KTH CSC

MPLS/RSVP/BGP lab

Juniper version

Group Nr	
Name1	
Name2	
Name3	
Name4	
Name5	
Grade	
Instructor's Signature	

Table of Contents

1 Goals

This lab gives an overview of signalled MPLS tunnels (LSPs) using RSVP (Resource Reservation Protocol) and traffic-engineering. It also shows how BGP can use MPLS for its transit traffic only.

The scenario is that you operate a backbone network with transit traffic that you wish to send through signalled LSPs. You will also steer the traffic along certain paths using traffic engineering and fail-over techniques.

When you have finished the lab, you will understand how signalled LSPs work and how it interacts with the IGP. You will also understand how to setup simple traffic engineering scenarios, where you can steer traffic in other directions than the shortest path computation offered by the IGP.

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 will need four routers and two hosts for this lab.

You should also have read the lecture notes and RFC pages about MPLS. You should also be familiar with ISIS [4], OSPF [5], static MPLS [6], and BGP[7]. Extra material from vendors published on the web are also useful.

Before you begin the lab, answer the following questions.

1. What does LSP stand for?

____

2. An LSP goes from an ingress router, via some transit router, to a pen-ultimate, and finally an egress router. Which label operations does a router make in the following roles?:
ingress: ____
transit: ____
pen-ultimate: _______________________________________
egress: ____

3. How large is the MPLS shim header, and which fields does it contain?

____
____
____

4. What is label stacking?

____

5. What happens when a packet with label 0 arrives at a router?

____

6. What happens when you assign label 3 as a swap operation?

____

7. Explain why you may have to set “icmp-tunnelling”?

____

____

8. What does RSVP-TE stand for?

____

9. What is loose and strict source routing?

____

10. What is the difference between E-BGP and I-BGP?

____

11. How do you configure next-hop self in BGP?

____

Milestone 1: Preparation questions.
Signature: ____________________

3 Initial configuration

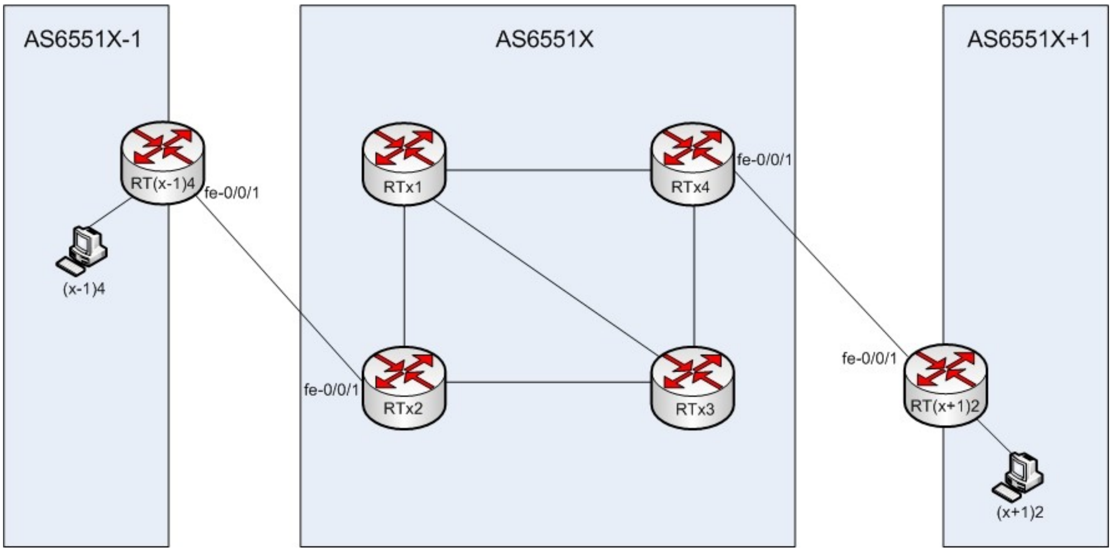

Configure two hosts, namely (x-1)4 and (x+1)2 which thus belong to other groups (see the Figure). For instance if you are group C your hosts will be b4 and d2. You will control hosts in the other groups in order to send transit traffic through your group. Configure eth0 and add static routes to the 10/8 and 192.168/16 networks as per usual.

Set up an IGP according to the netmap and the Figure by configuring all networks according to the topology map. Remember to add family iso on all interfaces if you are using IS-IS. Configure fe-0/0/0 and lo0 as passive IGP interfaces. Run the IGP as backbone. Configure loopback according to the netmap and assign it as router-id. Configure also fe-2/0/0 on RTX1 and RTX3 but choose a network address from your own address space (for example 192.168.X.52/30). Do *not* use the shared 192.168.15.0/24 network.

Check with ping and traceroute that intra-domain routing works and check the routing table.

Four routers will constitute an AS and every AS contains two border routers that connects to external peers. As can be seen in the Figure group B will peer with A and C, C with B and D and so on. The circle is closed with the peering between A and E. Use the AS numbers in the Figure.

Only two customer networks shall be set per group, 10.x.4.0/24 (at RTx2) and 10.x.12.0/24 (at RTx4). Do not bother adding the others as static routes. The non-border routers do not need any customer networks.

