TECHNISCHE UNIVERSITÄT BERLIN
Fakultät IV – Elektrotechnik und Informatik
Fachgebiet Intelligente Netze und Management verteilter Systeme
Prof. Anja Feldmann, Ph.D.
A. Mehmood, W. Mühlbauer, H. Schiöberg, A. Wundsam

Praktikum RouterLab SS 2008
Work Sheet 3: OSPF, Quagga

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	loadgen102	loadgen104
IP range	`10.1.0.0/16`, `10.2.0.0/16`	`10.3.0.0/16`, `10.4.0.0/16`
VLAN IDs	Each team can use up to 9 VLAN IDs, where last 2 digits reflect team ID *Example*: Team 09 has VLANs 109, 209, . . . , 909	

Question 1: (22 Points) *Quagga, Basic Configuration*

The goal is to setup the topology of Figure 1 that provides a basis for the remaining questions.

(a) To make yourself familiar with *Quagga*, have a look at the website: `http://www.quagga.net/docs/`.
Focus on the parts that are relevant to answer the following questions:

1. What are the routing protocols, currently supported by Quagga!
2. What is the task of *Zebra*?
3. What is *vtysh* and how does it provide communication between daemons?

(b) Now check if Quagga is already running on your *loadgen*! The following steps are needed to start a Quagga instance:

1. If it does not already exist on your *loadgen*, install Quagga: `apt-get install quagga`
2. The Quagga configuration files are usually placed in `/etc/quagga`:
   - `daemons`: Make sure that *zebra* and *ospfd* are activated.
   - `ospfd.conf` and `zebra.conf`: Make sure that both files exist and are empty.
3. Restart Quagga: `/etc/init.d/quagga restart`
   Check in the output whether the *zebra* and *ospfd* daemon have been started.
4. The actual configuration is done via the VTY shell which is started using the command `vtysh`. Afterwards, if you see only "(END)" on your screen, type `q`. Next, we recommend to enter the command `terminal length 0` for ease of convenience.

You will find that the Quagga is very similar to Cisco IOS, although there are slight differences.

(c) Which IP address is assigned to the loopback interface *lo* on your *loadgen* (ifconfig)?
In Linux it is possible to assign multiple IP addresses to the same interface. Add the IP address/prefix `10.2.200.1/24` (Muc-Cloud: `10.4.200.1/24`) to the loopback interface `lo`! For this purpose use vtysh (do not use ifconfig).

(d) Enable IP connectivity between adjacent routers in a similar manner as in the previous work sheet! For this purpose, please have a look at Figure 1. You can reuse existing configurations of the previous work sheets if you wish. However, this may still require certain modifications. Your configuration must satisfy the following conditions:

1. The following routers must be able to ping each other: (i) Router rc1 and rj1 (direct physical link) and (ii) all pairs of routers that are interconnected with an arc ("VLAN").
2. Assign IP addresses in a way such that there is no (!) overlap in IP address range. The IP address ranges used for different arcs ("VLANs") must be disjoint.
3. Exclusively take IP addresses from the range `10.1.0.0/16` (Muc-Cloud: `10.3.0.0/16`) for the left side in Figure 1, i.e., the box for area `0.0.0.0`. All interfaces on the "left side" must be configured with IP address from this range. For the "right side", use IP addresses from the range `10.2.0.0/16` (Muc-Cloud: `10.4.0.0/16`).
4. Configure separate VLANs for the following "logical networks":
   - rj1 to rj2: tagged VLAN, over switch sc2
   - rc1 to rj2: tagged VLAN over switch sc2
   - rj2 to *loadgen*: tagged VLAN over switch sc1
   - rj1 to *loadgen*: tagged VLAN over switch sc1
5. As the physical link from sc1 to *loadgen* is used for two separate VLANs, you need to configure two sub-interfaces for the phyiscal interface at the *loadgen*. The following example illustrates how to do this[1]:

```
root@loadgenxxx:~# apt-get install vlan
root@loadgenxxx:~# ifconfig ethX up
root@loadgenxxx:~# vconfig add ethX <vlan-id>
root@loadgenxxx:~# ifconfig ethX.<vlan-id> <ip-address> netmask <netmask>
```

