

ETSF10: Internet Protocol

Routing Project Assignment

1 Introduction

The goal of this project is to study two different routing protocols for intra domain Internet routing in two different error situations. The project will also give a hands-on experience in configuring and managing routers and also how to set up routing protocols. You will prepare and perform the practical exercises, and document your results in written reports.

In this project you will work in groups of one to two students.

1.1 Overview

The assignment consists of four sections:

- *Hands on*
 - Task 1 – 6
 - Section 2.1, pp 4 – 8
- *RIP*
 - Task 7 – 11
 - Section 2.2, pp 8 – 10
- *OSPF*
 - Task 12 – 18
 - Section 2.3, pp 10 – 14
- *Final report*
 - Task 19

The project has 5 milestones, M1 – M5, and one deliverable, D1. Each milestone finishes with an action or a short written report answering a few questions, while the deliverable is the project's final report. All milestones and the deliverable will be assessed passed or no passed.

- M1: Forming groups
- M2: Familiarise with the lab, status report
- M3: Hands on, answer to questions
- M4: RIP, answer to questions
- M5: OSPF, answer to questions
- D1: Final report: A comparison between two routing protocols in two error situations.

For deadlines see the course home page.

The carrying out of the practical assignments is performed on the department's router lab. The lab is accessible from a computer of your choice via Internet. The only demand on your computer is that it has an ssh client and an Internet connection.

You will practise on routers from Cisco and thus Cisco's router operation system IOS. As can be understood from the description of the router lab given in the Reference Guide the lab is fixed configured on the physical level. You change the outline of the lab networks by opening and closing individual router interfaces to build a network for your specific needs.

The preparations are a crucial part of the project. Without them there is no possibility to fulfil the practical elements within the specified time; trial-and-error is not an efficient way of using the router lab.

When you are ready for the practical sessions you book the lab; find more information on the course home page.

1.2 Project description

Each project section has these elements:

- Study the Router Lab documentation.
- Plan the laboratory moments.
- Perform the practical elements.
- Add the results and comments of the practical exercises to the reports.
- Complete the reports and hand in. It might be the case that you have to revise them after the teacher's assessment.

The router's physical interfaces are a bit special; most of them are actually ports on a switch. Read more about this in section 3.1.1 in the Reference Guide to the Router Lab.

In this project you will study two error conditions: loss of link and loss of neighbour, see Figure 1. If rtrA's interface experiences loss of link, Figure 1b), it can of course detect and react to that much in the same way as your computer. The fact that a network has disappeared can immediately be transferred to and included in the routing process. The case of loss of neighbour, Figure 1c), is somewhat more elaborate where rtrA experience link up and there is no way for rtrB to tell rtrA that the former has lost its connection. In this situation routers exchange keep-alive messages; lack of expected incoming keep-alive or hello messages can indicate loss of neighbour.

The loss of link error can be emulated by shutting down an interface. The loss of neighbour error is a bit trickier to emulate. In this project thos os done by filtering out ip packets.

1.3 Preparations

You must study the Router Lab Reference Guide as a first step in the preparation process. You must also make yourself familiar with the documentation that is given you as links on the Router Lab home page. And of course, you must study and understand appropriate theory elements of the course textbook. The FAQ found on the Router Lab web site is also a good source of information.




Figure 1: Loss of link and loss of neighbour

Please, make sure you have prepared yourself properly before using the lab. Only one group at a time can access the lab; thus available time in the lab is a limited resource. Even though most actions are specified in detail in this manual, you are advised to make notations beforehand on what steps you need to take and which commands to execute to fulfil the specified tasks.

1.4 The Reports – Milestones and Deliverable

The milestones M2 – M4 shall consist of answers to the tasks found in this manual and some printouts from commands. To ease the correcting process you are asked to add the tasks and its points – bold in this document – and enter your answers below each separate task. Failure to do this will lead to the milestone not being corrected.

In the deliverable *Final Report* you shall discuss the behaviour of the two routing protocols in the two different error situations.

Remember to mark all reports with group id and both student's names and mail addresses.

1.5 Cleaning Up

Before leaving the router lab, please clean up. If you want to save your configuration files, use ftp or tftp to copy the running configuration. If you have changed only the running configuration the cleaning up procedure is merely a matter of reloading the routers. Otherwise, you have to restore the start-up configuration with the default configurations. See the Router Lab Reference Guide for details.

