

# CS 268: Lecture 19 (Malware)

Ion Stoica  
Computer Science Division  
Department of Electrical Engineering and Computer Sciences  
University of California, Berkeley  
Berkeley, CA 94720-1776

(Based on slides from Vern Paxson and Stefan Savage)

## Motivation

- Internet currently used for important services
  - Financial transactions, medical records
- Could be used in the future for *critical* services
  - 911, surgical operations, energy system control, transportation system control
- Networks more open than ever before
  - Global, ubiquitous Internet, wireless
- Malicious Users
  - Selfish users: want more network resources than you
  - Malicious users: would hurt you even if it doesn't get them more network resources

2

## Network Security Problems

- Host Compromise
  - Attacker gains control of a host
- Denial-of-Service
  - Attacker prevents legitimate users from gaining service
- Attack can be both
  - E.g., host compromise that provides resources for denial-of-service

3

## Host Compromise

- One of earliest major Internet security incidents
  - Internet Worm (1988): compromised almost every BSD-derived machine on Internet
- Today: estimated that a single worm could compromise 10M hosts in < 5 min
- Attacker gains control of a host
  - Read data
  - Erase data
  - Compromise another host
  - Launch denial-of-service attacks on another host

4

## Definitions

- Worm
  - Replicates itself
  - Usually relies on stack overflow attack
- Virus
  - Program that attaches itself to another (usually trusted) program
- Trojan horse
  - Program that allows a hacker a back way
  - Usually relies on user exploitation
- Botnet
  - A collection of programs running autonomously and controlled remotely
  - Can be used to spread out worms, mounting DDoS attacks

5

## Host Compromise: Stack Overflow

- Typical code has many bugs because those bugs are not triggered by common input
- Network code is vulnerable because it accepts input from the network
- Network code that runs with high privileges (i.e., as root) is especially dangerous
  - E.g., web server


6

## Example

- What is wrong here?

```
// Copy a variable length user name from a packet
#define MAXNAMELEN 64
int offset = OFFSET_USERNAME;
char username[MAXNAMELEN];
int name_len;

name_len = packet[offset];
memcpy(&username, packet[offset + 1], name_len);
```




7

## Example

```
void foo(packet) {
#define MAXNAMELEN 64
int offset = OFFSET_USERNAME;
char username[MAXNAMELEN];
int name_len;

name_len = packet[offset];
memcpy(&username,
packet[offset + 1], name_len);
...
}
```




8

## Example

```
void foo(packet) {
#define MAXNAMELEN 64
int offset = OFFSET_USERNAME;
char username[MAXNAMELEN];
int name_len;

name_len = packet[offset];
memcpy(&username,
packet[offset + 1], name_len);
...
}
```



9

## Effect of Stack Overflow

- Write into part of the stack or heap
  - Write arbitrary code to part of memory
  - Cause program execution to jump to arbitrary code
- Worm
  - Probes host for vulnerable software
  - Sends bogus input
  - Attacker can do anything that the privileges of the buggy program allows
    - Launches copy of itself on compromised host
  - Spread at exponential rate
  - 10M hosts in < 5 minutes

10

## Outline



- > Worm propagation
  - Threat detection – content sifting

11

## Worm Spreading

$$f = \frac{(e^{K(t-T)} - 1)}{(1 + e^{K(t-T)})}$$

- $f$  – fraction of hosts infected
- $K$  – rate at which one host can compromise others
- $T$  – start time of the attack



12

## Worm Examples

- Morris worm (1988)
- Code Red (2001)
- MS Slammer (January 2003)
- MS Blaster (August 2003)

13

## Morris Worm (1988)

- Infect multiple types of machines (Sun 3 and VAX)
  - Spread using a Sendmail bug
- Attack multiple security holes including
  - Buffer overflow in fingerd
  - Debugging routines in Sendmail
  - Password cracking
- Intend to be benign but it had a bug
  - Fixed chance the worm wouldn't quit when reinfesting a machine → number of worm on a host built up rendering the machine unusable

14

## Code Red Worm (2001)

- Attempts to connect to TCP port 80 on a randomly chosen host
- If successful, the attacking host sends a crafted HTTP GET request to the victim, attempting to exploit a buffer overflow
- Worm "bug": all copies of the worm use the same random generator to scan new hosts
  - DoS attack on those hosts
  - Slow to infect new hosts
- 2<sup>nd</sup> generation of Code Red fixed the bug!
  - It spread much faster

