

Lab 7: 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
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. There is no pre-lab for Lab 7.
2. Complete the in-lab part of the exercise (see below), **before** the end of your lab period.
3. Complete the "Lab 7 Post-lab" exercise on OSPFv3 and submit to BrightSpace, **before** the assigned due date for your lab group.

Required Equipment:

- USB memory stick to save results for post-lab questions
- Hard-cover lab notebook, for reference during SBA at the end of the course.
- PC with internet access, a browser, Java, and terminal program (Provided in T108)

In-Lab Marks:

- 1 mark: Demo of successful OSPFv3 routing between CE and PE routers
- 1 mark: Demo of successful MP-BGP peering between adjacent PE routers
- 1 mark: Demo of successful IPv6 ping between adjacent (horizontal or vertical) CE-CE routers
- 1 mark: Demo of successful IPv6 ping across all pods between CE-CE routers

The in-lab is worth 1/3 of the overall lab mark; the post-lab is worth 2/3 of the overall lab mark.

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
- EDU Lab: remote-access lab facility hosted at the Nokia Kanata campus
- IPv6 "private" addresses = fd00::/8 see: https://en.wikipedia.org/wiki/Unique_local_address
- Wk07-IPv6-Reference.pdf document provided on course site

Addressing & Login Table

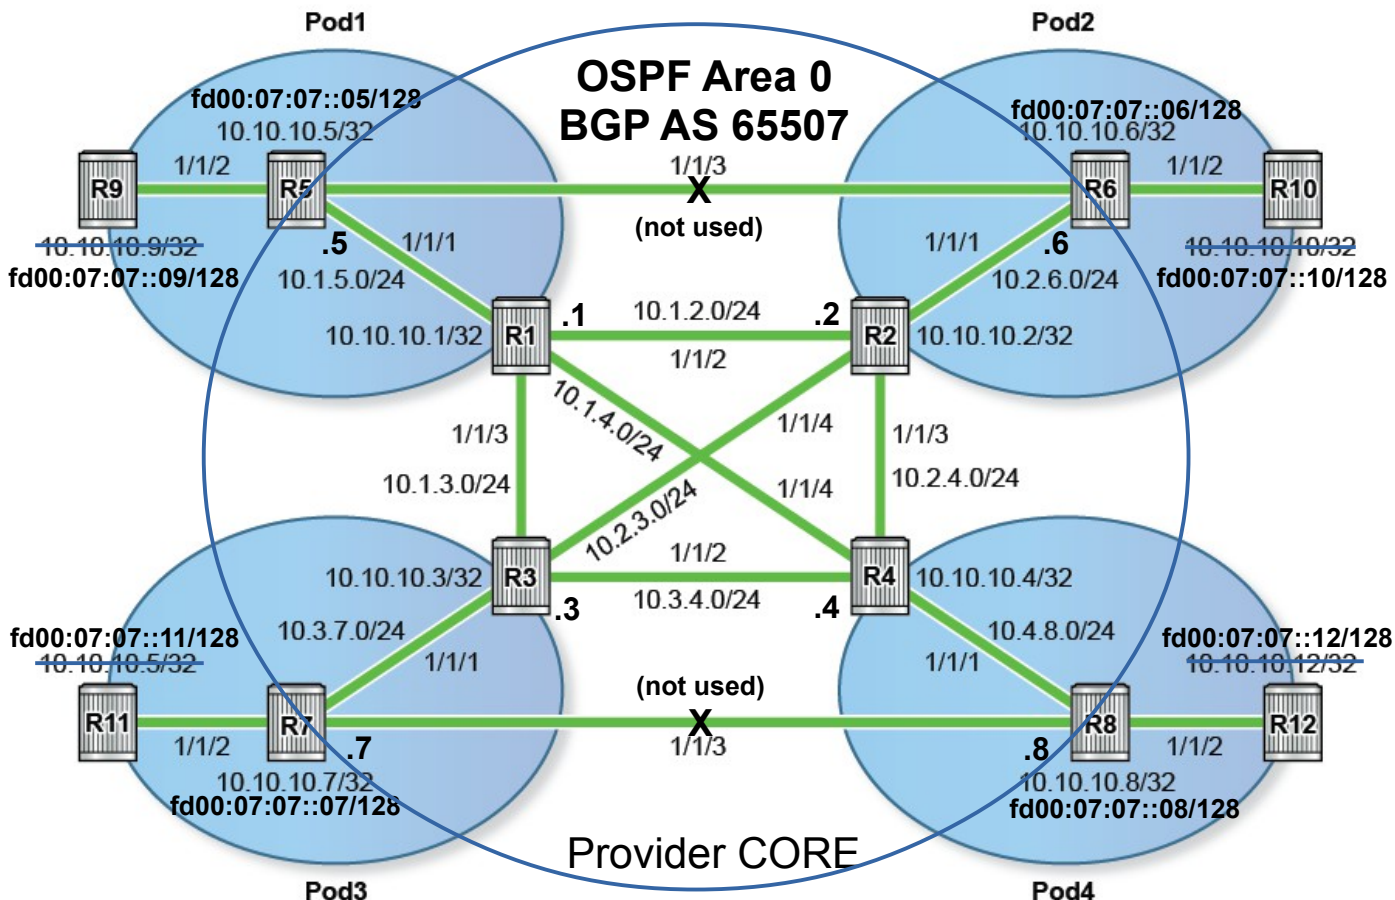
	Edu Lab 1	Edu Lab 2	Edu Lab 3
R1	.164	.196	.228
R2	.165	.197	.229
R3	.166	.198	.230
R4	.167	.199	.231
R5	.168	.200	.232
R6	.169	.201	.233
R7	.170	.202	.234
R8	.171	.203	.235
R9	.172	.204	.236
R10	.173	.205	.237
R11	.174	.206	.238
R12	.175	.207	.239