1.6 Bug reports

Please report any errors found in this manual or in the lab reference material to the author Jens_A.Andersson@eit.lth.se.

2 The Tasks

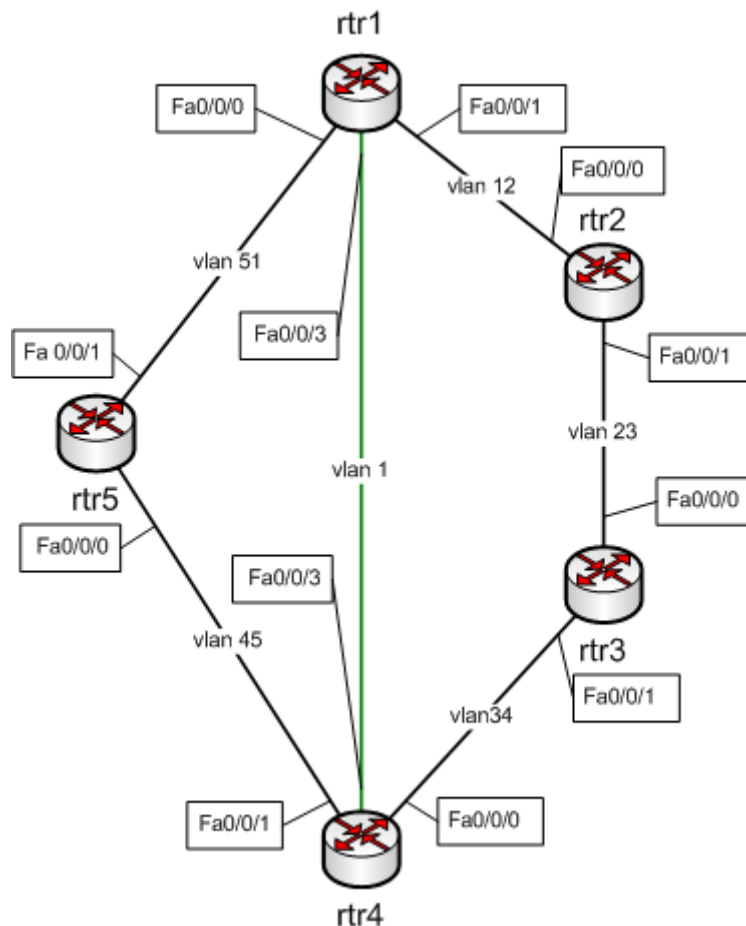


Figure 2: Lab network layout

2.1 Hands on

This section is about getting acquainted with the lab and some basic network configuration and management. Milestones M2 - section 2.1.1 - and M3 - sections 2.1.2 to 2.1.10 - are the reports from this section.

2.1.1 Preparations (Milestone M2)

Prepare yourself for this section by reading the Routerlab Reference Guide and perform the tips found in section 7 of the guide. Milestone M2 concludes the preparation with a short status report.

2.1.2 Login to the lab

You start this section with logging in to the front-end host. From there you connect to the different routers. Check that they are running the default configuration. If not, install the default configuration according to the Reference Guide. Also remember that it is a good advice to reload the routers so that the time marks in logs and debug printouts will be useful, see Reference Guide Section 6.5. See also chapter 7 in the Reference Guide for performing this section.

Hint: It is possible to create several sessions to the front-end host and then use each session for a unique connection to a router. You thereby have one terminal window for each router.

2.1.3 Configuration

Configure the lab according to Figure 2 and with ip addresses as per the Table 1 below. You do this by copying each routers lab configuration to that routers running configuration. Use the configure replace command as in this example for router rtr1:

```
Configure replace flash:rtr1-lab-config
```

For router rtr2 the configuration file is named rtr2-lab-config, and so on.

Router	Interface	IP address
Rtr1	Vlan 1	192.168.1.1
	Vlan 12	192.168.12.1
	Vlan 51	192.168.51.1
Rtr2	Vlan 12	192.168.12.2
	Vlan 23	192.168.23.2
Rtr3	Vlan 23	192.168.23.3
	Vlan 34	192.168.34.3
Rtr4	Vlan 1	192.168.1.4
	Vlan 34	192.168.34.4
	Vlan 45	192.168.45.4
Rtr5	Vlan 45	192.168.45.5
	Vlan 51	192.168.51.5

Table 1: IP addresses

Hint1: When performing the lab assignment the interface Fa0/1 is closed on the routers. You must open this whenever you need to transfer data between a router and the front-end host.

