Network Security: Intrusion Detection

Seungwon Shin, KAIST

most slides from Dr. Guofei Gu

Some Definition

- Intrusion
  - A set of actions aimed to compromise the security goals, namely
    - Integrity, confidentiality, or availability, of a computing and networking resource

- Intrusion detection
  - The process of identifying and responding to intrusion activities

Why Is Intrusion Detection Necessary?

- Protect your systems from intrusion

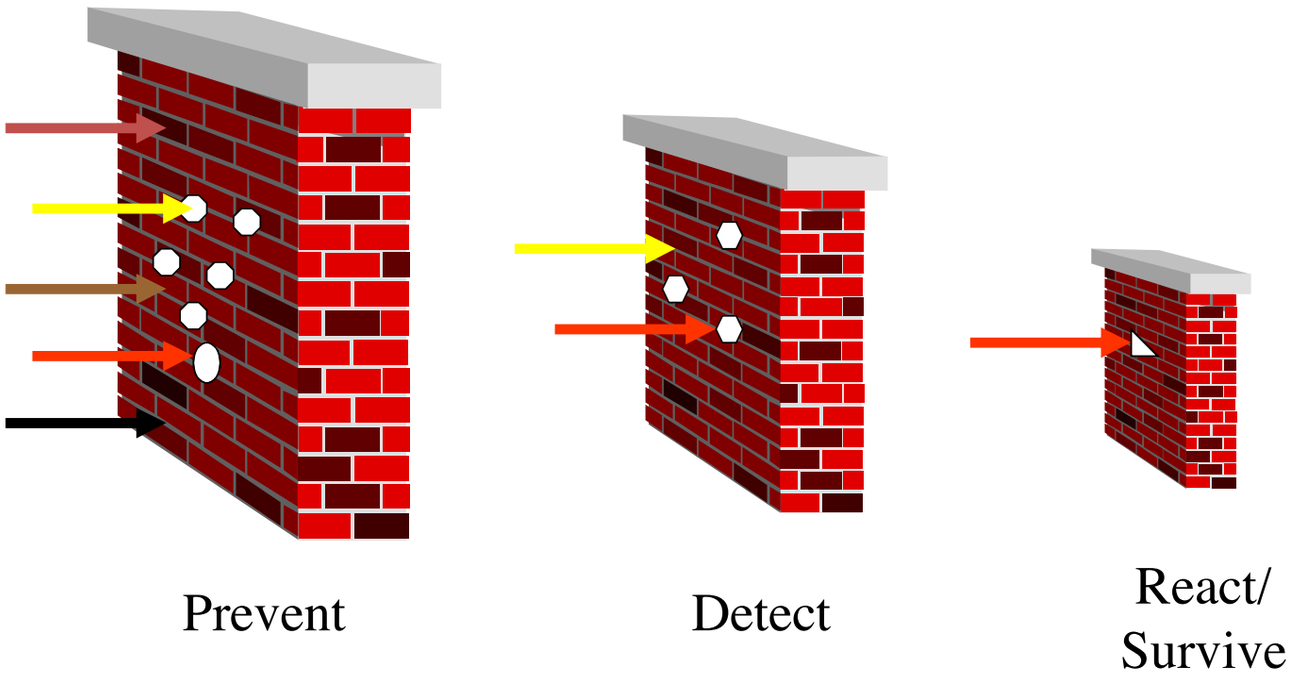

Security principles: layered mechanisms

Elements of IDS

- Primary assumptions:
  - System activities are observable
  - Normal and intrusive activities have distinct evidence
    - Components of intrusion detection systems:
- From an algorithmic perspective:
  - Features - capture intrusion evidences
  - Models - piece evidences together
- From a system architecture perspective:
  - Audit data processor, knowledge base, decision engine, alarm generation and responses

