6DISS

South America Workshop

WALC 2006 (Quito, Ecuador – 26-28 July '06)

Multiprotocol BGP configuration

Pedro Lorga (lorga@fccn.pt)

Miguel Baptista (miguel.baptista@fccn.pt)

Simon Muyal (muyal@renater.pt)

Laboratory Exercise: *Multiprotocol BGP configuration*

Objectives

In this laboratory exercise you will complete the following tasks:

- *Establish IPv6 unicast BGP sessions;*
- *Check the exchanged routing information and perform application-level tests towards addresses received from the BGP sessions;*
- *Filter routes announced and received through peerings.*

Visual Objective

The following figure shows the configuration of the current laboratory

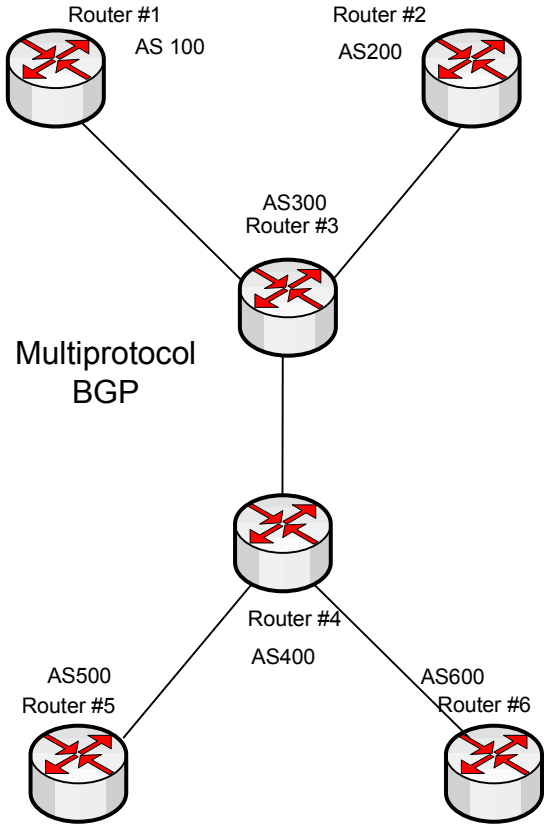

Figure 1 - Multiprotocol BGP

Setup/Scenario

There will be 4 students per router. The distribution will be the same as in the previous exercise.

All routers will run an eBGP process and each one will belong to a different AS, peering with two or three neighbours, according to the topology.

