

Using NetFlow Data

AcIS Network Systems
Columbia University

Daniel Medina
medina@columbia.edu

Overview

Identifying DoS Attacks

Identifying Traffic Trends

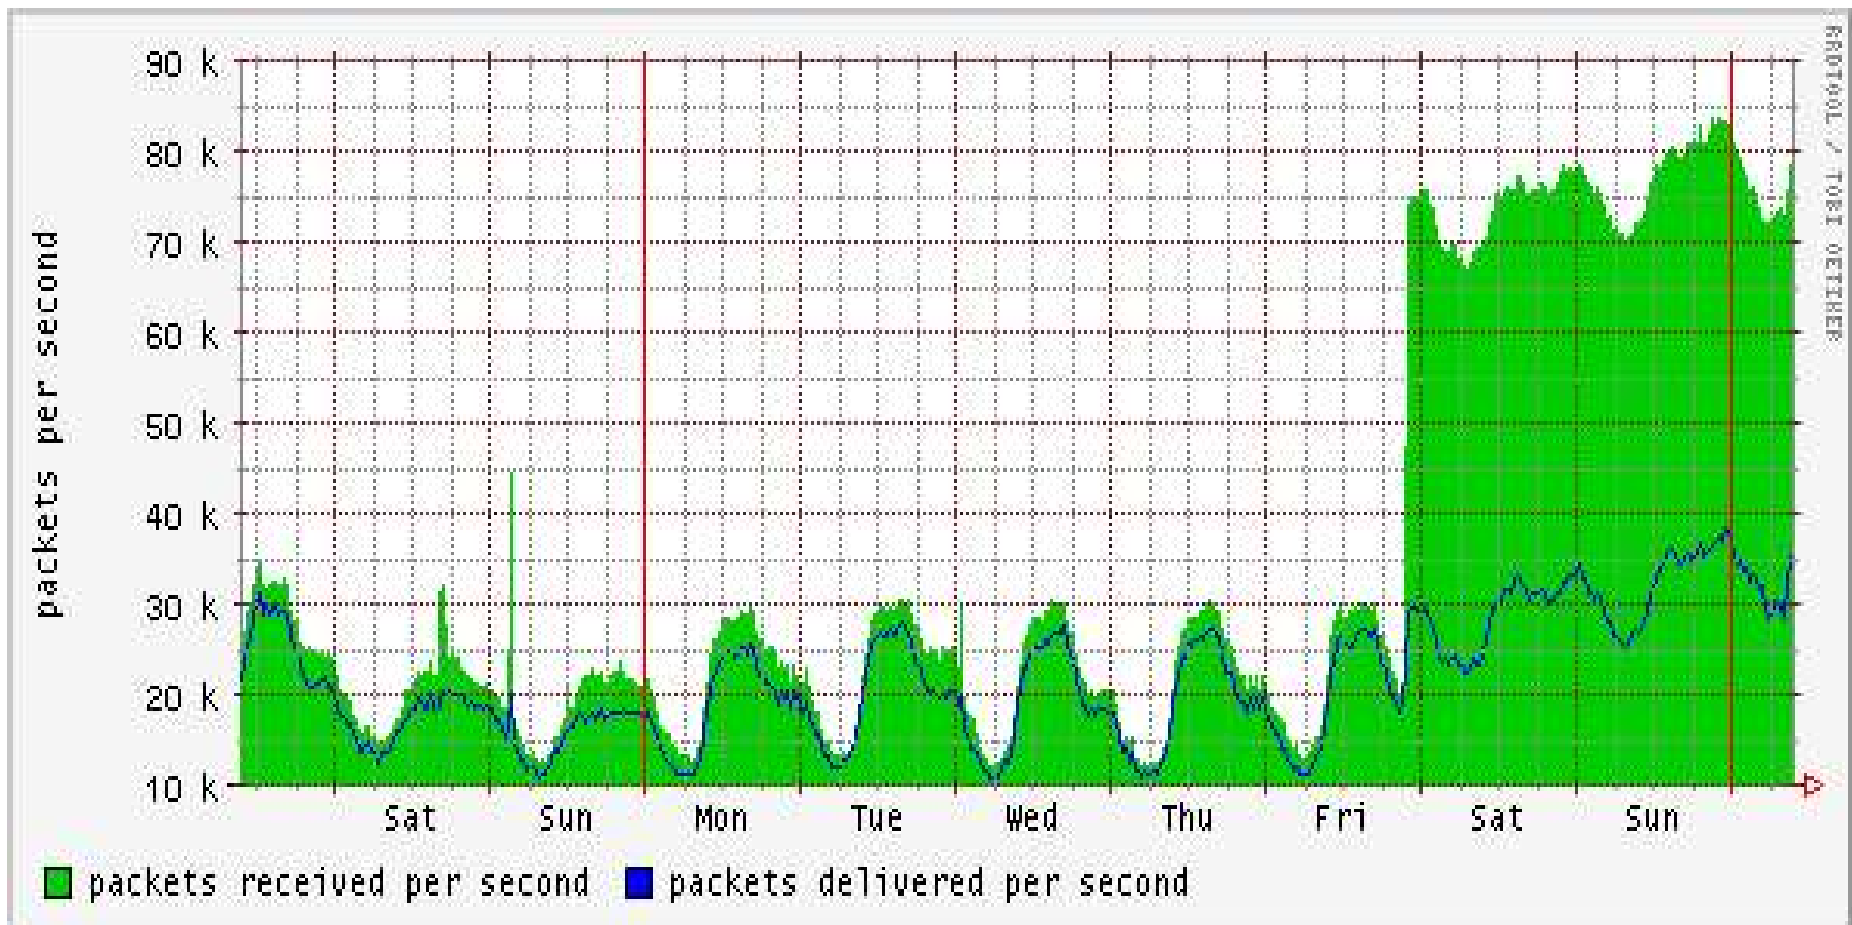
Identifying Infected Hosts

Network Bandwidth Quotas

Future Development

Identifying DoS Attacks

Late on Friday night, March 19th, Columbia was the target of a Denial-of-Service attack



Identifying DoS Attacks

Taking a peek at the router's NetFlow data...

```
$ flowdumper -se '$exporterip eq "128.59.0.4"' \  
> ft-v05.2004-03-22.013939-0500  
  
2004/03/22 01:39:21 68.94.5.31.2197 -> 156.111.252.210.135 6(SYN) 1 48  
2004/03/22 01:39:21 128.59.90.116.2869 -> 166.75.254.28.445 6(SYN) 1 48  
2004/03/22 01:39:21 209.104.48.3 -> 67.99.58.194 ICMP_ECHO 1 84  
2004/03/22 01:39:21 209.246.136.79 -> 67.99.58.194 ICMP_ECHO 1 64  
