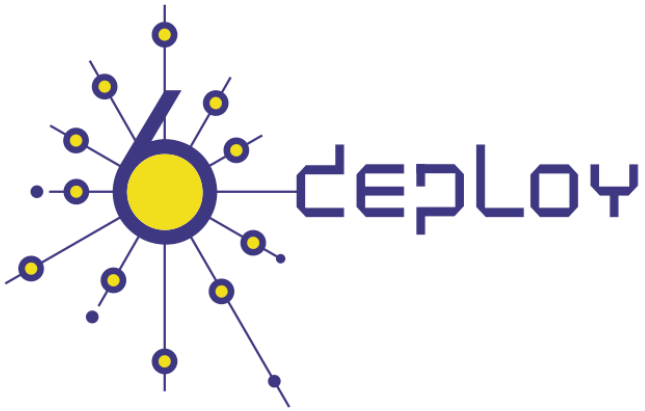

Title:	Document Version:
IPv6 OSPF and BGP Labs **IPv6 training material and related usage procedures**	0.3

Project Number:	Project Acronym:	Project Title:
223794	6DEPLOY	IPv6 Deployment Support

Contractual Delivery Date:	Actual Delivery Date:	Deliverable Type* - Security**:
30/05/2008	30/05/2008	R – CO

* Type: P – Prototype, R – Report, D – Demonstrator, O – Other
** Security Class: PU- Public, PP – Restricted to other programme participants (including the Commission Services), RE – Restricted to a group defined by the consortium (including the Commission Services), CO – Confidential, only for members of the consortium (including the Commission Services)

Responsible and Editor / Author:	Organization:	Contributing WP:
Consulintel	Consulintel	WP1

Authors (organisations):

Abstract:

This document describes the IPv6 OSPF and BGP Labs.

Keywords:

IPv6, Support, Training, Test-beds, Modules, 6Deploy, 6DEPLOY, Hands-on exercises

Revision History

The following table describes the main changes done in the document since created.

Revision	Date	Description	Author (Organization)
v0.1	12/06/2008	Document creation	Consulintel
v0.2	17/06/2008	Inputs from Paris, Sofia and APNIC Labs	
v0.3	17/06/2008	Inputs from Paris, Sofia and APNIC Labs	

Executive Summary

The goal of this document is to explain the IPv6 OSPF and BGP Labs.

It is based on “6Deploy Report on the Hands-on material”, September 2007.

Table of Contents

Figure Index

1. INTRODUCTION

This document is a report gathering the hands-on material used for the 6Deploy IPv6 technical workshops on Routing.

IGP (OSPF) and EGP (BGP) protocols are tested.

It includes information on the 6Deploy remote labs:

- RENATER, Paris, France
- BREN, Sofia, Bulgaria
- APNIC, Brisbane, Australia

2. ROUTING HANDS-ON

There are three remote laboratories available for the routing hands-on.

- Paris – 7 Routers
- Sofia – 6 Routers
- APNIC – 8 Routers

The distribution of the routers is as follow:

Laboratory	#	Router	Group
Paris	1	GSR-1	G1
	2	GSR-2	G2
	3	GSR-3	G3
	4	7200-1	G4
	5	7200-2	G5
	6	7200-3	G6
	7	7200-4	G7
Sofia	1	7200-1	G8
	2	7200-2	G9
	3	7500-1	G10
	4	1200-1	G11
	5	1200-2	G12
	6	1200-3	G13
APNIC	1	APNIC1	G14
	2	APNIC2	G15
	3	APNIC3	G16
	4	APNIC4	G17
	5	APNIC5	G18
	6	APNIC6	G19
	7	APNIC7	G20
	8	APNIC8	G21

