Inferring Internet Denial-of-Service Activity

Geoffrey M. Voelker

University of California, San Diego

Joint work with David Moore (CAIDA/UCSD)
and Stefan Savage (UCSD)

Simple Question

We were interested in answering a simple question:

How prevalent are
denial-of-service attacks in the Internet?

Anecdotal Data

Press reports:

Analysts: "Losses … could total more than $1.2 billion"
- *Yankee Group* report

Surveys: "38% of security professionals surveyed reported denial of service activity in 2000"
- *CSI/FBI* survey

Quantitative Data?

- Is not available (i.e., no one knows)
- Inherently hard to acquire
  - Few content or service providers collect such data
  - If they do, its usually considered sensitive
- Infeasible to collect at Internet scale
  - How can you monitor enough of the Internet to obtain a representative sample?

Our Contributions

- Backscatter analysis
  - New technique for estimating global denial-of-service activity
- First data describing Internet-wide DoS activity
  - ~4,000 attacks per week (> 12,000 over 3 weeks)
  - Instantaneous loads above 600k pps
  - Characterization of attacks and victims

- Paper appeared this August:
  - Moore, Voelker and Savage, *Inferring Internet Denial-of-Service Activity*, 2001 USENIX Security

Overview

- Describe backscatter analysis
- Experimental setup
- Series of analyses and attack characterizations
- Tracking the Code Red Worm

Key Idea

- Flooding-style DoS attacks
  - e.g. SYN flood, ICMP flood
- Attackers spoof source address randomly
  - True of all major attack tools
- Victims, in turn, respond to attack packets
- Unsolicited responses (*backscatter*) equally distributed across IP space
- Received backscatter is evidence of an attacker elsewhere

