CS259D: Data Mining for Cybersecurity

Cybersecurity

- "Information security as applied to computing devices such as computers, smartphones, as well as computer networks…"

- Information Security (InfoSec): "practice of defending information from unauthorized access, use, disclosure, disruption, modification, perusal, inspection, recording, or destruction".

Information security goals

- C-I-A triad
  - Confidentiality
    - Unauthorized disclosure of information
  - Integrity
    - Unauthorized modification of information
  - Availability
    - Unauthorized withholding of information or resources
- Others
  - Privacy
  - Authenticity
  - Non-repudiation
  - Accountability
  - Auditability

Risk management: Controls

- Administrative
  - Policies, guidelines
    - Password policies
    - Payment Card Industry Data Security Standard (PCI DSS)
    - Principle of least privilege
- Physical
  - Doors, locks, etc.
  - Principle of separation of duties
- Logical
  - Use software and data

Preventive measures

- Protocols
  - Secure Socket Layer (SSL): source authentication
- Host-based protections
  - Secure operating systems, Patching
- Access control
  - Identification: username
  - Authentication: Something you know/have/are
  - Authorization: File permissions, Kerberos, Need-to-know principle

Preventive measures

- Firewalls
  - Control inter-network traffic (e.g., from/to internet)
- Security by design
  - Principle of least privilege, Code reviews, unit testing, Defense in depth
- Secure coding
  - Buffer overflows, Format string vulnerabilities, Code/Command injection

Why is prevention not enough?

- Inherent weaknesses in increasingly more complex systems/networks
  - Poor Design
    - Software and hardware
    - Example: sendmail (race condition vulnerability, buffer overflow, group permission vulnerability, etc.)
  - Poor Implementation
    - Security an afterthought
    - Lack of personnel experience/training
    - Poor system configuration
      - Example: default firewall configurations with open insecure ports
  - Poor Management
    - Inadequate policies/procedures

Why is prevention not enough?

- Tradeoff between security and usability
- Non-tech tradeoffs in system engineering
  - social, organizational, economic, regulatory, legal
- Cost of prevention

Vulnerabilities

- Backdoors
  - Kleptographic attack
  - Rootkit
- Denial of Service
  - Resource exhaustion
  - Attack amplifiers (e.g., poorly designed FTP, DNS)
  - Application or OS exploit
- Eavesdropping
  - Listening to private communication on network
  - Monitoring hardware electro-magnetic transmissions
- Exploits
  - Gain control of a computer system, allow privilege escalation, or denial of service attack
  - Used in Trojan horses, viruses
- Social Engineering
  - Humans: the weakest link in security

Attack categories

- Probe
  - Information gathering (1:1, 1:m, m:1, m:n modes)
  - IPSweep, portsweep, nmap, etc.
- Denial of Service (DoS)
  - TCP SYN flood, Ping of Death, smurf, neptune, etc.
- Remote to Local attacks (R2L)
  - Brute force/Dictionary attack, buffer overflow, unverified input attacks
  - Social engineering, Trojans
- User to Root attacks (U2R)
  - Buffer overflow, rootkit, etc.
- Infections
  - Trojans/worms/viruses
  - Spreading attacks

Basic attack steps

- Prepare
  - Gather info: Valid IP addresses & ports, OS, software type & version
- Exploit
- Leave behind
  - Backdoors
- Clean up
  - Restart crashed daemons, clean registry/log files
- Variable order and duration
  - Attacker's skill level
  - Type of vulnerability to exploit
  - Prior knowledge
  - Starting location of attacker

Attack tools

- Information gathering
  - Sniffing: capture packets traversing network
    - Tcpdump, Ethereal, Gulp, Net2pcap, Dsniff, etc.
  - Network mapping/scanning/fingerprinting: hosts/IPs/ports, protocol details
    - Nmap, Amap, Vmap, Ttlscan, P0f, Xprobe, Queso, etc.
- Attack launching
  - Trojans
    - Danger, NukeNabbler, AIMSpy, NetSpy, etc.
  - DoS attacks
    - Targa, Burbonic, HOIC, LOIC, etc.
  - Packet forging tools
    - Packeth, Packit, Packet Excalibur, Nemesis, Tcpinject, Libnet, SendIP, etc.
  - Application layer tools
    - Code Red Worm, Nimda Worm, AppDDoS, RefRef, etc.
  - User attack tools
    - Ntfsdos, Yaga, etc.

