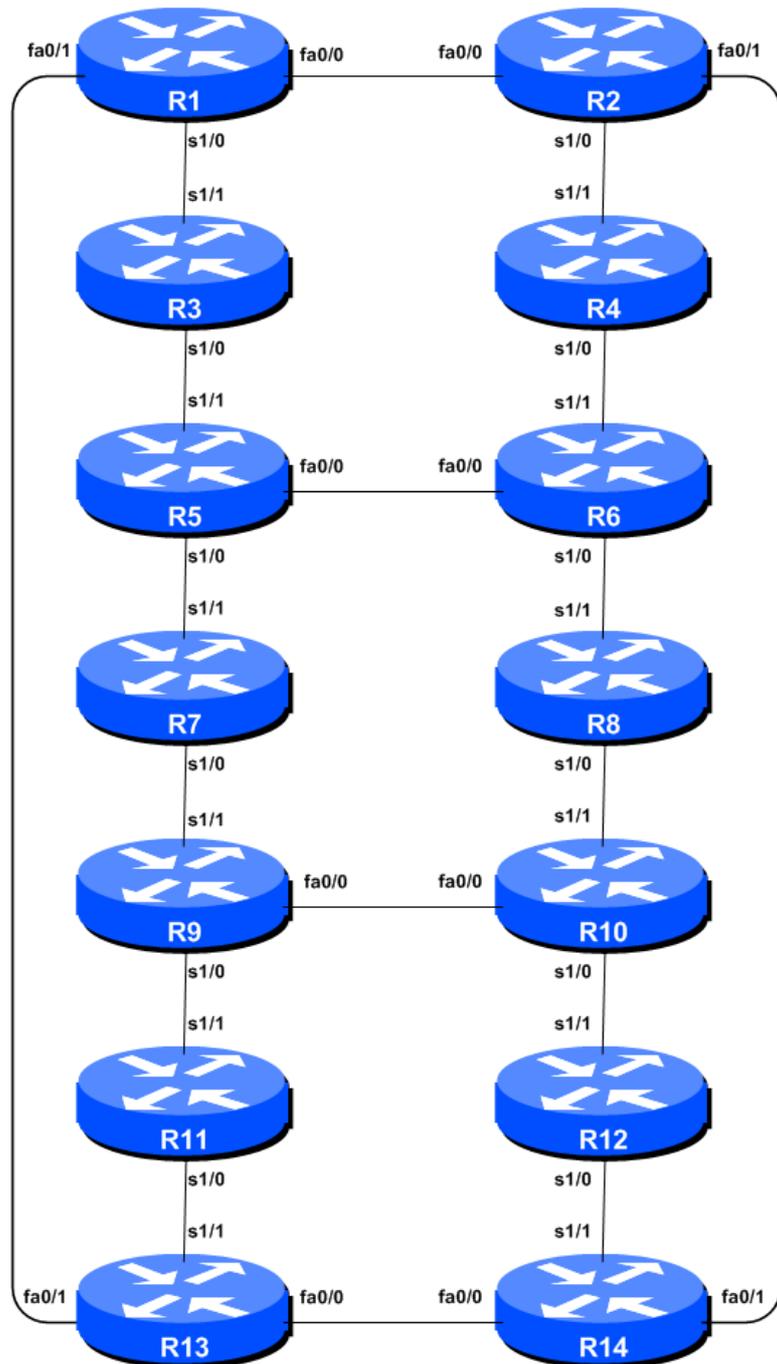


Figure 1 – ISP Lab Basic Configuration

Lab Notes

This lab continues from the previous one by adding iBGP to the configured and working IGP. Please refer to the IGP Module for further information and reference points about the purpose of Module 1.

Lab Exercise

1. **Setting up BGP.** We will be setting up BGP in our network using AS 10. To configure BGP on a Cisco router, we simply enter the following command, using Router2 as the example:

```
Router2(config)#router bgp 10
```

Cisco IOS assumes a default behaviour of all configured BGP neighbours exchanging IPv4 prefixes. The Internet today operates dual stack, supporting both IPv4 and IPv6. So we need to disable the default assumption that all neighbours will exchange IPv4 prefixes, leaving that determination to the operator of the network. The command we need is:

```
Router2(config)#no bgp default ipv4-unicast
```

From now on, as this is an IPv4 BGP lab, we need to enter the IPv4 address family configuration for any configuration. IOS will work out where the configuration belongs in the command line interface.

