Inferring Internet Denial-of-Service Activity

David Moore, Colleen Shannon, Douglas J. Brown, Geoffrey M. Voelker, Stefan Savage

Presented by
Thangam Seenivasan & Rabin Karki

Simple Question

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

Why is it important?

Loss could total more than $1.2 billion

-analysts

DDOS attacks have become common

Borrowed from G.Voelkar’s presentation

Recent DDOS attack

WIRED SUBSCRIBE » SECTIONS » BLOGS » REVIEWS » VIDEO » HOW

Sign In | RSS Fee

EPICENTER

MIND OUR TECH BUSINESS

Denial-of-Service Attack Knocks Twitter Offline (Updated)

By Eliot Van Buskirk August 6, 2009 | 10:06 am | Categories: Social Media

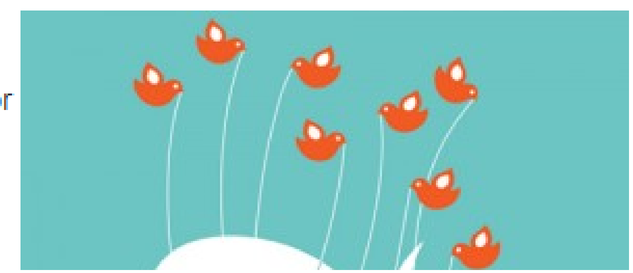

Twitter was shut down for hours Thursday morning by what it described as an "ongoing" denial-of-service attack, silencing millions of Tweeters. It was the first major outage the service has suffered in months and possibly the first ever due to sabotage. The outage appeared to begin mid-morning, EST, and affected users around the world. After about three hours, the service was coming

Challenges

- No quantitative data available about the prevalence of DOS attacks
- Obstacles gathering DOS traffic data
  - ISP consider such data private and sensitive
  - Need to monitored from a large number of sites to obtain representative data

Solution

- Backscatter Analysis
  - Estimate prevalence of worldwide DOS attacks
  - Traffic monitoring technique
  - Conservative estimate on the prevalence
  - Lower bound on the intensity of attacks

Outline

- Background
- Methodology
- Attack detection and classification
- Analysis of DOS

DOS attacks

- An attempt to make a computer resource unavailable to its intended users
- Classes of attacks
  - Logic attacks (exploits software flaws)
    - Ping-of-Death
  - Resource attacks
    - Sending a large number of spurious requests

This paper focuses only on resource attacks

Resource attacks

- Network
  - Overwhelm the capacity of network devices
  - Attacker sends packets as rapidly as possible
- CPU
  - Load the CPU by requiring additional processing
  - SYN flood
    - For each SYN packet to a listening TCP port
      - The host must search through existing connections
      - Allocate new data structures
    - Even a small SYN flood can overwhelm a remote host

Distributed attacks

- More powerful attacks
  - From multiple hosts

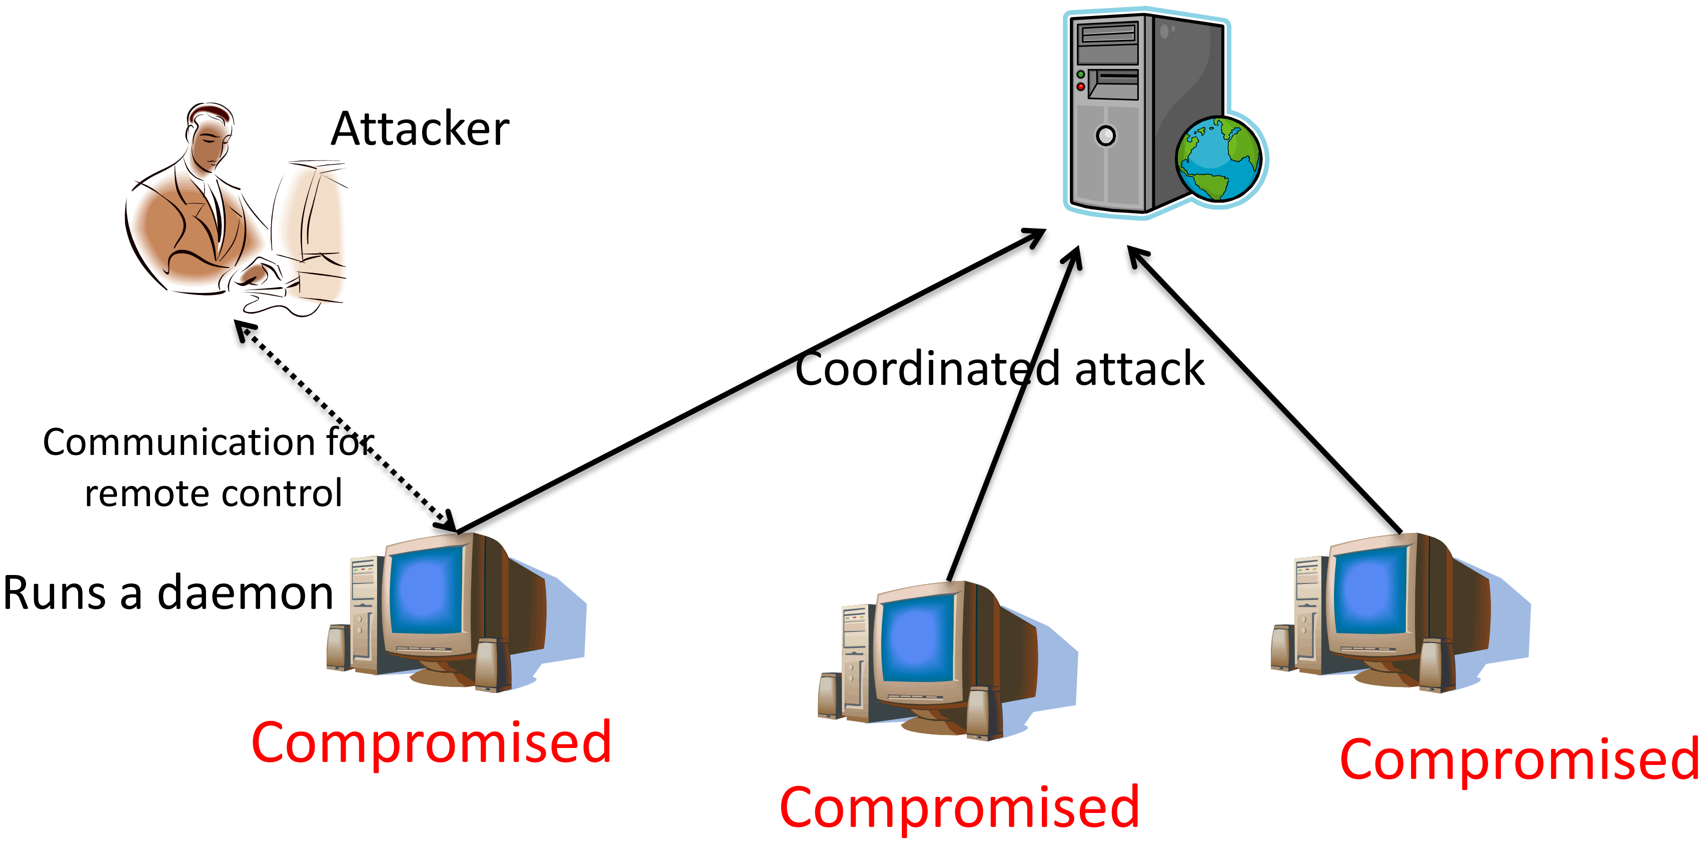

IP Spoofing

- Many attackers spoof IP source address
  – To conceal their locations
- Use random address spoofing
  – To overcome blacklisting/filtering

