

Module 2 – More iBGP, and Basic eBGP Configuration

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

Prerequisites: Module 1

Topology :

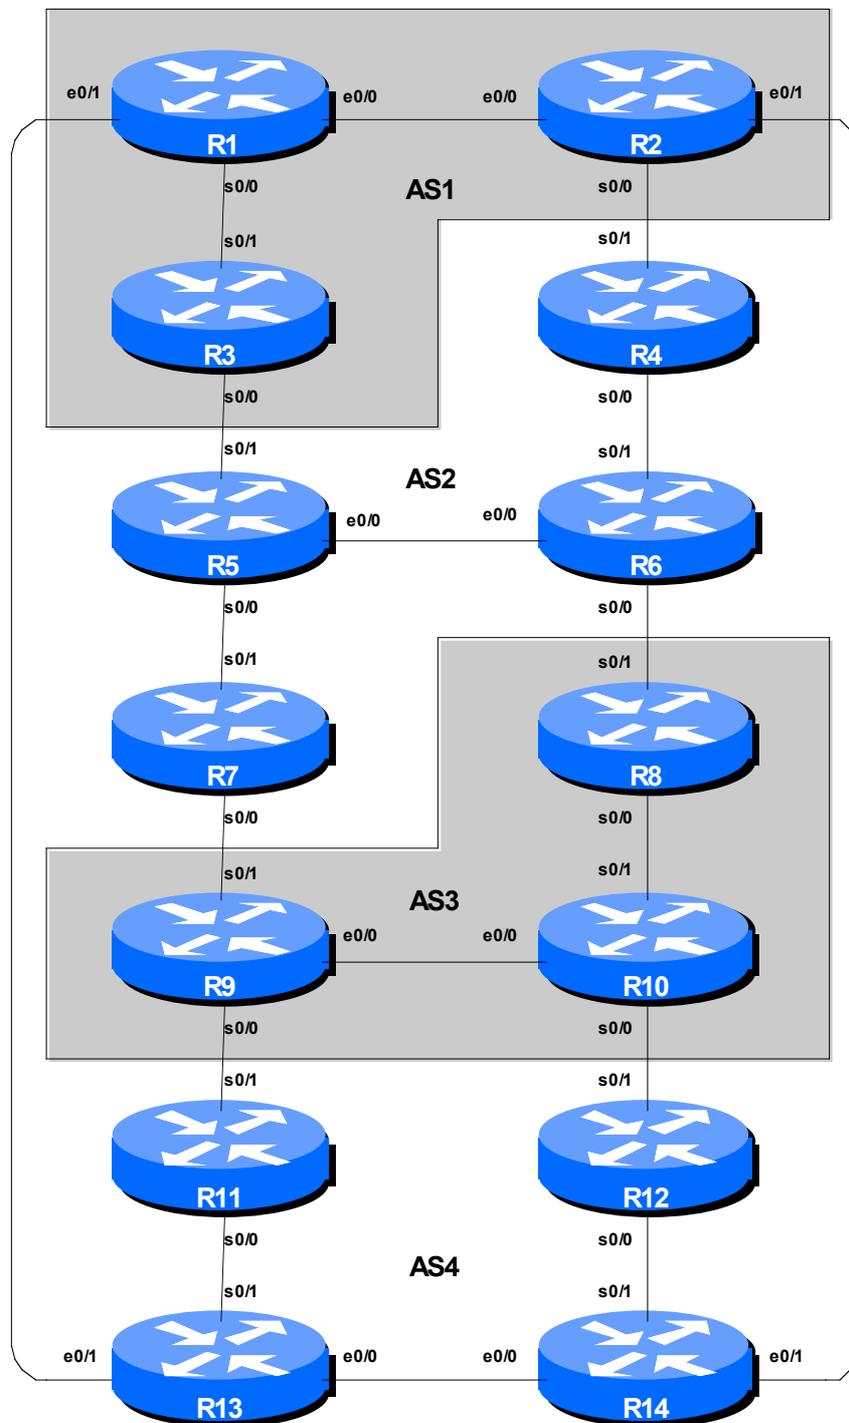


Figure 1 – BGP AS Numbers

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”.

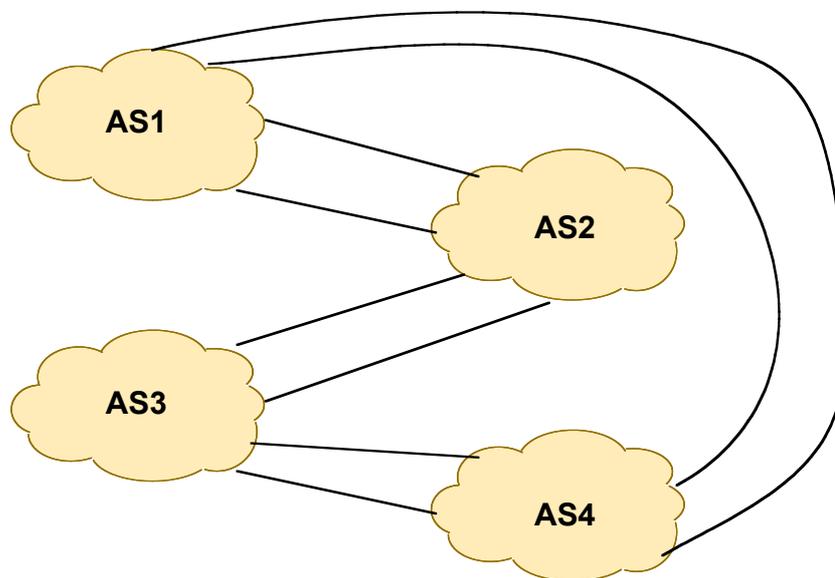


Figure 2 – AS relationship

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

```
Router1 (config)# no router bgp 10
Router1 (config)# no router ospf 41
```

This will clear the BGP and OSPF 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 configuration you saved at the end of Step 11, just prior to starting the OSPF configuration. Which

you probably forgot to do, even with both the lab instructor and these written notes telling you to do so.)