Backscatter Example

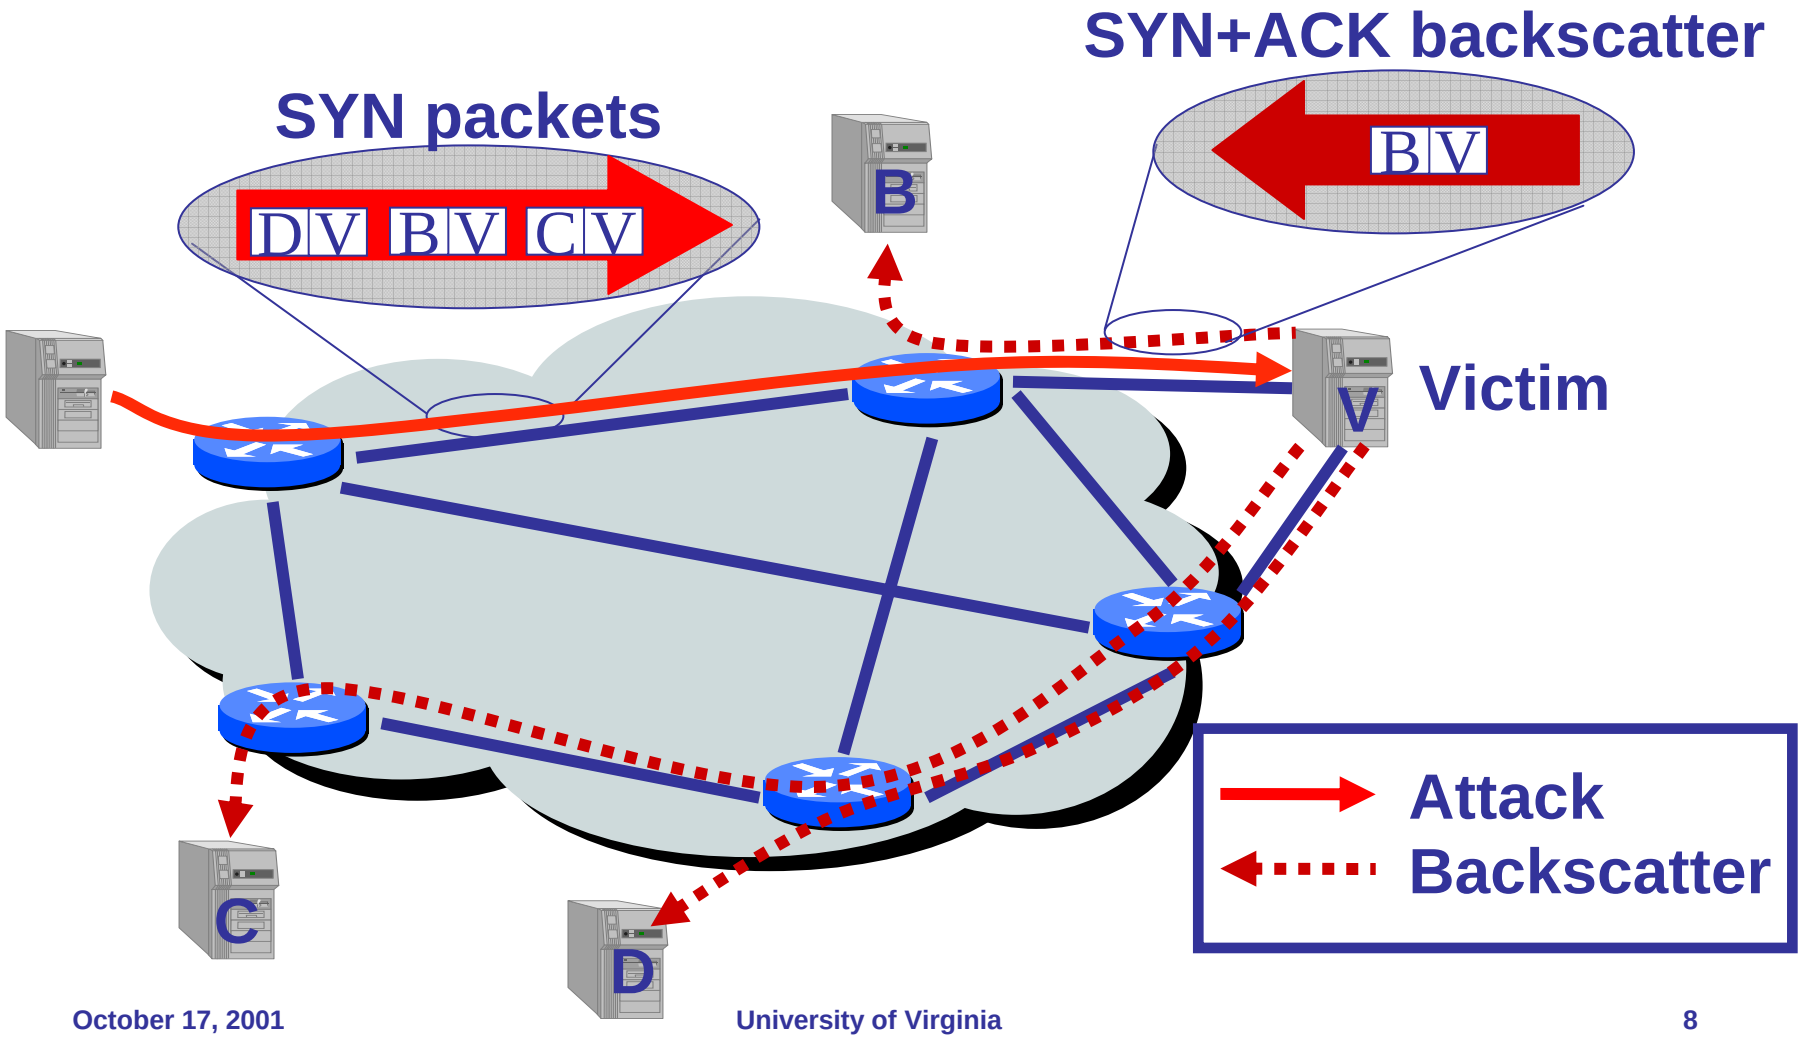

Backscatter Analysis

- Monitor block of *n* IP addresses
- Expected # of backscatter packets given an attack of *m* packets:

$$E(X) = \frac{nm}{2^{32}}$$

- Extrapolated attack rate R is a function of measured backscatter rate R':

$$R \geq R' \frac{2^{32}}{n}$$

Assumptions and Biases

- *Address uniformity*
  - Ingress filtering, reflectors, etc. cause us to underestimate # 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)
  - Do we detect backscatter at multiple sites?

Experimental Setup

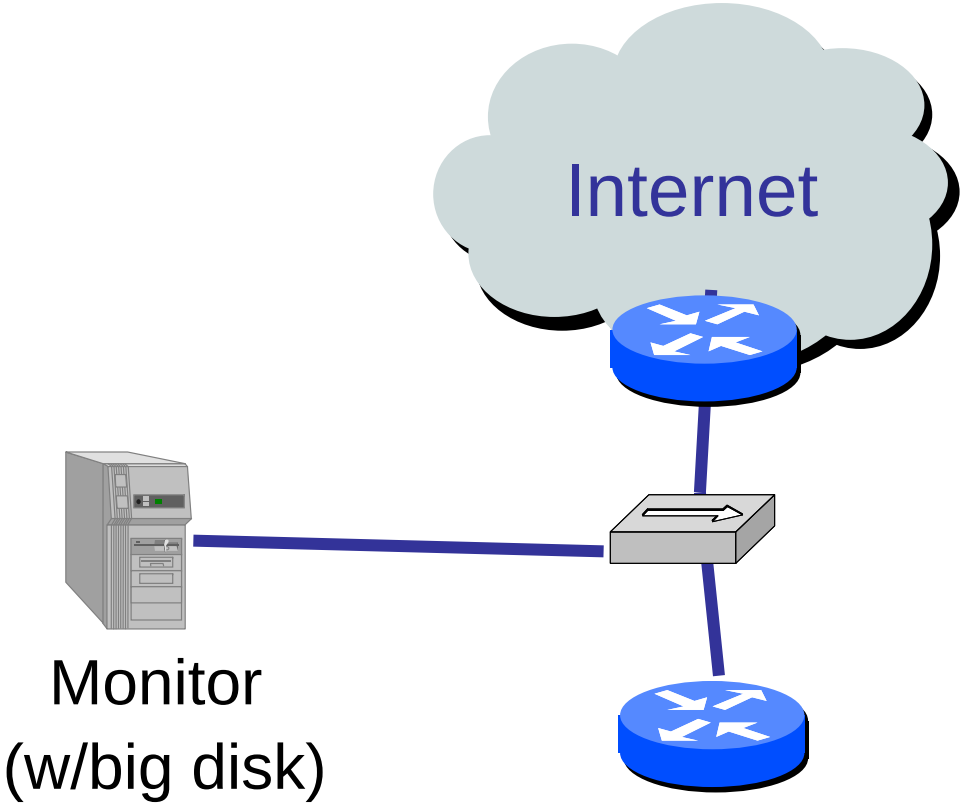

Methodology

- Collected three weeks of traces (February 2001)
- Analyzed trace data from two perspectives
- Flow-based analysis (categorical)
  - Number, duration, kinds of attacks
  - Keyed on victim IP address and protocol
  - Flow duration defined by explicit parameters (min threshold, timeout)
- Event-based analysis (intensity)
  - Rate, intensity over time
  - Attack event: backscatter packets from IP address in 1 minute window

Analysis

- Summary statistics
- Time behavior
- Protocol
- Duration
- Rate
- Victim categorization
  - DNS, top-level domain (TLD), AS
  - Popularity

Attack Breakdown

	Week1	Week2	Week3
Attacks	4173	3878	4754
Victim IP's	1942	1821	2385
Victim prefixes	1132	1085	1281
Victim AS's	585	575	677
Victim DNS domains	750	693	876
Victim DNS TLDs	60	62	71