Hint2: You open interfaces with the interface configuration command `no shutdown` and you close it with `shutdown`. Notice that the default state of interfaces varies. A good command to check status of all interfaces in a router is `show ip interface brief`.

Hint3: Keep in mind that you configure network layer parameters on the virtual interfaces, vlan 1, vlan 12, vlan 23 and so on, for all links *except LAN 1 and LAN 2 according to Figure 1 in the Reference Guide*. This is a little bit tricky, so discuss this with your supervisor before continuing if you feel that you do not completely understand.

2.1.4 CDP

Use CDP and check your configuration.

Task 1:

- a) Try the command `show cdp neighbors` and on some routers in the network.
- b) On `rtr1` try the command `show cdp neighbors detail`.
- c) Select what you think are the three most important parameters you get information about issuing these commands. Motivate from a network managing perspective why you have chosen these three parameters. (3p)

2.1.5 ping and traceroute

Make yourself familiar with IOS ping and traceroute commands. Both commands come in two versions, normal and extended. We are only examining the normal version.

Task 2:

- a) From router `rtr1`, ping a neighbour using the command `ping <host>`.
- b) How many packets are sent? (1p)
- c) What is the round trip time? (1p)
- d) How long time does the router wait until a ping is declared lost? Which events demarcate this time span? (2p)

Traceroute is not meaningful as of yet but you can try it anyway.

2.1.6 Check routing table

Check the routing table of some of the routers. For this you use the command `show ip route`. You will of course find entries for the networks that a router knows of from that router's configuration. But what about the networks at least one hop away?

Task 3:

- a) Describe which of all the network ids configured in the lab that are found in the routing table for `rtr1`, and why only these are found in this situation. For example, does `rtr1` know of network id `192.168.23.0/24`? How could `rtr1` get knowledge of all the networks in the lab? (5p)

2.1.7 Debug

Next we will look into the debug command. Make sure that the output from the debug process is written to your terminal by issuing the command `terminal monitor`.

To debug ICMP packets you can use the command `debug ip icmp`. Ping one of `rtr1`'s neighbours. Also, ping `rtr1` from one of `rtr1`'s neighbours.

Task 4:

- a) Describe shortly the information you get in this exercise. You only see one of the two icmp echo message types. Which? (2p)

To get a hint of the detailed information you can get from the debug command try adding `debug ip packet detail` and repeat the ping command.

Before you end this sub-section you should turn off debugging. Use the command `no debug all` to turn off all debugging.

2.1.8 Trouble shooting exercise

CDP is very useful in trouble shooting. To test this, and your abilities to solve problems ;-) we introduce an error by installing a faulty configuration in rtr4. The file rtr4-error-confg is found on flash disk. Use the configure replace command to install the error.

Your task is now to find out what the error is and correct it. **The physical configuration is such that rtr4 only has one connection and that is to rtr1.** The corresponding interface is configured with an IP address, and the interface is up. Even though it is impossible to connect to it over IP; It is not possible to ping rtr4 from rtr1 for example.

- 1) From rtr1 try to ping rtr4. See the Table 1 for addresses.
- 2) Check rtr1 interface's states. Interfaces in question are f0/0/3 and vlan 1.
- 3) Use `show cdp neighbor f0/0/3 detail` and try to find the cause of the error. Hint: It has to do with IP addresses.
- 4) In the error situation you cannot connect to rtr4 but it is truly possible to correct the error without cheating via the console. Instead, adopt rtr1 to the error situation so that it is possible to directly connect to rtr4 from rtr1. Note that you can only use the interfaces that directly connect rtr1 and rtr4! Try your rtr1 configuration changes with ping before trying to connect to rtr4.
- 5) If ping now gives a positive result use telnet in rtr1 to connect to rtr4 and correct the error in rtr4's configuration. In this situation it is very likely that you lose the connection to rtr4 again and that telnet hangs. Use the escape method (see paragraph 6.4 in the Reference Guide) to revert back to rtr1's prompt.
- 6) Lastly, change rtr1's configuration to the correct configuration and the connection to rtr1 should be restored.

