IAB Unwanted Internet Traffic Workshop

Session 2 - Sources of the Problem

Danny McPherson danny@arbor.net

March 9, 2006

Well, first…

- What's unwanted Internet traffic?
  - Control Plane (e.g., BGP route hijacking)
  - Data Plane (e.g., SYN Flood)
  - Application Data (e.g., SPAM)
  - Economics & other motivators
- Hrmm…
  - Deep or wide?

About this discussion

- Meant to get the gears churning…
- Little (no?) new data here
- Went wide (and a little deep in a couple places)
- My typical [overly?] broad perspective (I.e., broad v. deep - operational, research, commercial, SDOish)

Unwanted [Internet] Traffic?

Control Plane

- BGP
  - Route Hijacking - many motivations:
    - DOS
    - MITM [?]
    - SPAM
    - Sophisticated Phishing
    - Sharing Data Anonymously [?]
    - Misconfiguration
  - Route instability
    - routing table growth (e.g., multi homing, communities, other)
    - route oscillation
    - Sources of churn (e.g., MEDs, link flaps)
    - More specifics/route leaks
    - New AFI/SAFIs (IP-VPN, Label, Flow_Spec, etc..)
- IGPs
  - Gratuitous LSPs/LSAs (e.g., refresh defaults, lifetime :-)
- DNS
  - Caching poisoning
  - Reflective “stuff”
  - */. records

Route Hijacking

- NANOG 36: Short-lived Prefix Hijacking on the Internet:
  - http://www.nanog.org/mtg-0602/pdf/boothe.pdf
- ***"Result: between 26 and 95 successful prefix hijackings occurred in December of 2005"***
- Note: prefix hijackings do not include events which appear to be the result of misconfiguration

