Finding the Bad Guys on MITnet

James Kretchmar
MIT Network Operations

MIT Security Camp
July 20, 2001

Overview

What our network looks like

The tools we use to manage problems:

- IP Accounting
- MRTG
- Netflow
- Neo

The focus will be on Netflow and Neo

MIT Backbone Topology

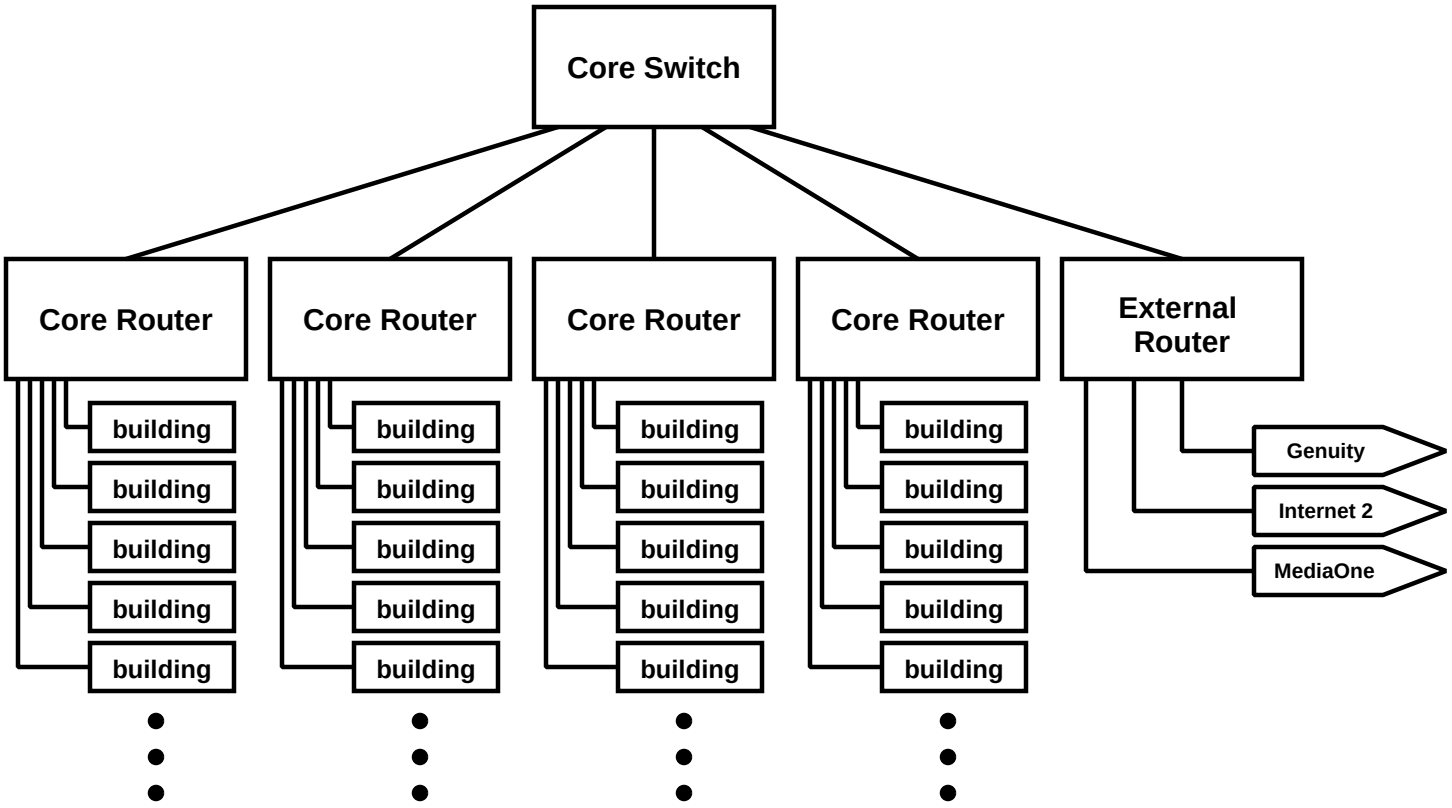

MIT Building Topology

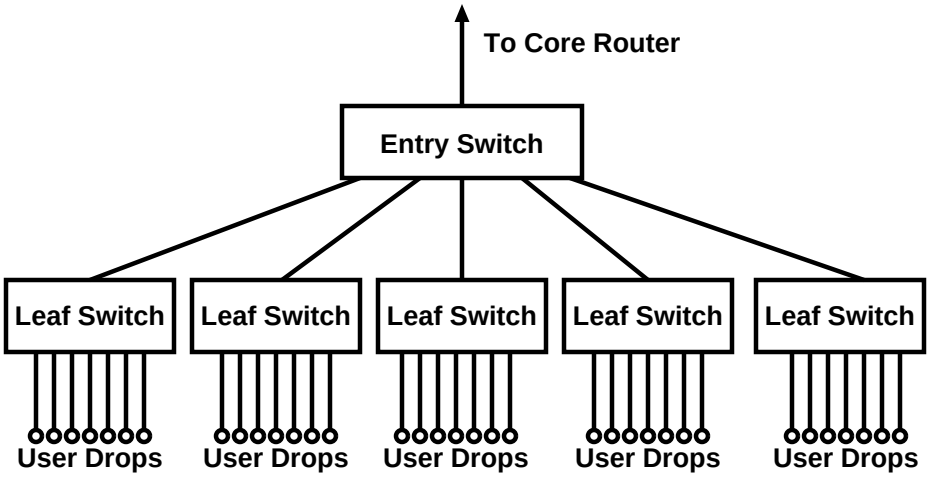

Key Network Design Points

The design is as consistent as possible

Every drop is managed

No firewalls

The exceptions are private networks

- LCS, AI, MediaLab, ILG T1s etc.
- Demarkation of service

So what's the problem?

Bad guys abound!

They break into machines anywhere on the network

From there they launch DOS attacks or other nastyness

Can affect:

- One machine
- An entire network (one building)
- All of MITnet

So how do we deal with these problems?

Tools!

Tools to find the problems:

- Is there a problem?
- Where is the problem?
- Who is the problem?

Tools to deal with the problems:

- Turn off a drop
- Block traffic

The Coolest Tools In Our Toolbox

IP Accounting

MRTG

Neo

Netflow

IP Accounting

For Cisco routers

Displays number of packets and bytes to src/dest IP pairs

For traffic out to the interface

IP Accounting Example

```
foo-rtr#show ip accounting
   Source            Destination              Packets                Bytes
 130.13.158.77     18.42.2.102                   41                 5602
 64.105.108.26     18.42.2.102                   27                 4246
 62.0.86.152       18.42.1.207                   72                 5491
 64.244.111.182    18.42.1.83                    12                  480
 140.192.41.4      18.42.1.83                    57                 2304
 24.42.251.131     18.42.2.102                    4                  326
 18.51.2.16        18.42.2.31                    36                 1656
 165.247.60.43     18.42.2.101                   15                 2328
 128.122.28.145    18.42.1.83                    15                  600
 24.190.250.41     18.42.1.83                    66                 2652
 61.153.19.98      18.42.0.204                   45                62968
 18.155.0.180      18.42.2.35                     2                   96