Failure of prevention

- Attacking constantly getting easier
  - Required expertise decreasing
  - Quality of attack tools increasing

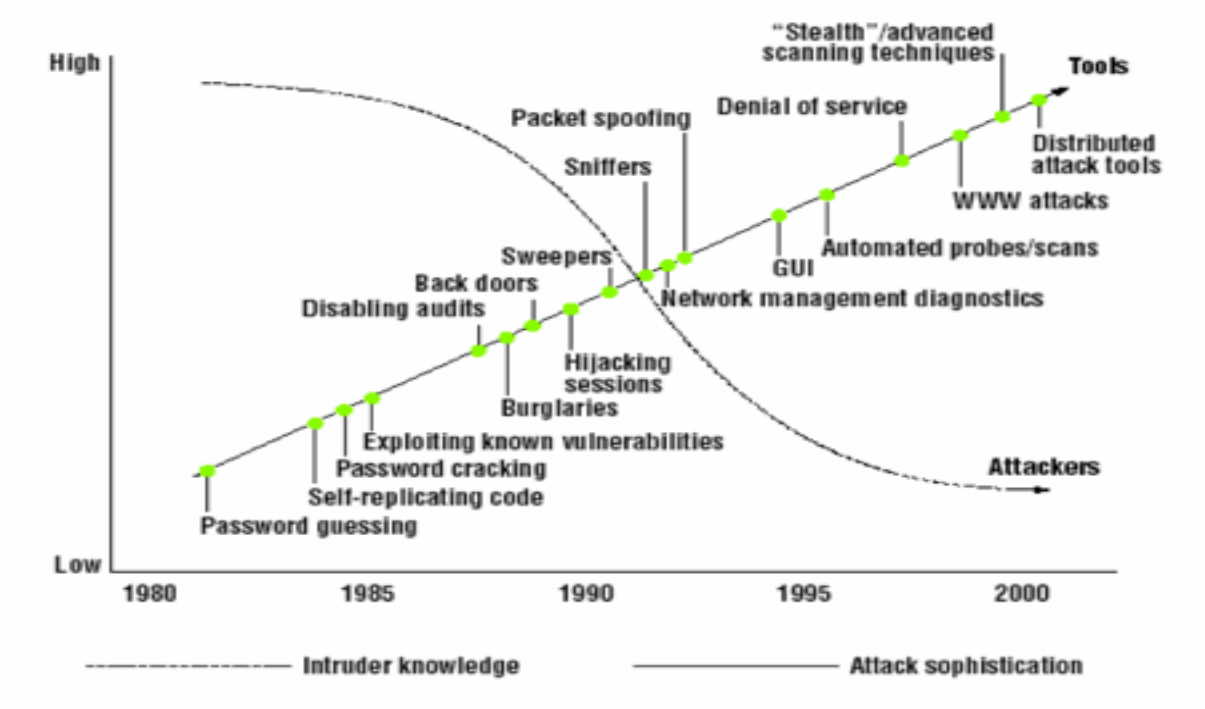

Failure of prevention

- Attack mechanisms constantly evolving/ mutating
- New connectivity options bypassing perimeter security
- Preventive measures getting obsolete
  - Polymorphic viruses resistant to current antivirus

