Network Telescopes: The FloCon Files

David Moore, Colleen Shannon

UCSD CSE

{dmoore,cshannon}@caida.org

www.caida.org

caida

Flocon Stream of Conciousness

- There are "reseachers" seriously interested in pieces of operational problems.
  - anomaly detection, early worm detection
  - flow aggregation, line-speed summarization
  - distributed data collection
  - modeling of "normal" traffic
- However, they can really use your help to understand the questions you currently ask and what you'd like to ask, but can't now.

What is CAIDA?

- Cooperative Association for Internet Data Analysis
- Goals include measuring and understanding the global Internet.
- Develop measurement and analysis tools
- Collect and provide Internet data: topology, header traces, bandwidth testlab, network security, DNS
- Visualization of the network

Current Project Areas

- Routing topology and behavior
- Passive monitoring and workload characterization
- Internet Measurement Data Catalog
- Bandwidth estimation
- Flow collection and efficient aggregation
- Security: DoS and Internet worms, syslog/SSH
- DNS performance and anomalies
- Visualization
- P2P traffic detection and modelling

Tools

- CoralReef, NeTraMet, cflowd – packet, flows
- Walrus & Otter, libsea, PlotPaths - visualization
- NetGeo – IP to geography (mostly defunct)
- Skitter – large scale traceroute
- Graph::Chart.pm, GeoPlot.pm – plotting
- ASFinder.pm – IP to prefix/AS from routing table
- Beluga, GTrace – user-level traceroute viz
- dnstat, dnstop – passive DNS analysis
- DBHost, OWL – historical network meta-data (whois, DNS)
- Collaborations:
  - RRDTool, AutoFocus, PathRate/PathLoad

What is a "Network Telescope"?

- A way of seeing remote security events, without being there.
- Can see:
  - victims of certain kinds of denial-of-service attacks
  - hosts infected by random-spread worms
  - port and host scanning
  - misconfiguration

Network Telescope

- Chunk of (globally) routed IP address space
- Little or no legitimate traffic (or easily filtered)
  - might be "holes" in a real production network
- Unexpected traffic arriving at the network telescope can imply remote network/security events
- Generally good for seeing explosions, not small events
- Depends on statistics/randomness working

Amount of Telescope Data

- Currently collecting 30G/day of compressed data, and this is not including NetBios.
- Some "real-time" web reporting.
- Keep packet headers for a couple days, more summarized data longer, everything automatically rolled off to tape archive system.

Flat File Compression

- Heard bzip2, gzip.
- We really like **lzop** for many things. It's close to gzip -1 size, but: faster, block-based, block checksums, …
- Both lzop, gzip -1:
  - Allows packet capture to disk at higher data-rates.
  - Allows faster wall-clock analysis on datasets.
- bzip always slow: compressing and decompressing.

Network Telescope: Denial-of-Service Attacks

- Attacker floods the victim with requests using random spoofed source IP addresses
- Victim believes requests are legitimate and responds to each spoofed address
- With a /8 ("class A"), one can observe 1/256th of all *victim responses* to spoofed addresses

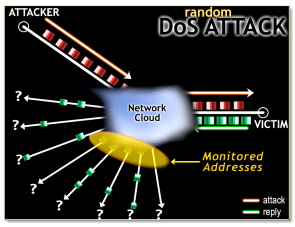

Assumptions and Biases

- *Address uniformity*
  - Ingress filtering, reflectors, etc. cause us to **underestimate** number of attacks
  - Can bias rate estimation (can we test uniformity?)
- *Reliable delivery*
  - Packet losses, server overload & rate limiting cause us to **underestimate** attack rates/durations
- *Backscatter hypothesis*
  - Can be biased by purposeful unsolicited packets
    - Port scanning (minor factor at worst in practice)
  - Can we verify backscatter at multiple sites?

Backscatter Hypothesis Busted?