- 4. 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 1, with two interfaces connecting to other routers in their AS, would have the following:

```
Router1 (config)# router ospf 1
Router1 (config-router)# network 100.1.0.0 0.0.0.3 area 0      ! eth0/0
Router1 (config-router)# network 100.1.1.0 0.0.0.3 area 0      ! ser0/0
Router1 (config-router)# network 100.1.2.0 0.0.0.3 area 0      ! eth0/1
Router1 (config-router)# network 100.1.15.224 0.0.0.0 area 0    ! loop0
Router1 (config-router)# passive-interface default
Router1 (config-router)# no passive-interface ethernet 0/0
Router1 (config-router)# no passive-interface serial 0/0
Router1 (config-router)# log-adjacency-changes
```

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.
 - 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 the convention used by a number of ISPs).
- 5. 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.
- 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
```

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 **classfull** boundary. As the Internet has been using classless routing protocols since 1994, auto summarisation needs to be disabled. For example, as this AS is using addresses from the former class A address space, BGP would announce the /8 prefix if auto summarisation was still enabled. Even though many ISPs are still being allocated address space from the former class C block and therefore the *no auto-summary* may not be required in their networks, it is still good practice to ensure that this command is part of all ISP BGP configurations.

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.

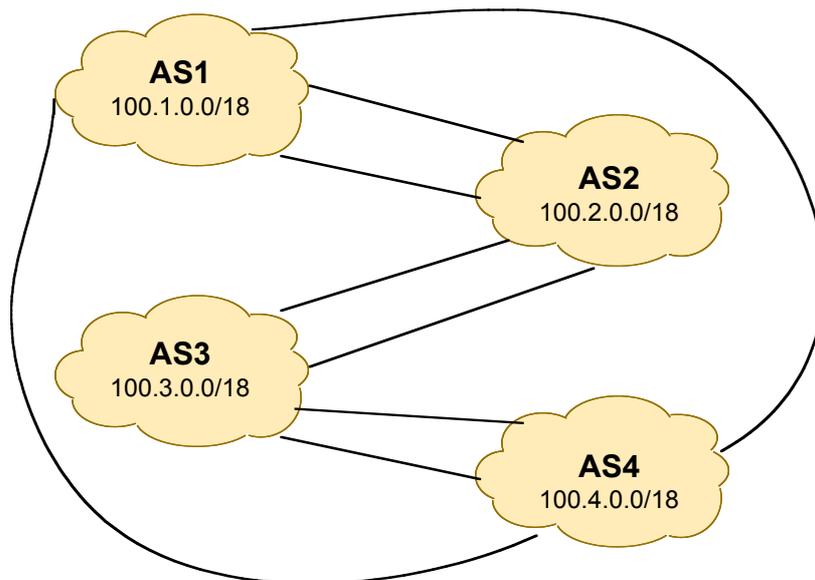


Figure 3 – Aggregates for each ASN

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.

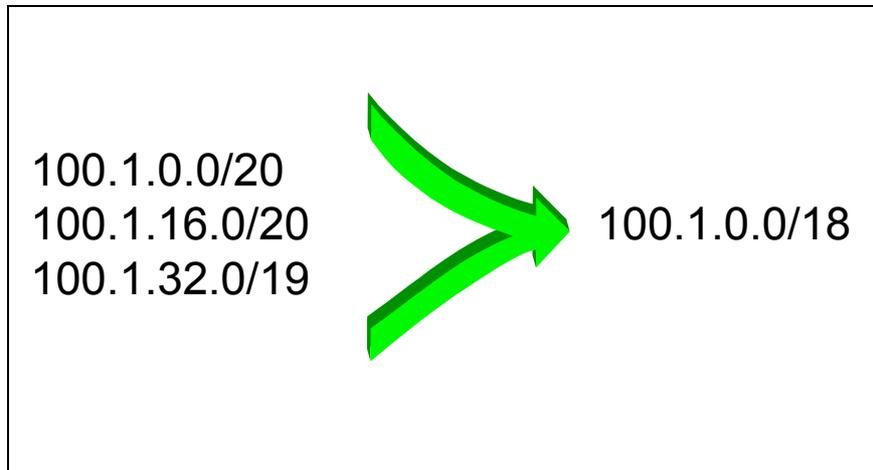


Figure 4 – Aggregating AS1 address space into a /18

13. Examine the *origin* of the network prefixes. What is the origin type for the aggregated prefixes? Write your answer here:

14. 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).

15. 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?

Sunday, July 17, 2005

CONFIGURATION NOTES

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