# ELEM011 Routing Experiment

**November 2010**

**Aim**: The purpose of this laboratory experiment is to familiarise yourself with the operation of Routing Information Protocol (RIP) and Open Shortest Path First (OSPF) routing protocols. This will involve examining a number of their configuration parameters and carrying out some performance tests. You are required to note down the results of various tests marked as 'Log *x.y*' throughout this lab-sheet.

**Background**:

Routers

This experiment uses 8 virtual routers running Junos (Juniper Network Operating System). Each router has up to four interfaces, labelled em1.0, em3.0, em4.0 and em5.0. The routers are connected in accordance with Figure 1. In addition to the routers, there are two virtual PCs which facilitate the monitoring of network traffic.

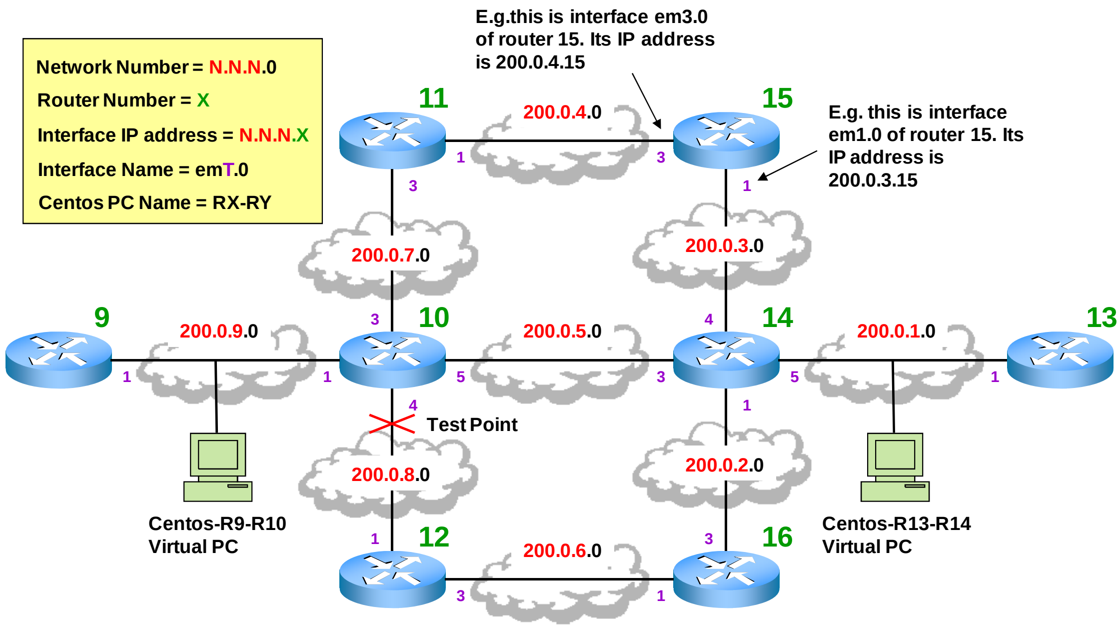


**Figure 1:- Router Topology, Subnet and Interface Assignments in the Virtual Juniper Cloud**

Figure 1 shows the router topology and the address scheme used throughout this lab-sheet. You can use one of two methods in order to login to the virtual routers and PCs. In both cases you first need to open a command window on your local machine. After this you can use either:

1) SSH (Secure SHell)

ssh *username@ip_address* -p *port_number*

At this point you may receive a warning with a prompt. Type "yes" and press [enter] in order to confirm you wish to continue. You will then be prompted for your password which is provided via email.

*2)* Telnet

telnet *ip_address port_number* (followed by [enter] twice)

When you do this you will be required to enter your username and password (which is provided via email).

The IP address and 10 port numbers (i.e. one per each virtual router/PC) are also provided to you via email. Furthermore, using SSH or Telnet it is possible to connect to multiple routers/PCs concurrently by opening separate console windows.

Once you are successfully logged in you will see a prompt similar to the following: "*username*@R*X*-Group*N*>", where *X* indicates the specific virtual router you are logged onto[1] and *N* indicates your group number. If you see the symbol '%' instead of '>' (i.e. "*username*@R*X*-Group*N*%") you MUST execute the command **`cli`** in order to obtain the desired prompt (i.e. "*username*@R*X*-Group*N*>").

