

Lab 7: OSPFv3

Or: When 4 Billion addresses just aren't enough

What You Will Do:

Working in pairs, combine your skills & knowledge to:

1. Configure multi-area OSPFv3 for IPv6
2. Confirm full end-to-end network reachability.
3. Capture and examine output from show commands which reveal the state of the network, especially the LSDBs and router states.

Things that you will need to know or learn:

1. How to set up multi-area OSPFv2 (from a previous lab)
2. As per the CLI command chart on the next page, OSPFv3 config is essentially *identical* to OSPFv2 config.

What you need to submit and when:

1. Complete the pre-lab quiz on BrightSpace, **before** the start of your lab period.
2. Complete the in-lab part of the exercise (see below), **before** the end of your lab period.
3. Complete post-lab exercise and quiz on BrightSpace, **before** your next lab.

In-Lab Marks:

[1 mark] Proof of multi-area OSPFv3 (Task 5 completed) with suitable "show" command(s)

[1 mark] Full end-to-end connectivity throughout your entire network.

The pre-lab is worth **24%**, the in-lab is worth **38%**, and the post-lab is worth **38%** of this lab, even though the number of points may differ between the three parts.

Required Equipment:

- A laptop and/or a USB memory stick to save results for post-lab questions
- **Hard-cover lab notebook**, for reference during **SBA** at the end of the course.

References and Resources:

- Slide deck on OSPFv3 networks, posted on course site
- OSPFv3 (free chapter) <https://www.ciscopress.com/articles/article.asp?p=2294214&seqNum=4>
- Packet Tracer ver 8 (available from NetAcad)

OSPFv3: Key config commands

	CLI context	Config Command	Show command
0	#	ipv6 unicast-routing interfaces have at least 1 IPv6 addr	ALWAYS needed!!!
1	# ipv6 router ospf {p-ID}	router-ID x.x.x.x !Still 32 bits!	show ipv6 protocols (+)
2		auto-cost reference-bandwidth {X}	show ipv6 ospf
3		passive-interface {i/f name}	show ipv6 protocols
4	(#) # ipv6 router ospf {p-ID}	ipv6 route ::/0 {ipv6-nextHop} default-information originate	show ipv6 route ospf
		clear ipv6 ospf {p-ID} process	
5	# interface {int-name}	ipv6 ospf {proc-ID} area {X}	show ipv6 ospf interface [i/f-name] [brief]
6		ipv6 ospf cost {value}	"
7		ipv6 ospf priority {0-255}	"
8		ipv6 ospf hello-interval {secs}	"
9		ipv6 ospf dead-interval {secs}	"
10		ipv6 ospf network point-to-point	"

Also: show **ipv6** ospf neighbor (neighbours) show **ipv6** ospf database (LSDB)

Generic: show **ipv6** route [ospf]

To be clear, here's a list of things that do not change going from OSPFv2 to OSPFv3:

Item	Changes?
0	Added; needed for any IPv6
1	Extra ipv6
2	NO changes
3	NO changes
4	NO changes
5	Just ip --> ipv6
6	Just ip --> ipv6
7	Just ip --> ipv6
8	Just ip --> ipv6
9	Just ip --> ipv6
10	Just ip --> ipv6

