Intrusion Detection

version : 2.0

Fakrul (pappu) Alam
fakrul@bdhub.com

Acknowledgement

- Original slides prepared by Steven M. Bellovin

Sometimes, Defenses Fail

- Our defenses aren't perfect
  - Patches weren't applied promptly enough
  - Antivirus signatures not up to date
  - 0-days get through
  - Someone brings in an infected USB drive
  - An insider misbehaves
- Now what?
- Most penetrations are never detected
  - This allows continuing abuse, and helps the attackers spread elsewhere

Unexpected Activity

- There could be an intruder even if you have security practice in place

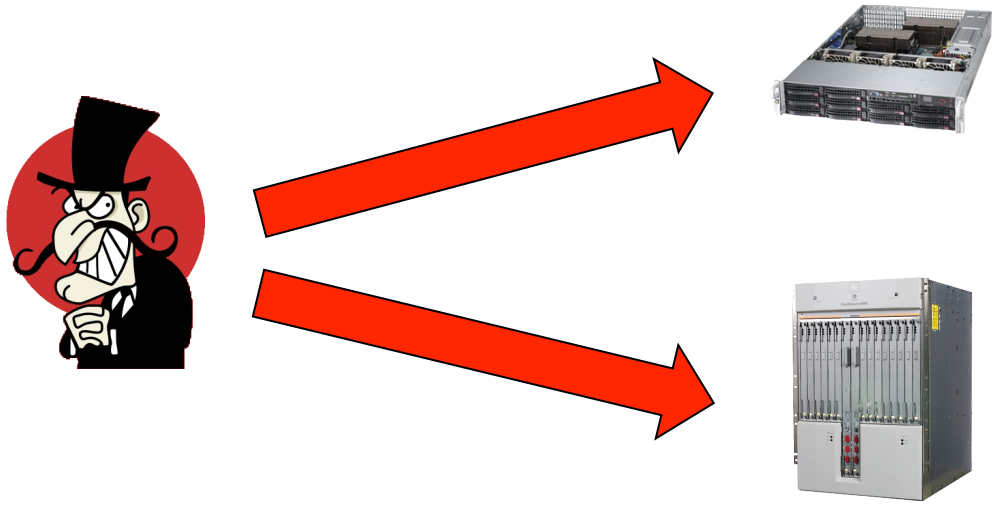

Additional Monitoring

- Activity in your network
- To confirm security

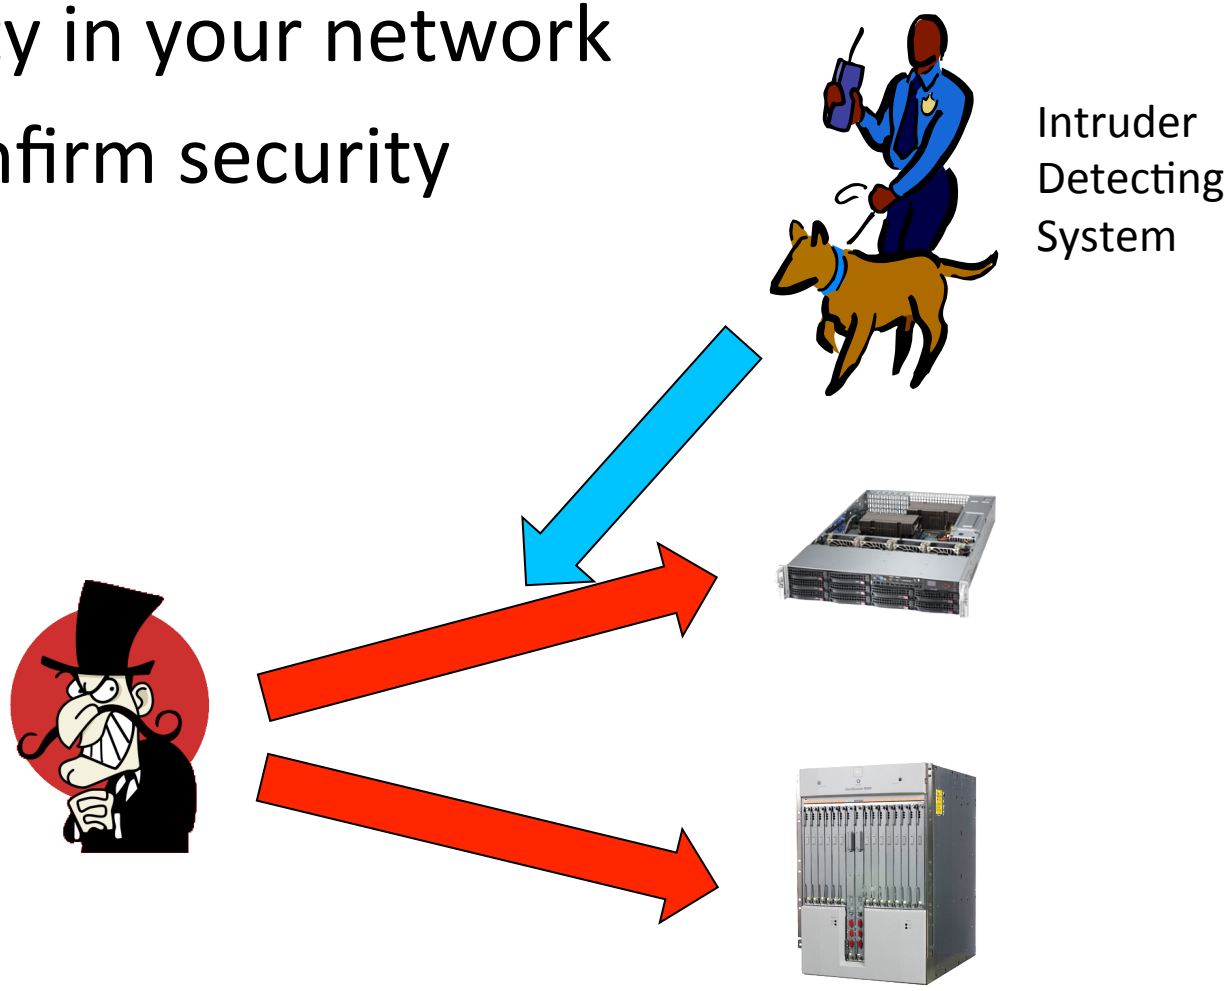

What can IDS realistically do

- Detect successful attacks
- Look for various things that shouldn't be there
  - Infected files
  - Attacks on other machines
  - Packets that shouldn't exist
  - Strange patterns of behavior
- Contain attacks before they spread further
- Clean up penetrated machines—because you'll know they're infected
- Recognition of pattern reflecting known attacks
- Statistical analysis for abnormal activites

What IDS can't do

- Compensate for weak authentication & identification mechanisms
- Investigate attacks without human intervention
- Guess the content of your organization security policy
- Compensate for weakness in networking protocols, for example IP Spoofing

Types of Intrusion Detection System

- Host-base IDSs
- Distributed IDSs
- Network-based IDSs

Types of Intrusion Detection System

- Host-based IDSs
  - Get audit data from host audit trails
  - Detect attacks against a single host
- Distributed IDSs
  - Gather audit data from multiple host and possibly the network that connects the hosts
  - Detect attacks involving multiple hosts
- Network-Based IDSs
  - Use network traffic as the audit data source, relieving the burden on the hosts that usually provide normal computing services
  - Detect attacks from network

Monitoring Point

- More specific rules can be applied for a point close to end nodes
- Generic activities can be found on network