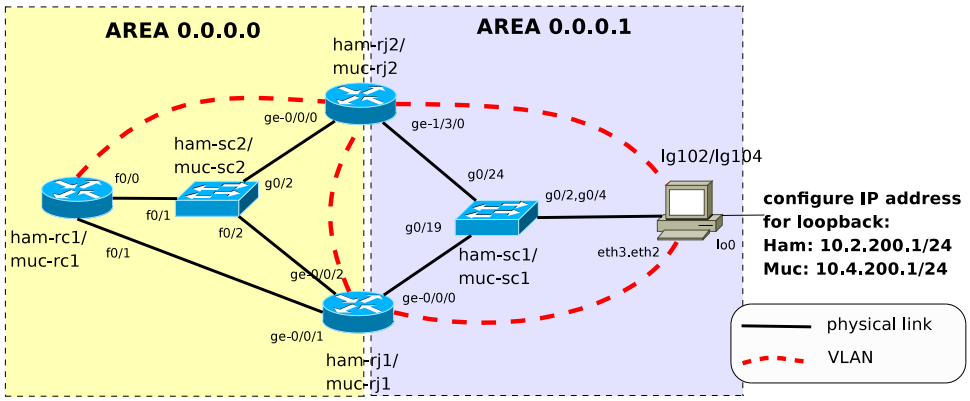

Figure 1: Topology – Basic Configuration.

Question 2: (22 Points) *OSPF: Basic Configuration*

(a) Have a look at the Cisco and Juniper OSPF manuals that are provided on the course website.

(b) Configure OSPF on all routers, including the *loadgen* (Quagga). All routers must have a route to the IP prefix that you assigned to the loopback interface of your *loadgen*. Check with show ip route (Cisco) and show route (Juniper) respectively.

Adhere to the following:

- All router and *loadgen* interfaces, for which you assigned IP addresses, must be OSPF-enabled. This also applies to the loopback interface of your *loadgen*.
- For now, all routers and their interfaces are in area `0.0.0.1`.[2]
- Do not assign OSPF costs yet.
- *Hint for Cisco*:

```
router ospf <random process id>
    network <interface-IPs> <netmask> area <area-id>
```

- *Hint for Juniper*:

[1] subinterfaces can be removed with vconfig rem, see man vconfig

[2] On Cisco areas can be specified as either a decimal value or as an IP address

```
protocols {
    ospf {
        area <area-id> {
            interface ge-x/x/x;
        }
    }
}
```

(c) Which routes does rc1 have in its routing table to reach `10.2.200.1` (Muc-Cloud: `10.4.200.1`)? What total costs do these routes have? Read Section 2.4 "Equal-cost Multipath" of RFC 2328 and explain!

Hint: To find out the total OSPF costs on a Cisco router, use the command show ip route. The following output would suggest a total cost of 1234:

```
O IA 1.1.1.0 [110/1234] via 10.1.100.2, 00:12:23, FastEthernet0/1.100
```

Question 3: (22 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).

(a) Read Section 4.3 "Routing Protocol Packets" and Section 4.4 "Basic Implementation Requirements" of RFC 2328 and answer the following questions:

1. Does OSPF rely on TCP or UDP?
2. What types of packets does OSPF define and when are they used?
3. What types of LSAs does OSPF define?

(b) Based on the OSPF configuration in the previous question, study the Link State Database of rc1 using the following command:

```
show ip ospf database
```

Which LSAs (ID and type) are currently in the database? Explain!

(c) 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 rc1 (shutdown). Wait a few seconds until rc1 and rj1 have noticed that the link is "broken".

Meanwhile, enable monitoring for OSPF traffic on rj1, using the following command:

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

Finally, activate the interface `f0/1` again and observe the debug output. Answer the following questions:

- Which OSPF packet types do you see in which sequence?
- Which LSA headers (LSA IDs) does rc1 send to rj1 within *database description packets*?
- Are complete LSAs exchanged between rc1 and rj1? If yes, which LSAs (LSA IDs)?

