

Cryptography and Network Security II

Second Course

Lecture 6: Intrusion Detection (Part 1)

Classes of intruders: criminals

- Individuals or members of an organized crime group with a goal of financial reward
 - Identity theft
 - Theft of financial credentials
 - Corporate espionage
 - Data theft
 - Data ransoming
- Typically young, often Eastern European, Russian, or southeast Asian hackers, who do business on the Web
- Meet in underground forums to trade tips and data and coordinate attacks

Classes of intruders: activists

- Are either individuals, usually working as insiders, or members of a larger group of outsider attackers, who are motivated by social or political causes
- Also known as hacktivists
 - Skill level is often quite low
- Aim of their attacks is often to promote and publicize their cause typically through:
 - Website defacement
 - Denial of service attacks
 - Theft and distribution of data that results in negative publicity or compromise of their targets

Intruders: state-sponsored

- Groups of hackers sponsored by governments to conduct espionage or sabotage activities
- Also known as Advanced Persistent Threats (APTs) due to the covert nature and persistence over extended periods involved with any attacks in this class
- Widespread nature and scope of these activities by a wide range of countries from China to the USA, UK, and their intelligence allies

Intruders: others

- Hackers with motivations other than those previously listed
- Include classic hackers or crackers who are motivated by technical challenge or by peer-group esteem and reputation
- Many of those responsible for discovering new categories of buffer overflow vulnerabilities could be regarded as members of this class
- Given the wide availability of attack toolkits, there is a pool of “hobby hackers” using them to explore system and network security

Skill level: apprentice

- Hackers with minimal technical skill who primarily use existing attack toolkits
- They likely comprise the largest number of attackers, including many criminal and activist attackers
- Given their use of existing known tools, these attackers are the easiest to defend against
- Also known as “script-kiddies” due to their use of existing scripts (tools)

Skill level: journeyman

- Hackers with sufficient technical skills to modify and extend attack toolkits to use newly discovered, or purchased, vulnerabilities
- They may be able to locate new vulnerabilities to exploit that are similar to some already known
- Hackers with such skills are likely found in all intruder classes
- Adapt tools for use by others

Skill level: master

- Hackers with high-level technical skills capable of discovering brand new categories of vulnerabilities
- Write new powerful attack toolkits
- Some of the better known classical hackers are of this level
- Some are employed by state-sponsored organizations
- Defending against these attacks is of the highest difficulty

Intruders: another classification

- Masquerader: unauthorized individuals who penetrates a system
- Misfeasor: legit user who accesses unauthorized data
- Clandestine: seizes supervisory control

User and software trespass

- User trespass: unauthorized logon, privilege abuse
- Software trespass: virus, worm, or Trojan horse

Example of intrusion

- Remote root compromise
- Web server defacement
- Guessing/cracking passwords
- Copying databases containing credit card numbers
- Viewing sensitive data without authorization
- Running a packet sniffer
- Distributing pirated software
- Using an unsecured modem to access internal network
- Impersonating an executive to get information
- Using an unattended workstation

Intruder behavior

- Target acquisition and information gathering
- Initial access
- Privilege escalation
- Information gathering or system exploit
- Maintaining access
- Covering tracks

Hacker behavior example

1. Select target using IP lookup tools
2. Map network for accessible services
 - study physical connectivity (via NMAP)
3. Identify potentially vulnerable services
4. Brute force (guess) passwords
5. Install remote administration tool
6. Wait for admin to log on and capture password
7. Use password to access remainder of network

Criminal intruder behavior

1. Act quickly and precisely to make their activities harder to detect
2. Exploit perimeter via vulnerable ports
3. Use Trojan horses (hidden software) to leave back doors for re-entry
4. Use sniffers to capture passwords
5. Do not stick around until noticed
6. Make few or no mistakes

Insider intruder behavior

1. Create network accounts for themselves and their friends
2. Access accounts and applications they wouldn't normally use for their daily jobs
3. E-mail former and prospective employers
4. Conduct furtive (covert) instant-messaging chats
5. Visit web sites that cater to disgruntled employees, such as f*dcompany.com
6. Perform large downloads and file copying
7. Access the network during off hours

Insider attacks

- Among most difficult to detect and prevent
- Employees have access & systems knowledge
- May be motivated by revenge/entitlement
 - When employment terminated
 - Taking customer data when move to competitor
- IDS/IPS may help but also need
 - Least privilege, monitor logs, strong authentication, termination process to block access & take mirror image of employee's HD (for future purposes)

Security intrusion & detection (RFC 2828)

- **Security intrusion:** a security event, or combination of multiple security events, that constitutes a security incident in which an intruder *gains, or attempts to gain*, access to a system (or system resource) without having authorization to do so.
- **Intrusion detection:** a security service that monitors and analyzes system events for the purpose of finding, and providing real-time or near real-time warning of attempts to access system resources in an unauthorized manner.