2004/03/22 01:39:20 221.216.65.139.4742 -> 128.59.1.200.3127 6(SYN) 3 144  
2004/03/22 01:39:21 63.252.188.119.192 -> 128.59.1.3.22 6(SYN) 2 80  
2004/03/22 01:39:21 63.252.188.119.192 -> 128.59.1.100.22 6(SYN) 2 80  
2004/03/22 01:39:21 128.59.77.27.53008 -> 239.255.255.253.427 17 2 182  
2004/03/22 01:39:21 63.252.188.119.192 -> 128.59.1.68.22 6(SYN) 2 80  
2004/03/22 01:39:22 209.104.46.3 -> 67.99.58.194 ICMP_ECHO 1 84  
2004/03/22 01:39:22 212.177.145.69 -> 67.99.58.194 ICMP_ECHO 1 64  
2004/03/22 01:39:22 160.39.248.92.2313 -> 224.0.1.76.2313 17 2 239  
2004/03/22 01:39:22 193.45.14.132 -> 67.99.58.194 ICMP_ECHO 1 64  
2004/03/22 01:39:21 63.252.188.119.5 -> 128.59.1.225.22 6(SYN) 2 80  
2004/03/22 01:39:21 63.252.188.119.20 -> 128.59.1.227.22 6(SYN) 2 80  
^C
```

