NTop, NetFlow, and Cisco Routers
Last update March 2004
Jonathan Feldman jf@feldman.org
Entre Solutions www.entresolutions.com

Why Use NetFlow?

NetFlow is a standard way for a router to export statistics about *socket pairs* that it is routing. (Remembering that a socket pair represents <address1>:<port1> and <address2>:<port2>.) Such a pair, in NetFlow parlance, is called a "flow."

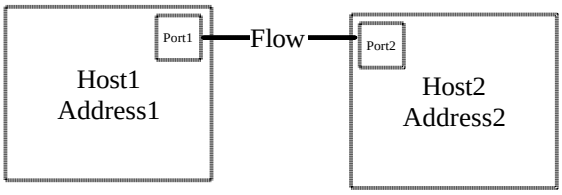

NTop, being majorly cool, has the ability to analyze such flows *without* router intervention – through the magick of sniffifying! So, why would you want the flows exported to your NTop server rather than collect them using your NTop server's promiscuous interface?

The answer is: in *certain* circumstances, it makes a lot of sense to let NTop do the gathering of flow stats instead of doing the sniffing. For example, here are some reasons where you might want to use NetFlow. (There are probably others.)

- You don't have a passive network tap (which fails "open") available for NTop to sniff data, tends to mean that, knowing that you only have a cheapo Ethernet hub to use, you decide that it's a bad idea for your $40,000 router to have a crappy $50 point of failure (there can also be problems related with certain "autosensing" hubs, which is out of scope for our discussion here)
- You are interested in wide area traffic and you don't have a (very expensive) WAN tap
- You don't care about OS fingerprinting or other swanky libpcap (sniff-like) features of NTop

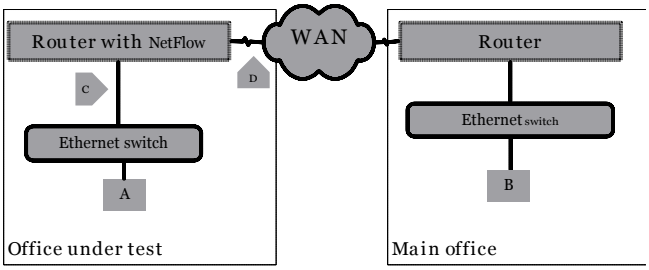

- You want physical access to the NTop box, but you want to collect data from elsewhere
- You are willing to put up with your traffic being *unclassified* if it is neither TCP nor UDP. NetFlow reports traffic that is neither TCP or UDP as 'port 0'. (That's a zero.) If you're looking to classify GRE, ESP, or other non TCP/UDP traffic, you're out of luck.
- You believe that the router itself is not malfunctioning or overloaded, so data from it may be trusted.

For example, in the WAN diagram, if you decide to use NTop with sniffing to collect your data, you must put an NTop (or NProbe) box at point "C" or point "D." You must also have a passive tap to do this (or a crappy hub if you are feeling cavalier). Both of these are at the remote office, which can be suboptimal for a variety of reasons.

However, if you decide to use NetFlow, you use point "A" or point "B" for your NTop collector – really, you can have the server located anywhere on the network that has IP reachability. Yow! Traffic characterization without taps! Yummy!

NetFlow, Cisco, and YOU!

Cisco routers are quite good at collecting the flow data without much of a performance penalty. It's not a big deal at all on WAN routers for sure. I have no data on NetFlow used at a network core (say, on a Cisco 6509.) I would really want to use Chariot or other "traffic stress" generation facility in a lab environment on big iron switches before I felt comfortable deploying NetFlow on such switches (unfortunately, I do not know of an open source equivalent to Chariot.) But – WAN links, you're in good shape. Monitoring CPU and memory on a T1 router with NetFlow shows hardly any difference in utilization.

NetFlow Support

The first and possibly most important question is: "does the IOS load I have on this router support CEF (Cisco Express Forwarding) and NetFlow." Really, the question is "does it support NetFlow," because NetFlow *requires* CEF – and therefore a load that supports NetFlow also supports CEF.