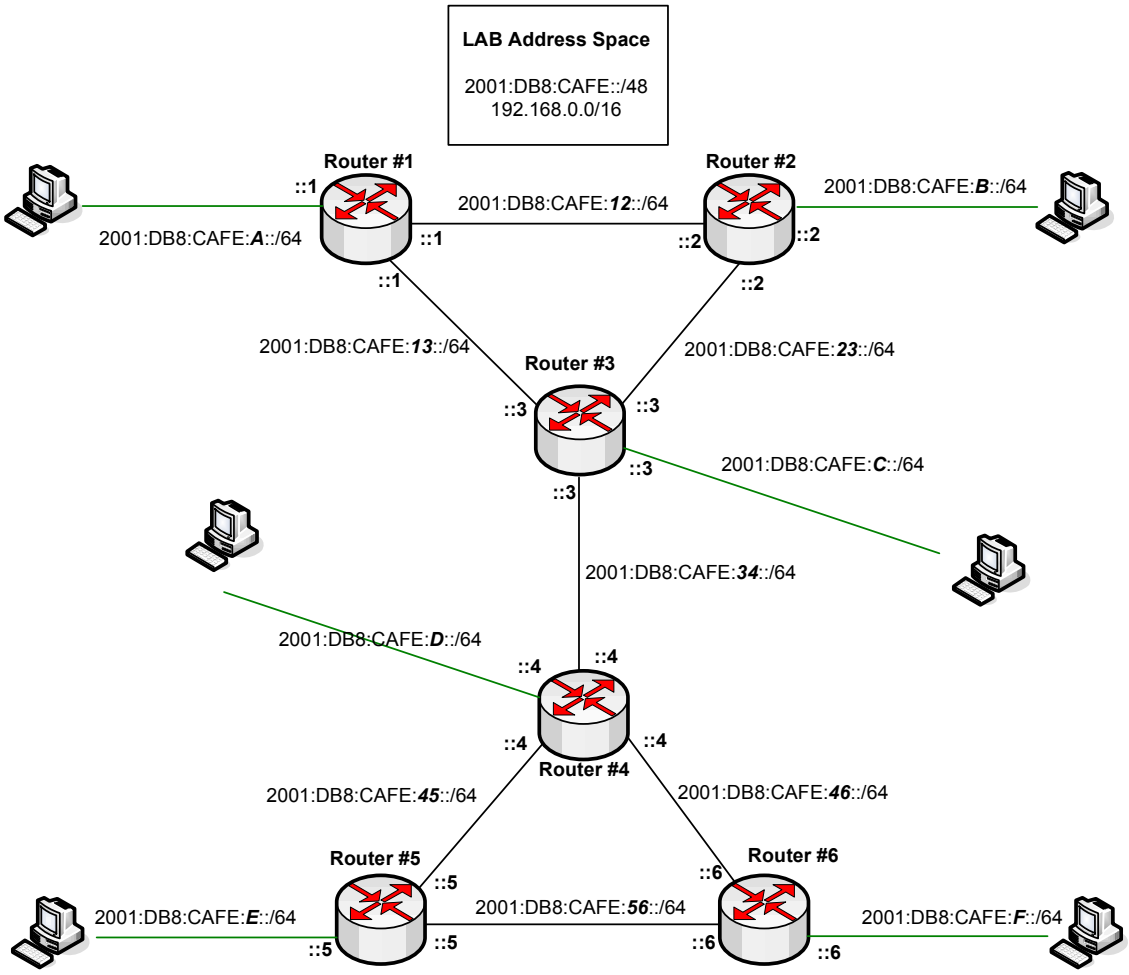

Figure 2 – Scenario Topology

Preparing the LAB

If your router is not yet configured, load the configuration from the flash memory:

```
RouterX# copy flash:initv6-config running-config
RouterX# wr
```

Bellow you'll find one table per router with the different addresses for each of the router's interface:

Router 1:

Interface	IPv6 Address
FastEthernet **0**	2001:DB8:CAFE:***12***::1 /64
FastEthernet **1**	2001:DB8:CAFE:***13***::1 /64
FastEthernet **5**	2001:DB8:CAFE:***A***::1 /64

Table 1: Addresses for each interface on router 1

Router 2:

Interface	IPv6 Address
FastEthernet **0**	2001:DB8:CAFE:***12***::2 /64
FastEthernet **1**	2001:DB8:CAFE:***23***::2 /64
FastEthernet **5**	2001:DB8:CAFE:***B***::2 /64

Table 2: Addresses for each interface on router 2

Router 3:

Interface	IPv6 Address
FastEthernet **0**	2001:DB8:CAFE:***13***::3 /64
FastEthernet **1**	2001:DB8:CAFE:***23***::3 /64
FastEthernet **2**	2001:DB8:CAFE:***34***::3 /64
FastEthernet **5**	2001:DB8:CAFE:***C***::3 /64

Table 3: Addresses for each interface on router 3

Router 4:

Interface	IPv6 Address
FastEthernet **0**	2001:DB8:CAFE:***34***::4 /64
FastEthernet **1**	2001:DB8:CAFE:***45***::4 /64
FastEthernet **2**	2001:DB8:CAFE:***46***::4 /64
FastEthernet **5**	2001:DB8:CAFE:***D***::4 /64

Table 4: Addresses for each interface on router 4

Router 5:

