CE 817 - Advanced Network Security

Lecture 5

Mehdi Kharrazi
Department of Computer Engineering
Sharif University of Technology

Acknowledgments: Some of the slides are fully or partially obtained from other sources. Reference is noted on the bottom of each slide, when the content is fully obtained from another source. Otherwise a full list of references is provided on the last slide.

Problem

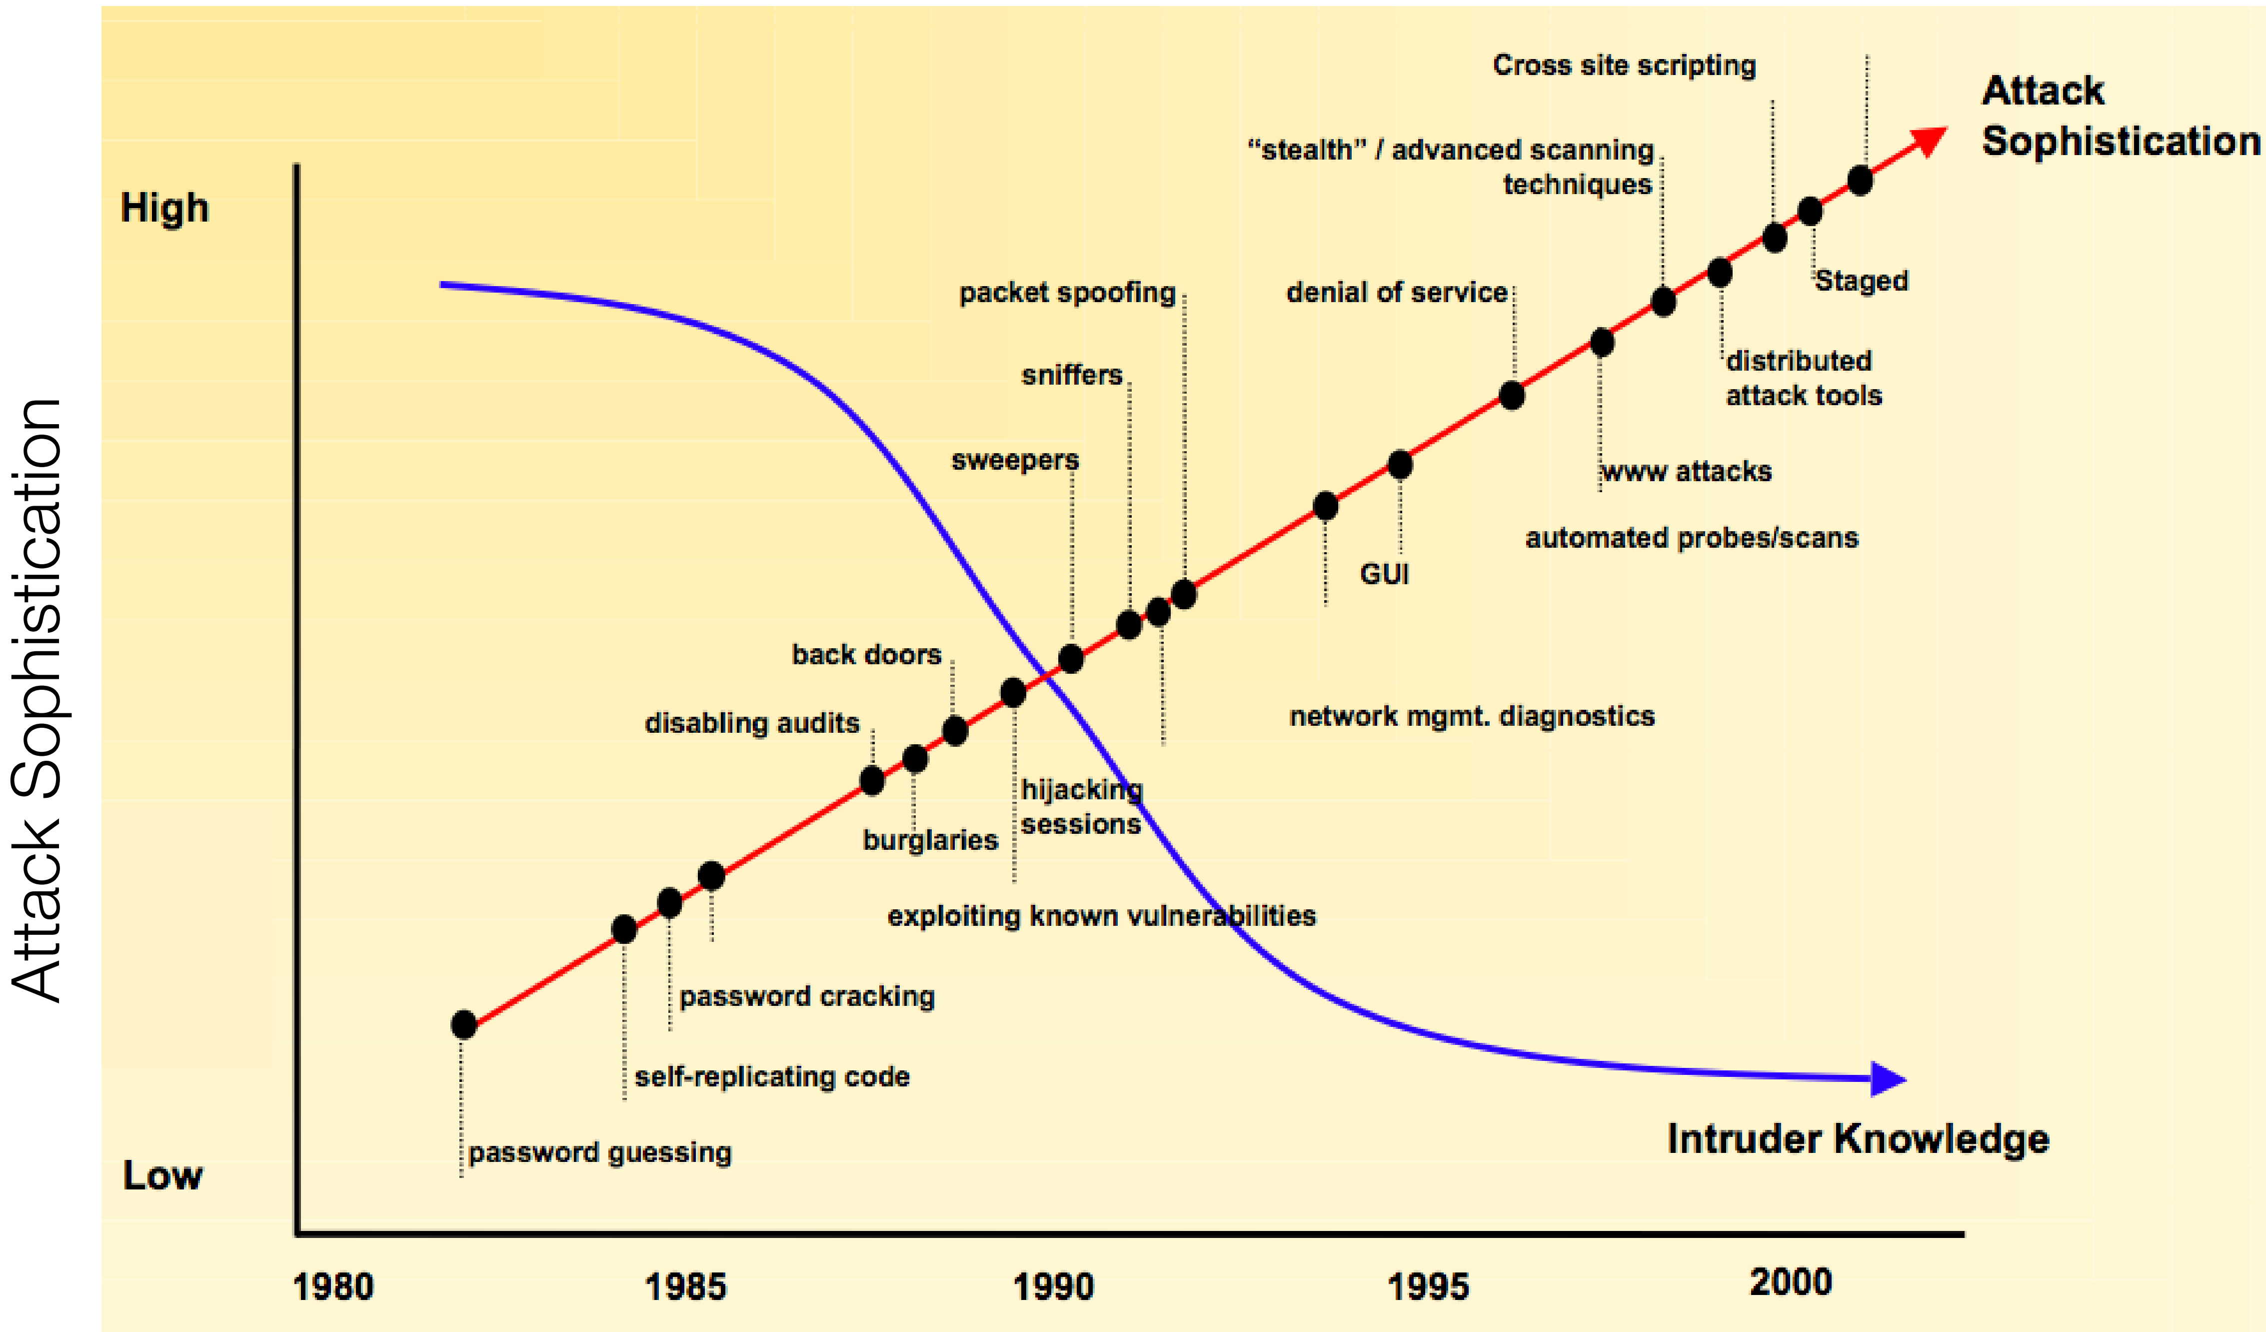

Attack Sophistication
High
Low
Attack Sophistication
Intruder Knowledge
Cross site scripting
"stealth" / advanced scanning techniques
packet spoofing
denial of service
Staged
sniffers
distributed attack tools
sweepers
www attacks
automated probes/scans
GUI
back doors
network mgmt. diagnostics
disabling audits
hijacking sessions
burglaries
exploiting known vulnerabilities
password cracking
self-replicating code
password guessing
1980
1985
1990
1995
2000

Fighting intrusion

- Prevention: isolate from network, strict authentication measures, encryption
- Preemption:
  - “do to others before they do to you”
- Deterrence: dire warnings,
  - “we have a bomb too.”
- Deflection: diversionary techniques to lure away
- Detection
- Counter attacks

Defense in Depth

- More generically, most single defenses can fail
- We always need defense in depth – multiple layers, of different designs and philosophies
- One such layer: Intrusion Detection Systems

What is IDS?

