




**2<sup>nd</sup> SEE 6DISS Workshop  
Plovdiv 27-29 June 2007**

***Routing Hands-on Session  
(OSPFv3 & MP-BGP)***

Athanassios Liakopoulos  
[aliako@grnet.gr](mailto:aliako@grnet.gr)

**Lab information**

The local IPv6 lab consists of multiple Cisco 1800 Series routers, already preconfigured according to the following diagram.



*IPv6 lab – Topology diagram & IPv4 addressing plan*

## Router login

1. Use *telnet* protocol to connect to the routers using the following IPv4 addresses:

<i>Router ID</i>	<i>IPv4 loopback address</i>
Router1	10.1.1.1
Router2	10.2.2.2
Router3	10.3.3.3
Router4	10.4.4.4
Router5	10.5.5.5
Router6	10.6.6.6

### *Routers connection information*

For all routers, use the following authentication/authorisation credentials:

Login: *6diss*  
 Password: *6diss*  
 Enable password: *6diss*

**Note:** Do not modify Vlan1 configuration. You might lose connectivity to router and the rest of the participants will not be able to complete their part.

## Activate IPv6 routing

2. Activate IPv6 routing.
  - (Hint: From a CLI, execute `ipv6 unicast-routing`)
3. Try to ping other router. Do all the routers reply?

## Addressing configuration

4. Configure the following addressing plan on the routers.

<i>Name</i>	<i>IPv6 Loopback address</i>	<i>IPv4 Loopback address (for router-ID)</i>
Router1	2001:4b58:27:aab1::1/64	10.1.1.1
Router2	2001:4b58:27: aab2::1/64	10.2.2.2
Router3	2001:4b58:27: aab3::1/64	10.3.3.3
Router4	2001:4b58:27: aab4::1/64	10.4.4.4
Router5	2001:4b58:27: aab5::1/64	10.5.5.5
Router6	2001:4b58:27: aab6::1/64	10.6.6.6



<i>Interconnection interfaces (Rx - Ry)</i>	<i>IPv6 Address Prefix</i>
Router1 - Router2	2001:4b58:27: aac1::1/64
Router2 - Router3	2001:4b58:27: aac2::1/64
Router2 - Router5 (VLAN10)	2001:4b58:27: aac3::1/64
Router2 - Router5 (VLAN20)	2001:4b58:27: aac4::1/64
Router3 - Router4	2001:4b58:27: aac5::1/64
Router3 - Router5	2001:4b58:27: aac6::1/64
Router5 - Router6	2001:4b58:27: aac7::1/64
Note: Rx interface address = <i>prefix::1</i> Ry interface address = <i>prefix::2</i>	

5. Verify that you can *ping* the address of the routers connected to the router that you manage.
6. Have a look at the IPv6 details of an interface. Write down the different addresses you observe and be ready to explain their types and usage.

## Part A : OSPFv3 IPv6 configuration

### Objectives

- *Enable OSPF routing in the lab routers*
- *Use basic troubleshooting techniques*

### Exercises steps

1. Enable OSPFv3 routing protocol for IPv6 on all routers.
  - (Hint: Activate OSPF on the interface. Select a name for *OSPFv3 processID*.)

```
RouterX# enable
RouterX# configure terminal
RouterX(config)# interface fastethernet[X]
RouterX(config-if)# ipv6 ospf processID area areaid
```
2. Enable CEF switching for IPv6 on lab routers
  - (Hint: Use the command

```
RouterX(config)# ipv6 cef
```

)
3. Enable the OSPFv3 process you have configured in step 1 on all interfaces of the lab (except loopback interfaces). Use area 0 for OSPFv3.
  - (Hint 1: You can verify that the routing process is already created using the following command:

```
Router1# show configuration | inc ospf
ipv6 ospf 1000 area 0
ipv6 router ospf 1000
```

There are two lines, the one you configured before and the routing process that was automatically created.

```
Router1(config)# ipv6 router ospf 1000
Router1(config-rtr)#router-id 10.1.1.1
```
  - (Hint 2: Be sure that you use correct router-id!)
4. Check OSPFv3 connections are established between routers.
5. Propagate the loopback addresses in OSPFv3.
  - Hints: There are several ways to achieve this:

**Redistribution**

```
Router1(config)# ipv6 router ospf 1000
Router1 (config-rtr)# redistribute connected
Router1 (config-rtr)# redistribute static
```

**Note:** The routes from an interface will only be announced if that interface is up, or if you add its address to the routing table, for example by introducing a static route:

```
Router1(config)# ipv6 route 2001:4b58:27:aab1::/64 null 0
```

**Including in OSPFv3 with passive**

```
Router1(config)# interface loopback0
Router1#(config-if)# ipv6 ospf 1000 area 0
Router1(config)# ipv6 router ospf 1000
```



```
Router1 (config-rtr)# passive-interface loopback 0
```

6. Check all routers in the labs receive all interconnection and loopback prefixes via OSPFv3.
7. Check reachability of all routers loopback addresses from your router using ping command.
  - Hints:

### Step 1: Check OSPFv3 interfaces

```
Router1# show ipv6 ospf
It is an autonomous system boundary router
Originate Default Route with metric 100 always
(...)
Number of areas in this router is 1. 1 normal 0 stub 0 nssa
Area BACKBONE(0)
    Number of interfaces in this area is 2
(...)
```

```
Router1# show ipv6 ospf interfaces
(...)
FastEthernet0 is up, line protocol is up
    Link Local Address FE80::216:C8FF:FE30:5FC4, Interface ID 2
    Area 0, Process ID 1000, Instance ID 0, Router ID 3.3.3.3
    Network Type BROADCAST, Cost: 1
(...)
Designated Router (ID) 1.1.1.1, local address FE80::7D2
Backup Designated router (ID) 3.3.3.3, local address FE80::FC4
(...)
```

### Step 2: Check OSPFv3 neighbors

```
Router3# show ipv6 ospf neighbor
Neighbor ID      Pri   State           Dead Time Interface ID Interface
-----
4.4.4.4          1    FULL/BDR        00:00:30          2      Vlan32
1.1.1.1          1    FULL/DR         00:00:37          2      FastEthernet0
```

### Step 3: Check the OSPFv3 database

```
Router1# show ipv6 ospf database
      OSPFv3 Router with ID (1.1.1.1) (Process ID 1000)
      Router Link States (Area 0)
ADV Router      Age      Seq#          Fragment ID Link count Bits
-----
1.1.1.1         81      0x80000047 0              1          E
3.3.3.3         76      0x80000040 0              1          E
(...)

      Net Link States (Area 0)
ADV Router      Age      Seq#          Link ID      Rtr count
-----
1.1.1.1         87      0x80000008 2              2
(...)

      Link (Type-8) Link States (Area 0)
ADV Router      Age      Seq#          Link ID      Interface
-----
1.1.1.1         1320    0x80000028 2              Fa0
(...)

      Intra Area Prefix Link States (Area 0)
ADV Router      Age      Seq#          Link ID      Ref-lstype Ref-LSID
-----
1.1.1.1         327     0x80000008 1002          0x2002 2
(...)
```

ADV Router	Type-5 AS External Link States	Age	Seq#	Prefix
1.1.1.1		563	0x80000006	2001:DB8:CAFE:A::/64
(...)				

## Step 4: Looking at the routes

```
Router1# show ipv6 route
IPv6 Routing Table - 5 entries
Codes: C - Connected, L - Local, S - Static, R - RIP, B - BGP
       U - Per-user Static route
       I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea, IS -
ISIS summary
       O - OSPF intra, OI - OSPF inter, OE1 - OSPF ext 1, OE2 -
OSPF ext 2
       ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2
OE2 2001:DB8:CAFE:A::/64 [110/20]
    via FE80::217:E0FF:FED6:7D2, FastEthernet0
C   2001:DB8:CAFE:13::/64 [0/0]
    via ::, FastEthernet0
L   2001:DB8:CAFE:13::3/128 [0/0]
    via ::, FastEthernet0
L   FE80::/10 [0/0]
    via ::, Null0
L   FF00::/8 [0/0]
    via ::, Null0
```