15

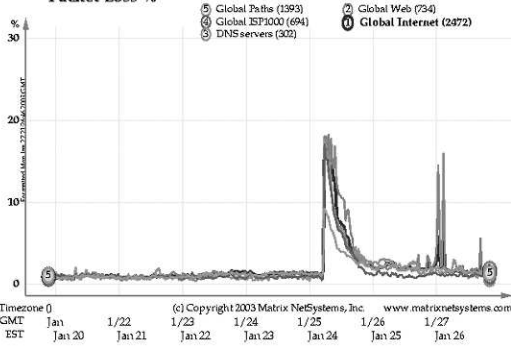
## MS SQL Slammer (January 2003)

- Uses UDP port 1434 to exploit a buffer overflow in MS SQL server
- Effect
  - Generate massive amounts of network packets
  - Brought down as many as 5 of the 13 internet root name servers
- Others
  - The worm only spreads as an in-memory process: it never writes itself to the hard drive
    - Solution: close UDP port on firewall and reboot

16

## MS SQL Slammer (January 2003)

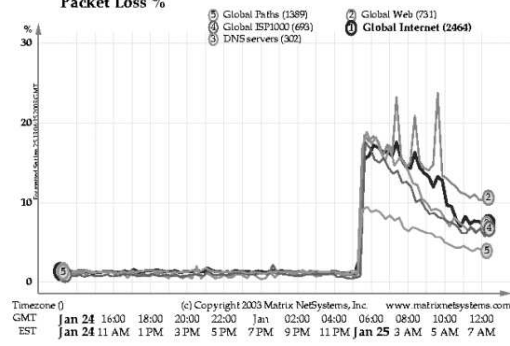
Packet Loss %



17

## MS SQL Slammer (January 2003)

Packet Loss %



18

## MS Blaster (August 2003)

- Exploit a buffer overflow vulnerability of the RPC (Remote Procedure Call) service
- Scan a random IP range to look for vulnerable systems on TCP port 135
- Open TCP port 4444, which could allow an attacker to execute commands on the system
- DoS windowsupdate.com on certain versions of Windows

19

## Hall of Shame

- Software that have had many stack overflow bugs:
  - BIND (most popular DNS server)
  - RPC (Remote Procedure Call, used for NFS)
    - NFS (Network File System), widely used at UCB
  - Sendmail (most popular UNIX mail delivery software)
  - IIS (Windows web server)
  - SNMP (Simple Network Management Protocol, used to manage routers and other network devices)

20

## Spreading faster—distributed coordination (*Warhol* worms)

- Idea 1: *reduce redundant scanning.*
  - Construct permutation of address space.
  - Each new worm instance starts at random point
  - Worm instance that "encounters" another instance re-randomizes
- Idea 2: *reduce slow startup phase.*
  - Construct a "hit-list" of vulnerable servers in advance
  - Then: for 1M vulnerable hosts, 10K hit-list, 100 scans/worm/sec, 1 sec to infect → 99% infection in 5 minutes.

21

## Spreading still faster — *Flash* worms

- Idea: use an *Internet-sized hit list.*
  - Initial copy of the worm has the entire hit list
  - Each generation, infects  $n$  from the list, gives each  $1/n$  of list
  - Need to engineer for locality, failure & redundancy.
  - But:  $n = 10$  requires, 7 generations to infect  $10^7$  hosts → tens of seconds.

22

## How can we defend against Internet-scale worms?

- Time scales rule out human intervention → Need automated detectors, response (And perhaps honeypots to confuse scanning?)
- Very hard research question!
- And it's only half of the problem . . .

