

Lab 7 Postlab: 6PE = IPv6 over MPLS (IPv4)

Or: Creating IPv6 tunnels over an IPv4 core using MPLS

What you will do:

1. Assign IPv6 addresses to interfaces (manual for system i/f; auto-generated for link-local)
2. Configure and verify OSPFv3 as an IPv6 routing protocol
3. Configure and verify LDP for MPLS transport across an IPv4 core
4. Configure export policies for route redistribution between OSPFv3 and MP-BGP
5. Configure and verify MP-BGP for distributing IPv6 routes across an IPv4 core
6. Verify end-to-end IPv6 connectivity via ping

Things that you will need to know or learn:

1. CLI command to enable IPv6 on an interface, to obtain a link-local address (NRS-II p. 285)
2. CLI command for configuring a routable IPv6 address on an interface (NRS-II p. 284)
3. CLI commands for configuring OSPFv3 (NRS-II p. 286)
4. CLI commands for configuring LDP (repeat from previous labs)
5. CLI commands for configuring an export policy for redistribution of OSPFv3 routes
6. CLI commands to configure & verify MP-BGP distribution of IPv6 routes (NRS-II p. 651)

What you need to submit and when:

1. Complete the “Lab 7 Post-lab” quiz on Blackboard with answers from your work, **before** your next lab period.

Required Equipment:

- Credentials for booking a MySRLab session
- An actual MySRLab booking
- A PC with internet access, Java, and terminal program; eg. Putty (provided by you)

Marks:

This post lab is worth 2/3 of the mark for Lab 7.

10% of your final mark is for labs done during the course of the semester.

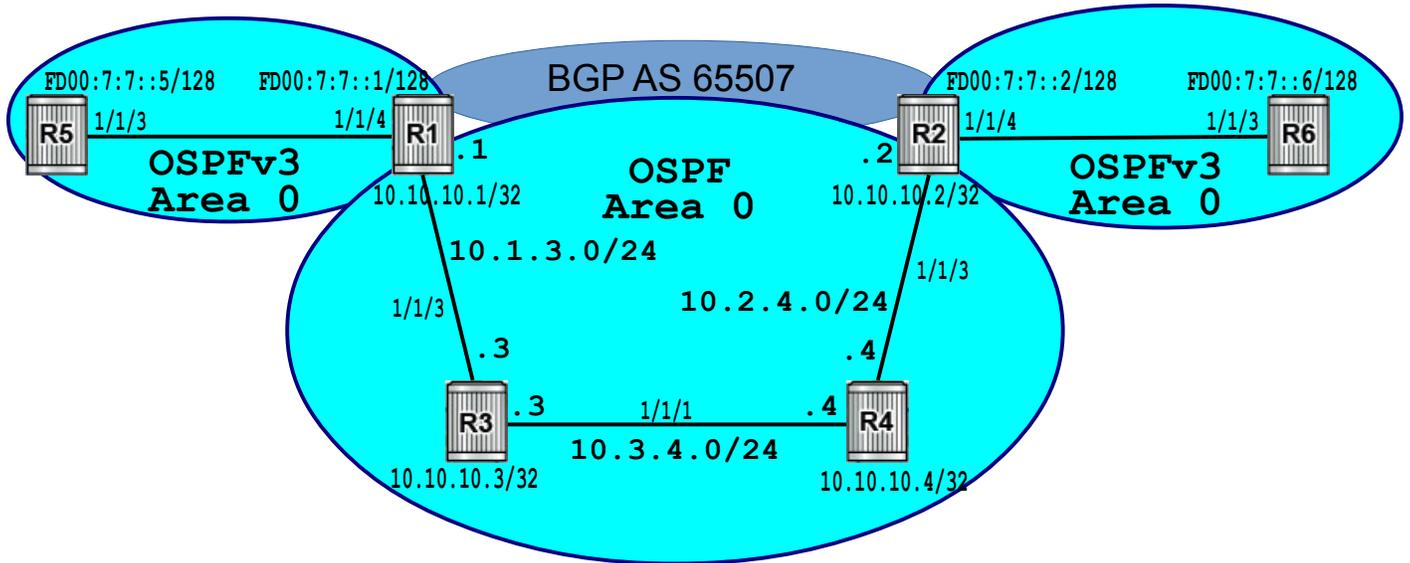
References and Resources:

- MPLS Module 5, section 5: notes on implementing 6PE using MPLS shortcuts
- Lab 13.4 in NRS-II p. 662-663; related material from p. 648-656
- OSPFv3: NRS-II Ch. 7, and p. 284-286 specifically
- IPv6 "private" addresses = fd00::/8 see: https://en.wikipedia.org/wiki/Unique_local_address
- IPv6-Reference.pdf document provided on course site
- **NOTE:** The topology shown in NRS-II Lab 13.4, page 662 is an exact match for MySRLab and thus this postlab exercise:

R5(CE) — R1(PE) – R3(P) – R4(P) – R2(PE) — R6(CE)

Topology

This is the MySRLab network topology that is used for this lab. It is similar to the top half of the Edu topology from the weekly in-lab. **NOTE:** that **only some** links are configured and activated!



