

Praktikum RouterLab SS 2009 Work Sheet 4: OSPF

The purpose of this work sheet is to introduce the Open Shortest-Path First (OSPF) as an example for a dynamic routing protocol. The goal is to understand the advantages and disadvantages of the link-state.

Important Notes:

The basic network topology for this worksheet is the same you developed in the previous worksheets, thus on the VLANs configuration of:

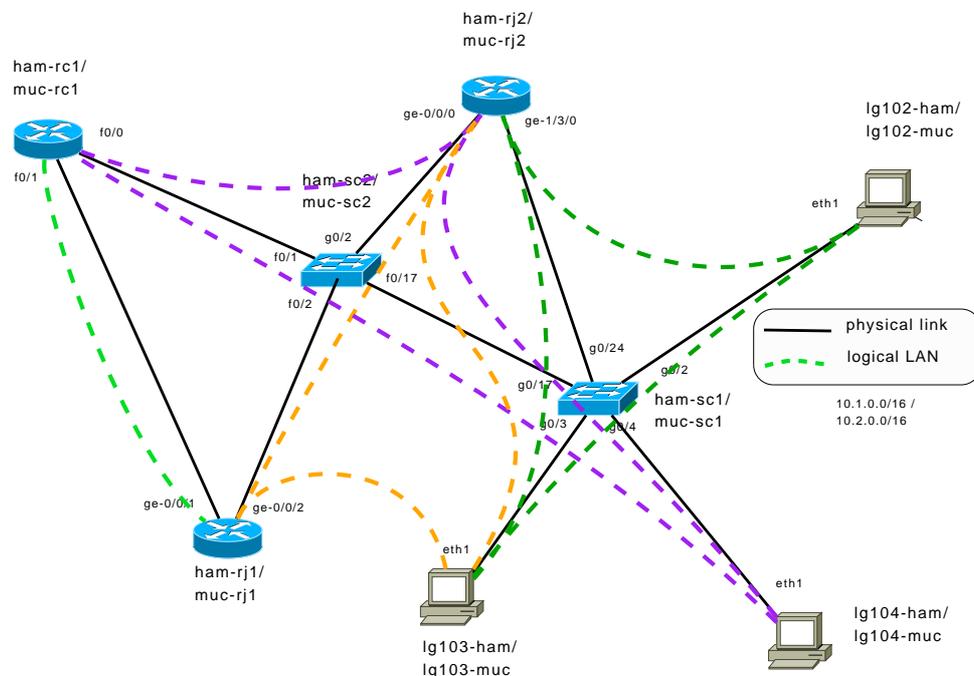


Figure 1: Topology

Table 1: Assignment of devices to groups

Group	Ham-Cloud	Muc-Cloud
Router	ham-rc1, ham-rj1, ham-rj2	muc-rc1, muc-rj1, muc-rj2
Switches	ham-sc1, ham-sc2	muc-sc1, muc-sc2
Loadgens	loadgen10[2/3/4]-ham	loadgen10[2/3/4]-muc
IP range	10.1.0.0/16	10.2.0.0/16
VLAN IDs	$x00 - x99$ where x is your team number	

Question 1: (10 Points) *OSPF: Calculating the Routing Table*

OSPF uses the information in the Link State Database to calculate the routing table using the Dijkstra algorithm. If you are not familiar with this algorithm: there is plenty of information in the web (even with cool animations).

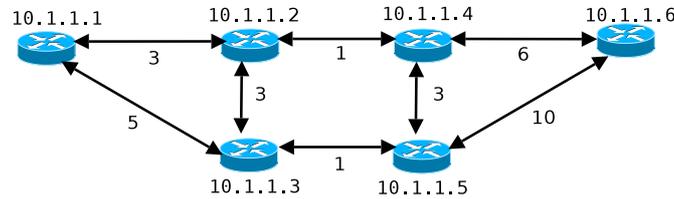


Figure 2: Dijkstra - Simple example

Figure 2 shows a simple topology where each router is identified by one IP address (for simplicity we don't consider interfaces). Compute the shortest-path tree for 10.1.1.4 toward every possible destination using Dijkstra.

- For this purpose, show each iteration of the algorithm with the same notation used in the tutorial.
- Once the algorithm is finished write down the routing table, again in the same format like in the tutorial¹.

Question 2: (30 Points) *OSPF: Basic Configuration*

Take the configuration of Question 3 of Work Sheet 3 (including configured VLANs and IP addresses), here we assume that you already created the same VLANs. Now, instead of using RIP to make these VLANs communicate we configure OSPF. The following documents explain how to configure OSPF for Cisco and Juniper (see website):

- CiscoRIP
- JunosRIP

Configure OSPF on all routers respecting the following:

- Create one single area with ID 0.0.0.0.²
- Do not assign OSPF costs yet.

Note that on *loadgen102-xxx*, *loadgen103-xxx*, and *loadgen104-xxx* you have to add static routes like in the worksheet 1 to reach all configured subnetworks.

After convergence, *xxx-rc1*, *xxx-rj1*, and *xxx-rj2* in the network should have a route towards each sub-network you configured.

- For each loadgen of your cloud: traceroute every other configured interface (the ones *you* have given IP addresses) of the routers and loadgens of your cloud to verify they're all reachable, and to see which path the packets are taking.
- Which routes does *xxx-rc1* have in its routing table to reach the subnetwork formed by *loadgen103-xxx/xxx-rj1/xxx-rj2*?
- What total costs do these routes have? Read Section 2.4 and 16.8 of RFC 2328 and explain (4-6 sentences).

