

PostLab 7: Pure OSPFv3

Plus: Tips on copying & pasting configs

What You Will Do:

1. Prove that OSPFv3 works **entirely using link-local addresses** by eliminating **global IPv6 addrs** from all inter-router links - interfaces are **not** pingable with regular IPv6.
2. Learn a pair of tips for making your configs more understandable and manageable.
3. Find out whether link-local addresses are routed by OSPFv3, and how Cisco handles link costs in OSPFv3.
4. Re-use config saved from previous labs to speed up the configuration for this lab.

Things that you will need to know or learn:

1. How to set up multi-area OSPFv3 (from the main lab exercise)
2. Distinguish (in a heart beat) a link-local IPv6 address from a globally routable IPv6 address

What you need to submit and when:

1. Complete post-lab exercise and quiz on BrightSpace, **before** your next lab.

This post-lab is worth **38%** of this lab, even though the number of points may differ between the three parts.

Required Equipment:

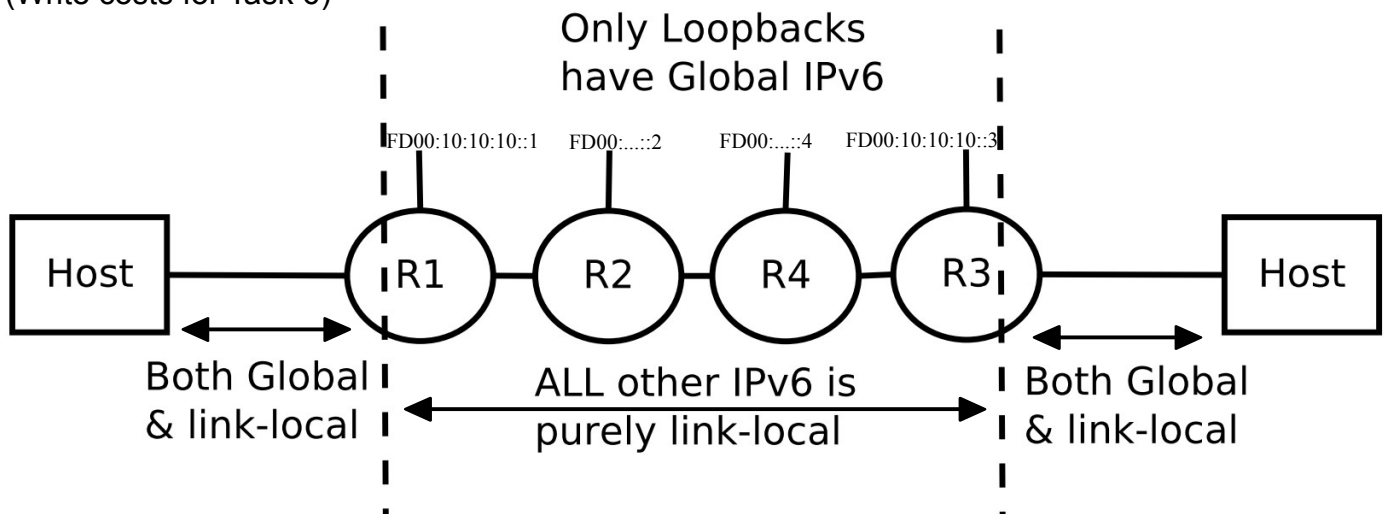
- A laptop or desktop computer, with telnet capability
- Access to Algonquin's NetLab facility

References and Resources:

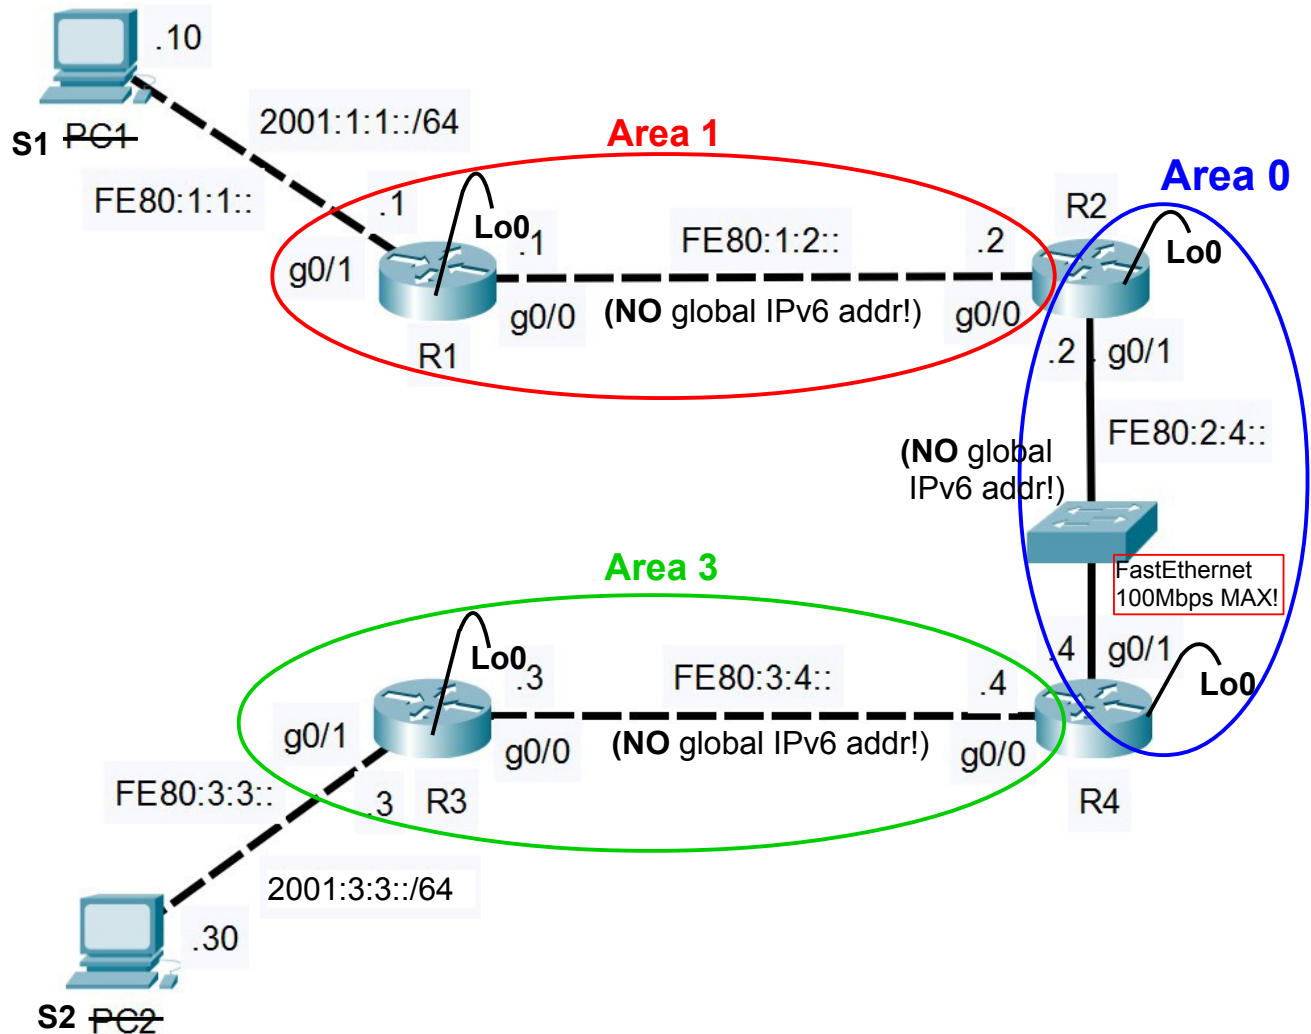
- Slide deck on OSPFv3 networks, posted on course site
- OSPFv3 (free chapter) <https://www.ciscopress.com/articles/article.asp?p=2294214&seqNum=4>
- I/F descriptions: https://petri.com/csc_configuring_router_interface_descriptions_on_a_cisco_router/