Intrusion techniques

- Objective to gain access or increase privileges
- Initial attacks often exploit system or software vulnerabilities to execute code to get backdoor
 - e.g. buffer overflow
- Or to gain protected information
 - Password guessing or acquisition (or via social engineering)

Intrusion detection systems

- **Host-based IDS:** monitor single host activity
- **Network-based IDS:** monitor network traffic
- **Distributed or hybrid:** Combines information from a number of sensors, often both host and network based, in a central analyzer that is able to better identify and respond to intrusion activity

Comprises three logical components:

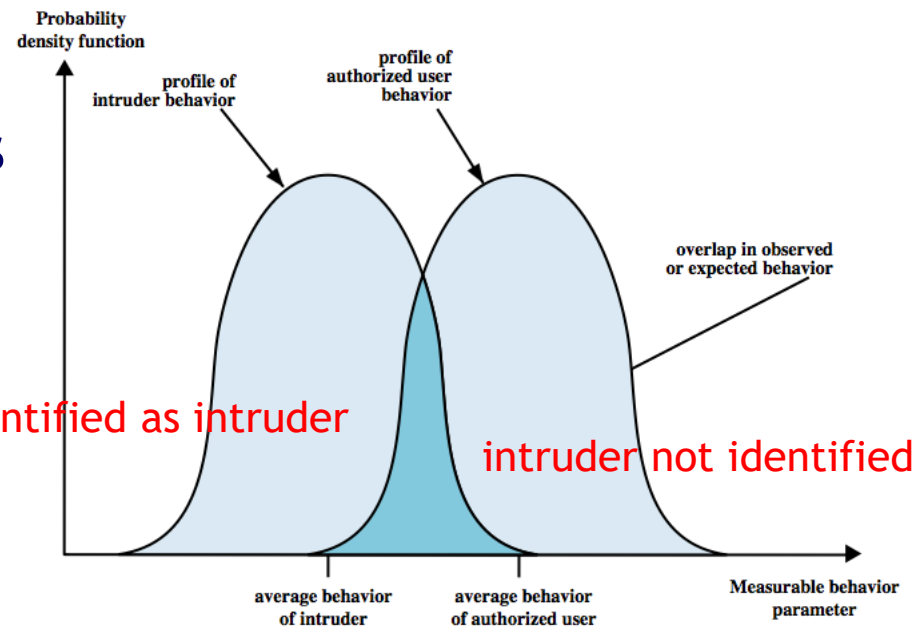
- **Sensors:** collect data
- **Analyzers:** determine if intrusion has occurred
- **User interface:** view output or control system behavior

IDS principles

- Assumption: intruder behavior differs from legitimate users

- Expect overlap as shown
- for legit users:
Observe major deviations from past history
- Problems of:
 - false positives
 - false negatives
 - must compromise

loose vs tight interpretation:
catch more (false +) or catch less (false -)



IDS requirements

Run continually	Be fault tolerant	Resist subversion
Impose a minimal overhead on system	Configured according to system security policies	Adapt to changes in systems and users
Scale to monitor large numbers of systems	Provide graceful degradation of service	Allow dynamic reconfiguration

IDS requirements

- Run continually with minimal human supervision
- Be fault tolerant: recover from crashes
- Resist subversion: monitor itself from change by intruder
- Impose a minimal overhead on system
- Configured according to system security policies
- Adapt to changes in systems and users
- Scale to monitor large numbers of systems
- Provide graceful degradation of service: if one component fails, others should continue to work
- Allow dynamic reconfiguration

Detection techniques

- Anomaly (behavior) detection
- Signature/heuristic detection

IDS: anomaly (behavior) detection

- Involves the collection of data relating to the behavior of legitimate users over a period of time
- Current observed behavior is analyzed to determine whether this behavior is that of a legitimate user or that of an intruder

Anomaly detection

- Threshold detection
 - checks excessive event occurrences over time
 - alone a crude and ineffective intruder detector
 - must determine both thresholds and time intervals
 - lots of false positive/false negative may be possible
- Profile based
 - characterize past behavior of *users/groups*
 - then detect significant deviations
 - based on analysis of audit records: *gather metrics*

Example of metrics

- **Counters:** e.g., number of logins during an hour, number of times a cmd executed
- **Gauge:** e.g., the number of outgoing messages [pkts]
- **Interval time:** the length of time between two events, e.g., two successive logins
- **Resource utilization:** quantity of resources used (e.g., number of pages printed)
- Mean and standard deviations