Identifying DoS Attacks

```
$ flowdumper -se '$dstport == 22' \  
> ft-v05.2004-03-22.013939-0500
```

```
2004/03/22 01:39:35 63.252.188.119.0 -> 128.59.210.116.22 6(SYN) 26 1040  
2004/03/22 01:39:38 63.252.188.119.0 -> 128.59.210.125.22 6(SYN) 40 1600  
2004/03/22 01:39:38 63.252.188.119.0 -> 128.59.210.157.22 6(SYN) 42 1680  
2004/03/22 01:39:35 63.252.188.119.0 -> 128.59.210.179.22 6(SYN) 23 920  
2004/03/22 01:39:38 63.252.188.119.0 -> 128.59.210.192.22 6(SYN) 34 1360  
2004/03/22 01:39:35 63.252.188.119.0 -> 128.59.210.188.22 6(SYN) 34 1360  
2004/03/22 01:39:38 63.252.188.119.0 -> 128.59.210.191.22 6(SYN) 36 1440  
2004/03/22 01:39:38 63.252.188.119.0 -> 128.59.210.225.22 6(SYN) 35 1400  
2004/03/22 01:39:38 63.252.188.119.0 -> 128.59.210.255.22 6(SYN) 39 1560  
2004/03/22 01:39:38 63.252.188.119.0 -> 128.59.210.114.22 6(SYN) 20 800  
2004/03/22 01:39:36 63.252.188.119.0 -> 128.59.46.33.22 6(SYN) 30 1200  
2004/03/22 01:39:36 63.252.188.119.0 -> 128.59.46.96.22 6(SYN) 28 1120  
2004/03/22 01:39:36 63.252.188.119.0 -> 128.59.46.95.22 6(SYN) 28 1120  
2004/03/22 01:39:36 63.252.188.119.0 -> 128.59.46.131.22 6(SYN) 29 1160  
2004/03/22 01:39:36 63.252.188.119.0 -> 128.59.46.136.22 6(SYN) 27 1080
```

```
^C
```

Identifying DoS Attacks

```
$ toptendst.pl -f ft-v05.2004-03-22.013939-0500
```

```
.....  
Processed 3225921 flows
```

```
Between Mon Mar 22 01:39:36 2004 and Mon Mar 22 01:44:37 2004
```

63.252.188.119	1679519	please.saveipv4.org
131.109.145.6	35545	(no hostname)
193.205.105.120	31405	(no hostname)

```
^C
```

Identifying Traffic Trends

New beasties constantly crop up on the network, most recently Welchia (Worm) and BitTorrent (P2P file-sharing)

