Lab 4: Advanced OSPFv2
Or: *Tweaking OSPF behaviour*

What You Will Do (WYWD):
- 0. **Work as a _team_ of 2**, to share the required typing and configuration
  e.g. student 1 configures R2, R3, PC1; student 2 configures R1, R4, PC2
- 1. Complete basic configuration including loopbacks, and enabling OSPF.
- 2. Use different methods to enable an interface in OSPF.
- 3. Configure and verify OSPF passive interfaces.
- 4. Manually set the router-id in an OSPF-enabled router.
- 5. Configure an interface priority to influence the DR election.
- 6. Change the default reference bandwidth of an OSPF-enabled router.
- 7. Manually set an interface cost in Cisco routers and observe the effect on the routing table.
- 8. Configure a router to propagate a default route
- 9. Use a variety of commands to **verify** basic and advanced **OSPF configuration**.
- 10. As always, clear the equipment of any configuration before leaving the lab.

Things that you will need to know or learn:
- 1. OSPF configuration and "show..." commands from Module 2 of the course materials
- 2. Essential device configuration – including basic security, SSH, vty access
- 3. The two methods of activating/including interfaces in OSPF for Cisco routers
- 4. How to transfer configs between different environments: physical lab, remote access, PT

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 "Lab 4 Post-lab" exercise and quiz on BrightSpace, **before** your next lab.

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.

In-Lab Marks:
[1 mark] Full end-to-end connectivity, including loopbacks, as shown in the routing table (Step 2)
[1 mark] Show the DR/BDR elections roles and explain how they were determined (Step 5)
[1 mark] Show the routing tables for all routers and identify the default route (Step 8)

The pre-lab is worth **33%**, the in-lab is worth **33%**, and the post-lab is worth **34%** of this lab, even though the number of points may differ between two parts.

References and Resources:
- All the materials used in your previous networking courses, including https://netacad.com
- Packet Tracer ver 8 (available from NetAcad)
- Algonquin's NetLab facility – http://netlab.algonquincollege.com
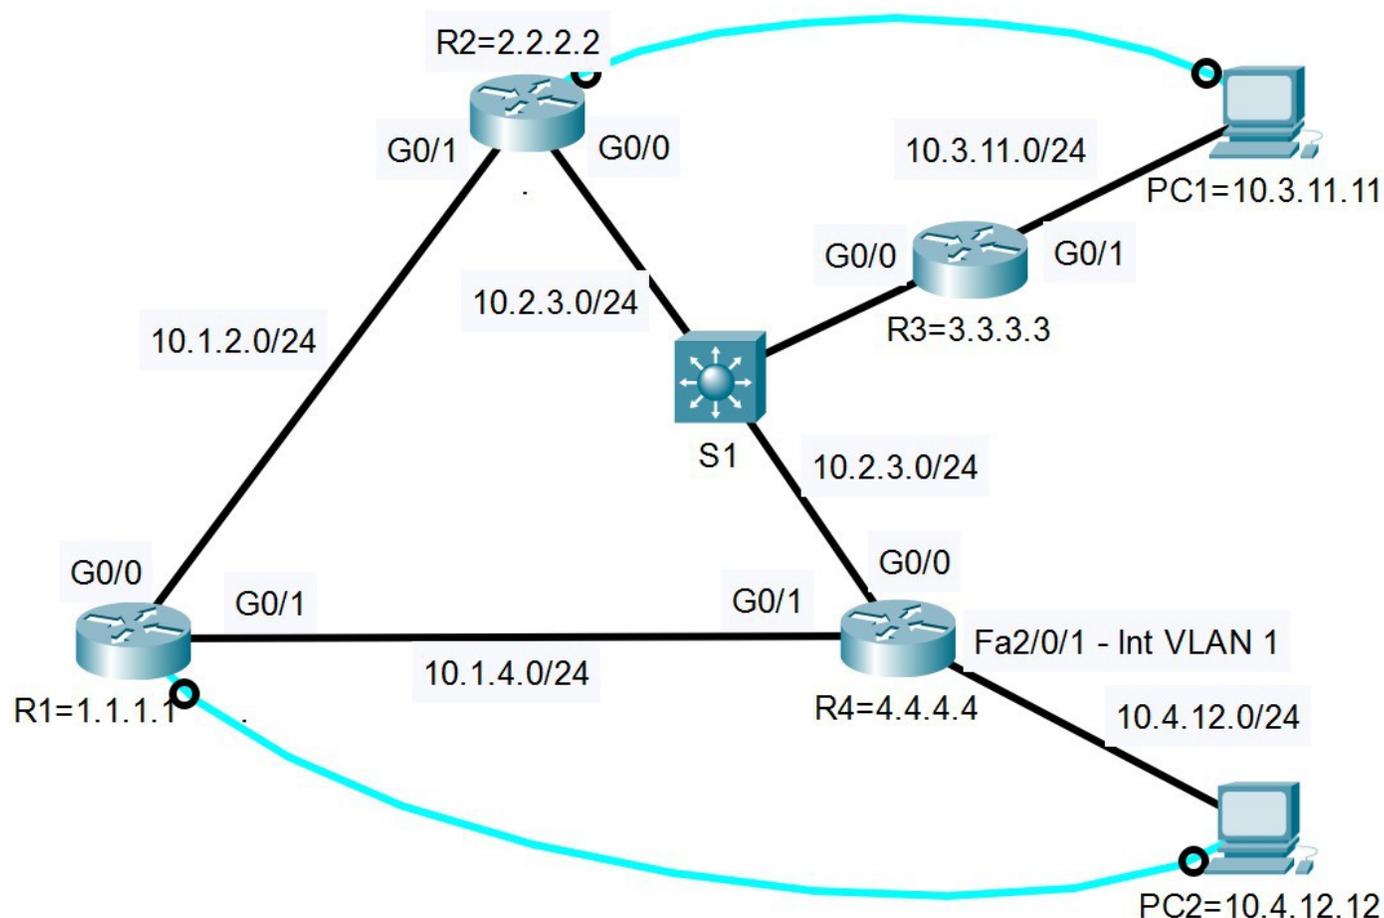
Addressing Table