Topology Concept for Lab 7 - Postlab

(Write costs for Task 6)



Topology and Addressing Diagram



Device	Interface	Address	Gateway
R1 (model 2911)	G0/0, G0/1	Global addr only on G01	(None needed since OSPF will distribute info of all subnets!)
	Lo0 – Loopback 0	FD00:10:10:10::1/128	
R2 (model 4331)	Gi0/0/0, G0/0/1	NO global IPv6 addrs!!	(None needed since OSPF will distribute info of all subnets!)
	Lo0 – Loopback 0	FD00:10:10:10::2/128	
S1 (a target to ping)	VLAN 1 (via G0/11)	Global IPv6 as per topology	?
R3 (model 2911)	G0/0, G0/1	Global addr only on G01	(None needed since OSPF will distribute info of all subnets!)
	Lo0 – Loopback 0	FD00:10:10:10::3/128	
R4 (model 4331)	Gi0/0/0, G0/0/1	NO global IPv6 addrs!!	(None needed since OSPF will distribute info of all subnets!)
	Lo0 – Loopback 0	FD00:10:10:10::4/128	
S2 (a target to ping)	VLAN 1 (via G0/11)	Global IPv6 as per topology	?

Task 1: Read this lab exercise carefully

In order to get correct answers to some of the key questions, you'll need to follow the instructions *fully and accurately*. For that reason, please study the topology diagram and read the entire lab carefully before starting any configuration!

Suggestion: Right at the beginning of your NetLab session, spend 5-10 mins reading the lab. When you're done, all devices will have booted so you won't need to wait for them to be ready!

Task 2: Create configs for Multi-area OSPFv3 for NetLab

The goal for this task: using an editor, create multi-area OSPFv3 configs like you did for Lab 7 (and very similarly to the OSPFv2 you configured in Lab 5/postlab 5). The config should be identical to Lab 7 except that all inter-router globally routable IPv6 addresses are removed.

Just like post-Lab 5, you'll use S1-S2 in place of PC1-PC2. If you haven't gotten into the habit of saving your config from all lab exercises, this is another reminder that you *should be* saving it!

- Just like the Lab 5 - Postlab 5 pair, the configs from Lab 7 for R1-R3 are almost *identical*, as are the configs for R2-R4. Ideally you can simply edit your configs from Lab 7 to get what you need for this lab.

Here are guidelines for what needs to be checked and/or modified:

- hostname
- interface names (either two digits or three: g0/0 vs **g0/0/0** for R2, R4)
- OSPF **area numbers** for all interfaces
- ALL Loopbacks: **must** have globally routable IPv6 addresses
- Inter-router links: **only** link-local addresses and **no globally routable** addresses for the links R1-R2, R2-R4, R3-4.
- ALL link-local addresses (S1-2, R1-4): should be manually set so that it's easy to distinguish one from another and there's no chance of confusing them.
- Host subnets: the "PC" subnets **must** have globally routable addresses on both ends (S1-R1, S2-R3).
- For NetLab, don't forget to *disable* all unused switch ports on S1, S2, and 4-port switch in R1, R3. See the last page of IntroToNetLab.pdf document for info.

Anyone coming along later will think you're crazy for omitting ("forgetting") the global IPv6 addresses. To prevent getting fired from your job later in life, it will be essential to *document* what you did and why, and do this directly in the config. (Kinda like commenting computer code!)

Here's how you do that:

! Best practice: Always comment your programming code and your configs for network devices
! On a Cisco device, this line is ignored because it starts with a "!" symbol
! This style of comment works for both PT and real equipment. Test it out!
The hash or number symbol is also a comment in many, many, many programming languages
It's also used as a comment by many network vendors: Cisco, Nokia, Juniper, MikroTik, etc
It's a comment for real equipment from Cisco, but unfortunately not in PT!