Security measures

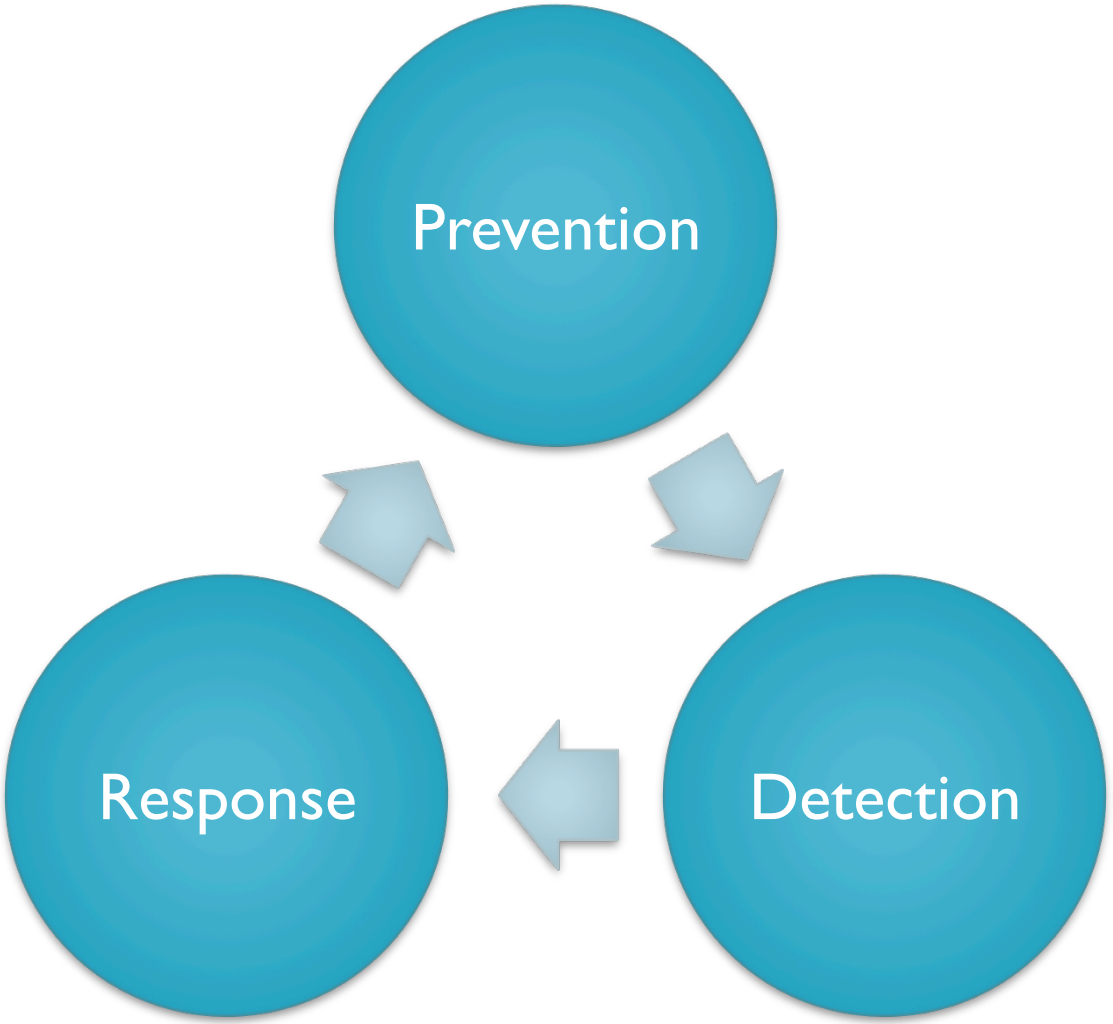

Defense in depth

- Layered approach
  - Separate systems into network sections
  - Place firewalls at section boundaries
  - Border router between ISP and firewall to filter traffic
  - Switches on each section to make sniffing less effective
  - Encryption
- Last layer of defense
  - Detection

Detective security

- 1st Generation: Intrusion detection systems (IDS)
  - 100% protection/prevention impossible
  - Layered security
- 2nd Generation: Security information and event management (SIEM)
  - Correlate alerts from different intrusion detection sensors
  - Present actionable information to security analyst
- 3rd Generation: Big Data analytics for security
  - Contextual security intelligence
  - Long-term correlations

Why now?

- Attack landscape
  - Attacks increasingly more sophisticated
  - Required attacker knowledge going down
  - Highly motivated attackers
    - Attacker needs to succeed only once, defense needs to be right every single time
- Current detection techniques failing
  - Polymorphic malwares
  - Zero-day attacks
  - APTs
- Network perimeter dissolving
  - Mobile/BYOD
  - Cloud

Why now?

- Big Data technology enable storage and analysis of higher volumes & more types of data
- 2010 Verizon data breach investigation
  - In **86%** of cases of breach, evidence was in the logs
  - Detection mechanisms failed to raise alerts
- How do we make sense of the data?

Attack landscape

"There are only two types of companies: those that have been hacked, and those that will be. Even that is merging into one category: those that have been hacked and will be again."