Device	Interface	IP Address	Subnet Mask	Gateway
R1	G0/0	10.1.2.1	255.255.255.0	–
	G0/1	10.1.4.1	255.255.255.0	
	Loopback1	10.10.10.1	255.255.255.**255**	
R2	G0/0	10.2.3.2	255.255.255.0	–
	G0/1	10.1.2.2	255.255.255.0	
	Loopback1	10.10.10.2	255.255.255.**255**	
R3	G0/0	10.2.3.3	255.255.255.0	–
	G0/1	10.3.11.254	255.255.255.0	
	Loopback1	10.10.10.3	255.255.255.**255**	
R4	G0/0	10.2.3.4	255.255.255.0	–
	G0/1	10.1.4.4	255.255.255.0	
	Fa2/0/1 – **Int VLAN 1**	10.4.12.254	255.255.255.0	
	Loopback1	10.10.10.4	255.255.255.**255**	
PC1	Realtek	10.3.11.11	255.255.255.0	?
PC2 **or** VM1	Realtek **or** Onboard	10.4.12.12	255.255.255.0	?

Task 1: Complete the Basic Network Setup

As we must do whenever working with real equipment, you'll need to cable the network and complete the basic setup. **You'll start to notice how much time you can save if you keep a copy of all required commands in a text file for ease of copying & pasting!! Please do so!** Note that you'll need to repeat steps 1-4 for each of the 4 routers. In this lab, we'll skip all configuration of S1 and use it as a simple unmanaged switch.

Step 1. **Fundamentals**: (a) hostname; (b) disable DNS lookup to prevent (long, slow) lookups from incorrect commands; (c) set console logging synchronous; (d) MOTD clearly restricting access to authorized users

Step 2. **Security**: (a) Exec ~~password~~ secret of **cisco**; (b) ensure plaintext passwords are all encrypted automatically; (c) console password of **cisco**; and (d) console **exec-timeout 0**

Step 3. **IP addressing**: (a) all physical interfaces; and then loopback interface Lo1

Step 4. **VTY access** via both SSH and telnet (to eliminate shuffling the console cable in later tasks). (a) SSH requires **5 additional commands** beyond what's already configured, and (b) telnet access requires modifying a single line for one of those commands.