Attacks Over Time

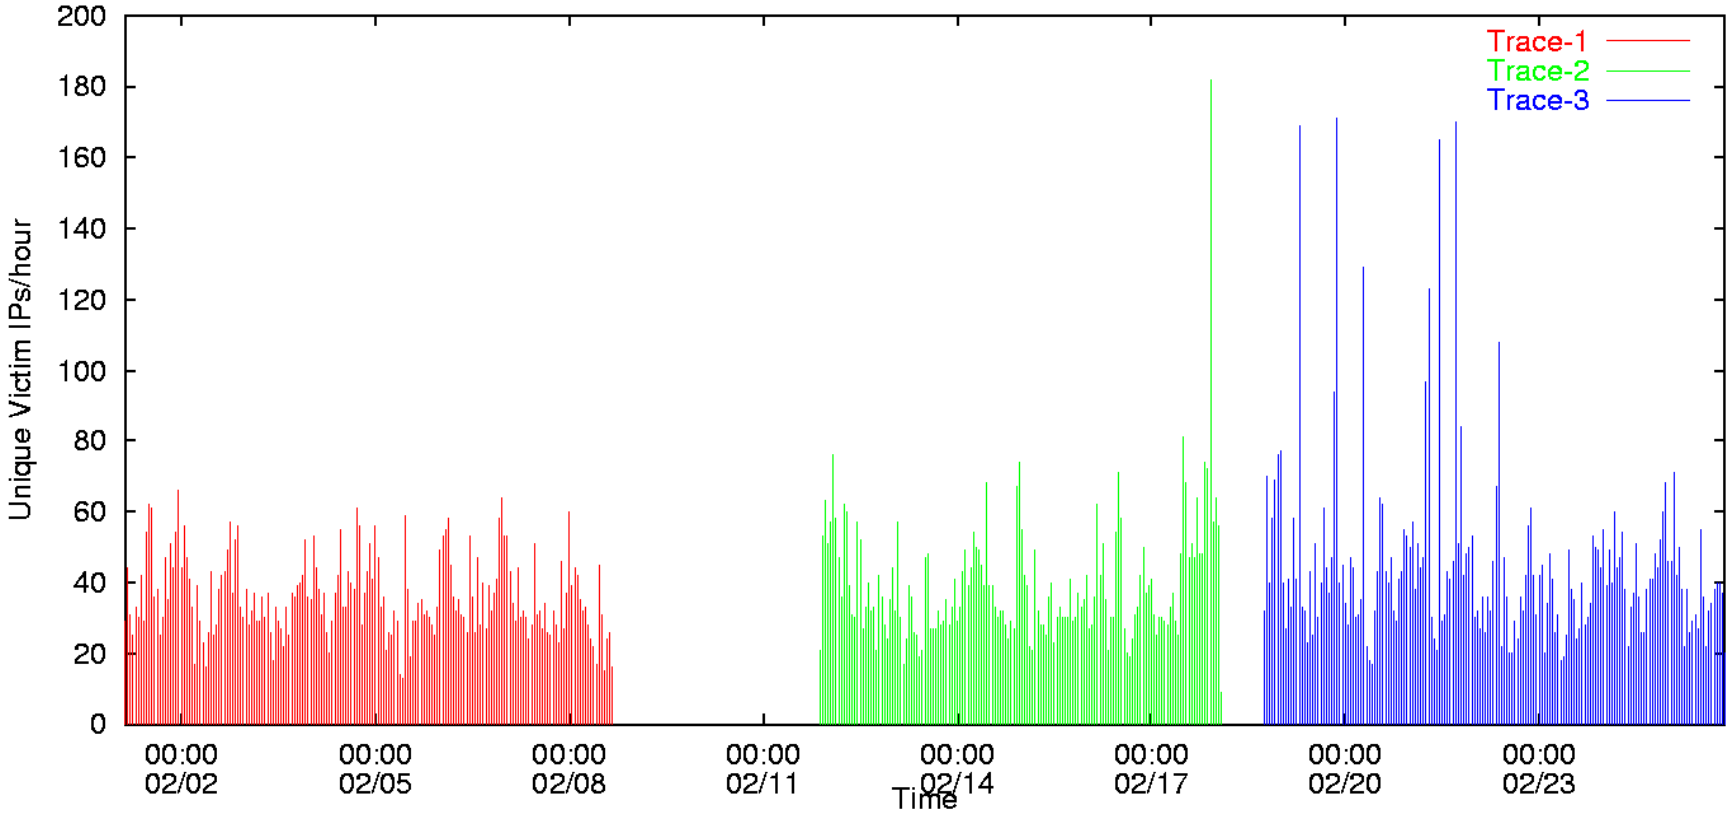

(Surprisingly uniform, no diurnal effects)

Periodic Attack (Daily)

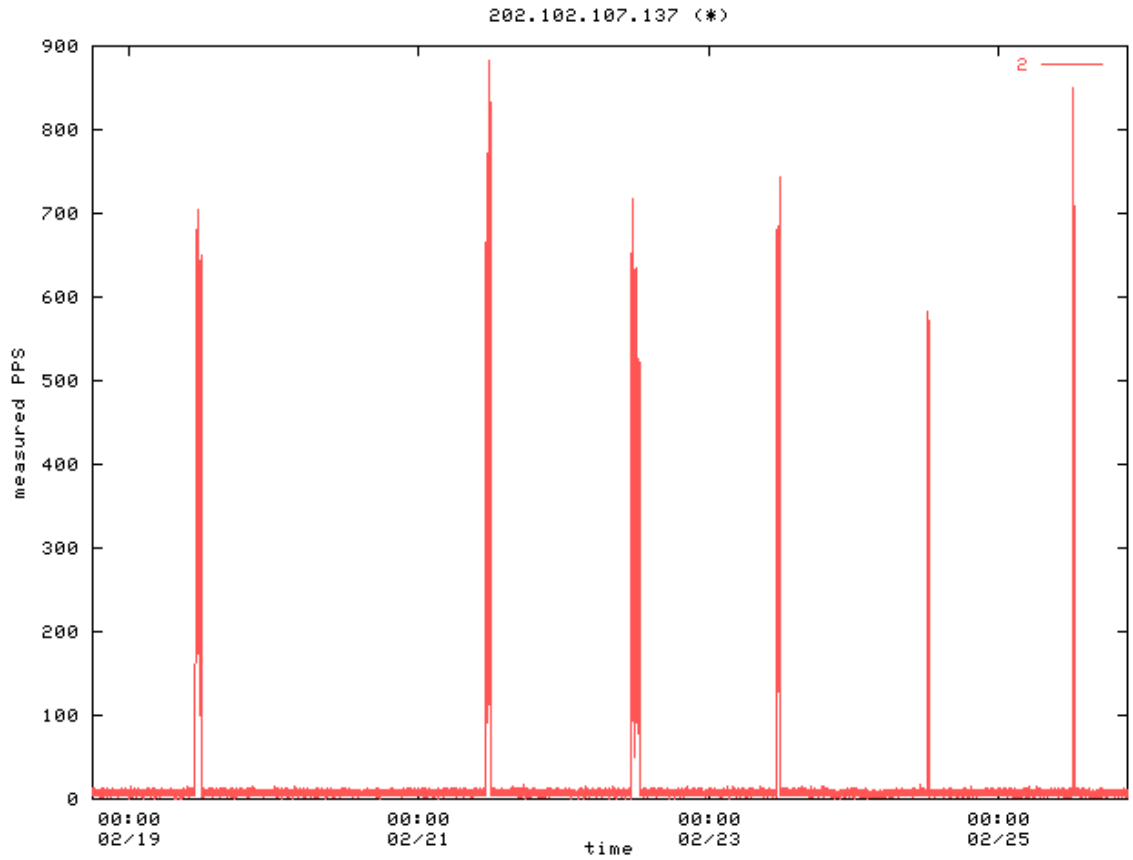

(Every day like clockwork)