If you wish to download a file from a virtual machine to your local PC you must first open a command window and type the following command:

scp –P {*port number*} {*username@i_ address*}:*src_filename destination_directory*

When prompted with a warning, type "yes" and press [enter]. Next type in your password followed by [enter]. The file will then be downloaded[2].

<u>Junos Primer</u>

Junos is an operating system for Juniper routers. It supports a rich-set of commands, some of which you will need to use, namely:

| | |
|---|---|
| **`?`** | Obtain a list of commands available within the current context. Can be repeatedly applied to obtain context-sensitive information. For example, in any router typing **`?`** at the prompt will display a base list of commands including **`show`**. If you then type **`show ?`** this reveals a variety of **`show`** options such as **`show route`** and **`show configuration`**. You can also type **`show route ?`** to obtain further options and so on. |
| **`show`** | Show system information |
| **`edit`** | Enter the edit mode |
| **`traceroute`** ***`Addr`*** | Locate the path taken by IP datagrams to the destination IP address given by the ***`Addr`*** field. Approximate timing information is provided. (*) means unreachable. |
| **`ping`** ***`Addr`*** | Test the reachability of the destination IP address given by the ***`Addr`*** field. "Host unreachable" means the target machine/network cannot be located. |

1 In the case of the virtual CentOS PCs the prompt will be of the format "[*username*@centos-R*X*-Group*N*-R*Y*-Group*N* ~]$" where X and Y are the two routers that it shares a network with, as illustrated in Figure 1.
2 The configuration file for each router is called "olive.conf" and can be downloaded if required.

In order to go back to the base configuration of a router at any point, enter the following commands in the **edit** mode:

```
load override /root/olive.external.conf
load merge /root/olive.conf
commit
```

Note: do NOT enter the **commit** command until you enter BOTH of the **load** commands.

Centos Primer

Centos is an operating system for PCs, derived from UNIX. Some commands you will need to use are:

| | |
|---|---|
| **ls** | List directory contents |
| **pwd** | Display current directory |
| **sudo tcpdump -*options*** | Capture packets. Typical options are –vv, -X and –i *interfacename*. To capture tcpdump to a file whilst watching the capture events, the following should be used: |
| | `sudo tcpdump -vv -X -i interfacename -l | tee filename` |
| **more *filename*** | Observe the contents of a file page by page |
| **man *command*** | Provides a manual page description of the command, if available. |

**Part 1: Routing Information Protocol (RIP) Basics**

To start with we will perform some tools familiarisation and make a number of basic RIP observations. All the routers are initially configured to run RIP version 2 and are arranged in accordance with Figure 1.

Typing **show configuration** displays the current router configuration. Note that OSPF is currently disabled on all routers.

> *Log 1.1:* Whilst logged onto R9, type **show route brief** to display the summarised routing table. Make a note of a typical line starting with "200.0…"and (later) try to explain the meaning of each of the fields.

To display information that has been specifically learnt through RIP updates the command **show route protocol rip** is used. However, note that this will omit entries for destination networks that the router is directly connected to. To obtain additional information it is possible to type **show route detail**. If appropriate, the [spacebar] can be used to display the next page of information.

> *Log 1.2:* Now enter the command **show route protocol rip detail** and examine the metric values for different destination networks including those that are directly connected. Check the Metric (hopcount) with Figure 1. Does the routing table give the complete path say from R9 to R16?

*Log 1.3:* Enter the command **`show route protocol rip detail range 200.0.8.0/24`** repeatedly (by using the up arrow) and observe the time value associated with the "Route learned from…" entry. Note the lowest and highest values you see.

*Log 1.4:* Whilst still logged into R9, type **`traceroute 200.0.6.16`** and note down the information provided. This may take a little while. Later (in the report) describe the meaning of the various information fields and explain how **`traceroute`** employs Internet Control Message Protocol (ICMP) messages to work out the path.

Note that you can copy and paste selected console output to a text editor[3].

*Log 1.5:* Type **`ping 200.0.6.16`** and observe what happens. To finish the information stream, press <Ctrl> and the "C" characters together. Note down a sample line of output and (later) try to explain the meaning of each of the fields.

