Intrusion Detection

Issues and Technologies

Julie J.C.H. Ryan, D.Sc.

Presented to the Department of Veteran Affairs

InfoSec2002 New Orleans May 2002

Detecting Intrusions

- Detection is one phase of security engineering
  - Protect --- Detect --- React/Correct
  - Important to detect problems in environment
    - Rarely do problems exist in only one domain
  - Detection plan must include all security domains
    - Including but not limited to networks
- Things to detect
  - Occurrence of problems
    - Both those that you protected against and those you didn't protect against
  - Failure of protection mechanisms
  - Use of counter-detection methods or techniques by adversaries

Intrusion Detection Systems

- Critical tool for detecting intruders in networks and systems
- Key things to consider:
  - What isn't a network today?
  - How much data is processed over networks?
  - If you tracked every single action in a computer system:
    - How much data would you collect?
    - How much processing overhead would you incur?
    - How would you analyze the data that you did collect?
  - If you decided not to track every single action:
    - Which ones would you track?
    - How would you protect the collection & storage processes?
    - How would you protect the analysis & reporting processes?

Types of IDSs

- Basically two types:
  - Manual
  - Automatic
- Manual
  - Someone sitting at a terminal monitoring the activities that are going on
  - Chances of success predicated on looking in the right place at the right time
  - Usually as a result of some sort of suspicion or tip
- Automatic
  - Collecting data on everything and running it through analysis tool

Manual Examples

- Cliff Stohl
  - "Cuckoo's Egg"
  - Used printers, beepers, and manual analysis
- Bellovin and Cheswick
  - "Firewalls and Internet Security"
  - Used manual traps, tricks and analysis
- Tsutomo Shimomura
  - "Takedown"
  - Used manual traps, tricks, and analysis augmented by specially designed technologies
- What's an ordinary net admin to do?

Essential Elements of IDS

- Success intrusion detection can be based on 1 or more of:
  - Audit trail processing
    - What data can be derived from an audit log of system activity to detect intrusions or access violations?
  - On-the-fly processing
    - What data is available right now to provide clues to behavior?
  - Profiles of normality
    - What's okay behavior, anyway? What predictive descriptions can be used to help identify abnormal behavior patterns?
  - Signatures of abnormal behavior
    - What's not okay behavior? How can it be spotted? Patterns?
  - Parameter pattern matching
    - Can patterns be derived that identify intrusions?

Audit Trail Processing

- System operators identify security critical events
  - When one of these events occurs, a record of that event is stored
    - Who, when, what, where, etc
  - Examples of candidate auditable events:
    - File open, file read, file write, file delete, file privilege modification
    - File creation, file removal, failed executions
    - Login attempts, unsuccessful logins, change of password
    - Adding a user, removing a user, changing a user's privileges
    - Adding a group, changing a group, deleting a group, changing a group's privileges
- Problem:
  - The more events are audited, the more data is generated
  - The more data is generated, the more data needs to be analyzed

Audit Log Analysis

- When looking at audit logs from a system, look for:
  - Users logging on at strange hours
    - Assuming you have some clue as to what normal is, and what time zone the user is currently in
  - Unexplained reboots
  - Changes to system clocks
  - Unusual error messages from mailers, daemons, or servers
  - Failed logins based on bad passwords
    - Especially multiple attempts
  - If Unix, unauthorized use of *su* command
  - Users logging in from unfamiliar sites

On-the-fly processing

- Fastest analysis of data collected in real time
  - What data is available right now to provide clues to behavior?
  - Data collected of security events is analyzed immediately
    - Limiting functions include processing speed and memory
- Attributes
  - Provides indications and warning of potentially harmful events, giving system time to protect against real harm or damage
  - Analysis is simplistic in nature in order to maximize speed
    - Complicated analyses can't be performed because they require too much data and take too long to run
  - Some data may be lost in rush to process
  - Capacity is much smaller than off-line analysis functions

Real Time Processing

- Data has to be gathered
  - Packet diversion to analysis function
    - By firewalls, routers, gateways, etc
    - Based on headers, routing tables, access lists, or other applications
  - Network sniffing
    - Tap the network by listening in promiscuous mode
      - Accepts all traffic along a certain pathway
    - Legal restrictions apply in some circumstances (like at an ISP)
- Data has to be analyzed, fast
  - Limited to looking for actual actions in real time
  - Can't detect subtle long term attacks that are in the noise