This paper focuses solely on attacks with random address spoofing

Outline

- Background
- Methodology
- Attack detection and classification
- Analysis of DOS

Key Idea

- Attackers spoof source address randomly
- Victim, in turn respond to attack packets
- Unsolicited responses (backscatter) equally distributed across IP address space
- Received backscatter is evidence of an attacker elsewhere

Backscattering

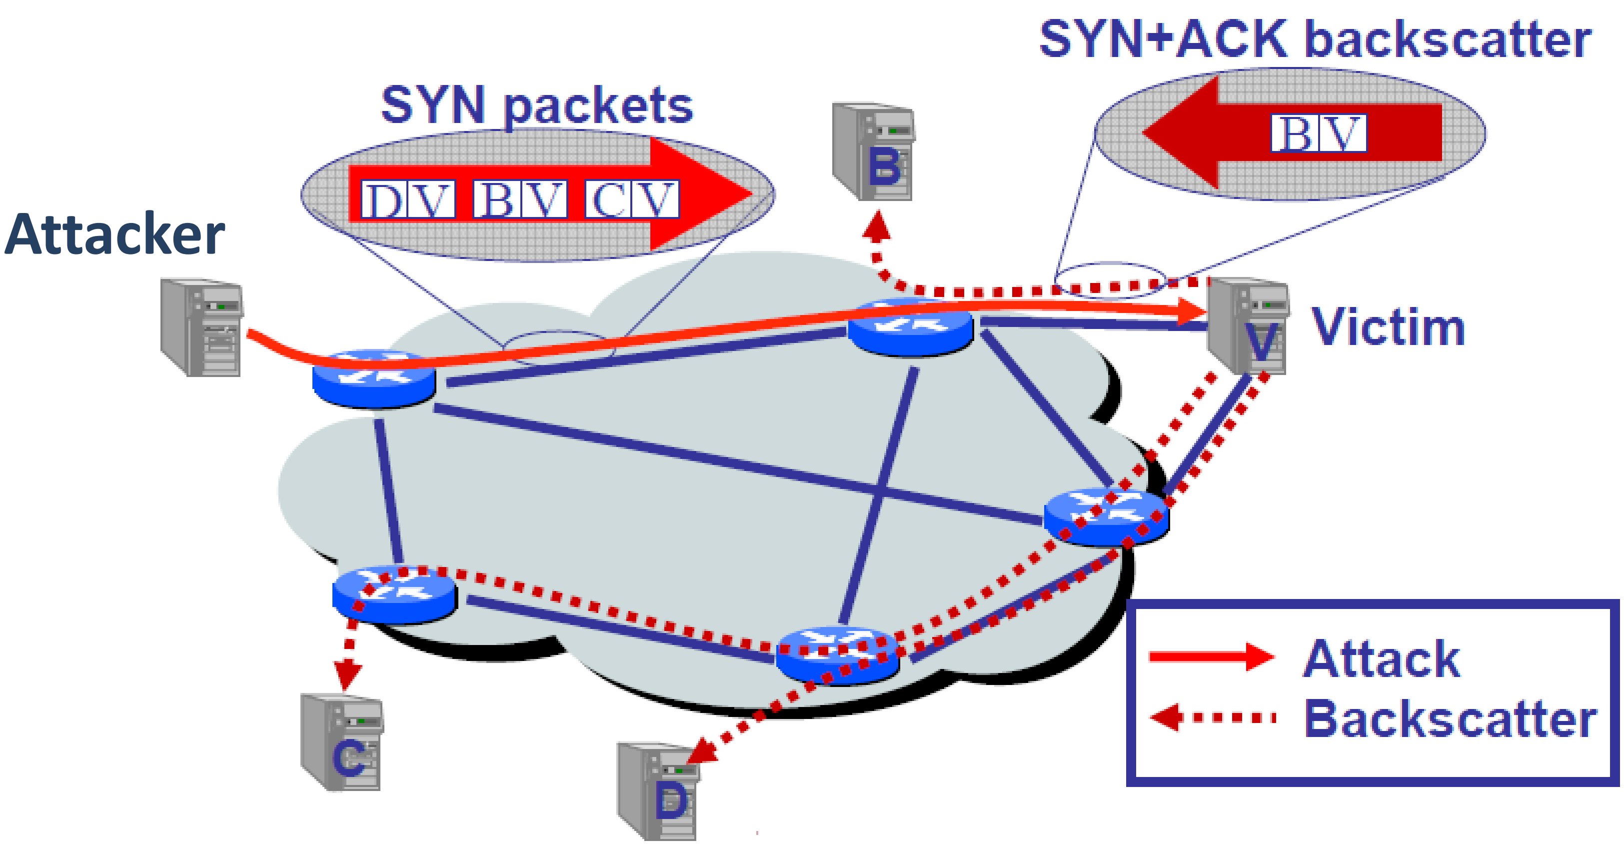

Borrowed from G.Voelkar's presentation

Typical victim responses

Packet Sent	Response from Victim
TCP SYN (to open port)	TCP SYN/ACK
TCP SYN (to closed port)	TCP RST (ACK)
TCP ACK	TCP RST (ACK)
TCP DATA	TCP RST (ACK)
TCP RST	no response
TCP NULL	TCP RST (ACK)
ICMP ECHO Request	ICMP Echo Reply
ICMP TS Request	ICMP TS Reply
UDP pkt (to open port)	protocol dependent
UDP pkt (to closed port)	ICMP Port Unreach
...	...

Backscatter Analysis

- Probability of one given host on the Internet receiving at least one unsolicited response during an attack of m packets

$$1 - \left(1 - \frac{1}{2^{32}}\right)^m$$

- Probability of n hosts receiving at least one of m packets

$$1 - \left(1 - \frac{n}{2^{32}}\right)^m$$

Backscatter Analysis

- Monitor from n distinct hosts
- Expected number of backscatter packets given an attack of m packets

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

- These samples contain
  - Identity of the victim
  - Timestamp
  - Kind of attack

Backscatter Analysis

- If arrival rate of unsolicited packets from a victim is R'
- Extrapolated attack rate R on the victim is

$$R \geq R' \frac{2^{32}}{n} \quad \textit{packets per sec}$$

Assumptions

- Address uniformity
  - attackers spoof source addresses at random
- Reliable delivery
  - Attack traffic and backscatter is delivered reliably
- Backscatter hypothesis
  - Unsolicited packets observed by the monitor represent backscatter