- An Intrusion Detection System (IDS) is a system that attempts to identify intrusions.
- Intrusion detection is the process of identifying and responding to malicious activity targeted at computing and networking resources.

Examples of IDS in daily life

- Car Alarms
- House Alarms
- Surveillance Systems
- Spy Satellites, and spy planes (U2 and SR-71)

Elements of Intrusion Detection

- 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 evidence from audit data
    - Models - piece evidence together; infer attack
  - From a system architecture perspective:
    - Audit data processor, knowledge base, decision engine, alarm generation and responses

Where Are IDS Deployed?

- Host-based
  - Monitor activity on a single host
  - Advantage: better visibility into behavior of individual applications running on the host
- Network-based (NIDS)
  - Often placed on a router or firewall
  - Monitor traffic, examine packet headers and payloads
  - Advantage: single NIDS can protect many hosts and look for global patterns

Host-Based IDSs

- 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 execution of system programs
  - E.G., Analyze system calls made by sendmail

 [Stolfo06]

Basic Audit Modules (Hosts)

- eventLog - Uses the windows Event Logging system to track entries into all three of the windows event logs: System, Security, Application
- netstat - Uses the information from the program netstat to provide information about network usage on the machine
- health - Runs the program health to give current information about the system (CPU usage, mem usage, swap usage)
- ps - Uses information from the /proc virtual file system as a data source

Network IDSs

- 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
- May be easily defeated by encryption
  - Data portions and some header information can be encrypted
- Other problems …

Architecture of Network IDS

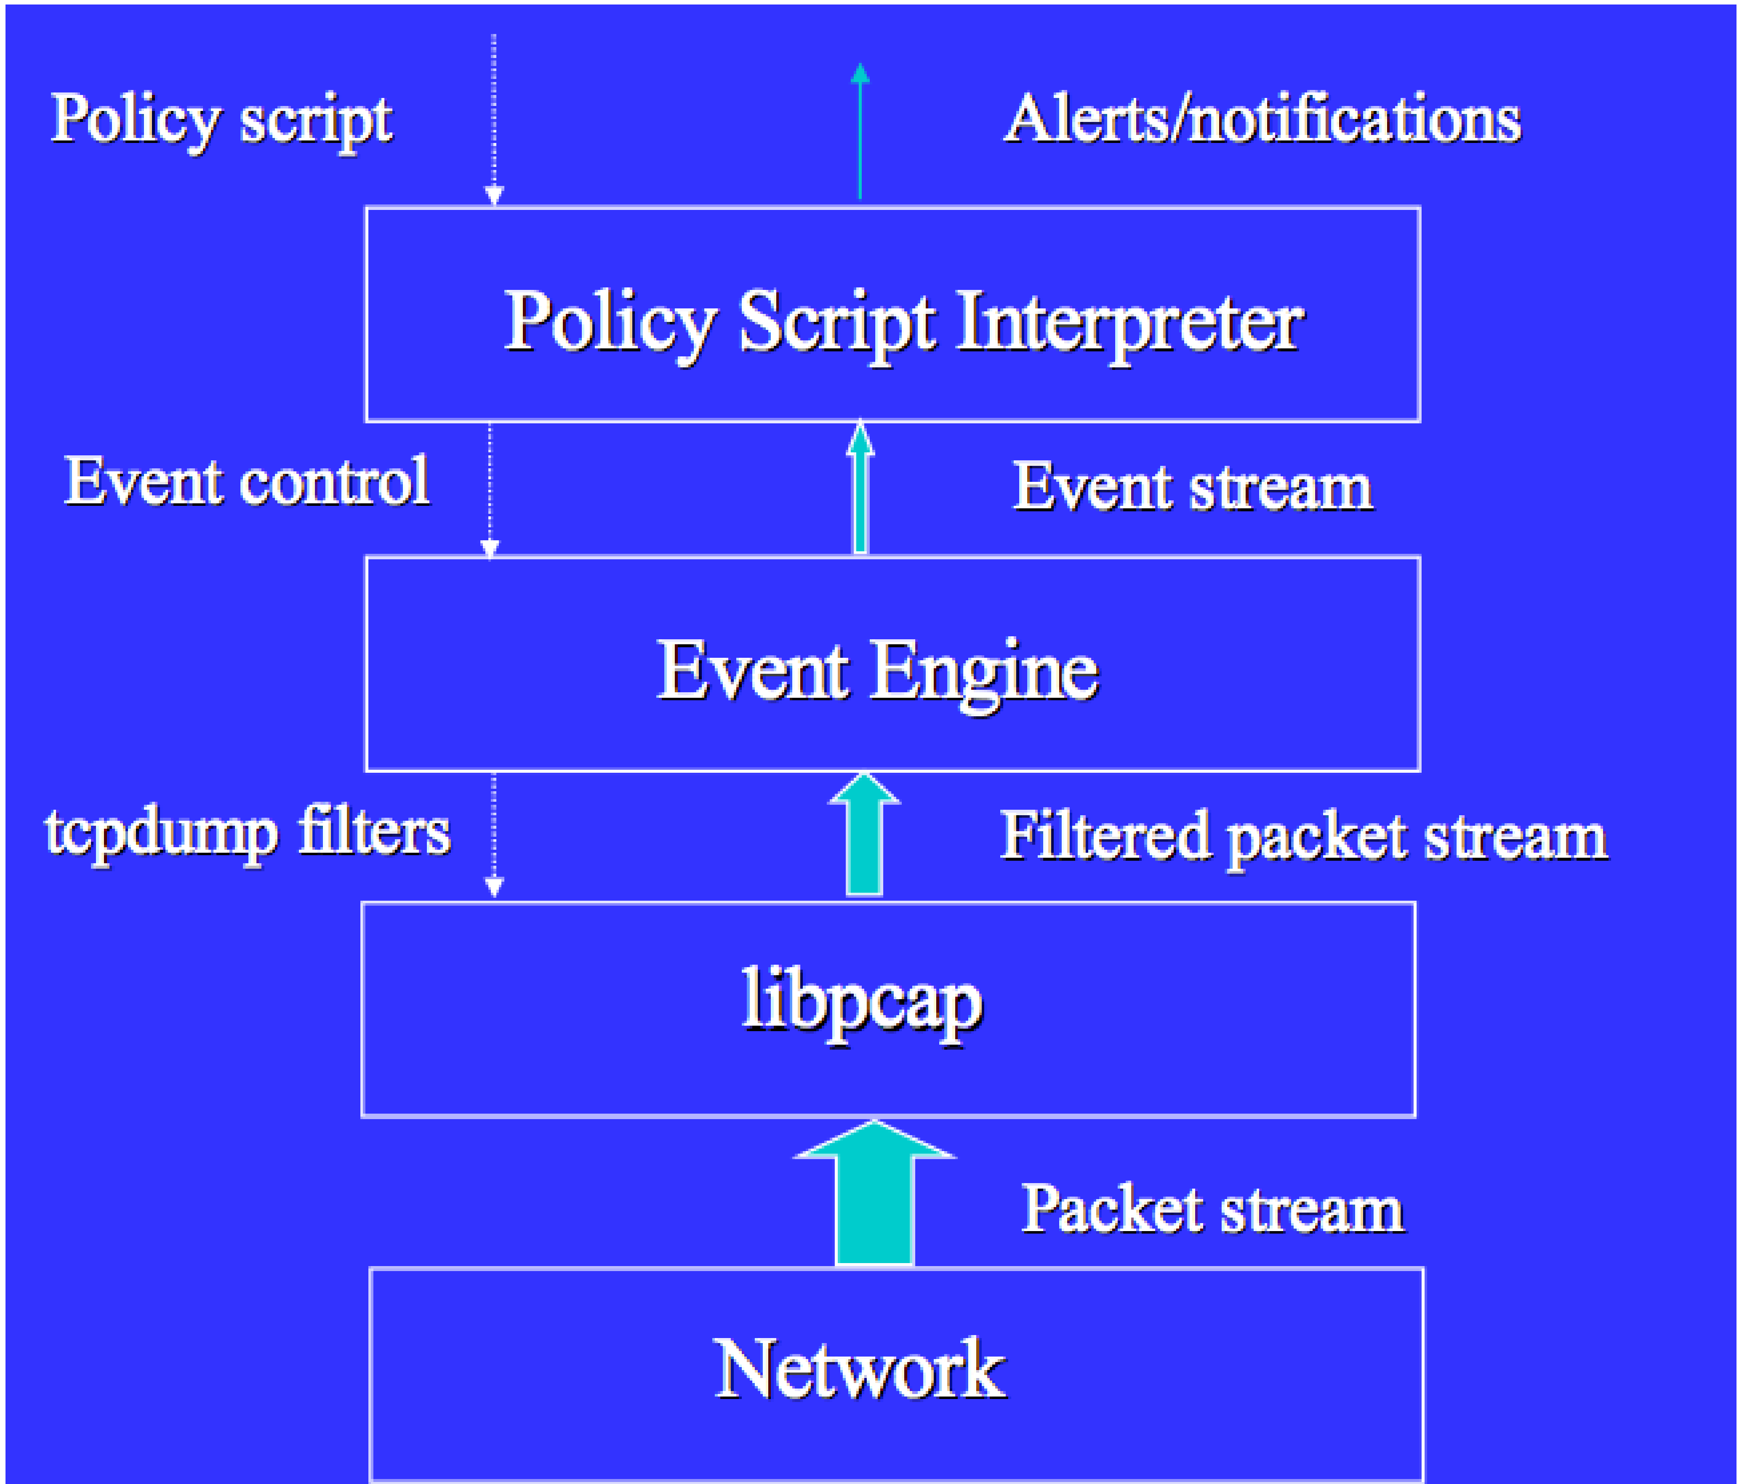

 [Stolfo06]