 18.100.0.50       18.42.1.89                     1                   84
 24.15.124.18      18.42.2.102                    5                  321
 18.155.0.181      18.42.2.35                     2                  117
 24.78.80.124      18.42.2.102                   30                 4583
 207.69.187.89     18.42.0.120                    2                   80
 212.41.206.69     18.42.2.102                   36                16724
 24.218.67.63      18.42.1.169                    4                  216
 205.188.136.217   18.42.0.40                     1                  693
 207.219.76.3      18.42.1.83                    18                  720
 212.170.161.231   18.42.0.53                     3                  120
```

More IP Accounting

Get selected inbound traffic with **"ip accounting-list"**

```
Enter configuration commands, one per line.  End with CNTL/Z.
foo-rtr(config)#no ip accounting-list
foo-rtr(config)#ip accounting-list 18.42.0.0 0.0.255.255
foo-rtr(config)#int f0/0
foo-rtr(config-if)#ip accounting
foo-rtr#show ip accounting
   Source           Destination              Packets               Bytes
 18.42.2.102      130.13.158.77                    9                3354
 130.13.158.77    18.42.2.102                      9                3485
 18.42.2.102      64.105.108.26                   15                2542
 64.105.108.26    18.42.2.102                     11                5938
 18.42.1.207      62.0.86.152                     15                3540
 62.0.86.152      18.42.1.207                     20                1647
 140.192.41.4     18.42.1.83                      22                 892
 18.42.1.83       140.192.41.4                    36               36484
 18.7.15.82       18.42.0.97                     164               65286
 18.42.0.97       18.7.15.82                     164               64629
 18.42.2.31       18.51.2.16                       6                 630
 18.51.2.16       18.42.2.31                       9                 432
 18.42.2.101      165.247.60.43                    5                 480
 165.247.60.43    18.42.2.101                      5                 728
 128.122.28.145   18.42.1.83                       1                  40
```