BGP/IGP Transport Protection

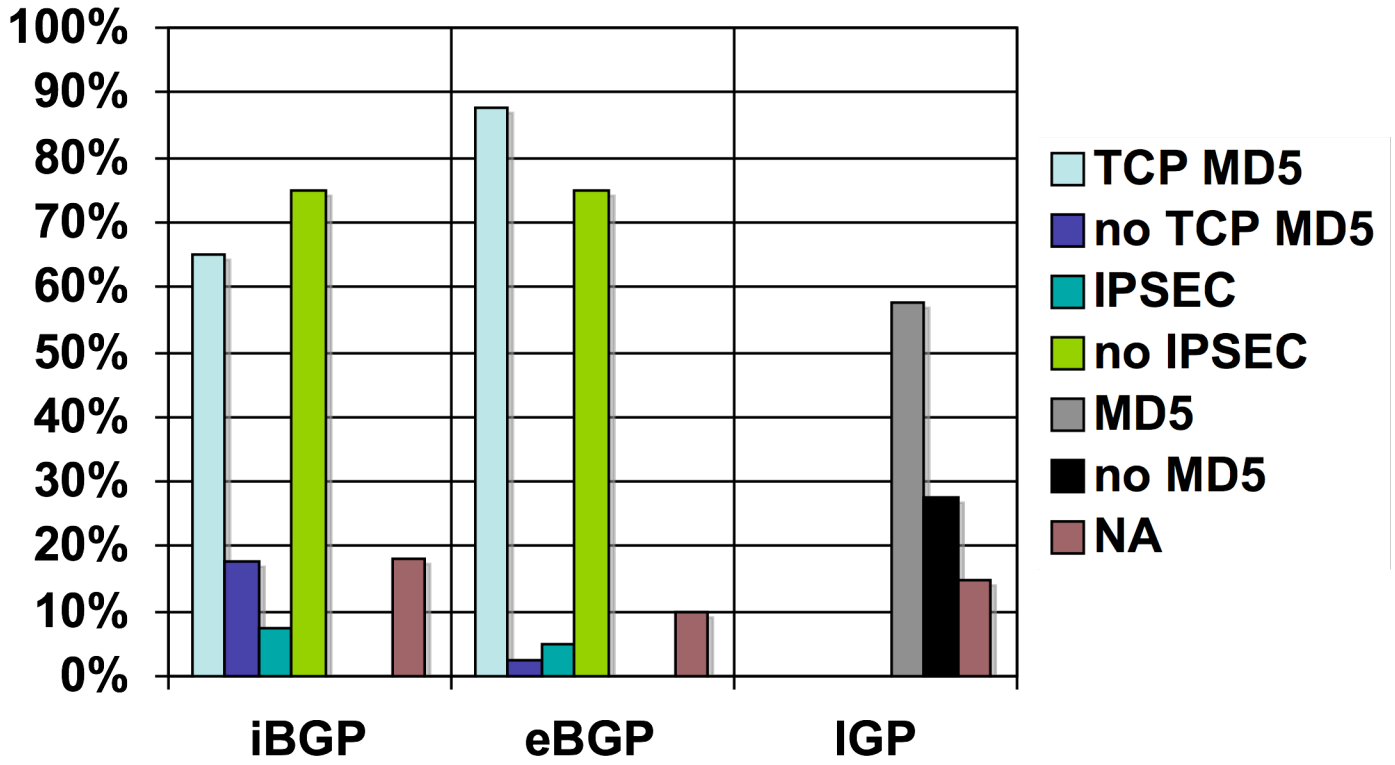

Detect Anomalous Events: SQL “Slammer” Worm

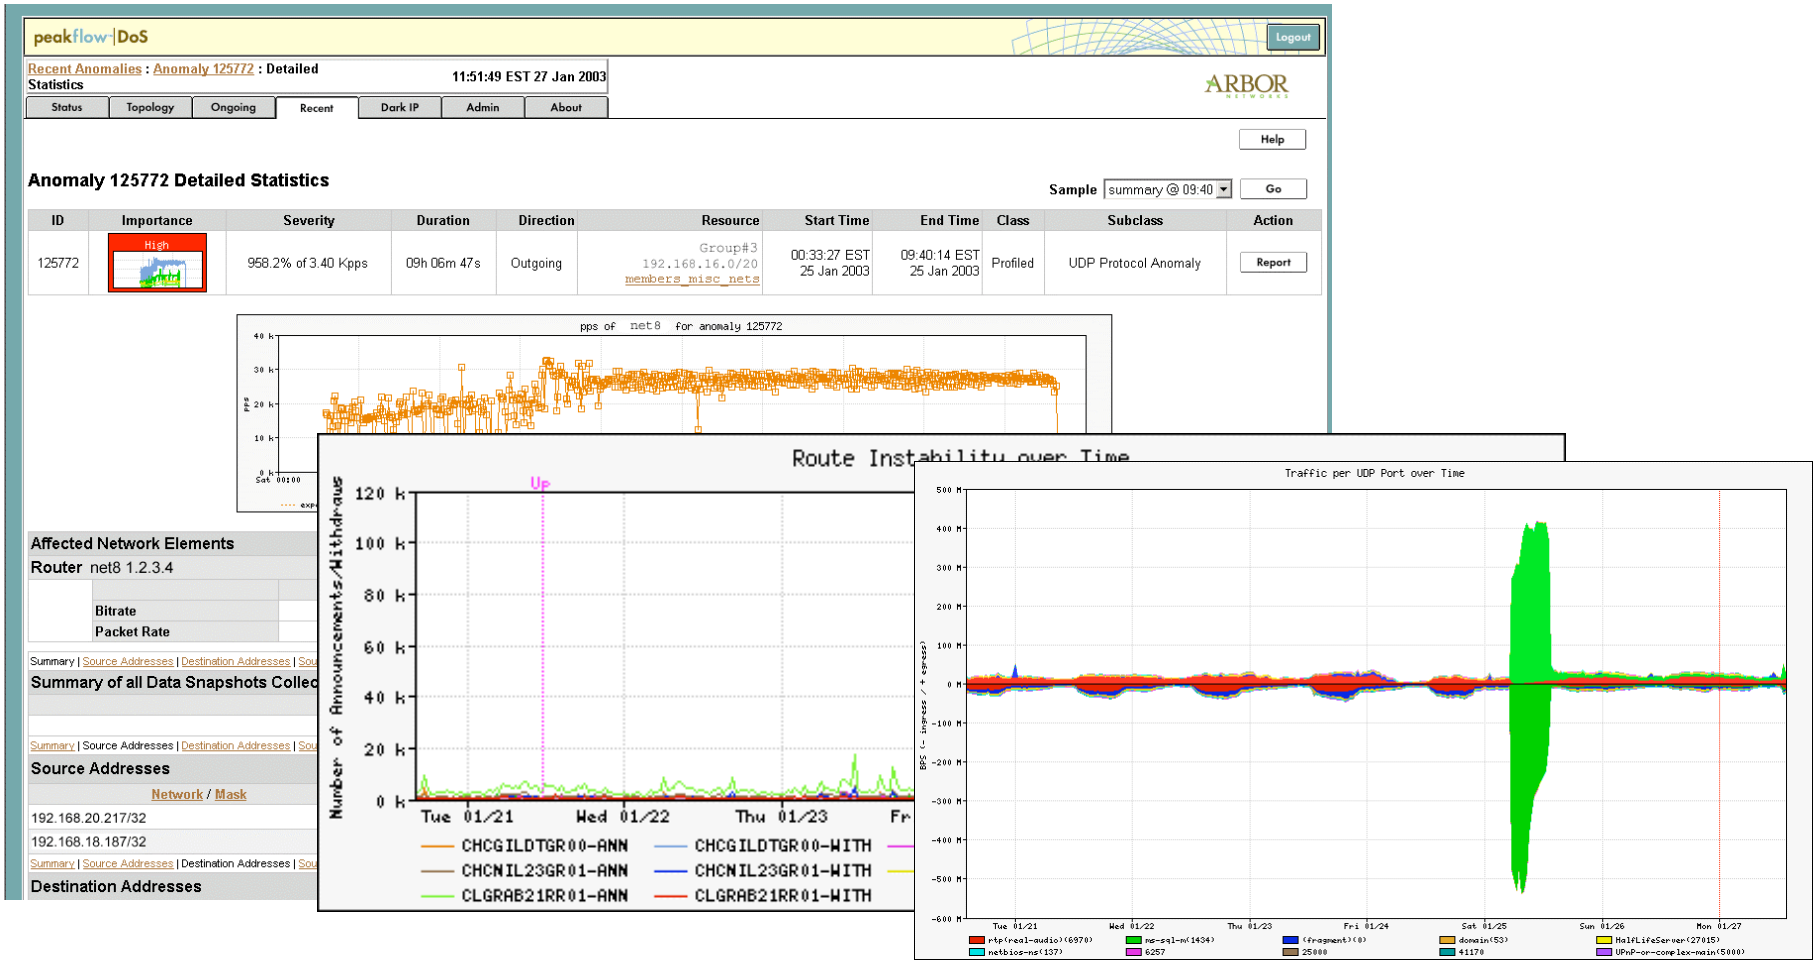

Slammer Data Plane Impact - A European SPs View

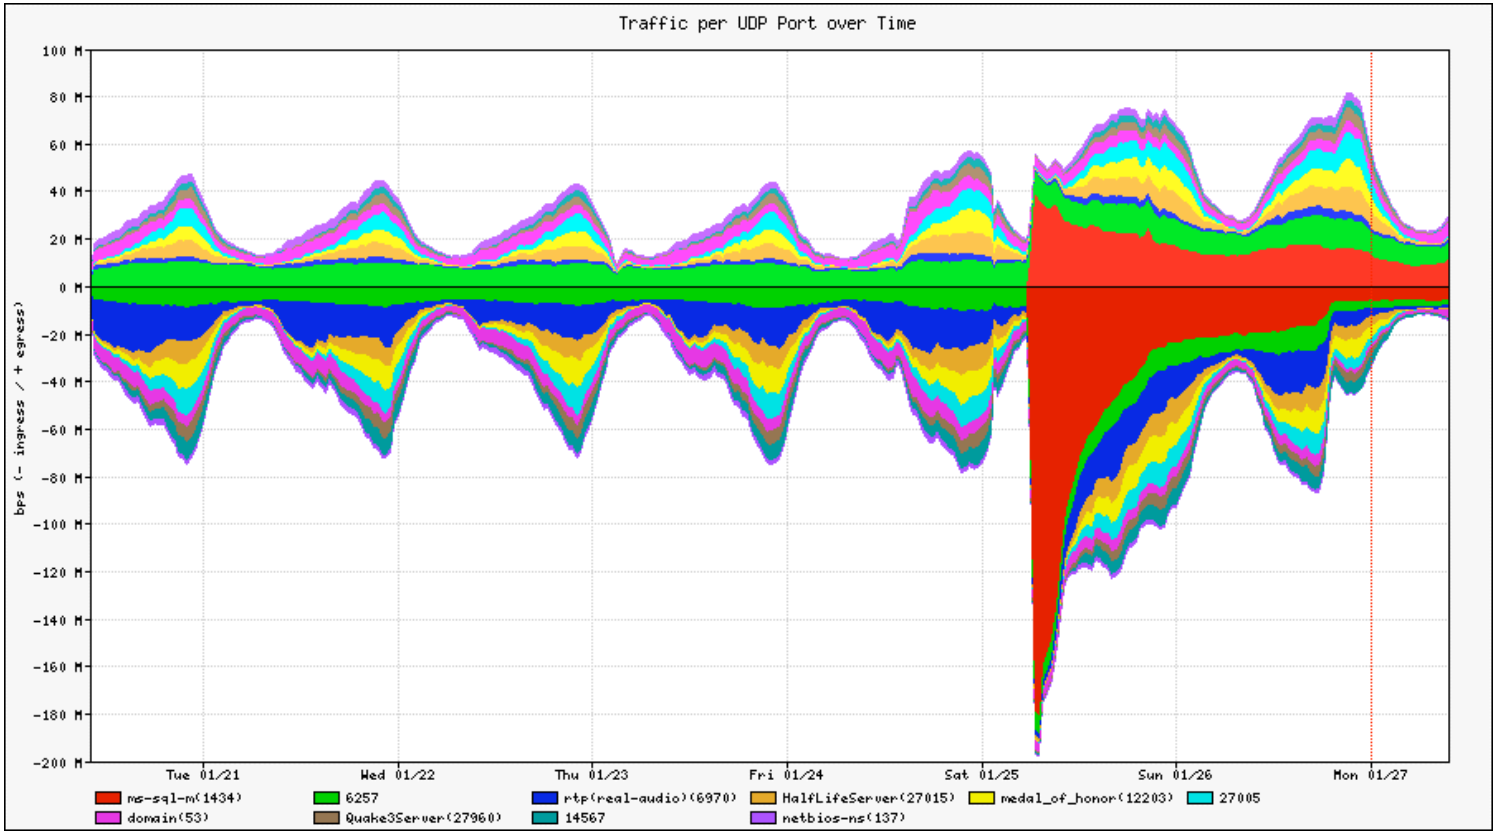

Slammer Control Plane Impact – THE BGP PICTURE

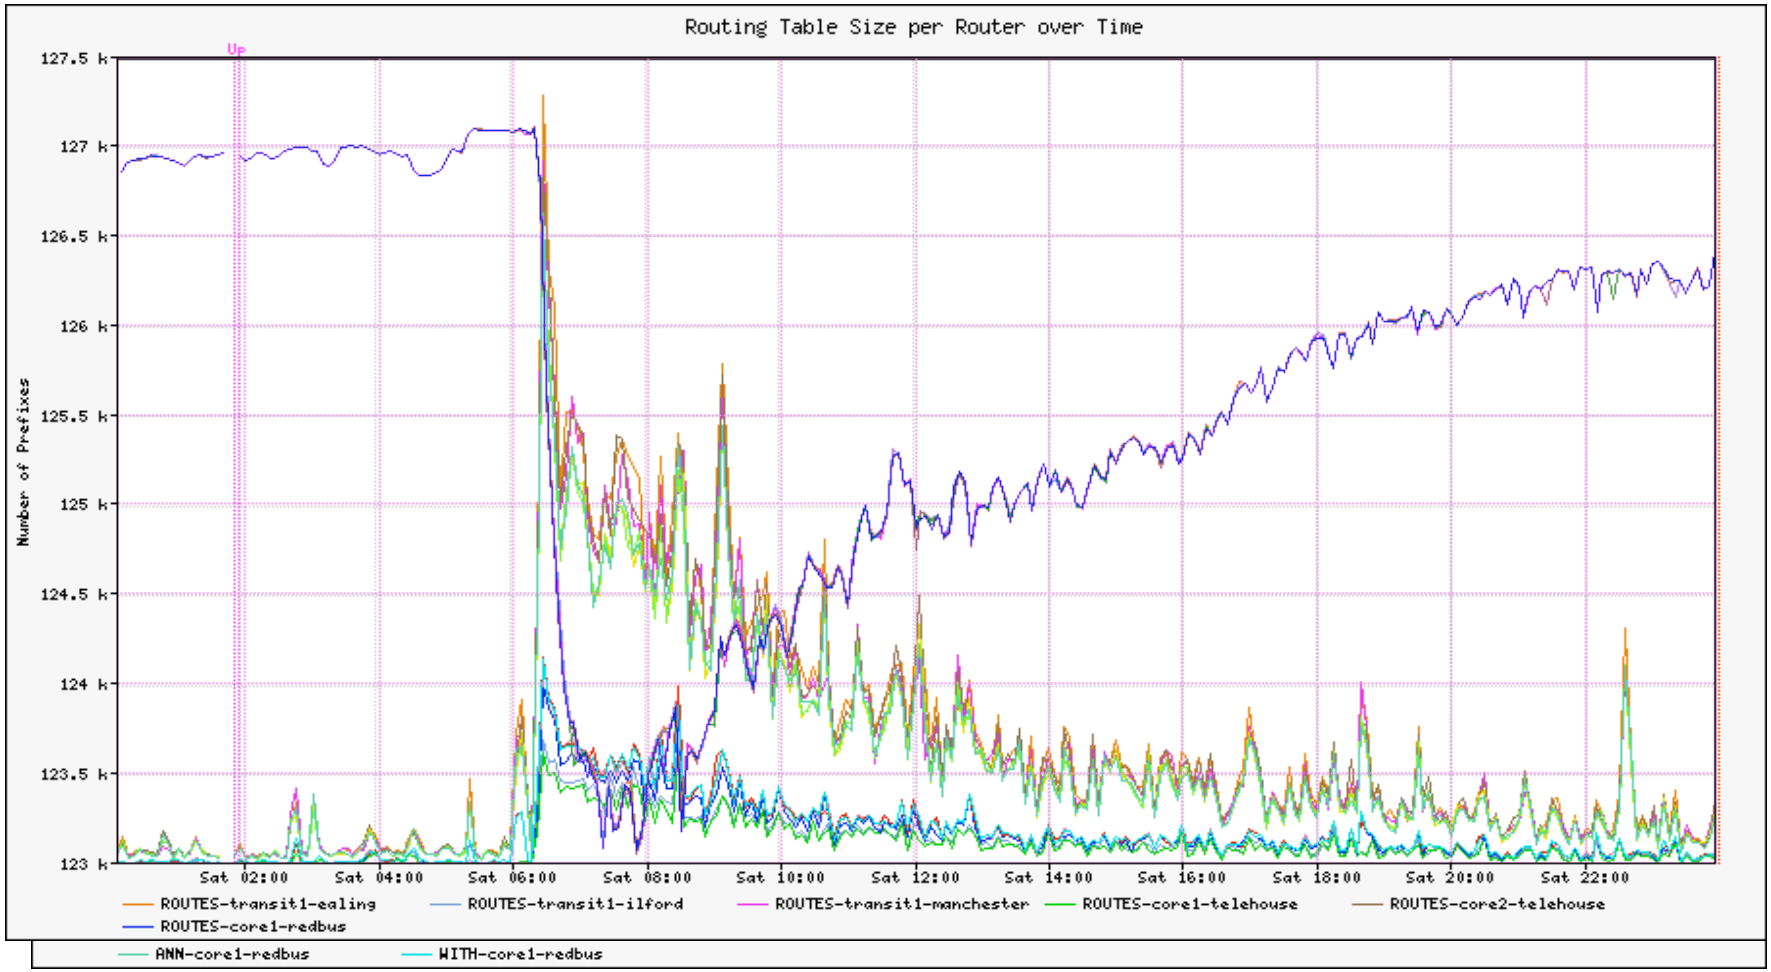

Operationalizing..

- Debugging capabilities
  - E.g., truly distributed nature of route hijacking and filtering - I may never see it..
  - IP-VPN.. Now add inter-provider w/non-congruent data/control planes
- Infrastructure compromise (lowly bot -> SP NMS/ bastion host -> high-end router -> .*)
- Crumbling network perimeters…

Data Plane

- Backscatter from spoofed attacks (and misconfiguration)
- Remote Infection attempts by bots/worms
- Misconfiguration: e.g. DNS/Network Management Apps
- Reflection/Amplification Attacks
  - Smurf
  - DNS
  - Etc…

About Spoofing

- Really only interesting for single-packet misuse functions..
- http://momo.lcs.mit.edu/spoofer
- ~23% of observed netblocks corresponding to ~24% of observed ASes allow spoofing

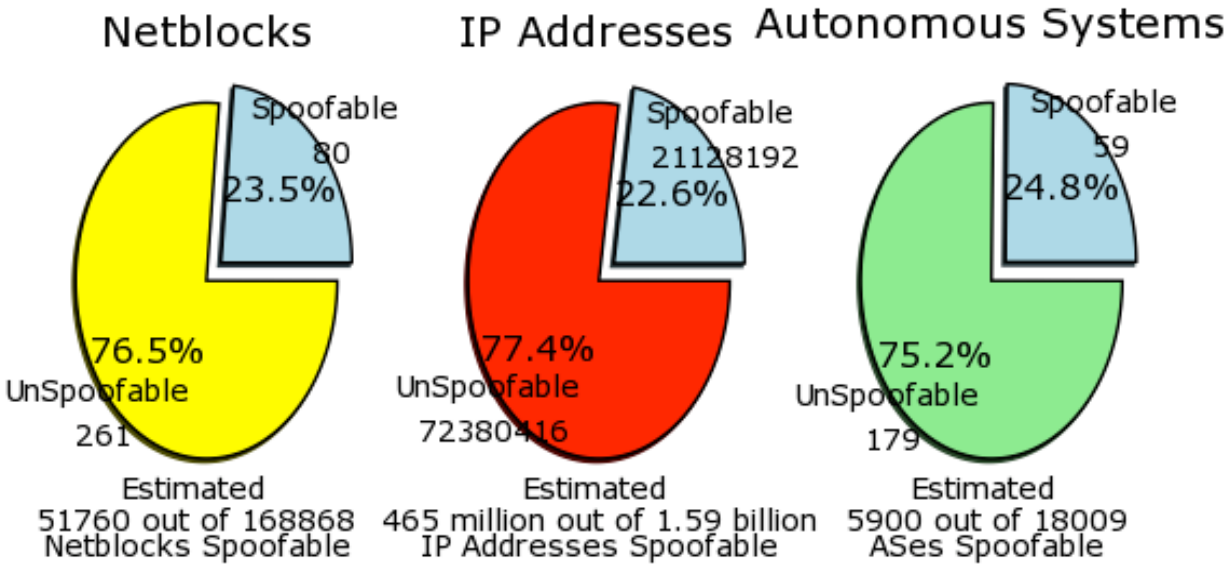

Spoofing out of Vogue?

- Why risk lowering firepower?
- Miscreants aren't as clueless as you might think (keen on uRPF/BCP38+)
- Ahh, and about those DNS attacks

Ingress Filtering Employment

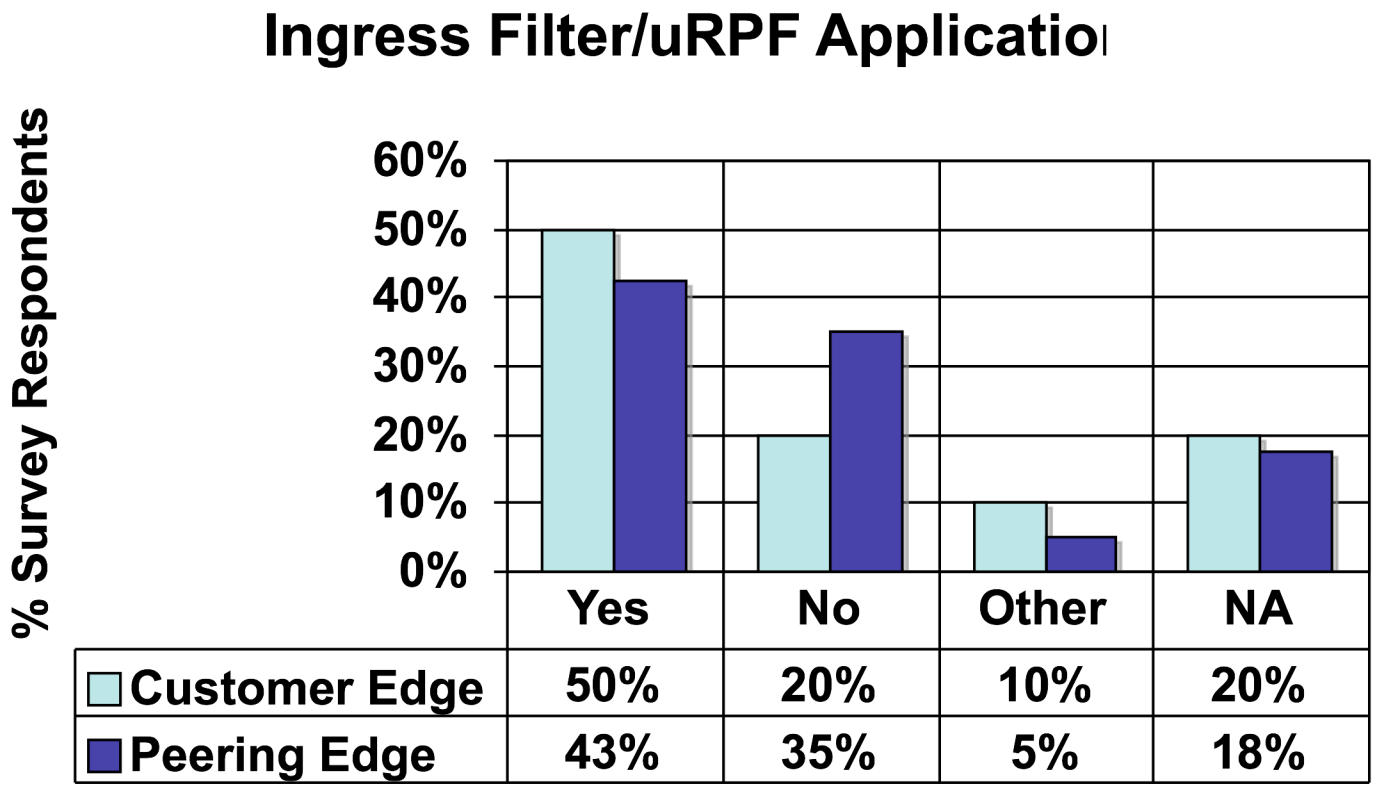

What kinds of spoofing?

