CS 526: Information Security

Physical Security

Most Forgotten Threats: Physical

- Examples:
  - NY investment bank premises were secure by day, but wide open to janitors' misbehaviors by night
  - Politician on a brief visit to data center complained that it was too cold, so the A/C was temporarily switched off during his speech. They forgot to turn it back on, and the heat cooked the equipment over the week-end
  - Lightning strike took out power, then backup generator failed to start.

Physical threats

- Heat/fire
- Cold/condensation
- EMP, power
- Vibration
- Smoke
- Dust
- Water
- Insects
- Magnetism
- Planes
- Natural disasters

- Theft
- Vandalism
- Insiders
  - Staff
  - Contractors
  - Temps
  - Suppliers

- Disease
- Strikes
- Political actions
- War

Often a Layered Structure (Perimeter mindset)

- Outermost layer: Low access requirements
  - Layer 1
- Innermost layer: Highest access requirements
  - Layer N (usually the number of layers is small)
- Must be compatible with layout of premises
  - Must be possible to walk into layer K–1 areas without walking through layer K areas
  - Typical: requirements for K–1 $\subseteq$ requirements for K

Example

- Layers are defined as physical areas
  - Layer 1 = Parking area
  - Layer 2 = Reception area, cafeteria
  - …
  - Layer N = Room containing corporate servers
- People categories
  - IT staff, employees, vendors, visitors, ...
  - Attributes: accompanied, authorized, approved, ...
  - Conditions: No cell phones, must wear badge, ...

Example (cont'd)

- Policy statements
  - "Layer K is accessible to employees, approved vendors, accompanied visitors"
  - "No visitor cell phones beyond layer K"
  - Sign at the door of a lab room of a place I visited:

    Authorized personnel only

    strictly forbidden to unauthorized personnel

    (not to mention unauthorized non-personnel)

Physical Controls for Sensitive Areas

- Barriers to physical entry
  - To people, but also to dust, pollutants, fire, …
  - Doorways, but also floors and ceilings (often forgotten)
- Office assignments that minimize risk
  - Place engineers with access to sensitive area close to it (minimizes traffic in and out of sensitive area)
  - “Need to know” policy for activities within layer N (no need to tell employees not allowed into layer)

Physical Controls … (cont'd)

- Prohibit 1 person being alone in a sensitive area
  - Require at least two (some malicious activities are easier to carry out when there is no witness)
  - Forces collusion between 2 employees (less likely)
  - When empty, must be locked and alarm-protected
- Audit log of all access authorizations granted
  - To visitors, vendors, employees, …
  - Record contains who made the decision, and why

Physical Controls … (cont'd)

- Record times of secure-area entries and departures of accompanied visitors and vendors
- Require wearing of visible ID badge
  - Require employees to challenge anyone without it
- Keep equipment managed by outside parties in a separate area
  - Not in same room as organization's own equipment
- Procedures for timely revocation of access
  - E.g., for someone who ceases to be an employee

Protection from Accidents

- Detection equipment for fire, flooding
  - Monitor temperature, smoke, humidity
  - Detection equipment can generate alarms
- Automate alarm-handling to minimize damage
  - Example: Automatically disconnect power to computers *before* the sprinkler system comes on (including power from the uninterruptable power supply), as damage from water is more extreme if the computer was on

Protection from Accidents (cont'd)

- Store supplies of flammables far from servers
  - Far = Separated by distance or fire-grade barrier
  - Includes printer paper, magnetic tapes, plastic, cleaning supplies, …
  - No more than 1 day supply of printer paper in the server room
- Limit electric power used in server rooms
  - Avoids buildup of heat and static electricity

What needs protection?