Components of IDS

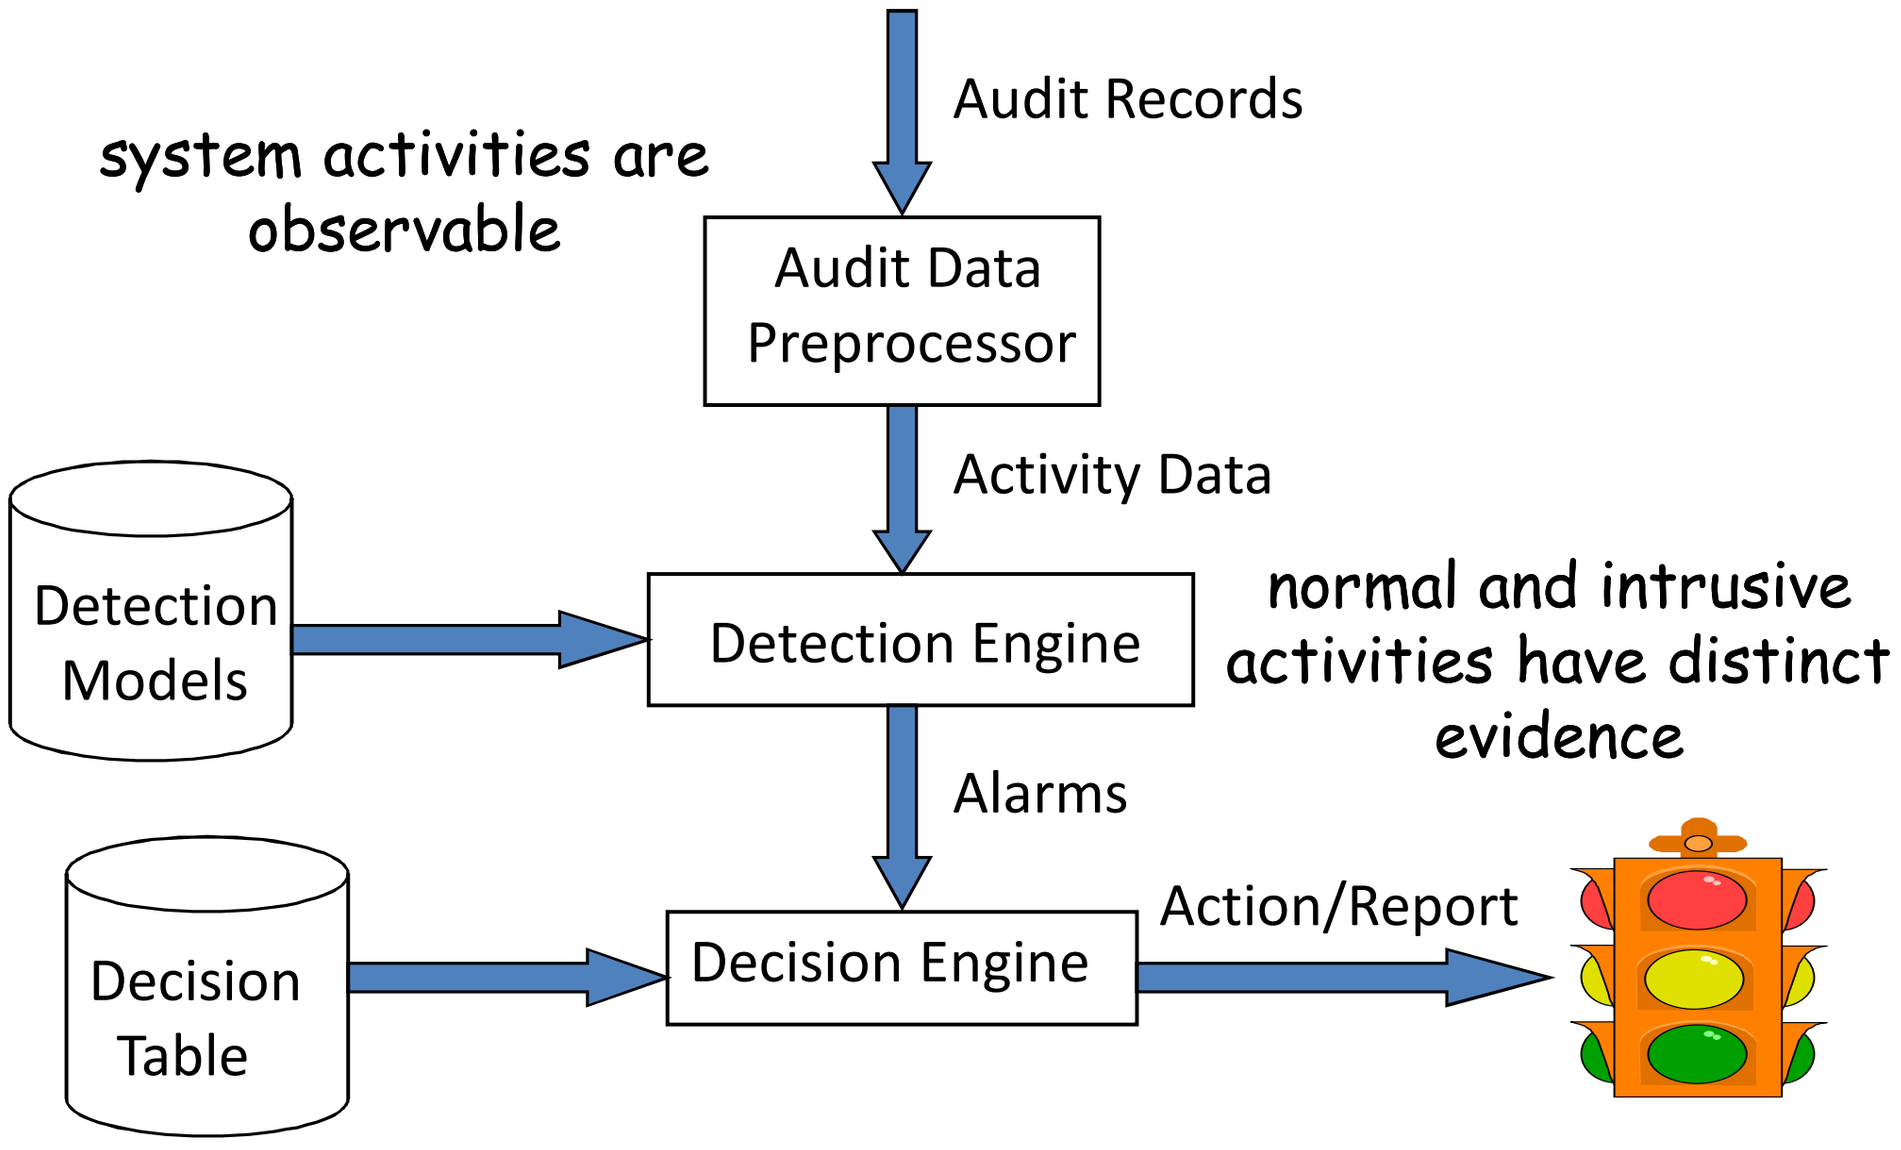

IDS Approachs

- Modeling
  - Features: evidences extracted from audit data
  - Analysis approach: piecing the evidences together
    - **Misuse detection (signature-based, e.g., Snort, Bro)**
    - **Anomaly detection (e.g., statistical-based)**

- Deployment
  - Network-based
  - Host-based

- Development and maintenance
  - Hand-coding of “expert knowledge”
  - Learning based on audit data

Misuse Detection

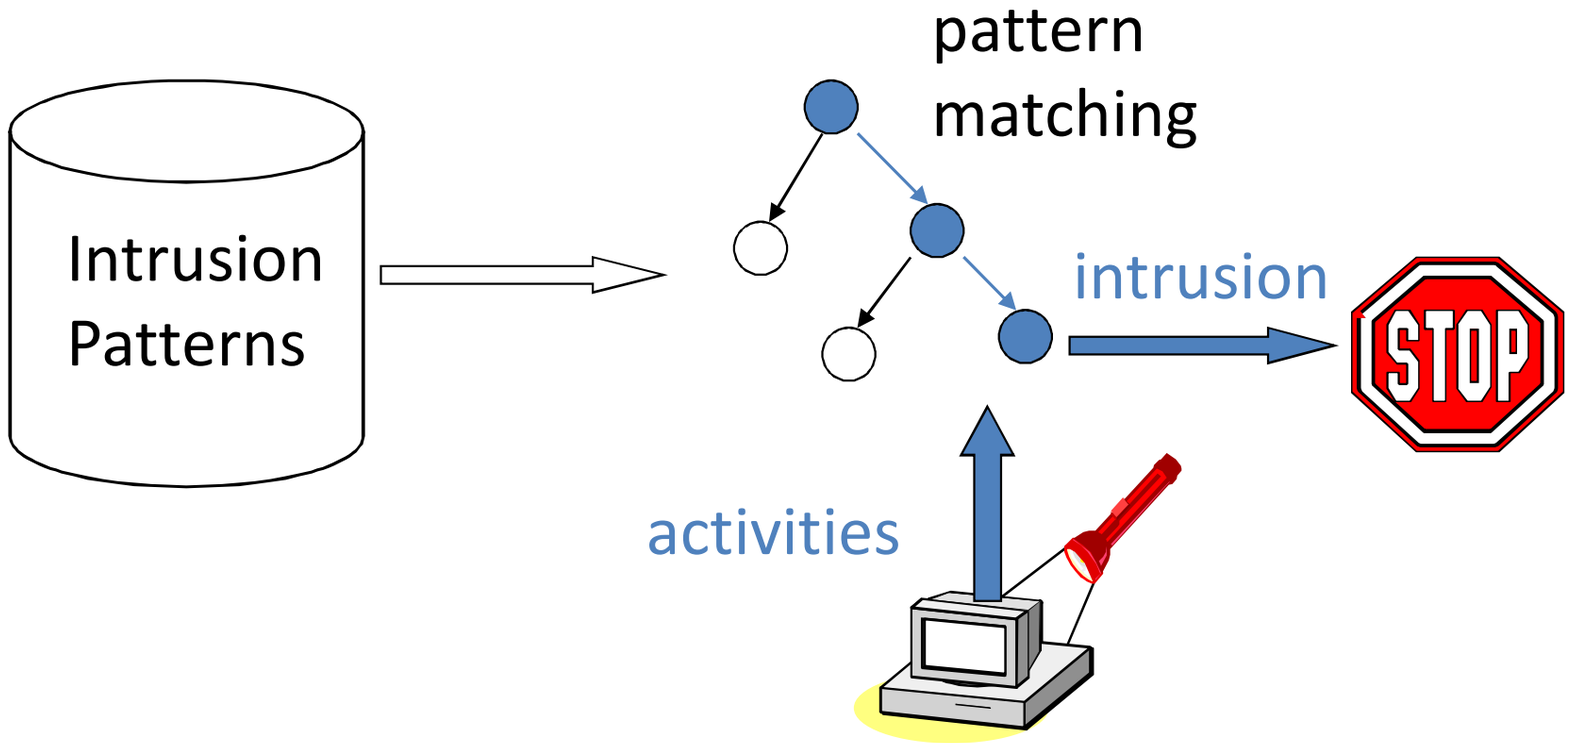

Example: ***if*** (src_ip == dst_ip) ***then*** “land attack”

