

Lab 4 Postlab: LDP and MPLS Shortcuts

Or. Using LSPs for iBGP full-mesh to transit an IGP

What you will do:

1. Configure iBGP and eBGP sessions, including route redistribution
2. Inspect network operation using debugging facilities to view ICMP packets
3. Configure BGP to use label switched shortcut paths: bridges across non-BGP speakers
4. Examine routing tables and BGP routes to determine reachability with/without shortcuts
5. Save configs & enter your answers for each task on BrightSpace

Things that you will need to know or learn:

1. Material in section 13.5 MPLS Shortcuts, **pp. 636-644 in NRS-II** textbook
2. CLI commands for configuring BGP (see previous labs)
3. CLI commands for debugging (and thus logging) (see next page)
4. CLI command for using shortcuts with BGP

What you need to submit and when:

1. Complete the “Lab 4 Post-lab” exercise on BrightSpace with answers from your work, **before** your lab section's assigned due date.

Required Equipment:

- Credentials for booking a MySRLab session
- An actual MySRLab booking
- A PC with internet access and terminal program; eg. Putty (provided by you)

Marks:

This post lab is worth 1/2 of the mark for Lab 4.
10% of your final mark is for labs done during the course of the semester.

References and Resources:

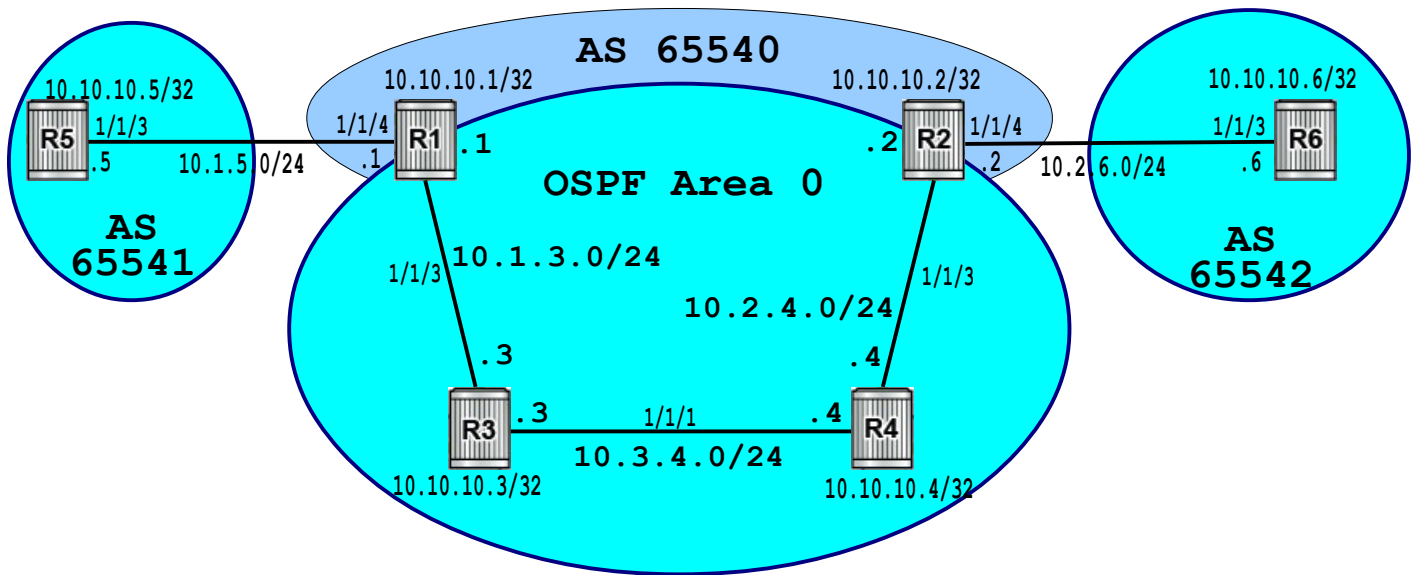
- Your textbook and all your notes on BGP from NET3008 – Advanced Routing
- Command references provided in previous labs
- BGP configuration commands provided in this lab (below)
- MySRLab: remote-access lab facility hosted at the Nokia Kanata campus
- 7750 SR OS Routing Protocols Guide Ver 12.0, see especially the sections on BGP configuration referenced in the previous post-lab

Tip: use google to search the terms:

7750 SR OS routing protocols guide 12.0 site:infoproducts.alcatel-lucent.com

Topology

This is the MySRLab network topology that is used for this lab. It is similar to the top half of the Edu topology from the weekly in-lab. **NOTE:** that **only some** links are configured and activated!




