CS530
Intrusion Detection

Bill Cheng

http://merlot.usc.edu/cs530-s10

Intrusion Detection

- Security enforcement mechanisms are not foolproof, so we need a way of knowing when they are not working
  - or even better, before they stop working
- We need ways to detect insider misuse
  - detect suspicious activities
    - e.g., is this employee selling information?

Taxonomy for Intrusion Detection

- What is detected
  - *misuse detection* - look for "bad" behaviors
    - e.g., virus checker, spam filters - need to download new "definition files"
  - *anomaly detection* - look at behavior and detect out of profile activities
    - need to compare against a *baseline*
- Where detected
  - network based
  - host based - system logs
  - application based
- When attack is detected
  - real time
  - after the fact (post mortem)

Basis for Detecting Attack

- Systems operating normally
  - activity conforms to statistically predictable patterns
  - actions do not include attempts to subvert policy
  - actions of processes conform to the policies regarding what they are allowed to do
  - e.g., when system is under attack, will see unusual amount of denied accesses

Rating ID Systems

- *False positives*
  - normal activity flagged as intrusion
    - affects administrator workload
      - e.g., port scanners - if you don't have the vulnerability, do not raise alarm
  - e.g., spam filtering
    - filter out all HTML-only e-mails
  - too many of them - denial of service on yourself
  - "the boy who cried wolf"
- *False negatives*
  - attacks that are not detected

Anomaly Detection

- How it works
  - analyze *baseline* characteristics of system or user behavior and record
    - need to have an abstraction or a model
  - compare current characteristics and behavior against baseline and determine if it's within tolerance
    - or is it just a statistical fluctuation
  - flag differences
- Why it is hard
  - deciding how to characterize behavior so that changes reflect intrusions and not normal changes in activities
- Credit card companies do this all the time

Metrics

- Threshold metrics
  - number of failed access attempts
    - e.g., confiscate ATM card after 3 bad PINs
  - bandwidth consumed
    - e.g., can be used to detect misuses from within
- State change probabilities (Markov models)
  - requires training by analyzing normal traces (system logs)
    - there are systems that can be trained while monitoring
  - looking for transitions that don't seem to follow the normal pattern

Misuse Detection

- Whether activities or code violate site policy
  - rule based
    - e.g., if A is followed by B and if B is followed by C, flag it
  - signature based
- Problems
  - can only detect attacks known in advance
  - virus checkers are usually signature based
    - can protect against write to boot sector
  - many more false negatives (subjective definition)
    - vendor's definition?
- Strengths
  - tend to have fewer false positives

Collecting Input Data

- Audit vs. Intrusion Detection
- Network based ID
- Host based ID
- Application based ID

Network Based ID

- Often based on network sniffing
  - listening to network traffic as it goes by a sensor node
    - could be placed in routers or other network components
    - e.g., SNORT - packet sniffer
  - issues
    - placement
      - be careful with switched Ethernet
      - wireless channel can be asymmetric
    - load
      - may log only summary information to reduce load
        e.g., IP traceback
    - encrypted traffic (such as IPSec)
    - (cont...)

Network Based ID (Cont...)

- Issues (cont..)
  - determining intent
    - e.g., if a message to port 24 (SMTP) does not look like e-mail, flag it
    - e.g., in HTTP, turn on encryption (but don't really encrypt) - ID will ignore these messages
      can use this "feature" for tunneling

Host Based ID

- We have better understanding of these
  - because hosts are usually not an open system (unlike networks)
  - but breakins can be covered up easier (unlike networks)
- Scan system and application logs
- Report on system state
  - e.g., load, who are logged in
- Report activity to ID system
- Issues
  - only get what applications already put into logs
  - might not understand the intent of an action

Application Based ID

- Application determines what to report to ID system
  - based on a policy
- Drawbacks
  - requires application involvement (some applications will not report)
    - authorization functions like GAA-API can help address this limitation
- Benefits
  - application understands the objects and entities to which policies apply

Issues In Intrusion Detection

- Collecting data on and reporting events
  - interoperability issues
  - languages, e.g. CIDF
- Reducing data
  - to reduce network traffic consumed
    - consider overhead
  - summarize data
    - e.g., "0 of the following messages have been seen
    - finding relationships
  - what have you filtered out that shouldn't be filtered out?

Components of ID Systems

- Collectors
  - gather raw data
- Director
  - reduces incoming traffic and finds relationships
- Notifier
  - accepts data from director and takes appropriate action

Advanced IDS Models

- Distributed detection
  - combining host and network monitoring (DIDS)
  - autonomous agents (Crosbie and Spafford)
  - GCSS/ACK project at USGISI - professor Papadopoulos

Intrusion Response

- Intrusion prevention
  - it's a marketing buzzword
- Intrusion response
  - how to react when an intrusion is detected (or an attempt of intrusion)

Possible Responses

- Notify administrator
- System or network lockdown
  - change firewall rules
- Place attacker in controlled environment
  - quarantine
    - done with worms - no outgoing traffic from this node
    - use a Honeypot to attract unsuspecting attacker
- Slow the system for commanding processes
  - commonly used for SMTP servers - if spam is detected, slow down the connection
- Kill the process
  - often it is more desirable to suspend the process so you can examine memory