Limitation - Address uniformity

- Many attacks do not use address spoofing
  - ISPs increasingly employ ingress filtering
- "Reflector attacks"
  - Source address is specifically selected
- Motivation for IP spoofing has been reduced
  - Automated methods for compromising host
  - DDOS attacks using true IP addresses

Each factor cause the analysis to underestimate the total number of attacks

Limitation – Reliable delivery

- Packets from attacker may be queued and dropped
- Filtered and rate limited by a firewall
- Some traffic do not elicit a response
- Responses may be queued and dropped

Causes the analysis to underestimate
the total number of attacks and attack rate

Backscatter hypothesis

- Any server in the Internet can send unsolicited packets
  - Possible to eliminate flows consistently destined to a single host
- Misinterpretation of random port scans as backscatters
- Vast majority attacks can be differentiated from typical scanning activity

Provides a conservative estimate of current denial-of-service activity

Outline

- Background
- Methodology
- Attack detection and classification
- Analysis of DOS

Attack detection and classification

- Identify and extract backscatter packets from raw trace
- Combine related packets into attack flows
  - Based on victims IP address
- Filter out some attack flows based on intensity, duration and rate

Extracting backscatter packets

- Remove packets
  - Involving legitimate hosts
  - Packets that do not correspond to response traffic
  - Remove TCP RST packets used for scanning
    - These scans have sequential scanning patterns
    - Remover RSTs with clearly non-random behavior
- Remove duplicate packets
  - Same <src IP, dst IP, protocol, src port, dst port> in the last five minutes

Flow-based classification

- Flow-based identification
  - Flow: Series of consecutive packets sharing the same victim IP address
  - Flow lifetime: Timeout approach
    - Defines when a flow begins and ends
    - Packets arrive within a fixed timeout relative to the most recent packet in the flow – same flow
    - More conservative timeout: long flows
    - Shorter timeout: large number of short flows

Flow timeout

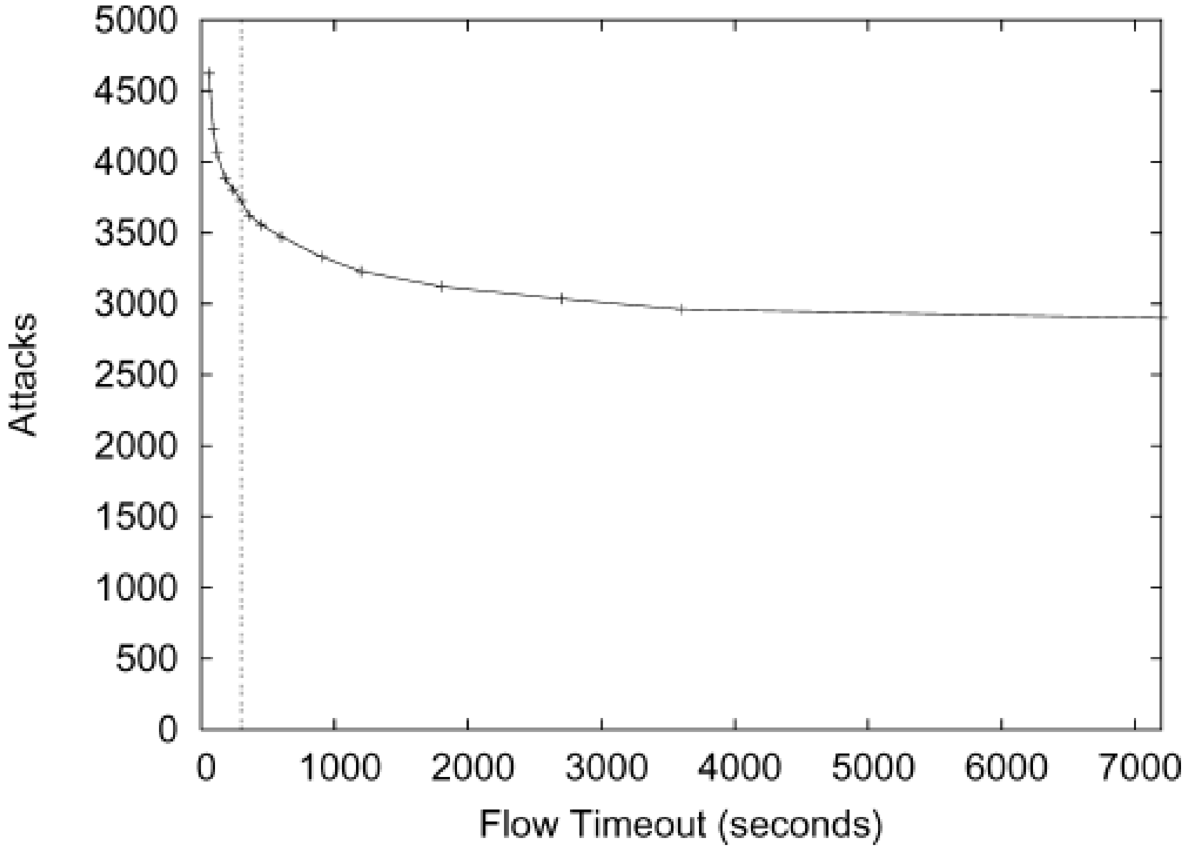

300 seconds (5 minutes)

Filtering attack flows

- Packet threshold
  - Minimum number of packets necessary to classify it to be an attack
  - Filter out short attacks which have negligible impact
- Attack duration
  - Time between first and last packet of a flow
  - Filter out short attacks
- Packet rate
  - Threshold for maximum rate of packet arrivals
  - Largest packet rate across 1-minute buckets