Addressing Table for MySRLab Practice Labs

ALL routers are IPv4 addresses in the range **192.168.X.Y** with **X.Y** as given in the table below.

	PL3	PL4	PL5	PL6	PL7	PL8
R1	188.4	188.36	188.68	188.100	188.132	188.164
R2	188.5	188.37	188.69	188.101	188.133	188.165
R3	188.6	188.38	188.70	188.102	188.134	188.166
R4	188.7	188.39	188.71	188.103	188.135	188.167
R5	188.8	188.40	188.72	188.104	188.136	188.168
R6	188.9	188.41	188.73	188.105	188.137	188.169
PC1	188.10	188.42	188.74	188.106	188.138	188.170

The Ixia tester PC is accessed via the Remote Desktop Protocol (RDP), with clients available for Windows (built-in as "mstsc.exe"), Linux (rdesktop), and OSX (<http://cord.sourceforge.net/>)

Practice Lab	User ID	Password
PL3	See pages 4-5 of the <i>MySRLab Technical Primer</i> and page 8 of the <i>MySRLab Getting Started Guide</i> (available on BrightSpace) for connectivity & login details.	
PL4		
PL5		
PL6		
PL7		
PL8		



Command Reference

```
configure • router • bgp • igp-shortcut mpls # (SR-OS ver 12 and earlier)
configure • router • bgp • next-hop-resolution • shortcut-tunnel • family ipv4 • resolution any

configure • log • log-id {log#} # Task 2: Apply on target router (eg. R1)
configure • log • log-id {log#} • from debug-trace
configure • log • log-id {log#} • to session
configure • log • log-id {log#} • exit

debug • router • ip • icmp # Task 2: Apply on target router (eg. R1)
no debug

ping {IP-addr} • count {#} # Task 2: Limit pings to generate manageable output

show • router • tunnel-table

environment • no more # Handy for displaying, copying, and saving your final config
```

Task 1A: Configure the base topology

Configure the base topology (above) so that R5-R1-R2-R6 have all 11 routes. (All of the configuration required is a repeat from previous labs and post-labs.)

- Do **not** include the links between R1-R5 and R2-R6 in OSPF.
- **Include** the system interfaces for R1-R2 in OSPF
- **Only** configure OSPF on R3-R4 and **not** BGP; those two are not BGP speakers!
- **Export** (or *share*) routes from direct and OSPF into BGP on R1 and R2
- **Remember** to use `next-hop-self` for the iBGP group on both R1-R2!
- **Export** (or *share*) routes from direct (ie. system interfaces) on R5 and R6 into BGP
- **Configure** LDP throughout Area 0

Verify your configuration:

- Test by pinging to ensure you have full IP reachability **within** each AS.
- Check the routing tables: R5-R1-R2-R6 should have **all 11 routes**; R3-R4 have **fewer**.
- Show the LDP active bindings and the tunnel-tables to ensure full reachability in Area 0
- Grab & save a copy of the full routing table as it appears on an OSPF router.

Questions to answer:

- How many routes are in the routing tables of R3, R4 ? (Hint: <11 ... and >6)
- On each of R1-R4, how many tunnels exist? How many active bindings?
- On R2, the tunnel to which router is "missing"? Why is this ok?

Task 1B: Experiment with the Router & LDP

Have a look at the routers and LDP and notice some of the interesting behaviors.

Step 1. On R1, start configuring an interface called "notWorking": first, give it a /32 address and then secondly, try to assign port 1/1/1. Copy the exact message you get. Make sure you understand it! Get rid of this experimental interface: shut it down and delete it.

Step 2. In the LDP interface-parameters, attempt to add the system interface. Copy the exact message you get. What's does this tell you about how Nokia implemented LDP?