- Dominated by spoofed TCP:

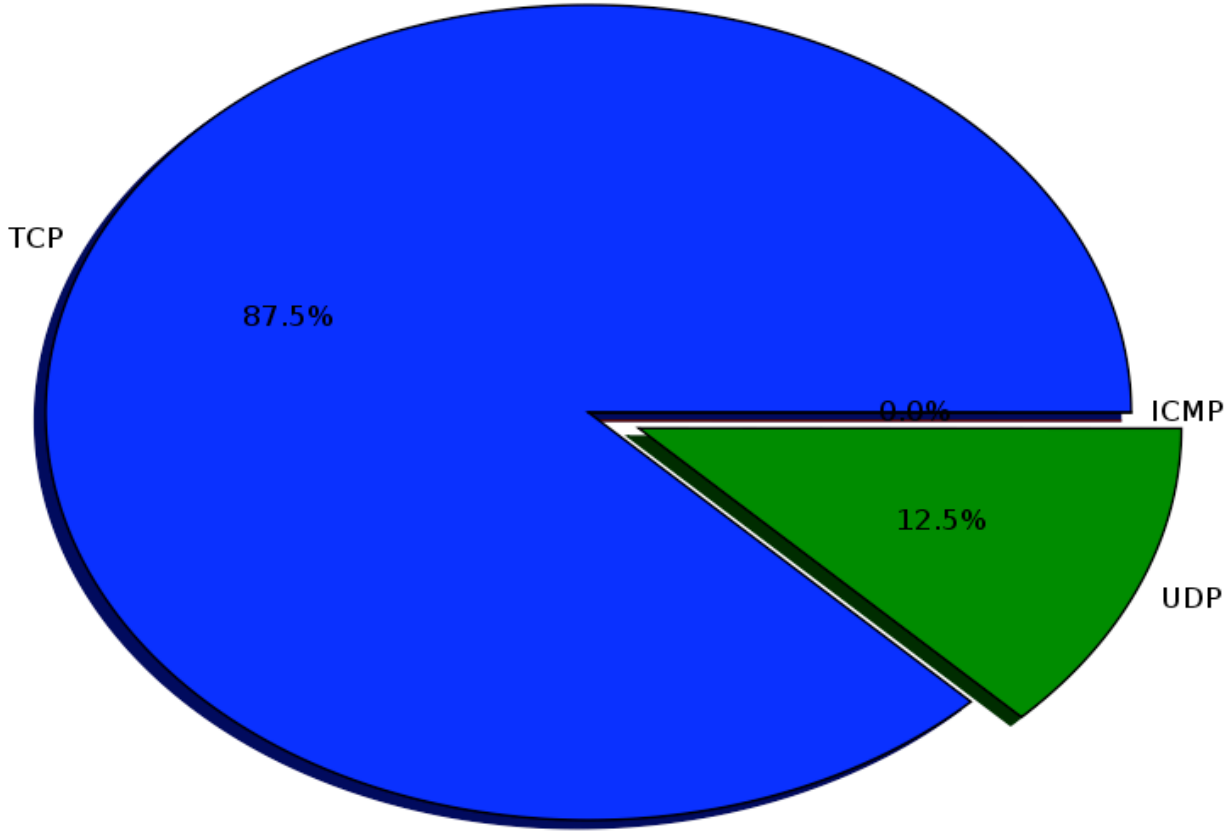

What kinds of spoofing?

- Dominated by spoofed SYNs:

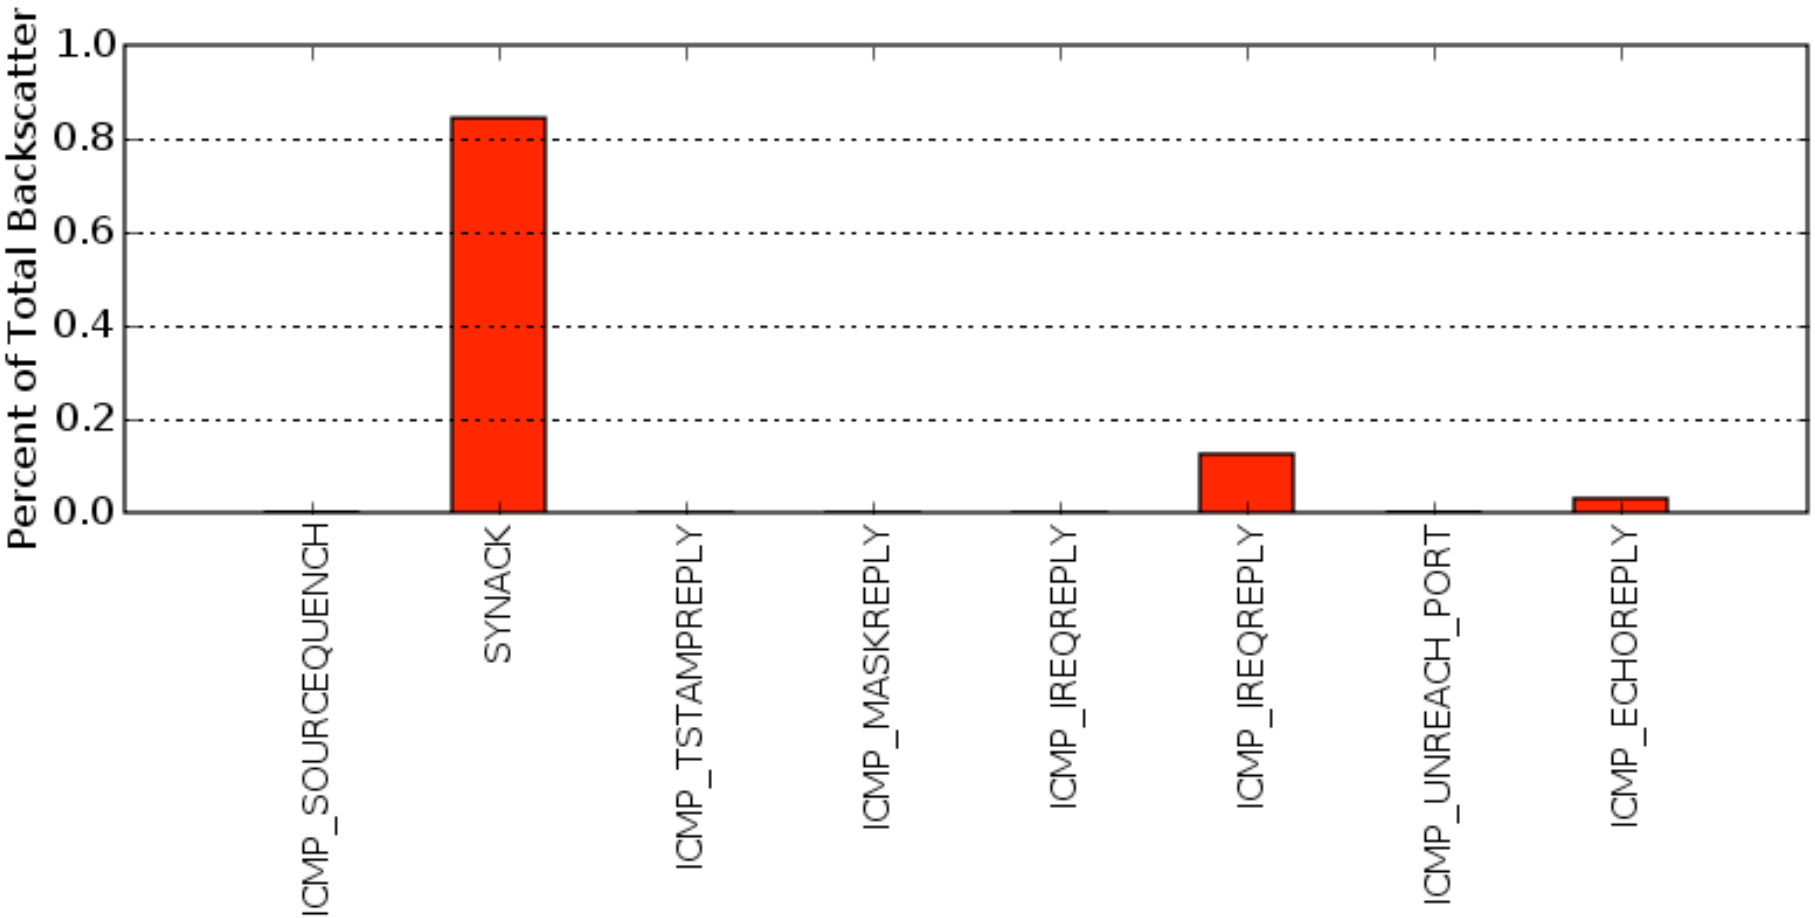

Spoofed TCP SYNs:
Top 10 Ports Targeted by

TCP Port	Service	Packets
80	HTTP (HyperText Transfer Protocol)	38805062
7000	W32.Gaobot, Spyboter, W32.Mydoom, W32.Mytob	2342659
6904	-	919211
300	-	828651
100	-	757745
25	SMTP (Simple Mail Transfer Protocol)	563894
6000	X11 - X-Windows	480937
3389	Microsoft Terminal Server (RDP)	391371
22	SSH	161991
7777	cbt/Oracle HTTP Server	155682

Ease of the Compromise

- Windows 2000/XP Honeypot
- Placed behind proxy:
  1. Rate limit traffic 12KB/s
  2. Disallow local network
  3. Log all traffic
- 12 experimental runs over a month:
  - 12-72 hour traces > 100MBs
  - Recruited into least **15 unique botnets**
  - **3 in one day!**
  - Bots used DCOM/RPC, LSASS

=> Bots are extremely prevalent

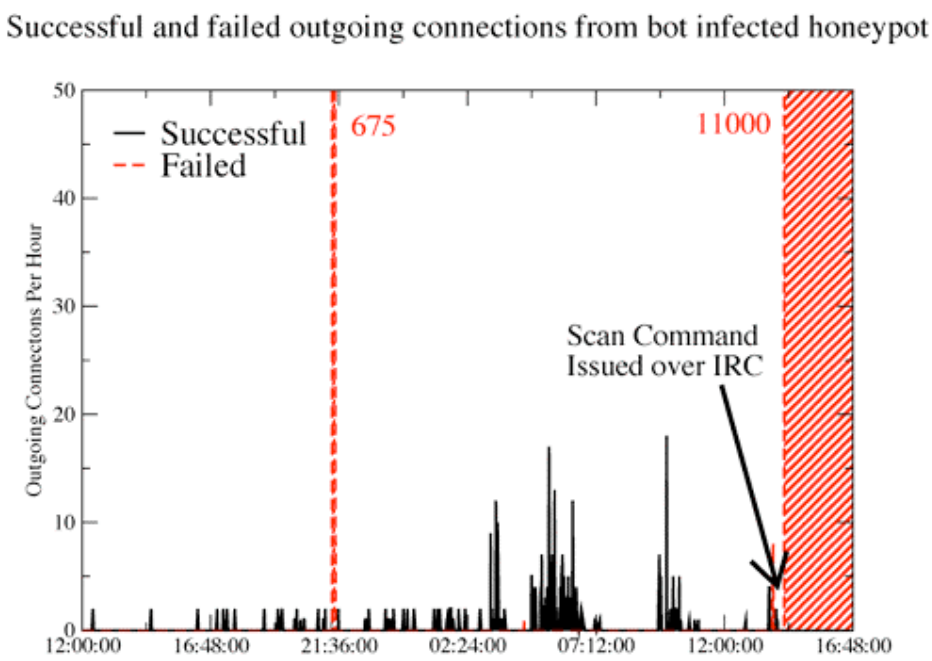

Just 2 worm infections during the experiment; (korgo, Windows LSASS TCP/445)