¹This does not mean to provide slides, the answers have to be provided in text format.

²Areas can be specified as either a decimal value or in IPv4 address representation format.

Question 3: (30 Points) *OSPF: Neighbor Discovery, Link-State Database*

In this question, we study how OSPF routers get to know their neighbors and how they learn about the topology. The topological information is kept in the *link-state database* in the form of so-called *Link State Advertisements* (LSAs).

Read Section 4.3 “Routing Protocol Packets” and Section 4.4 “Basic Implementation Requirements” of RFC 2328 and answer the following questions:

- a) Does OSPF rely on TCP or UDP?
- b) What types of packets does OSPF define and when are they used?
- c) What types of LSAs does OSPF define?

Based on the OSPF configuration in the previous question, study the Link State Database of *xxx-rc1* and *xxx-rj2*:

- d) Which LSAs (ID and type) are currently in the database? Explain (you have also to provide a dump of both databases).

Whenever a OSPF neighborhood between two routers is established, the two routers have to synchronize their LSAs. We will now study which packets are exchanged in such a situation. For this purpose, temporarily disable the interface *f0/1* on *xxx-rc1*. Wait a few seconds until *xxx-rc1* and *xxx-rj1* have noticed that the link is “broken”.

Enable detailed monitoring for OSPF traffic on *xxx-rj1*, using the following command:

```
monitor traffic interface ge-0/0/1 detail
```

and enable debugging mode for OSPF adjacencies on *xxx-rc1*. Finally, activate the interface *f0/1* again and observe both the debug output and the traffic monitor. Answer the following questions (dump of the debugging and monitoring sessions have to be submitted as well):

- e) Which OSPF packet types do you see in which sequence?
- f) Which LSA headers (LSA IDs) does *xxx-rj1* send to *xxx-rc1* within *database description packets*?
- g) Are complete LSAs exchanged between *xxx-rc1* and *xxx-rj1*? If yes, which LSAs (LSA IDs)?
- h) Read Section 10.3 “The Neighbor state machine” of RFC 2328. What is the evolution of the adjacency state that you can observe on the output of the debugging mode of *xxx-rc1* Explain.

When LSA information changes later on, the routing updates are propagated by a procedure called *reliable flooding*. Read the first 2 paragraphs of Section 13 in RFC 2328.

- i) How is flooding made reliable?

Question 4: (30 Points) *OSPF: Costs and Routing Hierarchy*

OSPF has a much more flexible management of link costs that do not create problem like the count-to-infinity in RIP. On *xxx-rc1*, configure OSPF for the link *xxx-rc1 – xxx-rj2* with a cost of 100.

- a) From *xxx-rc1*, what is the route to *loadgen102-xxx*? (dump the traceroute)
- b) From *loadgen102-xxx*, what is the route to *xxx-rc1*? (dump the traceroute to both interfaces of the router)
- c) Explain what you observe (4-6 sentences).

OSPF supports a hierarchical routing scheme through the use of OSPF areas. Each OSPF area is identified by a 32-bit Area ID. Read Section 3, 3.1, 3.2 and 3.3 of RFC 2328. Answer the following questions:

- d) Which ID is used for the backbone area?
- e) Is it possible that the path, a packet travels, covers 4 areas? Explain, why or why not.

- f) How can routers be classified in terms of their OSPF role?
- g) Is it possible to set the following configuration in our topology?
- Area 0.0.0.1 is formed by the network xxx-rc1/xxx-rj1.
 - Area 0.0.0.2 is formed by the network xxx-rc1/xxx-rj2/loadgen104-xxx.
 - All other networks are in Area 0.0.0.0.

If yes explain in few sentences how should the configuration look like. If no explain why.

Change your OSPF configuration such that the following three different areas are created (Note: maintain the cost of 100 from xxx-rc1 to xxx-rj2):

- Area 0.0.0.0 is formed by the network xxx-rc1/xxx-rj1 and xxx-rc1/xxx-rj2/loadgen104-xxx.
 - Area 0.0.0.1 is formed by the network xxx-rj1/xxx-rj2/loadgen103-xxx.
 - Area 0.0.0.2 is formed by the network xxx-rj2/loadgen103-xxx/loadgen102-xxx.
- h) Which route is selected by xxx-rc1 to loadgen102-xxx? Is it different from question 4.a?
- i) Why?
- j) What are the ABRs in this configuration?
- k) Study the routing table and Link State Database of xxx-rc1 (you have to provide a dump of both). Which LSA types can you see now (compared to question 3d) and why?

OSPF areas allow to reduce the routing state that needs to be kept in individual areas. The idea is that certain details of an area are not visible in other areas. However, there is another way to further reduce this state: Aggregation. *Area Border Routers* (ABR) can aggregate prefixes in order to reduce the number of prefixes announced. First configure area 0.0.0.2 as area 0.0.0.1 (so we have only two areas now). Then configure the ABRs such that only aggregated prefix ranges (i.e., 10.x.0.0/16 rather than 10.x.x.0/24) are announced from area 0.0.0.1 to area 0.0.0.0.³

- l) Which route is selected by xxx-rc1 to loadgen102-xxx? Is it different to question 4h? Explain the reason of what you observe (4-6 sentences).

For submission details please check the FAQ:

http://www.net.t-labs.tu-berlin.de/teaching/ss09/RL_labcourse/

Note also that topology maps have to be submitted as part of the answer for questions 2a, 4a, 4h, and 4l.

Due Date: Friday, May 29th, 2009, 08:00 am

³Hint: look for “area-range” in the documentation.