

Prolog to Lecture 12
CS 236
On-Line MS Program
Networks and Systems Security
Peter Reiher

Internet Intrusion Detection

- If intrusion detection is a great idea for my local network,
- Why isn't it a great idea for the Internet?
- Why can't I detect worms, botnets, denial of service attacks this way?

The Basic Idea

- Try to detect Internet-level problems
 - Not things just at one host or subnet
- But inherently these things relate to multiple places
- So gather data from many different sources
- Analyze the data to detect network-wide problems

An Example

- A new worm is trying to spread
- Some sites are susceptible to its techniques, some aren't
- Observe attempted spread
- Report the problem wherever it's seen
- Multiple reports of same problem allow us to deduce a worm

Why Is This Hard?

- Many reasons:
 - Do you trust all your alert sources?
 - Can you analyze fast enough?
 - Who gets to do the analysis?
 - Privacy issues
 - How much traffic will this generate?
 - Can attackers avoid detection points?
 - Who do you inform?
 - What do they do with the alerts?

Consider a Simple Case

- Assume 10,000 observation points
- Let's say a worm with a single known signature is spreading
- What if worm author prevents probing of these sites?
- What if worm is intentionally spreading slowly?
- How do you determine you've got a worm?
 - As opposed to just some bot activity?
- How many reports before you signal a problem?
- Who do you send signal to, and what to they do?

So, Good Idea or Bad Idea?

- Not really clear
- Informal, non-automated versions of this are widely used and vital
- Automated, trustworthy wide-scale system is challenging
- Not clear if we will ever get benefit from this kind of system