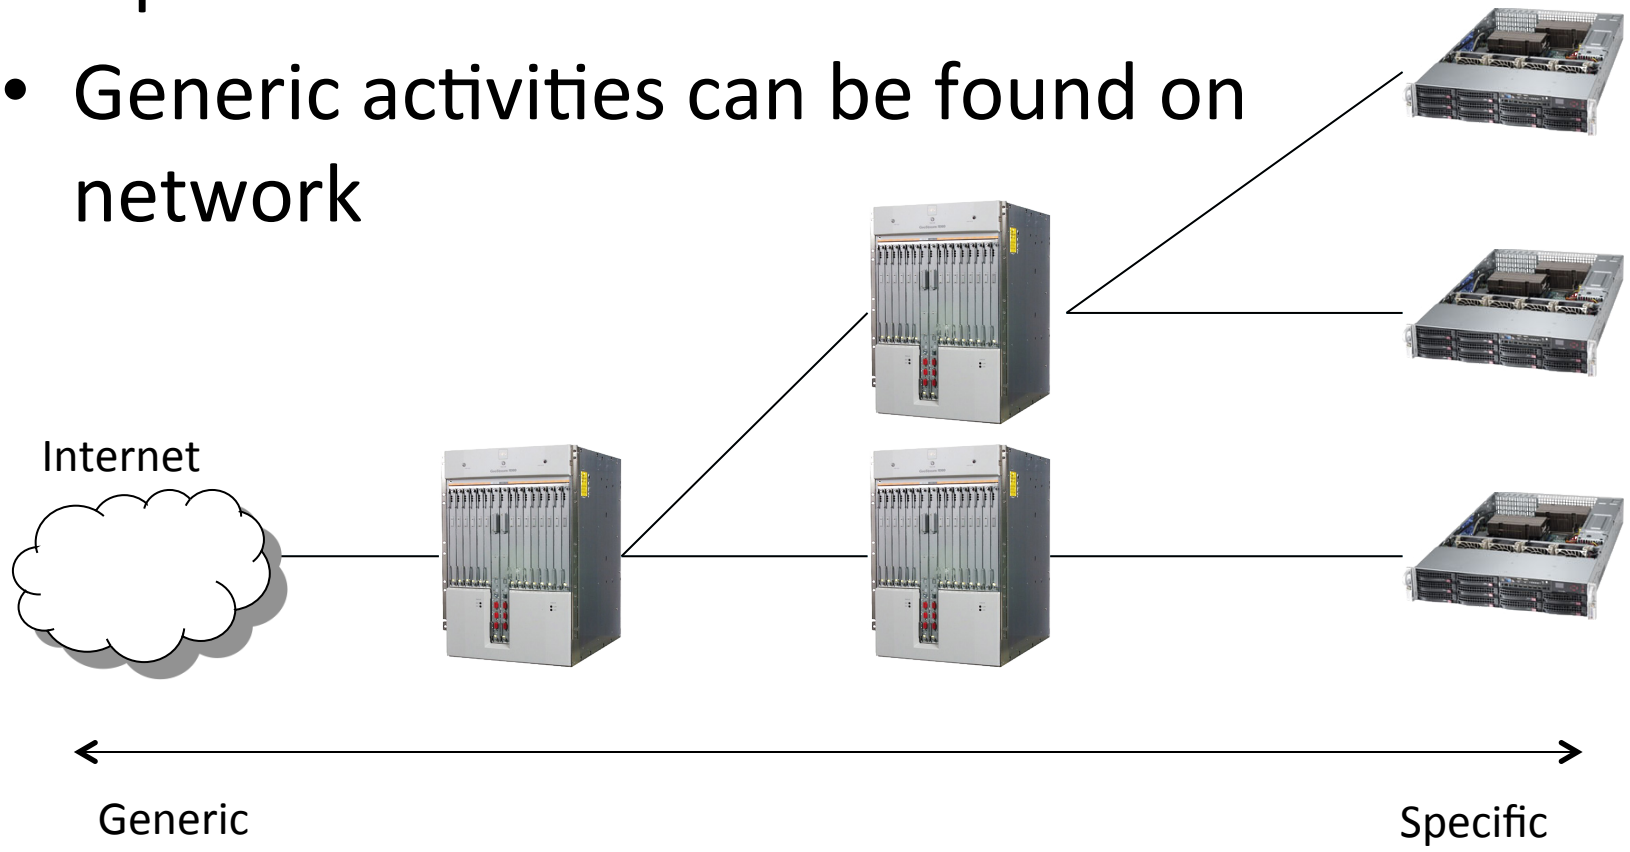

Network and Host IDS

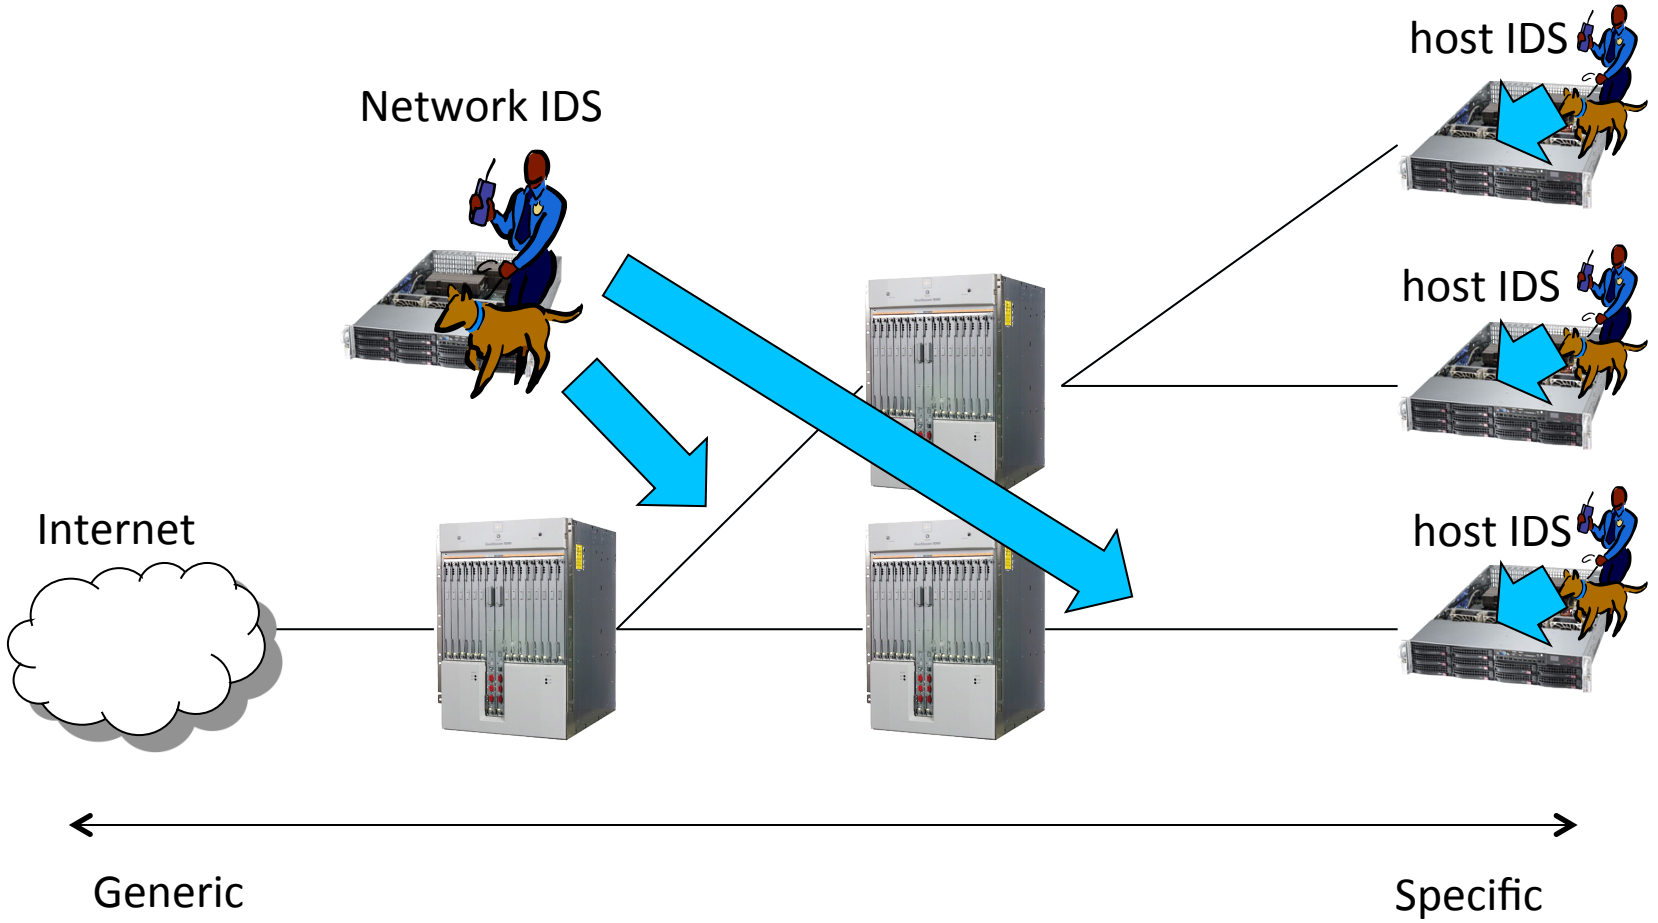

IDS Technology landscape