Task 5:

- a) Describe shortly what the error was, and the steps you took to fixed it. (5p)

This exercise might seem a bit weird, but it is actually based on real life experience.

2.1.9 Basic router behaviour

For the exercises in section 2.2 RIP (Milestone M4) and 2.3 OSPF (Milestone M5) we need some understanding of the basic behaviour of the router in two situations. If you read ahead you can see that rtr1 and rtr4 and networks 192.168.1.0 and 192.168.34.0 are involved in these exercises. Thus, you shall study the effect shutting down some interfaces has on the routing table of rtr1.

First, shut down **rtr1's interface f0/0/3** and inspect rtr1's routing table regarding the network 192.168.1.0. Open this interface again.

Next, shut down **rtr4's interface vlan 1** and repeat the routing table inspection. Open this interface again.

In both cases, wait until you have got a message that Vlan 1 has changed status before checking the routing table.

Task 6:

- a) What status has network 192.168.1.0 in the routing table of rtr1 when rtr1's interface f0/0/3 is shut down? (1p)
- b) What status has network 192.168.1.0 in the routing table of rtr1 when rtr4's interface vlan 1 is shut down? (1p)
- c) Explain the different outcome in a) and b) (1p)

2.1.10 Milestone 3

When you have finished the first six tasks it is time to hand in milestone M3. This milestone consists of the answers to task 1 – 6. These tasks give maximum 22 points. You must have 17 points or more to get acceptance on milestone M3.

2.2 RIP (Milestone M4)

In this section you will add RIP, a distance vector protocol, as routing protocol in your network. You are going to investigate RIP version 1 in two error conditions: loss of link and loss of neighbour.

First, make sure that all routers are configured with its lab configuration, see paragraph 2.1.3 if in doubt.

This section gives maximum 39 points. You must have 31 points or more to pass this section and milestone M4.

2.2.1 Start RIP

The first task is to start the RIP routing process on the routers. The configuration command `router rip` does this. Note that when issuing this command you enter the router rip configuration sub-mode. Configure `version 1`.

2.2.2 Add interfaces to RIP

Now you have to tell RIP which network ids that shall be involved in the RIP routing process. For this you use the router rip sub-mode command `network <network id>`. The network command takes an argument, the network id that shall be assigned to the RIP process. The network id is classful, so you must make sure that you in each router enter a network command for that router's local network ids.

With the command `show ip protocol` you get information about parameters and current state of routing processes running on the router. You also get information about which networks that are included in the routing process. Perform this command, and make sure that all network ids that shall be involved in the RIP process, and no other, are involved.

When you have configured all five routers, check the routing table in router rtr1. Make sure that the table contains a route to all active networks in the lab. The cost to reach each network is found inside brackets, e.g [120/1]. The cost is the second number. The first number is called the administrative distance, and is used to discriminate between information from different routing protocols about the same network.

Task 7:

- a) What is the cost to reach vlan 23, vlan 34, and 45 from rtr1? (1p)
- b) What is the ip address of the next hop to go from rtr1 to address 192.168.34.3? (1p)
- c) There are two paths to the network between rtr4 and rtr5. Why is that? Why isn't the path over Ethernet vlan 1 between rtr1 and rtr4, which has a capacity of 10 times that of the Ethernet link between rtr1 and rtr5, preferred? (3p)

2.2.3 Study RIP updates

Turn on debugging of RIP in rtr1. Use the command `debug ip rip`.

Task 8:

- a) How often does rtr1 receive updates from its neighbours? (1p)
- b) Which events triggers the sending of updates? (1p)
- c) Which routers announce network 192.168.34.0 to rtr1, and with which cost? (2p)
- d) Which networks are not announced to rtr1 by its neighbours? What do these networks have in common? What is this suppressing technique called? (3p)
- e) How often does rtr1 send updates to its neighbours? (1p)

The answer to task 7d) can also be found using the command `show ip protocol`. Turn off debugging again.

2.2.4 Check path

Now that you have full routing running in your network you can check if packets are forwarded the expected paths. You shall compare the routing table of rtr1 and the result of a traceroute.

Task 9:

- a) From the routing table in router rtr1 manually find the best path to interface vlan 34 of router rtr3. Which path will packets follow? Explain why? (2p)
- b) Do a traceroute to this interface. Which path did the packets follow? (1p)

2.2.5 Study convergence in error situations

In the last exercise on RIP you will study routing convergence when a failure occurs. For this you need two active terminal sessions, one to router rtr1 and one to router rtr4. We will use interface vlan 34 on rtr3 with ip address 192.168.34.3 as the target host in these exercises.

