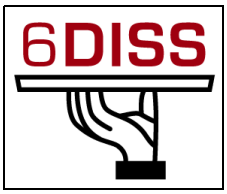

NIS Workshop - Ashgabad

25 - 27 April 2007

Routing hands-on: Local Lab

Piers O'Hanlon <p.ohanlon@cs.ucl.ac.uk>
Peter Kirstein <P.Kirstein@cs.ucl.ac.uk>
Janos Mohacsi <mohacsi@niif.hu>
Nikolay Pakulin <npak@ispras.ru>

Hands-on

Routing configuration / Local laboratory

Testbed setup

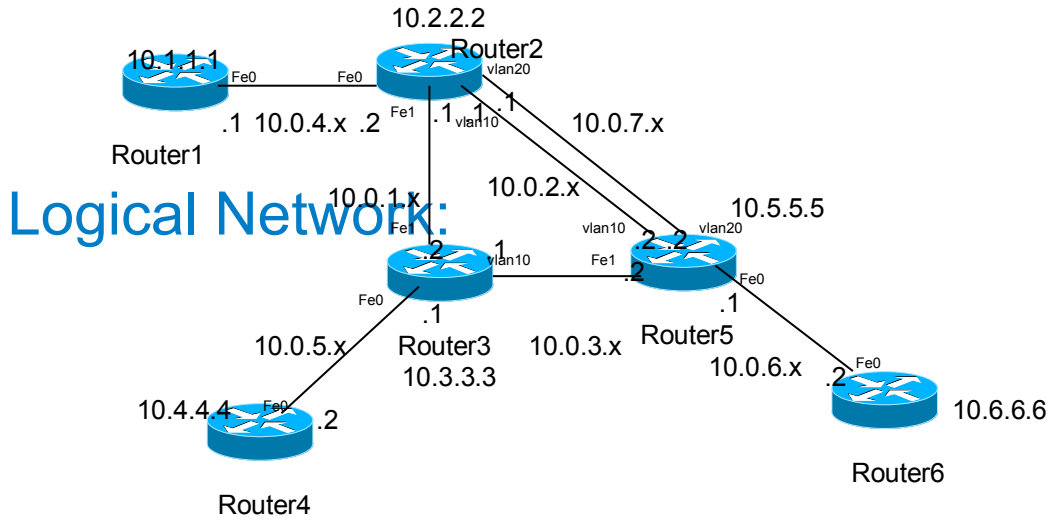

Testbed diagram

Routers login:

Use telnet protocol with the

Router	IPv4 loopback address	Ipv4 connection address
Router1	10.1.1.1	194.67.205.101
Router2	10.2.2.2	194.67.205.102
Router3	10.3.3.3	194.67.205.103
Router4	10.4.4.4	194.67.205.104
Router5	10.5.5.5	194.67.205.105
Router6	10.6.6.6	194.67.205.106

Routers connection information

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

Hints:
Connect to your router. Use the IPv6 auto configuration and plug your PC to any FastEthernet port on the router and telnet to it using the following authentication data:

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

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 to yours.

Note: do not modifiy Vlan1 configuration – you might lose connectivity to router

Addressing configuration

1°) Configure the following addressing plan on the routers.

Loopbacks:

Name	IPv6 Loopback address	IPv4 Loopback address (for router-ID)
Router1	2001:DB8:CAFE:8001::1/64	10.1.1.1
Router2	2001:DB8:CAFE:8002::1/64	10.2.2.2
Router3	2001:DB8:CAFE:8003::1/64	10.3.3.3
Router4	2001:DB8:CAFE:8007::1/64	10.4.4.4
Router5	2001:DB8:CAFE:8008::1/64	10.5.5.5
Router6	2001:DB8:CAFE:8009::1/64	10.6.6.6

Interconnections:

Interconnections (R1 - R2)	Prefix
router1 - router2	2001:DB8:CAFE:8101::/64
router2 - router3	2001:DB8:CAFE:8102::/64
router2 - router5 (VLAN10)	2001:DB8:CAFE:8103::/64
router2 - router5 (VLAN20)	2001:DB8:CAFE:8104::/64
router3 - router4	2001:DB8:CAFE:8105::/64
router3 - router5	2001:DB8:CAFE:8106::/64
Router5 – router6	2001:DB8:CAFE:8107::/64

R1 has address = prefix::1
R2 has address = prefix::2

2°) Check you can ping address of the routers connected to the router you manage.

3°) Take a look at the IPv6 details of an interface. Write down the different addresses you observe and give their types and usage.

OSPF configuration for IPv6

1°) Enable OSPFv3 routing protocol for IPv6 on all routers.
Hint:
Activate OSPF on the interface

```
RouterX# enable
RouterX# configure terminal
RouterX(config)# interface fastethernet[X]
RouterX#(config-if)# ipv6 ospf processID area areaid
```

Where process_ID is the specific name of the OSPFv3 process you will configure.

2°) Enable CEF switching for IPv6 on CISCO routers
Hint:
Activate CEF on router

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

3°) Enable the OSPFv3 process you have configured in question 1 on all interfaces of the lab (except loopback interfaces). Use area 0 for OSPFv3.
Hints:
If you look, you can see that the routing process is already created:

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

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

```
Router1(config)# ipv6 router ospf 1000
Router1(config-rtr)#router-id 10.1.1.1
```

Hint2: 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:
1. 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:DB8:CAFE:A::/64 null 0
```

2. 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
  (...)
                 Type-5 AS External Link States
ADV Router       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
```

8) Some useful commands

- To see the 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 - Try these commands and analyse their output.

```
debug ipv6 ospf packets
debug ipv6 events
debug ipv6 ospf adj
```

BGP configuration for Ipv6

0°) Remove OSPFv3 configuration between router2 router3 and router5:
tip: disable ipv6 ospf on the necessary interfaces

1°) Configure an eMBGP peerings between router2, router3 and router5. For this purpose, interconnection addresses must be used to setup the peerings. Also note that:

- *AS number of router2 is 65151*
- *AS number of router3 is 65152*
- *AS number of router5 is 65153*

2°) Configure an iMBGP peering between:

- *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.

Tips:
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.

3°) Check the status of the eMBGP and iMBGP peerings. They must be in established state before going to the next step.

Tips:

- 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?

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

5°) Check that you receive prefixes via the eMBGP peerings. Check they are properly propagated to the routers of the lab through iMBGP peerings.

6°) 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

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:
5°) Enforce policies on the eMBGP peerings to accept only one loopback prefix (e.g: 2001:DB8:CAFE:8007::1/64*).*

6°) 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.

7°) 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

- `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`