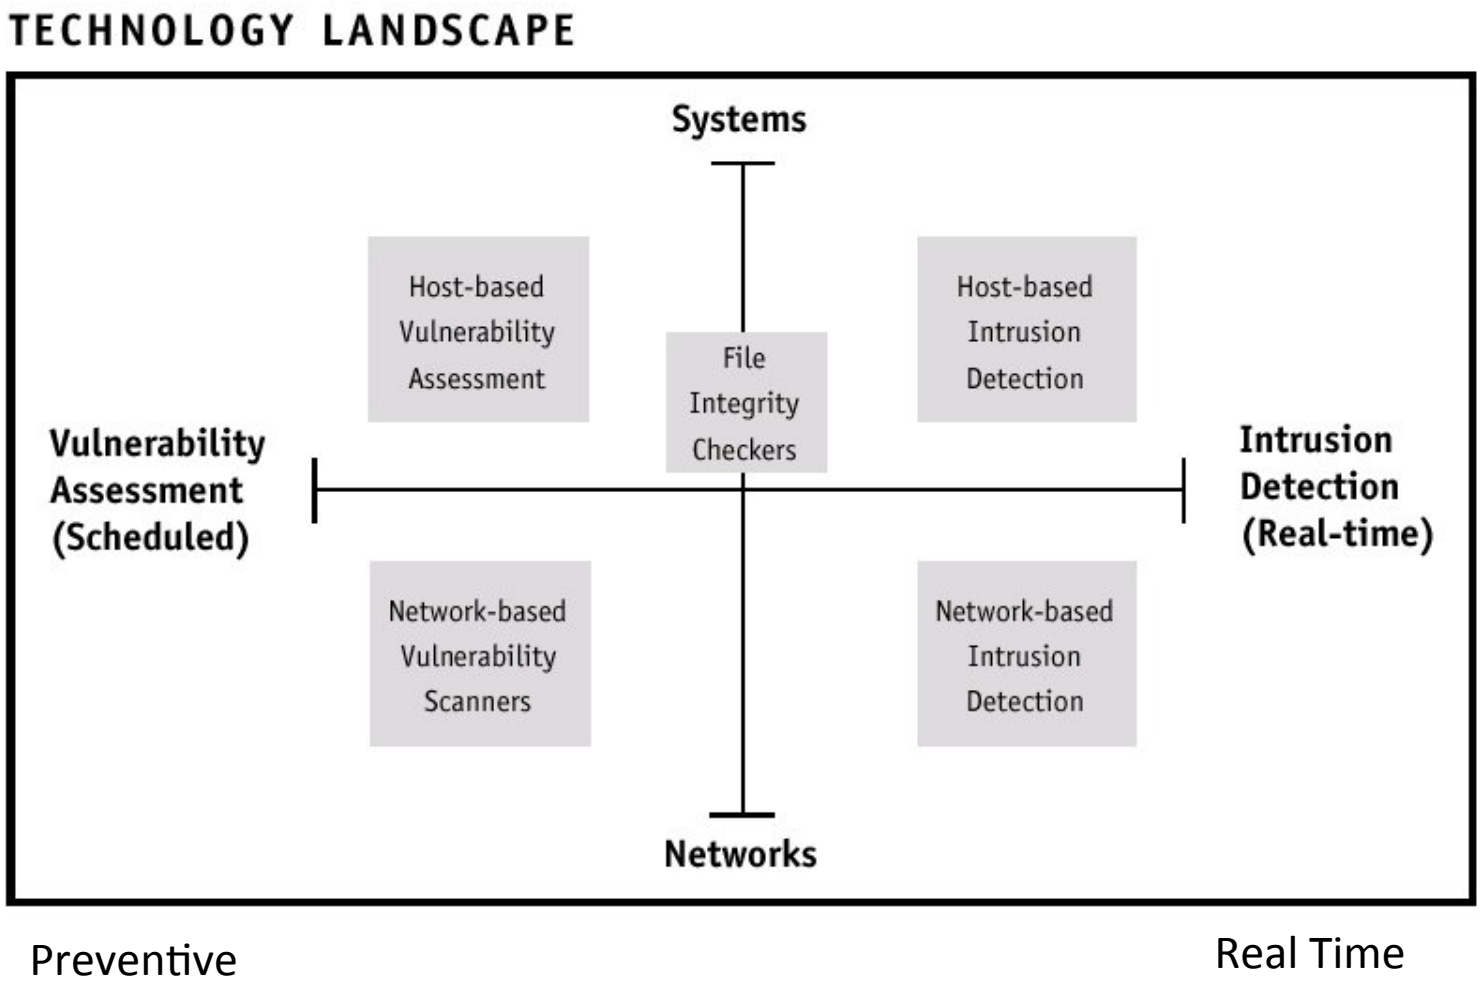

Network based IDS

- Monitors packets by sniffing
  - Using tapping device or port mirroring
  - IP address on the monitoring interface is *not* necessary
    - Difficult to detect by intruders