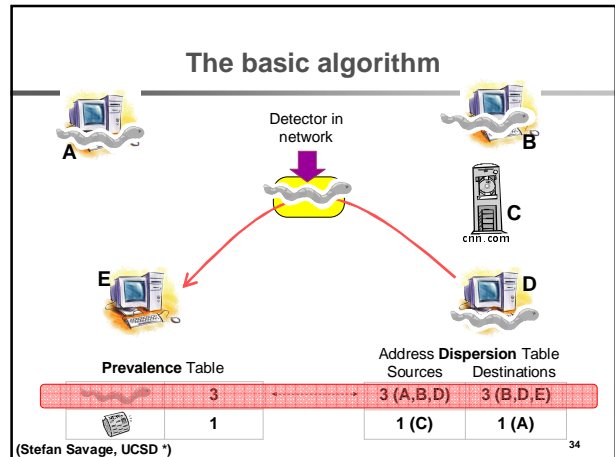
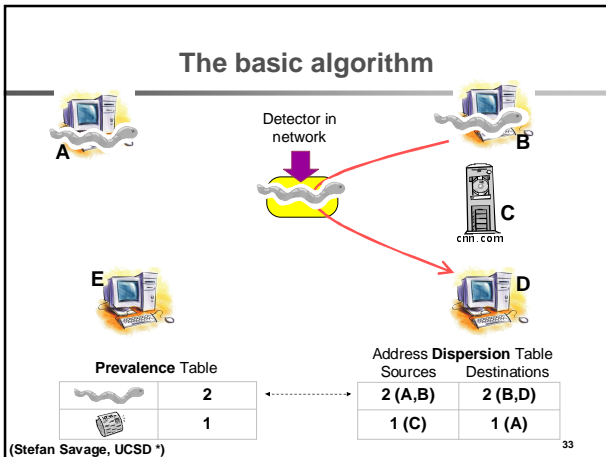
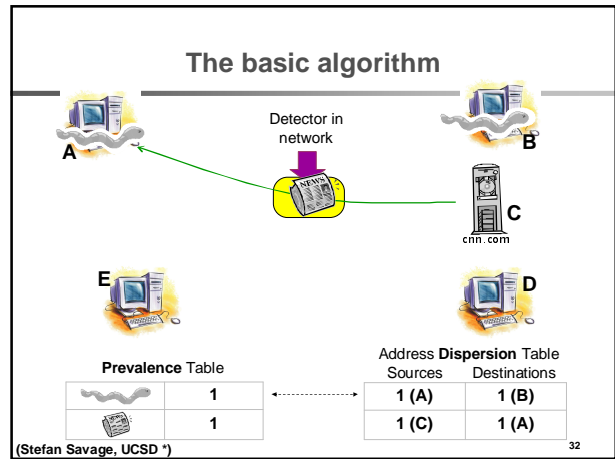
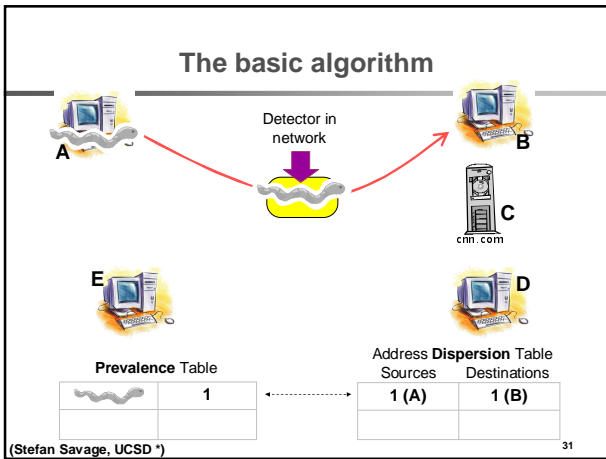
23

## *Contagion* worms

- Suppose you have two exploits: Es (Web server) and Ec (Web client)
- You infect a server (or client) with Es (Ec)
- Then you . . . wait (Perhaps you bait, e.g., host porn)
- When vulnerable client arrives, infect it
- You send over *both* Es and Ec
- As client happens to visit other vulnerable servers ) infects

24





### Challenges

- **Computation**
  - To support a 1Gbps line rate we have 12us to process each packet, at 10Gbps 1.2us, at 40Gbps...
    - Dominated by memory references; state expensive
  - Content sifting requires looking at *every* byte in a packet
- **State**
  - On a fully-loaded 1Gbps link a naïve implementation can easily consume 100MB/sec for table
  - Computation/memory duality: on high-speed (ASIC) implementation, latency requirements may limit state to on-chip SRAM

(Stefan Savage, UCSD \*) 35

### Which substrings to index?

- **Approach 1: Index all substrings**
  - Way too many substrings → too much computation → too much state
- **Approach 2: Index whole packet**
  - Very fast but trivially evadable (e.g., Witty, Email Viruses)
- **Approach 3: Index all contiguous substrings of a fixed length 'S'**
  - Can capture all signatures of length 'S' and larger

**A B C D E F G H I J K**

(Stefan Savage, UCSD \*) 36

## How to represent substrings?

- Store **hash** instead of literal to reduce state
- **Incremental hash** to reduce computation
- **Rabin fingerprint** is one such efficient incremental hash function [Rabin81,Manber94]
  - One multiplication, addition and mask per byte

