

## Lab 5: Multi-area OSPF

*Or. Digging into details of OSPF timers, passive interfaces, link costs*

### What You Will Do:

1. Re-use config saved from post-lab 4 and Lab 5 to complete this post-lab quickly & easily. (You did save your config, didn't you? **ALWAYS** save your config for future use!)
2. Observe (industry standard) timers for OSPF.
3. Observe and understand OSPF costs for different speed links.
4. Optimize OSPF through the use of passive interfaces.

### Things that you will need to know or learn:

1. How to make backups of your work, reliably and consistently, for future work.
2. A NetLab booking for a **full pod**, for a time when you have 2-3 hrs of time

### What you need to submit and when:

1. Complete post-lab exercise and quiz on BrightSpace, **before** Thu Oct 13 @ 11:59pm.

### Required Equipment:

- Access to Algonquin's NetLab facility

### Marks:

The post-lab is worth **35%** of the total lab mark.

### References and Resources:

- Slide deck on multi-area OSPF networks, posted on course site
- SPF throttling, to prevent overloading the CPU:  
[https://www.cisco.com/en/US/docs/ios-xml/ios/iproute\\_ospf/configuration/15-0s/iro-short-path.html](https://www.cisco.com/en/US/docs/ios-xml/ios/iproute_ospf/configuration/15-0s/iro-short-path.html)

### Command List

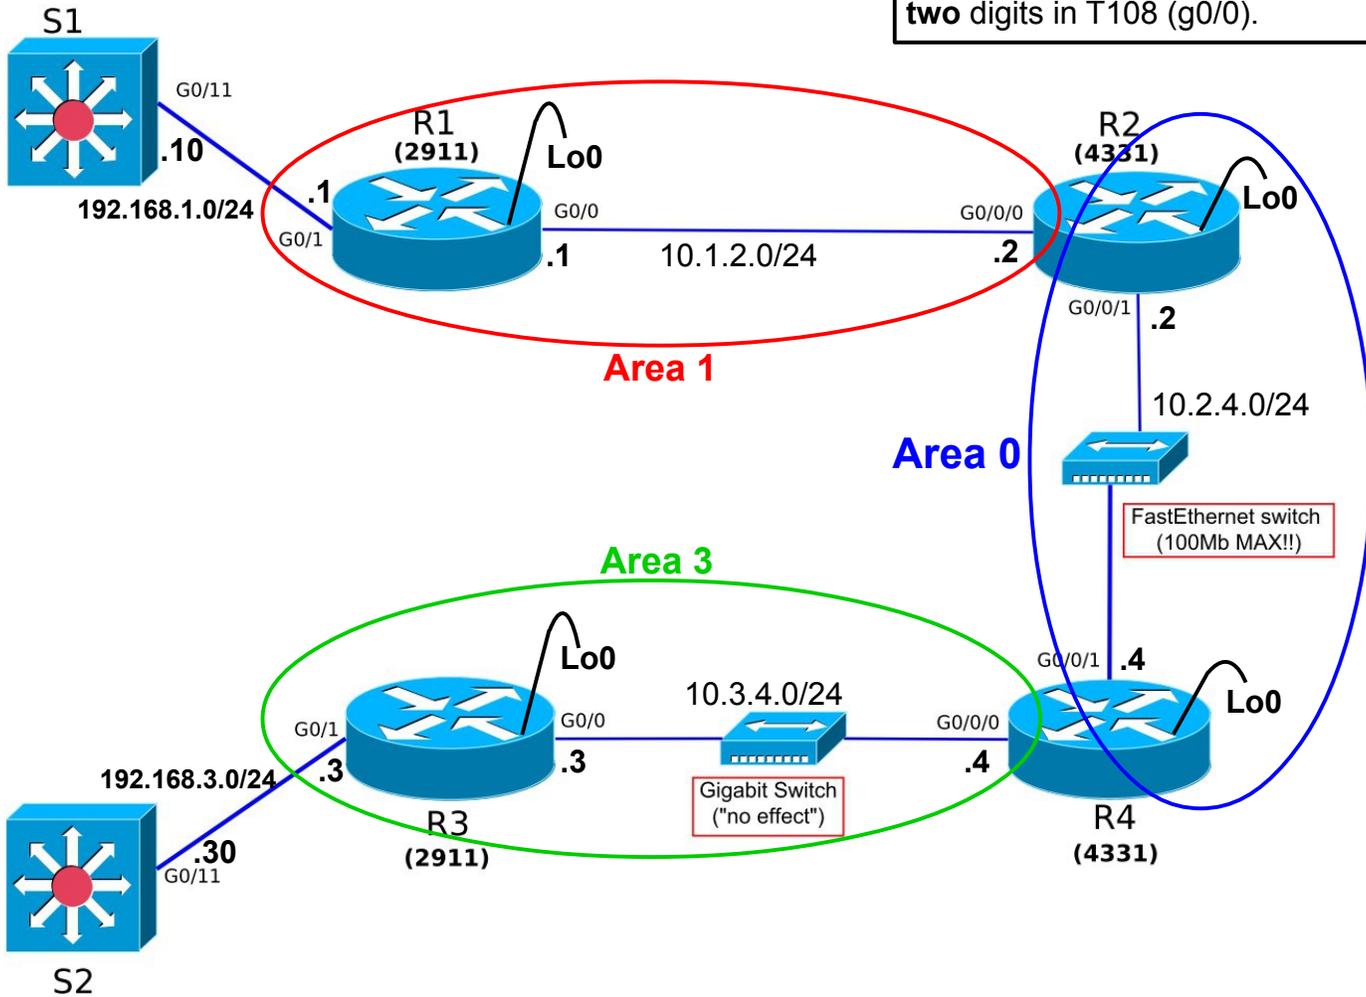
```
router ospf {pID}
router-ID {rID}
auto-cost reference-bandwidth
default-information originate
passive-interface {i/f-name}                (new)

ip ospf {pID} area {X}
ip ospf network point-to-point

show ip route [ospf]
show ip protocols
show ip ospf [...possibly with options...]
show ip ospf interface [brief]
show ip ospf neighbor [detail]
show ip ospf database
```