Signature/heuristic detection

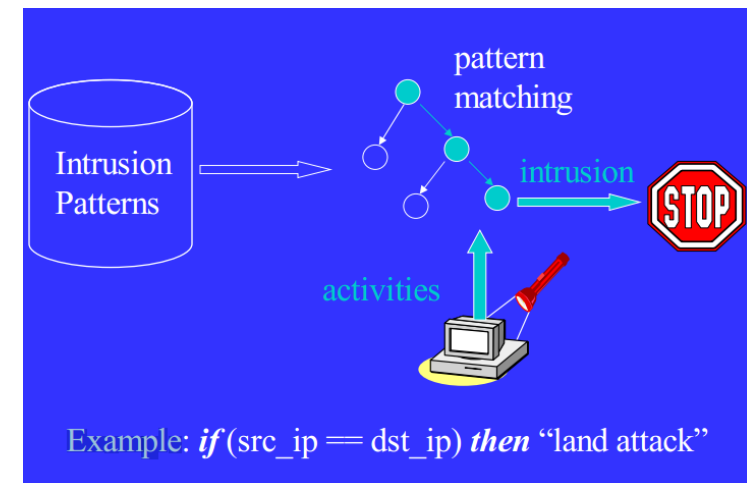
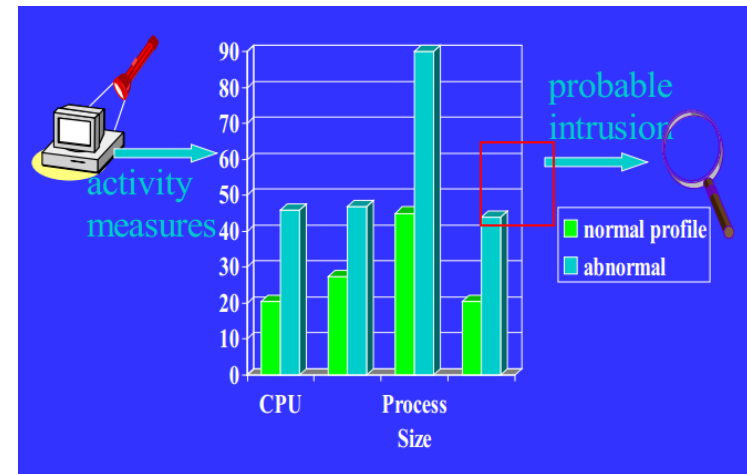
- Uses a set of known malicious data patterns or attack rules that are compared with current behavior
- Also known as misuse detection
- Can only identify known attacks for which it has patterns or rules (signature)
 - Very similar to anti-virus (requires frequent updates)
 - Rule-based penetration identification
 - rules identify **known** penetrations/weaknesses
 - often by analyzing attack scripts from Internet (CERTs)

Example of rules in a signature detection IDS

- Users should not be logged in more than one session
- Users do not make copies of system, password files
- Users should not read in other users' directories
- Users must not write other users' files
- Users who log after hours often access the same files they used earlier
- Users do not generally open disk devices but rely on high-level OS utils

Host-based IDS: signature vs anomaly detection

- Connection attempt from a reserved IP address
- Attempt to copy the password file
- Email containing a particular virus
- File access attack on an FTP server by issuing file and directory commands to it without first logging in



```
drop tcp $EXTERNAL_NET any -> $HOME_NET $HTTP_PORTS (msg:"Block Baidu Spider
```

Host-based IDS

- Specialized software to monitor system activity to detect suspicious behavior
 - primary purpose is to detect intrusions, log suspicious events, and send alerts
 - can detect both external and internal intrusions
- Two approaches, often used in combination:
 - **Anomaly detection:** consider normal/expected behavior over a period of time; apply statistical tests to detect intruder
 - threshold detection: for various events (#/volume of copying)
 - profile based (time/duration of login)
 - **Signature detection:** defines proper (or bad) behavior (rules)