Addressing Table for MySRLab Practice Labs

ALL routers are IPv4 addresses in the range **192.168.X.Y** with **X.Y** as given in the table below.

	PL3	PL4	PL5	PL6	PL7	PL8
R1	188.4	188.36	188.68	188.100	188.132	188.164
R2	188.5	188.37	188.69	188.101	188.133	188.165
R3	188.6	188.38	188.70	188.102	188.134	188.166
R4	188.7	188.39	188.71	188.103	188.135	188.167
R5	188.8	188.40	188.72	188.104	188.136	188.168
R6	188.9	188.41	188.73	188.105	188.137	188.169
PC1	188.10	188.42	188.74	188.106	188.138	188.170

The Ixia tester PC is accessed via the Remote Desktop Protocol (RDP), with clients available for Windows (built-in as "mstsc.exe"), Linux (rdesktop), and OSX (<http://cord.sourceforge.net/>)

Practice Lab	User ID	Password
PL3	See pages 4-5 of the <i>MySRLab Technical Primer</i> and page 8 of the <i>MySRLab Getting Started Guide</i> (available on Blackboard) for connectivity & login details.	
PL4		
PL5		
PL6		
PL7		
PL8		



Command Reference

```
configure • router • interface <name> • ipv6 # Start with link-local address, 1st addr
configure • router • interface <name> • ipv6 • address {IPv6-address/mask} # 2nd addr: manual
configure • router • ospf3 # configure OSPFv3
entry {#} • from • protocol bgp # policy statement to export from MP-BGP
entry {#} • from • protocol ospf3 # policy statement to export from OSPFv3
configure • router • bgp • group <name> • family ipv6 # xchange IPv6
configure • router • bgp • group <name> • neighbor {addr} • advertise-label ipv6 # signal a Label
show • router • route-table • ipv6 # Note the IPv6!
```

Task 1: Configure and Verify OSPFv3

The top-level exhaustion (ie. IANA's pool) of IPv4 addresses occurred on 31 January 2011: North American RIR exhausted their allocation on 24 September 2015, so customers are now having to use IPv6 addresses. The first task in this lab is to configure OSPFv3 between the CE and PE.

- Step 1. Start with a clean config: the Aventail login page should provide you with access to a "Deconfigure" option that you can click. You may need to reconfigure the cards and MDAs. The necessary commands are available in Lab 1 and Lab 7.
- Step 2. Configure the system interfaces on routers R5, R1, R2, R6 with the IPv6 addresses given in the topology diagram above: **FD00:07:07::xx/128** where xx = router number. Follow the method shown in listing 7.1 on NRS-II p. 284.
- Step 3. Configure a link local IPv6 address on both ends of the link between R5-R1 and R2-R6: this happens automatically when you simply enable IPv6 on an interface!
- Step 4. Don't forget to configure a Router ID! Use a 32-bit value identical to the conventional IPv4 system address (eg. 10.10.10.X)
- Step 5. Configure OSPFv3 between R5-R1 and R2-R6. Based on your knowledge of regular OSPF on 7750 routers, you should be able to configure it without extra help, given that the context begins at: **configure • router • ospf3**. Configure all four routers in OSPFv3 Area 0 and don't forget to include the system interface! In which context does the Router ID belong? If in doubt, consult the example in listing 7.3 on NRS-II p. 286.
- Step 6. If you've done everything correctly, you should be able to ping the system interfaces between R5-R1 and R2-R6. Try it. Also look in the IPv6 routing table for IPv6 addresses!
- Step 7. Execute and save the output for the following four "show" commands **on R1**:
(i) router interfaces; (ii) OSPFv3 neighbours; (iii) IPv6 routing table; (iv) an IPv6 ping
- Step 8. On R3 or R4, show the IPv6 routing table. How many entries do you expect?

Questions to answer:

- What is mask length for the IPv6 (auto-generated) link-local address?
- What are the first two octets in the IPv6 (auto-generated) link-local address?
- What is the exact, full command to display OSPFv3 neighbours?
- Assuming you followed the instructions correctly, in the IPv6/OSPFv3 neighbour listing, what is the exact value shown for R5?
- In the IPv6 routing table, what is the exact value for indicating a route is taken from the OSPFv3 RIB (ie. in the "proto" column)?
- In the output of a ping to an IPv6 address, what is the "name" given to the TTL field?

Task 2: Configure LDP throughout the Core

Configure OSPFv2 (ie. for IPv4) and LDP throughout the core, ie. routers R1-R4.