*Log 1.6:* Next, within a separate command window logon to Centos-R9-R10 and in order to capture the packet flow on network 200.0.9.0 type the following:

```
tcpdump -v -X -i eth1 –l | tee log16.txt
```

Note the IP protocol number(s) and the source and destination IP addresses being used. Also note down the time between several successive messages[4]. To view the captured file simply type **`more log16.txt`**. The left-hand column shows the offset in hexadecimal. The central portion shows the captured data in hexadecimal format and the right-hand column shows the ASCII equivalent of the captured data. After the Lab try and explain the meaning of the first 20 bytes of one of the captured IP packets. The first byte is "$45_{hex}$".

**Part 2: Examining Failure and Reconvergence with RIP**

In this part you will introduce a "link failure" while RIP is in operation.

*Log 2.1:* From R9 type **`traceroute 200.0.6.12`** and note the information obtained. You can copy this information to a file for later viewing.

*Log 2.2:* Then type **`ping 200.0.6.12`** and note the typical values of the TTL and Time fields. You can let **`ping`** continue for now.

3 To save the output you can right click on the menu bar and select edit → mark. Next highlight the information in the window that that you wish to copy and press [enter]. The copied data can then be pasted into any text editor.

4 A timestamp is displayed at the beginning of each message.

*Log 2.3:* In a separate command window logon to R9 enter the command **`show route protocol rip detail`**. Observe the entry for network 200.0.8.0 noting the metric.

Logon to R10 and type **`edit`**. This will cause the prompt to change to "*username*@R10-Group*N#*". In addition, above the prompt you will see [edit] which will change depending on where you are within the edit hierarchy.

Type **`edit interfaces`** and then type **`show`**. You should see the current configuration of the various router interfaces. To mimic the failure of a link we will disable em4.0. To do this type **`edit em4.0`** followed by **`set disable`**. If you then type **`show`** you should see a new line indicating that this interface is configured to be disabled. However, changes to the configuration file do not take effect until the **`commit`** command is entered. Once you enter the **`commit`** command, if successful you should see a "commit complete" acknowledgement. Finally, leave the edit mode by typing **`exit`** until you are back at the "*username*@R10-Group*N*>" prompt.

*Log 2.4:* Having now "failed the link", on R9 type **`show route protocol rip detail range 200.0.8.0/24`**. Repeat this command a number of times over a 5-minute interval (using the up arrow). Note the "Age" value once "Route learned from" expires.

*Log 2.5:* Continue this process until the ping messages start to flow again. How long does it take for reachability to be restored[5] and what else is different in the **`ping`** fields from before the failure and why? At R9 observe the value of the Metric field for network **`200.0.8.0/24`**.

*Log 2.6:* Also run **`traceroute 200.0.6.12`** and explain any differences that you see compared with before the failure.

While the ping utility is still running on R9; go to R10 and "restore the link" by re-enabling the interface. To do this, whilst in [edit interfaces em4 unit 0], type **`delete disable`** and **`commit`**. Type **`show`** to make sure that the disabled statement has been removed.

*Log 2.7:* With ping, observe any change(s). If any, how quickly did these change(s) take place?

*Log 2.8:* Rerun **`traceroute 200.0.6.12`**. Note the indicated path.

*Log 2.9:* Finally, whilst on R9 observe the value of the Metric field for network **`200.0.8.0/24`**.

[5] You can determine the time until reachability is restored using ping by noting the sequence numbers when the ping flow stops (i.e. failure) and when the flow is restored, as each ping attempt is1 second apart.

**Part 3: Open Shortest Path First (OSPF) Basics**

This time we will make a number of basic OSPF observations. Firstly, on R9 through to R16 perform the following setup:

- Enter the edit mode and type **`edit protocols rip`**
- Type **`set send none`**
- Type **`set receive none`**
- Typing **`show`** will display the intended settings
- Type **`exit`** to return to the top level of the edit hierarchy
- Type **`edit protocols ospf`**
- Typing **`show`** will display the current OSPF setting that will include the **`disable;`** statement
- Then type **`delete disable`**
- Type **`commit`**

Make sure all 8 routers have been configured as described above. Then, whilst logged on to R9, carry out the tasks *Log 3.1* to *Log 3.6*.

> *Log 3.1:* **`show route`**
> Make a note of a typical line that includes [OSPF/10] and (later) try to explain the meaning of each of the fields[6].

> *Log 3.2:* **`show ospf interface em1.0`**
> Note the state of the interface.

> *Log 3.3:* **`show ospf database detail advertising-router 200.255.0.10`**
> Note the point-to-point entries that are listed[7].

