



KTHNOC

OSPF-lab

Juniper version

Group Nr	
Name1	
Name2	
Name3	
Name4	
Date	
Grade	
Instructor's Signature	

Table of Contents

1 Goals.....	3
2 Lab report.....	3
3 Preparations.....	3
3.1 Preparations for area part.....	4
4 OSPF in pairs.....	4
4.1 Turn on OSPF.....	5
4.2 Monitor OSPF.....	6
5 Redistributing static routes.....	6
6 Redundant paths – square.....	7
7 Load balancing.....	8
8 Shared network – cross.....	8
9 Backbone.....	9
10 Areas.....	9
11 Areas in the world.....	10
12 Stub areas.....	10
12.1 Totally stub areas.....	10
13 NSSA.....	10
14 Cleanup.....	11
15 References.....	11

1 Goals

The goal with this lab is to introduce you to the link-state OSPF routing protocol and to understand how to configure OSPF for a variety of networks. The lab is relatively long and will take two days.

As previous labs, the scenario is an ISP where the workstations act as customers together with four virtual customer networks. You will start with single-area configurations, and will later proceed with multi-area networks. You will also connect all other groups into one OSPF network.

2 Preparations

Before you begin this lab, please consult documents [1] and [2]. Ensure that you have the network map of topology 1 [3]. You should also have read the lecture notes and RFC pages about OSPF[4]. Extra material from vendors published on the web are also useful.

Before you begin the lab, answer the following questions.

Which OSPF message types are there?

Describe LSA types 1 to 5.

What is an advertising router with respect to LSAs?

What is the role of a Designated Router?

What is the role of a Backup Designated Router?

How do you do the following in JunOS?

1. Check OSPF neighbours
2. Check the OSPF database
3. Check the OSPF (internal) routing table (NOTE: not the same as the global RIB).
4. Check which OSPF routes are installed in the (global) routing table

2.1 Preparations for area part

These questions need to be answered before you start the area (Section 9) part of the lab. Normally, you do not start with areas until the second lab.

What is the benefit of a regular area?


Describe the differences between a (regular) area, a stub area, an NSSA and a totally stub area? Explain the differences in terms of how LSA:s are filtered or mapped.

How do you do the following in JunOS?

1. Create an area
2. Create a stub area with a default route
3. Create a totally stub area
4. Create an NSSA

3 OSPF in pairs

First start working in pairs, RTX1 with RTX4 and RTX2 with RTX3.



Connect and bring up the network between the workstation and the router and configure the workstation with an IP address and static routes according to the topology map.

Do not configure any more router interfaces than required for a pair. You can re-use an old configuration from earlier labs.

Ensure that you can ping the router from the workstation, and between routers.

Investigate the state of the router. In particular note all routes and be sure that you have understood which routes exists in the routing table.

3.1 Turn on OSPF

Start with setting the router-id to 10.X.Y.1. Assign this address to the loopback interface (By default, however, OSPF will use the first lo0 address as router-id).

Turn on OSPF on the interfaces connecting you to your neighbour using area 0, the backbone. Also turn on OSPF on the interface to the work-station (fe-0/0/0). But since we do not expect an OSPF router to appear on the customer network – declare it as passive. That is, do not run the OSPF protocol, but announce it as an internal OSPF network.

Try to ping between your workstations. Also try to ping all the routers addresses, including the loopback addresses.

Check OSPF neighbours and interfaces to verify that OSPF runs correctly.

Check the OSPF database using the detailed (or extensive) option. Explain the entries and fields in the entries in the database. In particular, which different types of LSA:s exist? Which links? Who is advertising router? Explain the sequence and age values. What does the asterisk in front of some LSAs indicate?

Examine the OSPF internal routing table. For example, what does “path type” mean?

Check the routing table. Do you see any OSPF routes? Why/ why not?

Try disconnecting the cables between a work-station and a router and see how this influences the database and routing table. What happens (note that this will disconnect any ssh connections from that workstation)?

Try disconnecting the cable (or configuring the interface down) between the two routers. Study how the neighbour adjacency is reformed after the break. How long does it take before it is re-established?

Also examine which router is DR, and which is BDR. Does this change when you disconnect?

Milestone 1: Show a working pair in OSPF.

Signature: _____

3.2 Monitor OSPF

For debugging purposes, it might be useful to monitor the OSPF messages. Therefore, turn on traceoptions for hello, error and lsa-updates, and either monitor the results, or check the log now and then. If you monitor the log, it is recommended to do this in a separate ssh connection.

4 More routes

To create a more realistic scenario, you shall now add more routes to your network. These routes represent more customer networks, Consult the topology map[3] and create two of the 10.X.Y.0/24 networks as secondary subnetworks on the host network fe-0/0/0. Assign the .2 address of each of those networks to the interface (so that they do not collide with the router-id).

Then create the remaining two routes as static routes and redistribute them into OSPF using a policy. Set the .1 address of these two networks on the loopback (so that you can ping them later).

How do these new routes appear in the OSPF database?

Ensure that these networks appear also in the neighbour.


After this has been done, you should be able to ping the four .2 addresses from the other routers.

Can you also ping them from the workstations? If not, try adding a new static route to the 10.0.0.0/8 network on the workstation.

Examine the ospf neighbours, database, routing table, etc. How have they changed?

Try also restarting the routing processes and monitor with the show neighbour command the re-establishment. Which states does it go through and how long does it take to form adjacency?

5 Redundant paths – square



Now connect all four routers in a quad – a square network as shown in the Figure. Do not connect the shared network 192.168.15.0/24.