IP Accounting of an Attack

```
foo-rtr#show ip accounting
   Source            Destination              Packets              Bytes
 130.13.158.77     18.42.2.102                   41               5602
 64.105.108.26     18.42.2.102                   27               4246
 62.0.86.152       18.42.1.207                   72               5491
 64.244.111.182    18.42.1.83                    12                480
 140.192.41.4      18.42.1.83                    57               2304
 24.42.251.131     18.42.2.102                    4                326
 18.51.2.16        18.42.2.31                    36               1656
 165.247.60.43     18.42.2.101                   15               2328
 128.122.28.145    18.42.1.83                    15                600
 24.190.250.41     18.42.1.83                 566421              72652
 61.153.19.98      18.42.0.204                   45              62968
 18.155.0.180      18.42.2.35                     2                 96
 18.100.0.50       18.42.1.89                     1                 84
 24.15.124.18      18.42.2.102                    5                321
 18.155.0.181      18.42.2.35                     2                117
 24.78.80.124      18.42.2.102                   30               4583
 207.69.187.89     18.42.0.120                    2                 80
 212.41.206.69     18.42.2.102                   36              16724
 24.218.67.63      18.42.1.169                    4                216
 205.188.136.217   18.42.0.40                     1                693
 207.219.76.3      18.42.1.83                    18                720
 212.170.161.231   18.42.0.53                     3                120
```

Ups and Downs of IP Accounting

+ Very simple

+ Very effective for most problems

- Doesn't easily get all the traffic we might want

- Doesn't show us ports or protocols

But the **biggest** problem is ....

It's going away!

MRTG

The Multi Router Traffic Grapher **http://www.mrtg.org**

Great for looking at trends over a day:

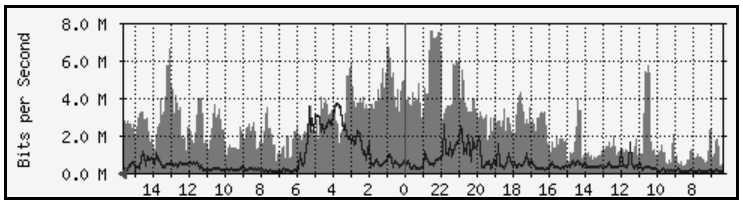

Over the week:

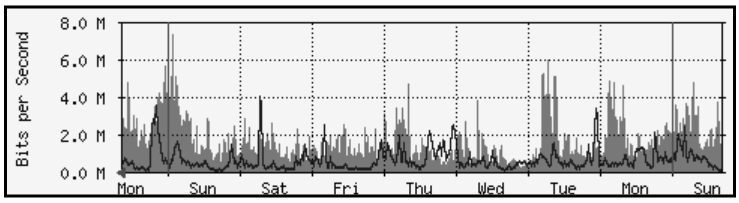

Or even over a year:

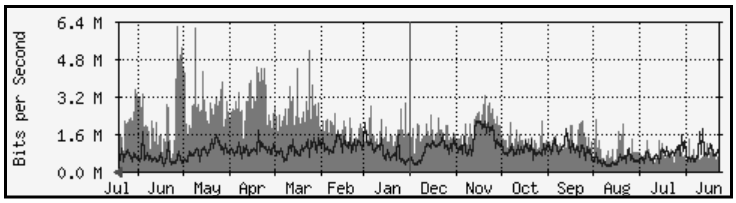

Ups and Downs of MRTG

+ It’s free

+ Makes it easy to see trends and notice problems

- Modifications can be messy

- End users mistake pictures for experience

Now, The Main Course

Netflow

- More detailed traffic monitoring

Neo

- Manipulate the Matrix at will

Netflow

It will give us detailed information on traffic

We can collect information to find problems

Also, it can improve network performance!