- Not all TCP RST packets are DoS backscatter.
- Have seen a distributed a *scan* using TCP RST packets spread over more than a month
  - "random" /25s (128 victim IPs) at a time, from a ~100 hosts, looking for a couple specific ports. TTL is not low. Seen at more sites than our /8.
- What were they trying to find? Current best guess, looking for differential ICMP error responses.

DoS Attacks over time

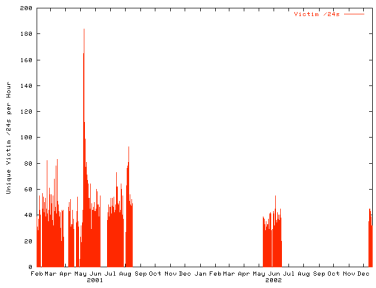

Our Telescope Data Analysis

- "Flow" based
  - Packets collected where possible, but most initial analysis is done with tools which work on flow-like aggregates.
- Eg, for backscatter
  - look at "outdegree" of victim IPs to telescope addresses

E.G. backscatter

- "Keys":
  - victimIP, protocols
- "Counters":
  - #pkts
  - #telescope IPs (also some distribution info)
  - #ports (also some distribution info) (for both src/dst)
  - are ports incrementing, decrementing (in little-endian byte order?)

Network Telescope: Worm Attacks

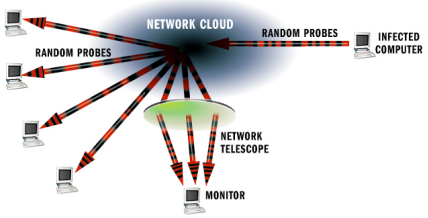

- Infected host scans for other vulnerable hosts by randomly generating IP addresses
- We monitor 1/256th of all IPv4 addresses
- We see 1/256th of all worm traffic of worms (when no bias or bugs)

Internet Worm Attacks: Code-Red

(July 19, 2001)

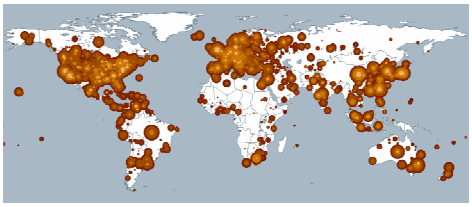

- 360,000 hosts infected in *ten hours*, 2,000 new per minute at peak
- No effective patching response
- More than $1.2 billion in economic damage in the first ten days
- Collateral damage: printers, routers, network traffic

Response to August 1st CodeRed

- CodeRed was programmed to deactivate on July 20th and begin spreading again on August 1st
- By July 30th and 31st, more news coverage than you can shake a stick at:
  - FBI/NIPC press release
  - Local ABC, CBS, NBC, FOX, WB, UPN coverage in many areas
  - National coverage on ABC, CBS, NBC, CNN
  - Printed/online news had been covering it since the 19th
- "Everyone" knew it was coming back on the 1st
- Best case for human response: known exploit with a viable patch and a known start date

Patching Survey

- How well did we respond to a best case scenario?
- Idea: randomly test subset of previously infected IP addresses to see if they have been patched or are still vulnerable
- 360,000 IP addresses in pool from initial July 19th infection
- 10,000 chosen randomly each day and surveyed between 9am and 5pm PDT

Patching Rate

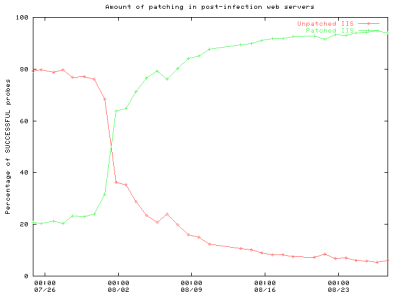

Dynamic IP Addresses

- How can we tell how when an IP address represents an infected **computer**?
- Resurgence of CodeRed on Aug 1st: Max of ~180,000 unique IPs seen in any 2 hour period, but more than 2 million across ~a week.
- This ***DHCP effect*** can produce skewed statistics for certain measures, especially over long time periods.
- Important to keep in mind if making big "bad lists".