Profiles of normality

- Based on understanding how users operate normally
  - What's okay behavior, anyway?
  - What predictive descriptions can be used to help identify abnormal behavior patterns?
- Implies
  - Database of user profiles
    - Which must be protected from modification
    - Which is resistant to "training"
  - Comparison of expected activity to actual activity
    - Some time delay in this analysis
    - The allowable deviance must be understood ahead of time

The IDES Model

- Published in mid 1980s by Dr. Denning
- An intrusion detection system is comprised of a six-tuple mathematical object
  - Subjects, objects, profiles, audit records, anomaly records, alarms
  - What these are
    - Subjects and objects are the initiators and targets of activity
    - Profiles are the models of behavior
      - Often based on statistical models of past behavior -- 'training'
    - Audit records are captured security events or observed behavior
    - Anomaly records are programmed decisions based on intrusion analysis
    - Alarms are how the potential problems are reported
- Example: detecting toll fraud

Detecting Toll Fraud

- Subjects and Objects
  - Caller and callee (be it a person, PBX, or other end process)
- Profiles of normal usage
  - Based on past behavior patterns, expected behavior patterns
  - Some calling patterns would be abnormal no matter what
    - Late night calls to certain numbers
- Audited activities
  - Calls
    - From, when, where, to
- Automated response patterns based on previous decisions
  - Reconfiguration of service, including outbound access

Signatures of abnormal behavior

- A method of detecting attacks rather than deviations from expected behaviors
  - What's not okay behavior? How can it be spotted? Patterns?
- Two variations
  - Attack signatures
    - A known set of activities that are highly correlated with attack activity
    - Usually a set of behaviors related in time sequence
    - Good source of these are hacking sites on the 'net
  - Select character strings
    - Anti-viral software makes use of these strings to detect viruses
    - Also includes certain activities that require a set of commands in a specific order
      - Firewalls use this to detect hostile intrusion attempts

Parameter pattern matching

- Very subtle
  - Can patterns be derived that identify intrusions?
  - Relies on monitoring a wide variety of system and network attributes
  - Useful when normal operational behavior cannot easily be characterized
- Most successful implementations
  - Based on humans viewing data presentations
  - Some success with data mining applications
  - Some limited success with AI in the research community

IDS Architectural Elements

- Knowledge base
  - Potential problems, etc
- Functional components
  - Target system
  - Feed
  - Processing
  - Knowledge Base
  - Storage
  - Alarms
  - Operator Interface
  - Comms infrastructure backbone