Firewall Versus Network IDS

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

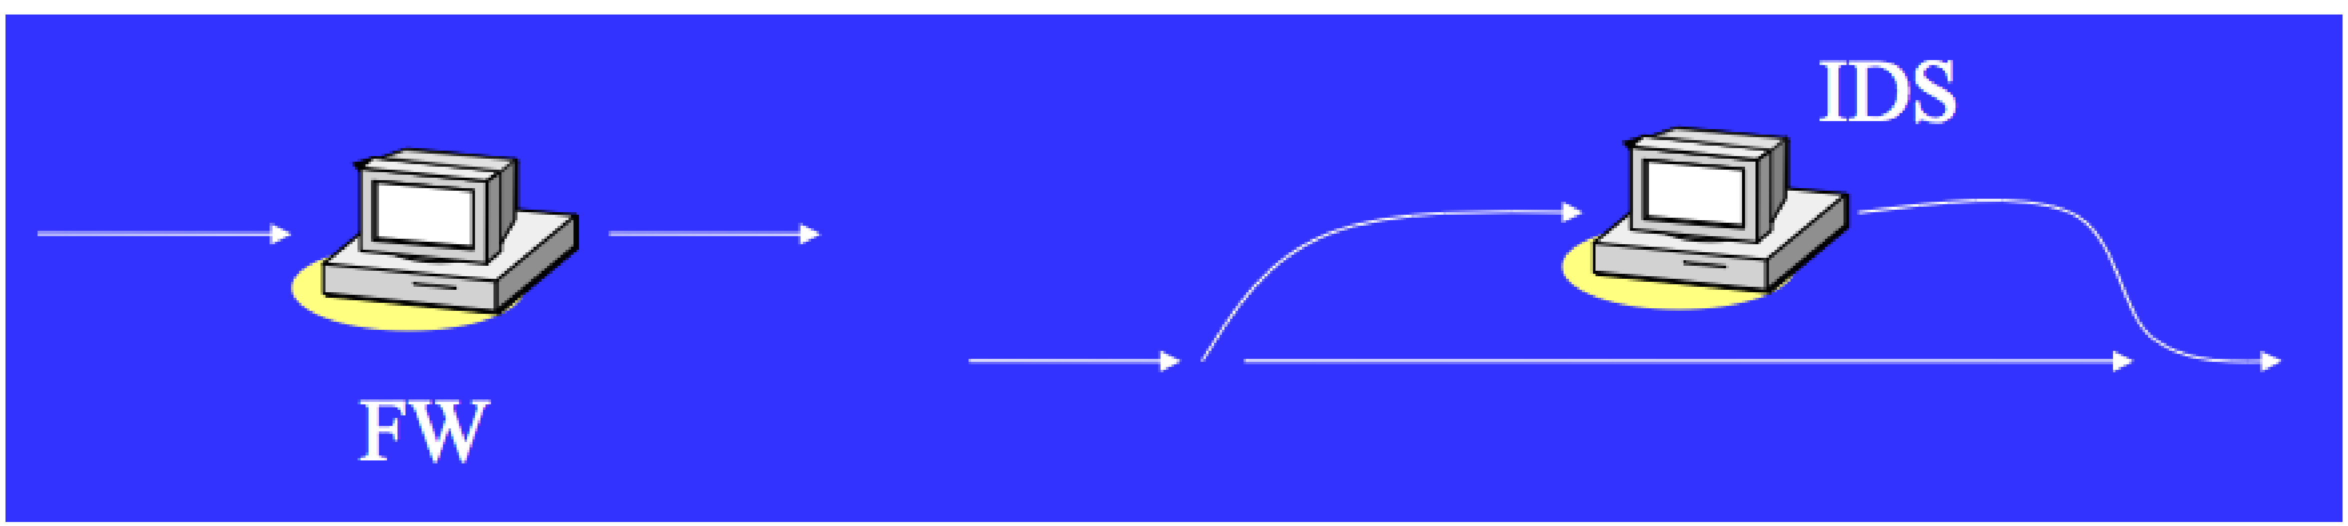

Requirements of Network IDS

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

Eluding Network IDS

- 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

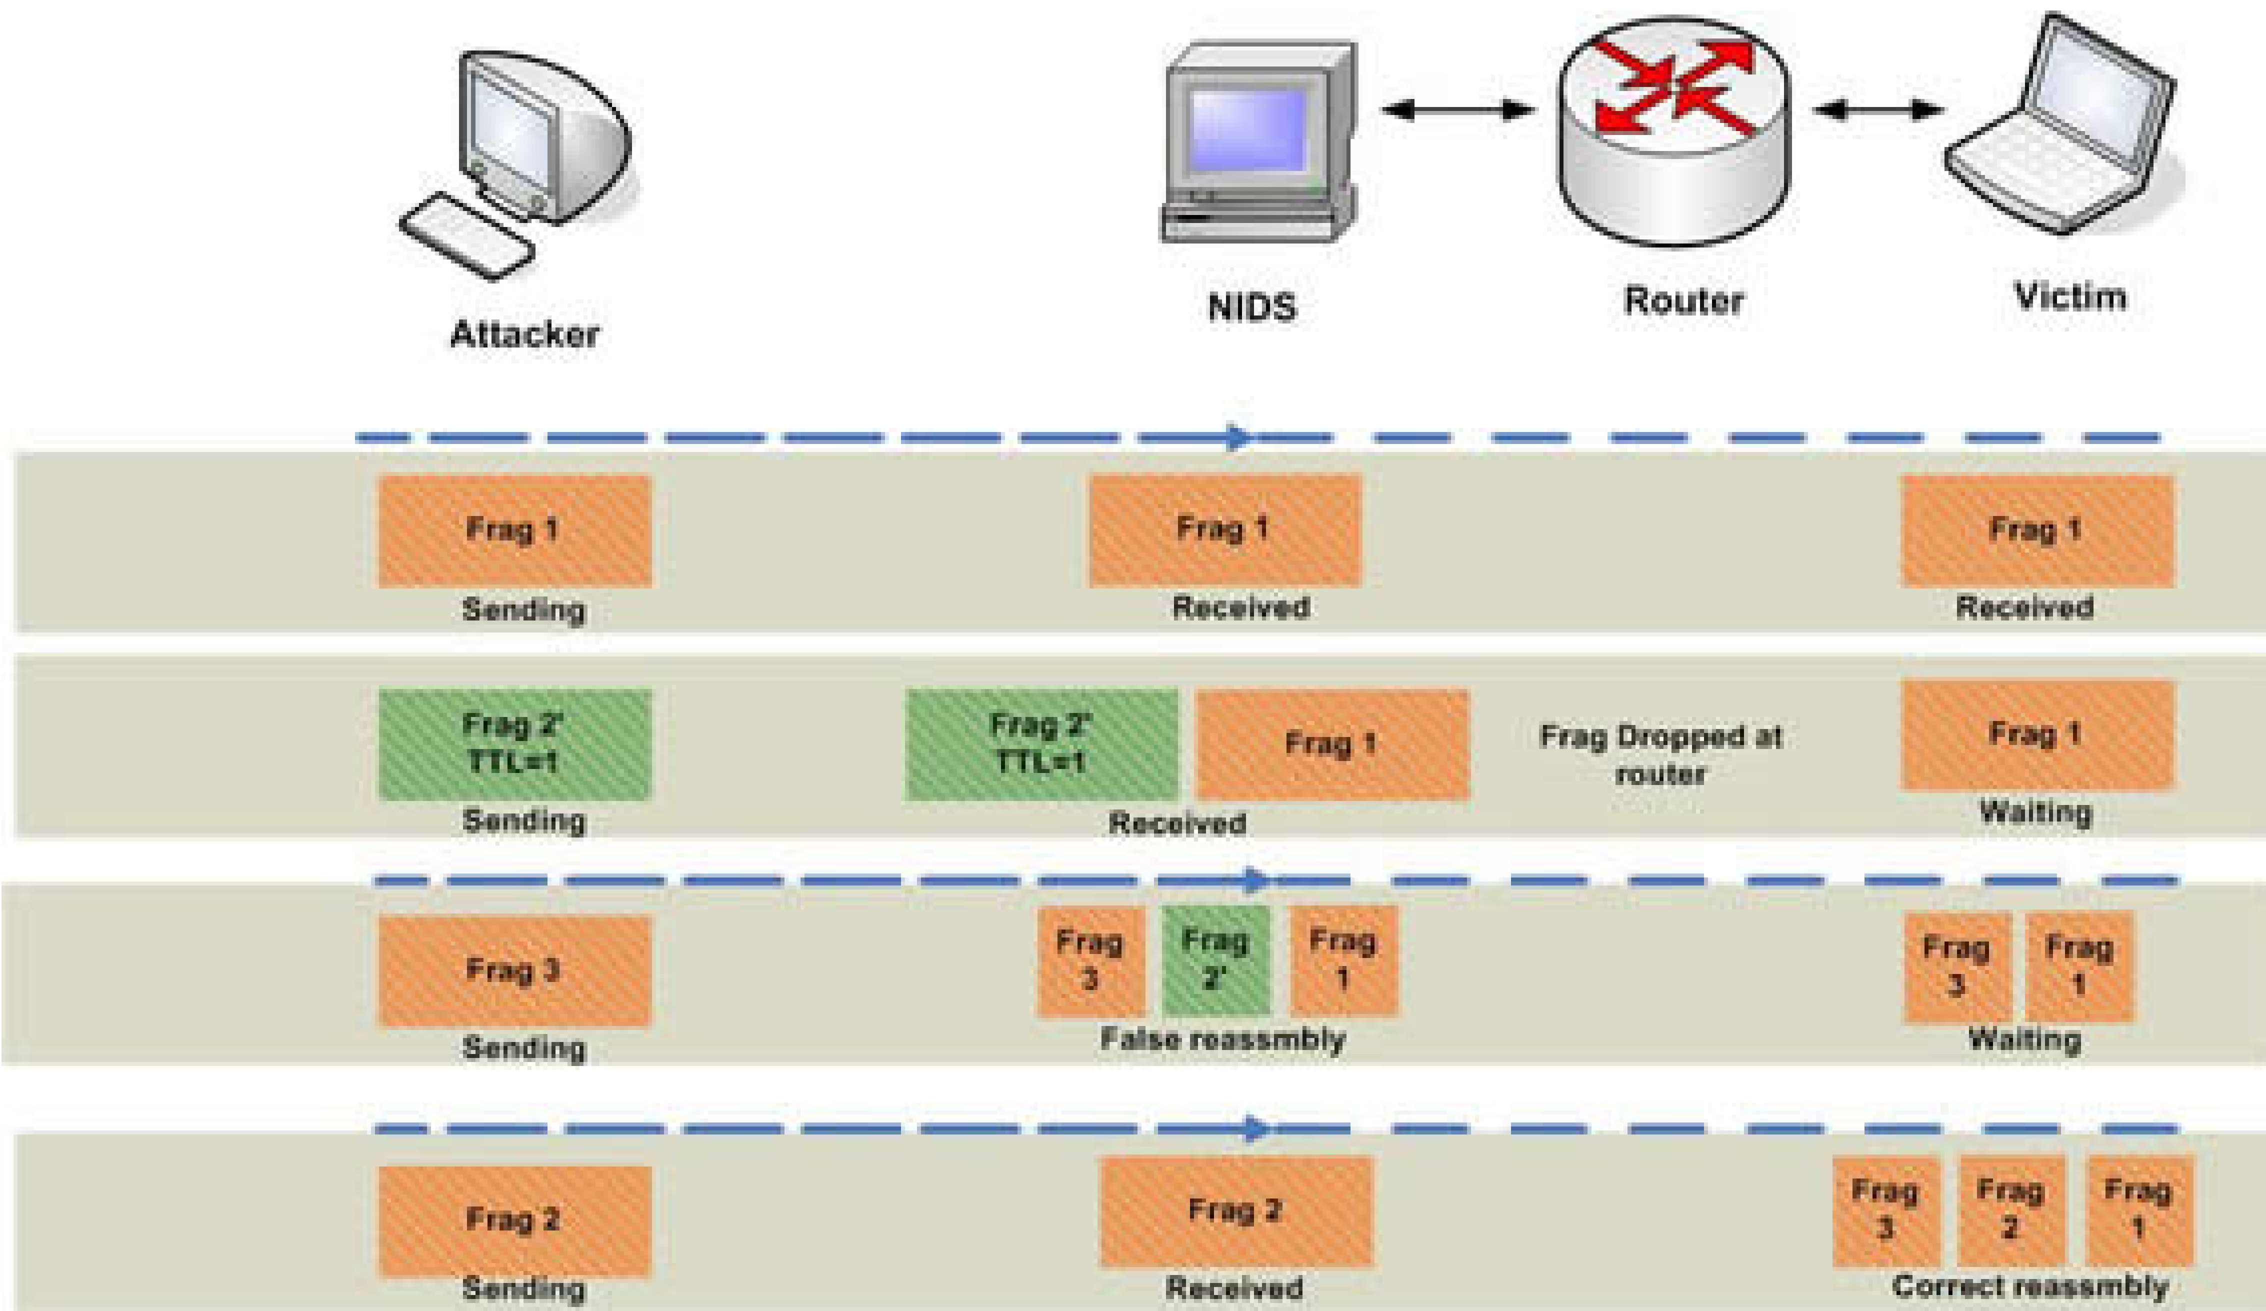

