

# Network DoS in 100x100: Should We Care?

And Other Inflammatory Opinions about 100x100

Mike Reiter

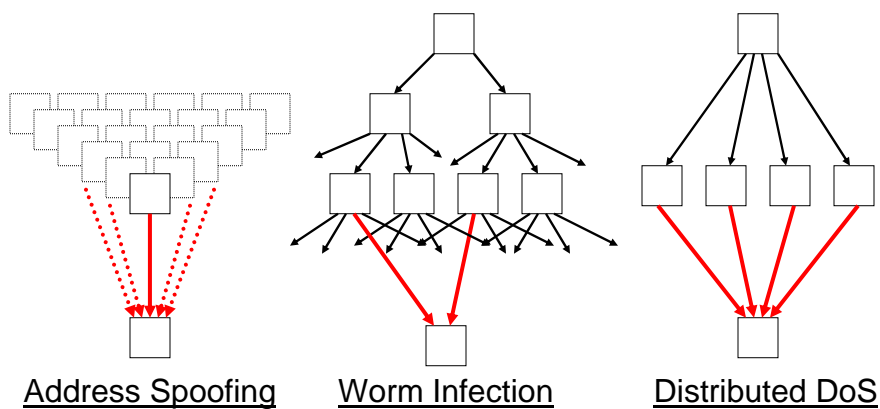
Professor of ECE and CS

reiter@cmu.edu

Copyright © 2003 by Michael Reiter  
All rights reserved.

CarnegieMellon

## Network DoS



Copyright © 2003 by Michael Reiter  
All rights reserved.

CarnegieMellon

## Root Causes

---

- For spoofing, the forgeability of source addresses
  - ▼ By some accounts, (fully) spoofed attacks now account for only 5% of attacks and are decreasing
  - ▼ Conjectures for decrease
    - ▼ More common ingress/egress filtering (leading to an increase in subnet spoofing)
    - ▼ Availability of more potent alternatives
  
- For worm infection and distributed DoS, the poor quality of consumer software
  - ▼ In fact, this is a common denominator among forged addresses, too

## Should We Care About DoS in 100x100?

---

- I'm serious in asking this question
- Perspective #1: everyone's knee-jerk reaction
  - ▼ DoS was arguably a very major oversight in the development of today's Internet
  - ▼ Higher bandwidth + more powerful end hosts will make DoS worse
  - ▼ We should learn from our mistake and build powerful controls into the network to defend against DoS
- Perspective #2: umm, on second thought ...
  - ▼ In the timeframe of 100x100, at least as much attention will be paid to improving software quality as to redesigning the Internet
    - ▼ General improvement in software quality will dramatically reduce the prevalence of worms and distributed DoS
    - ▼ This is the "morally right" fix
  - ▼ "Much DoS can largely be defeated 'outside' the network"

## Evaluating Perspective #2

- Contrary to most current thinking, and so we should try to evaluate it
- The “software will get better before 100x100” argument depends on mostly non-technical issues
  - ▼ Yes, better software engineering techniques are part of it ...
  - ▼ ... but this will come about iff legal and market pressures make it
  - ▼ One thing seems clear: The imperative for better software seems at least as urgent than for a network with the capacity of 100x100
- “We can defeat DoS ‘outside’ the net” is something we can evaluate technically, at least for today’s network
  - ▼ We’ve been trying to do so based upon data collected at the border of a very large network

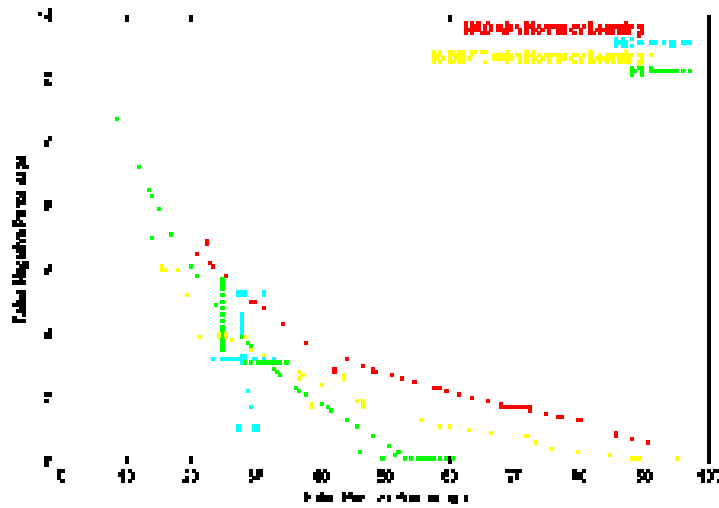
## Filtering DoS “Outside” the Net

[w/ Collins 2003]

- Evaluating various proposed target-resident filtering strategies
  - ▼ Static clustering
  - ▼ Network aware clustering [Jung et al. 2002]
  - ▼ Hop count filtering [Jin et al. 2003]
  - ▼ Path identifiers [Yaar et al. 2003]
    - ▼ Note: requires changes to routers
- Evaluating them under different learning assumptions
  - ▼ “Normalcy learning”: Filter is trained on normal data
  - ▼ “Attacker learning”: Filter is trained on both normal and attacker data
- Utilizing datasets collected at the border of a large network
  - ▼ Isolated DoS attacks of both spoofed and non-spoofed varieties
  - ▼ Utilized attacks to build model, and evaluated performance of filters against traffic generated from this model

## Filters Versus Spoofed Traffic

7




Copyright © 2003 by Michael Reiter  
All rights reserved.

Carnegie Mellon

## Filters Versus Non-Spoofed Traffic

8



Copyright © 2003 by Michael Reiter  
All rights reserved.

Carnegie Mellon

## What To Do With Source Addresses?

---

- Today, source (IP) addresses are a *hint*
- They are unreliable, and so can be forged by bad guys
  - ▼ Sometimes utilized in DoS, as we've seen
- Included by default, and so impinge on privacy for good guys
  - ▼ IP addresses provide index for building profiles from web logs
  - ▼ Research community works hard on overlays for providing anonymity, but anonymity is hard and expensive
- So, as hints, source addresses are detrimental in both respects
- What's the right answer in 100x100?