Unusual traffic must be monitored so we know where our bandwidth is going

Welchia traffic – 92-byte ICMP packets

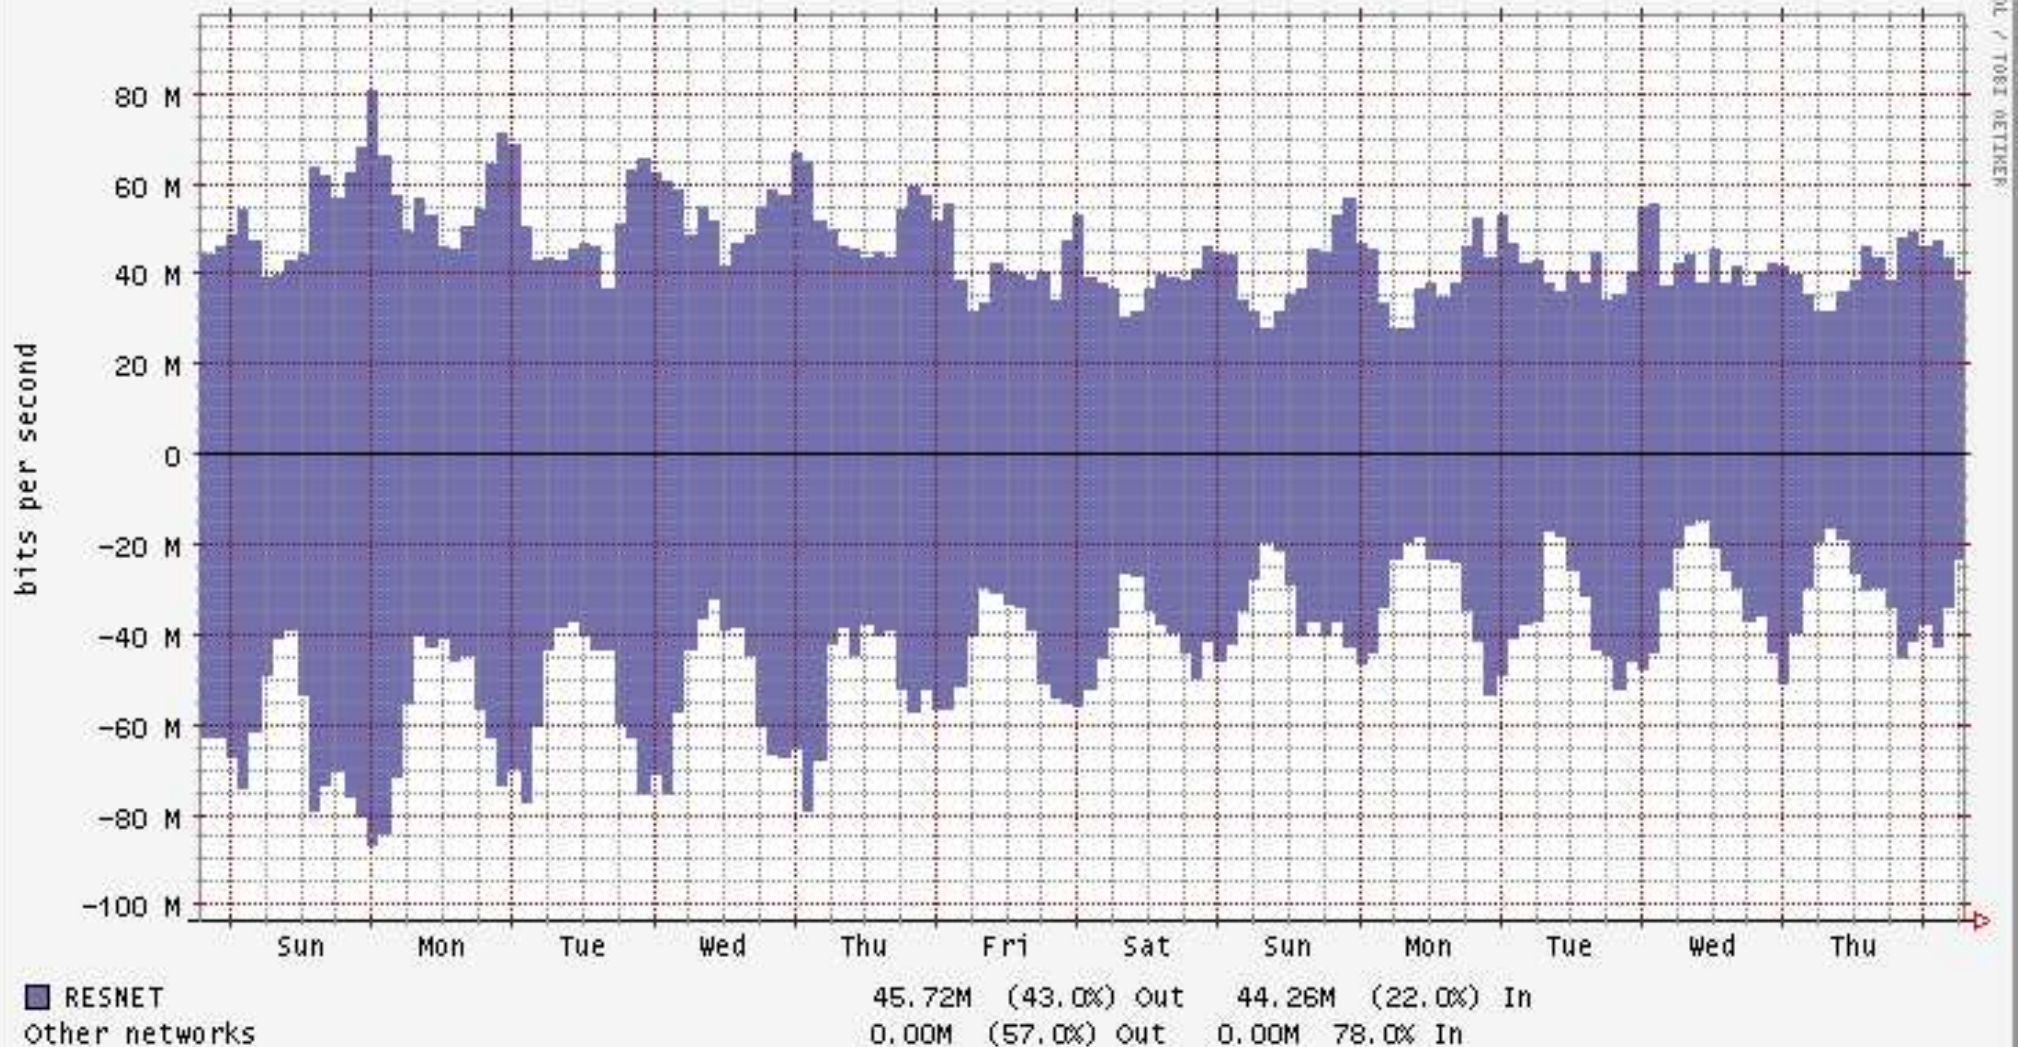
This worm-generated traffic consumed about 4 Mb/s