To get valid results from this task it is strongly advised to thoroughly study all the steps below before exercising them. You also need the result from Task 6 in section 2.1.9 Basic router behaviour as well as information from Task 7 and Task 8.

It is perfectly okay to write commands at the same time as you have debugging turned on. The only problem is that the print on the screen will be mixed up a bit.

It is advisable to log the output from the telnet session to rtr1 during these exercises.

1) Initial information

- a) First traceroute, from rtr1 and make sure that the connection is ok and that the path is over rtr4.
- b) Check the routing table in rtr1 and look for the target network id 192.168.34.0. Make a note of the cost and next hop to the target network.

2) Loss of link error

- a) Turn on debug routing with command `debug ip routing`.

- b) Shut down **interface vlan1 on rtr1**.
- c) Follow the debug output, best done in a log file, and try to estimate the time from the Line Protocol Down message for Vlan 1 until a new path for the target network has been established. Which events demarcate this time span? Look for 'RT: delete' and 'RT: add messages' regarding the target network, as well as the messages for routes via 192.168.1.4.
- d) After finalising 2c), restore the link on vlan1 on rtr1.
- e) Follow the debug output again and try to estimate the time from Line Protocol Up message for Vlan 1 until the original path to the target network is restored.
- f) Check the routing table and with ping and traceroute that the original path is restored.
- g) Repeat the emulated loss and restore of link as above, but this time without debugging. Instead periodically check the routing table with regard to the target network and ping the target host. From 2c) and 2e) you should have noticed interesting events. Compare your findings there with findings here.

3) Loss of neighbour error

- a) First, check routing is back to original state.
- b) Make sure you have `debug ip routing` and `debug ip rip` turned on.
- c) Next, emulate a loss of neighbour. This is done by denying any incoming ip traffic on rtr1's vlan. Add the configuration command `ip access-group 99 in` to **interface vlan 1 in rtr1**.
- d) In rtr1, keep track of the debug output, especially regarding the target network. This is easiest done by logging at least 5 minutes of debug output from rtr1, starting from where you apply the filter in c). In the log file, search for the string 'network 192.168.34.0 metric' and notice the cost. This is the cost which rtr1 advertise it has to reach the target network. In the normal state this cost should be 2. Notice the time for when this cost is set to 16 (infinity; unreachable) and 3. From this you can estimate the convergence time. Also notice that rtr1 receive updates from two neighbours during the whole process.
- e) After finalising 3d), restore the network by removing the filter from vlan 1 in rtr1 and observe as in point 3d). Estimate the convergence time also in this case.
- f) Repeat the emulated loss and restore of neighbour as above, but this time without debugging. Instead periodically check the routing table with regard to the target network and ping the target host. From 3d) and 3e) you should have noticed interesting events.
- g) When routing has converged note which path is the preferred path to the target.

Task 10:

- a) Give an explanation to the findings in the loss of link error exercise above. Recall the answers you gave to Task 6 and Task 8 above. How long does it take before rtr1 reacts to the error? How long does it take before the routing has fully adapted to the new situation? And how long does it take before it reacts when you restore the status of the interface you manipulated? (5p)
- b) Which path is the preferred path according to the routing table when routing has adapted to the error situation? And is this in accordance with your findings in Task 6 and Task 8? Motivate. (5p)

Task 11:

- a) Repeat Task 10 a) above but this time for the loss of neighbour case. Note that in this case you will be able to observe more events during adaption to the error. (5p)
- b) Repeat in Task 10 b) above for the loss of neighbour case. (5p)

- c) **Why aren't the alternative incoming announcements for the target network affecting rtr1's routing table? (3p)**

2.2.6 Remove RIP

Now you can remove all configurations concerning RIP. Use the configuration command `no router rip` in all routers, or reload the respective lab configurations.

2.2.7 Milestone M4

The milestone M4 consists of answers to the tasks 7 – 11.

2.3 OSPF (Milestone M5)

In this section you study how a link state based routing protocol, OSPF, will perform in the two error situations loss of link and loss of neighbour. You will use the same network layout and target host as in the RIP exercises.