Configure I-BGP between the two border routers. Set next-hop self. Next-hop self is normally done with a simple policy which is exported to the I-BGP group only. I-BGP is picky about local-address, it must be set.

Do not configure BGP on the non-border routers in your AS.

Configure E-BGP between the peers and announce your customer network towards your external peer. The announcement is done with a policy and a route-filter. Use the fe-0/0/1 interface for the BGP peering. Your own address for this interface is already in the configuration, but you must ask your peers for their addresses. The address range is 192.71.24.34 (RTA1) – 192.71.24.53 (RTE4).

Now you should be able to receive prefixes for the other groups (depending on their progress of course), you can check this with the *show route receive-protocol bgp* command. Also, make sure that you are announcing anything, check this with the *show route advertising-protocol bgp*.

In summary, your configuration should now be as follows:

1. Internally, in your AS, all interface and loopback addresses should be reachable.
2. The IGP should not be enabled on the DMZ.
3. Your host networks shall be seen in your AS border-routers.
4. No BGP on non-border routers.

Even though the opposite host networks is visible on the external border routers, you should not be able to ping them. That is, you shall not be able to ping between (x-1)4 and (x+1)2. Why?

____

Milestone 2: Show an initial IGP/BGP configuration.
Signature: _____________________

4 RSVP-signalled LSPs

Enable RSVP and MPLS on all internal interfaces (not DMZ and lo0).

Setup LSPs between the border routers in both directions using RSVP. The egress of the LSPs should be the loopback address of the border router. Disable constrained path LSPs – since you want to control the paths RSVP takes.

Examine the established MPLS label operations at the ingress, transit and egress. Which labels/label operations are used at each hop? (examine the mpls.0 routing table using the detailed option)

____

____

____

____

____

Set explicit-null on all routers. What happens to the MPLS label operations on the pen-ultimate router?

____

Examine the RSVP and MPLS state of the machine using the show-commands. Specifically, examine interface, and lsp state, with the detailed option set.

You can set traceoptions so that you are able to trace the RSVP operation.

5 Using LSPs for transit

You will now examine how the LSPs you set up in the previous section can be used for transit traffic.

Examine the routing tables. Where do the LSP routes appear (the routes created when the LSP was created)?

____

When you traceroute to the LSP endpoints internally, does traceroute use the LSP or just IP using the IGP?

____

Try tracerouting the external networks instead. Which route does it use? For traceroute to work you must enable *icmp-tunneling* on the routers.

____

The inet.3 table is used by BGP only when selecting paths for transit traffic.

Milestone 3: Show a working BGP/MPLS/RSVP setup.
Signature: ____________________

6 Traffic engineering: paths

6.1 Loose source routing

The LSPs in the previous section used the IGP to setup a shortest path. Use loose source routing in both directions to alter the automatic choice. Create an mpls path, and then use that in the LSP as a primary path.

It is generally advised to use loopback addresses when specifying loose source routing. Why?

____

Disable an interface so that an LSP goes through all internal routers. Check with *show mpls lsp detail.* How long does this failover take? Use ping or traceroute to determine this.

____

6.2 Strict source routing

Steer your LSPs using strict source routing in both directions instead. Let the LSP visit all internal routers.

Here it is advised to use interface addresses instead. Why?

____

7 Traffic engineering: failover

Make two paths for each LSP. The primary uses strict source routing along one of the internal routers. The secondary uses the alternative path and uses loose source routing.

Are both paths active?

____

Measure how long it takes for the fail-over to take place: Setup a continuous ping and traceroute between your (x-1)4 and (x+1)2 hosts, for instance between b4 and d2 (if you are group C). You may also need to enable RSVP tracing.

Disable an interface on the primary path and observe what happens:

When can you observe the LSP fail-over (how long is traffic lost)? How does this compare to the fail-over for loose source routing? Which is more effective? Motivate.

____

When you re-enable the interface, how long does it take for the LSP to resume the primary path (use traceroute or the *show lsp* command)?

____

Milestone 4: Traffic engineering: paths and failover
Signature: ______________________

8 Traffic engineering: reserving bandwidth

Set a bandwidth limit for the LSP. You will test this by using a ping flood from the host with decently large packets (ping -f -s 1500 <address>). Determine the amount of traffic you send and pick an appropriate limit.

What happens when the limit is exceeded?
____

9 References

[1] KTH/CSC Router lab Introduction - Juniper version
[2] KTH/CSC Router lab Reference - Juniper version
[3] KTH/CSC Router lab Netmap - Topology 1
[4] KTH/CSC Router lab ISIS - Juniper version
[5] KTH/CSC Router lab OSPF I & II - Juniper version
[6] KTH/CSC Router lab MPLS static – Juniper version
[5] KTH/CSC Router lab BGP Intro - Juniper version

