

Login: 6diss
Password: 6diss

Addressing configuration

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

Loopbacks:

Name	Loopback address
12000-1	2001:660:f182:1001::1/64
7200-1	2001:660:f182:1002::1/64
12000-2	2001:660:f182:2001::1/64
12000-3	2001:660:f182:2002::1/64
7200-2	2001:660:f182:2003::1/64
7500-1	2001:660:f182:2004::1/64

Interconnections:

Interconnections (R1 - R2)	Prefix
12000-1 - 12000-2	2001:660:f182:1011::/64
12000-1 - 12000-3 (POS)	2001:660:f182:1012::/64
12000-1 - 12000-3 (GE)	2001:660:f182:1013::/64
12000-1 - 7200-1	2001:660:f182:1014::/64
12000-2 - 12000-3	2001:660:f182:2011::/64
12000-3 - 7200-2	2001:660:f182:2012::/64
12000-2 - 7500-1	2001:660:f182:2013::/64

R1 has address = prefix::1

R2 has address = prefix::2

2°) Check you can ping the routers connected to the router you manage. Note you can only ping interconnection addresses at this stage.

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

4°) On 7200-1, a tunnel should be configured for connecting the lab to the IPv6 internet. The tunnel end point offering the internet connectivity is a tunnel broker located in RENATER backbone. If this tunnel is not configured, use the following parameters for the setup.:

- Tunnel interface number: Tunnel0
- Tunnel end point destination address: 193.51.182.109
- Tunnel source interface: FastEthernet1/0
- IPv6 address to configure on the tunnel interface: 2001:660:F001:5::2/64
- Tunnel mode: IPv6 in IPv4 (protocol 41 encapsulation)


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.*
- 8°) On 7200-1, configure an IPv6 default route toward the tunnel interface you have configured in the previous addressing exercise. Originate the IPv6 default route on 7200-1 and advertise it using OSPFv3.*
- 9°) Check that the route is received on the other routers of the lab. Note that you still won't have connectivity to the IPv6 internet as there is no static route at the other side of the tunnel but a BGP peering. We will get IPv6 internet connectivity when we finish the BGP exercise.*

BGP configuration for IPv6

1°) On 7200-1, remove the IPv6 default route toward the tunnel and stop advertising that route with OSPFv3.

2°) We will now split the lab in 2 distinct IPv6 autonomous systems (AS) as depicted in the figure below. The AS 65001 will contain 7200-1 and 12000-1 routers and the AS 65002 will contain the other routers. Remove OSPFv3 configuration on interfaces between the 2 autonomous systems so that there is no more route information exchanged via an IGP between domains.



3°) On 7200-1, configure a multi-hop eMBGP peering toward a router located in RENATER having the entire IPv6 BGP table. Use the following configuration parameters:

- Distant AS number: 1717
- Local AS number: 65001
- Neighbor address: 2001:660:3007:400:1::
- Source interface: Tunnel0

4°) Configure also an eMBGP peering between 12000-1 and 12000-2, and another eMBGP peering between 12000-1 and 12000-3. We remind you at this stage that the AS number of 12000-1 is 650001 and that the AS number of 12000-2 and 12000-3 is 65002. We remind also that eBGP peerings are configured using addresses of the links interconnecting the routers peering together.

5°) Configure an iMBGP full mesh in each AS. Note that the iMBGP full mesh is configured between loopback addresses of the routers.

6°) Check the status of the eMBGP and iMBGP peerings. They must be in established state before going to the next step.

7°) Check that you receive prefixes via the eMBGP toward RENATER. Check they are properly propagated to 12000-1 through the iMBGP peering. Check that 12000-2 and 12000-3 receive properly the prefixes through the eMBGP peering these routers have toward 12000-1. At last, check that 7200-2 and 7500-1 receive properly IPv6 prefixes from the iMBGP peerings they have to 12000-2 and 12000-3.

5°) Configure 12000-1 and 7200-1 to originate via BGP the prefix 2001:660:f182:1000::/52. Configure 12000-2, 12000-3, 7200-2 and 7500-1 to originate via BGP the prefix 2001:660:f182:2000::/52. Make sure these prefixes are exchanged on the eMBGP peerings.

6°) Configure 7200-1 to aggregate both /52 prefixes in the prefix 2001:660:f182::/48 that is the only prefix that is accepted by its BGP peer.

7°) Check the connectivity to the IPv6 internet. Use the ping / traceroute commands from the routers to some well known IPv6 web servers

- www.6diss.org
- www.renater.fr
- www.kame.net

8°) 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>

9°) 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. Do the symmetric operation on 12000-1.

10°) 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 always preferred.