More iBGP and Basic eBGP

- Objective: Connect your ISP to a Transit provider and the Internet Exchange Point using a combination of ISIS, internal BGP, and external BGP.
- Prerequisites: ISIS and iBGP modules

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 Figure 1.

Lab Exercises

1. Turning on neighbour authentication for ISIS – Part 1

ISIS supports neighbour authentication; this is considered more and more important inside ISP networks as attacks on infrastructure increase and ISPs seek to use all available tools to secure their networks. (While an attack on ISIS is harder as it runs on the link layer alongside IP rather than on top of IP like OSPF, some ISPs are still prudent and implement neighbour authentication.)

Each router team will now turn on neighbour authentication for ISIS. The first step is to set up the keychain to be used – we will use the key “cisco” for this lab:

```
Router1(config)# key chain lab-key
Router1(config-keychain)# key 1
Router1(config-keychain-key)# key-string cisco
```

2. Turning on neighbour authentication for ISIS – Part 2

Now that the keychain has been defined, we activate authentication within the ISIS processes. The first step is to enable MD5 for level-2 IS’s:

```
Router1(config)# router isis workshop
Router1(config-router)# authentication mode md5 level-2
```

And then associate the key-chain we defined earlier with the configured authentication:

```
Router1(config-router)# authentication key-chain lab-key level-2
```

Notice now that the ISIS adjacencies do not come up unless the neighbouring router has also entered the same configuration and key. Notice also how the ISIS adjacencies were reset as the configuration was entered – security is being introduced, so the adjacencies are reset.

3. Final ISIS check

Use the various *"show isis"* commands to see the ISIS status of the lab network now. Check the routing and the routing table. Make sure all the adjacencies have come back up again. If any adjacency has failed to come up, and you see several log messages saying:

```
*Mar 1 00:05:17.825: %CLNS-4-AUTH_FAIL: ISIS: LAN IIH authentication failed
```

you should reasonably expect that either you or your connected neighbour have forgotten to set up neighbour authentication.

Note: Wherever an ISIS session is configured from now on in the workshop, all Router Teams MUST use passwords on these ISIS sessions.

Checkpoint #2 : *call the lab assistant to verify the connectivity.*

4. Test internal BGP connectivity.

Remember that the loopback addresses are used as the end points of the BGP connections. If IS-IS is not working then the BGP sessions can't be established.

Use the BGP Show commands to ensure you are still receiving routes from within your AS.

5. 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 rs2-isp5-ixp2:

```
rs2-isp5-ixp2 (config)# router bgp 64025
rs2-isp5-ixp2 (config-router)# bgp deterministic-med
```

Note that it is unlikely that deterministic MED will have any impact on the path selection for this Module. However, it is an industry best practice now, and network operators should include it in their BGP configuration template by default.

6. 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 rs1-isp1-ixp1's peering with rs2-isp1-ixp1, with "cisco" used as the password:

```
rs1-isp1-ixp1 (config)# router bgp 64011
rs1-isp1-ixp1 (config-router)# neighbor 100.66.12.252 password cisco
```

NOTE: Don't use a password like this on a production server!

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 #3: *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.*

7. eBGP peering information.

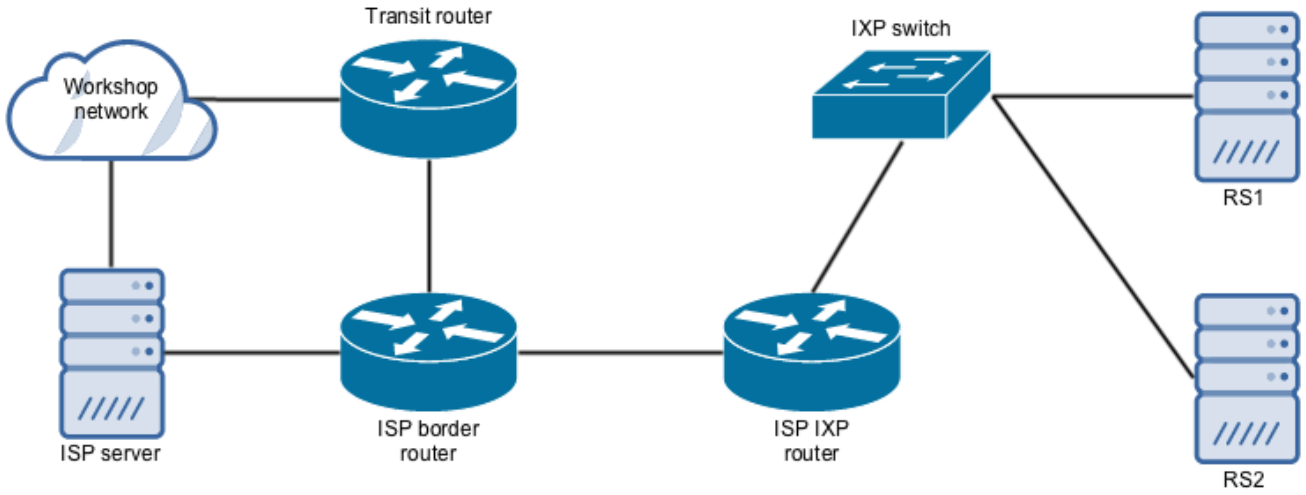

The diagram above shows the connections from your ISP routers to your Transit provider and to your IXP.