2.1 Routing hands-on: Paris

2.1.1 Testbed

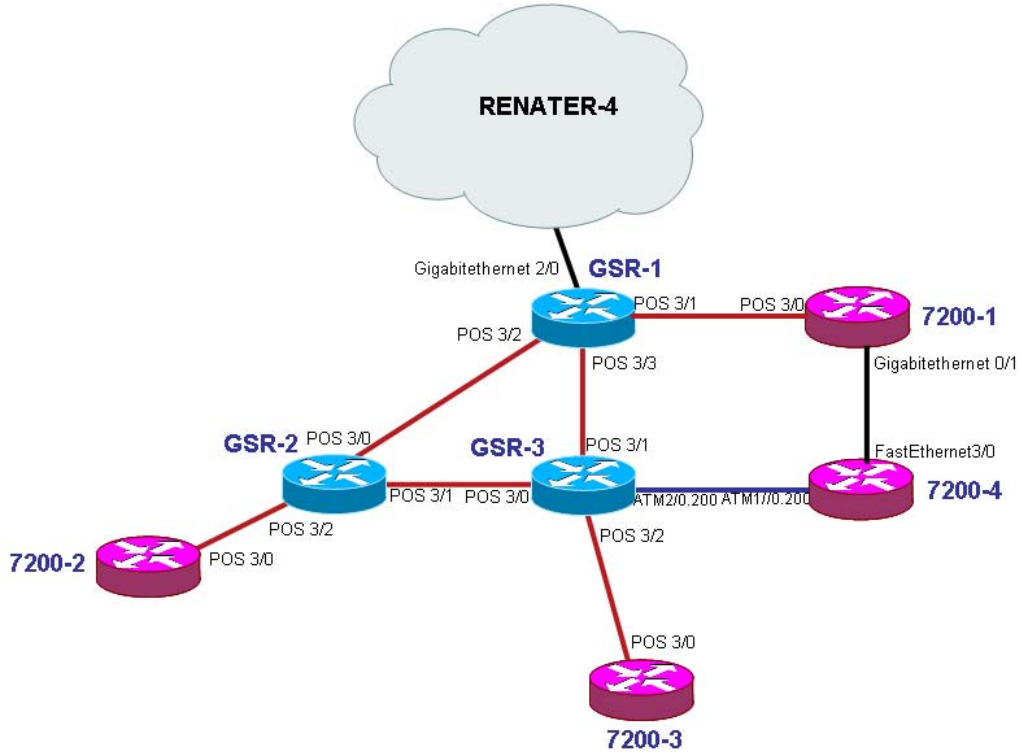

Figure 2-1: Paris Test-bed schema

Routers are configured beforehand to accept telnet connections from the training location.

The prefixes which can be used to configure the testbed are:

- IPv6 prefix: 2001:660:3008::/48
- IPv4 prefix: 194.254.101.128/25

These 2 prefixes are announced to the Internet by GSR-1.

2.1.2 Laboratory Exercise

Do NOT make any change on IPv4 protocols!

Do NOT set or change passwords!

2.1.2.1 Routers login

Use telnet protocol with the following addresses and login:

Router	IPv4 Address	Login
GSR-1	193.51.190.242	Login: 6deploy Password: 6deploy
GSR-2	193.51.190.246	Idem
GSR-3	193.51.190.247	Idem
7200-1	193.51.190.245	Idem
7200-2	193.51.190.249	Idem
7200-3	193.51.190.250	Idem
7200-4	193.51.190.253	Idem

2.1.2.2 Task 1: Addressing configuration

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

Loopbacks:

Name	IPv6 Loopback address	IPv4 Loopback address (for router-ID)
GSR-1	2001:660:3008:8001::1/64	194.254.101.130
GSR-2	2001:660:3008:8002::1/64	194.254.101.131
GSR-3	2001:660:3008:8003::1/64	194.254.101.132
7200-2	2001:660:3008:8007::1/64	194.254.101.133
7200-3	2001:660:3008:8008::1/64	194.254.101.134
7200-4	2001:660:3008:8009::1/64	194.254.101.135
7200-1	2001:660:3008:800A::1/64	194.254.101.136

Interconnections:

#	Interconnection (R1 - R2)	Prefix
I1	GSR-1- GSR-2	2001:660:3008:8101::/64
I2	GSR-1- GSR-3	2001:660:3008:8102::/64
I3	7200-2- GSR-2	2001:660:3008:8103::/64
I4	GSR-2- GSR-3	2001:660:3008:8104::/64
I5	GSR-3-7200-3	2001:660:3008:8105::/64
I6	GSR-3-7200-4	2001:660:3008:8108::/64
I7	GSR-1-7200-1	2001:660:3008:8109::/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.

2.1.2.3 Task 2: OSPF configuration for IPv6

1°) Enable OSPFv3 routing protocol for IPv6 on all routers.

2°) Enable CEF switching for IPv6 on CISCO routers

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.

4°) Check OSPFv3 connections are established between routers.