Cannot detect unknown attacks

Anomaly Detection

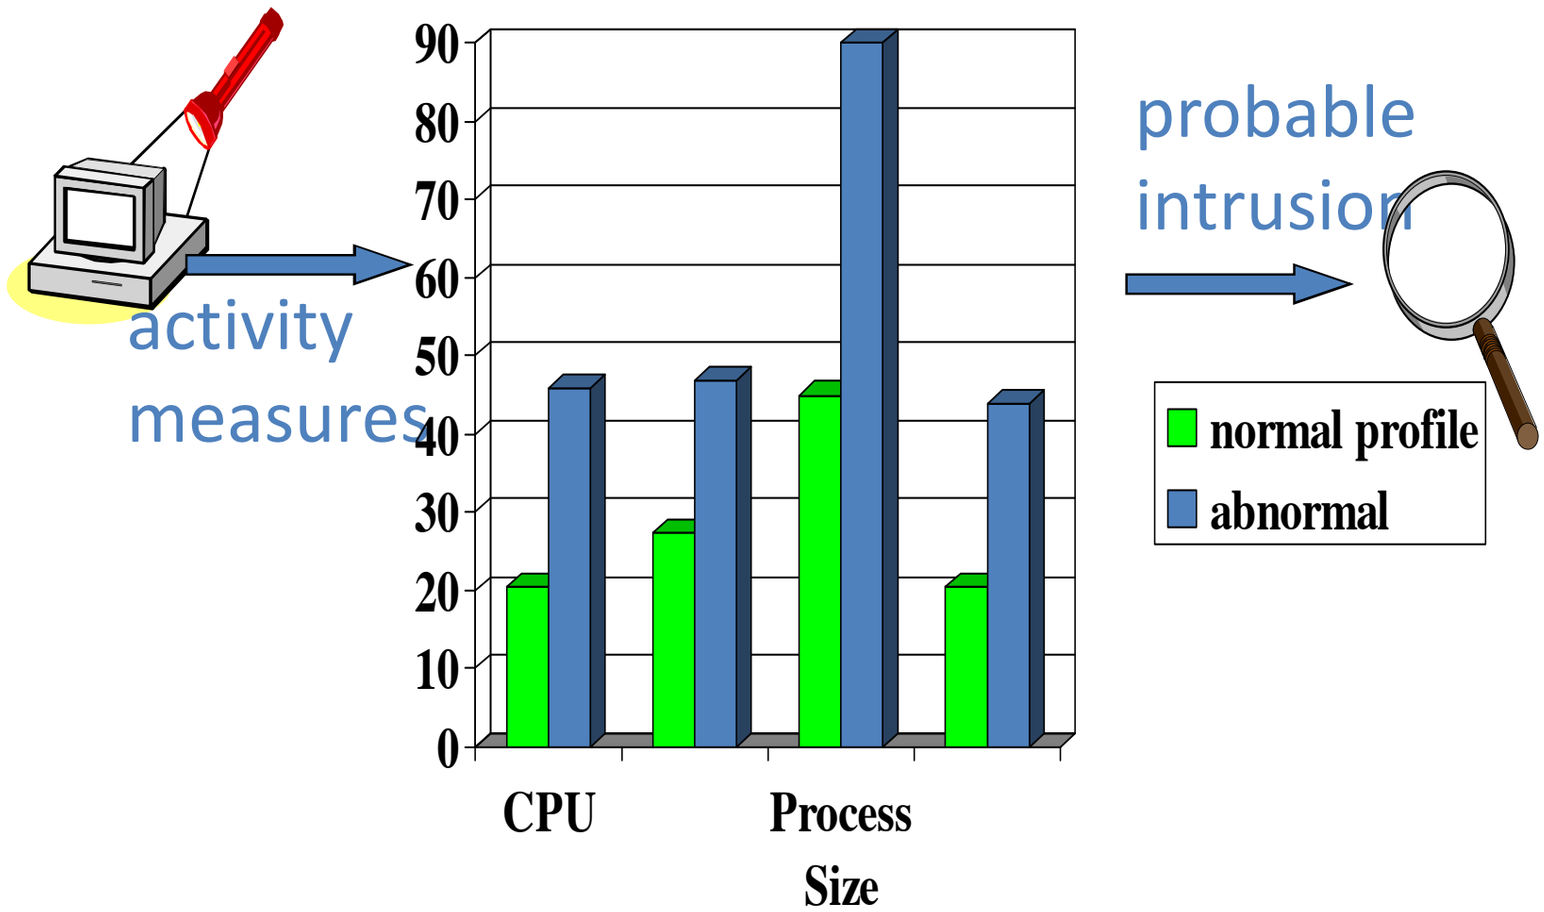

Relatively high false positive rate - anomalies can just be new normal activities.