Dynamic IP Addresses

- For each /24, count:
  - total number of unique IP addresses seen ever
  - maximum number seen in 2 hour periods
- On plot:
  - x-axis is total number of unique addresses seen ever
  - y-axis is maximum number for a 2 hour period
  - the x = y (total = max) line shows /24s that had all their vulnerable hosts actively spreading in same 2 hour period, and those hosts didn't change IP addresses
  - the space far below and to the right of the x = y line (total >> max) shows /24s that appear to have a lot of dynamic addresses
  - color of points represents density (3d histogram)

DHCP Effect seen in /24s

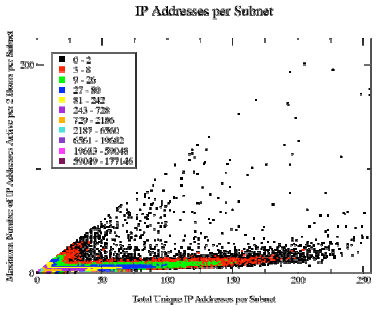

Internet Worm Attacks: Sapphire

(aka SQL Slammer) – Jan 24, 2003

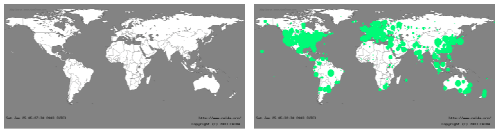

Before 9:30PM (PST) After 9:40PM (PST)

- ~100,000 hosts infected in *ten **minutes***
- Sent more than 55 million probes per second world wide
- Collateral damage: Bank of America ATMs, 911 disruptions, Continental Airlines cancelled flights
- Unstoppable; relatively benign to hosts

Spread of the Witty Worm

March 19, 2004

- First wide-spread Internet worm with destructive payload
  writes 64k blocks to disk at random location, repeatedly
- Launched from a large set of ground-zero hosts
  >100 hosts
- Shortest interval from vulnerability disclosure to worm release
  1 day
- Witty infected firewall/security software
  i.e. proactive user base
- Spread quickly even with a small population
  ~12,000 total hosts, 45 minutes to peak of infection

Early Growth of Witty

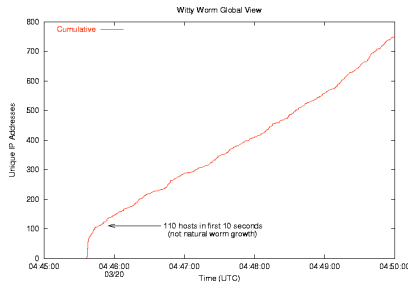

Geographic Spread of Witty

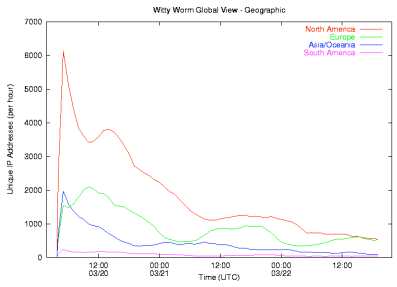

Passive Vulnerability Fingerprinting

- Really good idea, helps finds new devices/services on the network.
- Minor (?) downside: miss new services running which aren't actually used.
- Recent major downside: Miss many *passive* devices in the network.
  - transparent caches, proxies, BlackIce Defender…, AMP boxes

Conclusions

- Don't really have conclusions, but it seems like there are some good community opportunities out there. I don't see any transit ISP security folks here, some of them are currently using netflow (or passive devices).
- Watch out for DHCP effect.
- Watch out for passive devices in network.
- Academics generally don't understand operational needs, make lists.