Bots, Bots and More Bots…

- Spam Relay
- DDoS
- Phishing
- Reconnaissance (compromise, vulnerable population census, target lists for coming vulnerabilities)
- Distributed cracking systems (e.g., Brute Force SSH activity)
- Easy Pickens Spaces (e.g., 24/8, 64/8)
- ID Theft, Keyloggers, License Keys, etc..
- Open Proxy
- Rbot capabilities include using webcams to capture video and still images(!) - dates back to a variant of [how apropos] SpyBot
- Bots are a commodity - no significant resource constraints
- Better organized
- Broadband proliferation; Always-on, more firepower, more hosts
- Organic growth

Traceback: Manual

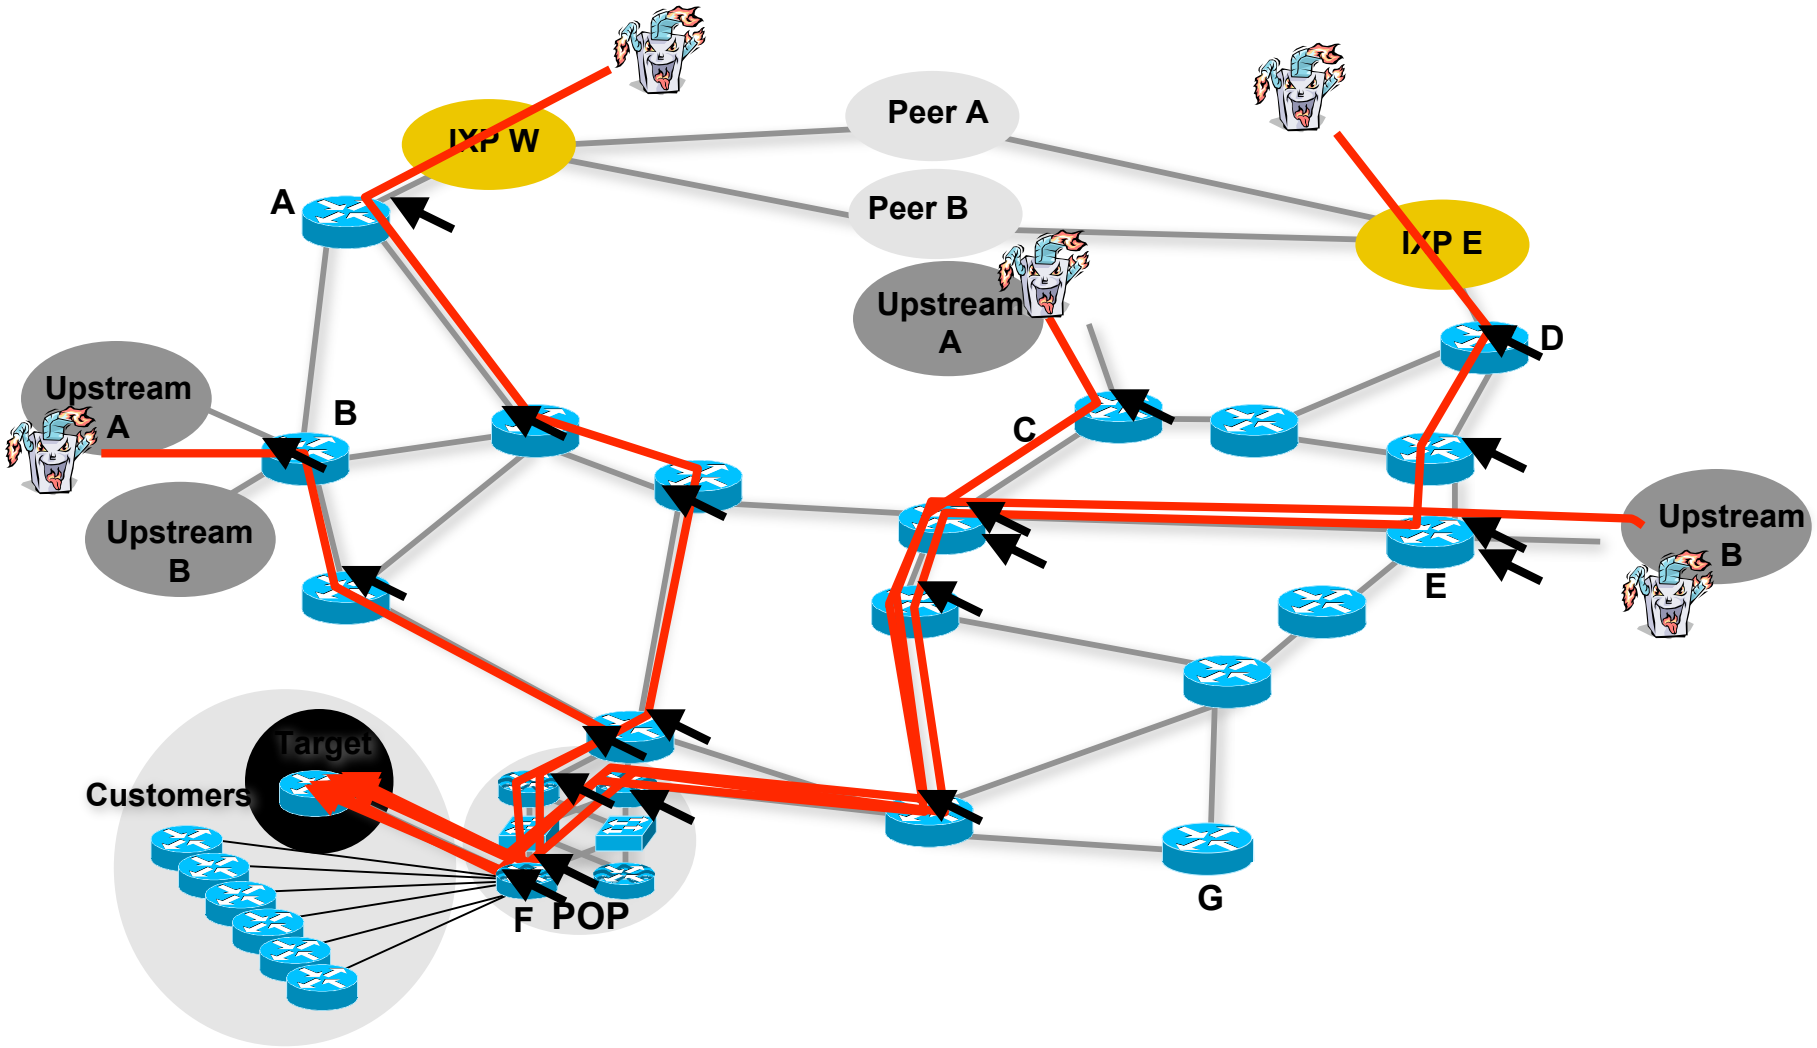

Traceback: Manual

- Classification ACL (cACLs) applied to customer interface:

```
access-list 101 permit icmp any any echo
access-list 101 permit icmp any any echo-reply
access-list 101 permit udp any any eq echo
access-list 101 permit udp any eq echo any
access-list 101 permit tcp any any established
access-list 101 permit tcp any any range 0 65535
access-list 101 permit ip any any

interface serial 10/1/1
ip access-group 101 out
```



```
router# sh ip access-list 101
Extended IP access list 101
        permit icmp any any echo (2 matches)
        permit icmp any any echo-reply (2171374 matches)
        permit udp any any eq echo
        permit udp any eq echo any
        permit tcp any any established (150 matches)
        permit tcp any any (15 matches)
        permit ip any any (45 matches)
```

- Once attack type is classified, Traceback ACL (tACLs) applied to egress then subsequent upstream interfaces back towards network ingress

```
access-list 170 permit icmp any any echo-reply log-input
access-list 170 permit ip any any

interface serial 10/1/1
   ip access-group 170 out
```



```
router# sh log
%SEC-6-IPACCESSLOGDP: list 170 permit icmp 1.1.1.1 (Serial0/1/1
*HDLC*) -> 192.168.1.1 (0/0), 1 packet
%SEC-6-IPACCESSLOGDP: list 170 permit icmp 2.2.2.2 (Serial0/1/1
*HDLC*) -> 192.168.1.1 (0/0), 1 packet
%SEC-6-IPACCESSLOGDP: list 170 permit icmp 3.3.3.3 (Serial0/1/1
*HDLC*) -> 192.168.1.1 (0/0), 1 packet
%SEC-6-IPACCESSLOGDP: list 170 permit icmp 4.4.4.4 (Serial0/1/1
*HDLC*) -> 192.168.1.1  (0/0), 1 packet
%SEC-6-IPACCESSLOGDP: list 170 permit icmp 5.5.5.5 (Serial0/1/1
*HDLC*) -> 198.168.1.1 (0/0), 1 packet
```

Escalation of Threats

- For example:
  - Slammer: Wide spread infection, induced network congestion
  - Code Red: DDoS against one IP
    - Changed IP/Null routed previous IP
  - Blaster: DDoS against hostname
    - Repeated DNS Shifts
    - Eventual NXDOMAINing of windowsupdate.com record
  - Deloder: Arbitrary DDoS toolkit
    - Hrmm…?
- Backdoors escalated from remote control (e.g., BO, NetBus) to harvesters and far more complicated
- Control channels include IRC commonly and other, encrypted mechanisms more and more.

Time to Market…..