5°) Redistribute the loopback addresses in OSPFv3.

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.

2.1.2.4 Task 3: BGP configuration for IPv6

1°) Configure an eMBGP peering between GSR-2 and GSR-1 and another peering between GSR-3 and GSR-1. For this purpose, interconnection addresses must be used to setup the peerings. Also note that:

- *AS number of GSR-1 is 65152*
- *AS number of GSR-2 is 65153*
- *AS number of GSR-3 is 65154*

Note that you have to disable OSPF in "external" interfaces:

- *for GSR1, OSPF must be disabled in POS3/2 and POS3/3*
- *for GSR2, OSPF must be disabled in POS3/0 and POS3/1*
- *for GSR3, OSPF must be disabled in POS3/0 and POS3/1*

2°) Configure an iMBGP peering between:

- GSR-1 and 7200-1

- GSR-2 and 7200-2

- GSR-3 and 7200-3

- GSR3 and 7200-4

Note: For iMBGP peerings, you have to specify the ipv6 address used for the BGP routing updates:

```
router bgp xxxx
...
...
address-family ipv6
...
...
neighbor X:X:X:X::X update-source Loopback 0
...
```

Note that the iMBGP full mesh is configured between loopback addresses of the routers. This is the reason why OSPF is needed to reach loopback addresses.

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

4°) Check that you receive prefixes via the eMBGP peerings. Check they are properly propagated to the routers of the lab through iMBGP peerings.

Bonus:

5°) Check the connectivity to the IPv6 internet. Use the ping / traceroute commands from the routers to some well known IPv6 web servers

- *www.6deploy.org*
- *www.renater.fr*
- www.ipv6tf.org
- www.kame.net

6°) Enforce policies on the eMBGP peerings to accept only legacy IPv6 prefixes. Some more details about this legacy prefixes and the way you can configure the policy can be found at http://www.space.net/~gert/RIPE/ipv6-filters.html

7°) Apply a policy to prefer the path between GSR-1 and GSR-2. For this purpose, configure on GSR-2 the local-preference 200 on prefixes received from GSR-1. Configure on GSR-3 the local-preference of 150 on prefixes received from GSR-1.

8°) Check the BGP details to make sure the policy is properly configured. Using traceroute command, make sure that the path between GSR-2 and GSR-1 is preferred.

