

```
% Monitoring Netflow with NfSen
%
% Network Monitoring and Management
```

Introduction

Goals

- * Learn how to export flows from a Cisco router

Notes

- * Commands preceded with "\$" imply that you should execute the command as a general user - not as root.
- * Commands preceded with "#" imply that you should be working as root.
- * Commands with more specific command lines (e.g. "RTR-GW>" or "mysql>") imply that you are executing commands on remote equipment, or within another program.

Export flows from a Cisco router

During this exercise we will ask that you export flows from your router to two PCs in the classroom. You should work together as a group. That is, for group 1, users of pc1, pc2, pc3, pc4 should work together and pick one machine where network flows will arrive.

In addition, you will export a second flow from your group's router to a PC in the group next to yours. That is, for example, if group 2 has chosen pc5 to be the PC that receives flows, then the second flow you export will go to pc5. And, if you chose pc1 to receive flows from router 1 (rtr1), then it should, also, receive flows from router 2 (rtr2):

These exercises work on the example of doing the following:

```
Group 1, Router 1
-----
rtr1 ==> pc1 on port 9001
rtr1 ==> pc5 on port 9002
```

```
Group 2, Router 2
-----
rtr2 ==> pc5 on port 9001
rtr2 ==> pc1 on port 9002
```

You may select the combination that works for your groups.

Here are the groups that should work together:

- * group 1 and 2
- * group 3 and 4
- * group 5 and 6
- * group 7 and 8

If there is a group 9 please see the instructors.

If you had three groups, you could do:

```
rtr1 ==> pc1 port 9001
rtr1 ==> pc5 port 9001
```

```
rtr2 ==> pc5 port 9002
rtr2 ==> pc9 port 9001
```

```
rtr3 ==> pc9 port 9002
rtr3 ==> pc1 port 9002
```

... just make sure to follow the rule:

- every router exports to two destinations: one in your group, one in another
- every designated PC in the group should receive two exports: one from your group, one from another

Configuring the routers

```
~~~~~
$ ssh cisco@rtr1.ws.nsrc.org
rtr1.ws.nsrc.org> enable
~~~~~
```

or, if ssh is not configured yet:

```
~~~~~
$ telnet 10.10.1.254
Username: cisco
Password:
Router1>enable
Password:
~~~~~
```

Remember - This is an EXAMPLE for the following situation:

```
rtr1 ==> pc1 on port 9001
rtr1 ==> pc5 on port 9002
```

Group 2, 3, 4, 5, 6, 7, 8 and 9 will do something different.

The following configures the FastEthernet 0/0 interface to export flows.

```
~~~~~
rtr1.ws.nsrc.org# configure terminal
rtr1.ws.nsrc.org(config)# interface FastEthernet 0/0
rtr1.ws.nsrc.org(config-if)# ip flow ingress
rtr1.ws.nsrc.org(config-if)# ip flow egress
rtr1.ws.nsrc.org(config-if)# exit
rtr1.ws.nsrc.org(config)# ip flow-export destination 10.10.1.1 9001
rtr1.ws.nsrc.org(config)# ip flow-export destination 10.10.2.5 9002
rtr1.ws.nsrc.org(config)# ip flow-export version 5
rtr1.ws.nsrc.org(config)# ip flow-cache timeout active 5
~~~~~
```

This breaks up long-lived flows into 5-minute fragments. You can choose any number of minutes between 1 and 60. If you leave it at the default of 30 minutes your traffic reports will have spikes.

```
~~~~~
rtr1.ws.nsrc.org(config)# snmp-server ifindex persist
~~~~~
```

This enables ifIndex persistence globally. This ensures that the ifIndex values are retained during router reboots - also if you add

or remove interface modules to your network devices.

Now configure how you want the ip flow top-talkers to work:

```
~~~~~
rtr1.ws.nsrc.org(config)#ip flow-top-talkers
rtr1.ws.nsrc.org(config-flow-top-talkers)#top 20
rtr1.ws.nsrc.org(config-flow-top-talkers)#sort-by bytes
rtr1.ws.nsrc.org(config-flow-top-talkers)#end
~~~~~
```

Now we'll verify what we've done.

```
~~~~~
rtr1.ws.nsrc.org# show ip flow export
rtr1.ws.nsrc.org# show ip cache flow
~~~~~
```

Note the packet size distribution - what are the two most common packet sizes ?

See your "top talkers" across your router interfaces

```
~~~~~
rtr1.ws.nsrc.org# show ip flow top-talkers
~~~~~
```

If it all looks good then write your running-config to non-volatile RAM (i.e. the startup-config):

```
~~~~~
rtr1.ws.nsrc.org#wr mem
~~~~~
```

You can exit from the router now:

```
~~~~~
rtr1.ws.nsrc.org#exit
~~~~~
```

Verify that flows are arriving from your router to the PC chosen to receive flows in your group:

```
~~~~~
$ sudo tcpdump -Tcnfp port 9001
~~~~~
```

Wait a few seconds and you should see something that looks like:

```
06:12:00.953450 IP s2.ws.nsrc.org.54538 > noc.ws.nsrc.org.9009: NetFlow v5, 9222.333 uptime, 135
  started 8867.952, last 8867.952
    10.10.0.241/0:0:53 > 10.10.0.250/0:0:49005 >> 0.0.0.0
    udp tos 0, 1 (136 octets)
  started 8867.952, last 3211591.733
    10.10.0.241/10:0:0 > 0.0.0.0/10:0:4352 >> 0.0.0.0
    ip tos 0, 62 (8867952 octets)
[...]
```

If you are using Netflow v9, do note that the above output may not be correct, as the tcpdump in this version of Ubuntu does not decode Netflow

v9 properly.

Verify that flows are arriving from the router in the group next to you to the PC chosen to receive flows in your group (you may have to wait until the group next to you is ready and exporting flows to your PC):

```
~~~~~  
$ sudo tcpdump -Tcnfp port 9002  
~~~~~
```

You are done for this lab.

Move on to exercise3-NfSen-PortTracker if NfSen is already installed.

Otherwise, go to exercise2-install-nfdump-nfsen.