Security - Intrusion Detection Systems

Note Title 4/27/2011

Announcements

- Lab due tomorrow

- Reading Assignment:
  - read first paper (overview of forensics)
  - pick one other paper

due next Thursday
if time, in class review
submit Monday or Tuesday (BEFORE the final)

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 a system in an unauthorized manner.

Two kinds:

- Host-based
- Network-based

Goal of IDS

Detect most intrusions while keeping false alarm rate low!

(Either extreme is bad.)

Any IDS has 3 components:

- Sensors : gather data
- Analyzer : computes reports, gives alerts...
- User interface : tailor rules

Goals of IDS

- Catch intrusions before any damage occurs
- Distinguish "legitimate" behavior from attacks

  • minimize false positives and false negatives

Generally, this is a balancing act.

Host based IDS

Add an extra layer of security to vulnerable systems.

Primary purpose is extra logging + sending alerts.

Nice feature: Will detect external and internal intrusions.

Auditing & host-based IDS

- Often can use built-in auditing on user activity.
  No extra resources, but might not be in correct format or have relevant data included.

- Some detection specific audit software can also be used.
  Downside: Slower machines

Two Types of host-based

① Anomaly detection: collect user data over a period of time.
Then apply one of two strategies:
- threshold detection
- profile based detection

② Signature detection

Define a set of rules or attack patterns to decide if a series of actions constitutes an attack.

Anomaly Detection

- Theshold detection generally raises many false positives or false negatives, depending on the # of occurances allowed.

Fairly inaccurate + crude, but can be used in conjunction with other methods.

Anomaly detection (cont.)

- Profile-based detection requires analysis of audit records before system is running.

  Ex: # of logins by a user, number of outgoing messages, time between logins or other actions, # of resources used.

Many different tests can then be used to define "normal".

- mean & std deviation
- Markov process model (probabilistic)
- time series
- multivariate

Signature Detection

Apply a set of rules to decide if some activity is suspicious.
Overall weakness: lack of flexibility

Types:

- Rule-based anomaly detection:
  use statistical data from previous audits to get rules, then compare to current behavior
  (again, don't need any knowledge of system weaknesses)

Signature Detection (cont)

- Rule based penetration identification:
  quantify known vulnerabilities & weaknesses

Example:
- no user-id concurrently on same machine
- no user should read another user's directory
- no user should copy a system program
- no user should directly open disk devices
- app specific

Distributed host-based ID

Host-based IDS can become stronger if systems on same LAN work together.

Major Issues:

- different audit record formats
- a node must serve as a collection point
  - so data will be transmitted across the network
  - and collection point is vulnerable

Network Intrusion Detection (NIDS)

- NIDS monitor traffic at selected points on the network
- Examine packets and aggregate data to detect suspicious patterns

Generally consists of:
- Sensors
- Servers
- management consoles

Types of sensors

- In line sensors: traffic passes directly through.
  (generally) combined with a firewall or switch)
  Advantage: no extra hardware, just software
  Disadvantage: packet delay

- Passive sensors: monitor a copy of the network traffic with no protocol interaction
  (the sensor connects to fiber optic cable by a direct physical tap)

Sensor deployment

- Sensor just "inside" the firewall and/or just "outside" firewall
  (different advantages either way)
- Sensors to protect major internal networks