# Topology and Addressing Diagram

For R2, R4 the interfaces have:  
**three** digits with NetLab (g0/0/0);  
**two** digits in T108 (g0/0).



Device	Interface	Address	Gateway
<b>R1</b> (model 2911)	G0/0, G0/1	As per topology	(None needed since OSPF will distribute info of all subnets!)
	<b>Lo0</b> – Loopback 0	<b>10.10.10.1/32</b>	
<b>R2</b> (model 4331)	G0/0/0, G0/0/1	As per topology	(None needed since OSPF will distribute info of all subnets!)
	<b>Lo0</b> – Loopback 0	<b>10.10.10.2/32</b>	
<b>S1</b> (a target to ping)	VLAN 1 (via G0/11)	192.168.1.10/24	?
<b>R3</b> (model 4331)	G0/0, G0/1	As per topology	(None needed since OSPF will distribute info of all subnets!)
	<b>Lo0</b> – Loopback 0	<b>10.10.10.3/32</b>	
<b>R4</b> (model 4331)	Gi0/0/0, G0/0/1	As per topology	(None needed since OSPF will distribute info of all subnets!)
	<b>Lo0</b> – Loopback 0	<b>10.10.10.4/32</b>	
<b>S2</b> (a target for ping)	VLAN 1 (via G0/11)	192.168.3.30/24	?

## Task 1: Configure/Complete Multi-area OSPF on NetLab

You'll need a booking on Algonquin's NetLab facility, for a **full pod**. Then as preparation, please configure complete multi-area OSPF as per Lab 5. Just like post-Lab 4, use S1-S2 in place of PC1-PC2. If you didn't save your configs from those two exercises, you'll need to re-do the config. (But that can just be extra practice for the SBA to make sure you're really quick with configuring!)

Note that configs for R1-R3 are almost *identical*, as are the configs for R2-R4. So you do *not* need your lab partner's configs since a few, quick edits are sufficient to (re-) generate them.

As a tip, first load up the configs in notepad (or any other text editor), and check / modify the following items so that they match the topology given on the previous page.

- hostname
- interface names (either two digits or three: g0/0 vs **g0/0/0**)
- ip addresses for each address (the "PC" subnet, g0/0, g0/1 interfaces)
- OSPF **area numbers** for all interfaces

Don't forget to *disable* all unused switch ports on S1, S2, and even the mini-switch within R1 (G0/1/x) and R3 (G0/3/x). If need be, see the last page of IntroToNetLab.pdf document for info.

## Task 2: Check the OSPF Interfaces

Since OSPF is a *discovery* protocol, perhaps the single most important element for proper configuration is the interfaces. They need to be included in OSPF, with the correct settings. Let's examine them more closely, using the command to show the list (in brief) of OSPF interfaces.