http://www.securityfocus.com/infocus/1852

Insertion Attack

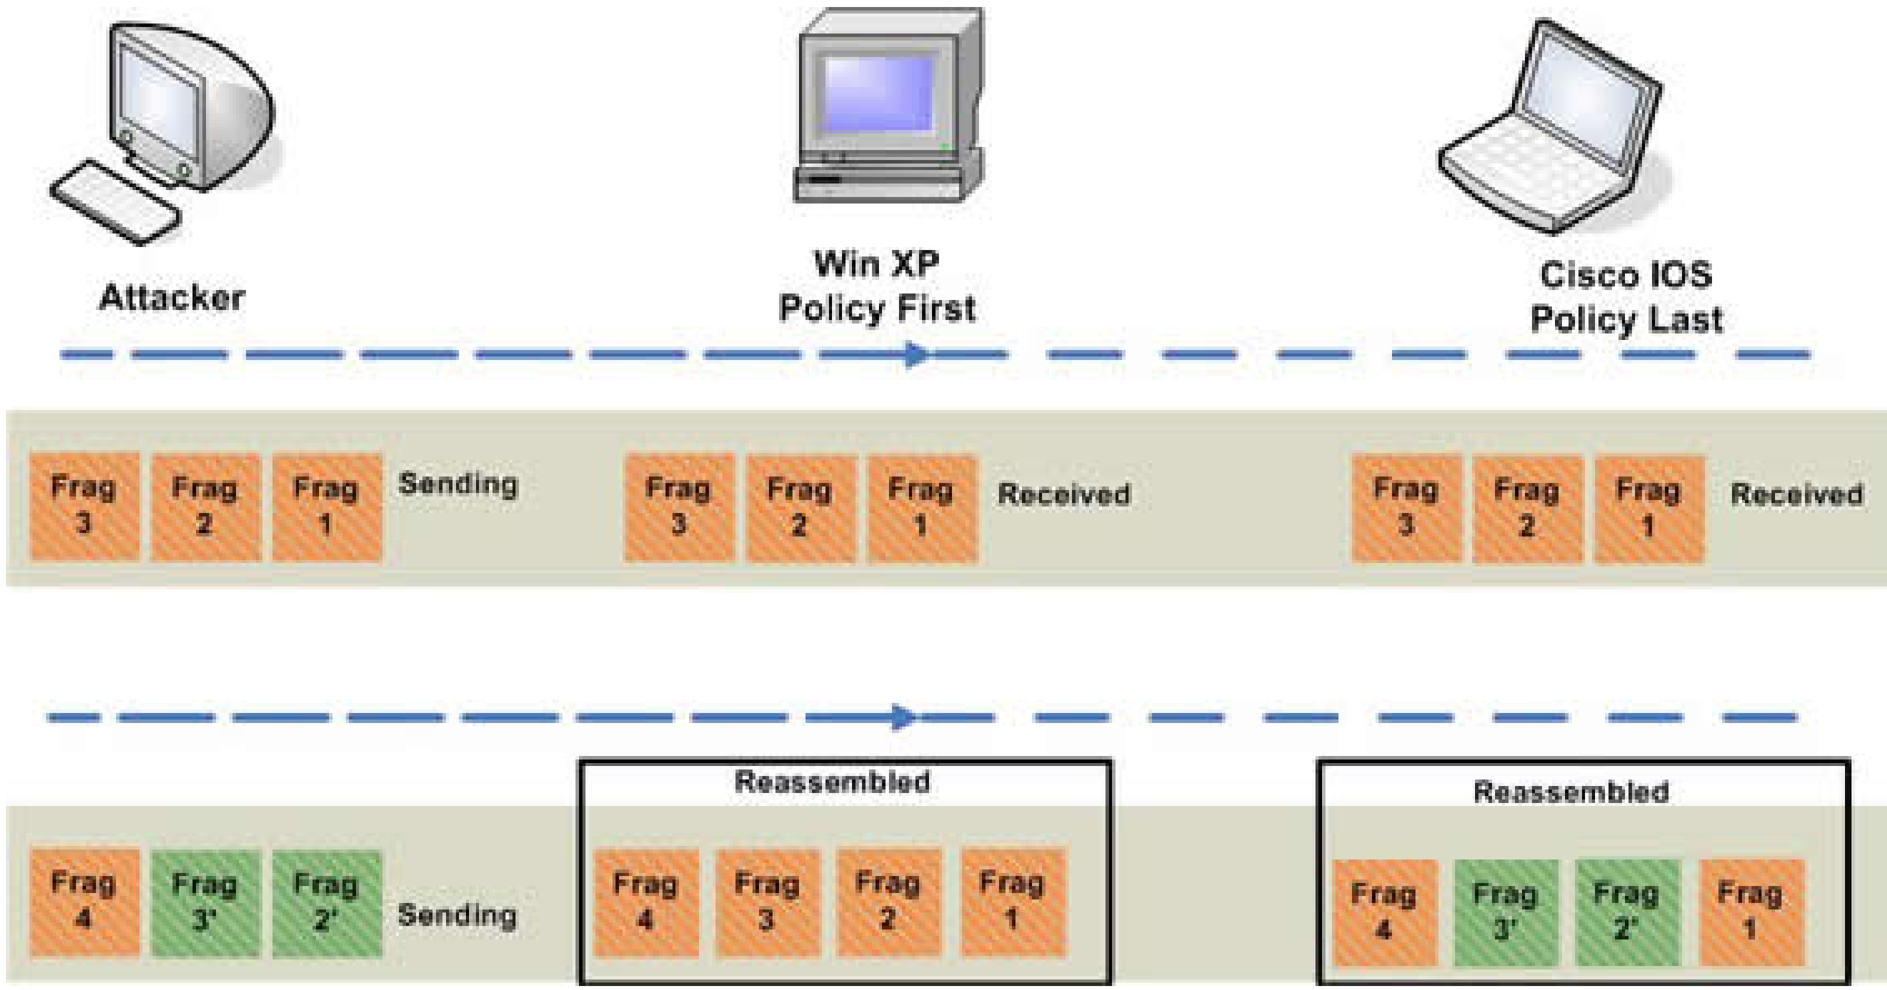

- **First**. This is where the operating System favors the original fragments with a given offset. For example, Windows 95/98/NT4/ME/W2K/XP/2003. **Last**. This is where the operating System favors the subsequent fragments with a given offset. For example, Cisco IOS.

