

Monitoring Netflow with NfSen

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. “rtrX>” or “mysql>”) imply that you are executing commands on remote equipment, or within another program.

Export flows from a Cisco router

We will ask that you work in pairs: e.g. for group 1 one pair will be pc1/pc2 and the other pair will be pc3/pc4. (If there are only three people in your group then the third person will do it all by themselves)

Because your Cisco router can only export flows to two destinations simultaneously, we will use the following configuration:

Group 1, Router 1

```
rtr1 ==> pc1 on port 9001  
rtr1 ==> pc3 on port 9001
```

Group 2, Router 2

```
rtr2 ==> pc5 on port 9001
rtr2 ==> pc7 on port 9001
```

etc. Therefore flows will only arrive at the first PC in each pair. However, when nfsen is installed, both people can point their web browser to the first PC.

Configuring the routers

```
$ ssh cisco@rtrX.ws.nsrc.org
rtrX> enable
```

or, if ssh is not configured yet:

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

The following configures the FastEthernet 0/0 interface to export flows. Replace 10.10.X.Y with the IP address of the PC in your pair which is going to receive them.

```
rtrX# configure terminal
rtrX(config)# interface FastEthernet 0/0
rtrX(config-if)# ip flow ingress
rtrX(config-if)# ip flow egress
rtrX(config-if)# exit
rtrX(config)# ip flow-export destination 10.10.X.Y 9001
rtrX(config)# ip flow-export destination 10.10.X.Z 9001
rtrX(config)# ip flow-export version 5
rtrX(config)# ip flow-cache timeout active 5
```

The last command 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.

```
rtrX(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:

```
rtrX(config)#ip flow-top-talkers
rtrX(config-flow-top-talkers)#top 20
rtrX(config-flow-top-talkers)#sort-by bytes
rtrX(config-flow-top-talkers)#end
```

Now we'll verify what we've done.

```
rtrX# show ip flow export
rtrX# 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

```
rtrX# show ip flow top-talkers
```

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

```
rtrX#wr mem
```

You can exit from the router now:

```
rtrX#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, 1359871921.013782000, #906334, 30 recs 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) [...]
```

These are the UDP packets containing individual flow records.

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.

You are done for this lab.

Go to `exercise2-install-nfdump-nfsen`.