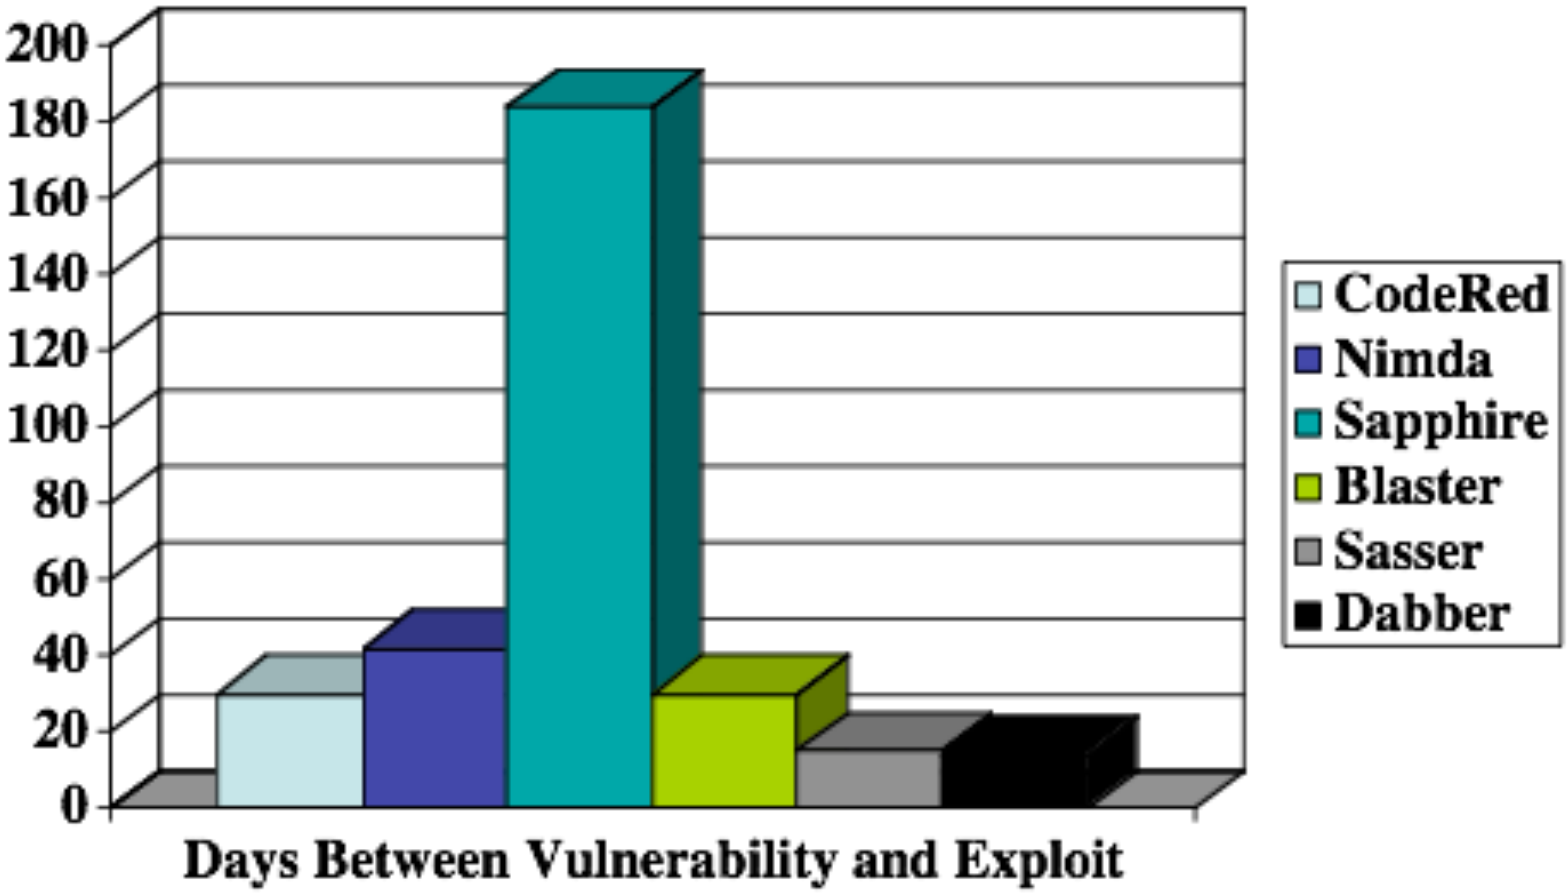

Persistence of compromised and vulnerable host population

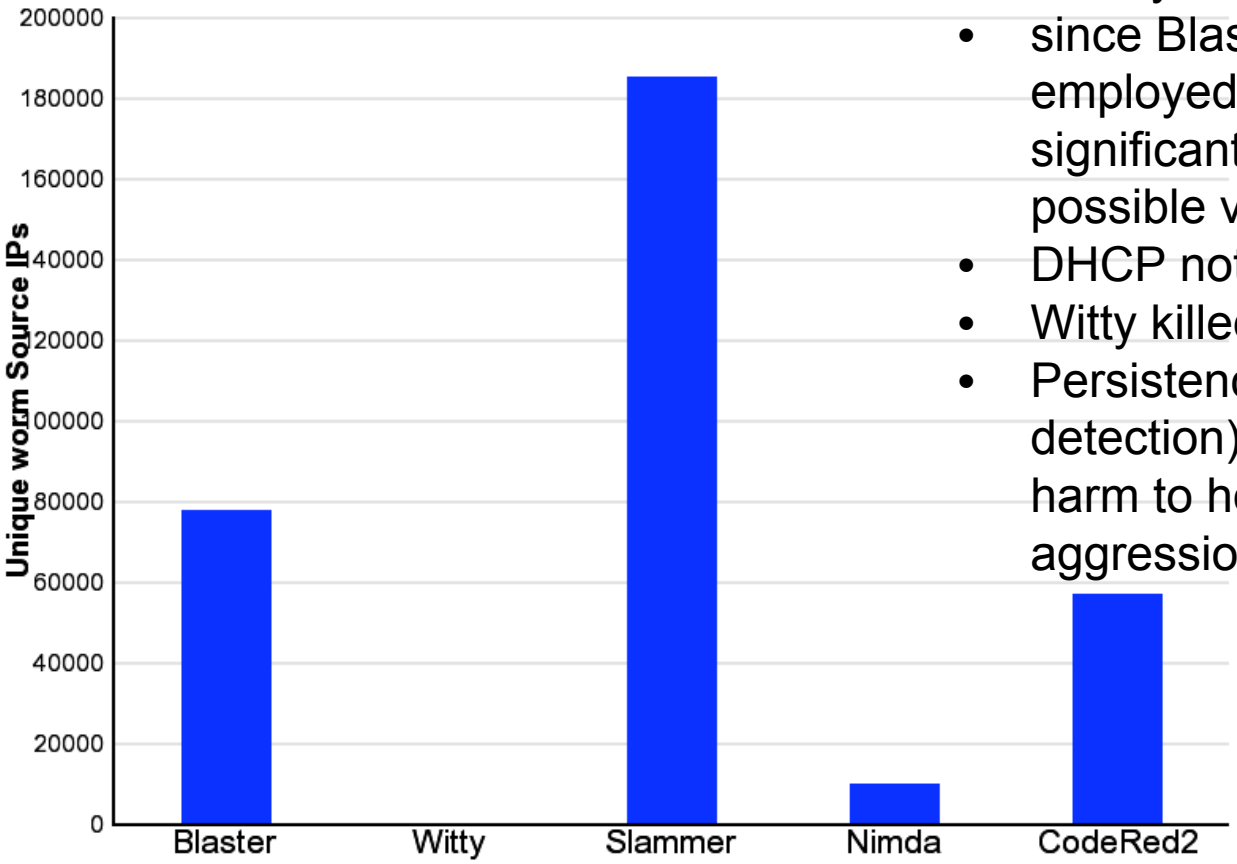

- 30 days on /8 darknet in late ‘05
- since Blaster TCP worm employed slow scans significantly under estimates possible vulnerable pop
- DHCP not accounted for
- Witty killed itself
- Persistence (and ease of detection) influenced by inflicted harm to host and scanning aggression

Worm Demographics

CodeRed ‘01

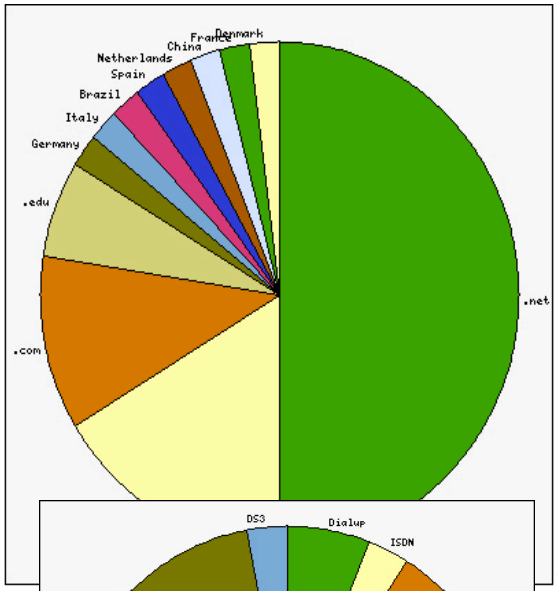

Nimda 9/01

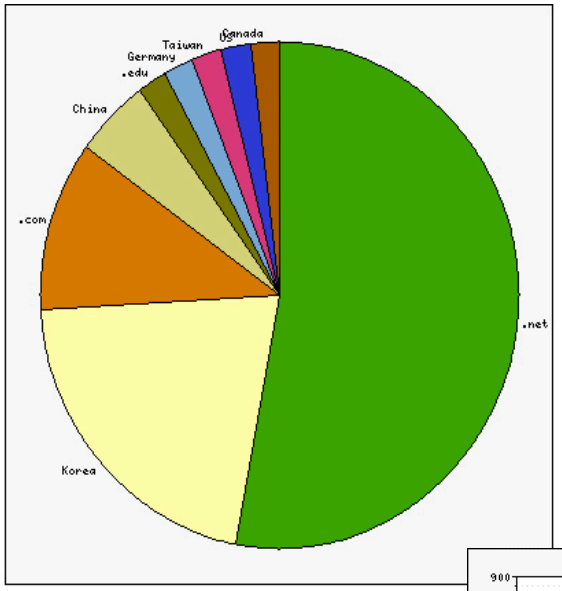

Blaster 8/03

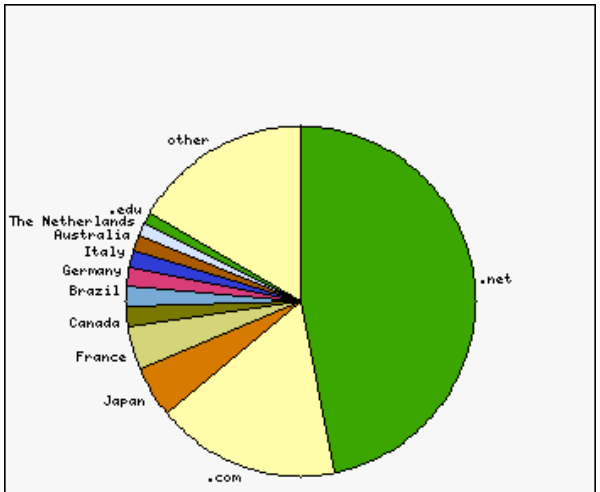

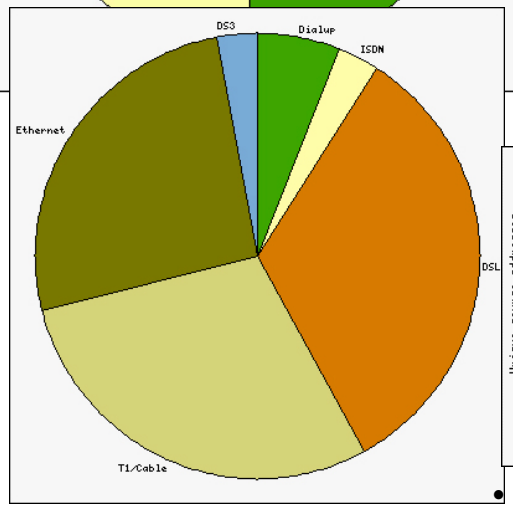

Nimda firepower

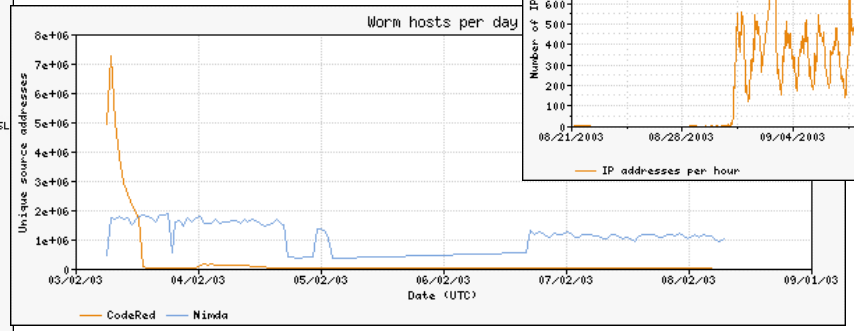

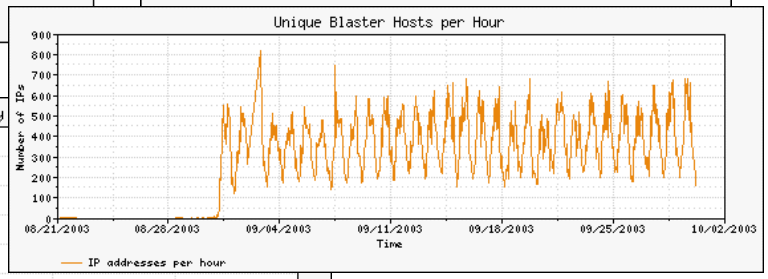

Nimda: Over one million unique hosts a day (August, 2003)

Today's Internet Locust?