Step 3. Grab & save a copy of the active LDP bindings table on each and every router in the

OSPF area. Find the range of values given out by LDP. Notice the highest and lowest label values.

Step 4. On each router, examine the pair of egress labels (push, swap) for each FEC.

Which labels, if any, match? Which ones, if any, don't match? Is this always the case?

Step 5. From within the LDP configuration context, issue the `info` command to see all the configuration. Grab & save a copy then examine the output carefully. In addition to the default FEC type (which you've just configured), what other FEC types can LDP generate labels for?

Questions to answer:

- If you associate a port with an interface having a /32 address, what message do you get?
- What's the exact message you if you add the system i/f into LDP?
- What is the *highest* (... and first!) label **given out** by LDP (on a Nokia router)?
In this exercise, what is the lowest (... and last!) label **given out**?
Is this range of values the exact same for every router? Why or why not?
- On any router, for a given non-local FEC, do the push & swap labels *always* match?
- For what additional FEC types can LDP generate labels?

Task 2: Inspect using log and debug facilities

This lab intentionally creates a (not unheard of) situation where internal routers are non-BGP speakers and thus lack routing information distributed in BGP (only). This task is a very brief introduction to log and debug facilities for the SR OS. The objective is to view ICMP packets associated with ping/traceroute, something that is *not* always easy to do on other vendors' equipment! If you are interested, you can get more info from the SR OS command reference guides.

- Step 1. Go back and check the Command Reference section above. On R1, configure a log (eg. #10) and then enable debugging for ICMP packets.
- Step 2. From R6, issue the command for a single ping: `ping 10.10.10.1 count 1`
- Step 3. Examine the output on R1 carefully. You should see **three** distinct debug entries. Each entry starts with an entry number, followed by a time stamp.
- Step 4. For the purposes of this lab, be sure you can identify three key fields:
 - the ingress / egress interface (... for the ICMP packets that we are debugging)
 - the source / destination ip address
 - the details of the packet (ie. ICMP message type)
- Step 5. **Capture** the debugging output for answering questions below.
- Step 6. Turn off debugging using `no debug` from the root context
Capture the output for answering questions below.

Questions to answer:

- Did the ping from R6 successfully reach R1? Can you prove it absolutely?
- Was a ping response sent out by R1? Can you prove it absolutely?
- What was the response of R1's neighbor in relaying pings? Why?
- What is the exact message you get when you turn off all debugging?
- Do you genuinely understand what is happening and why? If not, ask someone!!

Task 3: Create full-mesh BGP via MPLS shortcuts

The problem with the configured topology from Task 1 is that the iBGP peers are not fully meshed. Although there are a variety of solutions (confederations, route reflectors, etc), using label switched paths to create virtual links is extremely easy and effective.

- Step 1. Configure **both** iBGP peers R1-R2 to enable BGP to use a LSP. **Read** p. 636-644 in your NRS-II textbook to understand *why* we use "mpls" (or "any" in SR OS ver 13) in:
~~configure • router • bgp • igp-shortcut mpls~~ In SR OS 13, the command is now:
configure • router • bgp • next-hop-resolution • shortcut-tunnel • family ipv4 • resolution any
- Step 2. Examine the routing tables on R1-R4. Find the changes and determine whether the changes occur on all routers. **Capture** the output for answering questions below.
- Step 3. Repeat the pings from R6 to R1. Then run a full set of pings from every router to every other to determine what IP reachability exists. **Capture** the ping output.
- Step 4. Traceroute from R6 to R5. Are all hops visible? Why/why not? **Capture** the output.
- Step 5. Carefully examine R3-R4 to determine if there are any other changes.

Questions to answer:

- What change appears in the routing table(s)? On which routers does it occur?
- What other changes, if any, appear in R3-R4 routing and/or labels?
- Is there now full IP reachability everywhere? If not, what fails?

Task 4: Save all your configs

Really, after task 1, wasn't that easy?? As always, **save your final configs** so that you can refer to them later for practice and study.

save: admin save ftp://{mySRLabAlphaLogin}:{assignedPswd}@mysrlab/STD/Rx/L4.cfg
verify: file dir ftp://{mySRLabAlphaLogin}:{assignedPswd}@mysrlab/Rx/