- On each router, confirm that all the interfaces are active in OSPF. With our purpose-built topology, that should be three (3) interfaces for each router. If not, check your config!
- With this format of output, it's super easy to scan down different columns for verification. Are all interfaces in the same OSPF process? Is each interface in the correct area? Are all interfaces point-to-point? Do appropriate interfaces have neighbour(s)? ... The answer to each of these questions should be 'Yes' and it can be quickly confirmed!
- Normally we would check that the costs are correct, but we'll leave that for a later task.

Based on the verification above, you can troubleshoot if you don't have the expected results.

**Q.** Which column identifies whether interfaces are in Broadcast mode or point-to-point?

## Task 3: Confirm Full & Correct Operation

You can either consider this task as a simple confirmation, or you can consider it as "Establishing a Network Baseline" (reference: Wk05-Troubleshooting, Slide 8 in the section on Network Documentation). Hopefully it's obvious that you'll need the network fully functional or else you can't proceed to the remaining tasks!

There's usually multiple methods to confirm proper operation, and given this is a small network, let's do several of them since it's so quick!

Tip: if traceroute is slow, check that you remembered *no ip domain-lookup* .

- **Full, end-to-end connectivity.** Though we might like to use ping, Cisco routers don't provide the TTL, so use traceroute from S1-S2 instead to confirm the number of hops. Looking at the topology diagram, it's hopefully easy to see that there are five (5) hops. What's the single, most common delay for the traceroute responses? (Note that we're running a very small, very quiet network!)
- **Check the LSDB** on all four routers. With our topology, R1 should parallel R3, and R2-R4. How many routers are in Area 1 (or 3)? How many type 1 LSAs? Type 2's? Type 3's? Do the backbone routers have the same number of LSAs in their LSDB for areas 1 or 3? How many routers in Area 0? How many Type 1's? How many type 2, 3, 5 LSAs? (Tips: if you've configured each of R2, R4 correctly with a default route, there should be two type 5's. For today, we'll ignore any Type 4 LSAs - *Summary ASB Link States*.)
- **Check the full routing tables** on each of the four routers (not just OSPF routes). On each router, what's the total number of routes and how many are from OSPF? (A route always has a protocol identifier in the first few characters of the line.) Is the number & type of routes on R1 identical to R3? Is R2 identical to R4? On R1 or R3, how is the *default route* marked i.e. first 4 chars of line? Their meaning? How do the routing tables of Internal routers (R1, R3) compare to ABRs (R2, R4)?

Q: What is the most common delay for the traceroute responses?

Q: How many LSA's of each Type (1, 2, 3) are in the LSDB for each area?

Q: What is the total number of routes appearing on R1 or R3?

Q: What is the total number of routes appearing on R2 or R4?

Q: What is the meaning of the markings for the default route on R1, R3?

## Task 4: Examine Additional OSPF Status

Establishing a Network Baseline is essential for maintaining and troubleshooting any non-trivial network. Fortunately a single command gives us much of the basic status information for OSPF.

Everything, everything on a router happens at time intervals controlled by timers. (Did I mention everything?) It will be increasingly critical for your work in networking that you really get that concept. In many (most? all?) cases, these timer values can be seen using "show..." commands. For example, we'll look timers for determining how often the SPF calculation is run for OSPF.

For the overall status of OSPF, including timers, on R2 do: `show ip ospf` and then steps 1-4.

- Step 1. What is the Initial delay timer? i.e. After the initial start-up of OSPF or after a "long" period with no LSAs, how long does OSPF wait to collect LSAs before running **SPF**? (A single network failure might generate many LSAs; try to process them once, all together.)
- Step 2. What is the minimum (hold time) and maximum (wait time) time between re-running the **SPF** calculation? i.e. Since **SPF** can be CPU intensive, don't run it "all the time".
- Step 3. How many times has the **SPF** calculation executed in each area? If this number is stuck at zero, it certainly indicates a problem with the local router. On the other hand, if the number starts shooting through the roof, it indicates a larger problem in the network.
- Step 4. What is the number of External LSA? Does this correspond to the number of Type 5 LSA's observed in the previous task?

For other details of OSPF status, on R2 do: `show ip protocols` and then steps 5-7.

- Step 5. Find the line where the Router ID is displayed. Is it correct?
- Step 6. Near the Router ID, find what kind(s) of OSPF router R2 is acting as. Exactly how is it written out?
- Step 7. Again, within a couple of lines, how many areas is R2 participating in?