2.2 Routing hands-on: Sofia

2.2.1 Testbed

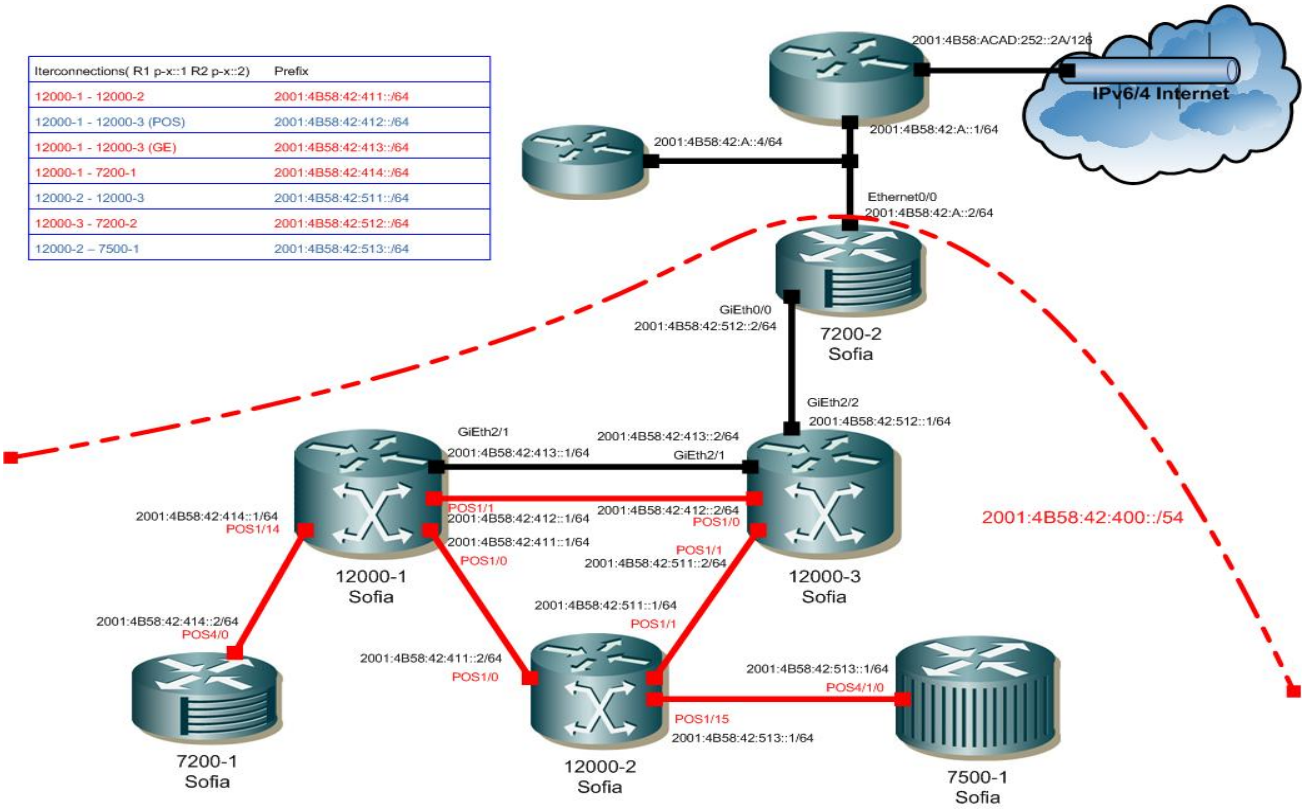

Figure 2-2: Sofia Test-bed schema

Routers are configured beforehand to accept telnet connections from the training location.

The prefixes which can be used to configure the testbed are:

- IPv6 prefix: 2001:4B58:42:400::/54

2.2.2 Laboratory Exercise

Do NOT make any change on IPv4 protocols!

Do NOT set or change passwords!

2.2.2.1 Routers login

Use telnet protocol with the following addresses and login:

Router	IPv4 Address - Port	Pass
7200-1	194.141.6.100 - 2005	6diss
7200-2	194.141.6.101 - 2005	7FP6diss
7500-1	194.141.6.102 - 2005	6diss
12000-1	194.141.6.103 - 2005	6diss
12000-2	194.141.6.104 - 2005	6diss
12000-3	194.141.6.105 - 2005	6diss

2.2.2.2 Task 1: Addressing configuration

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

Loopbacks:

Name	IPv6 Loopback address	IPv4 Loopback address (for router-ID)
7200-1	2001:4B58:42:701::1/64	194.141.6.130
7200-2	2001:4B58:42:702::1/64	194.141.6.131
7500-1	2001:4B58:42:703::1/64	194.141.6.132
12000-1	2001:4B58:42:707::1/64	194.141.6.133
12000-2	2001:4B58:42:708::1/64	194.141.6.134
12000-3	2001:4B58:42:709::1/64	194.141.6.135

Interconnections:

#	Interconnection (R1 - R2)	Prefix
I1	12000-1 - 12000-2	2001:4B58:42:411::/64
I2	12000-1 - 12000-3 (POS)	2001:4B58:42:412::/64
I3	12000-1 - 7200-1	2001:4B58:42:414::/64
I4	12000-2 - 12000-3	2001:4B58:42:511::/64
I5	12000-3 - 7200-2	2001:4B58:42:512::/64
I6	12000-2 - 7500-1	2001:4B58:42:513::/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.

2.2.2.3 Task 2: OSPF configuration for IPv6

1°) Enable OSPFv3 routing protocol for IPv6 on all routers.

2°) Enable CEF switching for IPv6 on CISCO routers

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.

4°) Check OSPFv3 connections are established between routers.

5°) Redistribute the loopback addresses in OSPFv3.

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.

2.2.2.4 Task 3: BGP configuration for IPv6

1°) Configure an eMBGP peering between 12000-2 and 12000-1 and another peering between 12000-3 and 12000-1. For this purpose, interconnection addresses must be used to setup the peerings. Also note that:

• AS number of 12000-1 is 65156

• AS number of 12000-2 is 65157

• AS number of 12000-3 is 65158

Note that you have to disable OSPF in "external" interfaces:

• for 12000-1, OSPF must be disabled in POS1/0 and POS1/1

• for 12000-2, OSPF must be disabled in POS1/0 and POS1/1

• for 12000-3, OSPF must be disabled in POS1/0 and POS1/1

2°) Configure an iMBGP peering between:

- 12000-1 and 7200-1

- 12000-2 and 7500-1

- 12000-3 and 7200-2

Note: For iMBGP peerings, you have to specify the ipv6 address used for the BGP routing updates:

```
router bgp xxxx
...
...
address-family ipv6
...
...
neighbor X:X:X:X::X update-source Loopback 0
...
```

Note that the iMBGP full mesh is configured between loopback addresses of the routers. This is the reason why OSPF is needed to reach loopback addresses.