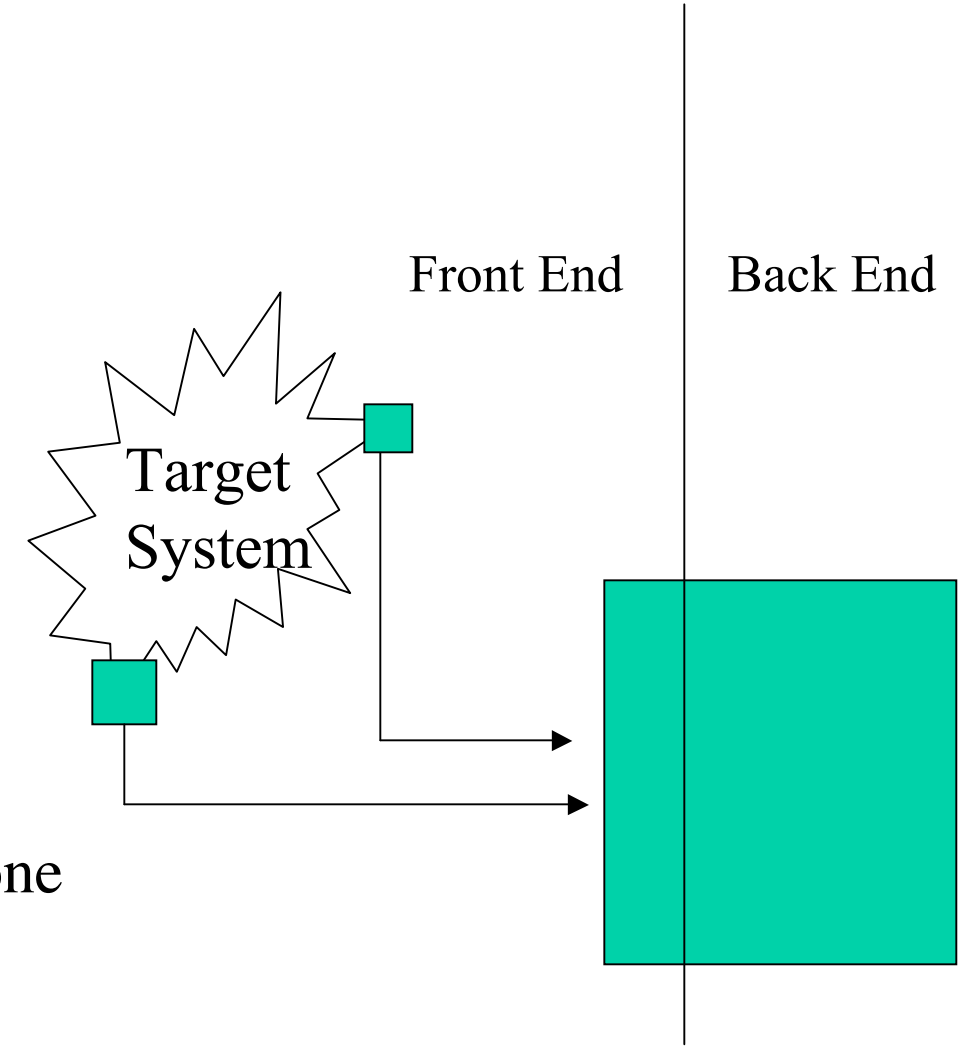

Back End Elements

- Architectural elements that support functionality
  - Knowledge base(s), Storage, Alarms, Interface, Comms Infrastructure