## Show, troubleshooting and debugging commands

- Number of routes by prefix  
Router1# show ipv6 route summary
- Forcing the SPF recalculation  
Router1#clear ipv6 ospf process
- If you want to recalculate the SFP algorithm again, clear the OSPF database. If you type clear ipv6 ospf force-spf instead, the database will not be cleared before you run the SFP algorithm.
- Authentication neighbors on an interface  
RouterX(config-if)# ipv6 ospf authentication ipsec spi 1000 md5  
12345678900987654321ascdefedcba0  
, where SPI value means Security Policy Index (a value between 256 and 4294967295) and the values after MD5 are the key in HEX format. One can also choose the SHA-1 algorithm instead of MD5.
- Authentication neighbors on an OSPF area  
RouterX(config-rtr)# area 0 authentication ipsec spi 1000 md5  
12345678900987654321ascdefedcba0  
, where SPI value means Security Policy Index (a value between 256 and 4294967295) and the values after MD5 are the key in HEX format. One can also choose the SHA-1 algorithm instead of MD5.
- Debug commands  
debug ipv6 ospf packets  
debug ipv6 events  
debug ipv6 ospf adj

## Part B: BGP configuration for IPv6

### Objectives

- *Enable BGP routing*
- *Use basic troubleshooting techniques*

### Exercises steps

1. Remove OSPFv3 configuration between router2 router3 and router5.
  - (Hint: Disable *ipv6 ospf* on the necessary interfaces)
  - In how many “domains” the network is split of?
2. Configure an *eMBGP* peerings between Router2, Router3 and Router5. For this purpose, interconnection addresses must be used to setup the peerings. Use the AS numbers:
  - *AS number of Router2 is 65151*
  - *AS number of Router3 is 65152*
  - *AS number of Router5 is 65153*
3. Configure an *iMBGP* peering between routers:
  - *Router1 and Router2 (AS65151)*
  - *Router3 and Router4 (AS65152)*
  - *Router5 and Router6 (AS65153)*

*Note that the iMBGP full mesh is configured between loopback addresses of the routers.*

- Hint: Configure the BGP 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.

```
router bgp <as_number>
address-fammily ipv6
neighbor <neighbor> remote-as <as_number>
```

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.".

```
bgp router-id <router_id>
```

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

4. Check the status of the *eMBGP* and *iMBGP* peerings. They must be in established state before going to the next step.
  - Hints:



## Check BGP Summary

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

```
show bgp ipv6 unicast summary
```

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

## Check advertised routes

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

```
show bgp ipv6 unicast neighbor <neighbor> advertised-routes
```

## Check received routes

```
show bgp ipv6 unicast neighbor <neighbor> routes
```

Verify the routes you are receiving from your peers. Are they correct?

Is the AS Path for each route correct?

## 5. Advertise your route.

Now advertise your routes to your peers.

(Tip: network ...)

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

<i>Group</i>	<i>Advertised Network</i>
1	2001:4b58:27: aad1:: /64
2	2001:4b58:27: aad2:: /64
3	2001:4b58:27: aad3:: /64
4	2001:4b58:27: aad4:: /64
5	2001:4b58:27: aad5:: /64
6	2001:4b58:27: aad6:: /64

## 6. Check that you receive prefixes via the eMBGP peerings. Check they are properly propagated to the routers of the lab through iMBGP peerings.

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

<i>Group</i>	<i>Advertised Network</i>
1	2001:4b58:27: aad7:: /64
2	2001:4b58:27: aad8:: /64
3	2001:4b58:27: aad9:: /64
4	2001:4b58:27: aada:: /64
5	2001:4b58:27: aadb:: /64
6	2001:4b58:27: aadc:: /64

*Check that you receive prefixes via the MBGP peerings.*

*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?*

*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
```

### **Bonus Exercises:**

8. *Enforce policies on the eMBGP peerings to accept only one loopback prefix (e.g: 2001:DB8:CAFE:8007::1/64).*
9. Apply a policy to prefer the path between router2 and router3 . For this purpose, configure on router3 the local-preference 200 on prefixes received from router2. Configure on router5 the local-preference of 150 on prefixes received from router2.
10. Check the BGP details to make sure the policy is properly configured. Using traceroute command, make sure that the path between router2 and router3 is preferred.

### **Debug commands**

Use to following command to debug the IPv6 routing updates

- terminal monitor
- debug bgp ipv6 updates
- debug bgp ipv6 neighbour 2001:4b58:27:<Y>::1 updates in
- debug bgp ipv6 neighbour 2001:4b58:27:<Y>::1 updates out