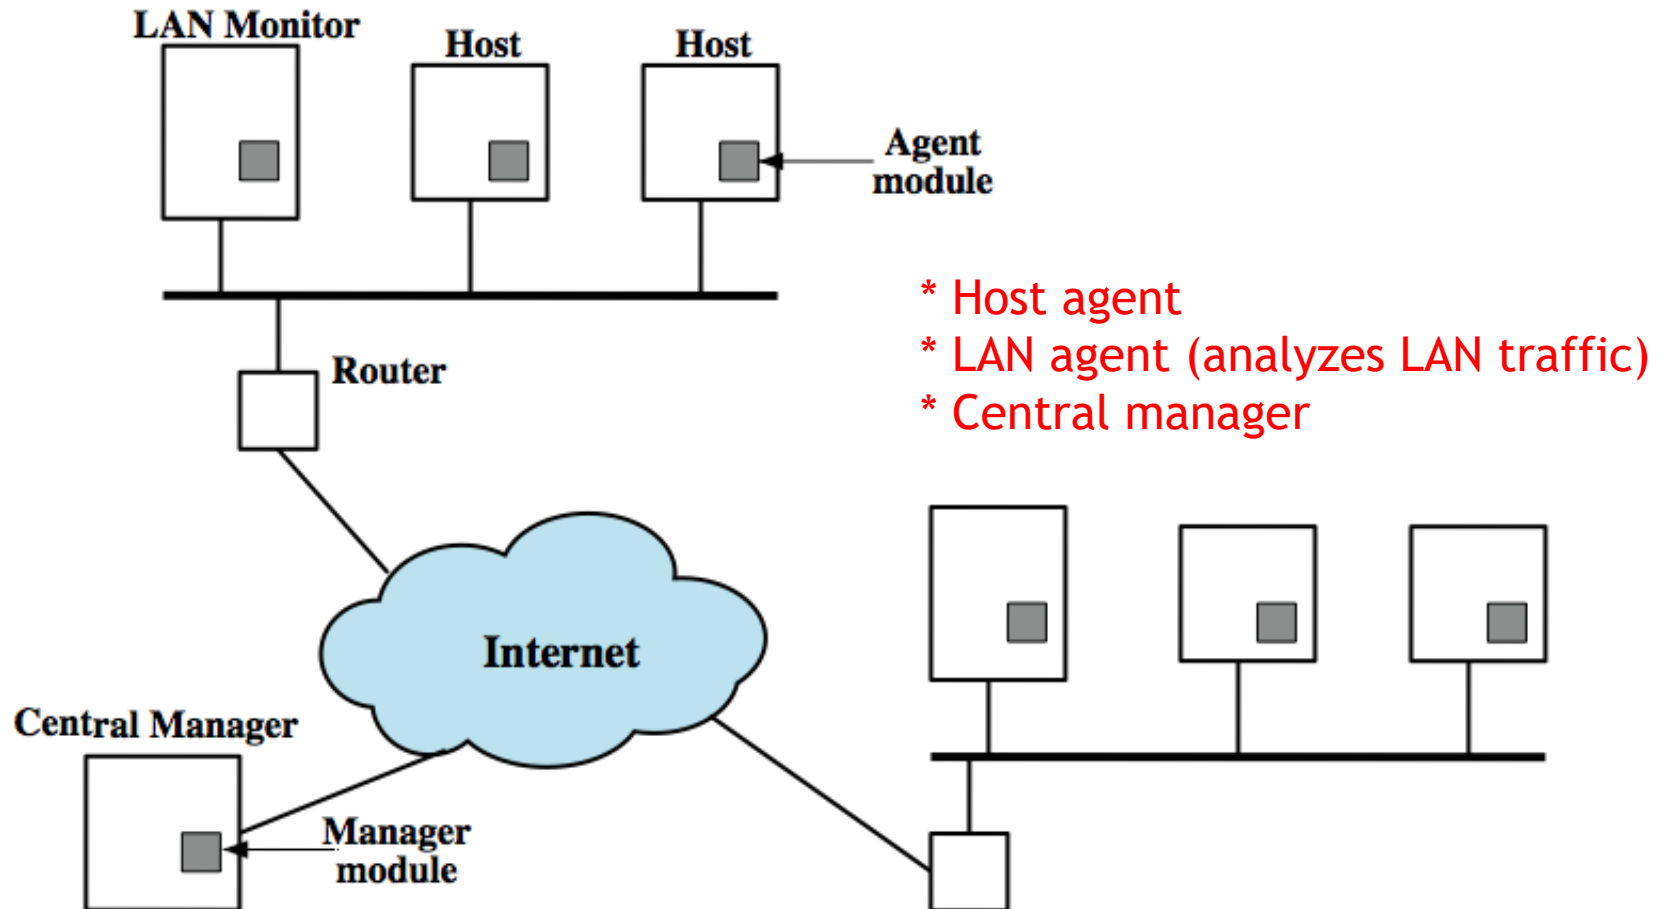
Audit records

- A fundamental tool for intrusion detection
- Two variants:
 - Native audit records: provided by O/S
 - always available but may not be optimum
 - Detection-specific audit records: IDS specific
 - additional overhead but specific to IDS task
 - often log individual elementary actions
