Measurement and Security

Steven M. Bellovin
smb@research.att.com
http://www.research.att.com/~smb
AT&T Labs Research

"If you can not measure it, you can not improve it."

"When you can measure what you are speaking about, and express it in numbers, you know something about it; but when you cannot measure it, when you cannot express it in numbers, your knowledge is of a meagre and unsatisfactory kind; it may be the beginning of knowledge, but you have scarcely in your thoughts advanced to the state of *Science*, whatever the matter may be."

—Lord Kelvin

Questions

- What can measurement tell us?
- What would we like it to tell us?
- What seems beyond reach?

Obvious Results

- If there's too much traffic all over, there's probably a worm outbreak
- If there's too much traffic towards a single point, it's probably a DDoS attack
- Or is it?

What is the Traffic Mix?

- Random protocols from random addresses?
  - ☞Probably an attack
- All TCP port 80 from plausible addresses?
  - ☞Maybe it's DDoS, maybe it's a flash crowd

Taking a Closer Look

- What is the historical behavior of this destination?
- If it's never received port 80 traffic before, it's probably an attack
- What is the distribution of source traffic, compared with historical patterns?
- If the pattern is very different, it's likely an attack

Controlling DDoS

- Not possible to block DDoS attacks
- Often, we can mitigate it if we can measure sources in real-time
- For sources with excess traffic over historical data, clamp to historical rate
- Note requirements: real-time data *and* historical archive

Early Warnings

- Monitor protocol distribution
- Unusual protocol types may indicate something suspicious
- AT&T detected Slammer and other attacks *before* they hit

Packet Telescopes

- Listen for “backscatter” from DDoS attacks on unused address space
- Detect attacks, learn incidence rate, discover targets
- Note: doesn’t detect attacks from non-forged addresses. Measure that at destination; compare rates.

Dark Address Space

- No one should be trying to talk to “dark” address space
- Probes to such addresses indicate compromised or evil machines
- Learn about new attack techniques

Botnet Command and Control

- Suppose you've identified several bots attacking some target
- To whom else do those bots speak?
- A common peer who is not a victim may be the control node

Detecting Routing Attacks

- It is not currently feasible to prevent BGP attacks
- Routes to major destinations tend to remain constant
- Changes indicate either an attack or a major outage

Counting Hosts Behind a NAT Box

- Observation: the `IPid` is usually implemented as a counter.
- By detecting *approximate sequences* of `IPid`, we can detect distinct hosts.
- Packets with the same IP address but belonging to different `IPid` sequences come from different hosts.