Step 5. Double-check that you've configured all routers correctly; no typos!

Step 6. **VERIFY** you can ping each adjacent router's physical interface before continuing!

Step 7. Determine the appropriate default gateway settings, then set up the IP addressing for both PCs. **VERIFY** that you can ping the gateway's loopback before continuing!

Step 8. Connect via SSH to the router connected directly to the PC; **be sure to set "terminal monitor"** from Privileged Exec mode so that you see log messages!

Step 9. Enable OSPF on every router. For now, don't bother setting the Router ID. **VERIFY** OSPF is running on every router. **Question**: What's the best command for verifying OSPF is running?

Question: Does the process ID need to be (a) identical on every router; (b) unique on every router; or (c) there's no interrelationship between different routers for the ID?

Task 2: Configure all interfaces to participate in OSPF

For full, end-to-end reachability throughout the network, _all_ interfaces will need to be activated in OSPF. We'll use a mixture of Network statements and interface configuration to do this.

Step 1. On each router, configure a **Network statement** to include the Loopback 1 in OSPF. Did you notice the subnet mask used for the IP configuration? If yes, then there's only one, single _wildcard mask_ that's an exact match for the IP subnet mask!

Step 2. On each router, do the configuration on **each physical interface** to participate in OSPF. For R4, don't forget to include the interface VLAN 1.
Hint: you can do this _very_ quickly if you make good use of the Up arrow key!

Step 3. After a few moments, you should have _full, end-to-end reachability_. From each PC, ping the other PC and all router loopbacks. **Record** the results and the TTL in particular, and be sure you understand the reason for each value. Troubleshoot if there's anything not reachable: start with the IP configuration, then proceed with the OSPF configuration.

Step 4. Show the routing table on each router. **Record** the output for later reference.

CHECK POINT #1: Have 4 windows open, showing the _hostnames_ and routing tables.

Task 3: Configure and verify Passive Interfaces

There's no need to send or listen for OSPF messages on (a) loopbacks, or (b) interfaces connecting to end-hosts. Set all interfaces of these two types to *passive mode* for OSPF.

Step 1.   For each router, identify the interface name(s) that should be made passive. On R1, R2: just the loopbacks; on R3, R4: the loopbacks and the interface connecting to the PC.

Step 2.   Within the **router ospf {process-ID}** context, configure the required interface(s) as passive. **Question**: Can you list multiple interfaces on a *single passive command*, or are *multiple passive commands* required?

Step 3.   Using the command **show ip proto**, verify that all interfaces are correctly set. **Record** the output of the command. (You'll want to be able to distinguish which interfaces were activated with Network statements, and which were done with Interface config!)

Task 4: Manually Configure and Verify the Router ID

Using any one of several different commands, you can verify the Router ID for each router. The value is almost certainly *not* the value shown on the *topology diagram* (if you accurately followed the instructions above). In this Task, you'll work through the requirements to change it.

Step 1.   Using an appropriate command, show and ***record*** the Router ID for each router. **Question**: Why did the router get assigned it's particular Router ID?

Step 2.   Check the course material and find the command to set the Router ID. Make sure you put it in your lab notebook, neatly organized with all the other OSPF commands you've learned so far.

Step 3.   For each router, set the Router ID as indicated on topology diagram. **Record** the message that's displayed on the terminal window.

Step 4.   Execute the additional command required to make the new Router ID take effect. **Question**: At what context (config level) must this command be executed?

Step 5.   Repeat step 1 to verify that the change has taken effect. If not, troubleshoot.

Task 5: Influencing and Verifying the DR election

The connection between R2, R3, R4 is a multi-access segment. It allows us to view a DR/BDR election. We can alter the results as we wish, but in a lab environment it requires some carefully timed steps to see the new results immediately. This is the *only* tweak to OSPF that requires careful timing; all other modifications in later tasks can be done at a leisurely pace.