Netflow in action

On July 11th we saw a network blip that we couldn't explain

We checked records of high-traffic flows in the last 30 minutes:

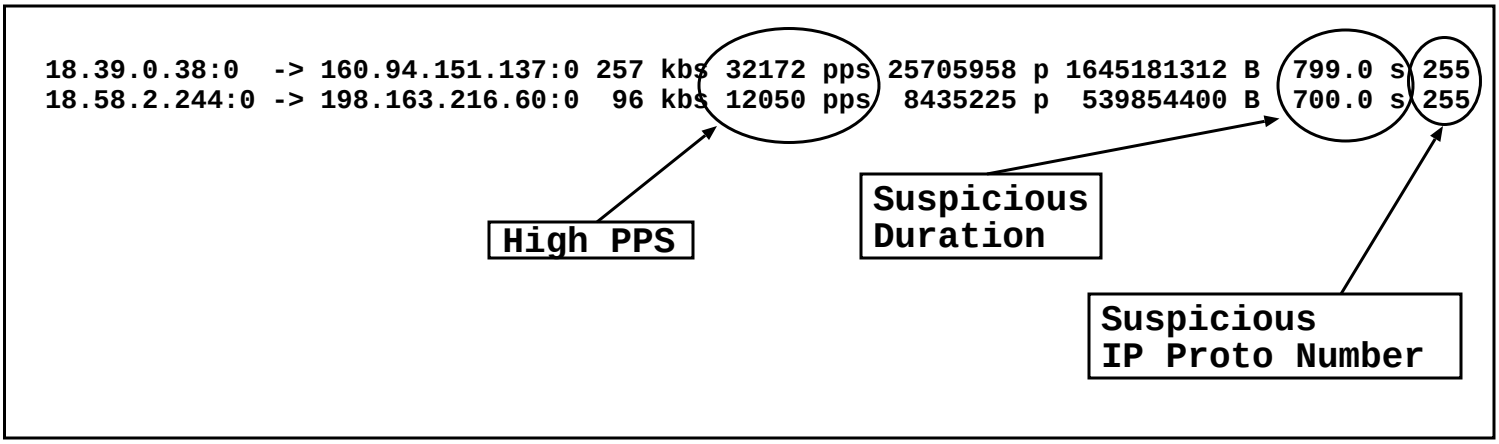

Logs from the previous 6 and 24 hours had similar entires

It's a pulsing zombie!

What else can we figure out?

Then we use netflow to monitor an affected host:

```
216.34.171.133:6667 -> 18.58.2.244:1613 0 kbs 0 pps 3 p 283 B 14.0 s TCP 18.168.0.17
```

Aha! A controlling IRC server

Finally we monitor connections to that address to get a list of compromised hosts

What is a flow?

Really it’s a routing optimization on Cisco routers

It’s expensive to go through the entire switching table
    Access-lists etc.

First packet sent normally, but then create a flow

A flow is a heuristic that corresponds to a real world flow of data
    A TCP connection is a flow
    UDP and other protocols are harder

Subsequent packets in the same flow take the same path

Examples of flows

A TCP connection to retrieve a web page

ICMP echo requests for the duration of the ping session

UDP packets for a TFTP download

So what hueristics define a flow?

At the beginning:

- Source addr/port
- Dest addr/port
- other?

At the end:

- TCP FIN
- 15 second inactivity timeout
- 30 minute active timeout

Remind me why we care?

The router can report on flows

Flows in progress can be seen on the router

Expired flow data sent out in a UDP packet

No more IP Accounting!