> *Log 3.4:* **`show ospf statistics`**
> Note the different types of OSPF messages that are listed. Explain their purpose.

> *Log 3.5:* **`show ospf neighbor`**
> Note the router state (the relationship with R10). "Full" refers to "Full adjacency". What activity is permitted when this state is reached? Repeat this command several times and note the highest and lowest values of "Dead" that you see.

> *Log 3.6:* **`show ospf interface em1.0 detail`**
> Note the values of Dead, Hello and Retransmit times and the cost.

---

6 In this experiment the configuration is such that all routers belong to area 0.0.0.0, the link cost is set to 1, the hello interval is set to 10 and the dead interval is set to 40. As default values, many of these may not be explicitly shown in the configuration file.

7 In JUNOS with routing protocols such as OSPF it is customary to identify each router by a router_ID (which is used in advertisements) rather than the IP address of individual interfaces. In this experiment each router has an ID of the format 200.255.0.X.

*Log 3.7:* From Centos-R9-R10, use **`tcpdump`** and note what messages are exchanged over the network **`200.0.9.0`**. For example, type:

**`tcpdump -vv -X -i eth1 –l | tee`** ***`log37.txt`***

Note the IP protocol number(s) and the destination address(es) being used.

**Part 4: Examining Failure and Reconvergence with OSPF**

In this part you will introduce a "link failure" while OSPF is in operation.

*Log 4.1:* From R9 type **`traceroute 200.0.6.12`** and note the information displayed.

*Log 4.2:* Then type **`ping 200.0.6.12`** and note the typical values of the TTL and Time fields.

*Log 4.3:* Whilst continuing with **`ping`** in a separate window on R9 type **`show ospf route`** and **`show ospf database detail advertising-router 200.255.0.10`**. Observe the Sequence number, Age and the point-to-point link information. You will need to repeat these commands a number of times over a 5-minute interval after you "fail the link" on R10 (see below).

On R10 "fail the link" following the same procedure as described in Part 2, remembering to **`commit`** to confirm the changes.

*Log 4.4:* On R9 watch the **`ping`** display and monitor the information received from R10 as previously explained (i.e via **`database detail`**). In particular, watch the Time field (Age) and note how the value changes. Also observe the value of the Metric field. Based on the **`ping`** information how fast is reachability restored and what is different from before the failure?

*Log 4.5:* Type **`traceroute 200.0.6.12`** and note the information obtained.

*Log 4.6:* Move to Centos-R9-R10 and use **`tcpdump`** to note what messages are exchanged over the network **`200.0.9.0`**. For example, type:

**`tcpdump -vv -X -i eth1 –l | tee`** ***`log46.txt`***

*Log 4.7:* Whilst tcpdump is running on Centos-R9-R10, return to R9 and restart **`ping 200.0.6.12`** (if it is not already running) noting down typical values that the ping utility is displaying, particularly for the TTL and Time fields.

*Log 4.7:* Whilst tcpdump is running on Centos-R9-R10, return to R9 and restart **`ping 200.0.6.12`** (if it is not already running) noting down typical values that the ping utility is displaying, particularly for the TTL and Time fields.

Return to Centos-R9-R10, stop the **`tcpdump`** capture and use **`more`** ***`log46.txt`*** to examine the captured file. What message types do you see and what is their purpose?

*Log 4.9:* Finally, on R9 type **`show ospf database detail advertising-router 200.255.0.10`**. Observe the Sequence number, Age and the point-to-point link information.

**Part 5: OSPF Parameter Negotiation and Cost Setting**

In this part you will experiment varying some OSPF parameters.

*Log 5.1:* From R9 type **`traceroute 200.0.1.13`** and note the information obtained.

*Log 5.2:* Go to R10 and type **`show ospf neighbor`** and note the "Full adjacency" entries.

Whilst on R10 type **`edit`**. Then type **`edit protocols ospf`** followed by **`show`** to display the current ospf configuration on R10. Next type **`edit area 0.0.0.0`**. Then type **`set interface em5.0 hello-interval 5`** to set the hello value to 5 seconds.

*Log 5.3:* Back on R9 repeat **`traceroute 200.0.1.13`** and note the path taken.

*Log 5.4:* Also on R10 once out of the **`edit`** mode, note any difference(s) when you type **`show ospf neighbor`**.