Punctuated Attack (1 min)

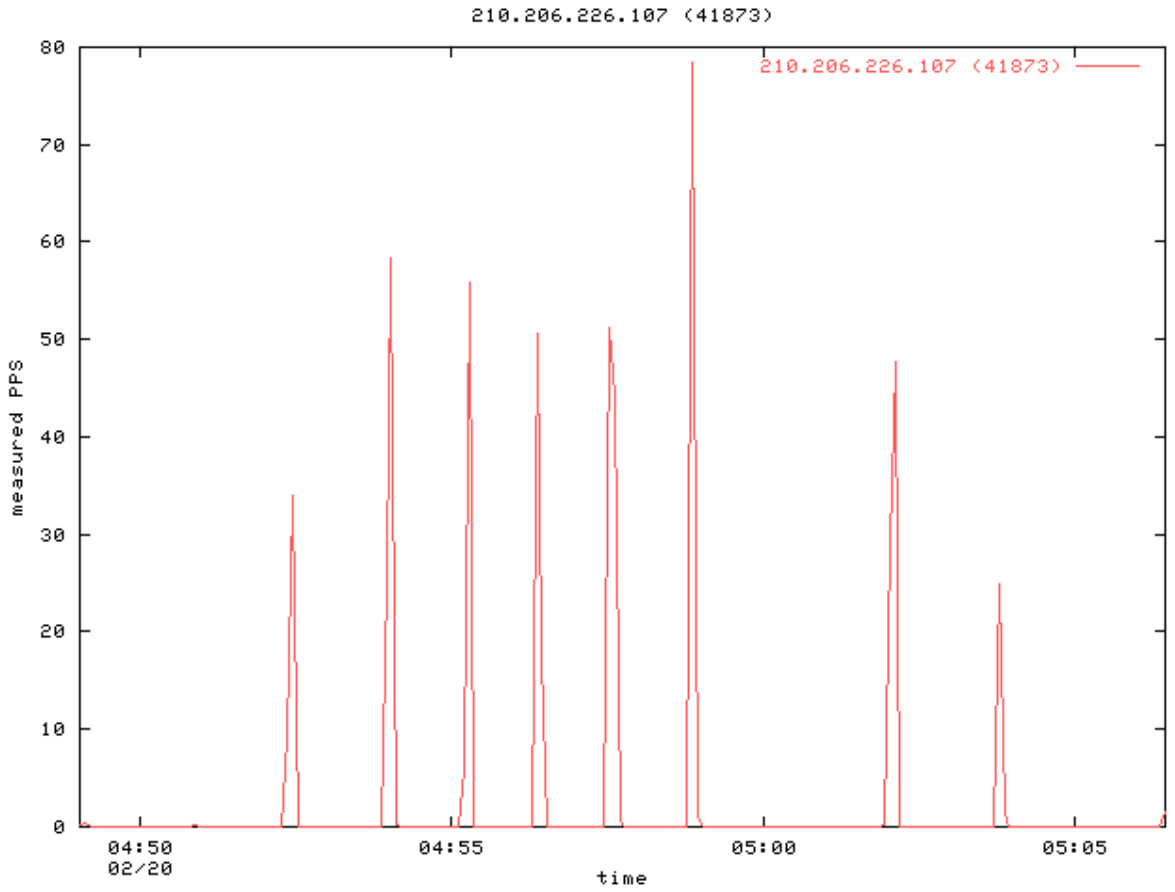

(Fine-grained behavior as well)

Attack Protocol/Services

- Protocols
  - Mostly TCP (90-94% attacks)
  - A few large ICMP floods (up to 43% of packets)
- Services
  - Most attacks on multiple ports (~80%)
  - A few services (HTTP, IRC) singled out

Attack Duration

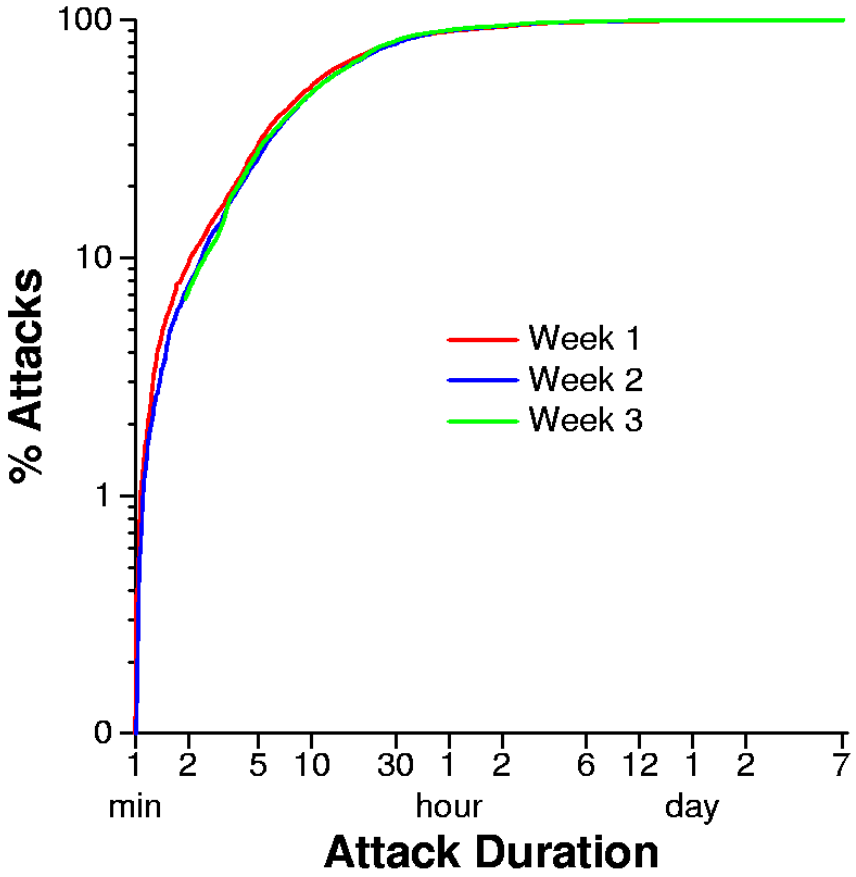

(50% > 10 mins)