- Was going to select a specific worm based on persistence, resources consumed

- Decided '**bots'** in general would be sufficiently fitting, as most other (beyond initial propagation and infection vector) is simply residual

Top Services Scanned

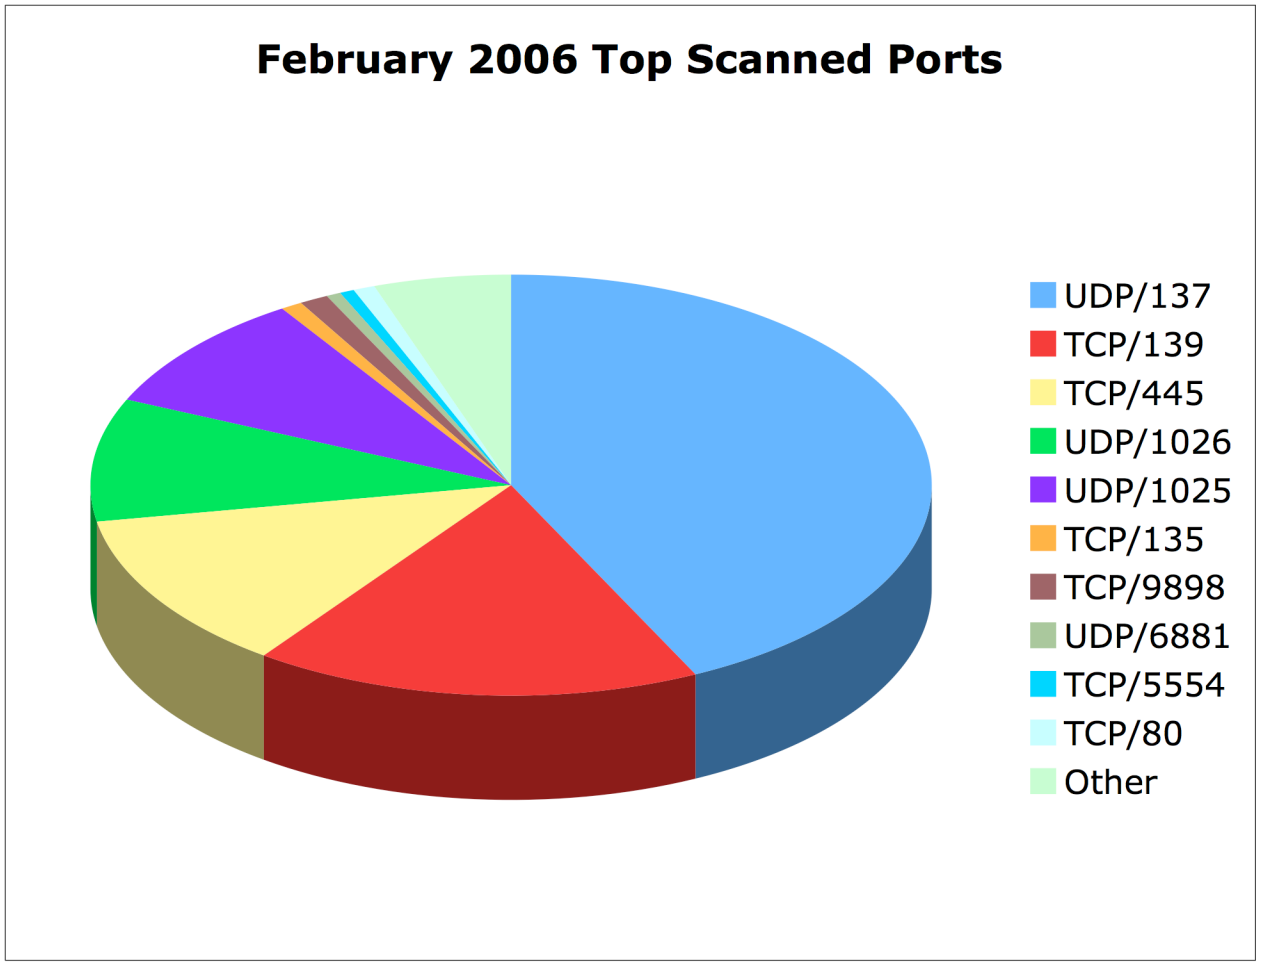

Command & Control

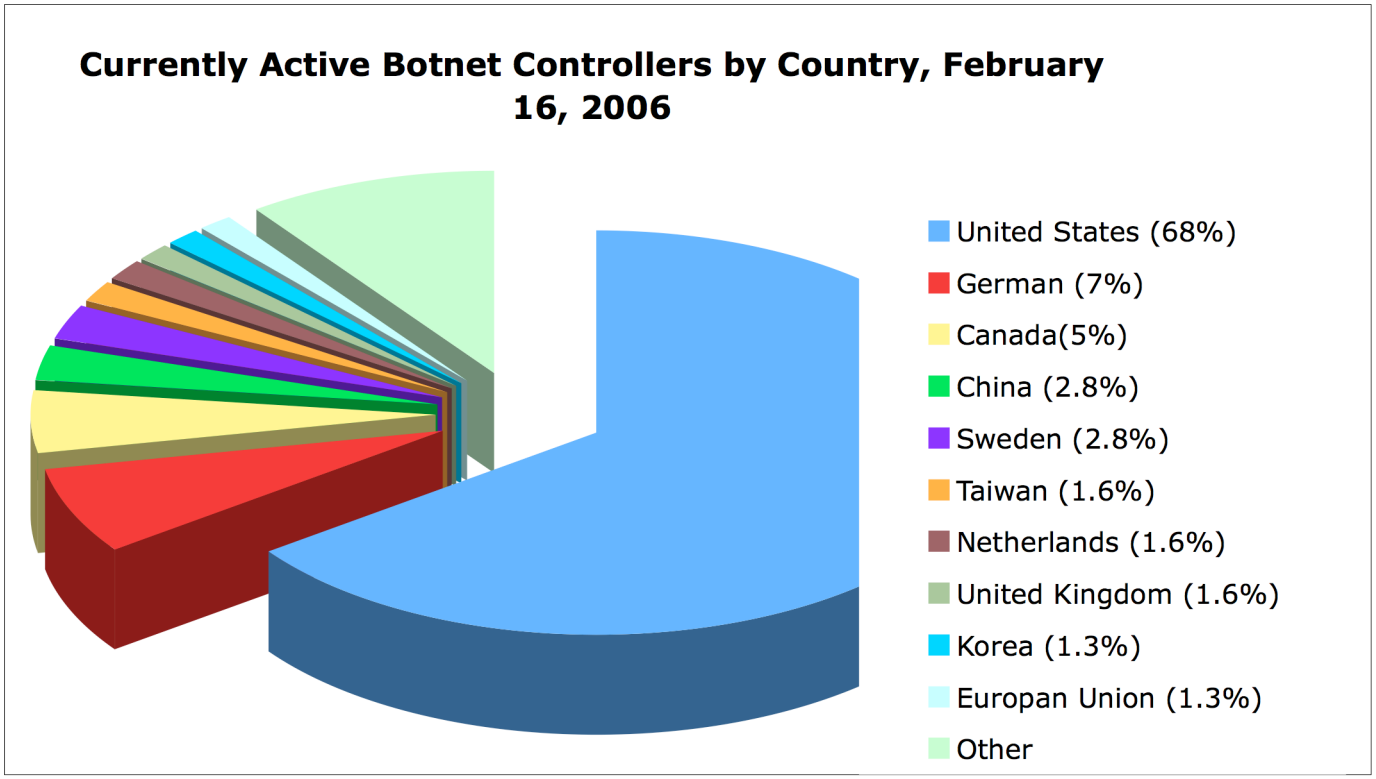

1398 active botnets evaluated at time of report generation

Infrastructure Security Survey

Survey Respondents

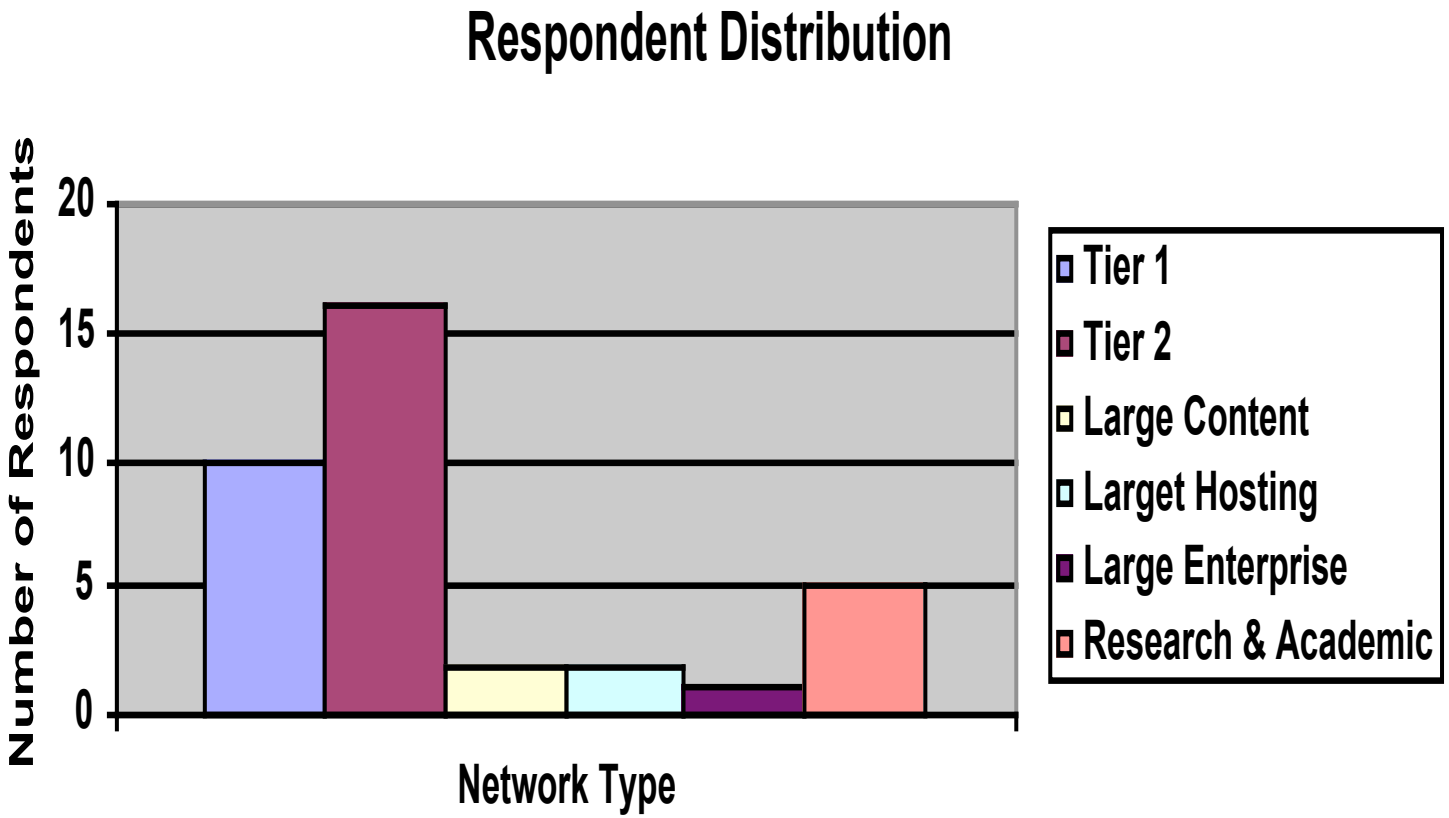

Primary Threat Concerns

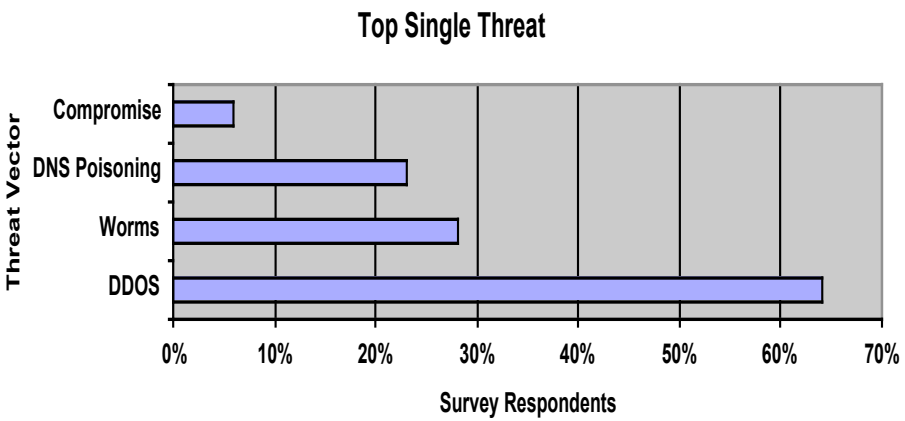

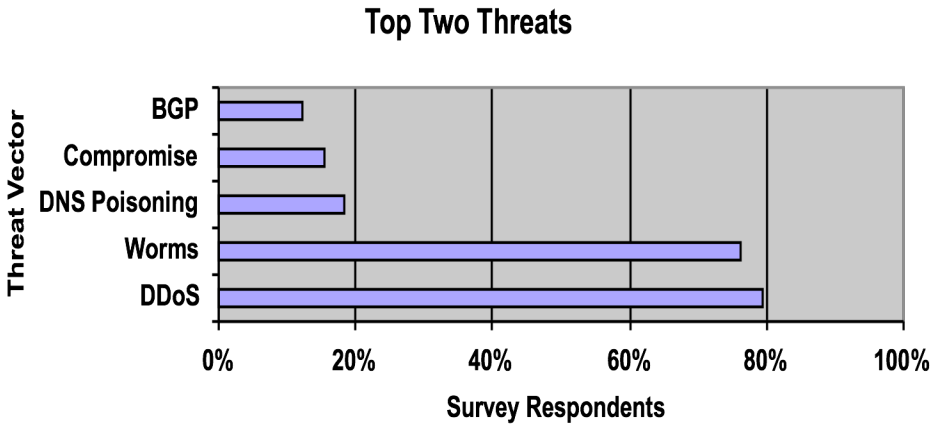

- DDOS was top concern, with worms coming in second
- Implicit DOS impacts of worm more concerning than worm payload itself
- BGP vulnerabilities weren't listed as anyone's **top** concern