- Computers
- Storage
- Printouts
- Communications lines
- People
- Meeting areas
- Power
- Cooling

Protecting Communications

- Encryption works for some content, but there is still a concern about disruption
- Armor around cables, alarms
- Protecting against interception of wireless signals

Disposal of Hardware

- Destroy the information on it (= sanitize it)
- The old days: Physical destruction
- Discontinued because of environmental rules:
  - Incineration
  - Acid bath
- Still allowed:
  - Crushing
  - Increasingly ineffective because of data density

Disposal of Hardware (cont'd)

- Today: Sanitization is mainly software-based
- Some areas of hard drive are reserved, hidden from user
  - Used for, e.g., testing, bad-block remapping, …etc
  - Need special software to get to them
  - Data in bypassed bad blocks can survive
- Printer memory must also be sanitized
  - Can contain confidential information

Disposal of Paper Documents

- The threat of “dumpster diving”
- Must use a verifiable destruction process
  - Certificate of Destruction (when and where info)
  - E.g., “sent at time t1 to approved and bonded shredder X, destroyed at time t2 by X”
- Collection process
  - Use special bins to collect
  - Mark the bins? (probably not)
  - Lock the bins? (probably yes)

Disposal of Paper Docs (cont'd)

- Destruction process
  - Recycle, shred, or burn
  - Use of bonded service-providers (recycler or paper mill, on-site or off-site shredder, …)
  - Contract with service-provider specifies method used, maximum time between collection and destruction (promptness), safeguards, penalties (require provider to have liability insurance), obligation to provide certificate of destruction

Protecting Employees at Physical Risk

- Employees with privileged access (physical or electronic) are at risk of physical coercion
- Physical switches for signaling duress situation
  - Physical button or switch (preferably foot activated, so it can be used without being noticed)
  - Duress-alarm use always alerts internal security and law enforcement
  - Should it be silent or loud? (probably silent)

Signaling Duress with Access Codes

- Duress access codes, e.g.,
  - Entering regular code causes door to open
  - Entering duress code also causes door to open, but silently alerts security and law enforcement
- By Kerckhoff's principle, adversary knows of existence of a duress code
  - Adversary may ask for both, and which one is duress
  - What should be the policy for providing an answer? (Probably "reveal both, randomize which is duress")

Signaling Duress with Biometrics

- Example of a bad duress signaling design
  - Normal: Finger is at 90 degress to pad's edge, or "apply normal pressure", or "use right hand"
  - Duress: Finger is at 110 degrees to pad's edge, or "apply extra pressure", or "use left hand"
  - Adversary knows system and demands use of normal
  - Angle or pressure have possibility of false alarms (e.g., if user is in a hurry), less so for right/left
- Somewhat better
  - Employee are randomly assigned a "left" or "right" label that determines which hand is duress
  - Adversary runs a 50% risk of triggering an alarm

Detecting Physical Intrusion

- Human guard patrol
  - Walk along perimeter, corridors
  - Video monitoring (alarm is raised by human watching the different video screens)
- “Burglar alarm” systems
  - Door/window/“break glass” sensors, motion detectors, pressure sensors for floors and stairs
- Special issues if multi-tenant building
  - Ceilings and floors lead to “outside the premises”

Extreme Events: Disaster Planning

- Total physical destruction of data center
  - Fire, earthquake, tsunami, meteor, sabotage, …
  - Probability p of occurrence is small (but positive)
  - If resulting loss is C, then expected loss = p * C
  - For many organizations, C is ∞ because it means cessation of operations and bankruptcy
- How to prepare for such a disaster?
  - Redundant mirror facilities (physically remote)

Disaster Planning (cont'd)

- Death or incapacitation of all key personnel
  - Airplane crash
  - Food poisoning at the same business dinner
  - lightning strike when playing golf
- To mitigate, buy insurance policies
  - All key and "difficult to replace" personnel
  - The policies are bought by the employer, and pay the employer in case of mishaps
- Promote health of employees

CS 526: Information Security

Intrusion Detection

Definitions

- **I**ntrusion **D**etection **S**ystem (IDS)
  - Like a burglar alarm for hosts and networks
  - Watches for signs of break-in and misuse
- Sources of information collected and used by IDS
  - Applications (e.g., DBMS)
  - Host (audit trails, system logs, system state)
  - Network (e.g., packet sniffing, network devices)
- IDS analyzes the information, can issue an alarm

Terminology

- Attack from insider is often called misuse
  - Insider attacks have different characteristics than attacks from outsiders
  - E.g., insider can spread attack over longer time span
- Audit log = a time-ordered sequence of events
  - A record of what happened
  - Different logs for OS, network, applications, …
  - Potentially huge amounts of data
    - Use reduction (summarization) techniques?