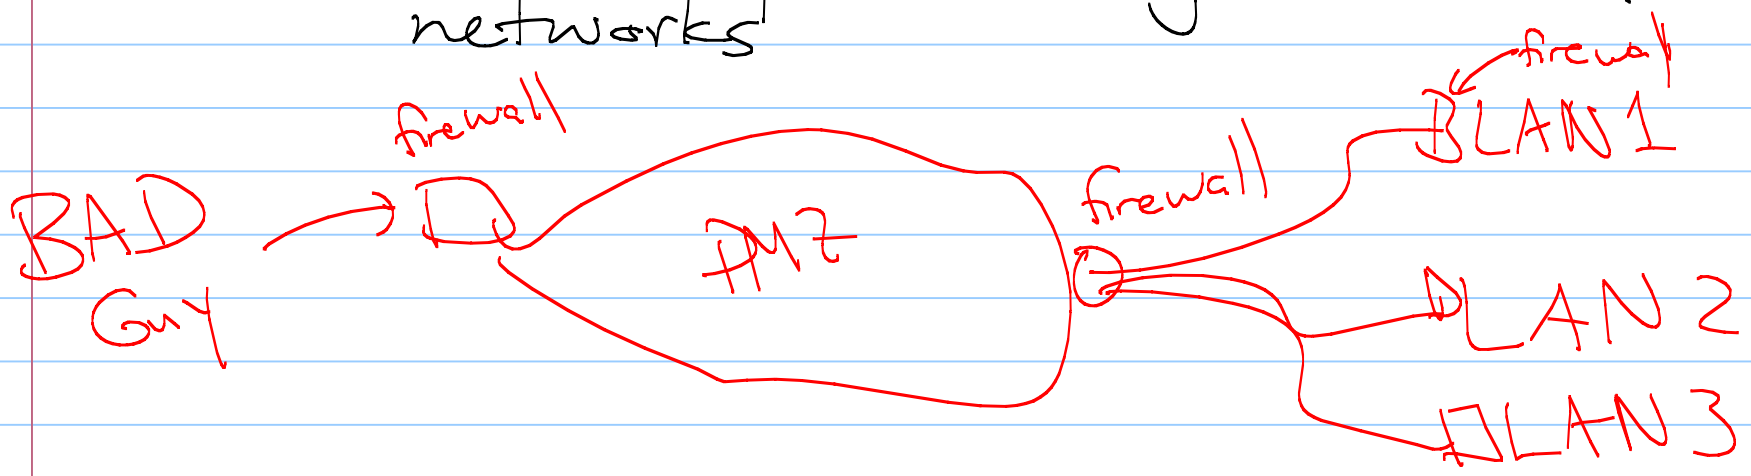

Signature Detection

- Monitor various protocols
  Ex: DHCP, DNS, finger, FTP, HTTP, IRC, POP, ...
  Look for known attack patterns or unusual behavior for these protocols.

- Transport layer attacks: TCP/UDP, etc.
  Ex: Scans of vulnerable ports, SYN floods, etc.

- Network Layer attacks:
  Ex: spoofed IPs, illegal IP headers

- Unexpected or forbidden applications

Anomaly Detection

- DOS attacks
- Scanning of network
- Worms propogated - eats up unusual amounts of bandwidth

NIDS reporting

As with host-based, the NIDS server must log relevant info

- time
- session ID
- event type
- network, transport, & app layer protocol
- IP addresses & ports
- Payload data

Honey pots

Decoy systems designed to lure attackers away from critical systems.

Also allows admin to collect data.

Any attack is made to seem successful, but system has no real valuable data + is loaded with sensors.

Note: If in internal network, this is a security risk!
Also greatly complicated network if internal.
(Caution: Also legal issues...)

Example IDS: Snort

- open source
- easily deployed on most nodes
- efficient
- easily configured

Can perform:
- real time packet capture
- protocol analysis
- content searching and matching

Snort architecture

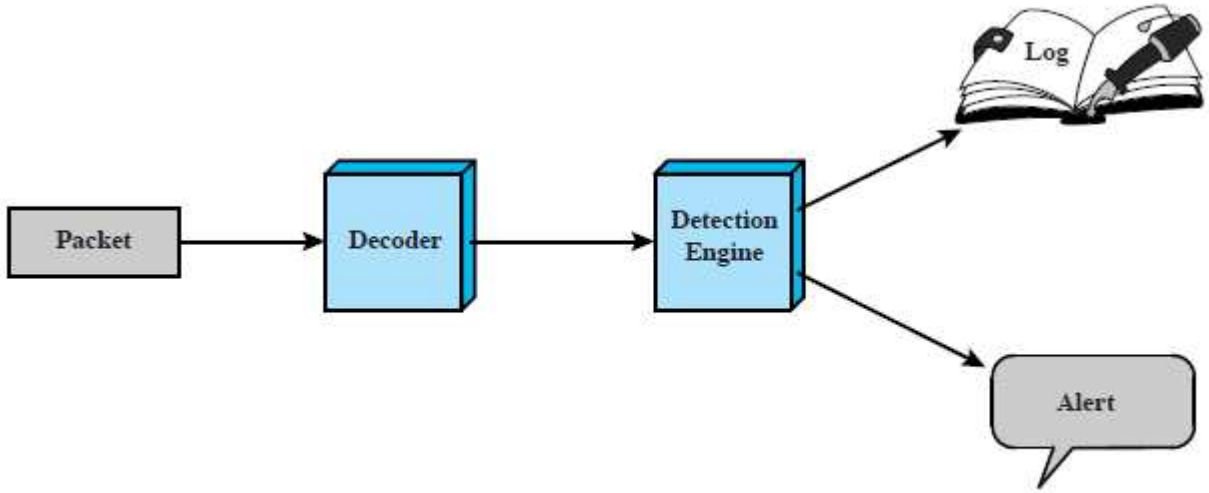

Figure 6.9 Snort Architecture

(Works for either passive or inline sensors.)

Snort Rules

Table 6.4 Snort Rule Actions

Action	Description
alert	Generate an alert using the selected alert method, and then log the packet.
log	Log the packet.
pass	Ignore the packet.
activate	Alert and then turn on another dynamic rule.
dynamic	Remain idle until activated by an activate rule , then act as a log rule.
drop	Make iptables drop the packet and log the packet.
reject	Make iptables drop the packet, log it, and then send a TCP reset if the protocol is TCP or an ICMP port unreachable message if the protocol is UDP.
sdrop	Make iptables drop the packet but does not log it.

Example:

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

meta-data
msg Defines the message to be sent when a packet generates an event.
reference Defines a link to an external attack identification system, which provides additional information.
classtype Indicates what type of attack the packet attempted.

non-payload
ttl Check the IP time-to-live value. This option was intended for use in the detection of traceroute attempts.
flags Test the TCP flags for specified settings.

(more options possible!)