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

4°) Check that you receive prefixes via the eMBGP peerings. Check they are properly propagated to the routers of the lab through iMBGP peerings.

Bonus:

5°) Check the connectivity to the IPv6 internet. Use the ping / traceroute commands from the routers to some well known IPv6 web servers

- *www.6deploy.org*
- *www.renater.fr*
- www.ipv6tf.org
- www.kame.net

6°) Enforce policies on the eMBGP peerings to accept only legacy IPv6 prefixes. Some more details about this legacy prefixes and the way you can configure the policy can be found at http://www.space.net/~gert/RIPE/ipv6-filters.html

7°) Apply a policy to prefer the path between 12000-1 and 12000-2. For this purpose, configure on 12000-2 the local-preference 200 on prefixes received from 12000-1. Configure on 12000-3 the local-preference of 150 on prefixes received from 12000-1.

8°) Check the BGP details to make sure the policy is properly configured. Using traceroute command, make sure that the path between 12000-2 and 12000-1 is preferred.

2.3 Routing hands-on: APNIC

2.3.1 Testbed

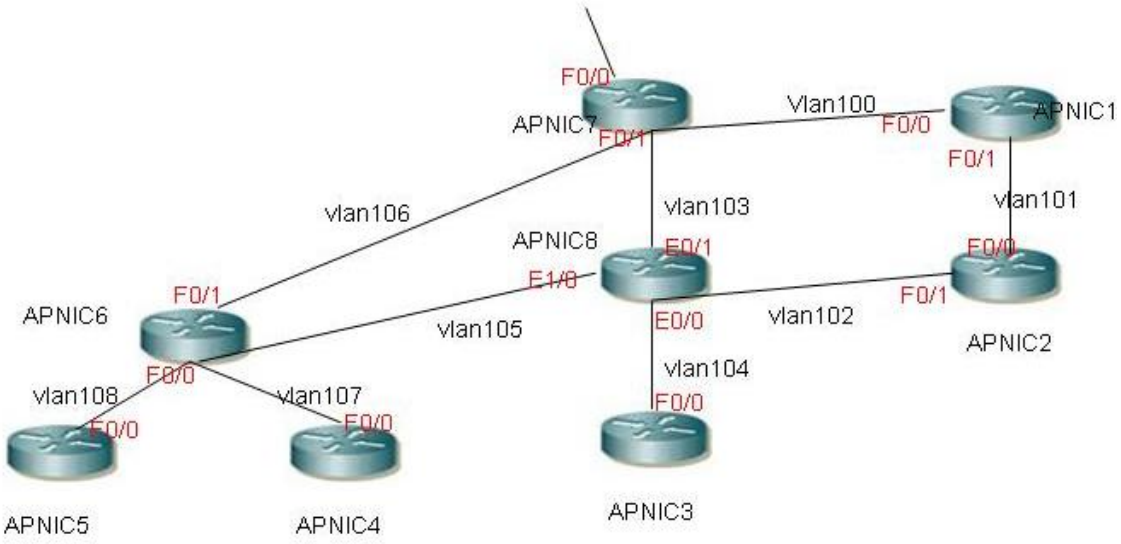

Figure 2-3: APNIC Test-bed schema

Routers are configured beforehand to accept SSH connections from the training location.

The prefixes which can be used to configure the testbed are:

- IPv6 prefix: 2001:DB8:3008::/48

2.3.2 Laboratory Exercise

Do NOT make any change on IPv4 protocols!

Do NOT set or change passwords!