Why IDS ?

- First (and preferred) line of defense is prevention
  - Better to prevent intrusion
  - Use the proper security controls
- Unfortunately, prevention sometimes fails
  - Software bugs
  - Malware
  - Misconfiguration
  - Human error
- IDS is a second line of defense

What IDS Does

- Detects unauthorized access to resources
- Detects violations of policies
- Detects placement or presence of malware
- Detects attacks (including denial of service)
- Detects abnormal patterns of activity
- Detects misconfigurations
- Facilitates security management

What IDS does (cont'd)

- Facilitates post-intrusion analysis of events
  - Damage assessment
  - Blame/credit assignment
  - Prevention of future re-occurrence of attack
- Gives customers and business partners enhanced sense of security
  - More likely to gain their trust
  - Less reluctance to share their info with you

IDS and Liability

- Intrusions can cause serious liability, e.g.,
  - Cause harm to users
  - Cause harm to customers
    - Steal credit card info, SSN, DoB, …
  - Cause harm to strangers
    - Attackers use system to attack others
- IDS reduces liability
  - Makes harm less likely
  - Having IDS is practicing “due diligence”

Some Caveats

- IDS is a supplement, not a substitute, for the other security techniques
- Human intervention still needed
  - Investigation
  - Identification of culprit(s)
  - Response
- Does not deal with all forms of attack
- Can give a “false sense of security”

Some Caveats (cont'd)

- Can be swamped by huge traffic levels
  - In which case adaptively respond by (e.g., give up exhaustive analysis, filter and prioritize, …)
- Problems if "garbage in"
  - If audit trail is corrupted by attacker
  - IDS should rely on multiple data sources
    - Redundancy

Interval-operation

- Less taxing on system than continuous (“real time”) operation
  - Judicious choice of timing
- Scanners
  - Pinpoint existing weaknesses
  - Determine possibility of a future attack, past occurrence of an attack
  - Do not (usually) detect an attack in progress
  - Check system response to intrusion scripts

Real-time operation

- Continuous monitoring
  - application, user, system, network, …
- Detect an attack in progress, and report it
- Respond in real-time?
  - error, liability, …
- Expensive
  - Performance and usability can degrade
  - Resource-intensive (esp. memory and CPU)

Detection engines

- Anomaly based
  - Notions of “normal”, “abnormal”
  - Need not know exact pattern of attack (works against new, unknown attacks)
- Signature based
  - Compares to known attack signatures
  - Need to update attack-signatures database

Detection engines (cont'd)

- Signature based is more common
  - Preferred by system administrators (no false alarms, easier to respond to an alarm)
  - Sometimes supplemented with anomaly-based detection
- Rarely is anomaly-based deployed without signature-based
  - Cannot replace signature-based

Anomaly Based IDS

- Relies on a set of variables
  - Normal and abnormal values
- How to give IDS notion of “normal”
  - Learned, based on history
  - Declared (no learning)
- Statistical techniques
  - Computes scores, compares to thresholds

Difficulties with Anomaly Based

- Alarm can be difficult to analyze
  - Is anomaly caused by intrusion, or accidental?
- Math issues
  - Statistical dependencies, curse of dimensionality
  - Make tractable => dubious statistical assumptions
- Computationally expensive
- Prone to false alarms (caused by variability)
- Can be circumvented

Difficulties with Anomaly Based (cont'd)

- Coarse (due to averaging)
- Time sequence of events is ignored
- Thresholds are hard to determine
  - And can often be circumvented (by, e.g., spreading the attack over time)
- Self-adaptability is double-edged sword
  - Insider can train it to learn a "new normal" in which misuse is not flagged