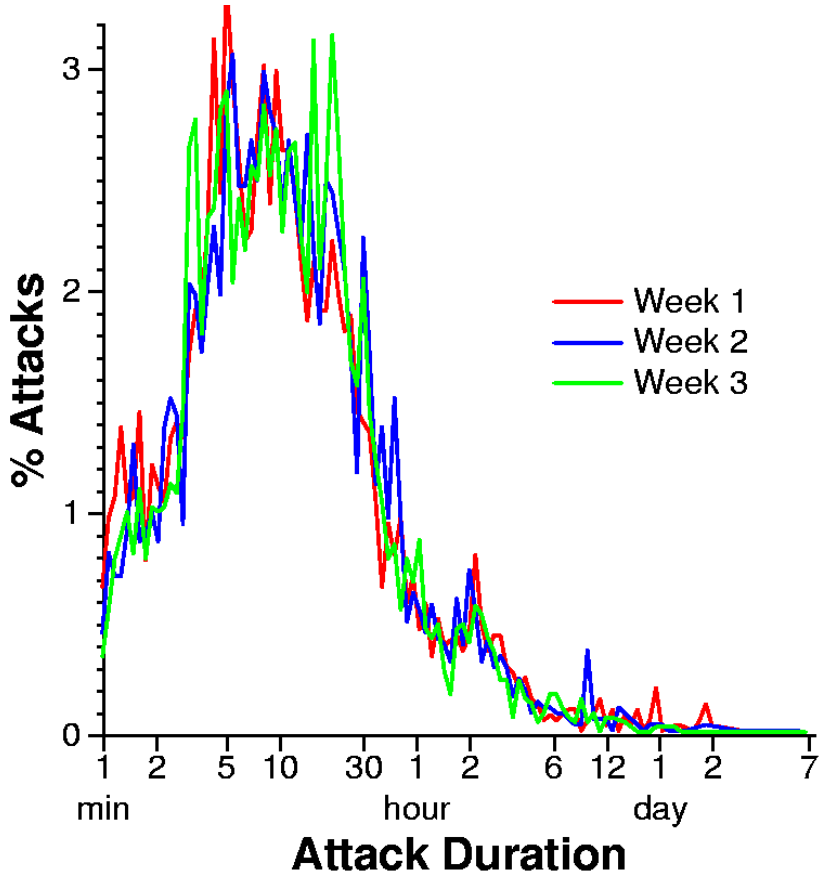

(Most between 3-30 mins)

Attack Rate

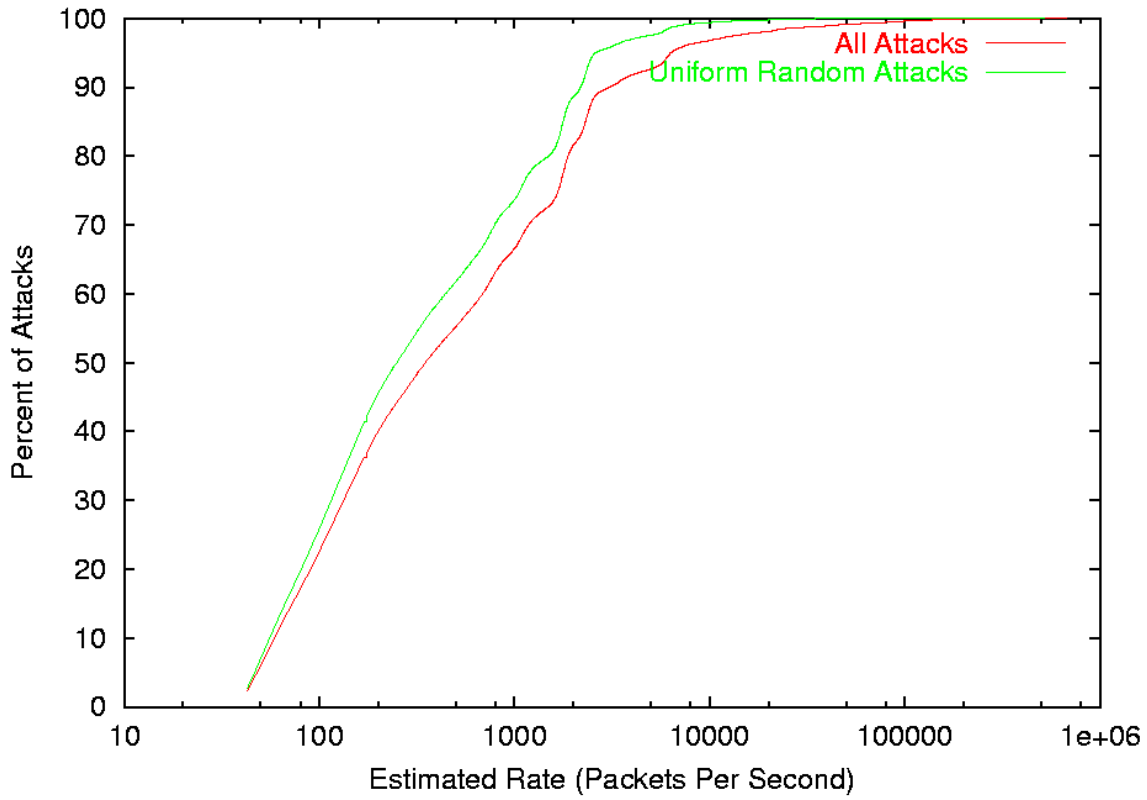

(50% > 350 pps/sec, most intense is 679,000 pps)

Victim Characterization (DNS)

