Detecting Security Problems and Evaluating Security Solutions
CS 239
Advanced Topics in Network Security
Peter Reiher
May 12, 2003

Internet Measurement

- The Internet as a whole is poorly measured
  - And, hence, poorly understood
- No existing network-wide infrastructure for measuring anything
- Ad hoc attempts to get some handle on what's going on in the network

Some Security Measurement Questions

- What fraction of all IP packets have spoofed addresses?
- How many DDoS attacks occur each day?
- How many compromised machines are there on the Internet?
- If I installed secure BGP at 200 chosen locations, how much better would things be?

So, How to Answer These Questions?

- Deduce based on the evidence available
- Obtain snapshots from some points in the network
- Use simulation techniques
- Use honeypots/honeynets to attract attacks for measurement and analysis
- Install serious measurement capabilities in the network

Inferring DoS Attacks

- An attempt to answer question of how common DoS attacks are
- How to answer that question?
  - Ask people to tell you when they're victims
  - Observe congestion and deduce when it's caused by DoS
  - Or, use backscatter

Idea Behind Backscatter Measurement Technique

- DoS consists of a stream of garbage packets to a single destination
- The victim doesn't know they're garbage, so it answers them normally
- Often, the attacker spoofs the source address of attack packets
  - So responses go to the real machines whose addresses were spoofed

Spoofing and DoS Attacks

- In principle, DoS attackers could spoof any source address
- Most often, they seem to spoof randomly from entire IP address space
  - Choose new address from 2^{32} possible addresses for each packet
- If enough packets are sent in attack, every machine on the Internet will see some responses

Using Backscatter in Practice

- Set up network of test machines
  - That send and receive no legitimate traffic
- Record every packet they receive
- Try to identify which of them seem to be legitimate responses to some packet
- Identify each such packet as a backscatter packet in a DoS attack

For Example,

I got a SYN/ACK from 67.123.55.40 when I didn't send him a SYN

Probably someone else sent him a SYN with my source address spoofed

Maybe there's a DoS attack on 67.123.55.40!

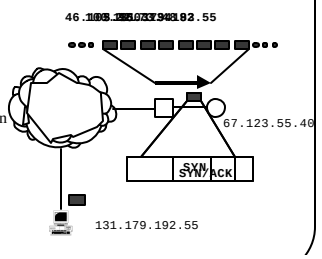

Practical Use of Backscatter

- Not definitive, since anyone could have sent this packet
  - Or could be a legitimate response to something other than DoS
- More accurate if you monitor lots of addresses
- Tricky to tell when attacks begin and end

CAIDA Backscatter Experiment

- Run during three week-long periods in 2000
- Using /8 network
  - So they control 2^{24} distinct IP addresses, or $1/256^{th}$ of all addresses
- Monitored all traffic arriving for any of these addresses

Results

- During one week, saw 12,805 attacks
- Over three weeks observed 200 million backscatter packets
  - Presumably out of around 50 billion such packets
- More than 5000 victim addresses in more than 2000 domains

Types of Attacks

- TCP dominated
  - 94% of all attacks were TCP
- Small number were ICMP
  - But they represented nearly half of the backscatter traffic
- Only 2% were UDP

Who Were the Victims?

- Victim's IP address is source address in backscatter packet
- Reverse lookup on that address to get victim's DNS name
- Failed 30% of time
- In other cases, .net and .com very popular for attack targets

How Long Were the Attacks?

- Typically one hour or less (90% of them)
- But 2% of attacks more than 5 hours
- 1% longer than 10 hours
- Dozens of attacks lasted for days

How Strong Are the Attacks?

- More than 90% were 10,000 pkts/sec or less
  - 500 SYNs per second overwhelms unprotected server
    - 46% of attacks were that strong
  - 14,000 SYNs overwhelms anti-DoS firewall
    - 2.4% of attacks were that strong

Other Approaches

- CERT
- Other network observation points
- Honeynets
- The grapevine

CERT

- Keeps reasonably close eye on the Internet
- Is extremely careful about issuing advisories
  - Avoids panic, but delays response
- Their staff observe and collect reports from other human sources

Other Network Observation Points

- CAIDA, at San Diego, measures and observes much Internet behavior
  - Including security-related behavior
- Other places observe some forms of behavior
  - E.g., Oregon Routeview project to collect router information from several ASs

Honeynets

- Certain people maintain networks just to watch for attack traffic
  - The backscatter project was a specialized version
  - More generally, draw attackers in by promise of unprotected machine
  - Use their behavior to learn about new problems
- Pretty much ad hoc research and volunteer activity, so far

The Grapevine

- Sysadmins and network administrators tend to know each other
- When they see problems, they talk to each other
- Word of new problems spreads quickly
  - Often using telephones, rather than the network
- Much of what CERT knows about originates this way

Commercial Players

- Companies like Network Associates and Symantec make it their business to know about certain kinds of problems
  - Typically viruses
- Smaller companies try to build reputation by finding and diagnosing problems

What Should We Do?

- Is the current approach to finding security problems in the Internet adequate?
- If not, what would be?
- What should a system for watching for Internet threats look like?
- Who would run it?
- How would they do it?