Having done this return interface em5.0 to a hello interval of 10 seconds. This can be done by entering the **`edit`** mode and then either:

1) **`set protocols ospf area 0.0.0.0 interface em5.0 hello-interval 10`** (followed by **`commit`**).
2) **`delete protocols ospf area 0.0.0.0 interface em5.0 hello-interval`** (followed by **`commit`**)[8].

*Log 5.5:* From R9 type **`traceroute 200.0.1.13`** and note the information obtained.

[8] This option works as the default hello interval is 10 seconds which does not need to be explicitly specified.

*Log 5.6:* Go to R10 and type **`show ospf neighbour`** and note the full adjacency entries. They should be restored to their original settings.

Still on R10 this time set the cost for interface **`em5.0`** to be 10. Whilst in the edit mode this can be done directly by typing: **`set protocols ospf area 0.0.0.0 interface em5.0 metric 10`** (followed by **`commit`**).

*Log 5.7:* Back R9 type **`traceroute 200.0.1.13`** and note the path taken.

Finally, you should restore the cost metric to its original value on interface **`em5.0`** remembering to commit the changes.

**Part 6: RIP / OSPF Route Redistribution**

In this final part of the experiment, separate RIP and OSPF parts of the network will be made to inter-work, sharing routing information using a mechanism called route redistribution. In this case, routers R13 and R14 will remain configured for OSPF-only. Routers R9 through to R12 and R15 will be reconfigured for RIP-only operation and R16 will be allowed to support RIP on interface **`em1.0`** and OSPF on interface **`em3.0`**.

For R9-12 and R15 do the following:

- Enter **`edit`** mode and to disable OSPF type **`set protocols ospf disable`**.
- As you are still at the top-level of the editor you can edit RIP directly. Type **`delete protocols rip send`** and **`delete protocols rip receive`** as these will have previously been set to send and receive none. Without specific entries JUNOS sends and receives RIPv2 messages by default (as long as the basic RIP configuration is present).
- Commit the changes.

R13 and R14 should be already configured for OSPF operation, therefore no changes are necessary.

For R16 do the following:

- Enter **`edit`** mode and type **`delete protocols rip send`** and **`delete protocols rip receive`** in order to re-enable RIP.
- Type **`delete protocols rip group rip neighbor em3.0`** in order to remove RIP from interface em3.0[9].
- Type **`delete protocols ospf area 0.0.0.0 interface em1.0`** in order to remove OSPF from interface em1.0.
- Type **`show protocols`** to check the configuration.
- Commit the changes.

9 Note that in the case of RIP, the term 'neighbor' is used to refer to interfaces.

*Log 6.1:* Go to R9 and R13 and examine the routing tables using **`show route`**. Note we have not enabled route redistribution at this time. This can be confirmed by going to R9 and performing a **`traceroute 200.0.1.13`**.

Go to R16 and enter edit mode. Type:

**`set protocols ospf export rip-to-ospf`** and
**`set protocols rip group rip export [ospf-to-rip redist-direct]`**.

If you wish to observe the policy options you can type (while in the edit mode) **`show policy-options`**. Once you are satisfied, type **`commit`** to accept the new configuration. On R16 this will permit the redistribution of routes learned from OSPF to be advertised within the RIP domain and conversely for the routes learned via RIP.

*Log 6.2:* To confirm the appropriate redistribution of routes go to R9 and R13 and examine the routing tables again. Also do this on R16 and comment on what you see.

*Log 6.3:* From R9 **`traceroute 200.0.1.13`** and explain what you observe.

*Log 6.4:* Go to R13. Explain why **`traceroute`** to **`200.0.3.15`** and to **`200.0.4.15`** (i.e. two interfaces of the same router) take markedly different paths.

**Part 7: Assessed report**

Your assessed report must include:

- Answers to all questions asked in part1 – part 6.
- All results (in the numbered sections as *Log x.y*)[10].
- An essay (of maximum 5 pages) addressing the following:

**Describe the operation of RIP and OSPF and discuss mechanisms that can be used to improve the reconvergence time of these interior gateway routing protocols. What are the consequences of least-cost routing in a single-service inter-network and how can source-routing mechanisms is used to improve overall network efficiency?**

The completed piece of coursework (parts1 through to 7) amounts to **25% of the total mark** for this course unit.

**Chris Phillips and Vindya Wijeratne, November 2010**
**Trial run by Ammar Lilamwala and Xian Zhang**

[10] In the Appendix.