Attack Detection Capabilities

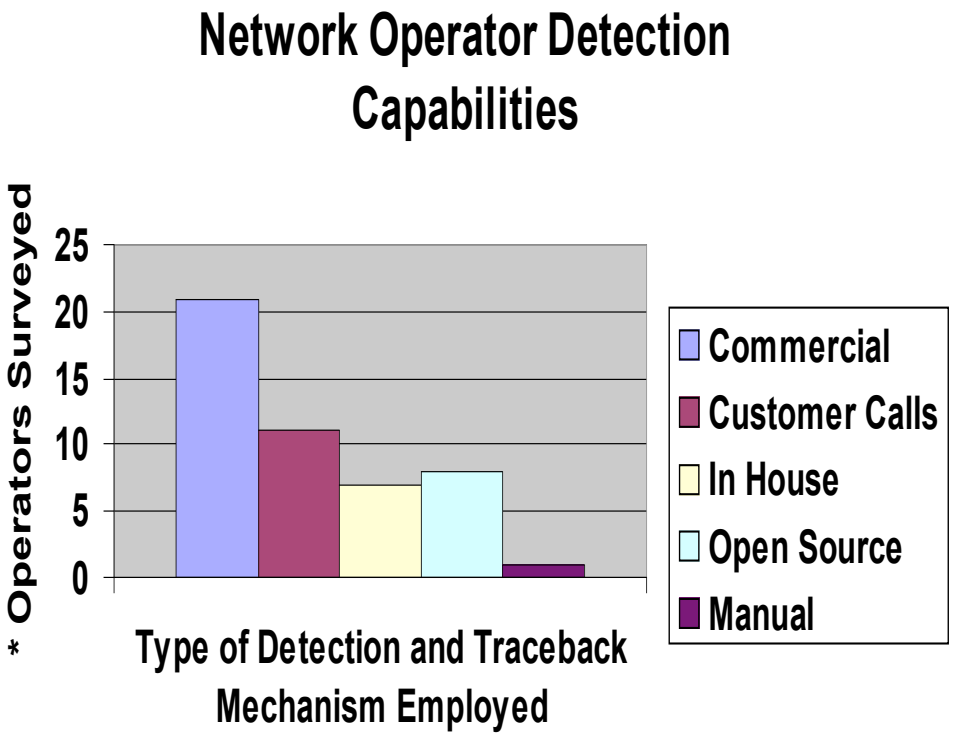

- Most operators had some commercial tools in place, though not covering the entire network perimeter
- Most provided employed multiple mechanisms for attack detection
- ISPs in wholesale/transit mostly rely on NOC trouble tickets (i.e., customer calls)

Customer Impacting Attacks

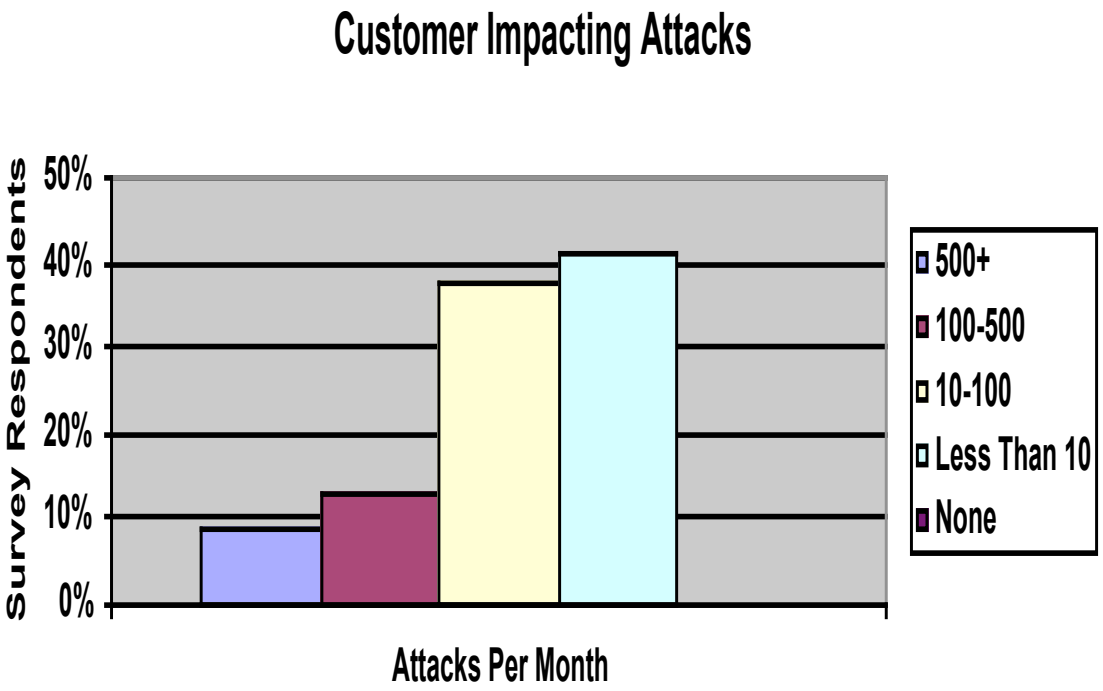

- An average of 40 actionable customer impact attacks per month were reported

Infrastructure Impacting Attacks

- Infrastructure impacting attacks were far less common, on the order of 1-2 per month on average
- These attacks were both directly at the infrastructure, as well as a result of collateral damage from customer attacks

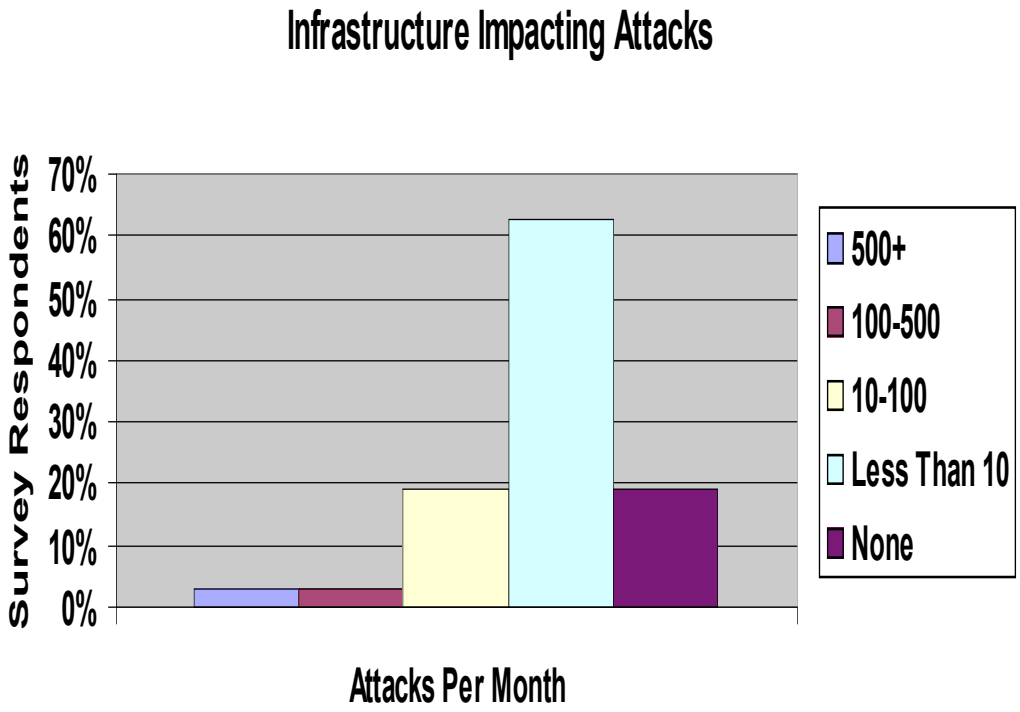

Largest Attacks Observed

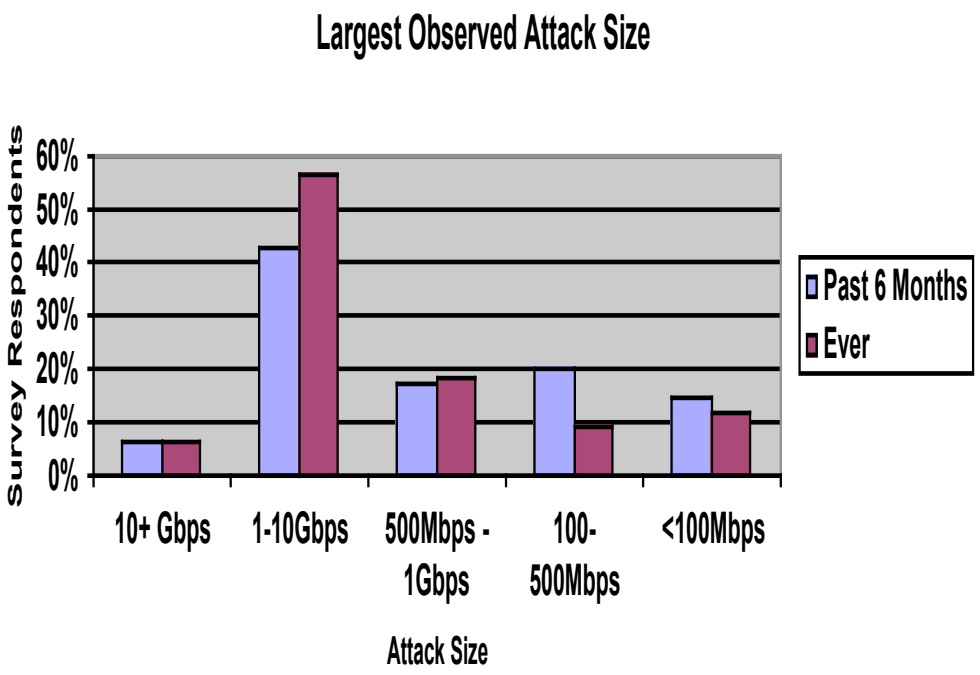

- Attacks greater than 10 Gbps sustained bandwidth were reported
- Not a large differential in largest attack ever v. largest in past six months - perhaps indicative of worsening problem

Attacks Reported to Law Enforcement

- Of actionable attacks, only ~1.5% are reported to law enforcement agencies
- Some of the reasoning provided:
  - Jurisdictional issue
  - Online gambling techniquely illegal is US
  - IRC users unloved
  - Customer profiles - they don't want attacks recorded
  - Lack of evidence and forensics data
  - Large amount of uncertainty from legal department

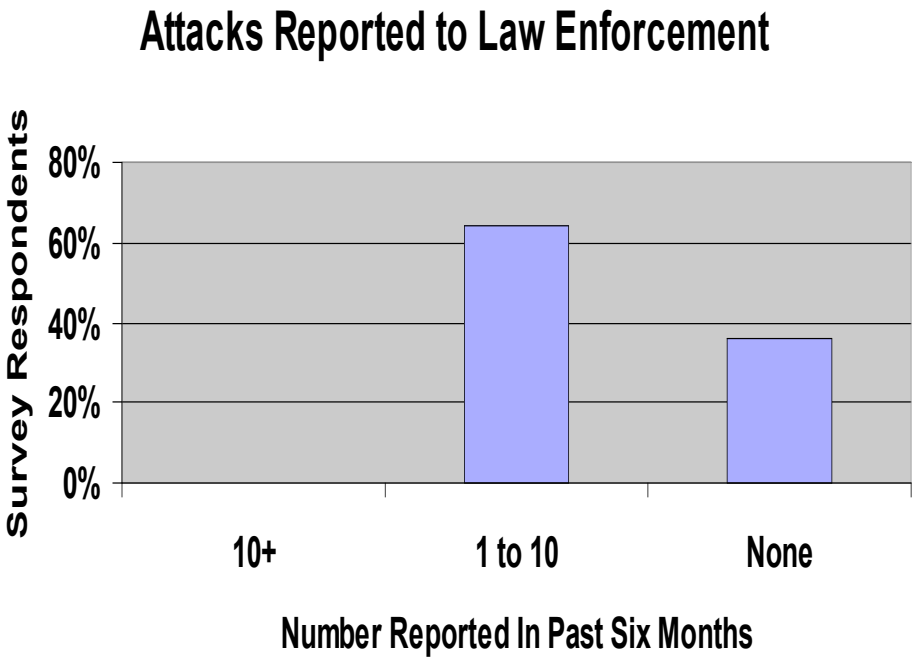

Most Common Mitigation Approaches

- Most common approaches are dst-based ACLs & BGP null-route destination
- BGP destination more scalable than ACLs and most common mitigation approach