Related CAIDA/UCSD Papers

- Inferring Internet Denial-of-Service Activity [MSV01]
  - David Moore, Stefan Savage, Geoff Voelker
  - http://www.caida.org/outreach/papers/2001/BackScatter/
- Code-Red: A Case Study on the spread and victims of an Internet Worm [MSB02]
  - David Moore, Colleen Shannon, Jeffrey Brown
  - http://www.caida.org/outreach/papers/2002/codered/
- Internet Quarantine: Requirements for Containing Self-Propagating Code [MSVS03]
  - David Moore, Colleen Shannon, Geoff Voelker, Stefan Savage
  - http://www.caida.org/outreach/papers/2003/quarantine/
- The Spread of the Sapphire/Slammer Worm [MPS03]
  - David Moore, Vern Paxson, Stefan Savage, Colleen Shannon, Stuart Staniford, Nicholas Weaver
  - http://www.caida.org/outreach/papers/2003/sapphire/

Additional CAIDA/UCSD Information

- Code-Red v1, Code-Red v2, CodeRedII, Nimda
  - http://www.caida.org/analysis/security/code-red/
- Code-Red v2 In-depth analysis
  - http://www.caida.org/analysis/security/code-red/coderedv2_analysis.xml
- Spread of the Sapphire/SQL Slammer Worm
  - http://www.caida.org/analysis/security/sapphire/
- Network telescopes
  - http://www.caida.org/analysis/security/telescope/

Using your own telescope: Effects of Size

- Larger telescopes are able to detect events that generate fewer packets, either because of short duration or low sending rate.
- Larger telescopes have better accuracy at determining the start and end times of an event.
- Using CIDR / notation on next few slides:
  - /8 = old class-A size, 16 million IP addresses
  - /16 = old class-B size, 65536 IP addresses

Detectable Events (95%)

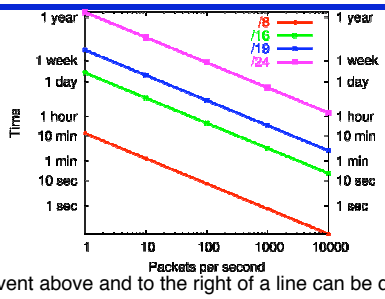

Any event above and to the right of a line can be detected (at least one packet seen) with at least 95% probability.

Detection Times - 10 pps events

(Code-Red approx. this rate)

Detection probability:	5%	50%	95%
/8	1.3 sec	18 sec	1.3 min
/14	1.4 min	19 min	1.4 hour
/15	3 min	38 min	2.7 hour
/16	6 min	1.3 hour	5.5 hour
/19	45 min	10 hour	1.8 day
/24	24 hours	14 day	58 day

Worm Spread – 10 probes/sec

(Code-Red approx. this rate)

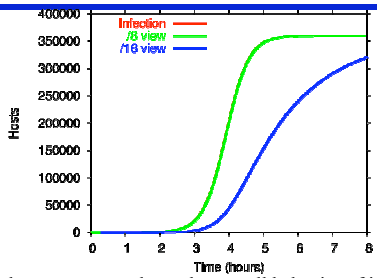

- /8 telescope accurately tracks overall behavior of infection
- /16 telescope lags behind in time and shape is misleading

Worm Spread – 10 probes/sec

(Code-Red approx. this rate)

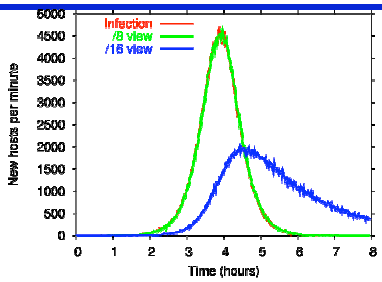

Smaller network telescopes can't accurately determine event start times (e.g. when a particular host is infected).

Organizational Telescopes

- Small telescopes may not be useful for observing external events
- However, setting up an internal facing telescope may help quickly identify internal problems
- With an internal facing telescope you can have /5 or better

Why have an internal telescope?