- Entire spectrum of commercial businesses
  - Yahoo, CNN, Amazon, etc. and many smaller biz
- Evidence that minor DoS attacks used for personal vendettas
  - 10-20% of attacks to home machines
  - A few very large attacks against broadband
  - Many reverse mappings clearly compromised (e.g. is.on.the.net.illegal.ly and the.feds.cant.secure.their.shellz.ca)
- 5% of attack target infrastructure
  - Routers (e.g. core2-core1-oc48.paol.above.net)
  - Name servers (e.g. ns4.reliablehosting.com)

Victim Top-Level Domains

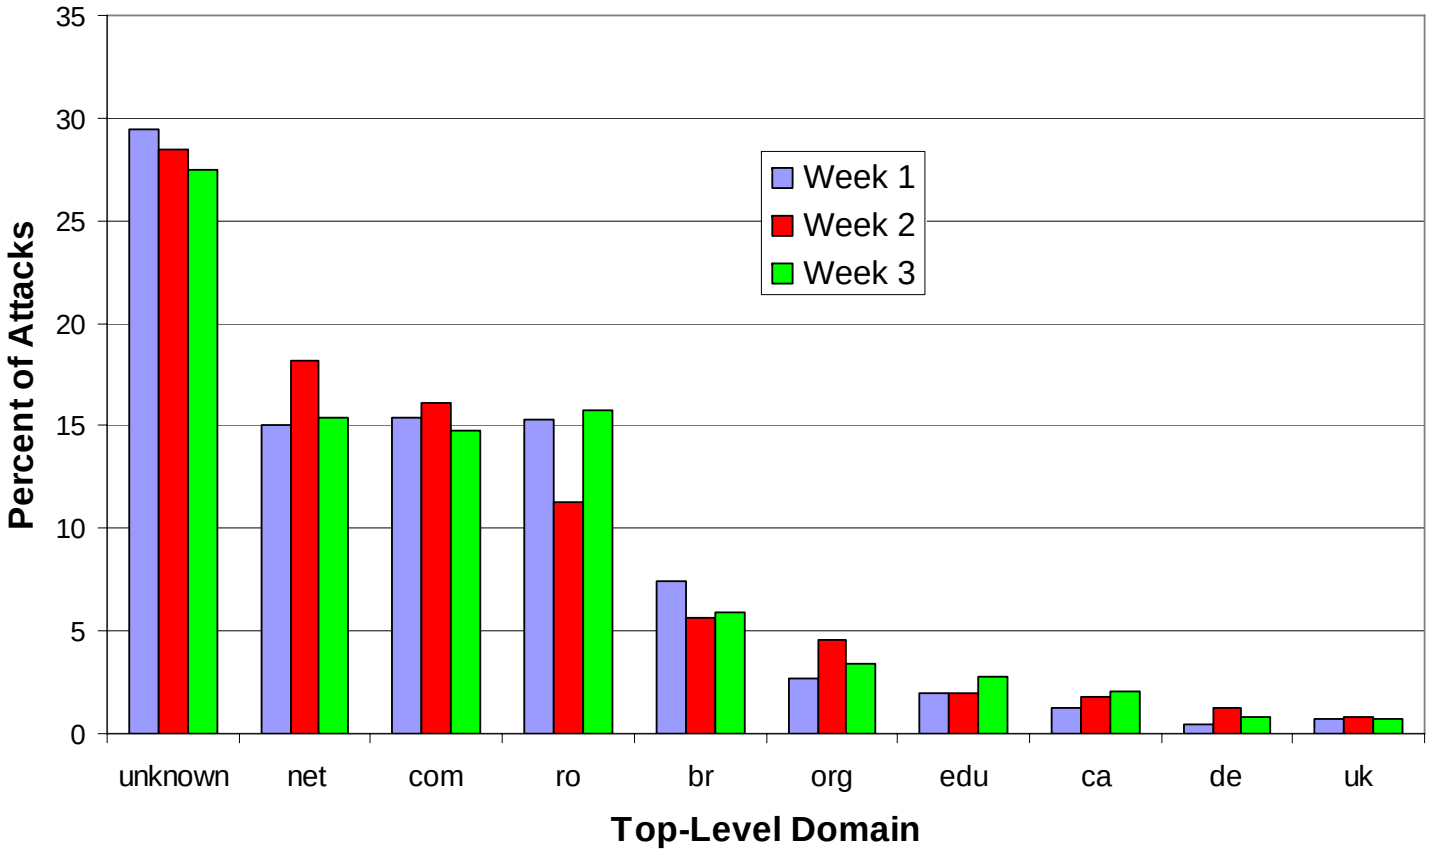

(net == com, edu small, ro and br unusual)