Packet threshold

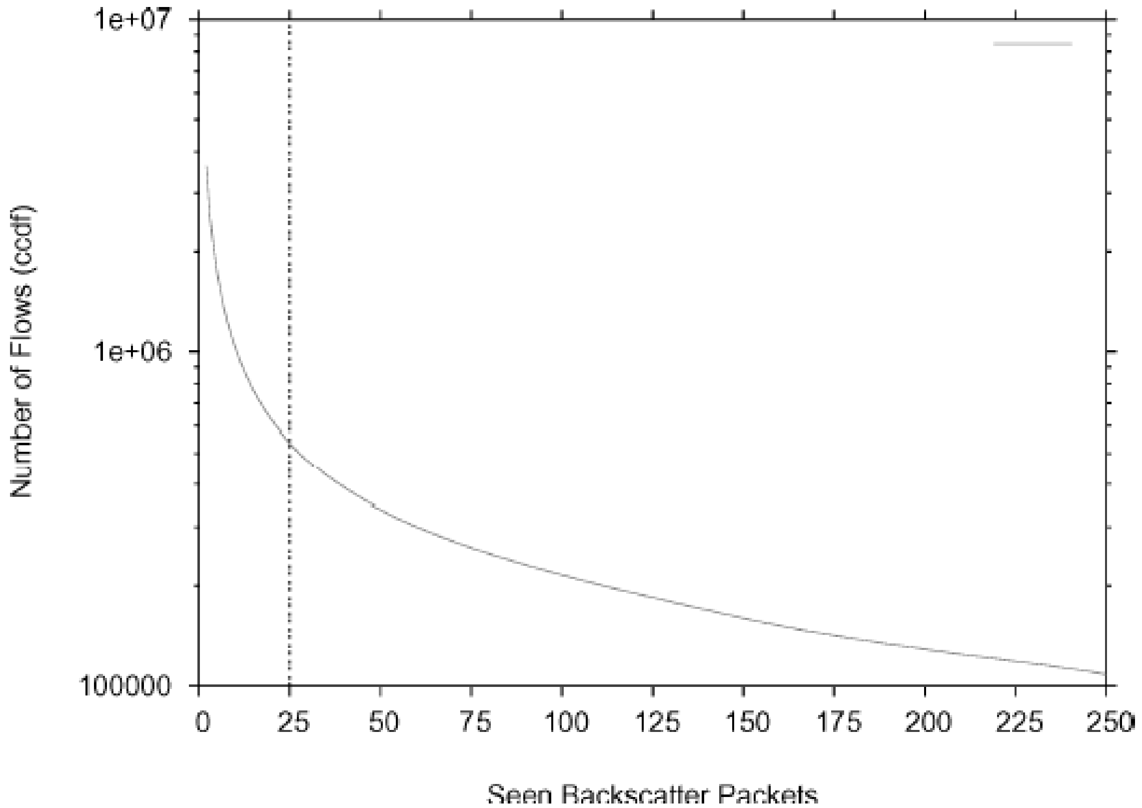

25 packets

Attack duration

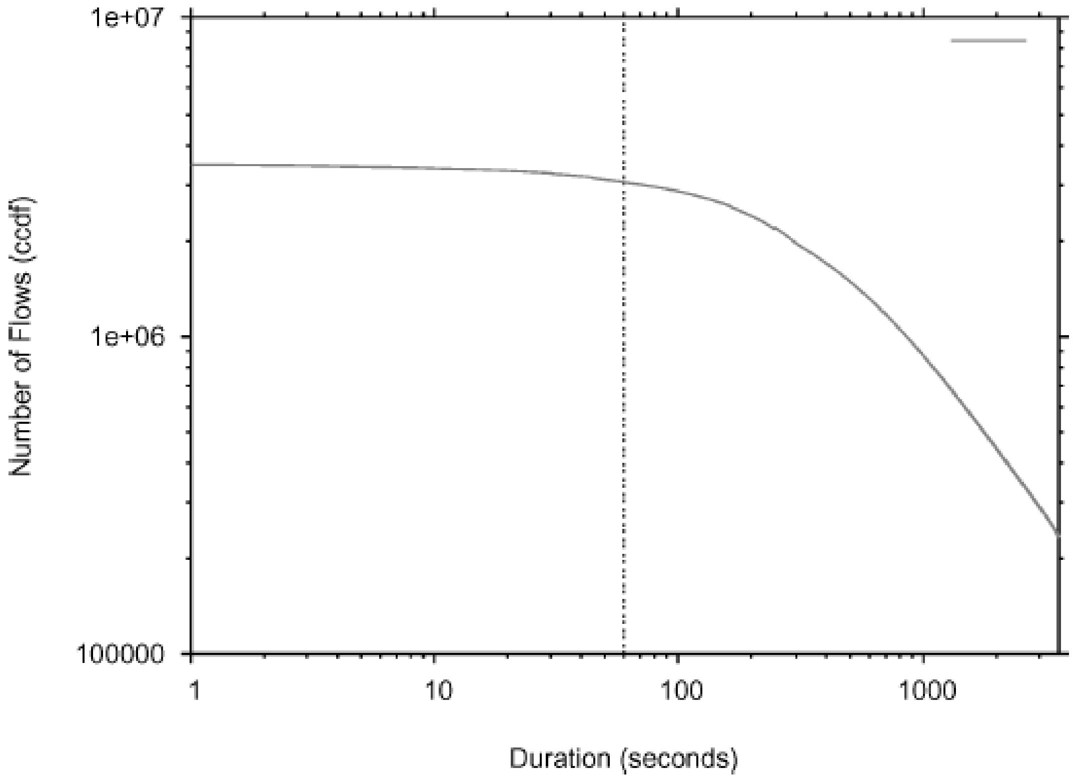

60 seconds

Packet rate

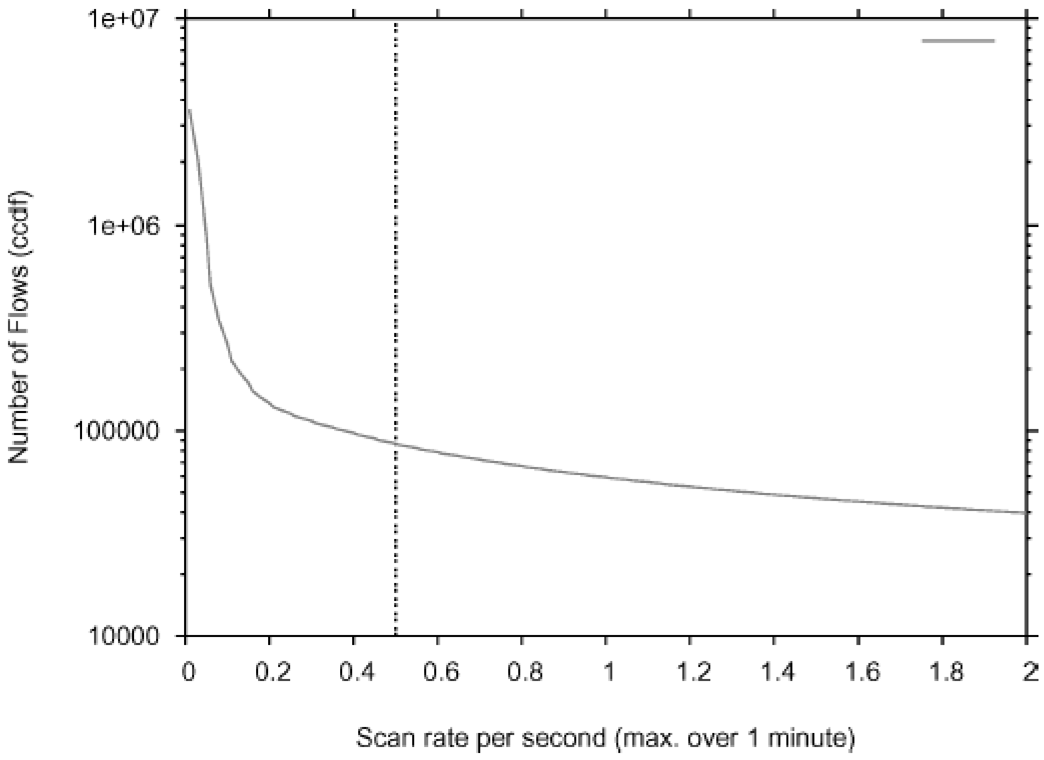

0.5 pps

Extracted Information

- IP Protocol (TCP, UDP, ICMP)
- TCP flag settings (SYN/ACKs, RSTs)
- ICMP payload (copies of original packets)
- Port settings (source and destination ports)
- DNS information

Outline

- Background
- Methodology
- Attack detection and classification
- Analysis of DOS