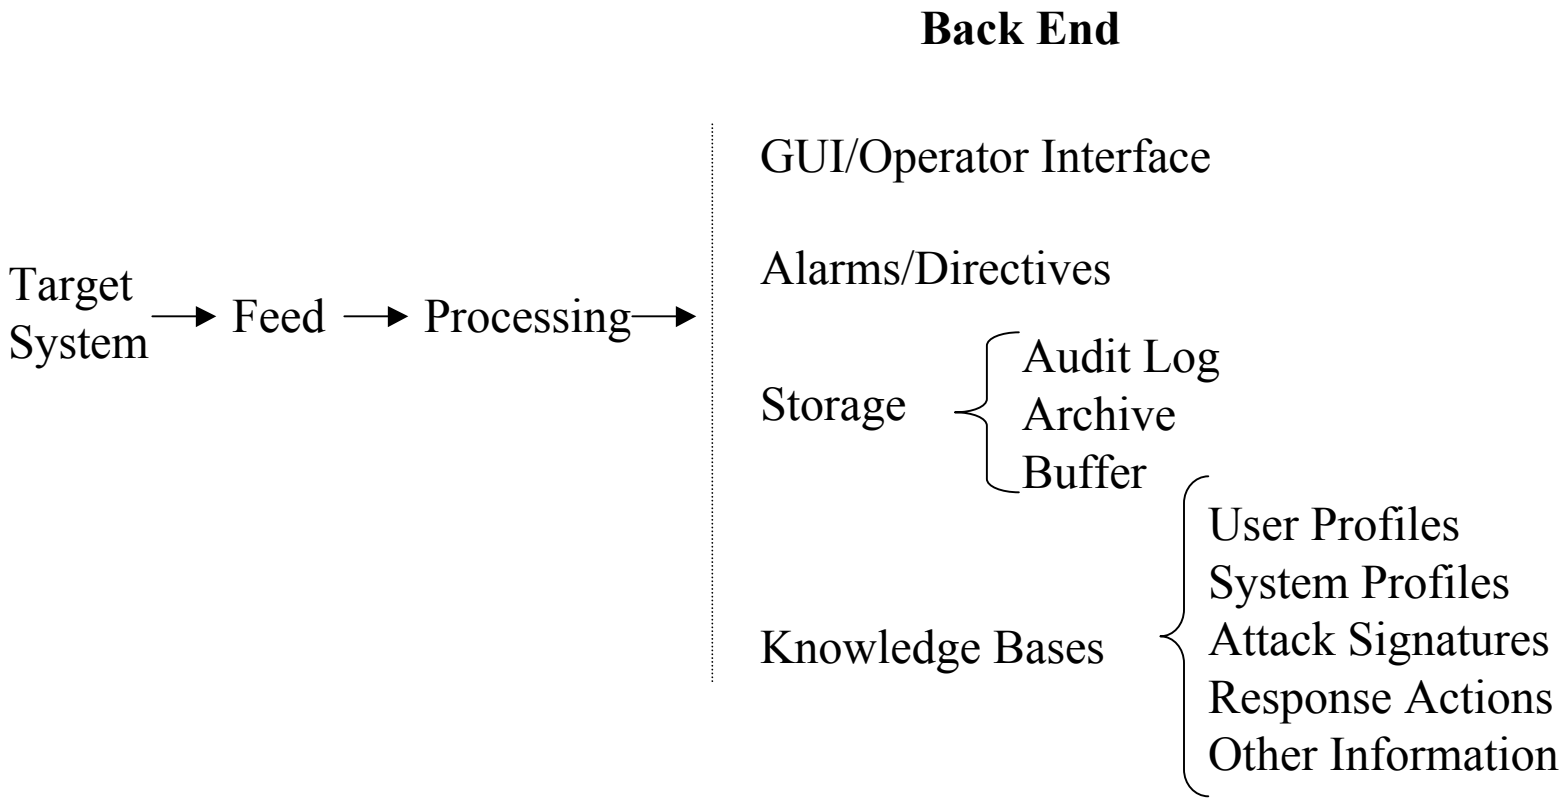

Back End Considerations

- Design Considerations
  - Knowledge Bases
    - Potentially contain normal and abnormal descriptions of user and/or system operational patterns
  - Alarms
    - Human Terminating Alarm
    - Automated Process Terminating Alarm
    - Hybrid
  - GUI/Operator Interface
    - Enable rapid operator response, present information clearly and unambiguously, avoid clutter, simple and easy response initiation
  - Comms Infrastructure
    - Reliable and secure both within the system and between components

Front End Elements

- Feeds and Processing Engines
  - Feeds transfer collected information to the processing component of the IDS
  - Processing suite contains the engines/filters, the system management, and the correlation elements