Example flows on the router

```
foo-rtr#show ip cache flow
SrcIf     SrcIPaddress    DstIf     DstIPaddress    Pr SrcP DstP Pkts B/Pk Active
Fd0/0     18.72.0.3       Et5/5     18.238.1.168    11 0035 103D    1  182    0.0
Fd0/0     18.32.0.110     Fa12/1    18.142.2.75     06 008B 0408    2   79    5.0
Fd0/0     18.181.0.23     Et4/3     18.56.0.60      11 1B5A 0AEC    6   73   11.2
Fd0/0     63.88.62.25     Et5/4     18.248.2.153    06 18CA 0FF6   83  153   19.4
Fd0/0     18.54.0.156     Et3/2     18.186.0.190    11 937F 82CF    1   39    0.0
Fd0/0     209.7.114.251   Se10/4    18.36.11.90     06 BCE0 0050    2   48    3.2
Fd0/0     217.224.219.202 Et5/4     18.248.2.56     06 0C16 18CA  214  484   60.4
Fd0/0     207.46.204.99   Se11/0/0  18.222.0.213    11 2329 2328  175   52  350.7
Fd0/0     207.46.204.104  Se11/0/0  18.222.0.213    11 2329 2328  769  149  238.6
Fd0/0     18.72.0.3       Et8/2     18.42.1.140     06 0025 0791    4   40    1.4
Fd0/0     18.72.0.3       Et8/2     18.42.1.140     06 0025 0793    4   40    1.4
Fd0/0     18.72.0.3       Et8/2     18.42.1.140     06 0025 0792    4   40    1.4
Fd0/0     213.132.197.200 Se11/0/0  18.224.0.124    06 0050 04E3    8 1098    0.4
Fd0/0     18.54.0.156     Et4/3     18.56.0.68      11 937B 82CF    1   39    0.0
Fd0/0     213.132.197.200 Se11/0/0  18.224.0.124    06 0050 04E4   11  868    0.5
Fd0/0     18.54.0.156     Et4/3     18.56.0.68      11 937D 82CF    1   39    0.0
Fd0/0     152.163.159.232 Et4/3     18.56.0.60      11 0035 0400    1  183    0.0
Fd0/0     18.181.0.31     Et5/5     18.238.2.8      06 0050 048C    4  210    8.1
```

Flows on the router

It’s an inhospitable place to look at flows

One trick is:

```
foo-rtr>show ip cache flow | include K
Gi0/0/0       18.24.10.20     PO10/0/0      216.27.195.241  06 0014 0408    20K
Gi0/0/0       18.157.5.58     PO10/0/0      24.40.29.11     11 6D38 6D38    14K
AT8/0.420     128.112.129.150 Gi0/0/0       18.75.3.101     11 1E6C 1B3A    29K
Gi0/0/0       18.24.10.28     AT8/0.420     192.26.210.165  06 0970 0077    10K
Gi0/0/0       18.241.2.224    PO10/0/0      198.234.217.139 06 0E12 0281    53K
Gi0/0/0       18.42.2.102     PO10/0/0      24.182.27.109   06 18D3 0BD5    11K
Gi0/0/0       18.42.2.102     PO10/0/0      24.45.23.7      06 18D3 06C9    18K
Gi0/0/0       18.63.2.247     PO10/0/0      212.155.223.3   06 1236 077D    38K
Gi0/0/0       18.170.2.161    PO3/0/0       66.31.90.53     06 0014 0577    22K
AT8/0.420     192.26.210.165  Gi0/0/0       18.24.10.28     06 0077 0717    99K
Gi0/0/0       18.24.10.28     AT8/0.420     192.26.210.165  06 0717 0077   179K
Gi0/0/0       18.24.10.28     AT8/0.420     128.223.220.30  06 0077 4C91    21K
AT8/0.420     128.223.220.30  Gi0/0/0       18.24.10.28     06 4C91 0077    44K
Gi0/0/0       18.244.0.188    PO10/0/0      63.137.60.100   06 0961 0E7E    18K
Gi0/0/0       18.238.2.151    PO10/0/0      203.164.216.156 06 157D 040C   123K
Gi0/0/0       18.181.0.26     AT8/0.420     128.169.76.138  06 CDCA 0077    73K
AT8/0.420     128.169.76.138  Gi0/0/0       18.181.0.26     06 0077 CDC5    37K
Gi0/0/0       18.250.1.160    PO10/0/0      195.13.204.67   06 0B32 0411    10K
Gi0/0/0       18.181.0.26     AT8/0.420     128.169.76.138  06 CDCC 0077    73K
AT8/0.420     128.169.76.138  Gi0/0/0       18.181.0.26     06 0077 CDCA    34K
```

Flows off the router

The UDP message includes:

- Source address/port
- Destination address/port
- Number of bytes, packets
- Length in seconds
- IP Prococol
- and more ...

```
18.21.0.103:0     -> 18.158.0.56:0      127 kbs  718 pps   718 p  1021542 B  1.0 s UDP 18.168.0.16
18.7.20.69:http   -> 18.56.1.34:49285   121 kbs  650 pps  1300 p  1944018 B  2.0 s TCP 18.168.0.15
18.7.21.72:1109   -> 18.184.0.39:42665  111 kbs  650 pps   650 p   891673 B  1.0 s TCP 18.168.0.11
18.155.0.236:1022 -> 18.159.0.37:515    111 kbs  840 pps   840 p   889785 B  1.0 s TCP 18.168.0.11
18.19.0.44:850    -> 18.159.0.27:515    110 kbs  599 pps   599 p   881868 B  1.0 s TCP 18.168.0.11
18.56.0.175:0     -> 18.179.0.43:0      109 kbs  619 pps   619 p   879176 B  1.0 s UDP 18.168.0.16
18.155.0.236:1023 -> 18.159.0.37:515    105 kbs  802 pps   802 p   846778 B  1.0 s TCP 18.168.0.11
18.7.21.71:1109   -> 18.58.0.43:1662    104 kbs  595 pps  2382 p  3332968 B  4.0 s TCP 18.168.0.17
18.58.1.47:0      -> 18.158.0.47:0       90 kbs  508 pps   508 p   722978 B  1.0 s UDP 18.168.0.16
```

A closer look at the exported flow

It’s a UDP datagram in a single packet

Contains multiple flows

Starts with a header

Header followed by flow data

Exported Flows, Cont.

Header contains a version number

Version number dictates packet format

Four versions currently in use (1,5,7,8)

Different devices support different versions

Different versions contain different info

Version 1

Source IP
Destination IP
IP of next hop router
Input interface number
Output interface number
Packets in the flow
Layer 3 bytes in the flow
SysUptime at start of flow
SysUptime at end of flow
Source port
Destination port
IP protocol number
TOS flags
TCP flags

Who could ask for anything more?

Other Versions

Version 5
Adds source and destionation AS numbers

Version 7
For the Cat5000

Version 8
Supports flow aggregation in several packet formats

Where do we use flows?

On the backbone interface on core routers

Grabs any traffic between routers

What about traffic between interfaces on a router?

Would have to enable flow caching on all interfaces

- Flow cashing everywhere is ok
- Exporting all the flows probably is not

When the need arises it can be enabled on specific interfaces

Configuring NetFlow

Turn it on on the interface:

```
foo-rtr#config t
Enter configuration commands, one per line.  End with CNTL/Z.
foo-rtr(config)#int f 0/0
foo-rtr(config-if)# ip route-cache flow
foo-rtr(config-if)# no ip route-cache optimum
```

Turn on exporting:

```
foo-rtr#config t
Enter configuration commands, one per line.  End with CNTL/Z.
foo-rtr(config)# ip flow-export destination 18.7.21.88 9995
foo-rtr(config)# ip flow-export source Fddi0/0
```

Later versions let you also specify some other options

What software is available?

Cisco will sell you a collector and an analyzer

Two free pieces of software from **http://web.mit.edu/ktools**

- Flowtee
- Flowmon

And more later

Flowtee

The router only lets you send flows to one host

Flowtee receives flows and sends them on to other hosts

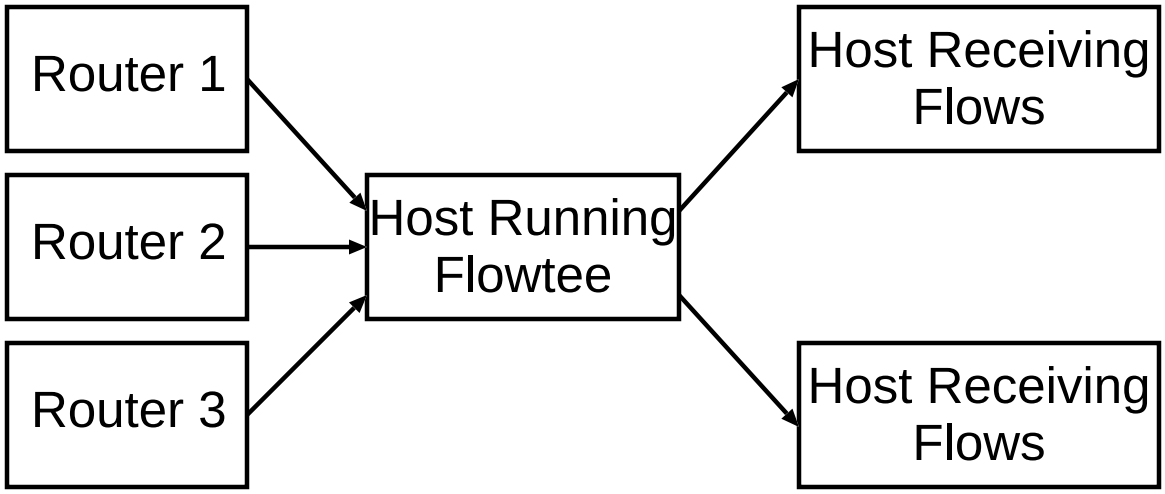

Source addresses are forged from original router

Future Flowtee Work

Flag to disable source forging **(in a few days)**

A service for remote reconfiguration **(in a few months)**

So you can quickly and easily redirect flows

It may just become a generic UDP port forwarder

(It very nearly already is)