DoS Attacks on Network IDS

- Resource exhaustion
  - CPU resources
  - Memory
  - Network bandwidth
- Abusing reactive IDS
  - False positives

Hybrid NIDS and HIDS

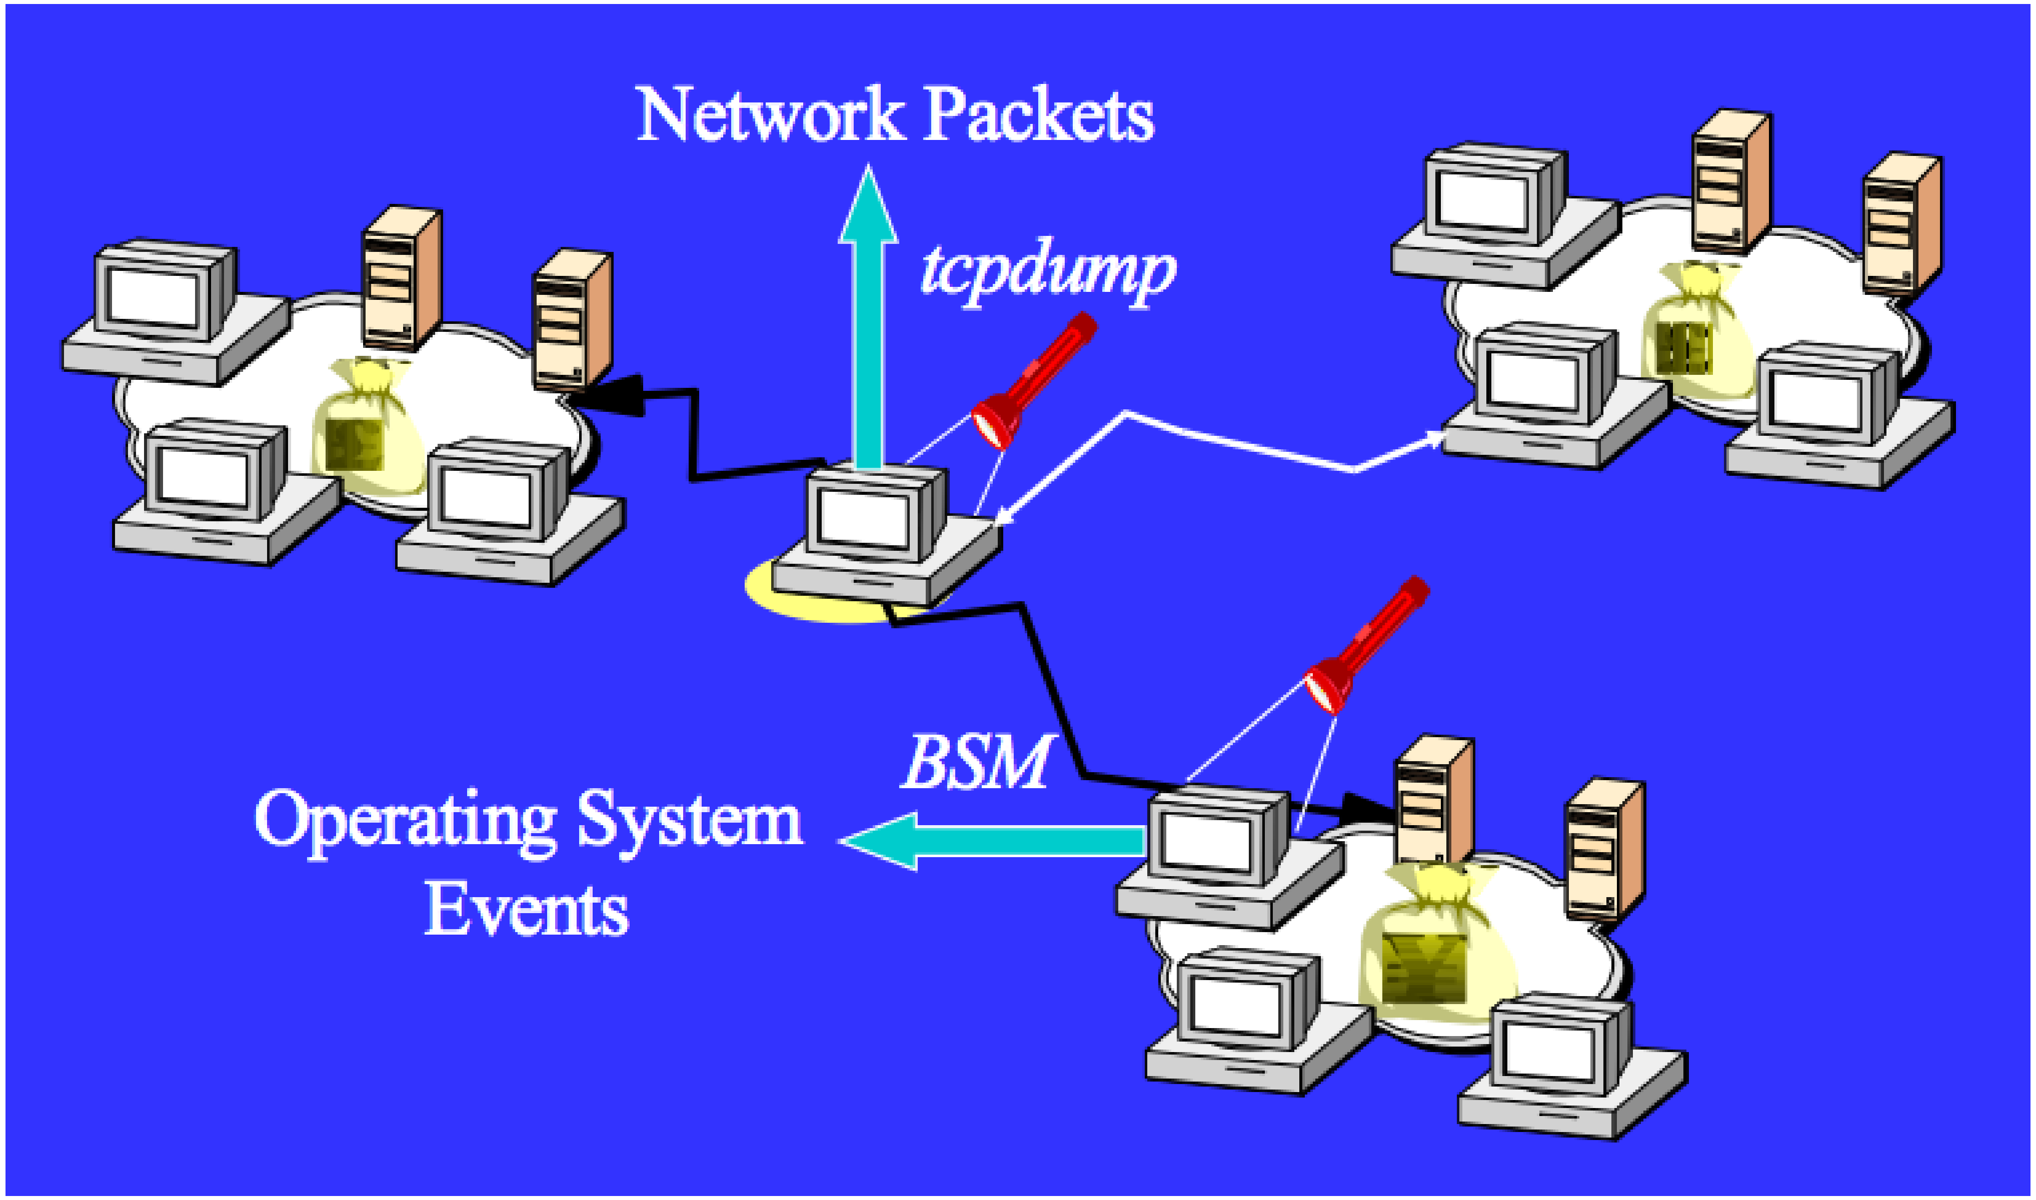

Hybrid NIDS and HIDS

- Correlate information from multiple sources
- How do you trust your sources?

Taxonomy of IDS’s

Intrusion Detection Approaches

- Modeling
  - Features: evidences extracted from audit data
  - Analysis approach: piecing the evidences together
    - Misuse detection (a.k.a. signature-based)
    - Anomaly detection (a.k.a. statistical-based)
- Deployment: Network-based or Host-based
- Development and maintenance
  - Hand-coding of “expert knowledge”
  - Learning based on audit data