The output of this monitoring sessions needs to be submitted!

(d) 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. How is flooding made reliable?

Question 4: (12 Points) *OSPF: Calculating the Routing Table*

Based on information in the Link State Database each router calculates its routing table using the Dijkstra algorithm. If you are not familiar with this algorithm: There is plenty of information in the web!

Figure 4 shows a simple topology where each router is identified by one IP address (for simplicity we don't consider interfaces). Compute the routing table for `10.1.1.3` using Dijkstra. For this purpose, complete the following table:

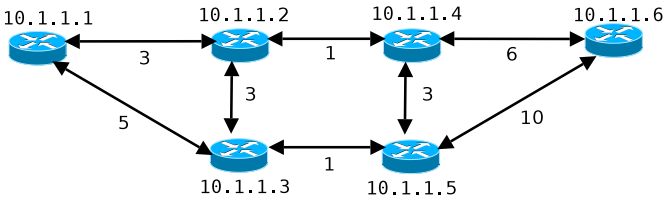

Figure 2: Dijkstra - Simple example

Iteration	Destination, added to shortest path tree	list of candidates `destination (cost, next hop)`
1	10.1.1.3	10.1.1.5 (1, 10.1.1.5) 10.1.1.2 (3, 10.1.1.2) 10.1.1.1 (5, 10.1.1.1)
2	10.1.1.5	10.1.1.2 (3, 10.1.1.2) 10.1.1.4 (4, 10.1.1.5) 10.1.1.1 (5, 10.1.1.1) 10.1.1.6 (11, 10.1.1.5)
...	...	...

Question 5: (22 Points) *OSPF: Costs and Routing Hierarchy*

OSPF supports a two-level hierarchical routing scheme through the use of OSPF areas. Each OSPF area is identified by a 32-bit Area ID and consists of a collection of network segments interconnected by the routers.

(a) To start with, configure OSPF costs for the following links:

- Traffic from rc1 to rj2: Cost 100
- Traffic from rj1 to *loadgen*: Cost 100

Which route is selected by rc1 to `10.2.200.1/24` (Muc-Cloud: `10.4.200.1/24`)?

Note that OSPF costs are associated with interfaces. If you assign OSPF costs to interface i of router A that is connected with interface j of router B, then these costs are applied for traffic from A to B (unidirectional). Use the following commands:

- Cisco: `(config-if)# ip ospf cost x`
- Juniper: `set protocols ospf area A interface ge-x/x/x metric X`

(b) Read Section 3, 3.1, 3.2 and 3.3 of RFC 2328. Answer the following questions:

1. Which ID is used for the backbone area?
2. Is is possible that the path, a packet travels, covers 4 areas? Explain, why or why not.
3. How can routers be classified in terms of their OSPF role?

(c) Change your OSPF configuration such that all interfaces of the "left part" are in OSPF area `0.0.0.0`, while the interfaces of the "right part" are in area `0.0.0.1`!

Which route is selected by rc1 to `10.2.200.1/24` (Muc-Cloud: `10.4.200.1/24`? What is the shortest path from rc1 to this prefix considering the OSPF costs that you assigned in Question 5a?

Compare this "shortest path" against the path that will be actually used for forwarding packets from rc1 to the *loadgen*. For this purpose, start from rc1 and walk along routers/hops by always continuing with the next hop, that you see in the current routing table. Do you see a difference or not? Explain!

Hint: OSPF prefers intra-area over inter-area routes!

(d) 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, this requires aggregation of routing information at *Area Border Routers* (ABR).

What are the ABRs in our example of Figure 1?

Configure both ABRs such that only aggregated prefix ranges (e.g., `10.2.0.0/16` rather than `10.2.1.0/24`) from the right into the left area. For this purpose, use the following commands on the two Juniper routers:

`set protocols ospf area <area-id> area-range <prefix>`

Study the routing table and Link State Database of rc1 (see Question 3b). Which LSA types can you see now and why?

For submission details please check the FAQ:
`http://www.net.t-labs.tu-berlin.de/teaching/ss08/RL_labcourse/`
Due Date: Friday, May 30th, 2008, 08:00 am