 - e.g. may contain fields for: subject, action, object, exception-condition, resource-usage, time-stamp
 - possible overhead (two such utilities)

Common data sources

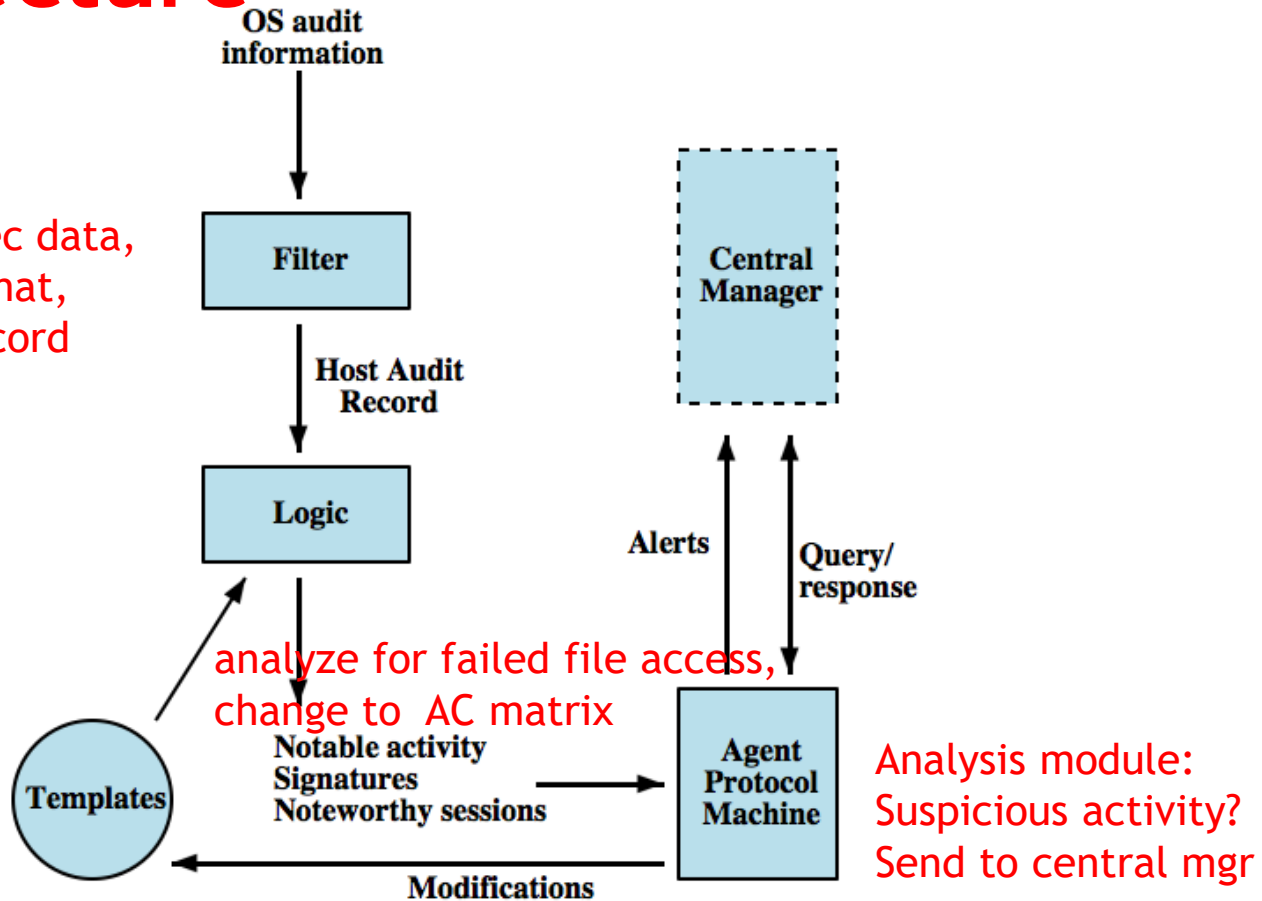
- Common data sources include:
 - System call traces
 - Audit (log file) records
 - File integrity checksums
 - Registry access