This section gives maximum 40 points. You must have 32 points or more to pass this section and milestone M5.

2.3.1 Start OSPF

In OSPF every router has a unique identity. Cisco routers use the highest IP address assigned to any interface in the router as the router id. The method works well until that interface goes down in which case the router id has to change and the OSPF process has to recalculate and resend its information to its neighbours. It is an advantage if the router id could be fixed, and independent of interface status. Therefore, Cisco IOS has given precedence to so called loop back interfaces in the selection of router id. If there is a loop back interface configured on the router, and the loop back interface has been assigned an IP address, this IP address will be used as the router id. A loop back interface has the benefit that it is very stable; it never goes down unless the whole router is going down. The conclusion is that it is good practise to always configure a loop back interface in every Cisco router that is to run OSPF.

Start with configuring loop back interface 0 on all routers. A loop back interface is created the first time you issue the configuration command `interface loopback interface number`. Give each loop back interface a unique IP host address. You can use any ip address for this but a suggestion is that you assign addresses 10.0.0.1 – 10.0.0.5 for the respective routers. Hint: A so called host route is assigned if you combine the IP address with a subnet mask with all ones, i.e. 255.255.255.255.

Now you can initiate OSPF on the routers. The configuration command is very similar to that you used when you started RIP. There is one difference: OSPF needs a process ID. Pick a number from 1 to 65535 of your choice. Use the configuration command `router ospf process ID`.

2.3.2 Add interfaces to OSPF

Adding network ids to the OSPF routing process is similar to that of RIP in that you use a `network` sub-mode command. But in the case of OSPF you have to state the network address space with a combination of a network id and a so-called wildcard mask. You also define

which area, or areas in the case of an area border router, a router's interfaces belong to with the same command.

The wildcard mask is the one's complement to that of a subnet mask. Any bit set in the wildcard mask indicates a position in the network id that has no significance for area assignment.

In our case we will only use one area, the backbone area, for all networks. We therefore can use the same network command in all routers. The general form of the network command is

```
network address wildcard-mask area area-id
```

Task 12:

- a) **What single command can you use to assign all interfaces in one router to OSPF backbone area? (1p)**
- b) **In this case we only want the networks in address space 192.168.0.0/16 to be part of OSPF routing in the backbone area. What single network command does this? (1p)**

You can check your OSPF configuration with command `show ip ospf` or `show ip ospf neighbor`.

When you have configured all five routers, check the routing information in router `rtr1`. Compare this output with the findings in exercise 2.2.2.

Task 13:

- a) **What is the cost to reach vlan 23, vlan 34, and 45? (2p)**
- b) **In what way does the cost differ from the cost that RIP uses? What network parameter does the OSPF cost reflect? (2p)**

Hint: Use the command `show ip ospf interface` to find more information about costs.

2.3.3 Check path

Repeat exercise 2.2.4, but now with OSPF as routing protocol. You shall compare the routing table of `rtr1` and the result of a `traceroute`.

Task 14:

- a) **From the routing table in router `rtr1` find the path to interface vlan 34 of router `rtr3`. Which path will packets follow? Explain why? (2p)**
- b) **Do a `traceroute` to this interface. Which path did the packets follow? (1p)**

2.3.4 Hellos and flooding

Go to router `rtr1` and turn on OSPF debugging. In this case you shall study one sub function of OSPF, so start debugging with the command `debug ip ospf hello`.

Task 15:

- a) **How often are hello events occurring? You should be able to identify two different hello events (2p)**
- b) **What is the main purpose of the hello messages sent and received? (2p)**

Turn off the debugging of OSPF hellos, and instead turn on debugging of OSPF flooding, command `debug ip ospf flood`.

Task 16:

- a) Do you see anything? (If you do not get any debug output you can temporarily change the status of for example vlan 34 in rtr4.)
- b) What kind of messages is flooded by OSPF? (1p)
- c) In which situations are these messages sent? (1p)
- d) What is the content of the information sent and received? (2p)

2.3.5 Study OSPF updates and convergence in error situation

We shall now repeat the two errors that we studied in the RIP section in exercise 2.2.5.

1) Initial information

- a) First traceroute, from rtr1 and make sure that the connection is ok and that the path is over rtr4.
- b) Check the routing table in rtr1 and look for the target network id 192.168.34.0. Make a note of the cost and next hop to the target network.