	My specific login information
EDU Lab # (1, 2, or 3)	
Individual login ID	
Corresponding password	

See BrightSpace for a list of login IDs and passwords; write **yours** in the space above.

Base IP address is: 192.168.206.0/24

Topology Diagram



- Note:**
1. CE routers (R9-R12) are IPv6 only
 2. PE routers (R5-R8) are dual-stack IPv4-IPv6 routers
 3. Core routers (R1-R4) are IPv4 only

Command Reference

```
show card # Step 1
configure • card {#} • card-type {card-type-from-show-command} # Step 2

show mda # Step 3
configure • card {#} • mda [#] mda-type {mda-name} # Step 4

configure • router • interface <name> • ipv6 • address {IPv6-address/mask}

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 between CE and PE routers

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 move to IPv6 addrs. The first task in this lab is to configure OSPFv3 between the CE and PE.

- Step 1. Start with a clean config: EDU labs: ask the lab Professor to deconfigure all the routers (For MySRLab, deconfigure them in the window from the additional login). You'll probably need to reconfigure the cards and MDAs. The necessary commands are repeated from Lab 1 in the command reference above.
- Step 2. Configure the system interfaces on your CE and PE routers 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 your CE and PE router: this happens automatically when you simply enable IPv6 on the interface!
- Step 4. Don't forget to configure a Router ID! (eg. 10.10.10.x)
- Step 5. Configure OSPFv3 between both your CE and PE router. 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 routers in 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 the CE and PE routers. Try it. Also look in the routing table for IPv6 addresses!

CHECK POINT #1: Prove, with output on your screen, that OSPFv3 is correctly configured.

Task 2: Configure LDP throughout the Core

Configure an IPv4 IGP and LDP throughout the core, on your P (R1-R4) and PE (R5-R8) routers. Simple link LDP is sufficient, so use the same method as in previous lab(s). Don't forget the IPv4 system interface addresses, even on the PE routers: they're **dual-stack** so need both addresses!

Verify: Use appropriate "show" commands to confirm that LDP is configured throughout the core.

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

The PE routers 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 throughout the network.

- Step 1. On your PE router, configure two export policies. They're very simple, with only a single entry required in each. Consult the command reference (above) to determine the keyword for each protocol. 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 learned from MP-BGP.
- Step 3. Hmm. An OSPF router is using an export policy to bring routes into Area 0. What else do you need to configure?

Task 4: Configure and verify MP-BGP on the PE router

The last requirement for 6PE is MP-BGP to carry 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. Pick a partner Pod. Confirm with them that you will configure mutual BGP sessions between your pods.
- Step 2. On your PE router, start by configuring an ordinary IPv4 BGP session to your partner pod's PE router.
- Step 3. **Verify** IPv4 BGP operation by confirming the neighbour state is (fully) Established. Do **not** proceed until you have this is working! If you don't get the peering, one troubleshooting option on SR OS, is to do a shut/no shut on BGP. If this fixed the problem, it's **important** you know the reason why (... it's not the router's fault; it's because you configured BGP in the **improper order**!!)
- Step 4. Add the two extra statements that turn BGP into MP-BGP for IPv6. Consult the command reference (above) for help with this step.