- Document your config to prevent confusion about why there's no global IPv6 addresses on the *inter-router links* i.e. this is a demo/learning exercise about OSPFv3 with link-local only. Use *at least 1 comment line* for every interface (minimum 6 comments for the 4 routers).

Task 3: Deploy & Troubleshoot Multi-area OSPFv3

You'll need a booking on Algonquin's NetLab facility, for a **full pod**. Copy & paste your configs, starting with simplest ones e.g. S1, S2, and then move along to R1-R4.

WATCH the output on the terminal screen from pasting your config. You need to spot and fix any errors or else you'll have no end of problems later on!

Spoiler Alert! Even without the global IPv6 addresses on the inter-router links, you should be able to ping end-to-end from S1-S2 if you've done everything correctly. In addition to getting pings through the routers, you should also be able to ping the routers themselves using their loopback addresses.

Test everything (it's only 5 pings in total: end-to-end + 4 loopbacks). If any ping doesn't respond, choose one of the troubleshooting methods from Ch 12 to help with finding & fixing the problem. Follow-the-path is probably a good one.

We'll use this working config in later tasks to examine the operation of OSPFv3.

Q. With IPv4 and OSPFv2, could you remove all the assigned addresses from inter-router links and still have end-to-end connectivity?

Q. With IPv4 and OSPFv2, could you remove all the assigned addresses from inter-router links and even reach the neighbour router at the far end of the link?

Task 4: Documenting your Configurations

There's one problem with the form of comments used in Task 2: they're not permanent. Let's have a look at fixing that.

- Display the config retained by the router with *show run* on each of the 4 routers. Look for the comments (with "#" and/or "!") that you created in Task 2. Notice they're now GONE. That makes sense if you consider that the "#" and "!" character means "ignore this line". That's exactly what happened when those lines were pasted: ignored and forgotten.

In Cisco's IOS, there's exactly and only two things that can have permanent comments attached:

1. Interfaces can be *permanently* commented by means of the 'description' command.
2. ACLs can be commented by means of 'remark' statements.

How long can an interface description be? Find out!

- In the interface configuration context (e.g. g0/0), issue the command 'description This is an fast interface'. Verify the description is persistent using 'show run'.
- Type a longer description. No, even longer than what you're thinking of. No, even longer than that! Keep trying with longer and longer descriptions until you get a message saying that characters in excess of XXX will be discarded. (Note: there's a different limit in PT).
- Add a description to all the other interfaces on that router. See what they look like:
show interface description

Q. What is the longest description that can be entered (on real hardware)?

Q. How practical are loooooong descriptions i.e. when printed out, do they stay in tidy columns or do they "break" the intended format of the 'show int desc' output?

Q. While we're on the topic, what is the character limit for comments on ACLs? Use "?" to find out: type **access-list 100 remark ?** to get the limit.

Q. Could commenting I/Fs and ACLs be useful with an automated system (think Dev Ops)?

Task 5: Routing, OSPFv3, and Link-local addresses

It's useful and important to understand the role of link-local addresses in OSPFv3. This task looks at what role link-local addresses play in the routing table, and can they be a (a) destination; (b) a next-hop; and (c) routable. You *must* have completed Task 3 successfully (end-to-end connectivity) in order to get the correct results for this task.

- Step 1. On R1 and then on R3, display the IPv6 routing table. For the next-hop address, what form of address(es) is/are used - globally unique or link-local? Or a mixture?
- Step 2. On R2 and then on R4, display the IPv6 routing table. For the next-hop address, what form of address(es) is/are used - globally unique or link-local? Or a mixture?
- Step 3. On R1-R4, how many link-local addresses/subnets are listed as a *destination*?
- Step 4. The link-local addresses for each interface are known (specified in topology and configured by you!) From R1, then R2 attempt to ping each of the other routers' link-local interfaces. (See Lab 3 for how to ping IPv6 link-local addresses!) Complete the chart.