- Q. On our Cisco routers, what is the default OSPF initial delay?
- Q. On our Cisco routers, what is the delay when there's lots of network changes?
- Q. Should you ever expect router config changes to be entirely completed instantly?
- Q. How many External LSA are reported in R2's overall OSPF status?
- Q. How does Cisco display info on what kind of OSPF router a device is?

## Task 5: Passive Interfaces

One way to optimize OSPF is to designate interfaces as passive if you (as the network designer) know that there won't be any OSPF neighbours. We'll make some interfaces passive and confirm that full reachability is maintained, while reducing the number of unnecessary Hello packets.

- Step 1. For comparison with later steps, grab some output from R1 and save it:  
`show ip protocols` and `show ip ospf int br`
- Step 2. R1 only has "PC" (as simulated by S1) attached to G0/1. They're not routers and don't run OSPF. Make G0/1 passive (check cover page or slide deck for help if needed). Are any status messages printed on the console? If yes, record them.
- Step 3. What, if anything has changed?? Re-run the two show commands from step 1 and compare the output. Find the new section from "ip protocols". What changes, if any, appear in the brief info on OSPF interfaces?
- Step 4. From any other device, ping to confirm that the host S1 is still reachable.
- Step 5. Experiment with an interface that has OSPF neighbours: make G0/0 passive. Are any status messages printed on the console? If yes, record them.
- Step 6. Make G0/0 active again by issuing a "no passive-interface ..." command. Are any status messages printed on the console? If yes, record them.
- Step 7. Now do some thinking. What other interface on R1 could be made passive without any ill effects? Make that interface passive and see what changes occur (console msgs, OSPF interface info, ip proto info, still reachable by ping).

Reflection: on Cisco routers, is there a whole lot of indication that an interface is passive, unless you specifically go looking for that info?

- Q. What status messages appear when an interface with no OSPF neighbors is set passive?
- Q. What is the title of the new section that appears in "ip protocols" when there's passive i/f's?
- Q. What changes are visible in the brief info on OSPF interfaces with passive interfaces?
- Q. What console message appear, if any, when an interface with OSPF neighbours is made passive?
- Q. What other interface on R1 has no neighbours and can be made passive?

## Task 6: Link Costs

NetLab provides a convenient way to have a look at link costs for different speeds of links, including loopbacks, gigabitEthernet, and fastEthernet links.

- Step 1. Do a few quick end-to-end pings to different interfaces to confirm that everything's working. If not, double-check that all interfaces are non-passive and included in OSPF.
- Step 2. On R1 or R3, display the brief OSPF interface info, to obtain a nice column of costs. What's the cost to reach the loopback link?
- Step 3. Change the reference bandwidth so that it's 1/10 the value i.e. 100Gbps. Display OSPF costs and check which costs have changed. There should be at least one! What is the new cost for the loopback?
- Step 4. Change the reference bandwidth so that it's 4x the original value i.e. 4 Tbps. Display OSPF costs and check which costs changed. Again, at least one should change.
- Step 5. Set the reference bandwidth back to it's original value i.e. 1 Tbps. Verify that the link costs are back to their original values as seen in step 2.
- Step 6. R2-R4 are connected via a 100Mbps switch. How does that affect the link speed and cost? On R2 or R4, display the OSPF costs. Compare the costs of G0/0 and G0/1. Are they the same or different? If different, confirm whether the values correspond to the formula:  $\text{cost} = \text{Ref-BW} / \text{link-speed}$  ? (Just cancel lots of zeros on the top & bottom.)
- Step 7. Back on R1, show the routing table and find the route to R3's loopback. The item in "[ ]" is the protocol preference (for Cisco, OSPF = 110) followed by "/" and the total cost.
- Step 8. Based on costs for links observed in steps 2 & 6, do a manual calculation of the total path cost R1-R3:  $\text{link R1-R2} + \text{link R2-R4} + \text{link R4-R3} + \text{loopback link} = ??$   
Does that match the cost in step 7 Repeat this test for several other paths.

Reflection: Is calculating OSPF path costs terribly difficult? Note that the choice of the Reference bandwidth can make quite a difference!!

- Q. What's the OSPF cost for reaching a loopback interface?
- Q. Does the loopback cost depend on the OSPF reference bandwidth?
- Q. What is the link cost for a 100Mbps link with a 1 Tbps reference bandwidth?
- Q. Is the link cost based on the interface type / speed, or the operating speed of the link?
- Q. What is the total path cost from R1-R3 for our NetLab topology?

## Task 7: 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**