- Step 5. 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. (... And tell the lab Professor if anything's missing from these steps!)
- Step 6. Examine the routing table on your **PE** router. What route(s) appear there? For example, confirm whether your PE routing table has an entry for your partner's PE IPv6 address. You must 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 7. Examine the routing table on your **CE** router. What route(s) appear there? You must be able to explain the presence or absence of all routes!

CHECK POINT #2: Prove, with output on your screen, that your PE routers have a successful peering for exchanging IPv6 routes.

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 your PE router, add an export policy in MP-BGP that allows routes learned from OSPFv3 to be redistributed to your partner's PE.
- Step 2. After a few moments, check the routing tables on your CE and PE routers. You must be able to explain the presence or absence of all routes!
- Step 3. Use ping to verify full IPv6 reachability to/from all your and your partners' CE & PE routers. Troubleshoot and correct any problems if necessary.
- Step 4. Verify the routing tables of the Core routers (R1-R4). Look carefully: can you find the system interface addresses of the CE routers? Explain why or why not.
- Step 5. [Optional] Shut down the port for the interface connecting your P router with your partner's P router. Wait a few moments. Can you still ping your partner's CE router from your CE and PE routers? Explain why or why not.
Remember re-enable the port once you are done.

CHECK POINT #3: Prove, with output on your screen, that you can successfully ping the IPv6 address of your partner pod's CE router.

Task 6: Add BGP peerings to provide full connectivity

Add additional MP-BGP peering sessions to the other three PE routers in your EDU lab. Then confirm full reachability between all CE and PE routers throughout the topology.

CHECK POINT #4: Prove, with output on your screen, that you can successfully ping the diagonally opposite CE router.

Task 7: Save your configs

As always, save the configs for your P, PE, and CE routers on the EDU ftp server:
use **edu1 / dEEtgx2Y** or **edu2 / dEEtgx2Y** for each 1/2 of the class; **Rxx = R01, R02, ...**
save: `admin save ftp://{acct}:{pswd}@192.168.181.81/Rxx/ande0001-L7.cfg`
check: `file dir ftp://{acct}:{pswd}@192.168.181.81/Rxx/`
restore: `exec ftp://{acct}:{pswd}@192.168.181.81/Rxx/ande0001-L7.cfg`

For additional practice: You will redo this lab in the post-lab. The topology shown in NRS-II Lab 13.4, page 662 is an exact match for MySRLab:

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

See important information on the next page for changing over the configuration!!

Reverting to a different configuration

Identify a configuration file, then revert to it:

```
configure system rollback rescue-location ftp://rollback:rollback@192.168.181.81/13.0.R5/Initial  
admin rollback revert rescue now
```

```
#  
show card          #2020: card-type = iom3-xp  
configure card 1 card-type {card-type-from-show-command}  
show mda           #2020: mda-type = m10-1gb-xp-sfp  
configure card 1 mda 1 mda-type {mda-type-from-show-command}
```

Manual De-configuration

```
exit all  
configure router  
ospf shutdown  
no ospf  
bgp shutdown  
no bgp  
ldp shutdown  
no ldp  
rsvp shutdown  
mpls shutdown  
no mpls  
ospf3 shutdown  
no ospf3  
no autonomous-system  
no router-id
```

Sample starting config to bring up all required ports

```
exit all  
environment no more  
configure system login-control idle-timeout disable  
configure port 1/1/[1..4] no shut  
configure router  
/configure system name Lab7-6PE-R1
```

...