Analysis: Experimental Platform

Captures all the inbound traffic via Hub

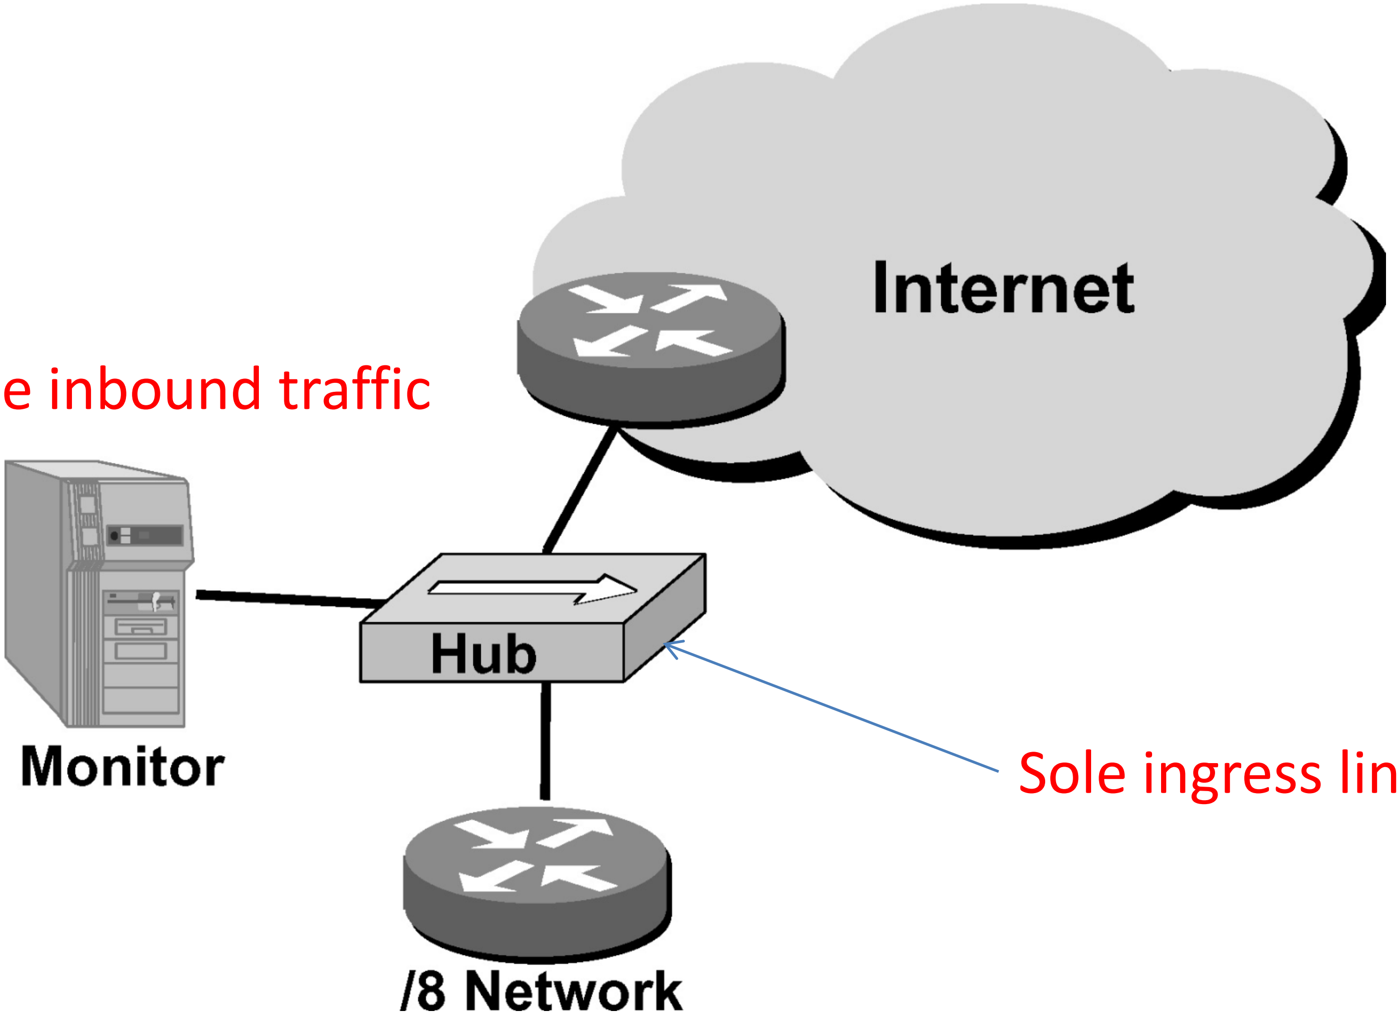

Sole ingress link

2^{24} distinct IPs, 1/256 of the total Ipv4 address space

Summary of Attack Activity

Table III. Summary of Backscatter Database

Starting Date	Duration	Attacks	Backscatter Packets	Unique Victim IPs	Unique Victim Domains	Unique Victim TLDs
2001-02-01	7.5 days	2,618	21,090,742	1,636	729	66
2001-02-11	6.2 days	2,242	30,222,201	1,510	659	63
2001-02-18	7.1 days	2,858	32,159,992	1,921	820	65
2001-02-25	8.9 days	3,346	49,449,404	2,050	677	62
2001-03-06	12.9 days	4,968	59,552,132	2,587	759	73
2001-03-19	8.2 days	2,635	23,588,586	1,618	506	60
2001-04-06	11.8 days	4,343	44,508,551	2,563	694	70
2001-04-22	5.4 days	1,944	14,386,681	1,197	398	55
2001-04-30	6.7 days	828	6,574,228	557	193	41
2001-05-07	14.1 days	4,990	60,647,948	2,933	774	80
2001-05-23	9.1 days	2,993	40,269,047	1,916	546	71
2001-06-01	8.5 days	3,026	47,508,181	1,930	575	60
2001-06-25	8.8 days	2,861	17,408,501	1,897	559	68
2001-07-04	15.8 days	5,666	52,882,496	3,102	747	79
2001-07-19	7.9 days	2,078	36,824,562	1,291	371	60
2001-08-01	7.0 days	974	16,420,358	670	248	47
2001-08-08	6.8 days	1,624	40,248,436	1,059	300	53
2002-05-09	17.5 days	4,820	69,933,861	2,855	681	82
2002-05-29	17.2 days	4,458	103,761,678	2,837	733	87
2002-12-11	7.3 days	2,340	31,139,696	1,016	296	46
2003-11-06	5.0 days	1,416	58,160,582	735	195	51
2004-02-25	10.0 days	5,692	210,181,843	3,088	531	63
Total	209.9 days	68,720	1,066,919,706	34,725	5,273	167

Summary of Attack Activity

- Collection done over a period of 3 years (Feb 1, 2001 – Feb 25, 2004).
- Captured 22 traces of DoS activity.
- Each trace roughly spans a week.
- Total 68,700 attacks to 34,700 unique victim IPs.
- 1,066 million backscatter packets ($\leq 1/256^{th}$ of the total backscatter traffic generated)