2.3.2.1 Routers login

Use SSH protocol with the following addresses and login:

Router	IPv4 Address	Login
APNIC1	203.119.0.145	Login: apnic1 Password: apnic1
APNIC2	203.119.0.145	apnic2:apnic2
APNIC3	203.119.0.145	apnic3:apnic3
APNIC4	203.119.0.145	apnic4:apnic4
APNIC5	203.119.0.145	apnic5:apnic5
APNIC6	203.119.0.145	apnic6:apnic6
APNIC7	203.119.0.145	apnic7:apnic7
APNIC8	203.119.0.145	apnic8:apnic8

You will see a terminal server prompt:

Terminal-Server>

At the prompt enter: ssh apnicx (where x is the router number you are accessing).

When prompted for a password enter apnicx (where x is the router number).

You will then see the router user mode prompt.

Just type “ena” to enter privileged mode. No password has been set.

2.3.2.2 Task 1: Addressing configuration

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

Loopbacks:

Name	IPv6 Loopback address	IPv4 Loopback address (for router-ID)
APNIC1	2001:DB8:3008:8001::1/64	10.10.10.1
APNIC2	2001:DB8:3008:8002::1/64	10.10.10.2
APNIC3	2001:DB8:3008:8003::1/64	10.10.10.3
APNIC4	2001:DB8:3008:8004::1/64	10.10.10.4
APNIC5	2001:DB8:3008:8005::1/64	10.10.10.5
APNIC6	2001:DB8:3008:8006::1/64	10.10.10.6

APNIC7	2001:DB8:3008:8007::1/64	10.10.10.7
APNIC8	2001:DB8:3008:8008::1/64	10.10.10.8

Interconnections:

#	Interconnection (R1 - R2)	Prefix
I100	APNIC1 - APNIC7	2001:DB8:3008:8100::/64
I102	APNIC2 – APNIC8	2001:DB8:3008:8102::/64
I103	APNIC7 – APNIC8	2001:DB8:3008:8103::/64
I104	APNIC3 – APNIC8	2001:DB8:3008:8104::/64
I105	APNIC6 – APNIC8	2001:DB8:3008:8105::/64
I106	APNIC6 – APNIC7	2001:DB8:3008:8108::/64
I107	APNIC4 – APNIC6	2001:DB8:3008:8107::/64
I108	APNIC5 – APNIC6	2001:DB8:3008:8108::/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.

2.3.2.3 Task 2: OSPF configuration for IPv6

(Is not possible to perform this task because the image of the routers, continue with Task 3 using eMBGP and iMBGP).

1°) Enable OSPFv3 routing protocol for IPv6 on all routers.

2°) Enable CEF switching for IPv6 on CISCO routers

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.

4°) Check OSPFv3 connections are established between routers.

5°) Redistribute the loopback addresses in OSPFv3.

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.

2.3.2.4 Task 3: BGP configuration for IPv6

1°) Configure an eMBGP peering between APNIC7 and APNIC6 and another peering between APNIC7 and APNIC8. For this purpose, interconnection addresses must be used to setup the peerings. Also note that:

• AS number of APNIC6 is 65166

• AS number of APNIC7 is 65167

• AS number of APNIC8 is 65168

Note that you have to disable OSPF in "external" interfaces:

• for APNIC6, OSPF must be disabled in F0/0 and F0/1

• for APNIC7, OSPF must be disabled in F0/0 and F0/1

• for APNIC8, OSPF must be disabled in E0/0 and E0/1

2°) Configure an iMBGP peering between:

- APNIC7 and APNIC1

- APNIC6 and APNIC5

- APNIC6 and APNIC4

- APNIC8 and APNIC3

- APNIC8 and APNIC2

Note: For iMBGP peerings, you have to specify the ipv6 address used for the BGP routing updates:

```
router bgp xxxx
...
...
address-family ipv6
...
...
neighbor X:X:X:X::X update-source Loopback 0
...
```

Note that the iMBGP full mesh is configured between loopback addresses of the routers. This is the reason why OSPF is needed to reach loopback addresses.

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

4°) Check that you receive prefixes via the eMBGP peerings. Check they are properly propagated to the routers of the lab through iMBGP peerings.

Bonus:

5°) (Cancelled, you don't have connectivity to the IPv6 internet, Why?) Check the connectivity to the IPv6 internet. Use the ping / traceroute commands from the routers to some well known IPv6 web servers

- *www.6deploy.org*
- *www.renater.fr*
- www.ipv6tf.org
- www.kame.net

6°) Enforce policies on the eMBGP peerings to accept only legacy IPv6 prefixes. Some more details about this legacy prefixes and the way you can configure the policy can be found at http://www.space.net/~gert/RIPE/ipv6-filters.html

7°) Apply a policy to prefer the path between APNIC7 and APNIC6. For this purpose, configure on APNIC6 the local-preference 200 on prefixes received from APNIC7. Configure on APNIC8 the local-preference of 150 on prefixes received from APNIC7.