Monitoring Network and Hosts

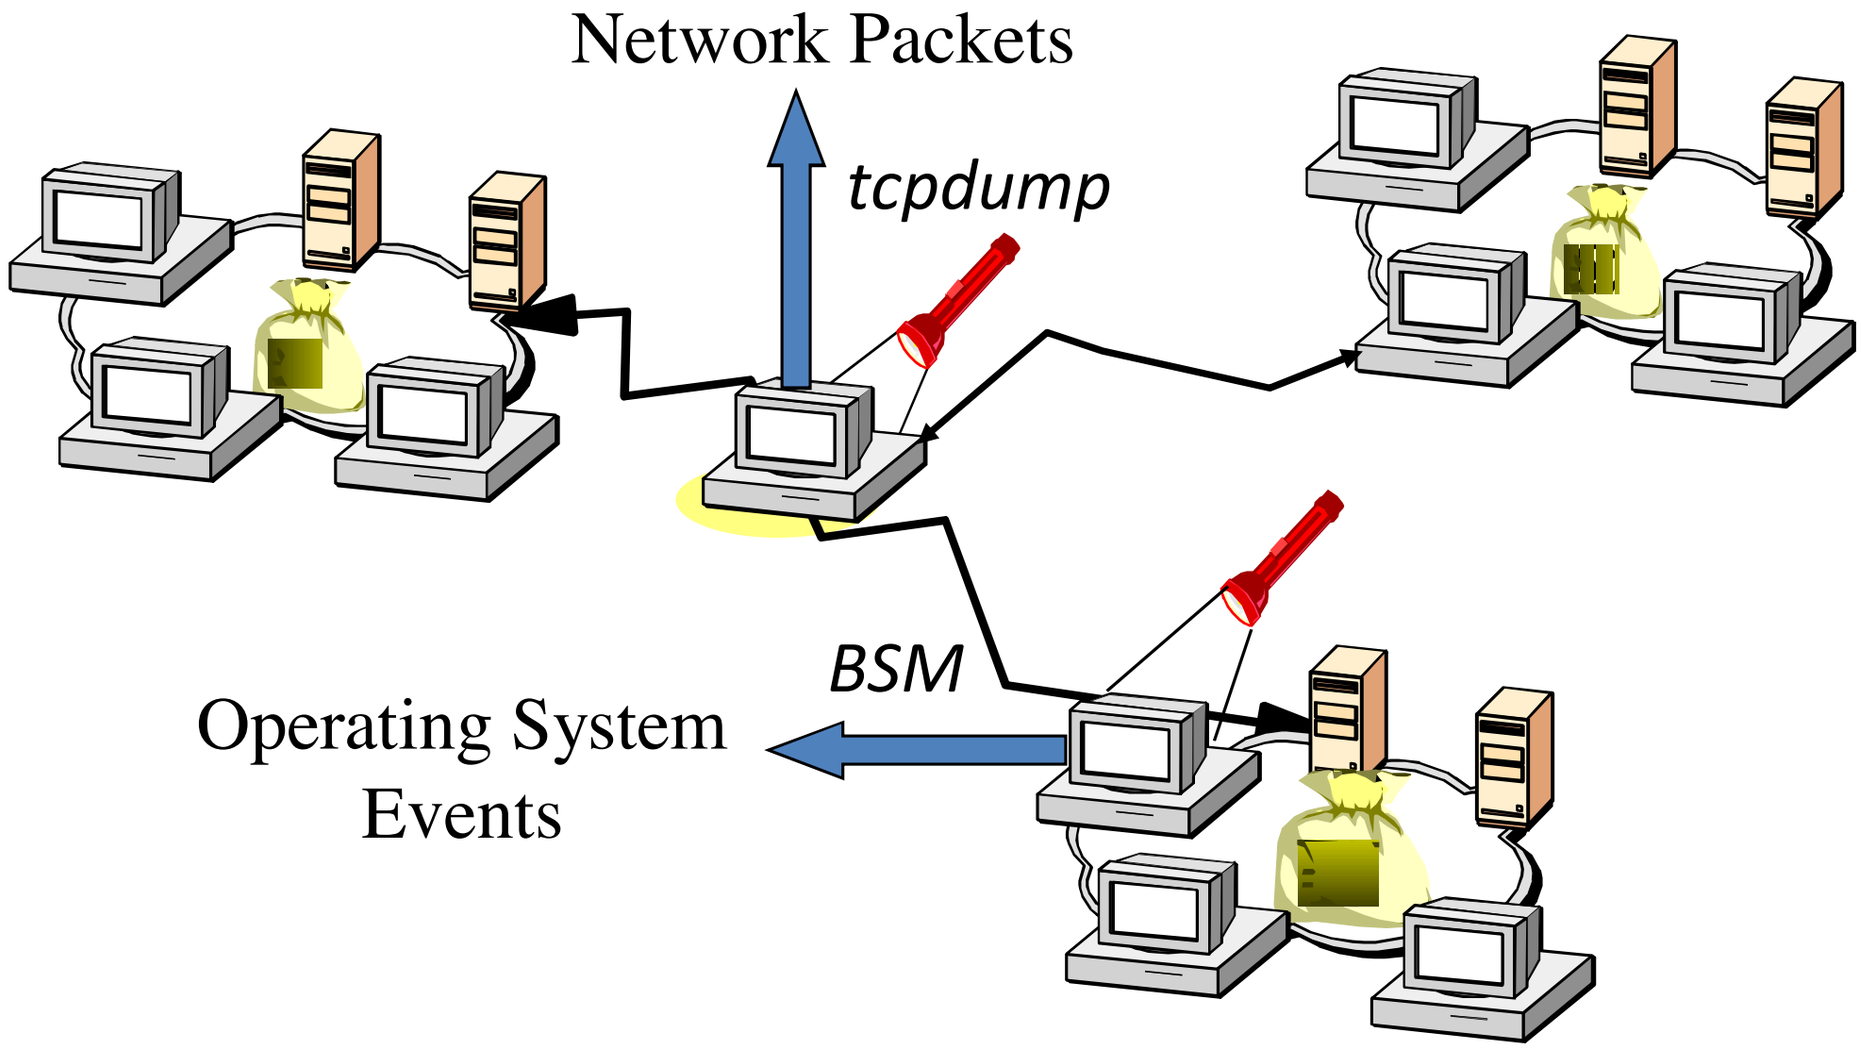

Performance Metric

- Algorithm
  - **Alarm: A; Intrusion: I**
  - Detection (true positive) rate: P(A|I)
    - False negative rate P(¬A|I)
  - False positive rate: P(A|¬I)
    - True negative rate P(¬A|¬I)
  - Bayesian detection rate: P(I|A)