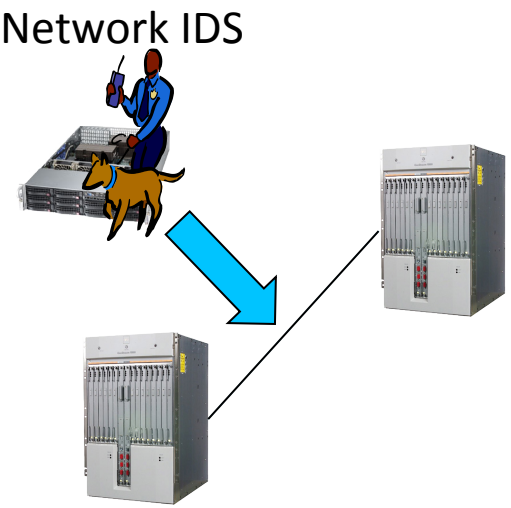

Network based IDS

- Can check if the packet:
  - Attacks on other machines
  - That shouldn't exist
  - Goes to unexpected destination IP/port
  - Has problematic header/payload
  - Or strange patterns of behavior, and so on
- Also may count or record packets

Host based IDS

- Monitors network connections, files, logs and activities on the host, and can check
  - Incoming connections
  - User logins
  - Infected files
  - Attacks on other machines
  - Or strange patterns of behavior, and so on

Network IDS vs. Encryption

- A network sniffer can't read encrypted traffic—but neither can an intruder
- Which is more important, intrusion detection or encryption?
- If you give the IDS all keys, it becomes a very tempting target for an attacker
  - If the protocol uses Diffie-Hellman, giving away the keys doesn't help
- But—encrypted traffic from a host that normally speaks plaintext is suspicious; so is plaintext traffic from a host that should use encryption
  - Use traffic analysis in your IDS…

Alert

- You may receive tons of millions of alerts
  - Depending on your detection rules
  - There are many suspicious activities in the Internet today
- You should notice a critical one at least
  - Detection rule is important!

Alert

- False Positive / Type I Error:
  - is the incorrect rejection of a true null hypothesis
  - is when a system raises an incorrect alert
- False Negative / Type II Error:
  - is the failure to reject a false null hypothesis
  - is when an attack pass undetected

Detection Rules

- Should always fit to your system
  - Installing new services
  - Changing network topology
  - Migrating to another system software
  - Terminating a service
- To get useful alerts

Types of Detection

- Signature Based
  - Match patterns against known attacks
  - Catch the intrusions in terms of the characteristics of known attacks or system vulnerabilities
- Anomaly Based
  - Look for unusual behavior
  - Detect any action that significantly deviates from the normal behavior