IXP1	IPv4	IPv6	AS number
Transit	100.66.101.254	FD90:100:66:101::254	64110
rs1	192.168.1.8	FD90:192:168:1::8	64010
rs2	192.168.1.9	FD90:192:168:1::9	64010
IXP2	**IPv4**	**IPv6**	**AS number**
Transit	100.66.102.254	FD90:100:66:102::254	64120
rs1	192.168.2.8	FD90:192:168:2::8	64020
rs2	192.168.2.9	FD90:192:168:2::9	64020

These tables show the relevant BGP parameters for the Transit and IXP BGP sessions. We'll configure the IPv4 sessions first and add the IPv6 connections later.

Q. When you make a BGP connection to the Transit provider which IPv4 address will you use as the **neighbor**? Which AS number will you use?

A.

```
IPv4 address:
AS number:
```

Q. When you make a BGP connection to the IXP route servers which IPv4 addresses will you use as the **neighbor**? Which AS number will you use?

A.

```
IPv4 addresses:
AS number:
```

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.)

8. Configure Transit peering.

The eBGP configuration on the Transit router is already configured. You need to configure your side of the connection to complete the setup.

For example on r2-isp4-ixp2:

```
router bgp 64024
  neighbor 100.66.102.254 remote-as 64120
  neighbor 100.66.102.254 description TRANSIT-V4

  address-family ipv4
    neighbor 100.66.102.254 activate
  exit-address-family
```

Now use **show ip bgp summary** to check your BGP session with the transit provider.

Q. How many prefixes are you learning from the Transit provider?

A. It will depend on whether you are the first person to get things working! If other people have things setup then you will see their routes.

Q. How would you look at any routes you are receiving?

A. show ip bgp and **show ip routes**

Checkpoint #1. *Call the lab assistant to verify the connectivity.*

9. Configure IXP peering.

The Internet Exchange Point offers a Route Server service. There are two Linux servers running the BIRD Routing Daemon software. The BIRD package talks a number of routing protocols but in this instance we will only be using BGP part of it. We run two servers to provide redundancy.

If each ISP peers with these servers then we will each be able to retrieve copies of each others routing announcements. Because we are all on the same Layer 2 switch we can get the routing information from the route servers but send the traffic directly from, for example, ISP1 to ISP4. Let's configure the connections and see how this works.

As we have two connections, we'll build one and then use it as model for the second. For example, on r1-isp5-ixp1:

```
router bgp 64015
  no bgp enforce-first-as
  neighbor 192.168.1.8 remote-as 64010
  neighbor 192.168.1.8 description ROUTE-SERVERS-V4 - rs1

  address-family ipv4
    neighbor 192.168.1.8 activate
  exit-address-family
```

This would add a peering session to **rs1** on **IXP1**.

Check that this is working using the commands you have already used for previous connections.

Now we can add the BGP session to **rs2**:

router bgp 64015

```
  neighbor 192.168.1.9 remote-as 64010
  neighbor 192.168.1.9 description ROUTE-SERVERS-V4 - rs2

  address-family ipv4
    neighbor 192.168.1.9 activate
  exit-address-family
```

Check that this session is working as well. As before the number of prefixes you see will depend on who else is peering with the route servers.

Q. Why did we add the command **no bgp enforce-first-as**?

A. A route-server mirrors routes from other eBGP peers in order to avoid a full mesh eBGP peer on the IX. It has an ASN (sometimes private, sometimes public) because it is mandatory to establish the session, but it does not modify anything in the route : the NEXT_HOP remains the same, and the AS_PATH is not changed at all.

By default, Cisco routers do not like this behaviour, and you have to disable the check of the AS_PATH in the global BGP configuration (**it will be disabled for all your peers**) with the command :

```
Router(config-router)# no bgp enforce-first-as
```

10. Configure passwords on the eBGP sessions.

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. You should use the example we had in the previous section as a template.

There is a problem however. In the previous section, you controlled both routers and could make changes on both.

Checkpoint #4: *Call the lab assistant and work with them to set the password as set on the eBGP sessions with the Transit router and the IXP route servers. Once confirmed by the lab assistant, move on to the next steps.*

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.

11. 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 /23 aggregate and a /26 *Customer* prefix from each AS . If this is not happening, work with your neighbours to fix the problem.

Checkpoint #5: *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 7 aggregate prefixes (one for each ISP) in the BGP table.*

12. The Importance of Aggregation.

Each AS was allocated a /23 address block. 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 (as we have done here) – but most operators consider leaking this subdivided address space out to the Internet at large antisocial and unfriendly.

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. The "aggregate-address" command is quite often used to achieve this. We are not doing any filtering or limitation of the announcements of the "customer" address blocks we have introduced into each ASN. This will be one of the goals of the next modules in the workshop.

13. 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?