- Architecture
  - Scalable
  - Resilient to attacks

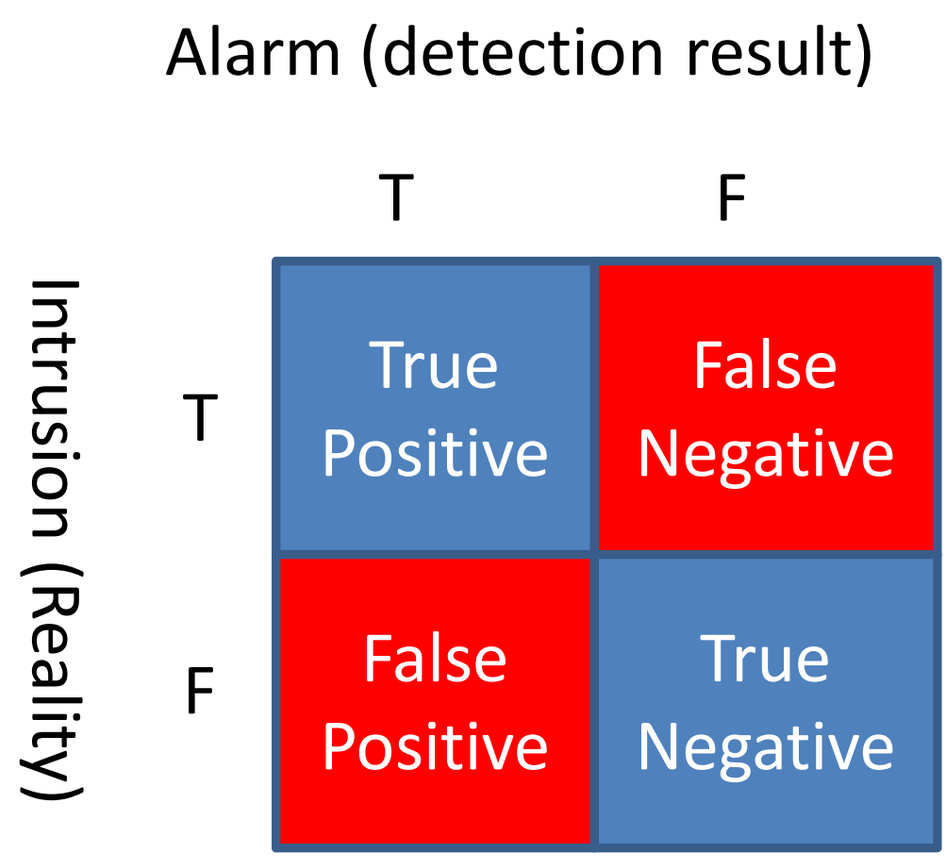

ROC Curve

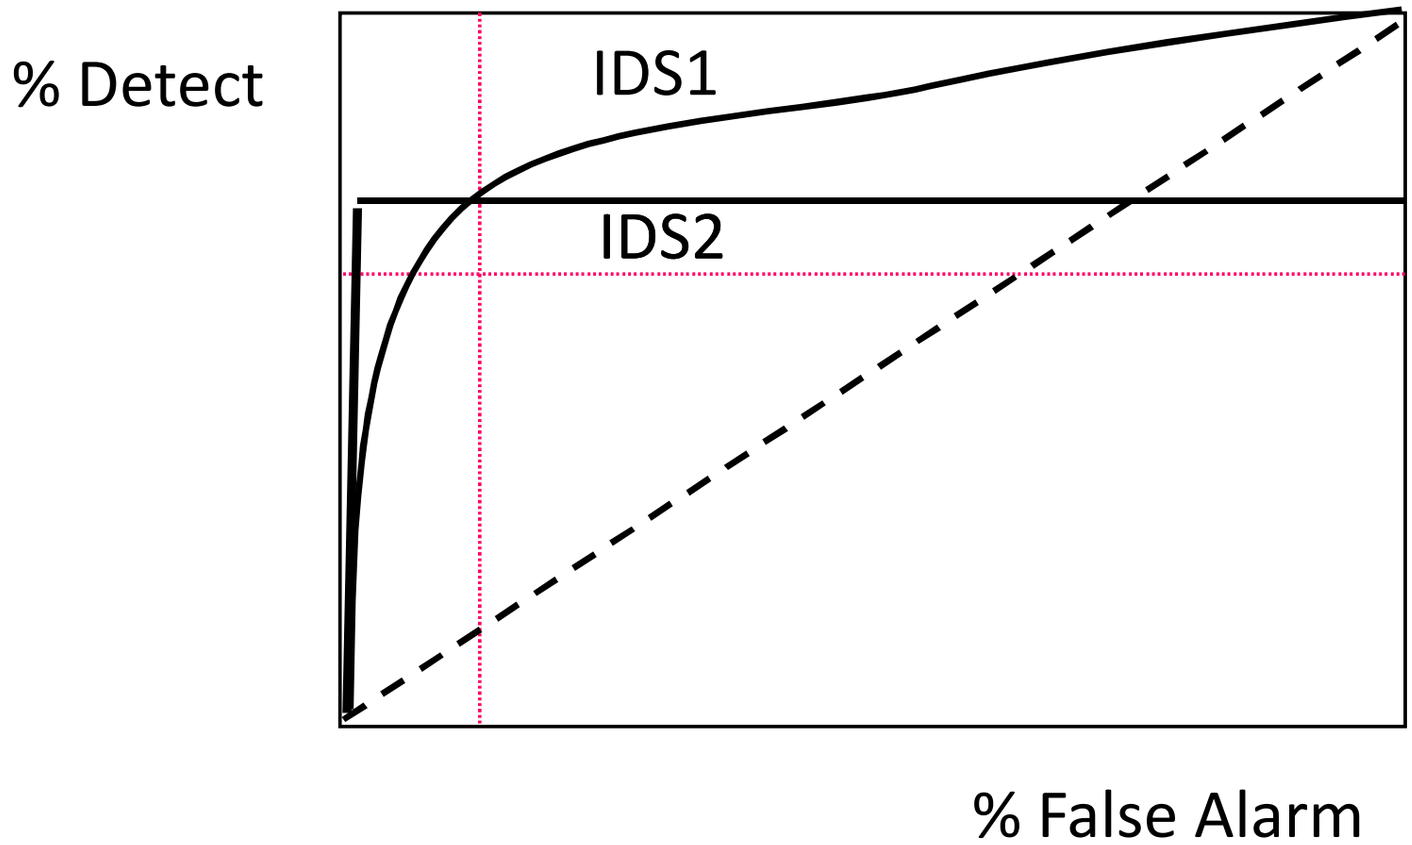

- Ideal system should have
  - 100% detection rate with 0% false alarm

HIDS

- Using OS auditing mechanisms
  - E.G., BSM on Solaris: logs all direct or indirect events generated by a user
  - ***strace*** for system calls made by a program
- Monitoring user activities
  - E.G., Analyze shell commands
- Monitoring executions of system programs
  - E.G., Analyze system calls made by sendmail

HIDS - Example

- A Sense of Self - Immunology Approach
  - Prof. Forrest at University of New Mexico
    - Anomaly detection
    - Simple and short sequences of events to distinguish “self” from not
    - Currently looking at system calls (strace)
    - Apply to detection of lpr and sendmail

Some More

- Anomaly detection for Unix processes
  - “Short sequences” of system calls as normal profile
    - (Forrest et al. UNM)

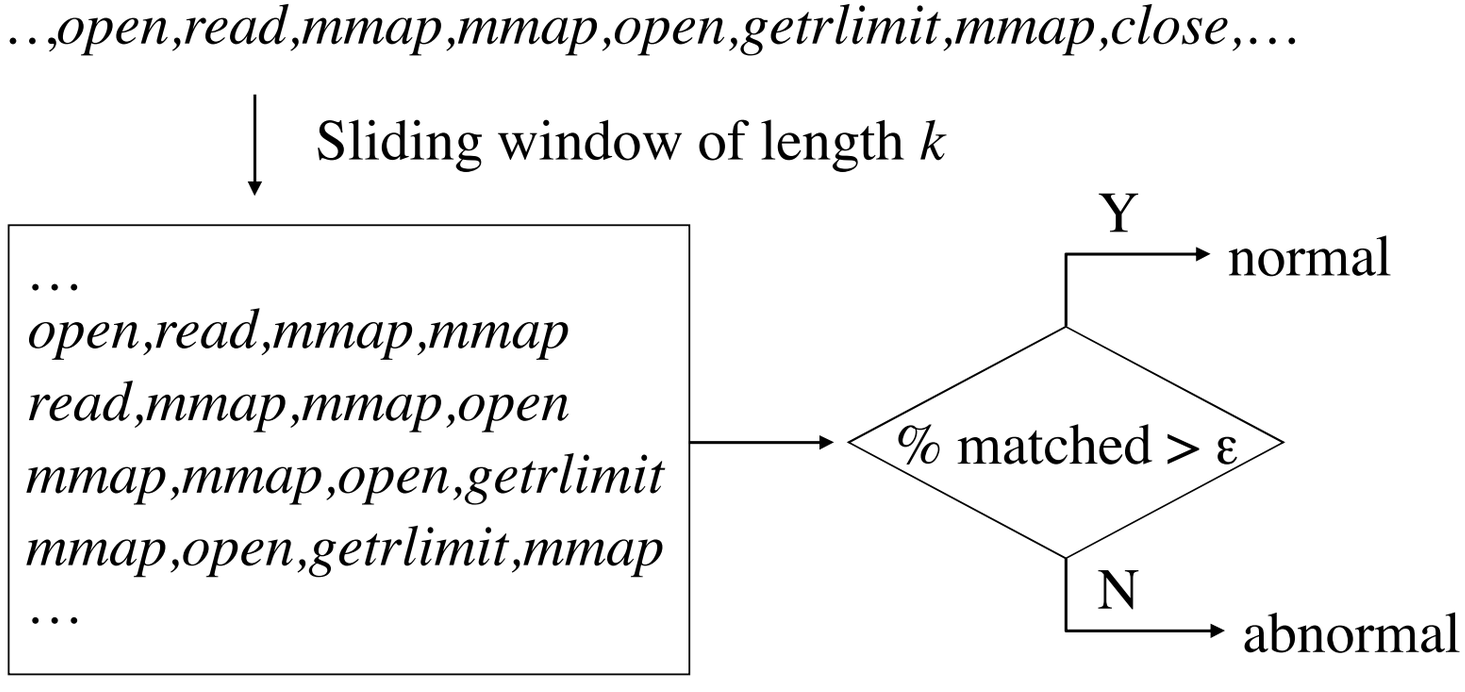

NIDS

- Deploying sensors at strategic locations
  - E.G., Packet sniffing via tcpdump at routers
- Inspecting network traffic
  - Watch for violations of protocols and unusual connection patterns
- Monitoring user activities
  - Look into the data portions of the packets for malicious command sequences
- Maybe easily defeated by encryption
  - Data portions and some header information can be encrypted
- Other problems...

Firewall vs. NIDS

- Firewall
  - Active filtering
  - Fail-close
- Network IDS
  - Passive monitoring
  - Fail-open

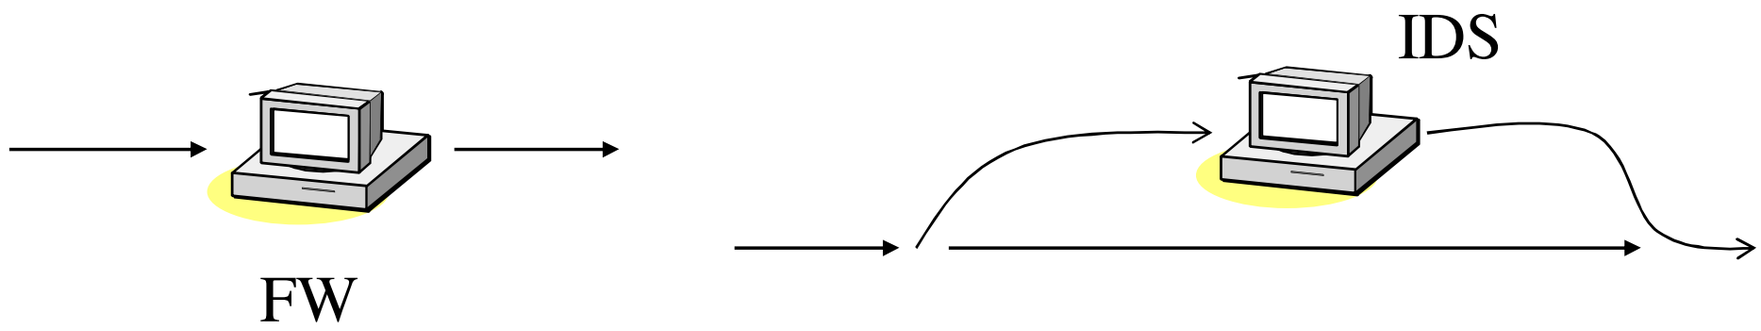

NIDS Requirements

- High-speed, large volume monitoring
  - No packet filter drops
- Real-time notification
- Mechanism separate from policy
- Extensible
- Broad detection coverage
- Economy in resource usage
- Resilience to stress
- Resilience to attacks upon the IDS itself!