Victim Autonomous Systems

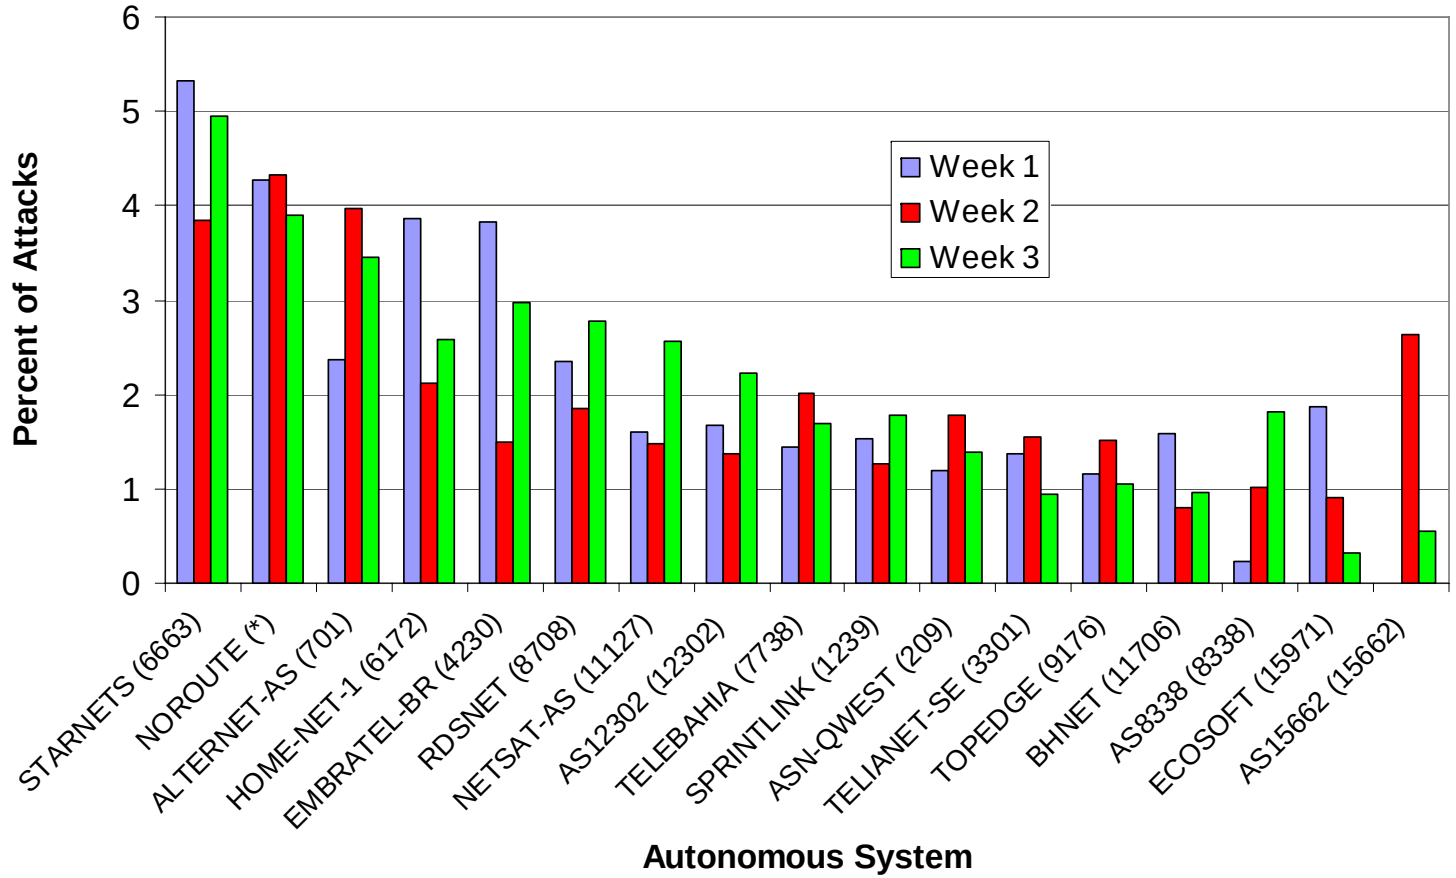

(No single AS/set of AS's are targeted (long tail, too))

Victim Popularity

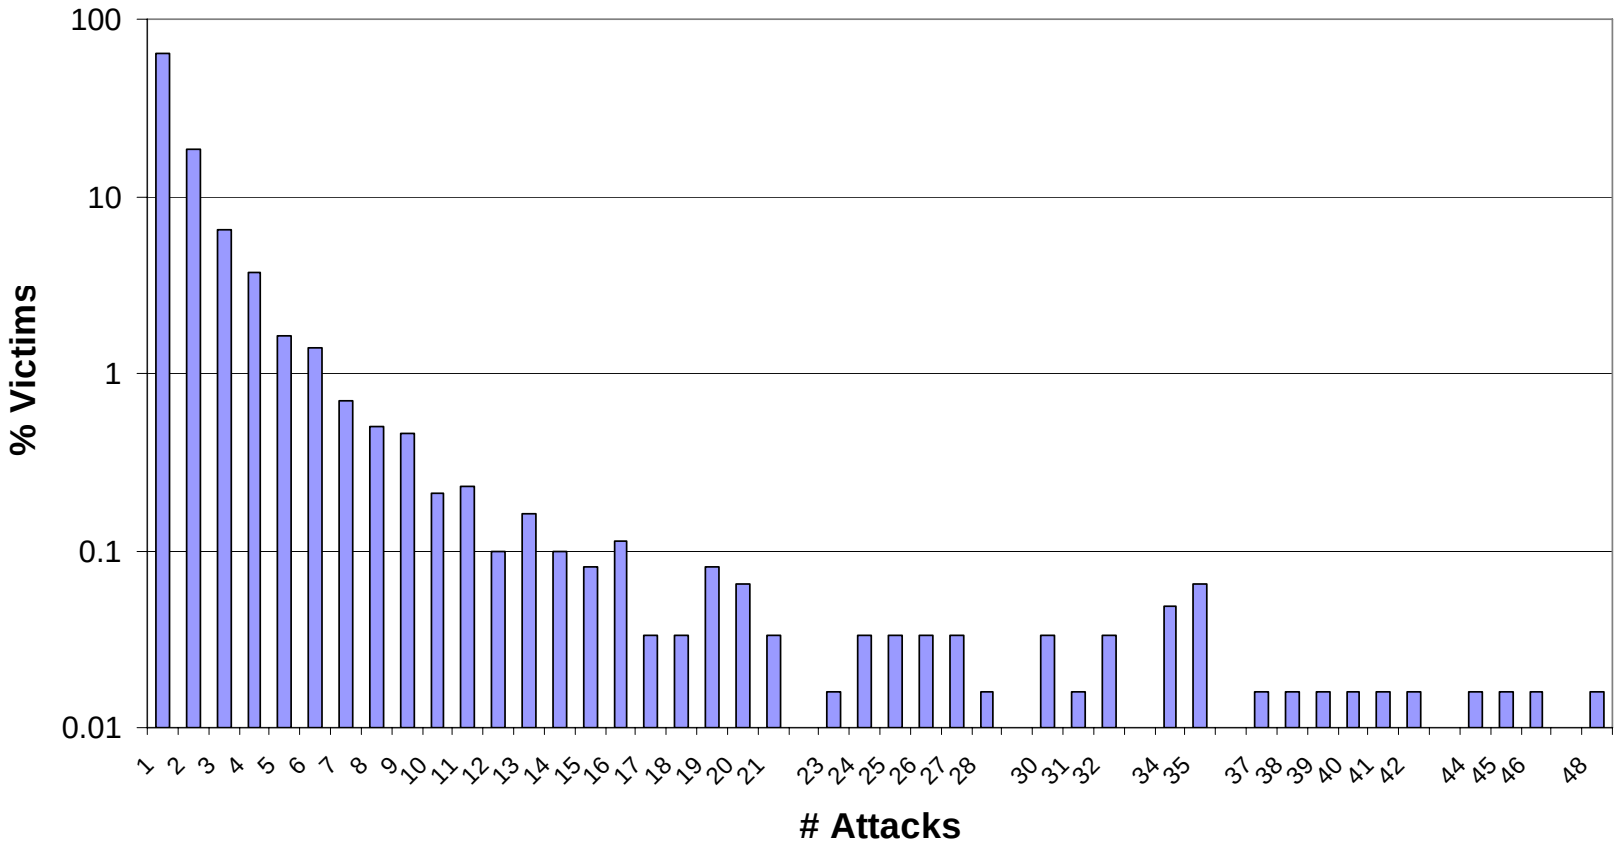

(Most victims attacked once, but a few are unfortunate favorites)