2. **Configuring iBGP.** Before we set up iBGP with our neighbours in our AS, we need to do some basic preparation on the router. 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 ISIS is 115. This means that there is a potential for a prefix learned by eBGP to override the identical prefix carried by ISIS. Recall from the Routing presentation that there is a distinct separation between BGP and ISIS processes – prefixes present in ISIS will never be found in BGP, and vice-versa. To protect against accidents¹, the eBGP distance is set to 200 also. The command to do this is the bgp distance subcommand, syntax is:

```
distance bgp <external-routes> <internal-routes> <local-routes>
```

Note: This should be included in all future BGP configurations in this workshop. For example, for Router2, the configuration might be:

```
Router2(config)# router bgp 10
Router2(config-router)# address-family ipv4
Router2(config-router-af)# distance bgp 200 200 200
```

3. **Logging of BGP Adjacency State.** Make sure that the logging of BGP neighbour changes is activated. This is so that a notification is generated every time the state of a BGP neighbour changes state, and is useful for debugging purposes:

```
Router2(config)#router bgp 10
Router2(config-router)# bgp log-neighbor-changes
```

¹ There have been several incidents in the past where denial of service attacks on ISP networks have been successful because ISPs have omitted basic routing protocol security. Setting the BGP distances to be greater than any IGP is one of the mitigation methods available.

- 4. Configuring iBGP Neighbours.** All Routers will be in the same Autonomous System (AS) 10 for this first lab. Use the *show ip bgp summary* to check the peering. The BGP peering will be established using the loopback interfaces' IP address.

```

Router2(config)#router bgp 10
Router2(config-router)# address-family ipv4
Router2 (config-router-af)#neighbor 10.0.15.241 remote-as 10
Router2 (config-router-af)#neighbor 10.0.15.241 update-source loopback 0
Router2 (config-router-af)#neighbor 10.0.15.241 description iBGP with Router1
Router2 (config-router-af)#
Router2 (config-router-af)#neighbor 10.0.15.243 remote-as 10
Router2 (config-router-af)#neighbor 10.0.15.243 update-source loopback 0
Router2 (config-router-af)#neighbor 10.0.15.243 description iBGP with Router3
Router2 (config-router-af)#
Router2 (config-router-af)#neighbor 10.0.15.244 remote-as 10
Router2 (config-router-af)#neighbor 10.0.15.244 update-source loopback 0
Router2 (config-router-af)#neighbor 10.0.15.244 description iBGP with Router4
Router2 (config-router-af)#
Router2 (config-router-af)#neighbor 10.0.15.245 remote-as 10
Router2 (config-router-af)#neighbor 10.0.15.245 update-source loopback 0
Router2 (config-router-af)#neighbor 10.0.15.245 description iBGP with Router5
Router2 (config-router-af)#
Router2 (config-router-af)#neighbor 10.0.15.246 remote-as 10
Router2 (config-router-af)#neighbor 10.0.15.246 update-source loopback 0
Router2 (config-router-af)#neighbor 10.0.15.246 description iBGP with Router6
Router2 (config-router-af)#
Router2 (config-router-af)#neighbor 10.0.15.247 remote-as 10
Router2 (config-router-af)#neighbor 10.0.15.247 update-source loopback 0
Router2 (config-router-af)#neighbor 10.0.15.247 description iBGP with Router7
Router2 (config-router-af)#
Router2 (config-router-af)#neighbor 10.0.15.248 remote-as 10
Router2 (config-router-af)#neighbor 10.0.15.248 update-source loopback 0
Router2 (config-router-af)#neighbor 10.0.15.248 description iBGP with Router8
Router2 (config-router-af)#
Router2 (config-router-af)#neighbor 10.0.15.249 remote-as 10
Router2 (config-router-af)#neighbor 10.0.15.249 update-source loopback 0
Router2 (config-router-af)#neighbor 10.0.15.249 description iBGP with Router9
Router2 (config-router-af)#
Router2 (config-router-af)#neighbor 10.0.15.250 remote-as 10
Router2 (config-router-af)#neighbor 10.0.15.250 update-source loopback 0
Router2 (config-router-af)#neighbor 10.0.15.250 description iBGP with Router10
Router2 (config-router-af)#
Router2 (config-router-af)#neighbor 10.0.15.251 remote-as 10
Router2 (config-router-af)#neighbor 10.0.15.251 update-source loopback 0
Router2 (config-router-af)#neighbor 10.0.15.251 description iBGP with Router11
Router2 (config-router-af)#
Router2 (config-router-af)#neighbor 10.0.15.252 remote-as 10
Router2 (config-router-af)#neighbor 10.0.15.252 update-source loopback 0
Router2 (config-router-af)#neighbor 10.0.15.252 description iBGP with Router12
Router2 (config-router-af)#
Router2 (config-router-af)#neighbor 10.0.15.253 remote-as 10
Router2 (config-router-af)#neighbor 10.0.15.253 update-source loopback 0
Router2 (config-router-af)#neighbor 10.0.15.253 description iBGP with Router13
Router2 (config-router-af)#
Router2 (config-router-af)#neighbor 10.0.15.254 remote-as 10
Router2 (config-router-af)#neighbor 10.0.15.254 update-source loopback 0
Router2 (config-router-af)#neighbor 10.0.15.254 description iBGP with Router14