Step 1.   First, practice the timed steps that are required:
– shut down all interfaces on that one segment: do a "shut" on G0/0 on R2, R3, R4
– wait at least 30 secs for the Dead timers to expire
– within a ~20 sec time span, do a no-shut on G0/0 on R2, R3, R4

Step 2.   If necessary, repeat Step 1 (use the Up arrow key) until you are able to re-enable all interfaces within 20 secs of each other.

Step 3.   Wait until you see the message that OSPF has gone from LOADING to FULL on all routers.

Step 4. Show OSPF neighbours on R2, R3, R4 to **verify** the role of each router. **Record** the results. Be sure you can correlate the output from all three routers. **Question**: Can you explain why each router obtained the role it did? **Question**: What abbreviations does Cisco use in the output to indicate each of the three possible roles?

Step 5. Use the command **ip ospf priority {value}** to configure the following priorities on G0/0 on each router: R2 = 20, R3=30, R4=1 (default priority)

Step 6. Repeat steps 1-3 to cause a force a new election where *all routers* are participating (all G0/0 interfaces **must be** re-enabled within the same short time span).

Step 7. Repeat step 4, including **recording** the results and answering the questions again. **Question**: What "show..." command(s) can be used to verify the (new) priorities?

CHECK POINT #2: Have three windows open, 1 per router, showing the *role* of the routers.

Task 6: Changing and Verifying the Default Reference Bandwidth

In the original days of OSPF, links speeds of dozens of Kbps were common. Today, optical fiber can reach or exceed 100 Gbps using techniques such as DWDM. To ensure that the *cost* of all links are properly calculated relative to one another, it's necessary to modify the reference bandwidth.

Step 1. Confirm and **record** the default reference value (i.e. original OSPF value) using the command **show ip ospf**. Ensure it's the same across all routers.

Step 2. Use the command **auto-cost reference-bandwidth {value}** to set a value of 100Gbps. **Question**: In what CLI context must this command be issued? **Question**: What numeric value is required? Note the *unit* that Cisco uses for this.

Step 3. Watch for and **record** the log message(s) that appear after issuing the command.

Step 4. To ensure neighbours are fully updated, either wait 1 minute or two or use the same command as when you changed the Router ID

Step 5. Ensure you make the same change across all routers. Repeat step 1 and **record**.

Task 7: Modify and Verifying Interface Cost

Imagine two links operating at the same speed, one being optical fiber and the other being a satellite link. The optical link is almost always preferable for multiple reasons: less expensive ($$$) per GB, not easily subject to eavesdropping, not affected by atmospheric conditions, and much lower latency. We want to translate our human knowledge of these factors into something OSPF can deal with. We do so by changing the *cost* parameter of an interface.

Step 1. First, confirm the path taken between PC1 and PC2 (in both directions!) Repeat for reaching each loopback from PC1 and PC2. Use the tools ping and traceroute and **record** your results. (In NetLab, only the traceroute command indicates the number of hops.)

Step 2. Determine and **record** the existing cost for the link from R4 to S1 (i.e. interface G0/0) using the command **show ip ospf interface g0/0**

Step 3. Calculate a new cost of 5x or even 10x higher and then assign that cost to the interface using the command: **ip ospf cost {value}** **Confirm** by repeating step 2.

Step 4. Wait a 10-15 secs and then repeat step 1 to determine the effect of the changed cost on the routes through the network. **Record** your results. **Question**: Is the new cost "one-way" or is it "bi-directional"?! You'll need to understand and use ping and traceroute carefully to get this correct. [Optional] **Can you explain the result**??

Task 8: Propagating and Verifying a Default Route

In order to be *the default gateway* for a larger network, a router must propagate a *default route*. This is easy to do in two steps in OSPF. This is covered in Module 2.5 of the course material.

Step 1.   Display and **record** the routing table on *each router*: **show ip route**
Step 2.   On **R2**, create a default static route; in this case we'll use Loopback 1 to simulate a connection to the global internet: **ip route 0.0.0.0 0.0.0.0 lo1**
Any log messages are interesting but not essential for this Lab exercise.
Step 3.   Repeat step 1 and find change(s), if any, in the routing tables.
Step 4.   On **R2**, configure OSPF to share the default route with it's neighbours –
in the router ospf {process-id} context type: **default-information originate**
Step 5.   Repeat step 3: display, **record**, and compare the routing tables.
Step 6.   Test it out using the PCs: use ping and traceroute on an address that exists in the network, and then on an address that does *not* exist in the network. **Record** your results.