Robert Mueller
FBI Director

Attackers and motivations

- Script-kiddies
  - Motivated by curiosity
- Cybercriminals
  - Motivated by profit
  - Typical demographics: east European, Brazilian
- Nation-state hackers
  - Motivated by power
  - Typical demographics: east Asian, middle eastern
- Hacktivists
  - Motivated by ideology
  - Typical demographics: north American, western European
- Cyber-mercenaries
  - Hired by to attack
- Insiders
  - Motivated by disgruntlement

Reactive defense

- Examples
  - Antivirus signatures for known malicious executables
  - Email filters for unwanted messages
  - Web filters for compromised websites
  - Sandboxes for malicious behaviors
- Median detection time between intrusion/breach to awareness of it: 300-400+ days
- Duration of zero-day attacks
  - 19 days to 30 months
  - Median of 8 months, Average of 10 months
- 61% of attacks discovered by a third party
- Businesses reluctant to disclose their breaches
  - Only 2%-30% do
- Porous perimeter
  - Cloud applications
  - Mobile/BYOD
  - Partner businesses

Explosion of malware

- 403 million new variants of malware created in 2011
- 100,000 unique malware samples collected daily by McAfee in 2012, Q1
- More than 100 million samples in McAfee's malware signature database by 2012 Q3
- Practically impossible to keep up with signatures

Targeted industries

Industry	% Advanced attacks
Aerospace and defense	17%
Energy, Oil & Gas	14%
Finance	11%
Computer software and hardware	8%
Legal and consulting services	7%
Media and Entertainment	7%
Telecommunications	6%
Pharmaceuticals	4%
Other	25%

Some recent attacks

- 2006: 10-20 TB data stolen from US NIPRNet Military Network
- 2007: Massive cyber-attack against Estonia
- 2008: ExxonMobil, ConocoPhillips, Marathon Oil; all unaware till alerted by FBI
- 2010:
  - Operation Aurora against Google & 20+ others, e.g., Yahoo, Morgan Stanley, Symantec, Northrop Grumman
  - Stuxnet: "The world's most advanced malware"
- 2011:
  - RSA SecurID breach, Lockheed martin attacked consequently
  - Comodo, DigiNotar certificate authorities breached
- 2012: Flame deemed most complex malware ever created, Red October
- 2013:
  - Adobe breach: 152 million customers' data including passwords stolen
  - Target breach: 40 million credit cards, 70 million addresses, phone numbers, etc, $61M spent to respond to breach, costs potentially in $Billions
- 2014:
  - eBay breach
  - Home Depot breach: 56 million credit and debit cards compromised
  - JPMorgan, 4 other banks: GBs of data including checking and savings account info stolen

Advanced Persistent Threats: APT

- Targeted attack against a high-value asset
- Low and slow
- Avoid alerts
  - Use stolen user credentials
  - Zero-day exploits
  - Low profile in network
  - Slow progress: Operating over months or years
  - Beyond limited correlation time windows of today's IDSs
- Multi-stage
  - Exploitation
  - Command and control
  - Lateral movement
  - Breach

Advanced Persistent Threats: APT

- Typical Goals
  - Steal intellectual property (IP)
  - Gain access to sensitive customer data
  - Access strategic business information
    - Financial gain, embarrassment, blackmail, data poisoning, illegal insider trading, disrupting organization's business
- Attackers
  - Well-funded
  - Highly skilled
  - Motivated
  - Targeted on specific data from specific organization

Administrativia

- Course URL: http://web.stanford.edu/class/cs259d/
- Instructor:
  - Bahman Bahmani
  - Email: bahman@cs
  - Office hours: TBD
- TA:
  - Dima Brezhnev
  - Email: brezhnev@cs
  - Office hours: TBD
- No textbook
- Guest speakers
- Week of October 13

Grading

- 4 Homework assignments
  - Individually or in pairs
  - Almost 2 weeks for each homework
  - Each 25% of the grade
  - Potential assignments:
    - Web attack detection
    - User profiling for authentication and authorization
    - Network profiling and intrusion detection
    - Botnet detection
    - Host-based insider threat detection
    - Deep packet inspection
    - Web proxy log analysis
    - Algorithmic alert correlation
- No exams
- No late days
- Honor Code

Detection Taxonomy

- Information source
- Analysis strategy
- Time aspects
- Activeness
- Continuality