8°) Check the BGP details to make sure the policy is properly configured. Using traceroute command, make sure that the path between APNIC7 and APNIC6 is preferred.

3. NOTES ON 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 before and 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
1234567890098765432lascdefedcba0
```

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
1234567890098765432lascdefedcba0
```

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
```

4. NOTES ON 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 localpreference 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
```

5. ROUTING HANDS-ON: HELP COMMANDS

A document including a commands glossary is distributed to the trainees for the routing hands-ons.

This document helps the people who are not very familiar with the routing configuration. The labs are mainly composed of Cisco routers. There is also a Juniper and an Alcatel router, so the commands glossary includes a section for each vendor.

5.1 Cisco commands

Enable IPv6 on an interface

```
interface xxxxx
ipv6 enable
```

Configure an address

```
interface xxxxx
ipv6 address X:X:X:X::X/<0-128> (general address)
ipv6 address X:X:X:X::X (link-local address)
ipv6 address autoconfig (auto-configuration)
```

Example (LAN interface)

```
interface Ethernet0/0
ip address 192.168.1.254 255.255.255.0
ipv6 address 2001:db8:123:1::2/64
```

Configure a tunnel

Configure an IPv6 in IPv4 tunnel

```
interface tunnel x
tunnel source interface
tunnel destination X.X.X.X
ipv6 address X:X:X:X::X/<0-128>
tunnel mode ipv6ip (for direct tunneling)
tunnel mode gre ip (for gre encapsulation)
```

Configure an IPv6 in IPv6 tunnel

```
interface tunnel x
tunnel source interface
tunnel destination X.X.X.X
ipv6 address X:X:X:X::X/<0-128>
tunnel mode ipv6 (for direct tunneling)
tunnel mode gre ipv6 (for gre encapsulation)
```

Enable IPv6 routing

```
ipv6 unicast-routing
```

Configure static routes

```
ipv6 route prefix/prefixlen next_hop
ipv6 route ::/0 2001:db8:10a:1001::1
```

Routing (OSPFv3)

```
interface Ethernet0/0
ipv6 address 2001:db8:1:1::1/64
ipv6 ospf 1 area 0
!
interface Ethernet0/1
```

```
ipv6 address 2001:db8:1:2::2/64
ipv6 ospf 1 area 1
!
ipv6 router ospf 1
router-id 2.2.2.2
```

Area 1

Area 0

ABR

Eth0/1

Eth0/0

Routing (BGP)

```
router bgp xxxx
no bgp default ipv4-unicast
bgp router-id a.b.d.f
neighbor X:X:X:X::X remote-as …
neighbor X:X:X:X::X …
address-family ipv6
neighbor X:X:X:X::X activate
neighbor X:X:X:X::X …
network 2001:db8::/32
no synchronization
exit
```

Routing policy filtering

```
ipv6 prefix-list bgp-in-6net seq 5 deny ::/0
```

Means filter ::/0 exactly

```
ipv6 prefix-list bgp-in-6net seq 10 deny 3FFE:300::/24 le 28
ipv6 prefix-list bgp-in-6net seq 15 deny 2001:db8::/35 le 41
ipv6 prefix-list bgp-in-6net seq 20 permit 2002::/16
ipv6 prefix-list bgp-in-6net seq 25 permit 3FFE::/17 ge 24 le 24
ipv6 prefix-list bgp-in-6net seq 30 permit 3FFE:8000::/17 ge 28 le 28
```

Means every prefix matching 3FFE:8000::/17 with length 28

```
ipv6 prefix-list bgp-in-6net seq 35 permit 3FFE:4000::/18 ge 32 le 32
ipv6 prefix-list bgp-in-6net seq 40 permit 2001::/16 ge 32 le 35
```

Means every 2001::/16 derived prefix, with length between 32 and 35

Access Control Lists

```
ipv6 access-list vty-ipv6
permit tcp 2001:db8:0:401::/64 any eq telnet
deny ipv6 any any log-input
```

