Customizing and Evolving Intrusion Detection

- A static, globally useful intrusion detection solution is impossible
  - Good behavior on one system is bad behavior on another
  - Behaviors change and new vulnerabilities are discovered
- Intrusion detection systems must change to meet needs

How Do Intrusion Detection Systems Evolve?

- Manually or semi-automatically
  - New information added that allows them to detect new kinds of attacks
- Automatically
  - Deduce new problems or things to watch for without human intervention

A Problem With Manually Evolving Systems

- System/network administrator action is required for each change
  - To be really effective, not just manual installation
  - More customized to the environment
- Too heavy a burden to change very often
- So they change slowly, akin to software updates

A Problem With Evolving Intrusion Detection Systems

- Very clever intruders can use the evolution against them
- Instead of immediately performing dangerous actions, evolve towards them
- If the intruder is more clever than the system, the system gradually accepts the new behavior
- Possible with manual changing systems, but harder for attackers to succeed

Intrusion Detection Tuning

- Generally, there's a tradeoff between false positives and false negatives
- You can tune the system to decrease one
  - Usually at cost of increasing the other
- Choice depends on one's situation

Practicalities of Operation

- Most commercial intrusion detection systems are add-ons
  - They run as normal applications
- They must make use of readily available information
  - Audit logged information
  - Sniffed packets
  - Output of systems calls they make
- And performance is very important

Practicalities of Audit Logs for IDS

- Operating systems only log certain stuff
- They don't necessarily log what an intrusion detection system really needs
- They produce large amounts of data
  - Expensive to process
  - Expensive to store
- If attack was successful, logs may be corrupted

What Does an IDS Do When It Detects an Attack?

- Automated response
  - Shut down the "attacker"
  - Or more carefully protect the attacked service
- Alarms
  - Notify a system administrator
    - Often via special console
  - Who investigates and takes action
- Logging
  - Just keep record for later investigation

Consequences of the Choices

- Automated
  - Too many false positives and your network stops working
  - Is the automated response effective?
- Alarm
  - Too many false positives and your administrator ignores them
  - Is the administrator able to determine what's going on fast enough?
- Logging
  - Doesn't necessarily lead to any action

How Good Does an IDS Have To Be?

- Depends on what you're using it for
- Like biometric authentication, need to trade off false positives/false negatives
- Each positive signal (real or false) should cause something to happen
  - What's the consequence?

False Positives and IDS Systems

- For automated response, what happens?
- Something gets shut off that shouldn't be
  - May be a lot of work to turn it on again
- For manual response, what happens?
- Either a human investigates and dismisses it
- Or nothing happens
- If human looks at it, can take a lot of his time

Consider a Case for Manual Response

- Your web site gets 10 million packets per day
- Your IDS has a FPR of .1% on packets
  - So you get 10,000 false positives/day
- Say each one takes one minute to handle
- That's 166 man hours per day
  - You'll need 20+ full time experts just to weed out false positives

What Are Your Choices?

- Tune to a lower FPR
  - Usually causing more false negatives
  - If too many of those, system is useless
- Have triage system for signals
  - If first step is still human, still expensive
  - Maybe you can automate some of it?
- Ignore your IDS' signals
  - In which case, why bother with it at all?

Intrusion Prevention Systems

- Essentially a buzzword for IDS that takes automatic action when intrusion is detected
- Goal is to quickly take remedial actions to threats
- Since IPSs are automated, false positives could be very, very bad
- “Poor man’s” version is IDS controlling a firewall

Sample Intrusion Detection Systems

- Snort
- Bro
- RealSecure ISS
- NetRanger

Snort

- Network intrusion detection system
- Public domain
  - Designed for Linux
  - But also runs on Windows and Mac
- Designed for high extensibility
  - Allows easy plug-ins for detection
  - And rule-based description of good & bad traffic
- Very widely used

Bro

- Like Snort, public domain network based IDS
- Developed at LBL
- Includes more sophisticated non-signature methods than Snort
- More general and extensible than Snort
- Maybe not as easy to use

RealSecure ISS

- Commercial IDS
- Bundled into IBM security products
- Distributed client/server architecture
  - Incorporates network and host components
- Other components report to server on dedicated machine

NetRanger

- Bundled into Cisco products
  - Under a different name
- For use in network environments
  - "Sensors" in promiscuous mode capture packets off the local network
- Examines data flows
  - Raises alarm for suspicious flows
- Using misuse detection techniques
  - Based on a signature database

Is Intrusion Detection Useful?

- 69% of CIS survey respondents (2008) use one
  - 54% use intrusion prevention
- In 2003, Gartner Group analyst called IDS a failed technology
  - Predicted its death by 2005
  - They're not dead yet
- Signature-based IDS especially criticized

Which Type of Intrusion Detection System Should I Use?

- NIST report[1] recommends using multiple IDSs
  - Preferably multiple types
    - E.g., host and network
- Each will detect different things
  - Using different data and techniques
- Good defense in depth

[1] http://csrc.nist.gov/publications/nistir/nistir-7007.pdf

The Future of Intrusion Detection?

- General concept has never quite lived up to its promise
- Yet alternatives are clearly failing
  - We aren't keeping the bad guys out
- So research and development continues
- And most serious people use them
  - Even if they are imperfect

Conclusions

- Intrusion detection systems are helpful enough that those who care about security should use them
- They are not yet terribly sophisticated
  - Which implies they aren't that effective
- Much research continues to improve them
- Not clear if they'll ever achieve what the original inventors hoped for