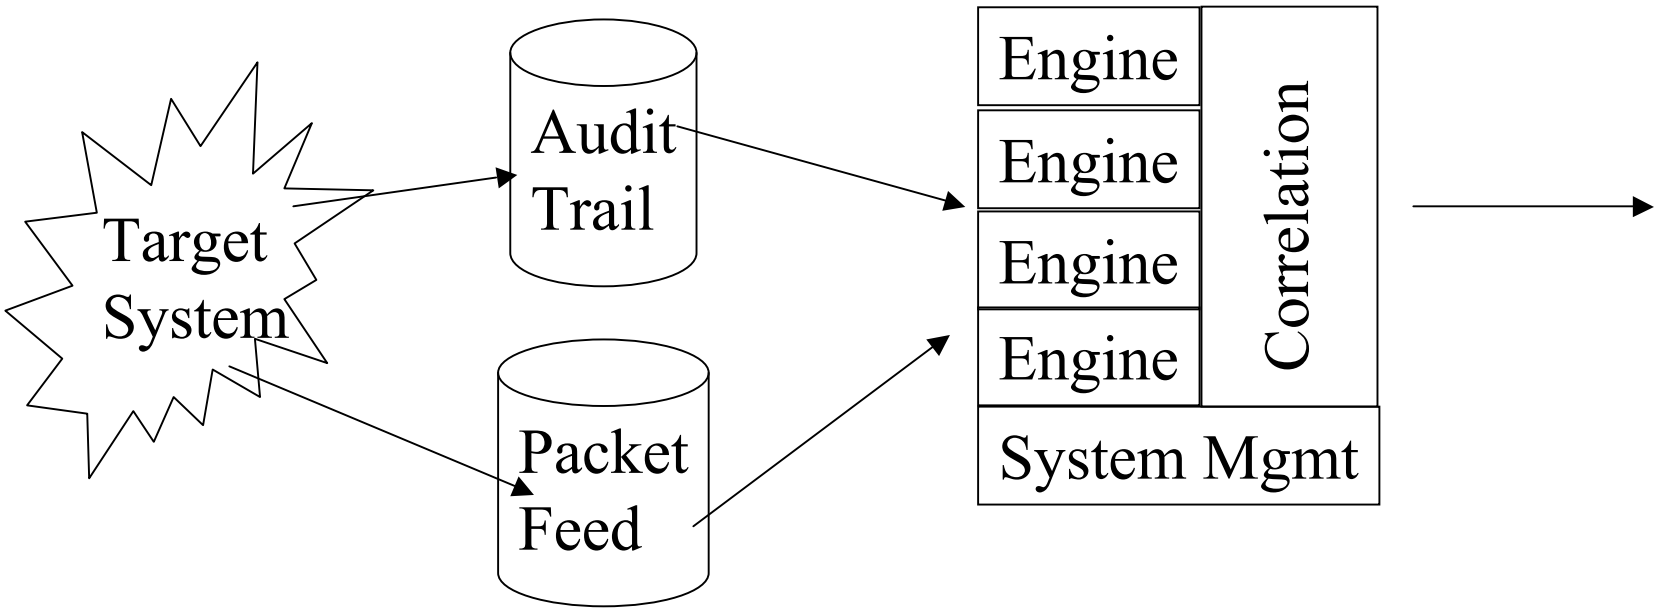

Front End Architecture

- Feed Considerations
  - Type(s) of target system
  - any real time response requirement
  - network capacity
  - target system activity
- Engines/Filters
  - Viewed as modular tasks
  - Enhances speed, enables complexity avoidance
  - Execute the philosophies of IDS discussed previously
    - Audit trail processing, On-the-fly processing, Profiles of normality, Signatures of abnormal behavior, Parameter pattern matching

Where to Put It?

- Network IDS
  - Principle activity to detect network activity that indicates unauthorized or malicious activity
  - Does nothing to detect issues that are not network oriented
- Host-based IDS
  - Principle activity is to detect actions on computers that indicate unauthorized or malicious activity
  - Does nothing to detect issues that are network oriented
- Enterprise IDS
  - An attempt to integrate intrusion detection efforts across host and network based activities to see the larger picture

Research Influences

- EMERALD Generic Monitor Architecture
  - Culmination of years of research by SRI, funded by DARPA
  - Heirarchically layered network surveillance in 3 layers
    - Lowest: Service specific analysis monitoring
    - Middle: Domain wide analysis monitoring
    - Highest: Enterprise wide analysis monitoring
  - Enables analysis that overarches implementation details
  - Features target specific optimization:
    - Configurable event structures
    - Engine configurations
    - 'Subscription' list
    - Event collection methods
    - Analysis unit configurations
    - Variable response methods
  - Lots of details available on the SRI web site
    - including technical publications

