eBGP using 4-byte ASNs

Configuring the link to the Transit Provider

The next step is to set up eBGP with the Transit Provider. The lab instructors will be running the routers for the two Transit Providers and will have already configured them.

Physical Link

Follow the guidelines in the [IP Address Plan](address-plan) document to configure the link to the upstream. Make sure that you can ping the upstream's router using both IPv4 and IPv6 - if it doesn't work, investigate why.

Here is an example configuration for AS70006:

```
interface fastethernet 0/0
 description Link to TR2
 ip address 100.122.1.10 255.255.255.252
 ipv6 address 2001:19:0:12::1/127
!
```

OSPF

Do not configure OSPFv2 or OSPFv3 towards the upstream provider! They are not part of your autonomous system.

External BGP

We now configure eBGP with the upstream. Again, the configuration on the two transit routers will have already been completed by the instructors, so once configured, the eBGP session should just come up and work.

Here is a configuration example for AS105:

```
router bgp 105
 address-family ipv4
  neighbor 100.122.1.5 remote-as 122
  neighbor 100.122.1.5 description eBGP with TRANSIT 2
!
 address-family ipv6
  neighbor 2001:19:0:11:: remote-as 122
  neighbor 2001:19:0:11:: description eBGP with TRANSIT 2
!
```

This set up is quite simple because the routers in AS105 and AS122 are not using 4-byte ASes and are running software which does not support 4-byte ASNs (for Cisco, this is IOS before 15.0 release).

External BGP with 2-byte and 4-byte ASNs

The routers in AS70002, AS70004 and AS70006 are running a more recent version of Cisco IOS. The transit router in AS70121 is also running a more recent version of IOS. Cisco's IOS gained support of 4 byte AS numbers in 2010, and all modern routers today now support 4-byte AS numbers.

However, there are still a few network operators who are running equipment with operating systems which are more than 5 years old (this is an incredibly bad idea in itself, at least from a security perspective).

It is quite simple to configure BGP to support routers which cannot handle the extended 4-byte AS number range. The BGP sessions we will set up between AS70121 and AS101/AS103, and AS122 and AS70004/AS70006 need to configure this compatibility mode.

Let's work through an example to show what needs to be done.

The router running the really old software which does not support 4-byte ASNs obviously cannot configure the extended 4-byte AS in the eBGP session. Instead, the BGP session is configured to peer with the remote AS23456, the transition AS. Here is an example for the AS101 router:

```
router bgp 101
 address-family ipv4
  neighbor 100.121.1.1 remote-as 23456
  neighbor 100.121.1.1 description eBGP with TRANSIT 1 (AS70121)
!
 address-family ipv6
  neighbor 2001:18:0:10:: remote-as 23456
  neighbor 2001:18:0:10:: description eBGP with TRANSIT 1 (AS70121)
!
```

Note the description field - we have documented that the actual AS is AS70121 just so we have documented on the router which 4-byte ASN peer we have connected to.

On the other side, the AS70121 router configures BGP as normal, talking to AS101. Here is the configuration example for TR1:

```
router bgp 70121
 address-family ipv4
  neighbor 100.121.1.2 remote-as 101
  neighbor 100.121.1.2 description eBGP with AS101
!
 address-family ipv6
  neighbor 2001:18:0:10::1 remote-as 101
  neighbor 2001:18:0:10::1 description eBGP with AS101
!
```

Once both ends are configured check and see how the BGP session has been established - for both

IPv4 and IPv6.

Run the *show ip bgp* and *show bgp ipv6 unicast* commands to see what the BGP table looks like now. Try it on AS101 router. Then try on a router which supports 4-byte ASNs (eg AS70102 router). Do you see the differences?

Status Check

Once this has been configured, you should now see routes coming from the upstream provider, and you should be able to see your aggregate being sent to your upstream. The commands to see what you are receiving are:

```
show ip bgp neigh 100.121.1.1 routes
show bgp ipv6 uni neigh 2001:18:0:10:: routes
```

and to show what you are sending:

```
show ip bgp neigh 100.121.1.1 advertised-routes
show bgp ipv6 uni neigh 2001:18:0:10:: advertised-routes
```

Confirmation

Check on the Core and Peering Router what you now see in the BGP table. Are there differences? Can you explain what they are, and why?

What does:

```
show ip bgp 0.0.0.0 0.0.0.0
show ip route 0.0.0.0 0.0.0.0
show ipv6 route ::/0
show bgp ipv6 uni ::/0
```

show you on the core router and on the peering router?

You will notice that the default route is being propagated by BGP throughout the AS.

- On the core router, it will complain of a **RIB failure** when you look at the BGP table, because OSPF has a default route with a lower protocol distance.

- On the border router, it will complain of a **RIB failure** when you look at the BGP table, because there is a static default route to the Null interface.

While there is nothing wrong with a **RIB failure**, we can just remove the default from being propagated by the iBGP. To do this, we go back to the eBGP session, look for the default, tag it with the *no-advertise* community, and then the border router will no longer announce the default by iBGP.

Here is a configuration example for AS70002:

```
ip prefix-list default-route permit 0.0.0.0/0
!
route-map tag-default permit 10
 match ip address prefix-list default-route
 set community no-advertise
!
route-map tag-default permit 20
!
ipv6 prefix-list default-v6route permit ::/0
!
route-map tag-v6default permit 10
 match ipv6 address prefix-list default-v6route
 set community no-advertise
!
route-map tag-v6default permit 20
!
router bgp 70002
 address-family ipv4
  neighbor 100.121.1.5 route-map tag-default in
 address-family ipv6
  neighbor 2001:18:0:11:: route-map tag-v6default in
!
```

To confirm, has the default route now disappeared from the BGP table on the Core and Peering routers? If not, did you remember to do route-refresh after you applied the above configuration?

Testing

Once this has been completed, test the connectivity. Can you reach the other groups in the class? You should be able to ping all the IPv4 and IPv6 loopbacks across the whole classroom.

Can you see the Internet too? The lab has IPv4 connectivity to the Internet - check that this works by trying a few pings or trace routes to well known destinations (e.g. to 8.8.8.8).

From:
https://www.bgp4all.com.au/pfs/ - **Philip Smith's Internet Development Site**

Permanent link:
https://www.bgp4all.com.au/pfs/training/cnnic/2-4byte

Last update: **2016/10/23 05:50**