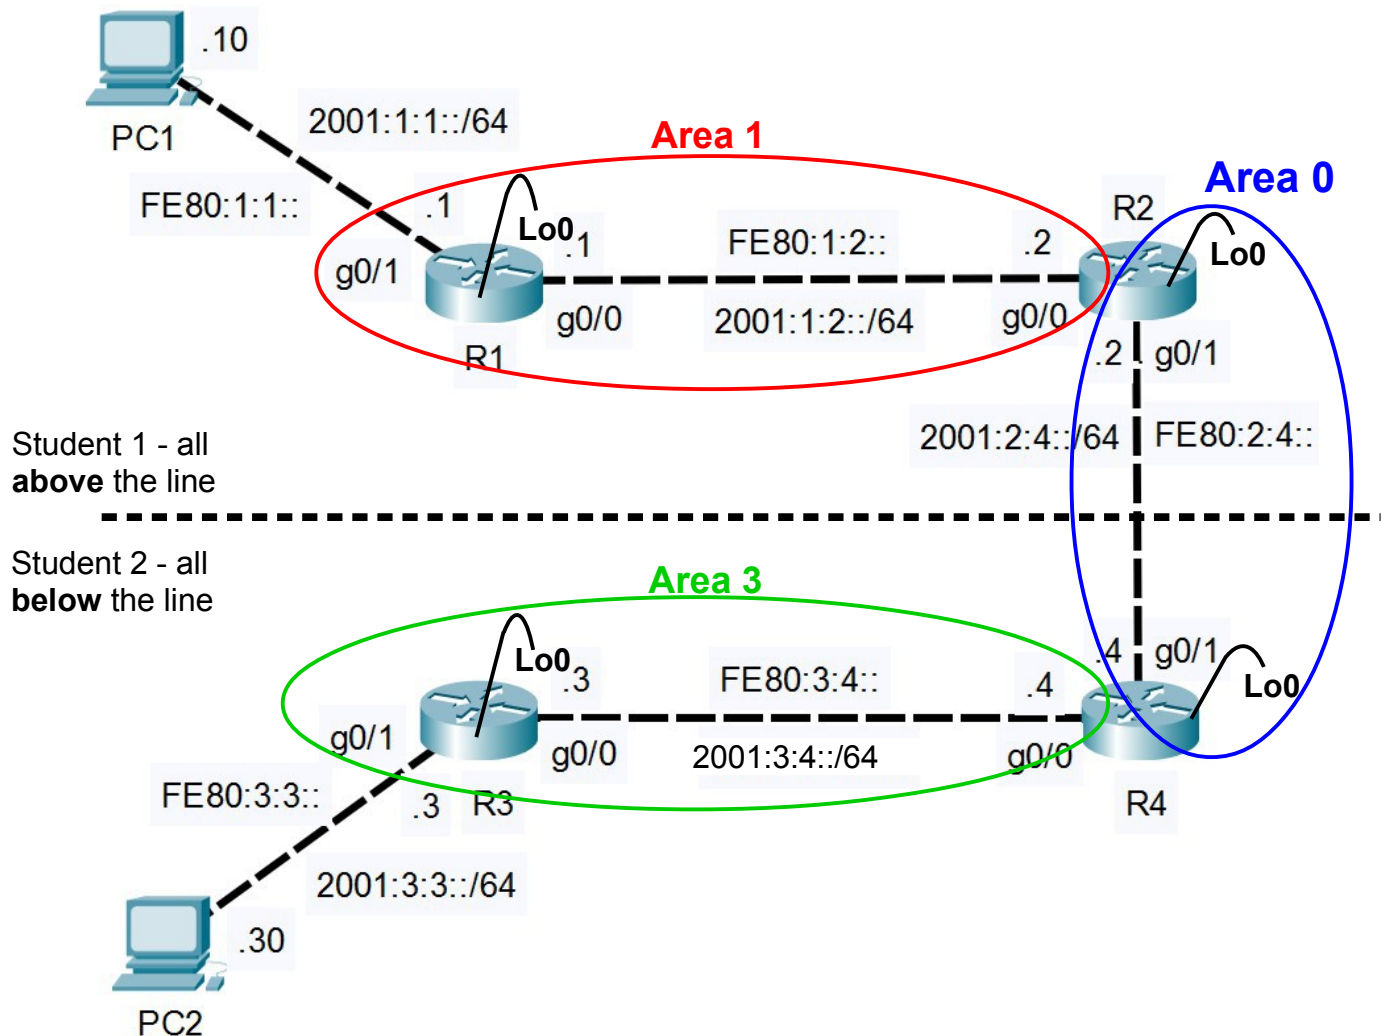
- Process ID is still a plain number in the range 1-65535.
- Router ID is still a plain, 32-bit number, still most conveniently written in dotted-quad decimal format.
- Area numbers are still the same 32-bit values: 0-4294967295.
- Reference bandwidth is the same units, specified the same way.
- Costs are still calculated (or assigned) the same way.
- DR/BDR priorities are still the same values: 0-255.
- Hello/Dead intervals still use the same range, in secs.
- ALL the same "must-match" items for neighbouring still apply!