Interface	IPv6 Address
FastEthernet **0**	2001:DB8:CAFE:***45***::5 /64
FastEthernet **1**	2001:DB8:CAFE:***56***::5 /64
FastEthernet **5**	2001:DB8:CAFE:***E***::5 /64

Table 5: Addresses for each interface on router 5

Router 6:

Interface	IPv6 Address
FastEthernet **0**	2001:DB8:CAFE:***46***::6 /64
FastEthernet **1**	2001:DB8:CAFE:***56***::6 /64
FastEthernet **5**	2001:DB8:CAFE:***F***::6 /64

Table 6: Addresses for each interface on router 6

Task 1: *Enabling the MBGP*

Complete the following steps for creating a MBGP process on your router.

Step 1: Testing connectivity

Connect to your router. Use the IPv6 auto configuration and plug your PC to a FastEthernet port on the router and *telnet* to it.

The first step you must do is to see if your router has IPv6 routing active. The global **ipv6 unicast-routing** command should appear in the configuration.

Try to *ping* another router that is not directly connected you yours.

Step 2: Configuring the BGP process and router ID

Configure the eBGP main process on your router. Remember that in the case of MBGP you will have to create an IPv6 address family and configure a BGP router ID.

(**Tip**: router bgp <as_number>
address-fammily …)

Also note, that if you don't have any IPv4 addresses on your router, you must configure a router ID, or your BGP process will not start and you'll get an error message "*%BGP-4-NORTRID: BGP could not pick a router-id. Please configure manually.*".

(**Tip**: bgp router-id …)

For your AS number, you should use the following peering table:

Group	AS Number	Router ID	Peer With AS
1	100	1.1.1.1	300
2	200	2.2.2.2	300
3	300	3.3.3.3	100, 200 and 400
4	400	4.4.4.4	300, 500 and 600
5	500	5.5.5.5	400
6	600	6.6.6.6	400

Table 7 – Peering configuration

Step 3: Peering with the neighbours

Use

Table 7 to check your peer routers.

(**Tip**: neighbour X:X:X:X::X/<0-128> …)

You should use the addresses directly connected to your interface. If it was iBGP you could use the loopback addresses as they would be known by you IGP.

Step 4: Advertise your route

Now advertise your routes to your peers.

(**Tip**: network …)

The network you should use is listed in the following table:

Group	Advertised Network
1	2001:DB8:CAFE:1::/64
2	2001:DB8:CAFE:2::/64
3	2001:DB8:CAFE:3::/64
4	2001:DB8:CAFE:4::/64
5	2001:DB8:CAFE:5::/64
6	2001:DB8:CAFE:6::/64

Table 8 – Advertised Routes

Step 5: Check BGP Summary

See the status of your BGP process and how many routes you are receiving.

(**Tip**: show bgp ipv6 unicast …)

Note: In case you are having trouble, look at your synchronization and auto-summary configuration.

Step 6: Check advertised routes

Look at the route you are advertising to your peer. Are they correct?

(**Tip**: show bgp ipv6 unicast …)

Step 7: Check received routes

Verify the routes you are receiving from your peers.

- Are they correct?
- Is the AS Path for each route correct?
- Elaborate a query using the *regexp* expression, to find routes from a neighbour not directly connected to you.

Step 8: Test Application level

Create a Loopback interface and configure an IPv6 address on it, according to the following table.

Router 1	Advertised Network
1	2001:DB8:CAFE:1::1/128
2	2001:DB8:CAFE:2::1/128
3	2001:DB8:CAFE:3::1/128
4	2001:DB8:CAFE:4::1/128
5	2001:DB8:CAFE:5::1/128
6	2001:DB8:CAFE:6::1/128

Table 9 – IPv6 address for loopback interface

Try to *ping* one of these addresses.

Did you succeed? Why?