2) Loss of link error

- a) Turn on debug routing with command `debug ip routing`.
- b) Shut down **interface vlan1 on rtr1**.
- c) Follow the debug output, best done in a log file, and try to estimate the time from the Line Protocol Down message for Vlan 1 until a new path for the target network has been established. Which events demarcate this time span? Look for 'RT: delete' and 'RT: add' messages regarding the target network, as well as the '%OSPF-5-ADJCHG' messages.
- d) After finalising 2d), restore the link on vlan1 on rtr1.
- e) Follow the debug output again and try to estimate the time from Line Protocol Up message for Vlan 1 until the original path to the target network is restored.
- f) Check the routing table and with ping and traceroute that the original path is restored.
- g) Repeat the emulated loss and restore of link as above, but this time without debugging. Instead, periodically check the routing table with regard to the target network and ping the target host. From 2c) and 2e) you should have noticed interesting events. Compare your findings there with findings here.

3) Loss of neighbour error

- a) First, check routing is back to original state.
- b) Make sure you have `debug ip routing` and `debug ip ospf hello` turned on.
- c) Next, emulate a loss of neighbour. This is again done by denying any incoming ip traffic on rtr1's vlan. Add the configuration command `ip access-group 99 in` to **interface vlan 1 in rtr1**.
- d) In rtr1, keep track of the debug output, especially regarding the target network. This is easiest done by logging at least 1.5 minutes of debug output from rtr1, starting from where you apply the filter in c). In the log file, follow the hello messages from rtr4. Also observe the change of neighbourhood state found in the '%OSPF-5-ADJCHG' messages as well as the 'RT: add' and 'RT: delete' messages.
- e) Restore the network by removing the filter from vlan 1 in rtr1 and observe as in point 3d) above. Estimate the convergence time is also in this case.
- f) Repeat the emulated loss and restore of neighbour as above, but this time without debugging. Instead periodically check the routing table with regard to the target network and ping the target host. From 3d) and 3e) you should have noticed interesting events. Also, you most probably have to be quick this time.

g) When routing has converged note which path is the preferred path to the target.

Task 17:

- a) Give an explanation to the findings in the loss of link error exercise. How long does it take before rtr1 reacts to the error? How long does it take before the routing has fully adapted to the new situation? How long does it take before it reacts when you restore the status of the interface you manipulated and which events demarcate this time span? (5p)
- b) Which path is the preferred path according to the routing table when routing has adapted to the error situation? (5p)

Task 18:

- a) Repeat Task 17 a) above but this time for the loss of neighbour case. Note that in this case you will be able to observe more events during adaption to the error. (5p)
- b) Repeat Task 17 b) above. (5p)
- c) Why aren't the alternative incoming announcements for the target network affecting rtr1's routing table? (3p)

2.3.6 Milestone M5: OSPF

This milestone consists of answers to the tasks 12 – 18.

3 Final Report (deliverable D1)

Now you have all the necessary information to write the deliverable D1 Final Report.

Task 19:

Describe and compare the behaviour of RIP v1 and OSPF in the two error situations lost link and lost neighbour. Consider such metrics as time to detect error and time to restoration of service during error condition as well as time recover from the error when going back to normal conditions. Also describe how a router gets information about the existence of a neighbour and how a neighbour is declared as dead in RIP and OSPF.

You are expected to describe the lab setup, the way you created the two different errors and the lab results that you use as basis for the discussion in task 19. You must also describe how you divided the work load i.e. who did what.

The target audience for the report is your class, that is knowledge of the theory for RIP and OSPF is expected.

The report must contain no less than 1000 words and no more than 1500 words.

Fail to fulfil these requirements and the deliverable will not be accepted.

As an appendix to the report you are asked to present how you as a group and as individuals fulfilled the project, who did what and why you chose to do as you did. Also please give your personal comments on this lab. How did you experience the lab? What was good, what was not so good? Was it relevant for the course? The appendix is not included in the word count.

4 Documentation

- URL <http://www.eit.lth.se/index.php?id=rtrlab>
- Reference Guide to the Router Lab (link on router lab web page)
- Course Textbook

5 Acknowledgements

Thanks to William Tärneberg, EIT, for proofreading and valuable comments. Thanks also to Anders Ahlberg, Genombrottet, LTH, for supportive discussions regarding the overall make-over of the project.