ALL commands except one are identical, with "ipv6" replacing "ip".

But please remember that IPv4 and IPv6 are still **DIFFERENT** and **completely separate!**

- areas are completely separate for OSPFv2 and OSPFv3
- values for process ID, router ID, reference bandwidth, costs, DR priority, Hello & dead intervals are all **completely separate!**

Topology and Addressing Diagram



Device	Interface	Address	Gateway
R1 (student 1) LHS rack (odd # pods): R1, or R3, or R5	G0/0, G0/1	As per topology	(None needed since OSPF will distribute info of all subnets!)
	Lo0 – Loopback 0	FD00:10:10:10::1/128	
R2 (student 1) RHS rack (even #pods): R1, or R3, or R5	Gi0/0, G0/1	As per topology	(None needed since OSPF will distribute info of all subnets!)
	Lo0 – Loopback 0	FD00:10:10:10::2/128	
PC1 (student 1)	Realtek NIC	Global: as per topology	?
R3 (student 2) LHS rack (odd # pods): R2, or R4, or R6)	G0/0, G0/1	As per topology	(None needed since OSPF will distribute info of all subnets!)
	Lo0 – Loopback 0	FD00:10:10:10::3/128	
R4 (student 2) RHS rack (even #pods): R2, or R4, or R6)	Gi0/0, G0/1	As per topology	(None needed since OSPF will distribute info of all subnets!)
	Lo0 – Loopback 0	FD00:10:10:10::4/128	
PC2 (student 2)	Realtek NIC	Global: as per topology	?

Task 1: Cable the Network

(same as Lab 5)

Cable the devices as shown in the diagram above, to create the required network topology.

Task 2: Configure basic *settings for PCs*

(same as Lab 5)

Working in parallel with your partner, determine the correct gateway address for each PC. Configure the necessary IPv6 settings: globally routable address, mask, and gateway.

Task 3: Configure your two Routers with the essentials

(same as Lab 5)

Still working in parallel with your partner, your two routers need to be configured with the four essential categories of basic config.

- As usual, starting with configs you saved from previous labs, edit them to use IPv6 addresses as shown on the topology diagram. You have a choice of either replacing the IPv4 addresses, or adding new IPv6 addresses.
- Hopefully you still have a "no ip domain-lookup" command to prevent long delays from a mis-typed command. Unlike most commands, that one applies to both IPv4 and IPv6.
- Verify IPv6 connectivity between your PC and its associated router. A successful ping from the PC to the router's loopback is rock-solid confirmation of full connectivity between PC + router.
- **Telnet or SSH** from your PC to its router, so that you can move the console cable to the right-side router. Yes, you can use IPv6 addresses with both Telnet and SSH. Like previous labs, you'll want to have the benefit of side-by-side simultaneous windows to both routers.

Task 4: Configure an OSPFv3 Internal Router

We'll begin with configuring OSPF on just the left-side router (attached to your PC). Prior to joining it the right-side router as a neighbour (next task), it behaves as a simple, single-area router. That means we can use it as a reference setup for comparison with a multi-area scenario.

Configure: Repeat the OSPF config process from the last couple of labs, but this time replacing "ip" with "ipv6" wherever it appears.

- Be kind to us humans: use a consistent process ID throughout. Use "7" since this is Lab7.
- Explicitly set the router ID to match the loopback IP address.
- Use the correct area when including interfaces: **area 1** for the top / **area 3** for the bottom.
- Configure all interfaces in OSPF as point-to-point. Remember the **link-local** addresses!!
- Remember to include the loopback interface!
- Configure the reference bandwidth to accurately calculate costs for speeds up to 1 Tbps.*

Verify:

- Confirm (prove!) OSPFv3 is running and all the interfaces are in OSPF: **show ipv6 protocols**
Compare with the IPv4 version you saved from Lab 5. Is it easy to distinguish the 2 versions?
- Compare the interface info from a more specific command: **show ipv6 ospf interface brief**
Compare with IPv4. What differences if any, can you spot, so you can distinguish the two vers?
- Grab a copy of the full routing table, for comparing after there's neighbour(s): **show ipv6 route**
Again, note differences, similarities, and distinguishing features compared with IPv4.
- Capture and **save** the output of all commands for later reference.