Step 9: Reset a BGP neighbour

Add another route to announce to your peer according to the following table:

Group	Advertised Network
1	2001:DB8:CAFE:11::/64
2	2001:DB8:CAFE:12::/64
3	2001:DB8:CAFE:13::/64
4	2001:DB8:CAFE:14::/64
5	2001:DB8:CAFE:15::/64
6	2001:DB8:CAFE:16::/64

Table 10 - Advertised Routes

See if you are advertising the route.

- Now reset the BGP process.
- How long does it take to have the peers exchanging routes again?
- Perform a soft reset to the BGP process. What is the difference?

Step 10: iBGP configuration

With the same topology, now there will be only two Autonomous Systems. So, only router 3 and 4 will run eBGP, all others will be running only iBGP.

In this case the distribution of routers to AS will be according to the following table:

Group/Router	Advertised Network
1, 2, 3	AS 100
4, 5, 6	AS 400

Table 11 – Configuration fot two AS

Which modifications do you have to do in your router's configuration?

Now place all routers in the same AS. Can you ping all routers?

Summary

After completing these exercises, you should be able to:

- *Configure the BGP Process*
- *Debug and analyse information from the process*
- *Reset the process*

Appendix

Step 1: Testing connectivity

From your PC ping the other routers or any PC from a computer not directly connected to your router.

Step 2: Configuring the BGP process and router ID

Configure the eBGP main process and router ID.

```
Router3# configure terminal
Router3(config)# router bgp 300
Router3(config-router)# address-family ipv6 unicast
Router3(config-router-af)# bgp router-id 3.3.3.3
```

Step 3: Peering with the neighbours

```
Router3(config)# router bgp 300
Router3(config-router)# address-family ipv6 unicast
Router3(config-router-af)# neighbor 2001:DB8:CAFE:13::1 remote-as 100
Router3(config-router-af)# neighbor 2001:DB8:CAFE:13::1 activate
Router3(config-router-af)# neighbor 2001:DB8:CAFE:34::4 remote-as 400
Router1(config-router-af)# neighbor 2001:DB8:CAFE:34::4 activate
```

Step 4: Advertise your routes

```
Router3(config)# router bgp 300
Router3(config-router)# address-family ipv6 unicast
Router3(config-router-af)# network 2001:690::/32
```

Step 5: Check BGP Summary

```
Router1# show bgp ipv6 unicast summary
BGP router identifier 3.3.3.3, local AS number 300
(…)
Neighbor V AS  MsgRcvd MsgSent TblVer InQ OutQ Up/Down  Ste/PfRd
2001:DB8:CAFE:13::1
         4 100 1117    1118    4       0   0    03:00:30 1
2001:DB8:CAFE:34::4
         4 400 148     150     4       0    0  02:23:48  1
```

Step 6: Check advertised routes

```
Router3#show bgp ipv6 unicast nei 2001:DB8:CAFE:34::4 advertised-routes
(...)
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete
Network       Next Hop            Metric LocPrf Weight Path
*> 2001:DB8:CAFE:1::/64
              2001:DB8:CAFE:13::1  0             0      100 i
*> 2001:DB8:CAFE:3::/64
                    ::             0             32768   i
*> 2001:DB8:CAFE:4::/64
                    2001:DB8:CAFE:34::4
                                   0             0       400  i
Total number of prefixes 3
```

You might be having trouble and find out that your route is not being announced. If this happens, try:

```
Router1(config-router)# no synchronization
Router1(config-router)# no auto-summary
```

Remember that you can only advertise routes that you are able to announce. So if the network you are advertising is not being used, you must force it to be up. For example, to force the route on router 3, do:

```
Router3#(config)# ipv6 route 2001:DB8:CAFE:3::/64 Null0
```

Step 7: Check received routes

```
Router4# show bgp ipv6 unicast nei 2001:DB8:CAFE:34::4 routes
BGP table version is 4, local router ID is 3.3.3.3
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

Network     Next Hop            Metric LocPrf Weight Path
*> 2001:DB8:CAFE:1::/64
              2001:DB8:CAFE:34::3               0      300 100 i
*> 2001:DB8:CAFE:3::/64
              2001:DB8:CAFE:34::3 0              0      300   i
Total number of prefixes 2
```

Here you can see the routes received, the AS-Path (in bold), and check the origin of the routes. You can also filter using regular expressions.