IMO, there's only one way to really *know* if your load supports NetFlow, and that is to try it. (For example, while 1721 routers with the IP load will support it, NetFlow didn't make it into the early releases – which are what are on most routers that have never been updated.)

The Software Advisor tool on Cisco's site can be of some help. With your TAC login in hand, you can check it out at http://www.cisco.com/cgi-bin/Support/CompNav/Index.pl. (If you don't have a TAC login, grab your SmartNet contract and call your Cisco rep, and get hooked up. If you don't have a SmartNet contract, you cannot legally upgrade your code, so stop here.)

If you find out that your load *doesn't* support NetFlow, and you want to upgrade, you want to check out the IOS planning tool at http://www.cisco.com/cgi-bin/Software/Iosplanner/Planner-tool/iosplanner.cgi.

Please take into consideration that IOS, like all other software, has bugs of both the benign and malignant variety. There are several fairly bad NetFlow bugs out there, and I highly recommend that you use the Bug Tool (http://www.cisco.com/cgi-bin/Support/Bugtool/launch_bugtool.pl) before you embark upon your NetFlow adventure.

Cisco Version Psychosis

Confused about all the darn versions that the various Tools are showing you? Well, you should be. ☺ There are about eleventy-seven bazillion Cisco IOS images out there, to fit each and every need under the sun. No worries, you *can* get familiar with the Cisco software release model. You will learn that a train doesn't always have tracks, among other incredible epiphanies. http://www.cisco.com/warp/public/620/1.pdf

Finally, Router Code!

Enabling NetFlow on the router in question is quite easy; just follow these simple steps and bake at 325 for however long it takes:

1. Make sure that CEF is enabled globally
2. Make sure that CEF is not disabled on the interfaces that you are interested in

3. Enable NetFlow globally, using the standard UDP port of 2055
4. Use version 5 unless you have a compelling reason not to
5. Enable NetFlow per interface that you are interested in. *This bears repeating—if you do not enable NetFlow across both interfaces that the traffic crosses, only half of the flow gets reported!*

Ok, so let's say you have "Router A" that needs to be configured to report flows to your NTop collector box at 10.0.0.1. All you need do is add the following to the router config:

```
#  ip flow-cache timeout active 5
#  ip cef
#  ip flow-export version 5
#  ip flow-export destination 10.0.0.1 2055
```

For *each* interface in question, you must enable NetFlow. For a router with interfaces eth0 and serial0, you might do this:

```
#  int eth0
#    ip route-cache flow
#  exit
#  int serial0
#    ip route-cache flow
#  exit
```

Oho, you ask, what about subinterfaces, like the ones you find with Frame Relay? Well, the short answer is "don't worry about them." Just configure the main interface. The long answer is, for more information (and if you are using anything more funky than Frame Relay like L2TP, you want more information, trust me) see Cisco's "Using NetFlow on Logical Interfaces" at http://www.cisco.com/warp/public/cc/pd/iosw/prodlit/ntflo_wp.pdf.

It's Not Working, Jonathan!

Well, yeah. Maybe. Here's what might be up:

- **NTop & router**: both have a propagation delay (well, NTop only has a propagation delay if you're looking at the trend data versus the live data – the live data should show up as soon as the router sends it). But in any case, chill out, dude. Go have a soda. If it's still not showing data in 5 minutes, keep troubleshooting.
- **NTop server**: Check to see whether UDP/2055 is listening (netstat –an | more). If it's not, you haven't started the plugin. Also, if the RECV-Q column of the netstat is showing a high number, it's likely that you need to restart NTop.
- **Router**: Check whether the router thinks it's seeing flows. But, also make sure that there *are* flows traversing the router, particularly if you are doing this at a quiet time! Generate some traffic that you know goes between the interfaces in question (ttcp if you are really geeky, or just go browse your favorite site if possible), then check what the cache looks like:

```
routerA#show ip  cache flow
IP packet size distribution (100917 total packets):
   1-32   64   96  128  160  192  224  256  288  320  352  384  416  448  480
   .000 .510 .005 .000 .001 .000 .001 .000 .000 .000 .000 .000 .003 .000 .000

    512  544  576 1024 1536 2048 2560 3072 3584 4096 4608
   .000 .000 .477 .000 .000 .000 .000 .000 .000 .000 .000

IP Flow Switching Cache, 278544 bytes
  4 active, 4092 inactive, 585 added
  10908 ager polls, 0 flow alloc failures
```

```
  Active flows timeout in 5 minutes
  Inactive flows timeout in 15 seconds
  last clearing of statistics never
Protocol         Total    Flows   Packets Bytes  Packets Active(Sec) Idle(Sec)
--------         Flows     /Sec     /Flow  /Pkt     /Sec     /Flow     /Flow
TCP-WWW              6       0.0      7825   308      1.6      25.1       1.2
TCP-other           49       0.0      1025   290      1.7      36.7      10.3
UDP-NTP            421       0.0         1    76      0.0       0.0      15.5
UDP-other           93       0.0         4   251      0.0       1.0      15.5
ICMP                12       0.0         2   119      0.0       0.2      15.5
Total:             581       0.0       168   297      3.4       3.5      14.9

SrcIf         SrcIPaddress    DstIf         DstIPaddress    Pr SrcP DstP  Pkts

SrcIf         SrcIPaddress    DstIf         DstIPaddress    Pr SrcP DstP  Pkts
Et0           10.0.24.150     Local         10.0.24.1       11 8006 00A1     4
Et0           10.0.24.150     Se1           192.168.0.2     06 815E 0050  1911
Se1           192.168.0.2     Et0           10.0.24.150     06 0050 815E  1904
Et0           10.0.24.150     Local         10.0.24.1       06 815B 0016     8
```

Long Term Trend Analysis of Flows

Once everything is working, you may crave historical data even between NTop restarts. Or, perhaps you might want a different time frame than NTop provides during a live run. To do this, see "RRD and NTop" at http://prdownloads.sourceforge.net/ntop/rrdandntop.pdf?download.

How do you harvest these flows and check out trends? It's pretty easy to write a script to do this. (Even I could do it and I make no claims about being a chi-chi programmer!) Remember, it's RRD, and RRD allows you to build whatever graphs you wish. One script that I wrote and that I also use ☺ may be found at: http://feldman.org/arc/000037.html. It produces layered graphs with different colors for each application type:

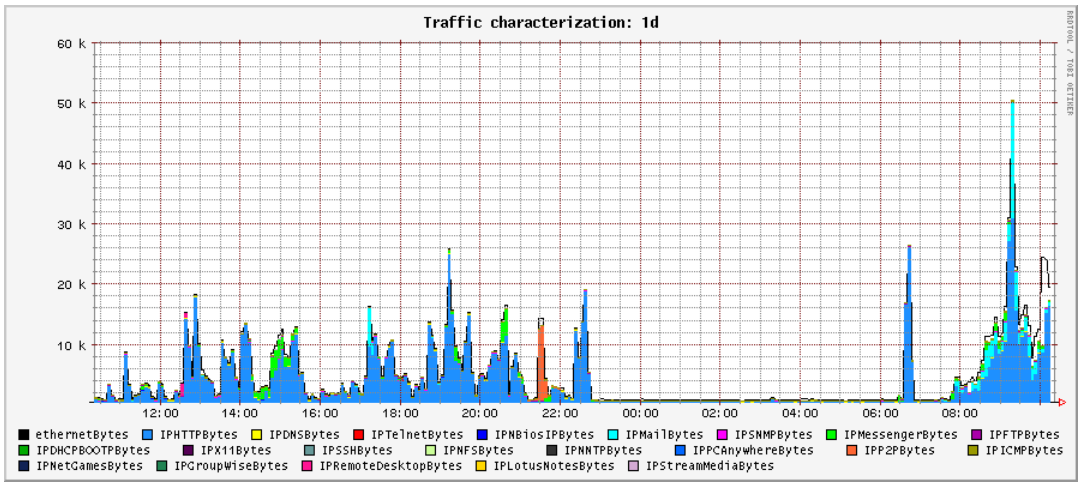

Have fun!

###