* Food for thought on Ref BW: https://flylib.com/books/en/2.286.1/adjusting_ospf_costs.html

Task 5: Configure a Backbone Router

The *right-side* router is both a backbone router and an ABR. (Be sure you know why!) Configuring it is almost an exact repeat of the previous task, except for which areas to which the interfaces are attached.

Configure:

- Configure OSPFv3; use the same/similar process that you used in the previous task. Remember the **link-local** IPv6 addresses.
- For the moment, **leave G0/1 shutdown** so that top and bottom halves remain separated. (You'll enable the G0/1 interfaces for end-to-end connectivity in Task 7.)
- Be mindful of the areas for the interfaces: G0/0 = **area 1 or 3**; G0/1 + loopback = **area 0**
- Don't forget the reference bandwidth and make all interfaces point-to-point

For troubleshooting, it makes sense to start by confirming that all interfaces including loopbacks are up/up, and then directly connected interfaces of neighbours are reachable via ping.

Verify:

- **Copy & save** OSPF state message(s) from when the OSPFv3 adjacency forms with R1 or R3. i.e. the message which includes the transition from "LOADING to FULL".
- **Copy & save** some output now that there's neighbours and compare with the previous task:
Left side router (R1/R3) – for OSPFv3 / IPv6: full routing table; ip protocols; LSDB
Right side router (R2/R4) – for OSPFv3 / IPv6: full routing table; ip protocols; LSDB

Q: Does the output from the right-side routers give any indication that they are ABRs??

Q: Are there new *types* of OSPF LSAs on the left/right routers? What are their titles?

Q: Is there anything new or noteworthy in the IPv6 routing table? What is the next-hop address for the OSPF routes? How does IPv6 make it easier to know the number of routes?

CHECK POINT: *Point out proof* to show that multi-area OSPFv3 is running successfully.

Task 6: Configure a Default Route

Most networks at the enterprise scale or smaller need a default route. In this task, the *right-side* router becomes an ASBR. (Be sure you know why!) Just like Lab 5, use two additional commands to add a default route to the right-side router (in area 0). Then you'll verify that it shows up in the routing table and LSDB of the left-side route.

Configure:

- Add a "dummy" (or *blackhole*) default route to loopback 0 on the right-side router (R2, R4) (see the pre-lab quiz)
- Add the command so that the router shares the newly installed default route.

Verify:

- Grab (and save!) the (a) full routing table and (b) LSDB for both the left and right-side routers.

Q: Can you identify: the new default route and new LSA(s)? What LSA *type* is the default route? What is similar, and what is different from OSPFv2?

Q: Compared to OSPFv2, what new route markings/types exist in the first column?

Task 7: Join the Networks for End-to-End Connectivity

For full end-to-end connectivity, all that's required is to ensure that G0/1 on both routers (the vertical interfaces) are added into OSPF in area 0 and enabled (no shut).

Configure:

- Ensure the G0/1 interfaces are added in OSPF as point-to-point links. Remember the **link-local** IPv6 addresses.
Q: What command (issued in previous tasks!) confirms this?
- Enable the interfaces on both routers, then wait a dozen secs or so for convergence.

Verify:

- Grab (and save!) the (a) routing table and (b) LSDB for all four routers. What changes do you see? Are they what you expected? Why or why not?
- End-to-End connectivity: use both ping and traceroute.

CHECK POINT: Have a CMD window open showing the number of hops between PCs.

Task 8: Backups and Clearing the equipment

You'll want/need a backup of the config for all four routers for the post-lab. Make backups of your two routers, then exchange configs with your lab partner.

Be sure your equipment is cleared of all config before de-cabling and powering down.

show startup-config If the startup-exists, then: **erase startup-config**

... Did this lab feel smoother or easier than Lab 5? Hopefully the answer is yes. :-)