Types of Detection

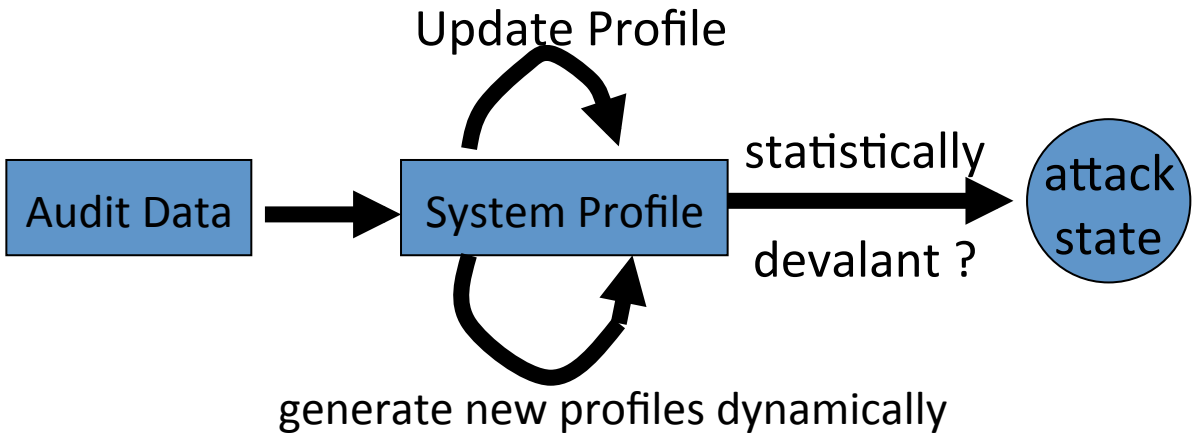

< Signature Based

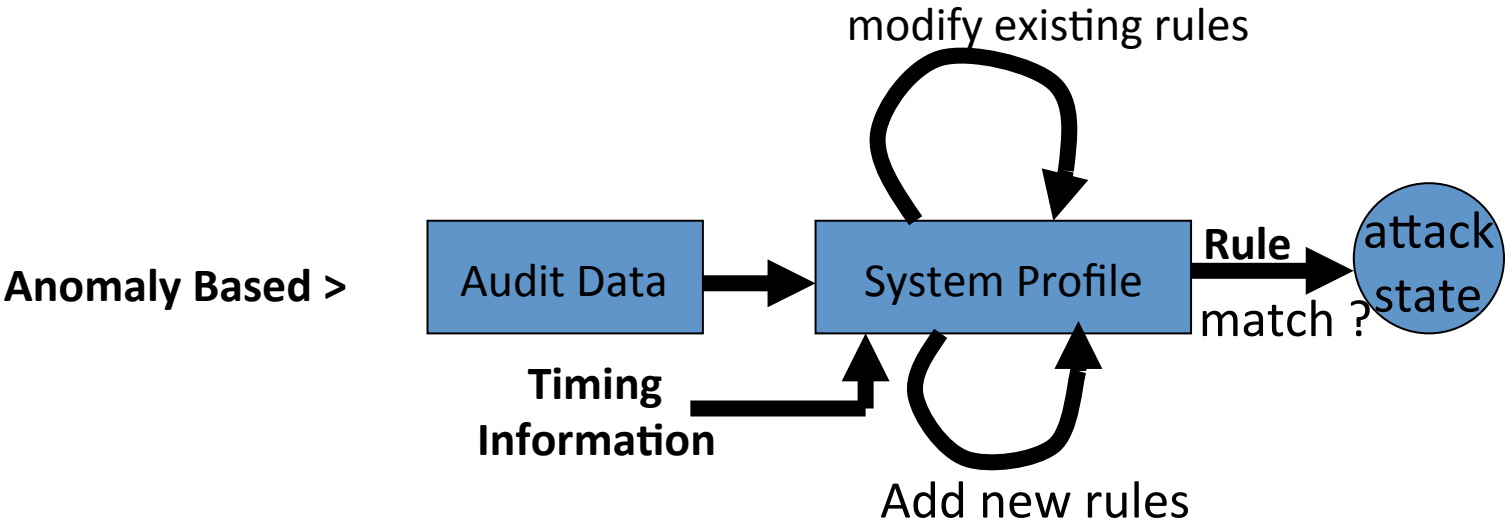

Properties

Signatures	Anomalies
• Require an up-to-date database of attack symptoms • Can do very well against attacks in the database • Completely useless against attacks not listed • Including, but not limited to, 0-days • Most antivirus software is signature-based	• Doesn't need an attack database • Does need to know what normal—uninfected—behavior is like • Can detect 0-day attacks • More susceptible to false positives

Signature Based vs Anomaly Based

	Advantage	Disadvantage
Signature-based	Accurately and generate much fewer false alarm	Cannot detect novel or unknown attacks
Anomaly-based	Is able to detect unknown attacks based on audit	High false-alarm and limited by training data.

Signature Detection

- Look for known-bad types of traffic coming from your customers
  - Example: Connection attempts to your dark space
  - Example: Connections to your email submission server from too many strange places
  - Example: Connections to known botnet controller

Anomaly Detection

- Could monitor upstream links for odd traffic
- However—a lot of misbehavior shows up in traffic metadata (even if you're not the NSA)
- Use Netflow to spot oddities or *changes* in customer behavior
- But—watch out for new applications, or new-to-this-customer applications

Intrusion Detection for ISPs

- Monitor your own network—but that's no different than any other enterprise
- Monitor your customers
  - Good: you can help them by detecting problems
  - Good: you can prevent them from clogging your infrastructure
  - Bad: it can be privacy-invasive