Eg: filter the ones from AS 300 appears:

```
Router4# show bgp ipv6 unicast regexp 300
(...)
Network          Next Hop            Metric LocPrf Weight Path
*> 2001:DB8:CAFE:1::/64
                 2001:DB8:CAFE:34::3                  0      300 100 i
```

Step 8: Test Application level

Try to *ping* to one of the addresses. If you can't reach it, try debugging one IPv6 packet on the destination. Can you see why you are not reaching?

Now, try an extended IPv6 ping using the loopback address as your origin IPv6 address. Can you ping now?

Eg:

```
Router4# ping ipv6
Target IPv6 address: 2001:DB8:CAFE:4::1
Repeat count [5]:
Datagram size [100]:
Timeout in seconds [2]:
Extended commands? [no]: y
Source address or interface: 2001:DB8:CAFE:1::1
UDP protocol? [no]:
Verbose? [no]:
Precedence [0]:
DSCP [0]:
Include hop by hop option? [no]:
Include destination option? [no]:
Sweep range of sizes? [no]:
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2001:DB8:CAFE:4::1, timeout is
2 seconds:
Packet sent with a source address of 2001:DB8:CAFE:1::1
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max =
0/0/4 ms
```

Step 9: Reset a BGP neighbour

```
Router4# clear bgp ipv6 unicast 2001:DB8:CAFE:34::3
Router4# clear bgp ipv6 unicast 2001:DB8:CAFE:34::3 soft
```

Step 10: iBGP configuration

With iBGP you should not calculate the next-hop. So all iBGP neighbours must be configured with *next-hop-self* option.

```
Router3(config-router-af)# neighbor 2001:DB8:CAFE:13::1 next-hop-self
Router3(config-router-af)# neighbor 2001:DB8:CAFE:34::4 next-hop-self
```

Now try putting all routers in AS300. See if in your router you see al the routes? Do you know why? Your "transit" routers must reflect the addresses that they are receiving. They must be reflector routers. Eg:

```
Router3(config-router-af)# neighbor 2001:DB8:CAFE:13::1 route-reflector-client
Router3(config-router-af)# neighbor 2001:DB8:CAFE:34::4 route-reflector-client
```

Some useful commands

```
Router1(config)# alias exec bgp6 show bgp ipv6 unicast summary
```

Debug commands

- **debug bgp ipv6 updates**
- **debug bgp ipv6 neighbour 2001:DB8:CAFE:<Y>::1 updates in**
- **debug bgp ipv6 neighbour 2001:DB8:CAFE:<Y>::1 updates out**

How to Filter Announcements

When receiving nothing from your peer, first you have to create the prefix list and then, apply this filter to the neighbor you wish.

Eg:

```
Router1# config term
Router1(config)# ipv6 prefix-list NOTHING seq 5 deny ::/0 le 128
Router1(config)# router bgp 300
Router1(config-router)# address-family ipv6 unicast
Router1(config-router-af)# neighbor 2001:DB8:CAFE:<X>::1 prefix-list NOTHING in
```

If you have time, create a prefix-list to guarantee that you are only announcing your network *2001:DB8:CAFE:<router_number>::/64* to your neighbors.

```
Router1(config)# ipv6 prefix-list <router_number>ANNOUNCE seq 5 accept 2001:DB8:CAFE:<VLAN>::/64
Router1(config)# ipv6 prefix-list < router_number >ANNOUNCE seq 10 deny ::/0 le 128
Router1(config)# router bgp <AS>
```

```
Router1(config-router)# neighbor 2001:DB8:CAFE:<X>::1 prefix-list <router_number >
ANNOUNCE out
Router1(config-router)# exit
Router1(config)# exit
Router1# exit
```