Summary of Attack Activity

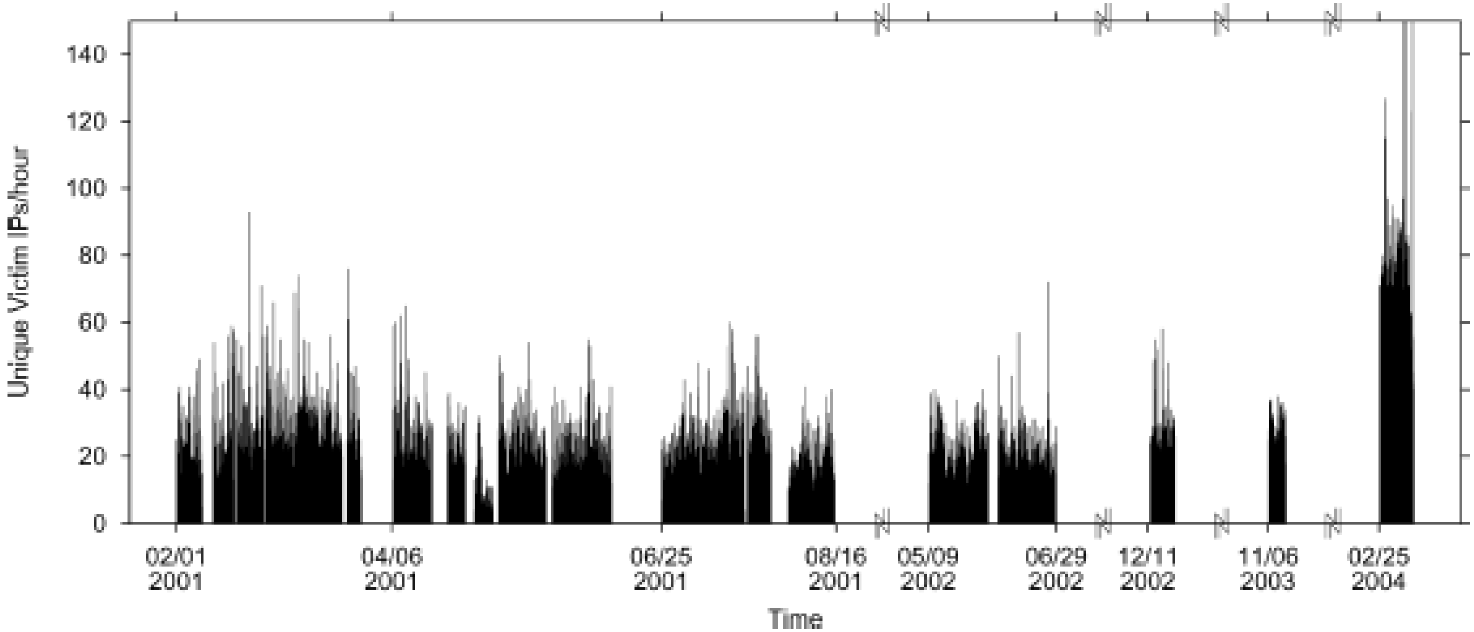

Fig. 7. Estimated number of victims per hour as a function of time (UTC).

- No strong diurnal patterns, as seen in Web or P2P file sharing.
- Rate of attack doesn't change significantly over the period of time.
- Attacks were not clustered on particular subnets.

Summary of Attack Activity

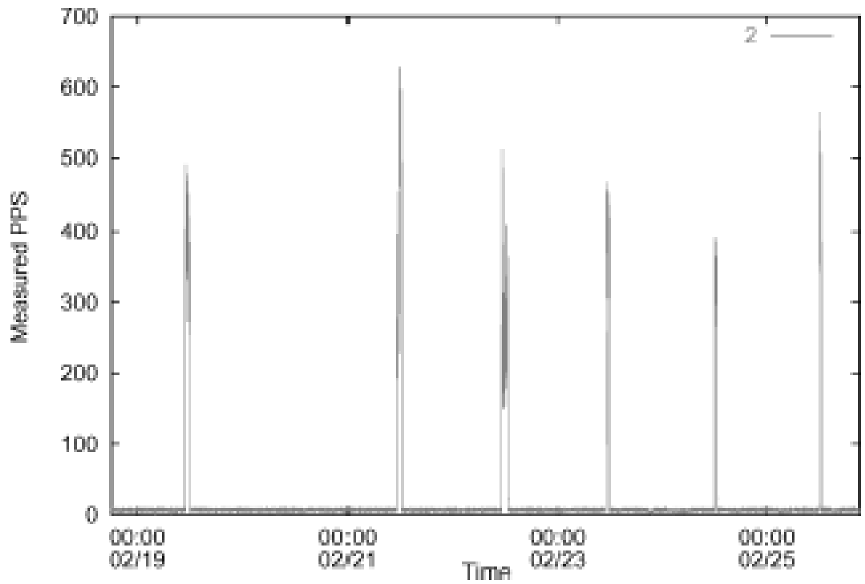

Fig. 8. The measured intensity of an attack to one particular host during the week of February 18, 2001. The spikes occur at noon local time and last for an hour. The attack skipped February 20, 2001, which was a Tuesday.

- Exhibits daily periodic behavior.
- At the same time everyday, attack increases from est. 2,500 pps to 100,000-160,000 pps.
- Attack persists for one hour before subsiding again.
- Tuesdays off (suggests attacks are scripted).

Attack Classification: Protocol

Table IV. Breakdown of Protocols Used in all Attacks Across all Traces. An Entry with Multiple Protocols Indicates an Attack Consisting of a Combination of Packets from Each of the Protocols Listed. "Other" Indicates that the Attack Contained Packets with One or More Miscellaneous Protocols Other Than Those Named in the Table

	Total					
Kind	Attacks (%)		Packets × 1000 (%)		Victims (%)	
TCP	64,952	(95)	949,373	(89)	32,275	(93)
ICMP	1,797	(2.6)	24,567	(2.3)	1,334	(3.8)
TCP/UDP	696	(1.0)	8,526	(0.80)	566	(1.6)
UDP	466	(0.68)	723	(0.07)	312	(0.90)
ICMP/TCP	441	(0.64)	63,728	(6.0)	356	(1.0)
ICMP/IGMP/TCP/UDP	118	(0.17)	342	(0.03)	104	(0.30)
ICMP/TCP/UDP	87	(0.13)	18,865	(1.8)	64	(0.18)
IGMP/TCP/UDP	27	(0.04)	42	(0.00)	22	(0.06)
Other	21	(0.03)	22	(0.00)	10	(0.03)
Other/TCP	18	(0.03)	62	(0.01)	18	(0.05)
ICMP/UDP	16	(0.02)	38	(0.00)	15	(0.04)
ICMP/IGMP/Other/TCP/UDP	16	(0.02)	368	(0.03)	13	(0.04)
IGMP/Other/TCP/UDP	10	(0.01)	56	(0.01)	8	(0.02)
IGMP/TCP	9	(0.01)	32	(0.00)	8	(0.02)
ICMP/IGMP/TCP	7	(0.01)	4	(0.00)	7	(0.02)
ICMP/Other/TCP	6	(0.01)	13	(0.00)	3	(0.01)
ICMP/Other	6	(0.01)	3	(0.00)	4	(0.01)
IGMP/Other/TCP	5	(0.01)	145	(0.01)	5	(0.01)
Other/TCP/UDP	5	(0.01)	2	(0.00)	5	(0.01)
IGMP/Other	5	(0.01)	3	(0.00)	4	(0.01)