Examples of anomalies

- Abnormally high rate of password failures
  - In a specific account, or in the system as a whole
- Unusual login times for a user (3am), and unusual system usage thereafter
  - E.g., excessive browsing of directories and executing system status commands, and none of the usual editing and compiling (user account may have been compromised, or the user may be misbehaving)
- Burst of re-writing of executable files
  - Could be indicative of malware spreading

Examples of anomalies (cont'd)

- Abnormal sequence of system calls made by a process
  - Build a database of sequences of system calls made by process in normal behavior (using a fixed-width sliding window, e.g., of width 6)
  - When a security hole in the program is exploited, sequences not in the database start appearing

  [Successfully detected intrusions involving older versions of many Unix utilities (sendmail, lpr, etc)]

Signature Based

- Uses pattern matching
  - Alarm gives a precise, useful diagnostic
- Finite state machines
  - Fast (even for large patterns)
- Colored petri nets
  - Can be slow for large patterns
- Flexibility
  - Look only for what you care to detect

Signature Based (continued)

- Pattern specification in a formal language can be difficult to automate
  - Because descriptions of known attacks are documented and described using natural language
- New attacks mean a growing database of attack patterns
    - Can get expensive
    - Impact performance

Example of a Signature in an IDS

- "k or more failed login attempts in t seconds

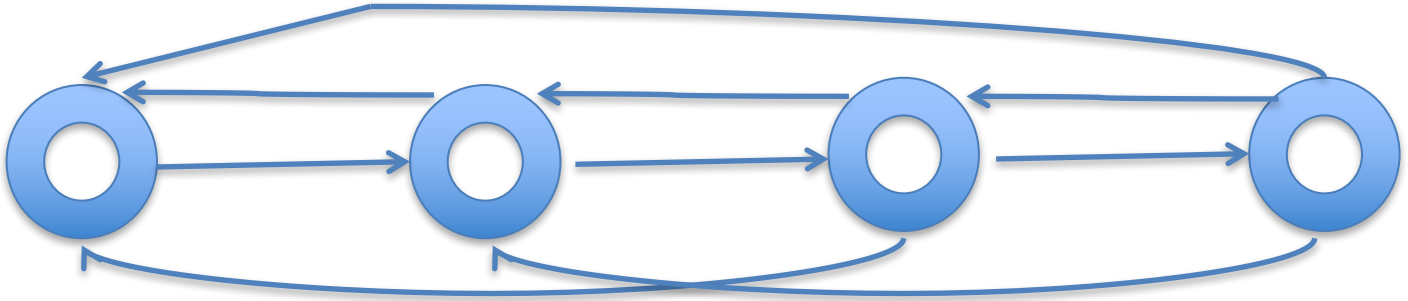

Finite state machine with k = 3

- Transitions are caused by …
  - events
  - the passing of time

Example of a Signature (cont'd)

- ***if*** (source_ip == destination_ip)

  then issue a "LAND attack" alarm

LAND attack = DoS attack, consists of sending a TCP connection-initiation packet with the target host's IP address as both source and destination. Can cause host to keep sending acks to itself (in some implementations)

Target-based monitoring

- Does not need signatures
- Some similarity to anomaly detection but doesn't require a baseline

- Key concept is to monitor targets for change or access.
  - Tripwire is a canonical example

IDPS = IDS + Automatic Response

- **I**ntrusion **D**etection and **P**revention **S**ystem
- Response examples:
  - “Terminate network connection”
  - “Change access control list on router to block an IP address”
  - “Freeze user account”
- Immediate automatic response
  - Can result in errors (and liability issues)

IDS Issues and Tradeoffs

- Security vs performance & usability
- Security vs cost of keeping IDS running
- Vendor lock-in
  - Constant upgrades (e.g., signatures database)
- Monitoring vs privacy
  - Inform employees of monitoring
- Real-time response vs possibility of error