P1 R A N D A B C D O M  
Fingerprint = 11000000

P2 R A B C D A N D O M  
Fingerprint = 11000000

(Stefan Savage, UCSD \*)

37

## How to subsample?

- **Approach 1: sample packets**
  - If we chose 1 in N, detection will be slowed by N
- **Approach 2: sample at particular byte offsets**
  - Susceptible to simple evasion attacks
  - No guarantee that we will sample same sub-string in every packet
- **Approach 3: sample based on the hash of the substring**

(Stefan Savage, UCSD \*)

38

## Value sampling [Manber '94]

- Sample hash if last 'N' bits of the hash are equal to the value 'V'
  - The number of bits 'N' can be dynamically set
  - The value 'V' can be randomized for resiliency

A B C D E F G H I J K

Fingerprint = 1100000000000000


SAMPLE

- $P_{\text{track}}$  → Probability of selecting at least one substring of length S in a L byte invariant
  - For 1/64 sampling (last 6 bits equal to 0), and 40 byte substrings
  - $P_{\text{track}} = 99.64\%$  for a 400 byte invariant

(Stefan Savage, UCSD \*)

39

## Observation: High-prevalence strings are rare



(Stefan Savage, UCSD \*)

40


## Efficient high-pass filters for content

- Only want to keep state for prevalent substrings
- Chicken vs egg: how to count strings without maintaining state for them?
- **Multi Stage Filters:** randomized technique for counting "heavy hitter" network flows with low state and few false positives [Estan02]
  - Instead of using flow id, use **content hash**
    - Rabin Fingerprints with Mandber's Value sampling
  - **Three orders of magnitude** memory savings

(Stefan Savage, UCSD \*)

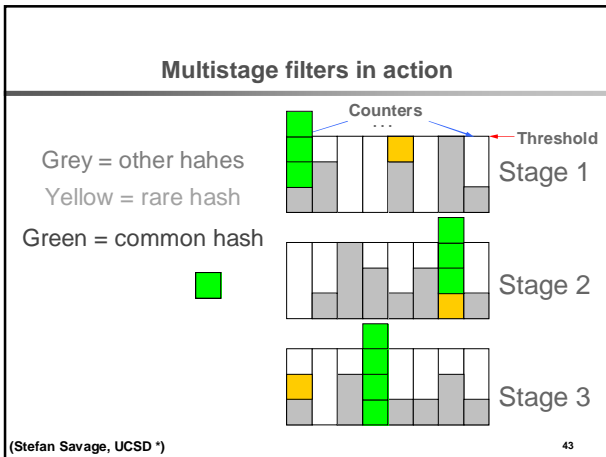
41

## Finding "heavy hitters" via Multistage Filters



(Stefan Savage, UCSD \*)

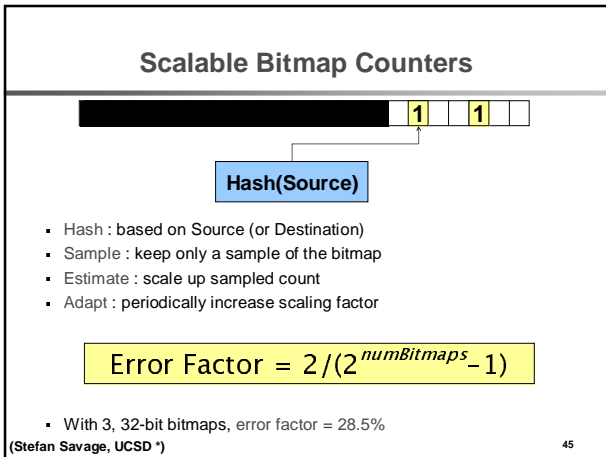
42



### Observation: High address dispersion is rare too

- Naïve implementation might maintain a list of sources (or destinations) for each string hash
- But dispersion **only** matters if its over threshold
  - Approximate counting may suffice
  - Trades accuracy for state in data structure**
- Scalable Bitmap Counters
  - Similar to multi-resolution bitmaps [Estan03]
  - Reduce memory by 5x for modest accuracy error

(Stefan Savage, UCSD \*) 44



### Content sifting summary

- Index fixed-length substrings using incremental hashes
- Subsample hashes as function of hash value
- Multi-stage filters to filter out uncommon strings
- Scalable bitmaps to tell if number of distinct addresses per hash crosses threshold
- Now** its fast enough to implement

(Stefan Savage, UCSD \*) 46