Experimental Graph

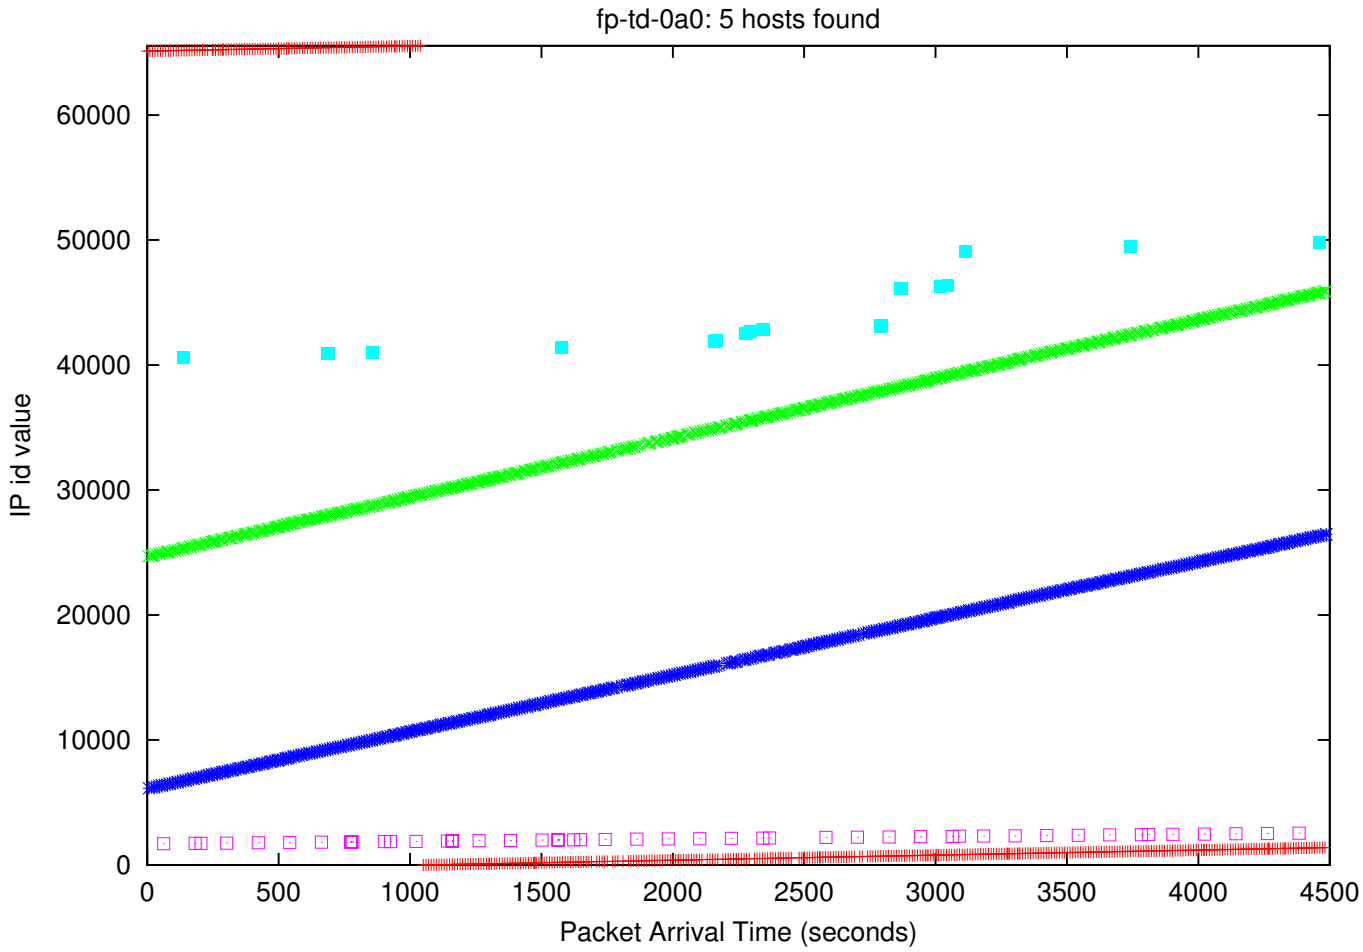

Traffic Analysis

- Many traffic patterns show through encryption
- Look at packet sizes, timing, and direction
- Very hard to hide

How Secure is a Network? A Host?

- That's the question of most interest!
- Unfortunately, we can't answer it yet
- But measurements can help

Monitor Outbound Traffic

- *Know* what a host should be sending
- Watch for differences
- Deviations may indicate misbehavior
- Outbound traffic often more significant than inbound

Do Host and Network Scans Help?

- Risk is proportional to number of exposed ports (see Microsoft's RASQ)
- Actually, proportional to weighted sum of exposed ports — some services are riskier than others
- Get weights from historical data

Measuring Client Risk

- Measure outbound traffic type
- Look for signature of particular implementations
- Again, weight sum appropriately

Is this Doable?

- Some of this is being done today
- Most is clearly feasible
- The trick is integrating all of the results

Conclusions

- Measurements can help us run a secure network
- Historical data archive almost as important as real-time measurement
- Need to develop proper statistical models