Detection taxonomy: Information source

- Host-based
  - system calls, system logs
- Network-based
- Wireless Network
- Application logs
  - DB logs, web logs
- IDS sensor alerts
  - Lower level sensor alarms

Detection taxonomy: Detection strategy

- Misuse detection
  - Premise
    - Knowledge of attack patterns provided by human experts
    - Signature matching
    - Data mining using labeled data sets
  - Benefit: high accuracy in detecting known attacks
  - Drawbacks:
    - Ineffective against novel attacks
    - Signatures need updates with each new discovered attack
- Anomaly detection
  - Premise
    - Build profiles of normal behavior (users, hosts, networks)
    - Detect deviations from normal profiles
  - Benefit: detect novel attacks
  - Drawback: Possible high false alarm rate

Detection taxonomy: Time aspects

- Real-time
  - Analyze live data (e.g., session data)
  - Raise alert immediately if attack detected
- Offline
  - Analyze data offline
  - Useful for forensics

Detection taxonomy: Activeness

- Passive reaction
  - Only generate alarms
  - Benefit: Human in the loop
  - Drawback: Alert may go unnoticed
    - Example: Target breach
- Active response
  - Corrective response (e.g., reconfigure firewalls)
  - Proactive (e.g., log out attacker)
  - Benefit: Speed
  - Drawback: May turn into DoS attack against

Detection taxonomy: Continuality

- Continuous monitoring
  - Continuous real-time analysis
  - Collect information about actions immediately
  - Higher deployment effort
- Periodic analysis
  - Take periodic snapshots of the environment
  - Lower security: Exploitation between two snapshots

Example data mining method

- PageRank
  - Developed by founders of Google
  - Used for search ranking, recommendation systems, etc.

Example: Detect malware-infected hosts

- Build host-domain access graph from web proxy logs
- Seed the analysis using minimal ground truth
  - a blacklist & a whitelist of known bad & good sites
- Belief-propagation (PageRank-style): estimate likelihood of a host/domain to be malicious
- An instance of semi-supervised learning

Example data mining methods

- Hubs and Authorities
  - Similar to PageRank
- Clustering

Example: P2P botnet detection using netflow modeling

- Represent each host (IP address) as a node in a graph
- Edge (A,B) if & only if a flow from A to B exists
- Compute Hubs and Authorities scores
- Bots show similar hub-authority characteristics
- Do clustering in the 2-dimensional space (hub, authority)

Example data mining method

- Frequent itemset mining

Example: Beaconing pattern detection

- Web proxy alerts repeated access to suspicious IP
- Perform frequent itemset mining on sets of events prior to those IP accesses
- Frequent authentications to critical DB detected

Example: Behavior profiling for APT detection

- Anomaly sensors to detect specific deviations
  - Unusual connections from a host to external IPs
  - Profile set of machines each user logs into
  - Model each user's normal working hours
  - Model flow of data between internal hosts (e.g., to detect staging servers before data exfiltration)

Syllabus (tentative)

- **Introduction:** Introduction to Information Security, Introduction to Data Mining for Information Security
- **Malware Detection:** Obfuscation, Polymorphism, Payload-based detection of worms, Botnet detection/takedown
- **Network Intrusion Detection:** Signature-based solutions (Snort, etc.), Data-mining-based solutions (supervised and unsupervised), Deep packet inspection
- **Host Intrusion Detection:** Analysis of shell command sequences, system call sequences, and audit trails, Masquerader/Impersonator/Insider threat detection
- **Web Security:** Anomaly detection of web-based attacks using web server logs, Anomaly detection in web proxy logs
- **Email:** Spam detection, Phishing detection
- **Social network security:** Detecting compromised accounts, detecting social network spam
- **Authentication:** Anomaly detection of Single Sign On (Kerberos, Active Directory), Detecting Pass-the-Hash and Pass-the-Ticket attacks
- **Automated correlation:** Attack trees, Building attack scenarios from individual alerts
- **Issues:** Privacy issues, Adversarial machine learning (use of machine learning by attackers, how to make ML algorithms robust/secure against adversaries)
- **Other potential topics:** Fraud detection, IoT/Infrastructure security, Mobile/Wireless security

Job market

- Security analyst
- Security engineers
- Security architect
- Security administrator
- Chief Information Security Officer
- Security consultant
- **Security Data Scientist**