- Quickly detect internal machines infected with worms, certain kinds of misconfigurations, and potentially hacked machines.
- Capture data for hosts connecting to unallocated IP address space by:
  - if you use BGP (default-free) to all providers, you can point a default route at a monitor box
  - enable flow collection on your edge routers
  - announce a couple unallocated networks, but be careful if they ever get allocated by IANA (least desirable)

Extending it

- Combine a telescope watching traffic to unallocated IP addresses with monitoring all outbound traffic
  - you may notice anomalous behavior like a spam relay
  - verify that your firewall seems to be doing what you think
- Watch all *inbound* ICMP error messages, in particular HOST/NETWORK UNREACHABLE
  - evidence of scanning behavior
  - may show external connectivity & performance problems before users pick up the telephone

Tools to use

- Flow data (Cisco NetFlow, Juniper cflow, others):
  - FlowScan: http://net.doit.wisc.edu/~plonka/FlowScan
- Packet data
  - CoralReef report generator: http://www.caida.org/tools/
- Either
  - AutoFocus: http://ial.ucsd.edu/AutoFocus/
- Not an exhaustive list ☺

AutoFocus example

- Sapphire/SQL Slammer worm
  - Find worm port & proto automatically

Source IP	Destination IP	Protocol	Source Port	Destination Port	bytes	Unexpectedness(%)
*	*	6	highports	highports	827M	77.7
*	*	17	highports	1434	10.5G	112.6
*	152.249.0.0/16	*	*	*	604M	100
138.0.0.0/9	*	*	*	highports	3.66G	99.4
138.0.0.0/10	*	*	highports	*	3.68G	99.9
138.54.3.58	*	17	3341	1434	2.14G	672.5
138.54.11.4	*	17	7062	1434	950M	1551.3
152.249.56.0/22	*	*	highports	highports	723M	103.4
152.249.191.120	*	17	1959	1434	1.78G	810.0
152.249.191.121	96.0.0.0/8	17	1531	1434	645M	39523.7
152.249.210.3	*	17	4315	1434	2.36G	609.5
152.249.254.152	*	17	3787	1434	1.53G	941.8

AutoFocus example

- Sapphire/SQL Slammer worm
  - Can identify infected hosts

Source IP	Destination IP	Protocol	Source Port	Destination Port	bytes	Unexpectedness(%)
*	*	6	highports	highports	827M	77.7
*	*	17	highports	1434	10.5G	112.6
*	152.249.0.0/16	*	*	*	604M	100
138.0.0.0/9	*	*	*	highports	3.66G	99.4
138.0.0.0/10	*	*	highports	*	3.68G	99.9
138.54.3.58	*	17	3341	1434	2.14G	672.5
138.54.11.4	*	17	7062	1434	950M	1551.3
152.249.56.0/22	*	*	highports	highports	723M	103.4
152.249.191.120	*	17	1959	1434	1.78G	810.0
152.249.191.121	96.0.0.0/8	17	1531	1434	645M	39523.7
152.249.210.3	*	17	4315	1434	2.36G	609.5
152.249.254.152	*	17	3787	1434	1.53G	941.8

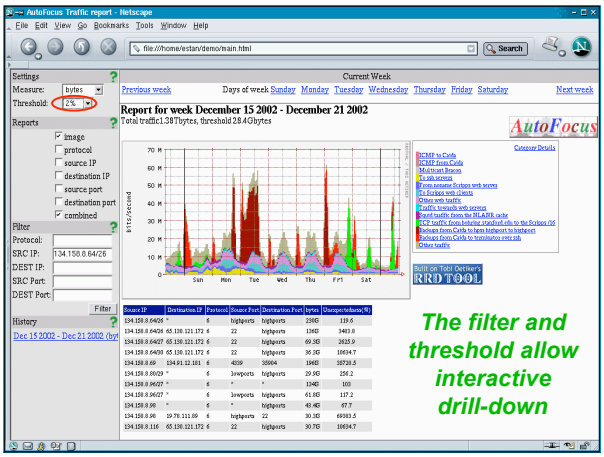

The filter and threshold allow interactive drill-down

Conclusions

- Network telescopes provide insight into non-local network events
- Larger telescopes better capture the behavior of events and can see smaller events
- Build your own internal telescope – it's fun AND easy.