Two Well-known NIDS

Bro

Alternative?

Bro

Vern Paxson at ICSI

remember TRW?

Bro: A System for Detecting Network Intruders in Real-Time - USENIX Security 1998

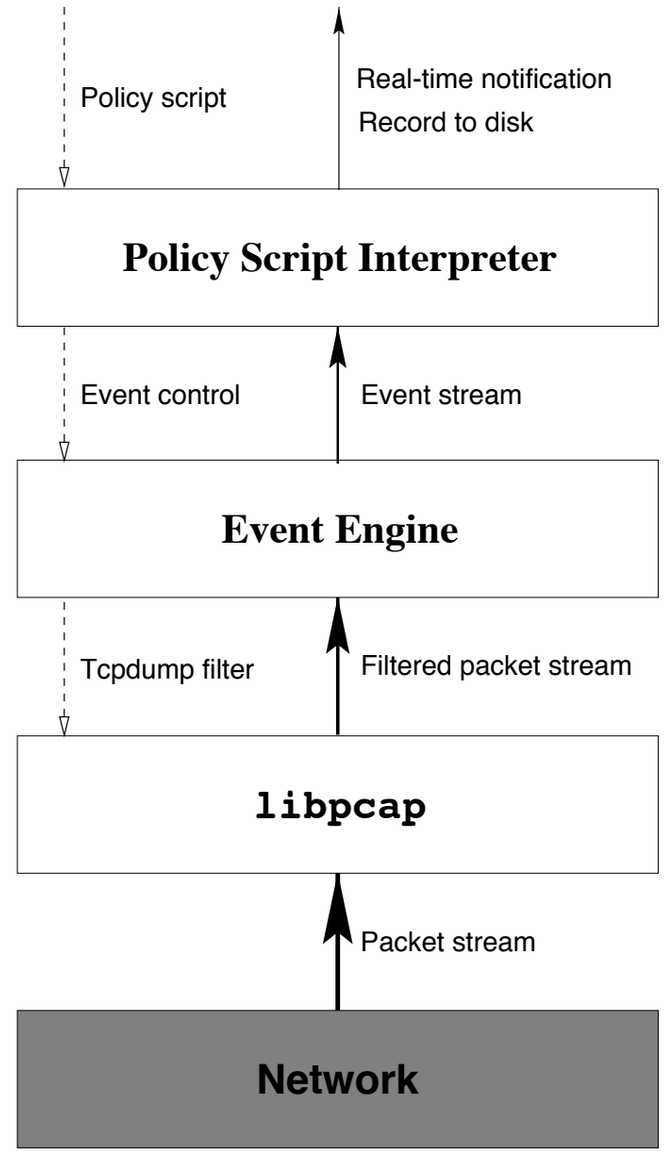

Bro: How it works

- Taps GigEther fiber link passively, sends up a copy of all network traffic.

Bro: How it works

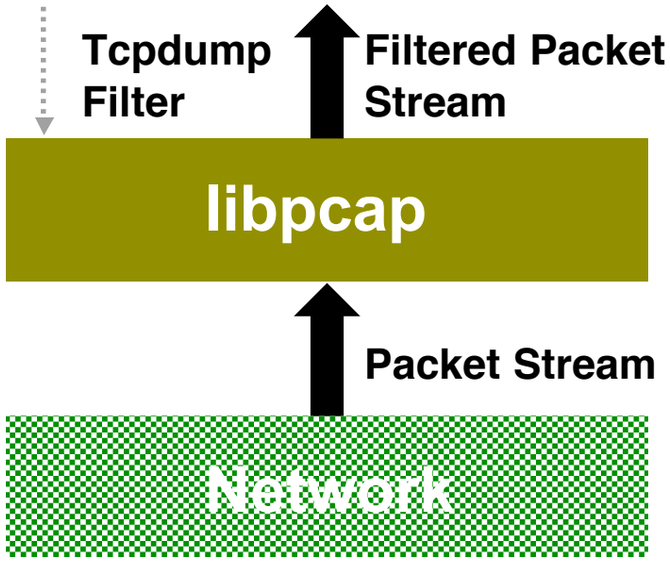

- Kernel filters down high-volume stream via standard *libpcap* packet capture library.

Bro: How it works

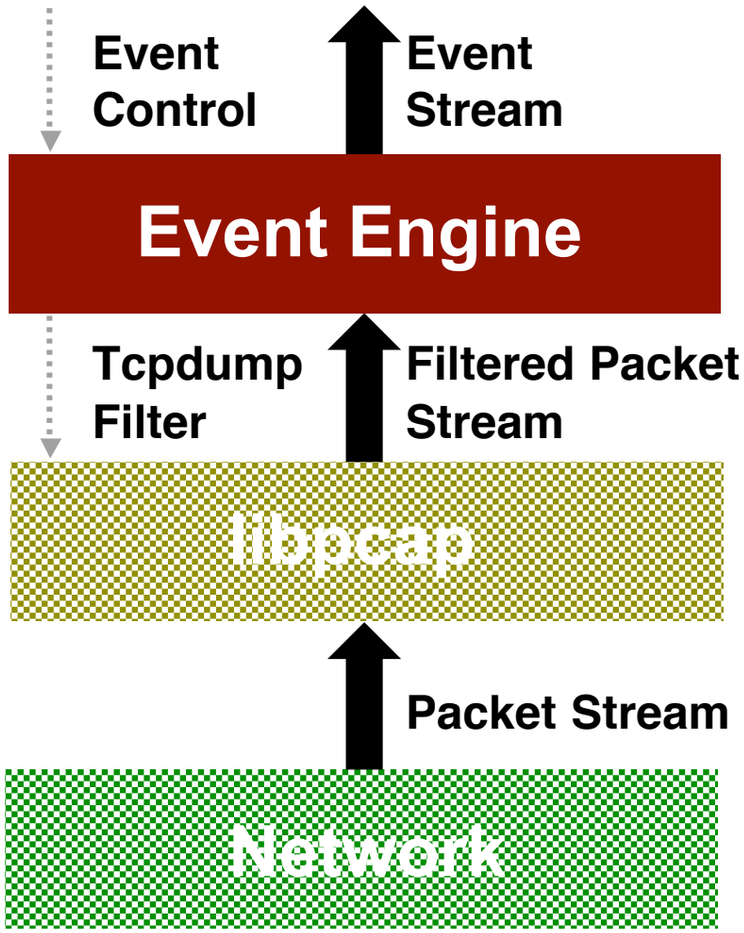

- "Event engine" distills filtered stream into high-level, *policy-neutral* events reflecting underlying network activity
  - E.g. Connection-level:
    - connection attempt
    - connection finished
  - E.g. Application-level:
    - ftp request
    - http_reply
  - E.g. Activity-level:
    - login success