EMERALD Architecture

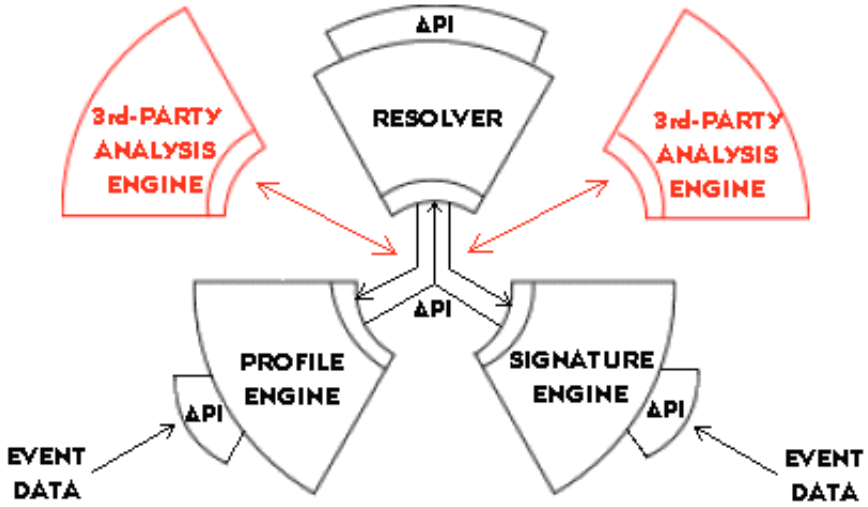

Source: http://www.sdl.sri.com/emerald/concepts.html, 22 May 2000

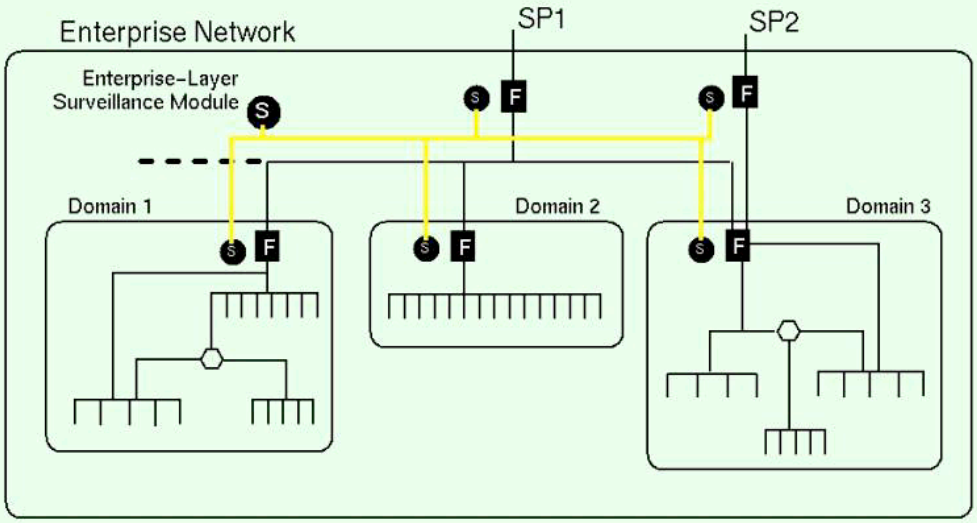

Source: http://www.sdl.sri.com/emerald/live-traffic.html, 22 May 2000

Common Intrusion Detection Framework

- CIDF began in 1997; government sponsored effort
- Includes:
  - Set of architectural conventions for how IDSs can be modeled
    - Event generators, analysis engines, storage mechanisms, response components
  - Specification for messaging
    - Generalized Intrusion Detection Objects (GIDO)
  - Specification for how to move GIDOs through system components
  - Interoperability protocols
- Enables
  - Multi-vendor solutions
  - Distributed IDS architectures

Challenges in IDS

- Correlating data from many types of targets
  - Personnel security systems (badges, etc)
  - Physical security systems (door alarms, etc)
  - Help desk
  - Network security systems (firewalls, etc)
- Distributed enterprises with integrated IT
  - Correlating data from many different physical locations
- Maintaining security attributes of IDS elements and data
  - C, I, A
  - Auditing to ensure correct operations
  - Support to forensics processes

Caveats

- This subject matter is incredibly complex
  - This presentation has been a very high level look
- A couple of good texts on the subject include:
  - Edward Amoroso's "Intrusion Detection Systems"
  - Paul Proctor's "Practical Intrusion Detection Handbook"
- The technology is not mature and research is on-going
- Understanding how to integrate an IDS into a security architecture requires both technical and business process analysis skills

Contact Information

Julie J.C.H. Ryan, D.Sc.

1776 G. Street NW #110

Washington DC, 20052

jjchryan@seas.gwu.edu

http://www.seas.gwu.edu/~infosec/

The George Washington University is an NSA Certified Center of Academic Excellence in Information Assurance Education and meets the Federal Training Standards for Information Systems Security Professionals (NSTISSI 4011). We offer Graduate Certificate, Master's, and Doctoral level education in Infosec for professionals from all educational backgrounds. GWU is located in the heart of Washington DC very near the White House and other government offices.