Distributed host-based IDS

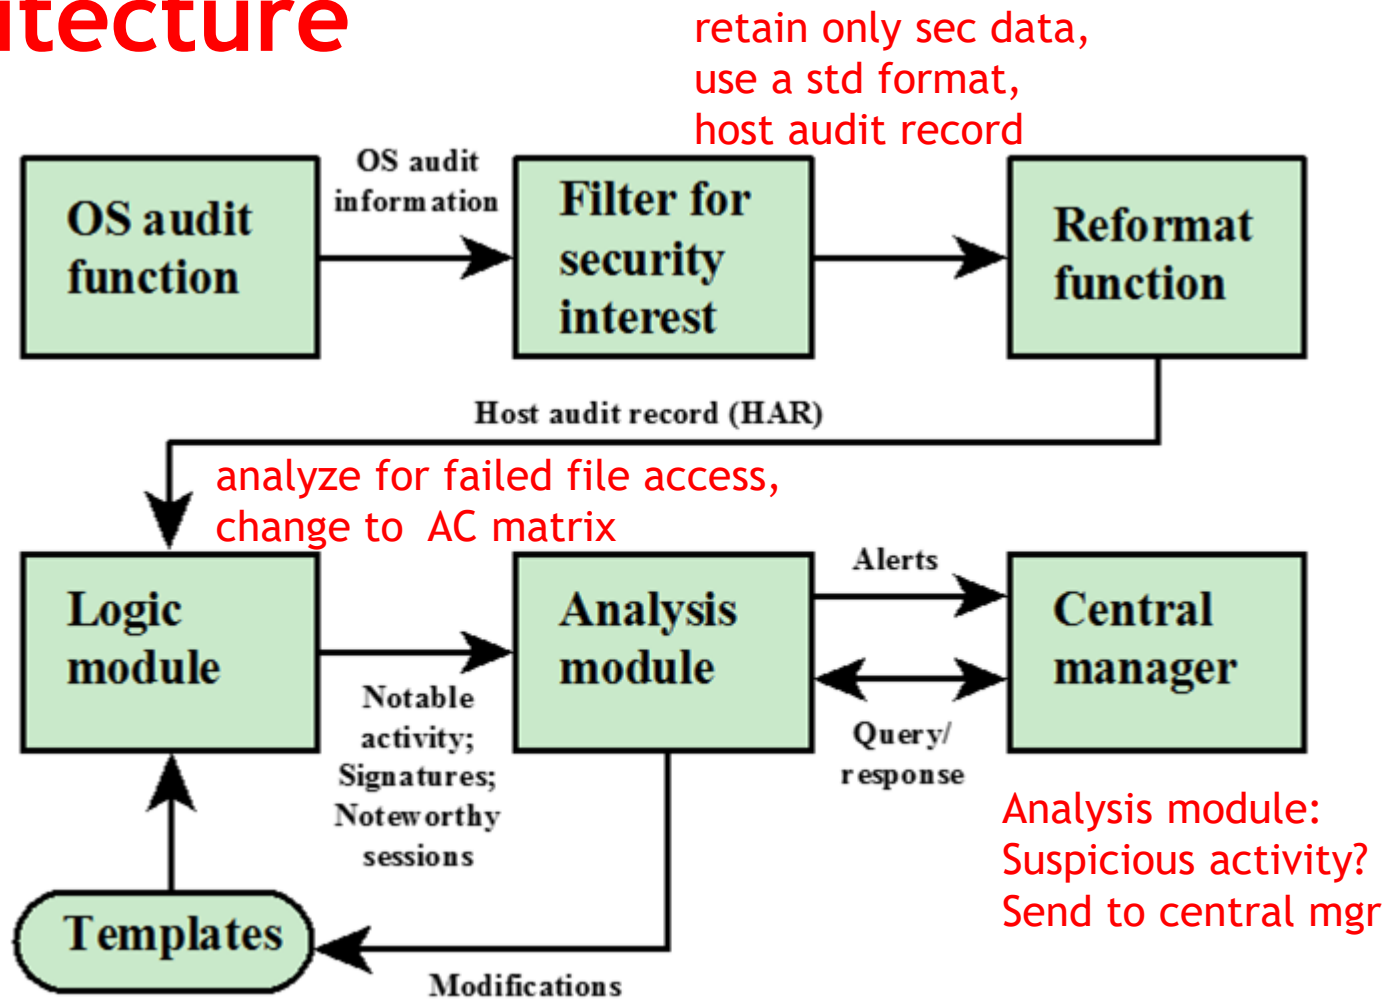


Distributed host-based IDS: agent architecture

retain only sec data,
use a std format,
host audit record




Distributed host-based IDS: agent architecture



Network-Based IDS

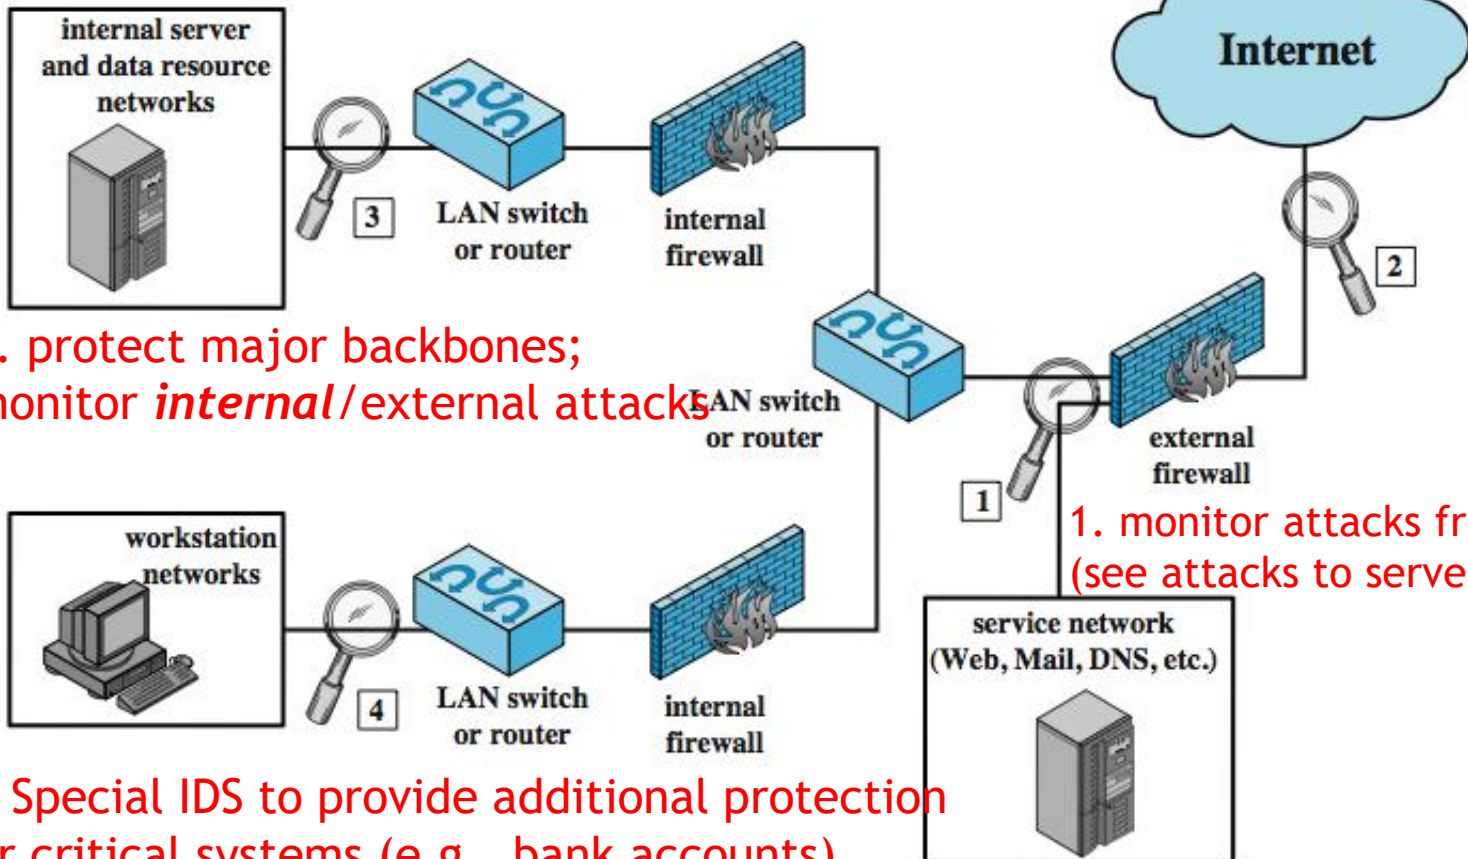
- Network-based IDS (NIDS)
 - *Monitor traffic at selected points on a network* (e.g., rlogins to disabled accounts)
 - In (near) real time to detect intrusion patterns
 - May examine network, transport and/or application level *protocol* activity directed toward systems
- Comprises a number of sensors
 - Inline (possibly as part of other net device) - traffic passes thru it
 - Passive (monitors copy of traffic)