Bro: How it works

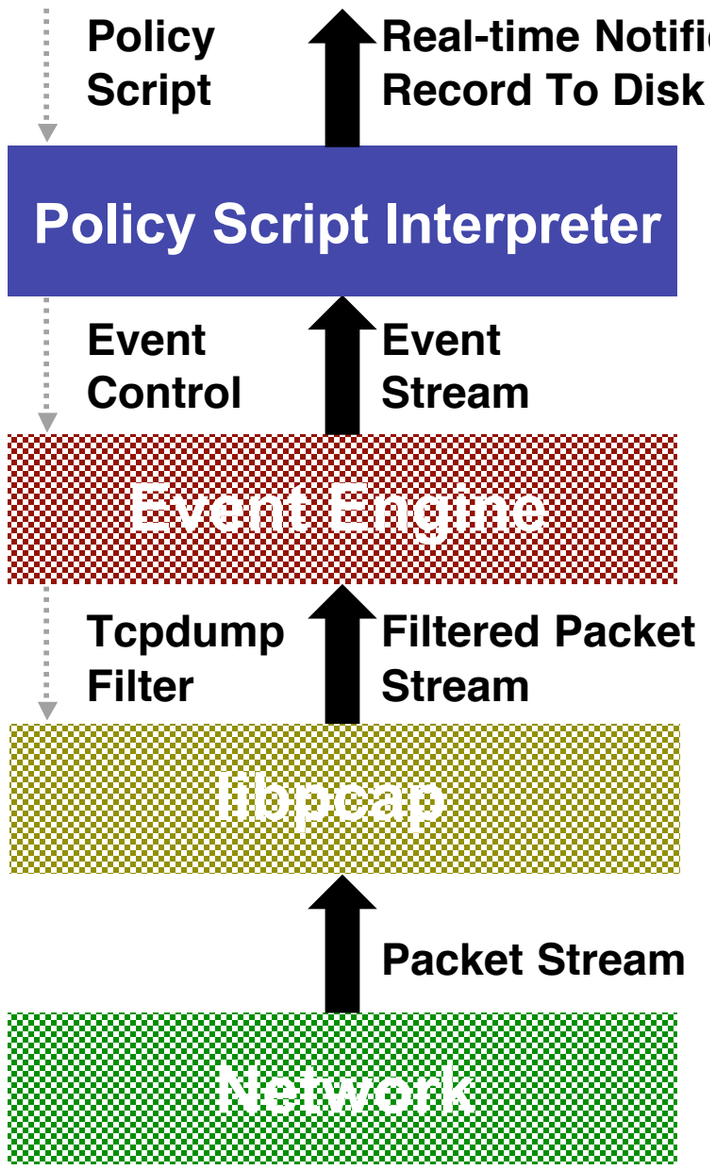

- "Policy script" processes event stream, incorporates:
  - Context from past events
  - Site's particular policies

Snort

open source
SourceFire leads this project
now commercial??

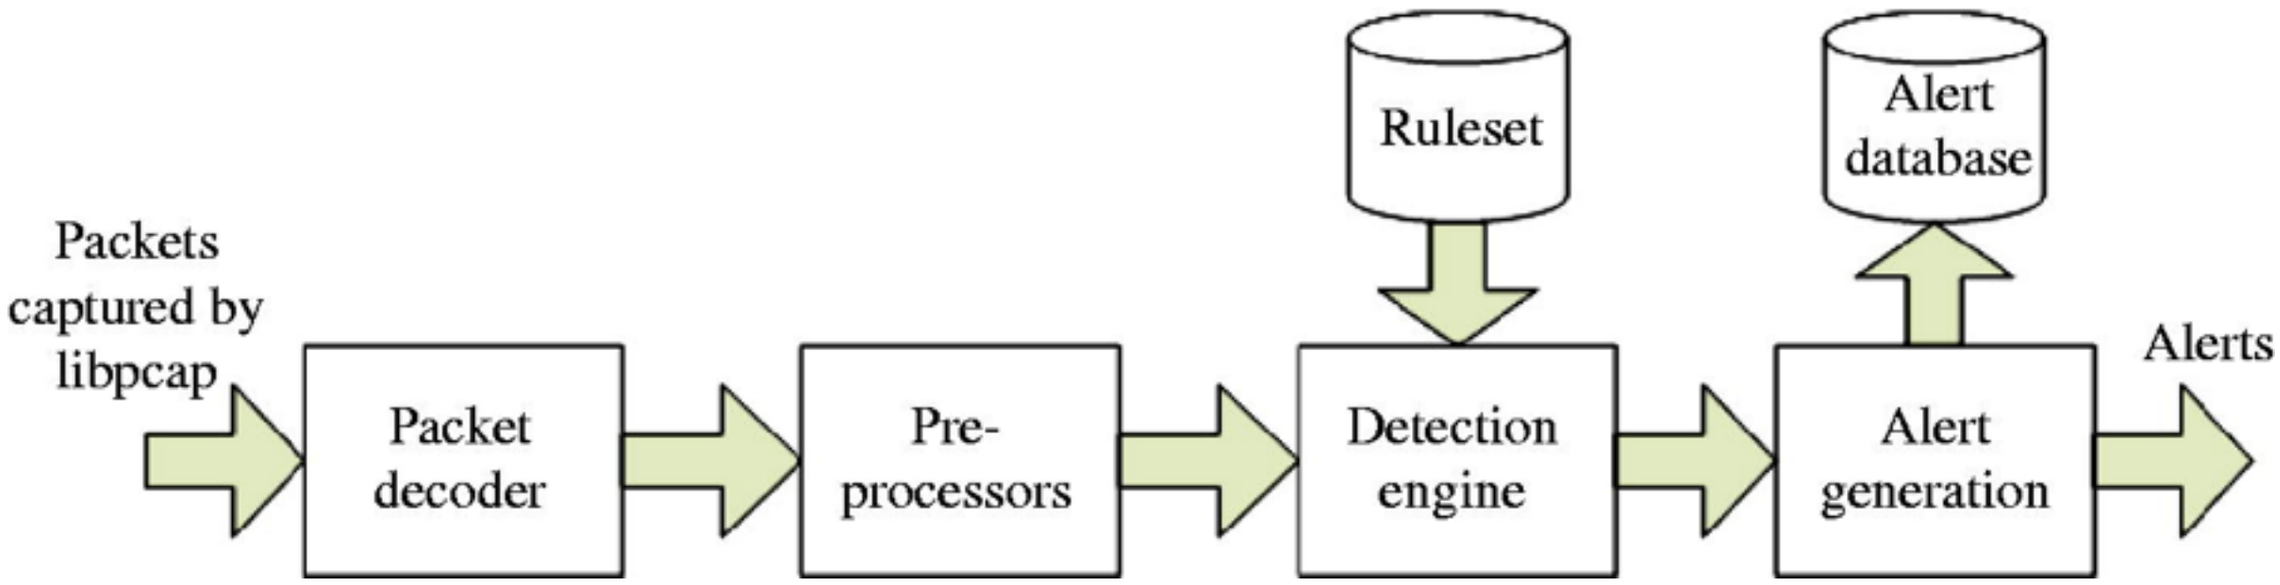

Snort: Rule

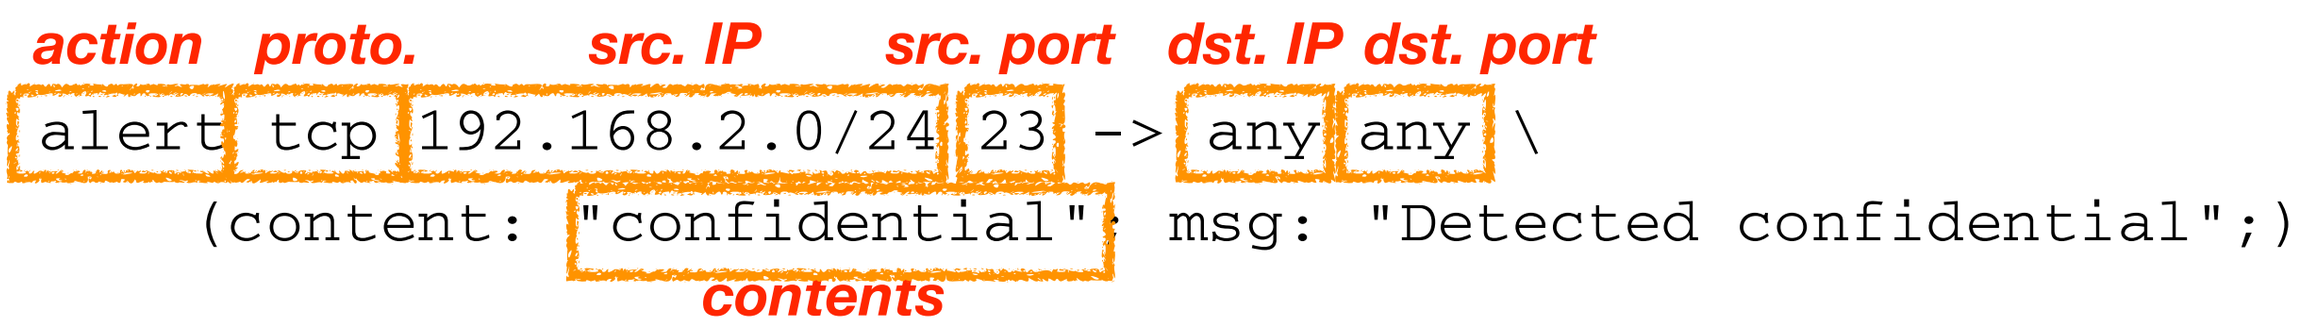
action
proto.
src. IP
src. port
dst. IP
dst. port
alert tcp 192.168.2.0/24 23 -> any any \
(content: "confidential"; msg: "Detected confidential";)
contents

Eluding NIDS

- What the IDS sees may not be what the end system gets.
  - Insertion and evasion attacks.
    - IDS needs to perform full reassembly of packets.
  - But there are still ambiguities in protocols and operating systems:
    - E.G. TTL, fragments.
    - Need to “normalize” the packets.