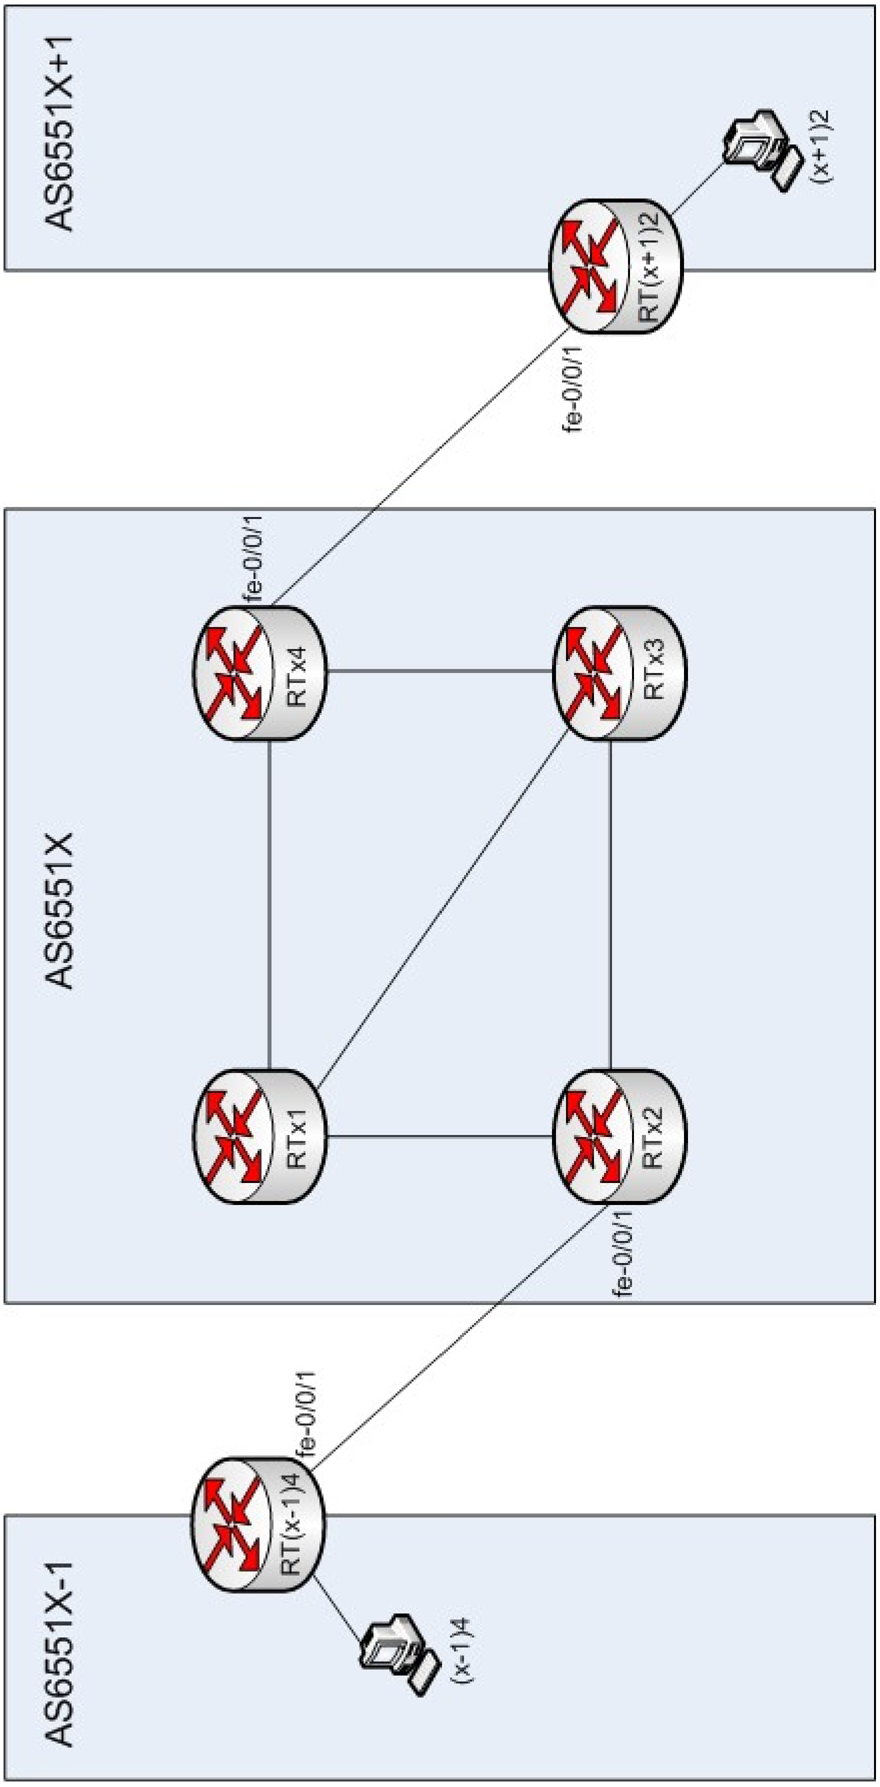
AS6551X-1
RT(x-1)4
(x-1)4
fe-0/0/1
fe-0/0/1
AS6551X
RTx1
RTx2
RTx4
RTx3
fe-0/0/1
fe-0/0/1
AS6551X+1
RT(x+1)2
(x+1)2