Examine the OSPF database and routing table and verify that you can ping from all routers and workstations to all other routers and

workstations. In particular, ensure that you can reach the diagonal neighbour and workstation.

Now ping and traceroute “diagonally” (eg, RTX1 to RTX3) – both to the other workstation and to the local addresses on the 10/8 network.

Examine if all packets follow the same path or if they choose different paths. In particular, ping the four different loopback addresses on the diagonal router with the record-route option. Explain the result.

Now examine how long OSPF takes to make failover by pinging to a neighbouring workstation (eg wsX1 to wsX2) and pulling the cable between the closest routers (eg RTX1 and RTX2). How long does it take before the ping resumes. Which path does the traffic now take?

Study how the LSAs are affected by the change.

In particular, you should be able to examine all entries in the link-state database (`show ospf database`), and explain all OSPF routes (`show route protocol ospf`).

Milestone 2: Show a working square in OSPF.

Signature: _____

6 Load balancing


Use the same square configuration as in the previous part of the lab. Try now to load balance traffic to the same address on a diagonal workstation or router.

In order to achieve load-balancing you need to create a “per-packet” policy and export this policy into the *forwarding-table*. This will cause layer-3 load-balancing which causes address-pairs (and pings) to take separate paths. See [3] for details.

Milestone 3: Load balancing.

Signature: _____

7 Shared network – cross



Connect fe-2/0/0 to the hub at the bottom of your rack to study OSPF over a transit network with more than two routers. Use the addresses assigned in the topology map.

For this exercise, you shall connect one of the workstations to the shared network and start Ethereal on that network. You will need to temporarily change network interface address on the workstation's interface.

When you configure OSPF on the routers' interfaces study (and list in report) the output on Ethereal. Describe the messages being transferred. Describe which states the routers go through before they reach a stable state.

Disconnect the cable connecting the DR to the 192.168.15.0 net.

- Which machine becomes the next DR?
- How long does it take for the new DR to be assigned?

Reconnect the old DR to the net.

- Does the old DR become DR again? If yes, how long does it take? Otherwise, why does it not become DR again?

Illustrate with Ethereal output.

Milestone 4: Transit link.

Signature: _____

8 Backbone

In this part, you will continue the previous section by connecting your square to the ``world'' by connecting the switches into one shared network¹. The goal is that everyone should be able to ping every other workstation and loopback address.

A helpful attitude helps in this part of the lab. :-)

¹ Depends on the progress on the other groups.

Study the database, routing tables, and the DR election on the shared network: What happens when two shared medium networks (with two DRs) are connected to each other?

Use Ethereal to see what happens.

Milestone 5: Backbone.

Signature: _____

9 Areas

Disconnect the hub and configure a square again (see Section 6). Keep the secondary, and static routes and loopback addresses. It is recommended that you restart your routers, or at least restart the routing process, so that you do not inherit old OSPF state.

Disconnect the cable between RTX1 and RTX2 so that your topology looks like a “U” (lying down).

Define two areas per group: RTX1 and RTX4 shares one area and RTX2 and RTX3 another. Call the areas 0.0.X.1 and 0.0.X.2, respectively, where X is 1 for group A, 2 for group B, etc. Let RTX4 and RTX3 communicate over the backbone area.

Ensure now that your network works by pinging between all workstations and routers including the loopback interfaces.

Examine your OSPF database to see if you can detect the new LSAs. Which new LSAs do you see?

In particular, study the OSPF link-state database on RTX3 and RTX4. Explain the LSAs in the different area databases. Why do some seem to appear in both areas? What are their differences?

Milestone 6: Areas in OSPF.

Signature: _____

10 Areas in the world¹

Connect RTX4 to the world, that is, the Ethernet switch, and declare the interface (fe-2/0/0) as area 0.

¹ Depends on the progress of other groups

The goal is to try to create one large OSPF network using areas. Try to ping all other groups.

You can here also use Ethereal to see what is going on.

11 Stub areas

Continue with the configuration from the section 9. But declare your sub-areas as *stub* areas, and inject a default route into the areas.

Ensure now that your network works by pinging between all workstations and routers. Do this especially for the networks that are treated as external by OSPF.

Examine the OSPF database and routes. What happens to the static routes inserted by RTX1 and RTX2? Why? Notice the default summary route in the stub-areas.

As in the previous section, study the states of the routers. What are the differences between the regular areas?

11.1 Totally stub areas

Make your two stub areas totally stub, that is, filter LSA-3's as well. What is the difference?

Milestone 7: Stub areas.

Signature: _____

12 NSSA

In this part, you shall create Not-so-stubby areas. That is, you should be able to distribute the static routes from RTX1 and RTX2 again.

Redefine your sub-areas as nssa areas.

Again, go through the state of the routers, list the output in the report, and describe what happens.

What new LSAs appear inside the sub-areas? How do the external routes appear outside of your area. What LSA are they?

Milestone 8: NSSAs.

Signature: _____

13 Cleanup

To reset the machines to their initial state, follow the following steps for all routers.

1. Reset the configuration by loading (using override) the `~root/baseconf` configuration.
2. Halt the routers using `request system halt`. After the system has stopped, you can power-down the routers.
3. Disconnect the cables and ensure the equipment is in the same state (or better) as when you started.

Also do not forget to log out from your workstations.

14 References

- [1] KTHNOC Router lab Introduction - Juniper version
- [2] KTHNOC Router lab Reference - Juniper version
- [3] KTHNOC Router lab Netmap - Topology 1
- [4] IETF RFC 2328, OSPF version 2