Insertion Attack

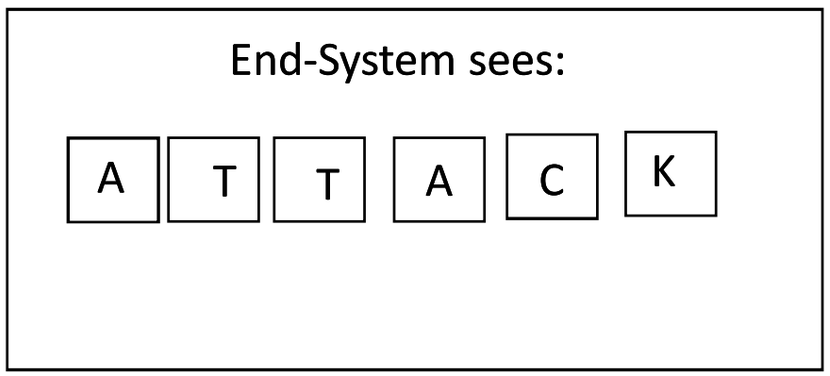

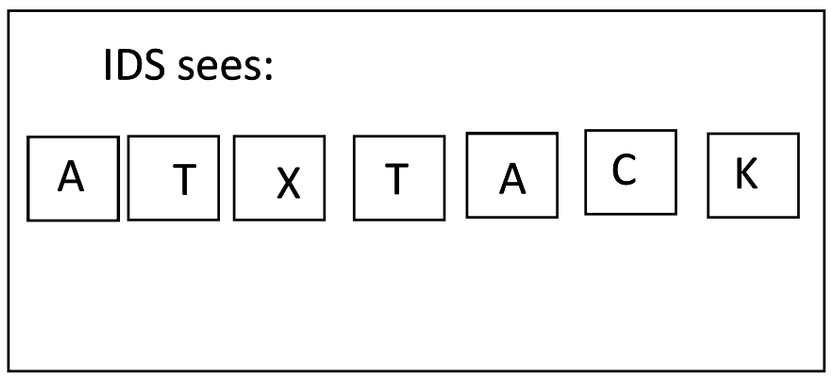

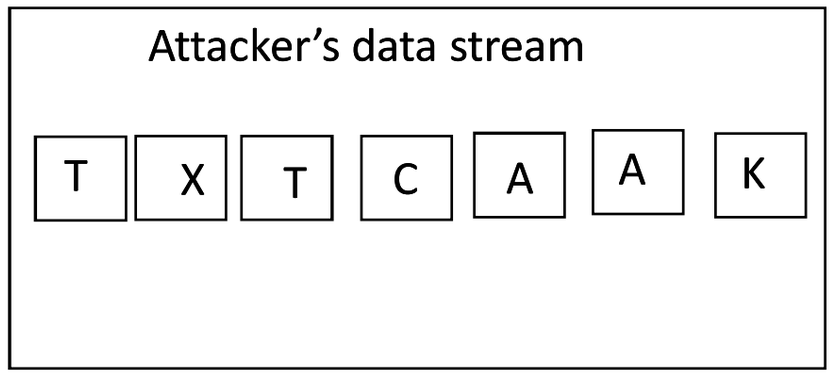

Examples: bad checksum, TTL.

Evasion Attack

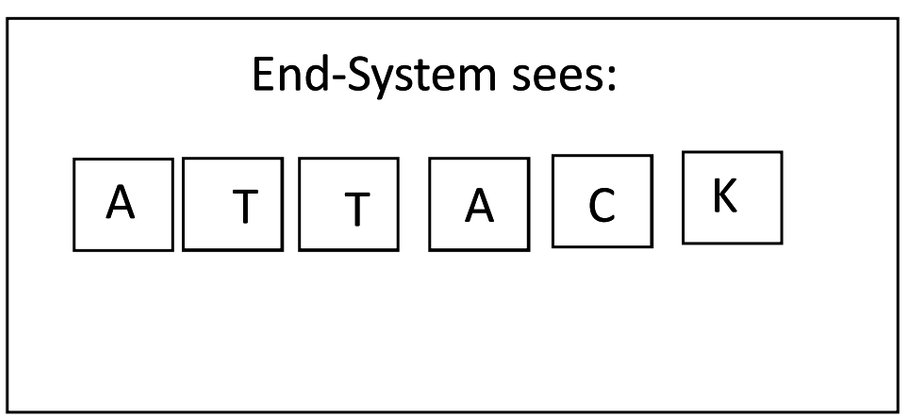

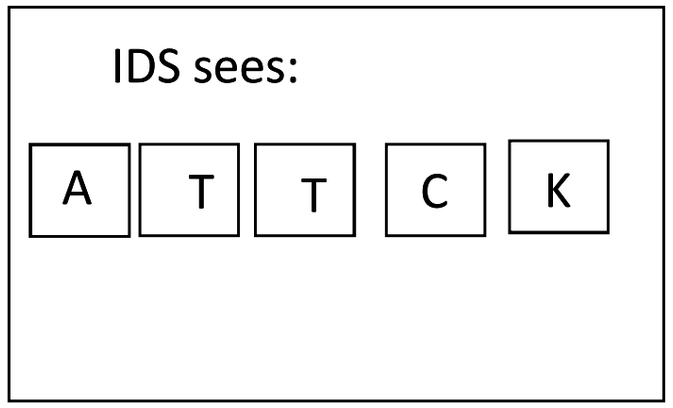

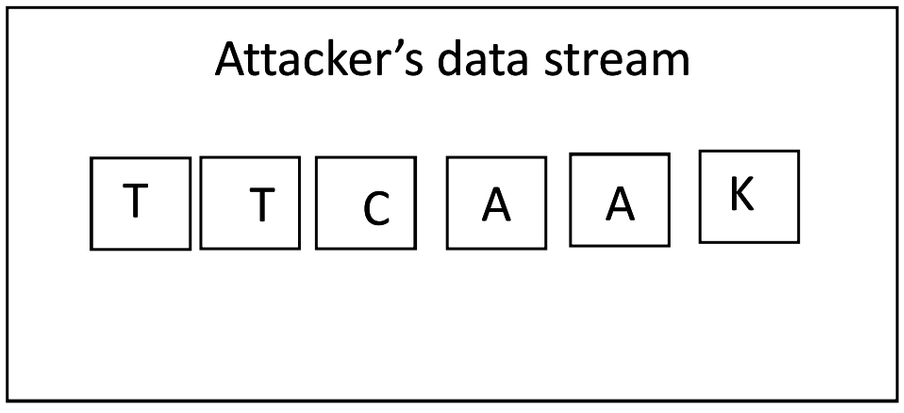

Example: fragmentation overlap

Summing up

- Network intrusion
  - A set of actions aimed to compromise the security goals, namely
    - Integrity, confidentiality, or availability, of a computing and networking resource
- Detecting network intrusion
  - Method
    - misuse vs. anomaly
  - Placement
    - Host level
      - AV-tools
    - Network level
      - Snort, Bro