BitTorrent traffic – 6881-6889,6969/tcp

Consumed 10+ Mb/s, before being rate-limited

Identifying Traffic Trends

Estimated Columbia University Campus Well Known Protocols/Services, Bits, +out/-in



Identifying Infected Hosts

Currently, we monitor CERT warnings and Anti-Virus Vendor alerts

Simple signatures are compiled and hosts flagged

```
# A Signature 3-tuple
# protocol, port, bytes
$SIG {6} {6667} {-1} = "IRC";
$SIG {17} {1434} {404} = "SQL Slammer";
$SIG {6} {445} {-1} = "445/tcp";
```

Identifying Infected Hosts

Incident page <http://netflow.cc.columbia.edu/incidents.txt>

Src Addr	Incident	Count	Src MAC	Hostname
2004-03-25 11:27:37				
128.59.86.172	6667/tcp	5	0003BA10AF44	tamil.cisl.columbia.edu
128.59.86.172	445/tcp	1581	0003BA10AF44	tamil.cisl.columbia.edu

NetFlow host

```
$ flowdumper -se '$srcip eq "128.59.86.172" && $dstport == 6667' \  
> ft-v05.2004-03-25.113235-0500  
2004/03/25 11:32:22 128.59.86.172.4218 -> 65.201.175.144.6667 6 3 144  
2004/03/25 11:35:00 128.59.86.172.4830 -> 10.10.10.10.6667 6 7 336
```

```
$ flowdumper -se '$dstip eq "65.201.175.144"' \  
> ft-v05.2004-03-25.113235-0500  
2004/03/25 11:32:22 128.59.86.172.4218 -> 65.201.175.144.6667 6 3 144  
2004/03/25 11:33:17 128.59.86.152.4756 -> 65.201.175.144.6667 6 3 144
```

Network Bandwidth Quotas

We track outbound commodity bytes per hour per IP address (packaged as FlowMonitor)