Applying an ACL to an interface

```
ipv6 traffic-filter <acl_name> in | out
```

Restricting access to the router

```
ipv6 access-class <acl_name> in | out
```

Applying an ACL to filter debug traffic

```
debug ipv6 packet [access-list <acl_name>] [detail]
```

Show commands

```
show bgp
show bgp ipv6 unicast/multicast/all summary
show bgp ipv6 neigh <addr> routes
show bgp ipv6 neigh <addr> advertised-routes
show bgp ipv6 neigh <addr> received-routes
show ipv6 route
show ipv6 interface
show ipv6 neighbors
```

5.2 Juniper commands

Interface configuration

```
interfaces {
name of interface {
unit x {
family inet {
address X.X.X.X/prefixlength;
}
family iso {
address Y.Y.Y.Y.Y.Y;
}
family inet6 {
address Z.Z.Z.Z::Z/prefixlength;
}
```

Router advertisements (stateless autoconfiguration)

```
protocols {
router advertisement {
interface interface name {
prefix IPv6_prefix::/prefixlength;
}
```

Configure tunnel (with Tunnel PIC)

```
interface {
ip-x/x/x {
```

```
tunnel {
source ipv4_source_address;
destination ipv4_destination_address;
}
family inet6 {
address ipv6_address_in_tunnel/prefixlength;
gr-x/x/y/z {
unit 0 {…}}
}
```

Static routes

```
Routing options {
rib inet6.0 { -> Means IPv6 unicast routing table
static {
route IPv6_prefix next-hop IPv6_address;
}
Routing options {
rib inet6.0 {
static {
route IPv6_prefix discard; -> Useful to originate a network
}
```

Routing (OSPFv3)

```
protocols {
ospf3 {
preference 20;
area 0.0.0.0 {
interface ge-0/3/0.808 {
metric 100;
}
interface lo0.0 {
passive;
}}}}
```

Routing (BGP)

```
protocols {
bgp {
local-as local_AS_number;
group EBGP_peers {
type external;
family inet6 {
(any | multicast | unicast) }
```

```
neighbor neighbor_IPv6_address;
peer-as distant_AS_number;
import in-PS;
export out-PS; }
```

Policy routing

```
policy statement in PS {
term from_outside_accept {
from {
route-filter 2002::/16 exact;
route-filter 3FFE::/17 prefix-length-range /24-/24;
route-filter 3FFE:8000::/17 prefix-length-range /28-/28;
route-filter 3FFE:4000::/18 prefix-length-range /32-/32;
route-filter 2000::/3 prefix-length-range /16-/16;
route-filter 2001::/16 prefix-length-range /29-/35; }
then {
accept; }
then reject; }
```

Show commands

```
show bgp summary
show route advert bgp <addr>
show route rece bgp <addr>
show route table inet6.0 (terse)
show interfaces
show ipv6 neighbors
```

5.3 Alcatel commands

Enable IPv6 on a VLAN interface

```
vlan "number"
show vlan
vlan "number" port default a/b-c
ipv6 interface "name" vlan number
ipv6 address "2001:XXXX::3/prefix" "name"
```

Tunnels configuration

```
vlan "number"
vlan "number" port default a/b
ip interface "name-v4" vlan "number"
ip interface "name-v4" address D.E.F.G mask H.I.J.K
ipv6 interface "name-v6" tunnel "number"
```

```
ipv6 interface "name-v6" tunnel source "@v4" dest. @v4"
ipb6 address "2001:XXXX::3/prefix" "name-v6"
```

Router Advertisements and auto-configuration

```
show ipv6 interface "name"
ipv6 interface "name" parameter numerical-value/yes/no
ipv6 interface name ra-send no
ipv6 interface ns-interval value
ipv6 interface ra-interval value
...
```

Starting routing

```
ipv6 route IPv6_prefix/length IPv6_address
```

Starting and configuring RIP routing

```
ipv6 load rip
ipv6 rip status enable
ipv6 rip interface "name of the IPv6 VLAN"
show ipv6 routes
ipv6 rip interface "name of the IPv6 VLAN" send-version\
[v1|v2|v1compatible|none]
ipv6 rip interface "name of the IPv6 VLAN" receive-version\
[v1|v2|v1compatible|none]
```

ipv6 rip interface "name of the IPv6 VLAN" metric [1-15]