Attack Classification: Protocol

Table shows –

- 95% of attacks and 89% of packets use TCP protocol.
- Distant second is ICMP with 2.6% of attacks.
- Breakdown of TCP attacks shows most of the attacks target multiple ports.
- Most popular individual target ports: HTTP (80), IRC (6667), port 0, Authd(113)

Attack Classification: Rate

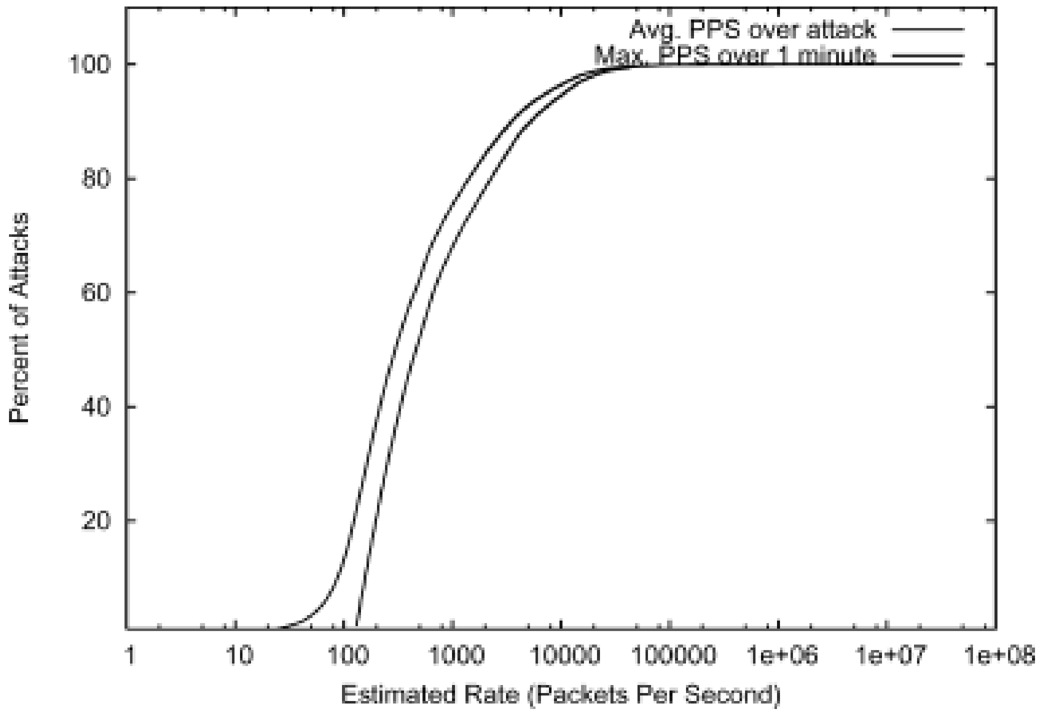

Cumulative distributions of estimated attack rates in packets per second.

- 500 SYN pps are enough to overwhelm a server.
- 65% attacks had 500 pps or higher.
- 4% attacks had ≥ 14,000 pps, enough to compromise attack-resistant firewalls.

Attack Classification: Duration

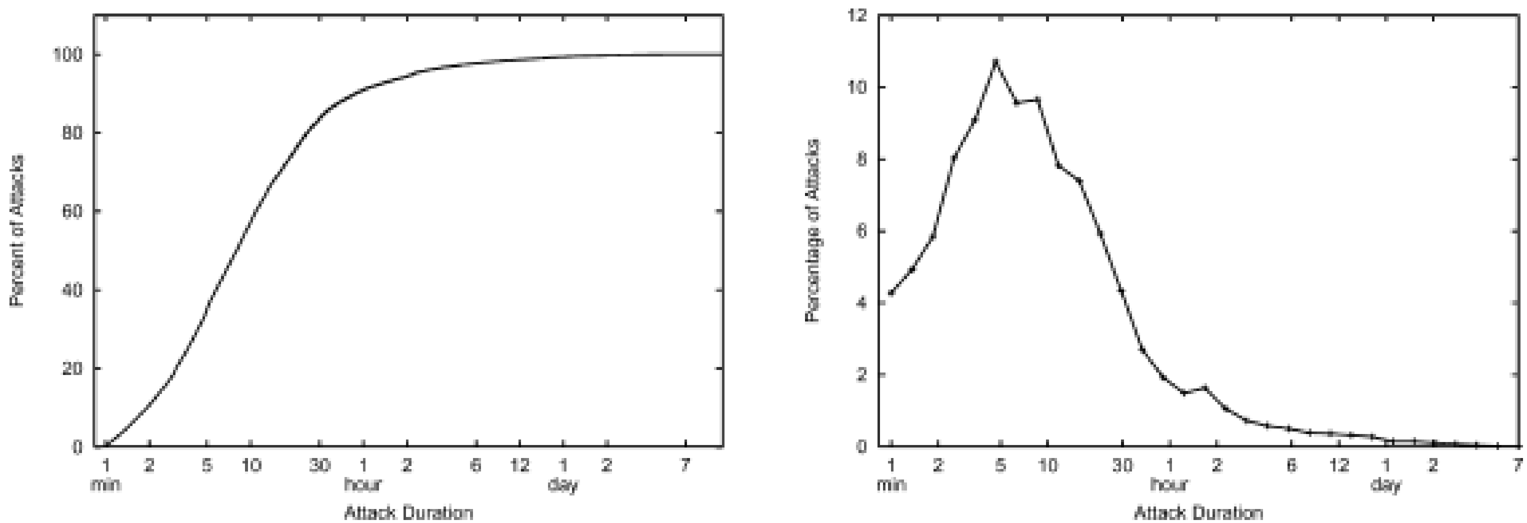

(a) Cumulative distribution (b) Probability density

Fig. 10. Cumulative and probability distributions of attack durations.