A Generic IDS

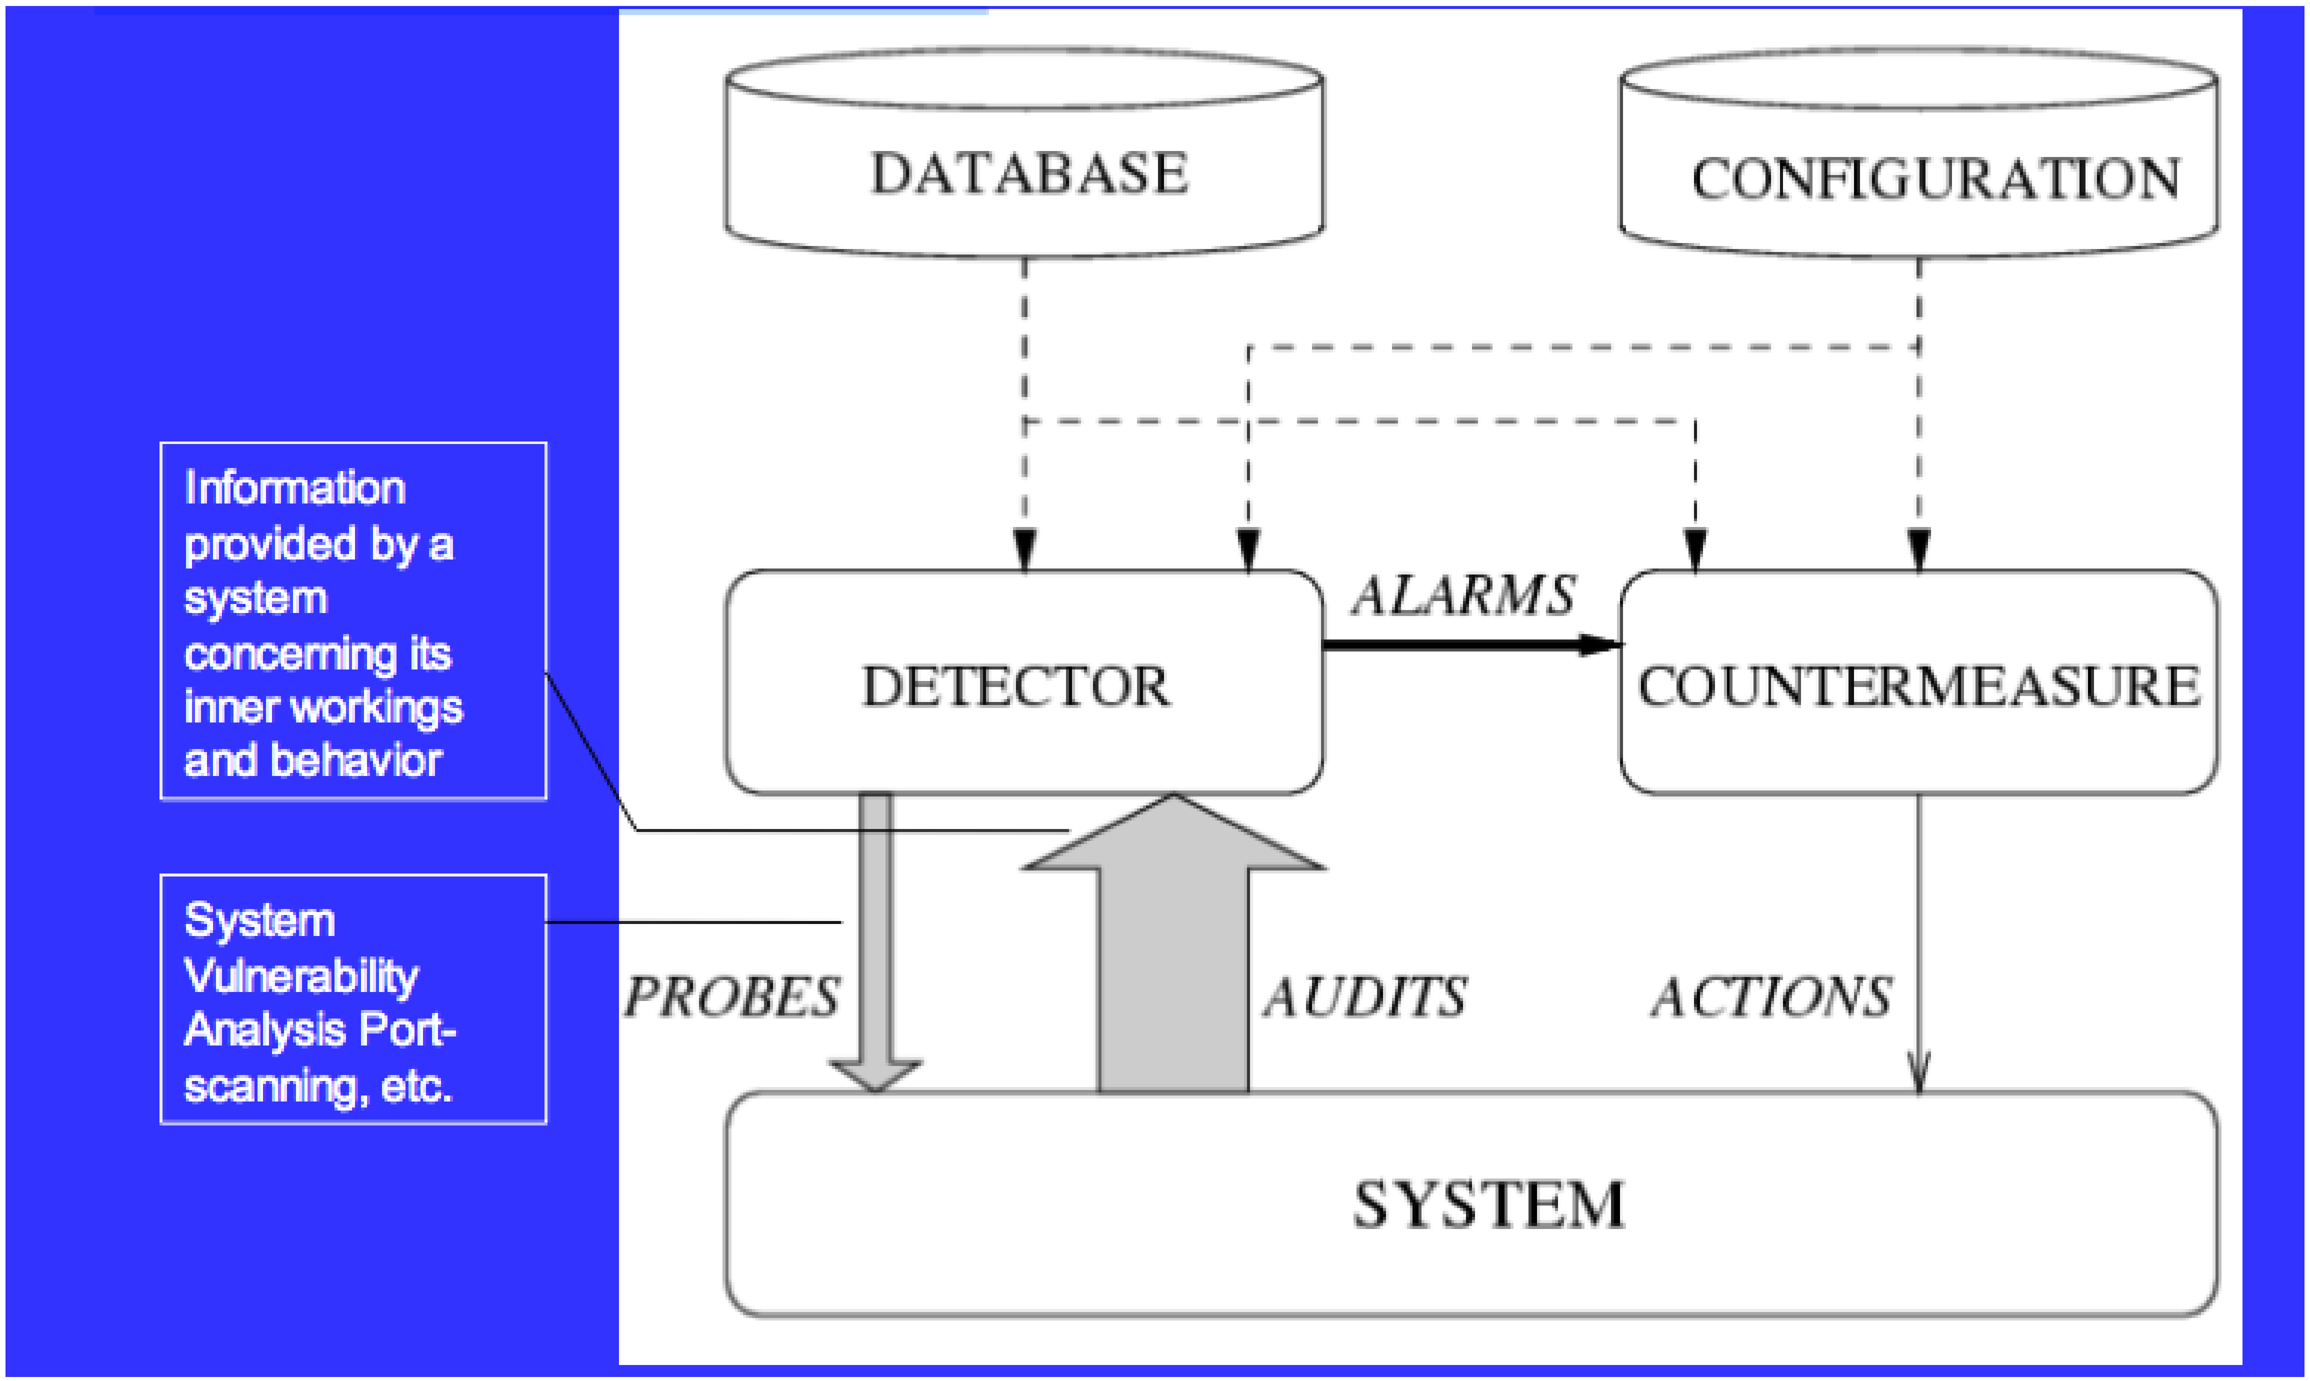

Characteristics of IDS

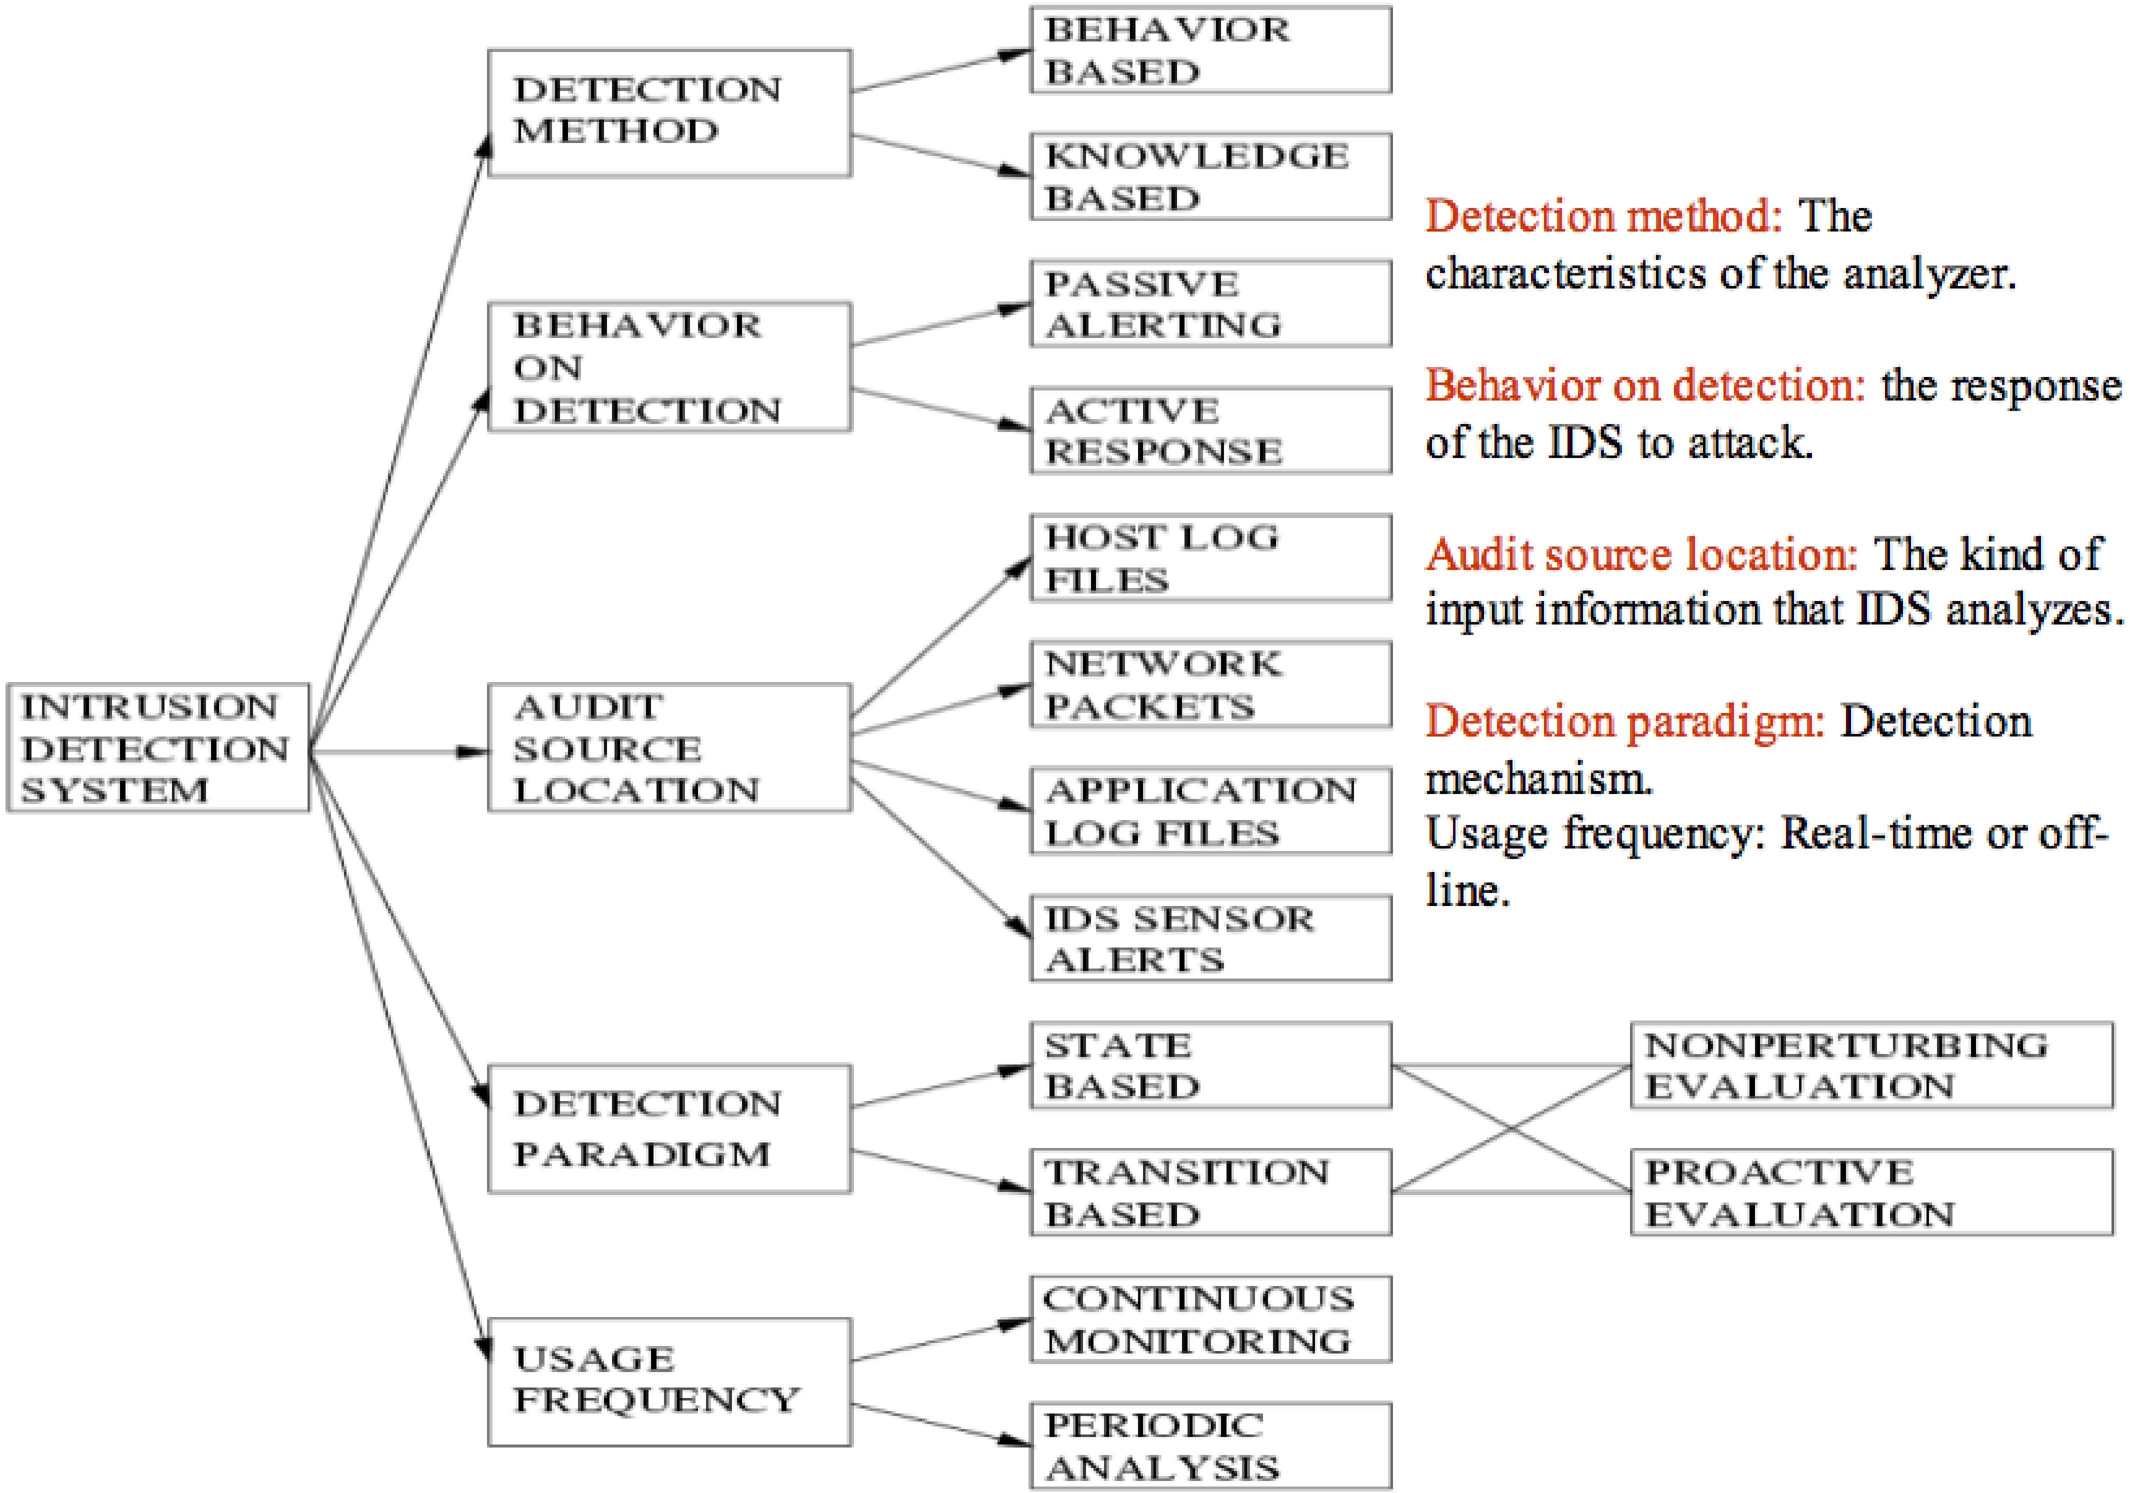

Detection method: The characteristics of the analyzer.

Behavior on detection: the response of the IDS to attack.

Audit source location: The kind of input information that IDS analyzes.

Detection paradigm: Detection mechanism.
Usage frequency: Real-time or off-line.

Detection Paradigm

- State-based versus transition-based IDS
  - State-based: Identifies intrusions on the states
  - Transition-based: Watches events that trigger transition from one state to another
- Non-perturbing versus pro-active analysis of state or transition
  - Non-perturbing: Acquire information transparently
  - Pro-active: Analysis by explicitly triggering events

IDS: Time aspect

- Real-time IDS
  - Analyzes the data while the sessions are in progress
  - Raises an alarm immediately when the attack is detected
- Off-line IDS
  - Analyzes the data after the information has been already collected
  - Useful for understanding the attackers' behavior

Knowledge-based IDS

- Good accuracy, bad completeness
  - Drawback: need regular update of knowledge
  - Difficulty of gathering the information
  - Maintenance of the knowledge is a time-consuming task
- Knowledge-based IDS
  - Expert systems
  - Signature analysis
  - State-transition analysis

Misuse Detection

- The system is equipped with a number of attack descriptions (“signature”). Then matched against the audit data to detect attacks.
- Pro: less false positives (But there still some!)
- Con: cannot detect novel attacks, need to update the signatures often.
- Approaches: pattern matching, security rule specification.