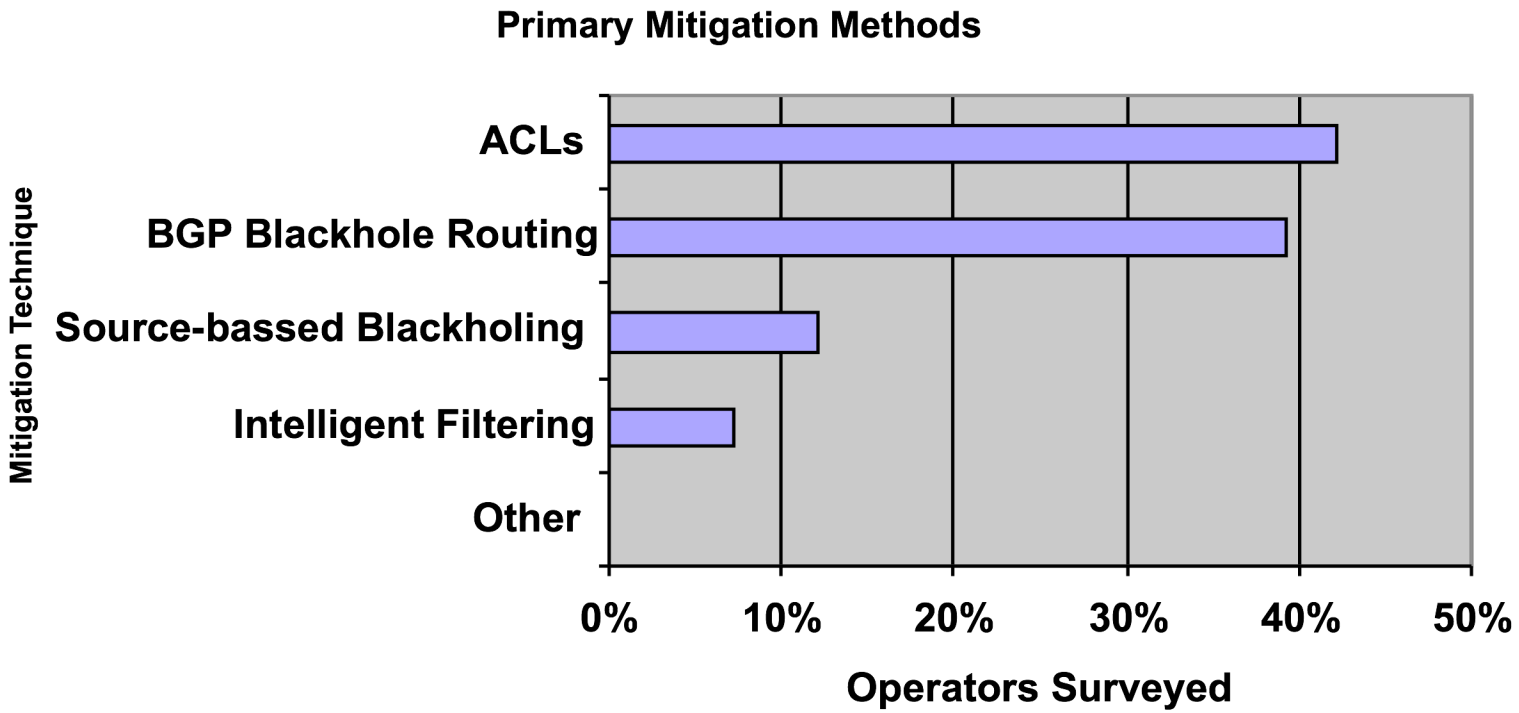

About the Survey

- Conduct bi-annually
- Hope to get more details and pose less ambiguous questions in future revisions
- Full survey report can be found here:
  - http://www.arbor.net/sp_security_report.php
  - Or send me email…
- Second edition due out in a couple weeks

Effects of Mitigation

- Majority of techniques today effectively complete DOS attack!
- Contacting source/upstream networks perceived not worth the effort - then hosts resurface
- Overt mitigation attempts are immediately noticed by miscreants

Sample Worm Activity Graph on 3/8/06 from Regional NSP

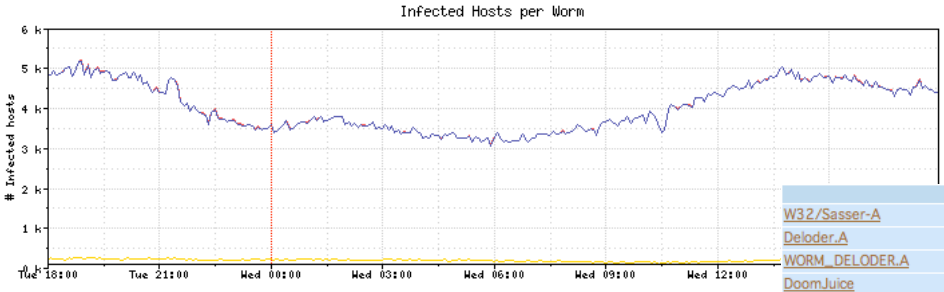

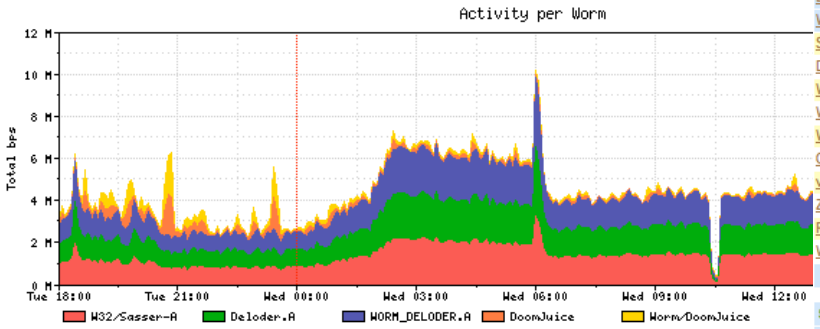

Name	Current	Average	Max ▼
W32/Sasser-A	4,483	4,034	5,198
Deloder.A	4,469	4,028	5,188
WORM_DELODER.A	4,469	4,028	5,188
DoomJuice	249	210	285
Worm/DoomJuice	249	210	285
SQL Slammer	67	76	102
Dabber.A	37	30	98
Worm/Dabber.A	37	30	98
Worm/Witty.A	28	25	44
Witty.A	28	25	44
Code Red or Nimda	27	26	46
various Windows worms	9	3	9
Zotob	6	4	9
Rbot variants	0	0	1
WINS scanning	0	0	1

Showing 15 Items

Name	Current	Average	Max ▼
W32/Sasser-A	1.64 Mbps	1.36 Mbps	3.27 Mbps
WORM_DELODER.A	1.64 Mbps	1.36 Mbps	3.27 Mbps
Deloder.A	1.64 Mbps	1.36 Mbps	3.27 Mbps
Dabber.A	140.00 Kbps	55.47 Kbps	2.77 Mbps
Worm/Dabber.A	140.00 Kbps	55.47 Kbps	2.77 Mbps
DoomJuice	74.05 Kbps	196.62 Kbps	2.10 Mbps
Worm/DoomJuice	74.05 Kbps	196.62 Kbps	2.10 Mbps
Dark IP	69.00 Kbps	63.33 Kbps	887.00 Kbps
various Windows worms	2.60 Kbps	1.34 Kbps	13.02 Kbps
Witty.A	1.95 Kbps	8.45 Kbps	114.00 Kbps
Worm/Witty.A	1.95 Kbps	8.45 Kbps	114.00 Kbps
Code Red or Nimda	0.00 bps	878.00 bps	5.12 Kbps
Zotob	0.00 bps	542.00 bps	6.90 Kbps
SQL Slammer	0.00 bps	248.00 bps	55.00 Kbps
WINS scanning	0.00 bps	4.00 bps	1.29 Kbps

Accidental?

DoCoMo and thttpd: i-mode DDoS attack!

```
Date: Thu, 02 Aug 2001 11:22:14 -0700
>From: Jef Poskanzer <jef@acme.com>
 To: thttpd@bomb.acme.com
 Subject: [THTTPD] DoCoMo and thttpd

 Hey, is anyone on the list familiar with DoCoMo?  Apparently it's a type
 of cell-phone / web browser device from Japan.  I have suddenly started
 getting a [whole] lot of hits to http://www.acme.com/software/thttpd/ with
 various versions of DoCoMo in the user-agent field.  Unfortunately the
 referrer field is blank, which makes it difficult to figure out why this is
 happening.  Current working theory is that some server run by the DoCoMo
 company switched over to using thttpd, and I'm getting the usual spillover
 from any 404 pages on their site.  I've seen this effect before with large
 ISPs, but never with such a high volume of hits.  My bandwidth is pegged
 to the throttle right now, and they're not even fetching the inline images
 (which by the way means I'm not getting any ad impressions from these
 hits, which is somewhat annoying).  [...]
 Jef Poskanzer  jef@acme.com  http://www.acme.com/jef/
```

Museum of Broken Packets

http://lcamtuf.coredump.cx/mobp/

Exhibit 2: lost its head

```
10:06:19.235208 213.76.114.40.18245 > 212.244.100.102.21536: SE 795438439:795438776(337) ack
                794976622 win 12147 urg 28261 <[bad opt]> (DF)

        0000: 4500 017d 5d03 4000 7906 22a8 d54c 7228 E..}].@.y."..Lr(
        0010: d4f4 6466 4745 5420 2f69 6d67 2f62 616e ..dfGET /img/ban
        0020: 6572 2f73 7964 6e65 7932 3030 302e 6a70 er/sydney2000.jp
        0030: 6720 4854 5450 2f31 2e31 0d0a 4163 6365 g HTTP/1.1..Acce
        0040: 7074 3a20 2a2f 2a0d 0a52 6566 6572 6572 pt: */*..Referer
```

This little poor thing looks just odd. When you take a closer look, you will see that port numbers and other TCP parameters of this packet are actually... constructed of what should be the packet payload! Source port and destination port, two 2-byte values that start every TCP header, are 18245 and 28261 - 0x4745, 0x5420 in network endian order. This translates to ASCII string 'GET ', a beginning of a HTTP request. This kid has lost its TCP header, but IP header (with protocol type set to TCP) and TCP payload are still there... We started to see thousands of packets just like this one somewhere in the middle of 2000, coming from many locations in Poland. After some time, we realized that all were generated by a badly broken Nortel CVX access servers deployed country-wide by the Polish Telecom. Firmware was fixed within a month or so, but this priceless packet dump will live forever.

Exhibit 3: espionage

```
+ TCP 0x14 64.4.14.250:80 -> 193.0.67.34:62990 ttl=1 off=0x0 id=0x7529 tos=0x0 len=40 phys=46

        45 00 00 28 75 29 00 00 01 06 F1 86 40 04 0E FA
        C1 00 43 22 00 50 F6 0E 00 00 00 00 00 D0 79 FE
           50 14 00 00 EB 82 00 00 00 00 00 00 00 00
```

This packet arrived from www.law10.hotmail.com, one of web servers handling Hotmail traffic, to a valid address in a cyber-cafe during its hours of operation. It looks legitimate - ah, just a typical RST|ACK packet. What is interesting is TTL, set to 1. This behavior, observed there and in few other places, might be a result of some fairly uncommon routing problems, but our paranoid minds kept telling us this must be something more. What if, within an established TCP connection, you started sending any of your responses several times, with increasing TTL, starting from something pretty low? Well, you get just another version of traceroute, and should receive ICMP responses from subsequent routers on the way to your target. But unlike the traditional, "standalone" version, this embedded traceroute that lives within a legitimate session established by the customer, will cut thru almost every stateful firewall and address translation, allowing you to obtain a valuable information about their internal network structure, depth, target distance and such. All this without risking being detected, really. Well done, packet! You looked so innocent...

Sources of the Problems?

- Economics have changed the dynamics of the game
- Anonymization factor with global Internet
- Prosecution - unlikely
  - Better follow the money, following the bits is futile
- Host sanitization
  - Automate walled garden and clean-up models
  - But careful there; 911/000/etc..
  - and lest we not forget - operator profitability
- Persistence of compromised & vulnerable host population
- Lack of network and security clue? - nah, seems as though all communities are generally clueless [REF danny@ xNANOG, xRIPE, xNetworkers, xIETF, xxMISC, /'06 et al.] - and this is the network clueful population!
- IETF results in security bolted on after the fact (implementation->development/WG->community->security->/)
- No secure central authoritative database for address and associated policy ownership

Acknowledgements

- Jose Nazario, Dug Song, Evan Cooke Rob Thomas and Michael Bailey, all of whom I rely on quite regularly