Ping successful?	R2 – fe80:1:2::2	R4 – fe80:2:4::4	R3 – fe80:3:4::3	S2 – fe80:3:3::30
R1	?	?	?	?

Ping successful?	R1 – fe80:1:2::1	R4 – fe80:2:4::4	R3 – fe80:3:4::3	S2 – fe80:3:3::30
R2	?	?	?	?

- Q. What form of next-hop addresses are used in an IPv6 routing protocol like OSPFv3?
- Q. Which link-local addresses/subnets are listed as a *destination* in the routing table?
- Q. Did you capture the complete set of ping results for the two tables above?
- Q. What is the *reason* a ping will succeed or fail in the above tests?
- Q. Can link-local addrs be: the destination subnet for a route; the next-hop; and/or routable?

Task 6: OSPFv3 Path Costs

It's easy to show the cost for each link and make a map (diagram) in order to calculate the cost for a complete path. With one exception, Cisco has implemented everything exactly as you would expect. If we're going to say we "know Cisco equipment", it's necessary to know even the crazy details of exceptions. In this task, you'll find that Cisco's implementation of OSPFv3 has one "irregularity", which fortunately isn't present in their implementation of OSPFv2.

- Step 1. On each router, use the appropriate 'show...' command to display the table of OSPFv3 interfaces which includes the link cost. For each link, confirm that the routers on both ends agree on the cost of the link. Don't forget the loopbacks!
Did you **notice** that R2-R4 are joined by a 100Mbps switch, making a 100Mbps link?
- Step 2. Using some handy diagram (i.e. cover page), make a map of all the costs. Every straight line on the diagram should have a cost (including loopbacks).
- Step 3. Complete the tables below of costs to all destinations, originating from R1 and R2.

Path Cost (Calculated)	R2 (to loopback)	R4 (to loopback)	R3 (to loopback)	S2
From R1 to ...	?	?	?	?

Path Cost (Calculated)	R1 (to loopback)	R4 (to loopback)	R3 (to loopback)	S2
From R2 to ...	?	?	?	?

Step 4. On R1, R2, display the full IPv6 routing table. If you've completed Task 3 successfully, you should see routes to all the loopbacks and the two "PC" subnets. Recall that each route has preference (AD) and metric info in "[...]". Copy the displayed values into the charts below.

Path Cost ('show ...')	R2 (to loopback)	R4 (to loopback)	R3 (to loopback)	S2
From R1 to ...	?	?	?	?

Path Cost ('show ...')	R1 (to loopback)	R4 (to loopback)	R3 (to loopback)	S2
From R2 to ...	?	?	?	?

Step 5. Compare the tables from the previous two steps. Are the values identical? If not, can you identify which link cost isn't what you expected (i.e. is it off by 1, 1000, or 10,000)?

Step 6. Hopefully you're keeping your results from all labs and post-labs for study purposes but in case you haven't, the relevant results from Lab 5 for OSPFv2 are included on the next page. Do the OSPFv2 results match up with your calculated results?

Step 7. The basic SPF math must be correct, so verify what information is actually being shared in the relevant LSA. To see the details, on R1 do:

```
show ipv6 ospf database prefix
```

Scan the output to find the line containing the "Metric" (i.e. cost). What value do you see? That line, and the preceding line, contain the info for which interface/subnet is being advertised. Make sure you are clear about which one it is.

Just to make sure there's no mistakes, repeat this step on R3.

Don't forget to save the results from both routers.

Step 8. We've covered how to change the cost for an interface in OSPF. Manually change the cost value to 1234 (very recognizable). Wait 10 secs for various timers to expire, then repeat the previous step to check the cost being advertised in the LSA. Any change(s)?