Passive sensors



NIDS Sensor Deployment

2. monitor and documents unfiltered packets; more work to do



3. protect major backbones; monitor *internal*/external attacks

1. monitor attacks from outside (see attacks to servers)

4. Special IDS to provide additional protection for critical systems (e.g., bank accounts)

NIDS intrusion detection techniques

- Signature detection
 - at application (*FTP*), transport (*port scans*), network layers (*ICMP*); unexpected application services (*host running unexpected app*), policy violations (*website use*)
- Anomaly detection
 - of denial of service attacks, scanning, worms (*significant traffic increase*)
- When potential violation detected, sensor sends an alert and logs information
 - Used by analysis module to refine intrusion detection parameters and algorithms
 - by security admin to improve protection

Distributed hybrid intrusion detection

(host-based, NIDS, distributed host-based)


Issues:

1. Tools may not recognize new threats
2. Difficult to deal with rapidly spreading attacks

Solution:

Distributed Adaptive IDS thru Peer-to-peer gossip and cooperation

One developed by Intel



PEP = policy enforcement point
DDI = distributed detection and inference

Logging of alerts (for all types)

- Typical information logged by a NIDS sensor includes:
 - Timestamp
 - Connection or session ID
 - Event or alert type
 - Rating
 - Network, transport, and application layer protocols
 - Source and destination IP addresses
 - Source and destination TCP or UDP ports, or ICMP types and codes
 - Number of bytes transmitted over the connection
 - Decoded payload data, such as application requests and responses
 - State-related information

Intrusion detection exchange format

To facilitate development of a distributed IDS

Not a product, but a proposed IETF standard

Key elements

Data source: raw data from an IDS

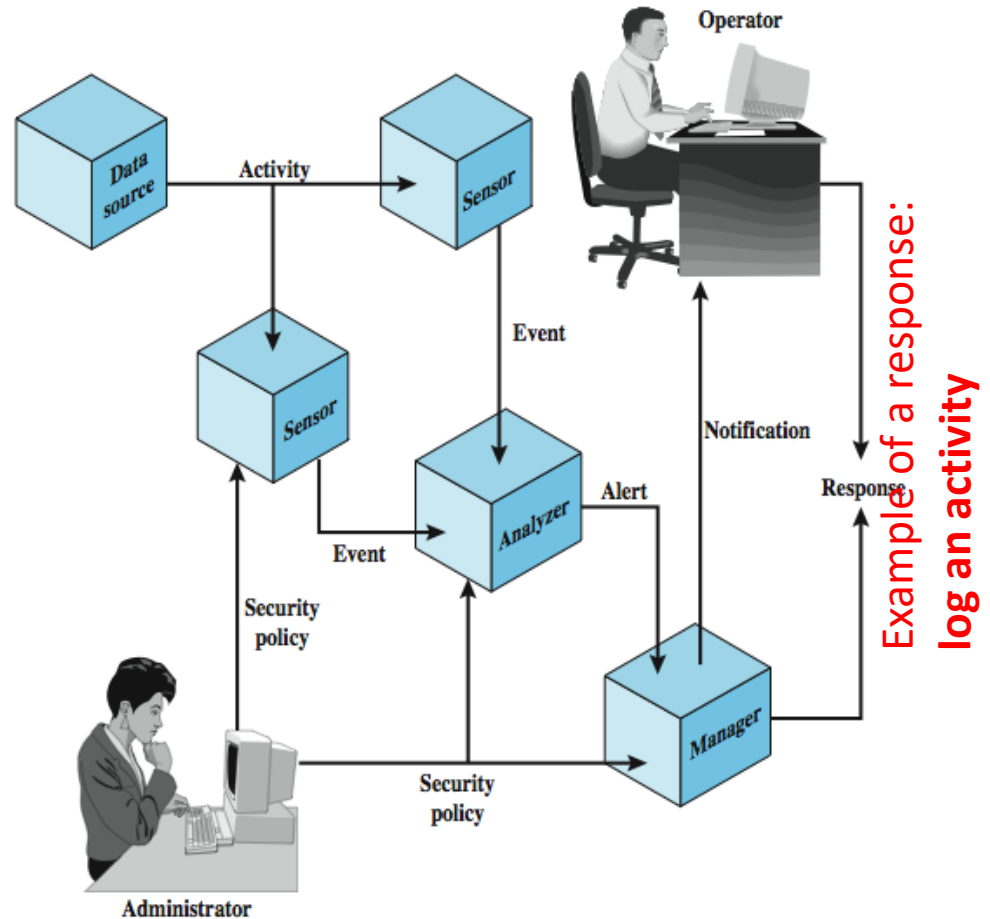
Sensor: collect and forward events

Analyzer: process data

Administrator defines sec policy

Manager: a process for operator to manage the IDS system

Operator: the user of the Manager



Honeypots

- Decoy systems
 - Filled with fabricated info and instrumented with monitors/event loggers
 - Lure a potential attacker away from critical systems
 - Collect information about the attacker's activity
 - Encourage the attacker to stay on the system long enough for administrators to respond
 - Divert and hold attacker to collect activity info without exposing production systems
- Initially were single systems
- More recently are/emulate entire networks