Flowmon

A flow monitoring package

"flowmon" stores highest talkers

"flowprint" prints all flows to the screen

Useful for real time debugging

Future Flowmon Work

Detect port scanning

Detect other interesting attack patterns?

Better formatting

Ups and Downs of Netflow

+ It gives you lots of useful data

+ Without IP accounting it’s your only hope

- Lousy interface on the router

- Flows are exported inconsistently (this may get better)

- You have to give more $$$ to Cisco to turn it on

+ But you can collect the data with free software

Neo

Manages switches and repeaters and routers. Oh my!

Uses SNMP

It’s a generic framework for interacting with network devices

Designed to support many different kinds of hardware

Modular design makes adding new hardware pretty easy

Neo Location Syntax

port @ device	**3@w92-165t-sw-1**
board **/** port @ device	**2/14@w92-entry**
board **/** @ device	**2/@w92-entry**
allports @ device	***@w92-165t-sw-1**
board **/** allports @ device	**2/*@w92-165t-sw-1**
allboards **/** @ device	***/@w92-entry**

Device Summary

```
neo: device summary w92-165t-sw-1
Port summary:
 p type           u   lnk adm ap
------------------------------
 1 100TX              -  On
 2 100TX              -  On
 3 100TX              -  On
 4 100TX              -  On
 5 100TX              -  On
 6 100TX              -  On
 7 100TX              -  On
 8 100TX              -  On
 9 100TX            100  On
10 100TX            100  On
11 100TX              -  On
12 100TX              -  On
13 100TX              -  On
14 100TX              -  On
15 100TX            100  On
16 100TX              -  On
17 100TX              -  On
18 100TX              -  On
19 100TX              -  On
20 100TX              -  On
21 100TX              -  On
22 100TX            100  On
23 100TX              -  On
24 100TX              -  On
25 100?X          *   100  On
26 100?X          *     -  On
27 loop              10  On
```

Device Summary for Device with Boards

```
neo: dev sum w7-entry
Board summary:
 b type            sg   ul ultype
-------------------------------------
 1 Mgmt-3           0
 2 12 100FX         0
 3 12 100FX         0
 4 12 100FX         0
 5 12 100FX         0

neo: dev sum 2/*@w7-entry
Port summary:
 p type           u  lnk adm ap
------------------------------
 1 100F(mm)         100  On
 2 100F(mm)         100  On
 3 100F(mm)         100  On
 4 100F(mm)         100  On
 5 100F(mm)         100  On
 6 100F(mm)         100  On
 7 100F(mm)         100  On
 8 100F(mm)         100  On
 9 100F(mm)         100  On
10 100F(mm)         100  On
11 100F(mm)         100  On
12 100F(mm)         100  On
```

Locating a Host

```
neo: arpfind plugs-out nw12-rtr-bb
18.18.0.1 says 18.18.1.101 is 08:00:20:7D:48:1B

neo: locate 08:00:20:7D:48:1B @k:18.18
Found on 2/14@W92-165T-SW-ENTRY
Found on 3@W92-165T-SW-13
```

The Keyfile

A very simple database of hosts

Maps a string to a group of hosts

```
neo: location print @k:18.18
Devices (32) are:
  W92-264T-SW-7
  W92-264T-SW-8
  W92-165T-SW-16
  W92-170-AP-1
  W92-126-AP-1
  W92-149-AP-1
  W92-199-AP-1
  W92-165T-SW-1
  W92-165T-SW-ENTRY
  W92-165T-SW-2
  W92-165T-SW-3
  etc.
```