Step 9. Try another manual value for the cost of 43210. Wait 10 secs and check the LSA. What change, if any, do you see?

Step 10. Now here's the glitch in Cisco's implementation of OSPFv3: it's not an issue with the *type* of interface (loopback vs hardware), it's an oddity associated with any interface that is a 'host' route (ie. a mask length of /128).

Step 11. Isn't Cisco fun?! But wait before you say that; every vendor of networking equipment has oddities and hiccups in their implementations, and so this is the stuff of networking! What will make you valuable is your (deep!) knowledge of these oddities and your ability to overcome any challenges they present.

Q. Did you complete both sets of tables above?

Q. For OSPFv2, are the calculated and actual path costs identical?

Q. What is the value of the advertised cost that's unexpected?

Q. Can you manually override the cost to get a different advertised cost?

If yes, by how much can you change it?

Task 7: Backups

Is it even necessary to suggest that it's useful and important to have full backups of all your configs and output from 'show ...' commands? :-)

Routing tables from Lab 5 - multi-area OSPFv2

R1(config)#do show ip route

Gateway of last resort is 10.1.2.2 to network 0.0.0.0

```
O*E2 0.0.0.0/0 [110/1] via 10.1.2.2, 00:00:04, GigabitEthernet0/0
      10.0.0.0/8 is variably subnetted, 8 subnets, 2 masks
C      10.1.2.0/24 is directly connected, GigabitEthernet0/0
L      10.1.2.1/32 is directly connected, GigabitEthernet0/0
O IA   10.2.4.0/24 [110/11000] via 10.1.2.2, 00:05:57, GigabitEthernet0/0
O IA   10.3.4.0/24 [110/12000] via 10.1.2.2, 00:05:41, GigabitEthernet0/0
C      10.10.10.1/32 is directly connected, Loopback0
O IA   10.10.10.2/32 [110/1001] via 10.1.2.2, 00:05:57, GigabitEthernet0/0
O IA   10.10.10.3/32 [110/12001] via 10.1.2.2, 00:05:41, GigabitEthernet0/0
O IA   10.10.10.4/32 [110/11001] via 10.1.2.2, 00:05:41, GigabitEthernet0/0
      192.168.1.0/24 is variably subnetted, 2 subnets, 2 masks
C      192.168.1.0/24 is directly connected, GigabitEthernet0/1
L      192.168.1.1/32 is directly connected, GigabitEthernet0/1
O IA   192.168.3.0/24 [110/13000] via 10.1.2.2, 00:05:41, GigabitEthernet0/0
```

R2(config)# do show ip route

Gateway of last resort is 0.0.0.0 to network 0.0.0.0

```
S*    0.0.0.0/0 is directly connected, Loopback0
      10.0.0.0/8 is variably subnetted, 9 subnets, 2 masks
C      10.1.2.0/24 is directly connected, GigabitEthernet0/0/0
L      10.1.2.2/32 is directly connected, GigabitEthernet0/0/0
C      10.2.4.0/24 is directly connected, GigabitEthernet0/0/1
L      10.2.4.2/32 is directly connected, GigabitEthernet0/0/1
O IA   10.3.4.0/24 [110/11000] via 10.2.4.4, 00:07:08, GigabitEthernet0/0/1
O      10.10.10.1/32 [110/1001] via 10.1.2.1, 00:07:24, GigabitEthernet0/0/0
C      10.10.10.2/32 is directly connected, Loopback0
O IA   10.10.10.3/32
      [110/11001] via 10.2.4.4, 00:07:08, GigabitEthernet0/0/1
O      10.10.10.4/32
      [110/10001] via 10.2.4.4, 00:07:08, GigabitEthernet0/0/1
O      192.168.1.0/24 [110/2000] via 10.1.2.1, 00:07:24, GigabitEthernet0/0/0
O IA   192.168.3.0/24 [110/12000] via 10.2.4.4, 00:07:08, GigabitEthernet0/0/1
```