- Assign an IPv4 address to all system interfaces. Include the system interface in OSPFv2. Note that R1 and R2 are **dual-stack**: their system interfaces have both an IPv4 and an IPv6 address, and they're running both OSPFv3 and OSPFv2. R3-R4 are IPv4 **only**.
- Simple link LDP is sufficient, so use the same method as in previous lab(s).
- Run show commands (eg. on R1) and capture the output to prove that everything is configured and working properly.

Questions to answer:

- How many routes are in the core IPv4 routing table?
- On R1, when showing all LDP bindings, what value is displayed for the "No. of prefix bindings"?
- On R1, when showing active LDP bindings, what value is displayed for the "No. of active prefix bindings"? Tip: did you notice that the value given for active bindings is greater than the value given for all bindings?!!

Task 3: Configure export policies for redistribution between OSPFv3 and MP-BGP

Routers R1 & R2 are dual stack, so they'll be the cross-over point for IPv6 between OSPFv3 and MP-BGP. You'll need export policies for both directions, so the IPv6 addresses can flow between protocols, throughout the network.

- Step 1. On both R1 & R2, configure two export policies. They're very simple and similar to previously configured policies. Consult the command reference (above) to determine the keyword for each protocol. For the export policy from OSPFv3, you should also include and entry for directly connected interfaces. Refer to previous labs for example(s) of the complete method.
- Step 2. Since you already have OSPFv3 configured, apply the export policy which provides access to the IPv6 routes that will be learned from MP-BGP.
- Step 3. Hmm. An OSPF router is using an export policy to bring routes into Area 0. What else needs to be configured (in the OSPFv3 section)?

Task 4: Configure and verify MP-BGP on PE routers R1-R2

The last requirement for 6PE is MP-BGP which carries IPv6 routing information across the IPv4 core. Peering is iBGP, so the basic configuration should be familiar to you from previous post-labs, and then just two extra statements are sufficient to turn ordinary BGP into MP-BGP.

- Step 1. On R1 & R2, start by configuring an ordinary IPv4 BGP session between routers. **Verify** IPv4 BGP operation by showing BGP neighbours and checking for a state of (fully) "Established". This **must** be working in order to distribute IPv6 routes via MP-BGP.
- Step 2. Add the two extra statements that turn BGP into MP-BGP for IPv6. Consult the command reference (above) for help with this step.

Step 3. Wait a few moments for the protocol to converge. Then determine a suitable "show" command that proves the two PE routers have successfully established an "MP-BGP session". Troubleshoot if necessary. Remember, there's only 4 basic components and you already have experience with 3 of them: IGP, LDP, and BGP.

Step 4. On R1, show the BGP neighbours and capture the output for questions below.

Step 5. Examine the routing table on a PE router, R1 or R2. What route(s) appear there? For example, confirm whether your PE routing table has an entry for your partner's PE IPv6 address. You should be able to explain the reason for the presence of any routes that appear, and the reason for the absence of any routes which don't appear.

Step 6. Examine the routing table on a CE router, R5 or R6. What route(s) appear there? You should be able to explain the presence or absence of all routes!

Questions to answer:

- In show router bgp neighbor, what protocol is shown for the "Local Family"?
- In show router bgp neighbor, what protocol is shown for the "Remote Family"?
- In show router bgp neighbor, what protocol is shown for the "Advertise Label"?
- In show router bgp neighbor, what value is relevant to this lab in Local Capability"?
- In show router bgp neighbor, what's the relevant value in Remote Capability"?

Task 5: Redistribute the IPv6 routes from OSPFv3 into MP-BGP

Complete the final step of configuring IPv6 route redistribution from OSPFv3 into MP-BGP.

Step 1. On the PE routers R1-R2, add an export policy in MP-BGP that allows routes learned from OSPFv3 to be redistributed to the opposite PE.

Step 2. After a few moments, check the routing tables on the PE routers (R1-R2) and CE routers (R5-R6). You should be able to explain the presence or absence of all routes!

Step 3. Use ping to verify full IPv6 reachability between all CE & PE routers. Troubleshoot and correct any problems if necessary.

Step 4. Run the following commands and save the output for answers to questions below:

on R5: show the IPv6 routing table

on R1: show the IPv6 routing table; and **show router bgp routes ipv6 detail**

Step 5. Verify the routing tables of the Core routers (R3-R4). Look carefully: can you find the system interface addresses of the CE routers? Explain why or why not.

Questions to answer:

- On R5, what's the preference value (ie. Admin Distance) for the route to R2 or R6?
- On R5, what's the preference value (ie. Admin Distance) for the route to R1? Why?
- In R1's IPv6 route table, what extra marking exists for the routes to R2 or R6?
- In the BGP route details on R1, what indicates that the route to R2's or R6's IPv6 address is tunneled? (Hint: look at the "res[olved] NextHop".)
- Again in the BGP route details, what's the value for the "6PE Label"?

Task 6: Save your configs

As always, save the configs for all routers locally as well as in MySRLab.