CHECK POINT #3: Have 4 windows open, show the routing tables and identify the default route

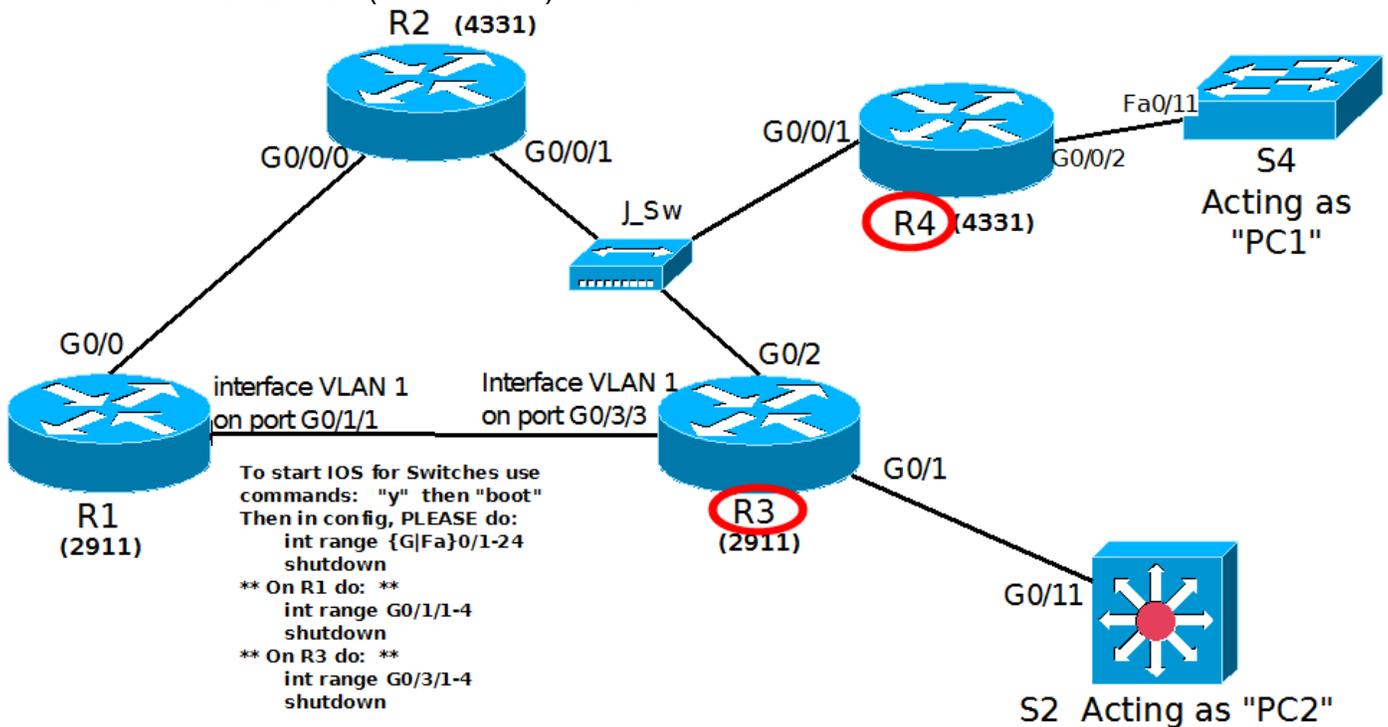## Task 9: Verify Basic and Advanced OSPF Configuration

Hopefully you've noticed that we've been verifying the configuration *and* the results of the OSPF configuration throughout the previous steps. In any case, pat yourself on the back because you're almost done this lab! :-)

Task 10: Clearing the equipment

Hopefully you know how to clear Cisco equipment, but just in case:
Router: **erase startup-config**   Switch: **erase startup-config** then **del vlan.dat**

Here's the simplified NetLab topology to repeat (or complete) this using NetLab; **note** interfaces and the **router numbers** (for the RHS) are different:

NetLab2 - Wiring/Topology