Specification-based Detection

- Manually develop specifications that capture legitimate (not only previous seen) system behavior. Any deviation from it is an attack
- Pro: can avoid false-positive since the specification can capture all legitimate behavior.
- Con: hard to develop a complete and detailed specification, and error-prone.
- Approach: state machines

Today's IT Security Tools

- We make lists of bad behavior
  - Virus definitions
  - SPAM filters and blacklists
  - IDS signatures
  - Policies
- We distribute the lists to applications and detection systems
- They flag behavior that fits the pattern
- The system is about to collapse
  - Delays
  - Administrative Overhead
  - False positives

Behavior-based IDS

- Good completeness, bad accuracy
- Detect intrusion by observing a deviation from the normal or expected behavior of the system or the users
- Can detect attempts to exploit new and unforeseen vulnerabilities
- Behavior-based IDS
  - Statistics
  - Expert systems
  - Neural networks
  - User intention identification

[Stolfo06]

Anomaly Detection

- Build models of “normal” behavior of a system using machine learning or data mining. Any large deviation from the model is thought as anomaly.
- Pro: can detect previous unseen attacks
- Con: have higher false positives, and hard to train a system for a very dynamic environment.
- Approaches: statistical methods, clustering, outlier detection, SVM

 [Stolfo06]

Anomaly Detection

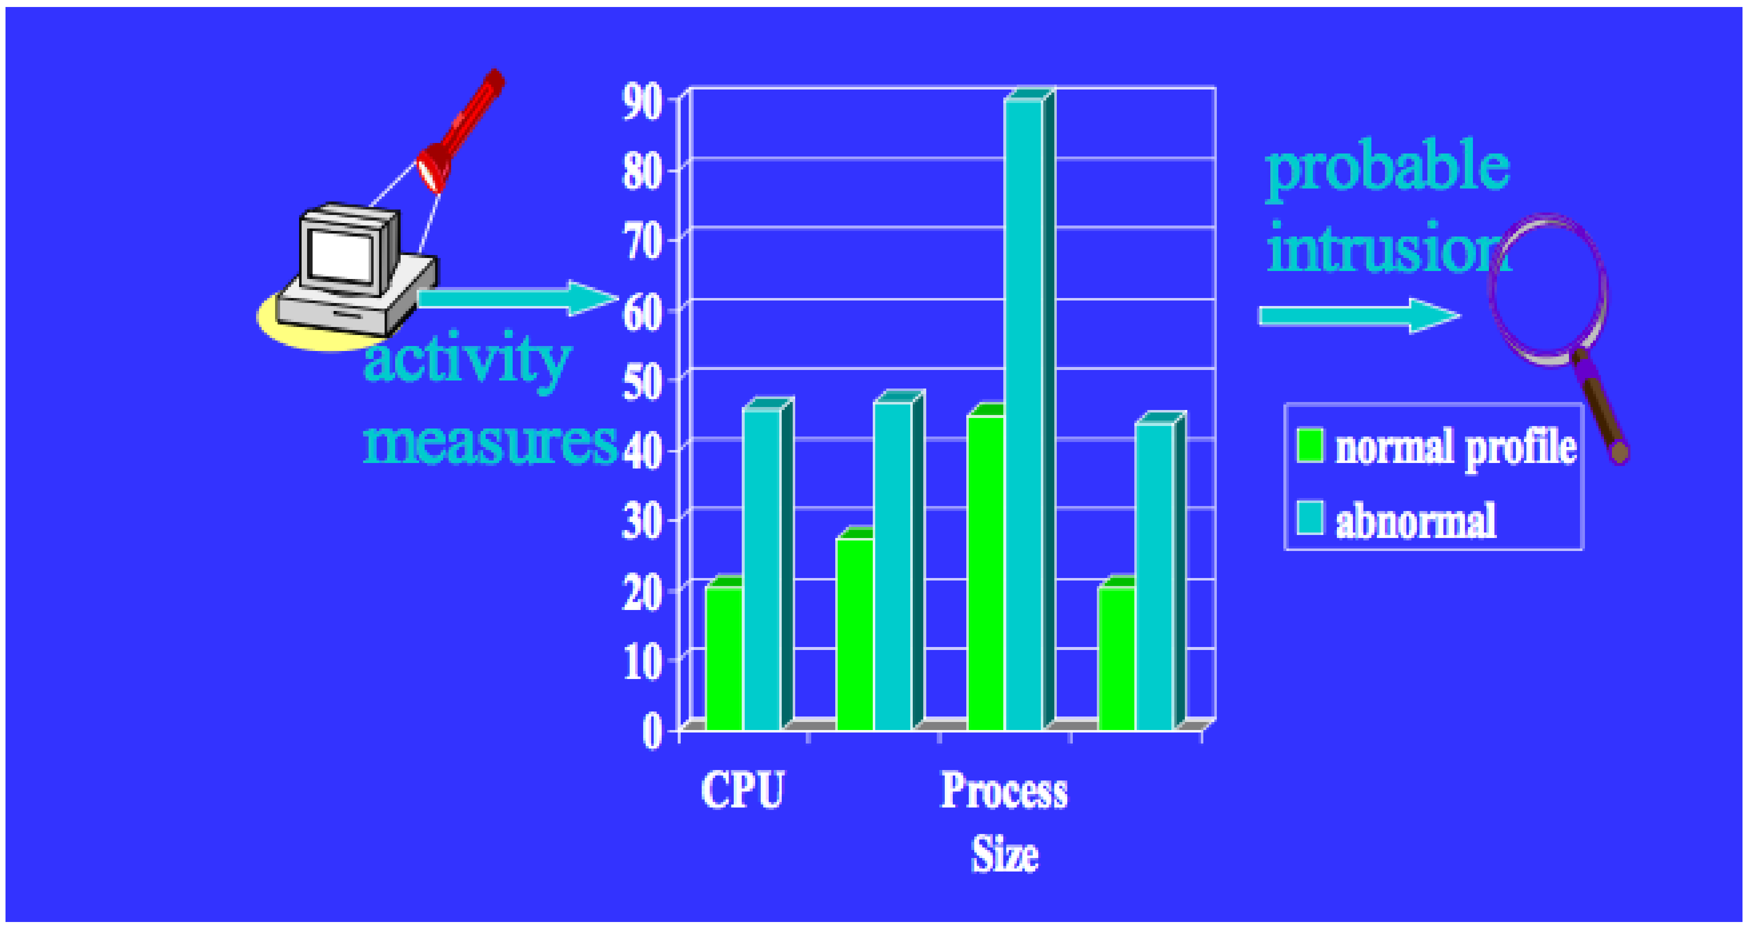

- Relatively high false positive rate - anomalies can just be new normal activities.

Anomaly Detection

- Algorithm
  - Supervised / unsupervised
  - Compute online?
- Data source / feature selection
  - Depends on expert knowledge now
- Cost
  - Computation cost
  - Feature audit and construction cost
  - Damage cost
- Goal: detect attacks accurately and promptly

Data sources

- Single packet
  - src and dst ip, port (most commonly used)
  - All packet header fields
- A sequence of packets
  - Follow the automaton for the protocols (specification-based)
- Reconstructed connections
  - Connection status, frequency (commonly used)
- Traffic flows
  - Volume / velocity.

Learning

- Supervised
  - Statistical tests
    - Build distribution model for normal behavior, then detect low probability events
  - SVM
- Unsupervised
  - Outlier detection
  - Clustering
  - OCSVM

 [Stolfo06]

Examples of IDS

- Misuse detection
  - SNORT : signature based commercial IDS
  - STAT: real time IDS using state transition analysis, attack scenarios specified by STATL. (Higher level signature, abstract from raw packet) [Vigna 03]
  - Bro: real time, events driven, security policy written in a specialized script language. [Paxson 99]
- Anomaly detection
  - MADAM ID
  - ADAM: mining association rule + Bayes classifier
  - Specification-based detection [Sekar 02]

IDS Evaluation

- Accuracy: false positives and false negatives should be minimized.
- Performance: the rate at which audit events are processed.
- Completeness: to detect all attacks.
- Fault tolerance: resistance to attacks.
- Timeliness: time elapsed between intrusion and detection.

Key Performance Metrics

- Algorithm
  - Alarm: A; Intrusion: I
  - Detection (true alarm) rate: $P(A|I)$
    - False negative rate $P(\neg A|I)$
  - False alarm rate: $P(A|\neg I)$
    - True negative rate $P(\neg A|\neg I)$
- Architecture
  - Scalable
  - Resilient to attacks
- Which is a bigger problem?
  - Attacks are fairly rare events
  - IDS often suffer from base-rate fallacy

Base-Rate Fallacy

- 1% of traffic is SYN floods; IDS accuracy is 90%
  - IDS classifies a SYN flood as attack with prob. 90%, classifies a valid connection as attack with prob. 10%
- What is the probability that the connection flagged as a SYN flood by IDS is actually valid?

$$\Pr(\text{valid} \mid \text{alarm}) = \frac{\Pr(\text{alarm} \mid \text{valid}) \cdot \Pr(\text{valid})}{\Pr(\text{alarm})}$$

$$= \frac{\Pr(\text{alarm} \mid \text{valid}) \cdot \Pr(\text{valid})}{\Pr(\text{alarm} \mid \text{valid}) \cdot \Pr(\text{valid}) + \Pr(\text{alarm} \mid \text{SYN flood}) \cdot \Pr(\text{SYN flood})}$$

$$= \frac{0.10 \cdot 0.99}{0.10 \cdot 0.99 + 0.90 \cdot 0.01}$$

= 92% chance raised alarm is false!!!

Problems with (Commercial) IDS

- Cost of update and keeping current is growing
  - Organizations lack internal expertise
- Knowledge based IDS systems suffer from False Negative Problem
  - New augmented IDS with Anomaly Detectors are appearing in the commercial market
- IDS are inherently noisy and chatty and suffer from the False Positive problem
  - Volumes of alerts are crushing
  - Zooming in on most serious threats is hard
- NIDS positioned at the perimeter
  - The most serious/predominant threat is the insider
  - Host and LAN-based IDS now more crucial

 [Stolfo06]

What new solutions are needed for these problems?

- Maintenance problem – Automatic Update
- Data Reduction problem – Human can’t be in the loop
- Insider problem – Look inward, not only outward

Next Generation Detection Systems

- Behavior-based (like credit card fraud):
  - Automated analysis
  - Learn site specific characteristics (e.g., outbound traffic) and prioritize attacks per cost modeling
  - Reduce time to update and deploy
  - Increase analyst/security staff productivity
  - Discover New Attacks
- Offload and load balance detection tasks among separate specialized modules
- Correlation among distributed sites provides new opportunities for
  - Real-time global detection (early warning)
  - Detecting attackers

Acknowledgments/References

- [stolfo06] Part of COMS W4180 - Network Security. Salvatore J. Stolfo, Columbia University.
- [shmatikov] CS 378, Network Security and Privacy, Fall 2007. Vitaly Shmatikov. University of Texas at Austin.
- [memon03] Network Security. Cs 682, Fall 2003. Nasir Memon. Polytechnic University.