Simple:

```
while ( <FLOWS> ) {
    $bytes_used{$ip_addr} += $bytes
    if &CUhost{$ip_addr};
}

foreach $ip_addr (keys %counter) {
    print "$ip_addr\n" if
        $bytes_used{$ip_addr} > $quota;
}
```

Future Development

Package some useful scripts up

```
$ topports.pl ft-v05.2004-03-25.120728-0500
```

```
.....
```

```
Between Thu Mar 25 12:07:24 2004 and Thu Mar 25 12:12:26 2004
```

Port/Protocol	Bytes	Percent of Total
80/6	556,474,096	(4.8 %)
61552/17	471,420,382	(4.1 %)
6881/6	429,949,029	(3.7 %)
6882/6	325,036,287	(2.8 %)
1947/6	256,937,604	(2.2 %)
4662/6	242,402,761	(2.1 %)
6714/6	177,608,976	(1.5 %)
119/6	175,771,917	(1.5 %)
6883/6	174,254,622	(1.5 %)

Future Development

```
$ countips.pl ft-v05.2004-03-25.123223-0500
  Total          13508
    tc           430
  lamont        298
    news         1
  resnet       3529
shared-buslaw   95
  business     380
    akamai       6
    cpmc        3285
    law          223
  wireless     234
  barnard     1129
```

Future Development

Darkspace monitoring

```
$ darkspace.sh ft-v05.2004-03-25.123223-0500
2004/03/25 12:33:42 69.140.112.58.3711 -> 160.39.4.24.3127 6 1 48
2004/03/25 12:33:44 63.199.201.86.1116 -> 160.39.17.158.3127 6 2 96
2004/03/25 12:33:48 69.11.207.54.2298 -> 160.39.5.207.3127 6 2 96
2004/03/25 12:33:51 200.164.6.68.2011 -> 160.39.0.236.3127 6 2 96
2004/03/25 12:33:50 81.79.49.211.3516 -> 160.39.30.127.3127 6 3 144
2004/03/25 12:33:42 195.132.163.224.1761 -> 160.39.22.196.3127 6 2 96
2004/03/25 12:33:46 62.99.91.215.2262 -> 160.39.5.223.3127 6 1 48
2004/03/25 12:33:41 63.199.201.86.4635 -> 160.39.17.157.3127 6 3 144
2004/03/25 12:33:54 24.6.208.67.1270 -> 160.39.14.178.3127 6 3 144
2004/03/25 12:33:41 82.48.233.181.2012 -> 160.39.26.105.3127 6 3 144
2004/03/25 12:33:50 200.216.43.90.4227 -> 160.39.22.22.3127 6 2 96
2004/03/25 12:33:52 81.57.29.31.4301 -> 160.39.19.7.3127 6 1 48
2004/03/25 12:33:52 24.7.195.47.2337 -> 160.39.4.88.3127 6 2 96
...
2004/03/25 12:34:08 63.199.201.86.1661 -> 160.39.17.160.3127 6 3 144
^C
```

Future Development

Shows off Dave Plonka's flowdumper

```
$ cat darkspace.sh
#!/bin/bash

flowdumper \
  -I 'use Net::Patricia;
      $pt = Net::Patricia->new;
      map { $pt->add_string($_, 1) } qw( 160.39.0.0/19
                                          160.39.80.0/20 )' \
  -s \
  -e
  '$pt->match_integer($dstaddr) ||
  $pt->match_integer($srcaddr)' $@
```

man flowdumper **and** man Cflow for more...

Future Development

Work by Matt Selsky and Johan Andersen:

Distributed processing hosts

Improved per-AS statistics

Improved Grapher

Work by Daniel Medina

Packaging incident-reporting

Questions?

For More...

Consult your local man pages (of course)

CUFlow

<http://www.columbia.edu/acis/networks/advanced/CUFlow/>

FlowMonitor

<http://www.columbia.edu/acis/networks/advanced/FlowMonitor/>

FlowScan

<http://net.doit.wisc.edu/~plonka/FlowScan/>

Installing it all

<http://www.linuxgeek.org/netflow-howto.php>