- 60% attacks less than 10 min
- 80% are less than 30 min
- 2.4% are greater than 5 hrs
- 1.5% are greater than 10 hrs
- 0.53% span multiple days
- PDF graph shows peak is at 5 min (10.8%), 10 min (9.7%)

Victim Classification: Type

Table VI. Breakdown of Victim Hostnames

	Total			
Kind	Attacks (%)		Packets×1000 (%)	
In-Addr Arpa	28,547	(42)	498,775	(47)
Unclassified	25,216	(37)	404,111	(38)
Broadband	5,520	(8.0)	31,006	(2.9)
Dial-Up	4,864	(7.1)	39,479	(3.7)
IRC Server	1,156	(1.7)	49,950	(4.7)
Nameserver	1,141	(1.7)	17,685	(1.7)
Web Server	996	(1.4)	11,968	(1.1)
Router	885	(1.3)	11,148	(1.0)
Mail Server	377	(0.55)	2,501	(0.23)
Firewall	18	(0.03)	297	(0.03)

Victim Classification: TLD

Table VII. Breakdown of Victim Top-Level Domains (TLDs). The "arpa" TLD Represents Those Attacks for which a Reverse DNS Lookup Failed on the Victim IP Address

Kind		Total					
		Attacks (%)		Packets×1000 (%)		Victims (%)	
arpa		28,547	(42)	498,775	(47)	14,513	(42)
net		9,291	(14)	150,339	(14)	5,113	(15)
com		7,721	(11)	162,539	(15)	4,046	(12)
	ro	7,235	(11)	33,661	(3.2)	3,031	(8.7)
	br	2,822	(4.1)	22,286	(2.1)	1,228	(3.5)
edu		1,219	(1.8)	13,258	(1.2)	659	(1.9)
	ca	1,167	(1.7)	5,307	(0.50)	636	(1.8)
org		890	(1.3)	26,340	(2.5)	431	(1.2)
	it	638	(0.93)	5,843	(0.55)	424	(1.2)
	mx	610	(0.89)	1,793	(0.17)	375	(1.1)
	nl	566	(0.82)	1,857	(0.17)	306	(0.88)
	jp	520	(0.76)	14,467	(1.4)	154	(0.44)
	de	435	(0.63)	3,114	(0.29)	247	(0.71)
	no	429	(0.62)	4,422	(0.41)	220	(0.63)
	uk	409	(0.60)	3,510	(0.33)	221	(0.64)
	be	405	(0.59)	1,516	(0.14)	177	(0.51)
	pl	383	(0.56)	1,794	(0.17)	188	(0.54)
	au	378	(0.55)	7,710	(0.72)	244	(0.70)
	se	346	(0.50)	11,548	(1.1)	216	(0.62)
	fr	313	(0.46)	1,083	(0.10)	145	(0.42)

- Over 10% targeted *com* & *net*
- 1.3-1.7% targeted *org* & *edu*
- 11% were targeted to *ro*
- 4% to *br*

Victim Classification: Repeated Attacks

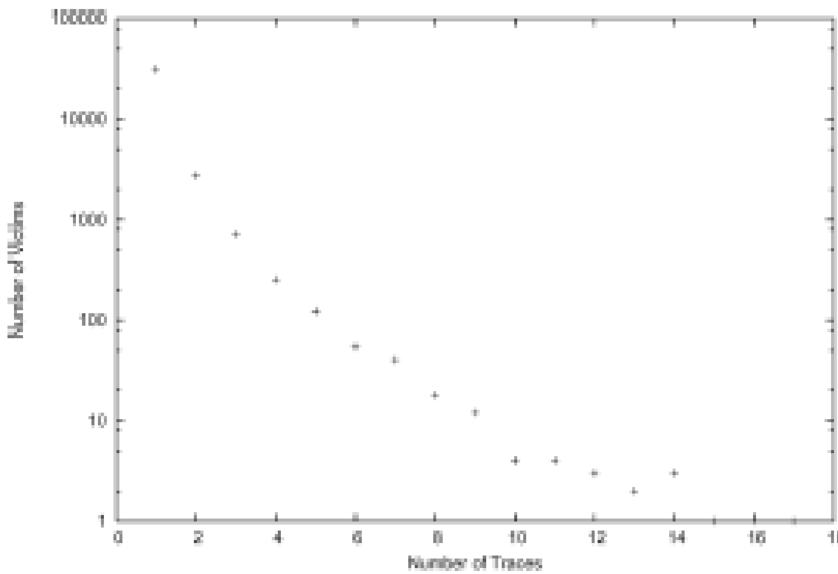

(a) Histogram showing the number of traces DoS victims were attacked in.

- Most victims (89%) were attacked in only one trace.
- Most of the remaining victims (7.8%) appear in two traces.
- Victims can appear in multiple traces because of attacks that span trace boundaries.
- 3% victims appear in more than 3 traces, nevertheless.

Victim Classification: Repeated Attacks

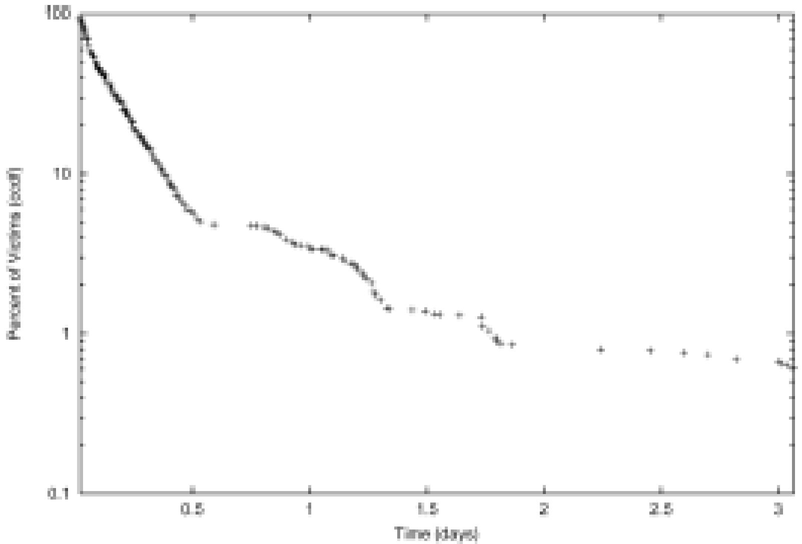

(b) CCDF showing the length of time between the first and last traces in which we observed an attack targeting a victim. Only victims attacked in more than one trace are shown.

Table VIII. The Host Types of the 15 Most Frequently Attacked Victims

Host Type	Number of Victims
Nameservers	5
IRC servers	3
Broadband	4
Education	2
No Hostname	1

Table IX. The Countries in Which the 15 Most Frequently Attacked Victims are Located

Country	Number of Victims
United States	6
Romania	4
Norway	2
Japan	1
France	1
Austria	1

15 victims that appear in 10 or more traces

Validation

- Nearly all of the packets attribute to the backscatter do not provoke a response, so these packets could not have been used to probe the monitored network.
- Anderson-Darling test (*a statistical test of whether there is evidence that a given sample of data did not arise from a given probability distribution*) to determine if the distribution of destination addresses is uniform. Validated for most attacks at the 0.05 significance level.

Validation cont'd…

- Duplicated portion of the analysis using data taken from several university-related networks in California.
  - Although this is a much smaller dataset; for 98% of the victim IP recorded in this dataset, corresponding record was found at the same time in larger dataset.
- Data from Asta Networks describing DoS attacks detected also qualitatively confirms the data in this paper.

Conclusions

- Presented new technique called “backscatter analysis” for estimating DoS attack activity on the Internet.
- Observed widespread DoS attacks distributed among many domains and ISPs.
- Size and length of attacks were heavy tailed.
- Surprising number of attacks directed at a few foreign countries. (or as we non-US citizens call them – home countries).
- Witnessed over 68,000 attacks during 3 years, with little signs of abatement.

Questions?