Stats

```
neo: stats 2/*@w92-entry
Probing devices ...
Getting device summary ...
Getting first set of stats...
Getting second set of stats...
Port statistics:
 p type          u  lnk adm ap   kbs   ikbs   okbs   pps  ipps  opps ierps oerps
-------------------------------------------------------------------------------
 1 10/100T       100  On      655    530    125    171   107    64     0     0
 2 10/100T       100  On      186     53    133     97    37    60     0     0
 3 10/100T       100  On      387     24    363    123    27    96     0     0
 4 10/100T       100  On       84      9     75     32     5    27     0     0
 5 10/100T       100  On       49      3     46     28     5    23     0     0
 6 10/100T       100  On      108     43     65    109    46    63     0     0
 7 10/100T       100  On     1035    432    603    293   139   154     0     0
 8 10/100T       100  On       20      0     20     15     0    15     0     0
 9 10/100T       100  On       30      4     26     25     5    20     0     0
10 10/100T       100  On      159      3    156     37     5    32     0     0
11 10/100T       100  On       27      4     23     29     7    22     0     0
12 10/100T       100  On       62      4     58     26     4    22     0     0
13 10/100T       100  On      111      9    102     42    10    32     0     0
14 10/100T       100  On      802     29    773    160    35   125     0     0
15 10/100T       100  On      183      6    177     39    10    29     0     0
16 10/100T       100  On       17      0     17     15     0    15     0     0
17 10/100T       100  On      360     29    331     89    28    61     0     0
18 10/100T       100  On     8575   8366    209   1178   787   391     0     0
19 10/100T       100  On       19      2     17     19     3    16     0     0
20 10/100T       100  On       72     52     20     26     7    19     0     0
21 10/100T       100  On     6789   6584    205    956   613   343     0     0
22 10/100T       100  On       27     10     17     28     8    20     0     0
23 10/100T       100  On      278    189     89     65    28    37     0     0
24 10/100T       100  On       42     21     21     43    15    28     0     0
```

Port Disabling

```
neo: set writecom
Write community:
neo: port disable 3@w92-165t-sw-1
3@w92-165t-sw-1 disabled

neo: port status 3@w92-165t-sw-1
3@w92-165t-sw-1 disabled

neo: dev sum *@w92-165t-sw-1
Port summary:
 p type          u  lnk adm ap
------------------------------
 1 100TX           -  On
 2 100TX           -  On
 3 100TX           - Off
 4 100TX           -  On
 5 100TX           -  On
 6 100TX           -  On
 7 100TX           -  On
 8 100TX           -  On
 9 100TX         100  On
 (etc)

neo: port enable 3@w92-165t-sw-1
3@w92-165t-sw-1 enabled
```

Port Search

```
neo: port search 3@W92-165T-SW-13
08:00:20:7D:48:1B

neo: port search 2/14@W92-165T-SW-ENTRY
00:00:1D:FA:D3:B2
00:00:1D:FA:D3:CA
00:03:BA:04:AE:F0
00:C0:4F:A3:21:59
08:00:20:76:F9:65
08:00:20:7D:48:1B
08:00:20:A0:3A:A6
08:00:20:E7:7F:34
08:00:46:0C:96:0C
08:00:69:0A:DE:A1
```

Bells and Whistles

```
neo: dev info w92-165t-sw-1
w92-165t-sw-1
  Device type: C2200
  Contact    : network@mit.edu
  Name       : w92-165t-sw-1
  Location   : w92-165T
  Uptime     : 39 days 23:46:25
  ObjectID   : .1.3.6.1.4.1.52.3.9.3.4.84
  Descr      : Cabletron Systems, Inc. 2H253-25R Rev 04.00.08 ...

neo: dev info weather 64.106.2.1
Temp Alarm:    No Alarm
Temperature:   31 C
Fan 1:         Running
Fan 2:         Running
Fan 3:         Running
Fan 4:         Running
Fan 5:         Running

neo: dev info power 64.106.2.1
Primary Power:     Running
Redundant Power:   Installed and running
Power Alarm:       No Alarm
```

Where do I get neo?

Also from **http://web.mit.edu/ktools**

Putting it All Together

We frequently use more than one tool at a time

Example:

MRTG to the find the problem network
IP Accounting to find the problem IP address
Neo to locate the host and disable the drop

Or:

Neo to find high traffic port and IP address
Netflow to monitor traffic
Netflow to search for similarly compromised hosts

The End

Neo, Flowmon, Flowtee and more:

http://web.mit.edu/ktools

MRTG:

http://www.mrtg.org

Me:

James Kretchmar
jk@mit.edu