Honeypot classification


- Low interaction honeypot
 - Consists of a software package that emulates particular IT services or systems well enough to provide a realistic initial interaction, but does not execute a full version of those services or systems
 - Provides a less realistic target
 - Often sufficient for use as a component of a distributed IDS to warn of imminent attack
- High interaction honeypot
 - A real system, with a full operating system, services and applications, which are instrumented and deployed where they can be accessed by attackers

Honeypot deployment

1. Tracks attempts to connect to an unused IP address; can't help with inside attackers

3. Full internal honeypot; can detect internal attacks

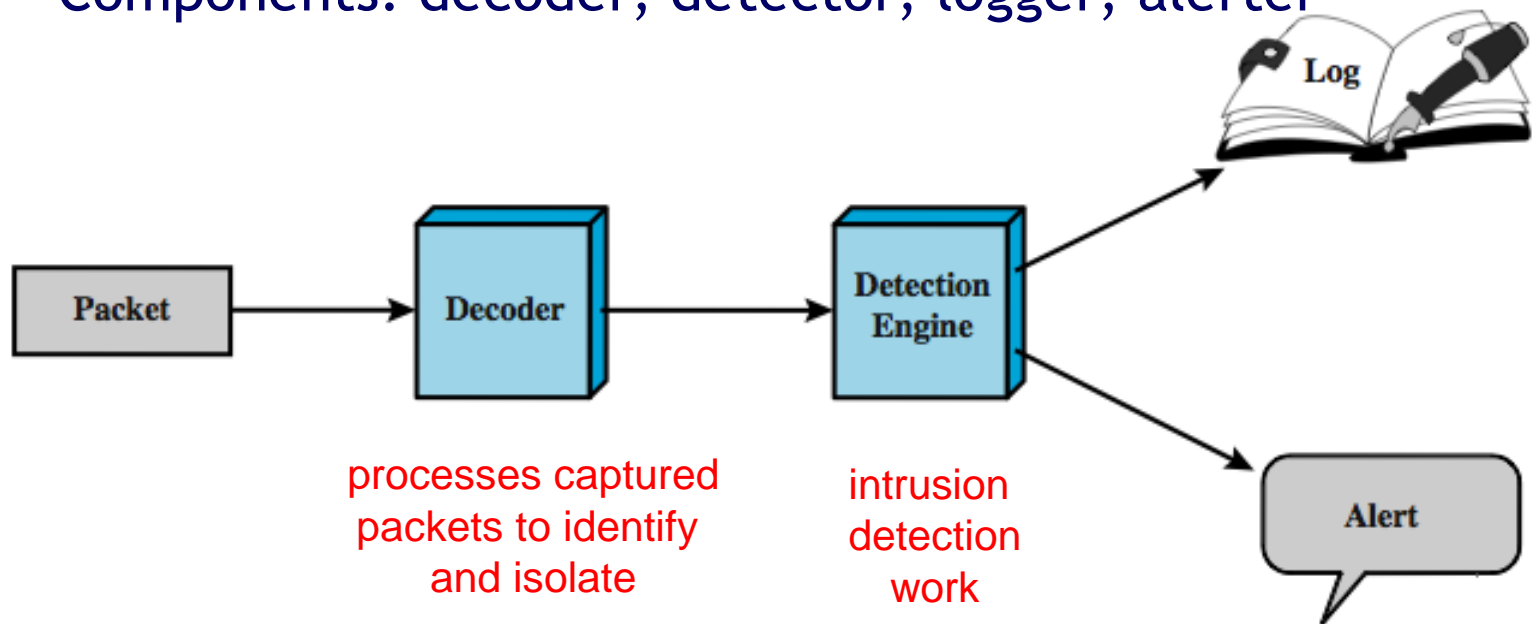
2. In DMZ; must make sure the other systems in the DMZ are secure; firewalls may block traffic to the honeypot



Snort IDS

- Lightweight IDS

- Open source (rule-based)
- Real-time packet capture and rule analysis
- Passive or inline
- Components: decoder, detector, logger, alerter



SNORT Rules

- Use a simple, flexible rule definition language
- Fixed header and zero or more options
- Header includes: action, protocol, source IP, source port, direction, dest IP, dest port
- Many options
- Example rule to detect TCP SYN-FIN attack:

```
alert tcp $EXTERNAL_NET any -> $HOME_NET any \  
(msg: "SCAN SYN FIN"; flags: SF, 12; \  
reference: arachnids, 198; classtype: attempted-recon;)
```

- detects an attack at the TCP level; \$strings are variables with defined values; any source or dest port is considered; checks to see if SYN and FIN bits are set

Summary

- Introduced intruders & intrusion detection
 - Hackers, criminals, insiders
- Intrusion detection approaches
 - Host-based (single and distributed)
 - Network
 - Distributed adaptive
- Honeypots
- Snort example