```

Q. Why is *update-source loopback 0* necessary on iBGP?

Use *show ip bgp summary* to check the status of the iBGP neighbour connections. If the iBGP session is not up and/or no updates are being sent, work with the Router Team for that neighbour connection to troubleshoot the problem.

5. Sanity Check. Remember to use the following commands to ensure you are getting the information you are suppose to be getting:

```
show ip bgp summary      : see a list of BGP peers that the router sees
show ip bgp              : see a list of BGP paths that the router sees
show ip route            : see all the routes that the router has installed
```

Q. Are there routes seen via *show ip bgp*? If not, why not? Are there any routes tagged "B" when you do a *show ip route*?

6. Add Networks via BGP. Each Router Team will use BGP to advertise the address block used for the Module. For example, Router Team 1 would add:

```
Router1 (config)#router bgp 10
Router1 (config-router)# address-family ipv4
Router1 (config-router-af)#network 10.0.0.0 mask 255.255.240.0
```

Use *show ip bgp* on neighbour's router to see if you are advertising your network via BGP.

Q. Does the network show up via BGP? If not, why?

Enter a static route for the CIDR block. For example, Router 1 would use:

```
Router1 (config)#ip route 10.0.0.0 255.255.240.0 Null0
```

Q. Does the network show up via a neighbour's BGP? Use the command *show ip bgp neighbor <neighbour's IP address> advertised-routes* to see what you are exporting to the other router. Physically go to one of your neighbour's routers and check their BGP Table. Explain what you see.

Q. Does the network appear in the router's forwarding table? Use the command *show ip route* to check the local forwarding table. If not, why not?

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

7. Adding a "customer" route into BGP. We are now going to add a "customer" route into BGP on each router. Now in the lab we don't have any "customers" as such connected to our routers, so we are going to simulate the connectivity by simply using a Null0 interface.

In real life, the configuration to connect a customer would look something like this.

```
ip route 172.16.4.0 255.255.255.128 Serial 0/5/2 permanent
!
router bgp 64509
 network 172.16.4.0 mask 255.255.255.128
!
```

This would add a static route pointing 172.16.4.0/25 towards Serial 0/5/2 – the latter interface would be a fixed link connecting to the customer site. 172.16.4.0/25 would be the address space that the ISP had assigned to the customer. The BGP network statement would then add the customer address block into the ISP’s iBGP.

Note: the **permanent** keyword ensures that the static route is always in the routing table, even if the interface physically goes down. Many ISPs use this to ensure they don’t have iBGP churn when their customer links go down.

The “customer” address space that each router team will introduce into the iBGP is listed below – we will each use a /26, for simplicity’s sake.

R1	10.0.0.0/26	R8	10.0.1.192/26
R2	10.0.0.64/26	R9	10.0.2.0/26
R3	10.0.0.128/26	R10	10.0.2.64/26
R4	10.0.0.192/26	R11	10.0.2.128/26
R5	10.0.1.0/26	R12	10.0.2.192/26
R6	10.0.1.64/26	R13	10.0.3.0/26
R7	10.0.1.128/26	R14	10.0.3.64/26

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 Router11:

```
Router11 (config)# ip route 10.0.2.128 255.255.255.192 Null0
Router11 (config)# router bgp 10
Router11 (config-router)# address-family ipv4
Router11 (config-router-af)# network 10.0.2.128 mask 255.255.255.192
```

- 8. 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 as well as the fourteen /26s introduced in the previous step. If this is not happening, work with your neighbours to fix the problem.

Checkpoint #2: call the lab assistant to demonstrate the current BGP table.

9. **Other Features in iBGP.** Review the documentation or use command line help by typing ? to see other *show* commands and other BGP configuration features.
10. **Advanced Configuration.** Those router teams who have completed this module should refer to Module 11 of the Advanced BGP Workshop. The set-up steps have been extended to include all the basic requirements of a router being used in an ISP backbone. While waiting for the module to complete, now would be a good time to review the advanced Module and incorporate the additions to the configuration used here.

Review Questions

1. What IOS show command(s) will display the router's BGP route table?
2. Why change the eBGP distance from the default 20 to 200?
3. Why is the static pull up route necessary in relation to inserting a prefix into iBGP?