Validation

- How do we know we are seeing backscatter from attacks, and not just funky traffic to our network?
- Backscatter not explained by port scanning
  - 98% of backscatter packets do not cause response
- Repeated experiment with independent monitor (3 /16's from Vern Paxson)
  - Only captured TCP SYN/ACK backscatter
  - 98% inclusion into larger dataset
- Matched to actual attacks detected by Asta Networks on large backbone network

Summary

- Lots of attacks – some very large
  - >12,000 attacks against >5,000 targets in a week
  - Most < 1,000 pps, but some over 600,000 pps
- Everyone is a potential target
  - Targets not dominated by any TLD, 2LD or AS
    - » Targets include large e-commerce sites, mid-sized business, ISPs, government, universities and end-users
  - Something weird is happening in Romania
- New attack “styles”
  - Punctuated/periodic attacks
  - Attacks against infrastructure targets & broadband

Code Red Worm

- In July, David Moore used the same technique to track the Code Red Worm
  - While collecting backscatter data (no way to predict)
- Code Red
  - Infects MS IIS Web servers via security hole
  - Once infected, victim tries to infect other hosts
  - Culminates in a coordinated attack against whitehouse.gov
- Impact
  - Tremendous amount of popular press
    - » FBI warning on second round of Code Red Worm

Monitoring Code Red

- Victims randomly choose an IP address to infect
  - Try to establish a HTTP connection to that address
  - 1/256th of connection requests in our /8 (our looking glass)
  - Easy to distinguish from backscatter
- As with backscatter, can determine
  - Who: Set of IP addresses of victims infected
    - » Breakdown by DNS, TLD, AS, etc.
  - Infection rate: Real-time spread of worm across Internet
  - Patch rate: Real-time patching, shutdown of infected hosts

Rate of Infection

359,104 hosts were compromised in approximately 13 hrs.

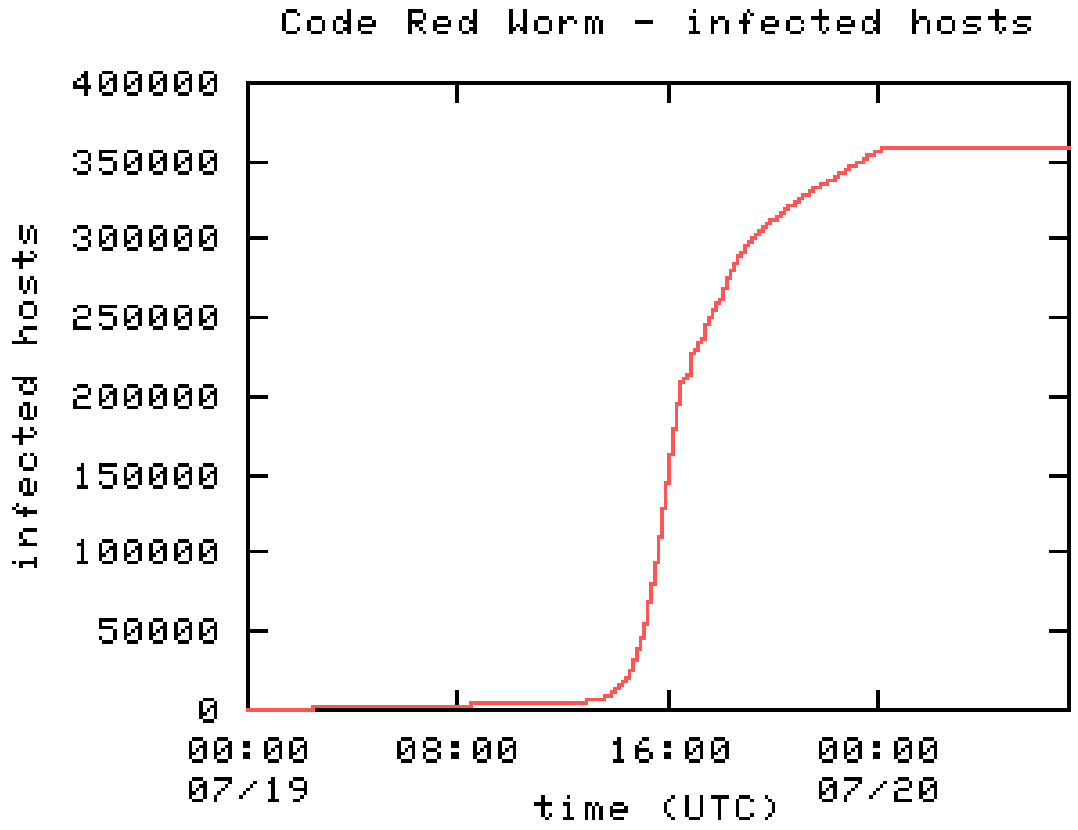

More Info

- Backscatter
  http://www.caida.org/outreach/papers/backscatter/
- Code Red
  http://www.caida.org/analysis/security/code-red/

Protocol Breakdown (1 week)

Backscatter protocol	Attacks	BS Packets (x1000)
TCP (RST ACK)	2027 (49)	12,656 (25)
ICMP (Host Unreachable)	699 (17)	2892 (5.7)
ICMP (TTL Exceeded)	453 (11)	31468 (62)
ICMP (Other)	486 (12)	580 (1.1)
TCP (SYN ACK)	378 (9.1)	919 (1.8)
TCP (RST)	128 (3.1)	2,309 (4.5)
TCP (Other)	2 (0.05)	3 (0.01)

Attack Protocol Breakdown

Attack Protocol	Attacks	BS Packets (x1000)
TCP	3902 (**94**)	28705 (56)
UDP	99 (2.4)	66 (0.13)
ICMP	88 (2.1)	22,020 (**43**)
Proto 0	65 (1.6)	25 (0.05)
Other	19 (0.46)	12 (0.02)