CS 4910: Intro to Computer Security

Network Security V: Intrusion Detection Systems (IDS)

Instructor: Xi Tan

Next

- Network Security
  - Network Firewalls
  - Intrusion Detection Systems

Intrusion Detection

- Intrusion
  - Actions aimed at compromising the security of the target (confidentiality, integrity, availability of computing/networking resources)
- Intrusion detection
  - The identification of intrusions and report of intrusion activities
- Intrusion prevention
  - The process of both detecting intrusion activities and managing automatic responsive actions throughout the network

Intrusion Detection System (IDS) Components

- The **IDS manager** compiles data from the IDS sensors to determine if an intrusion has occurred.
- This determination is based on a set of **site policies,** which are rules and conditions that define probable intrusions.
- If an IDS manager detects an intrusion, then it sounds an **alarm**.

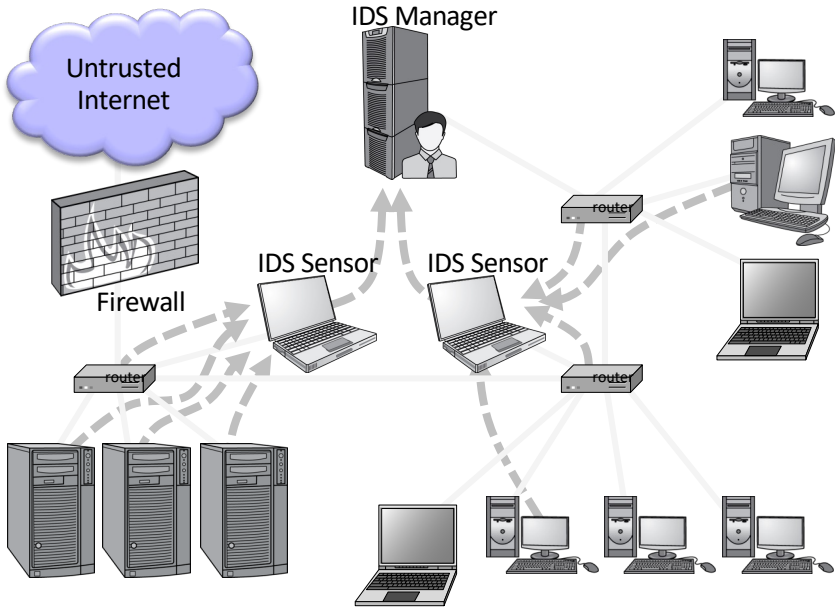

Intrusion Detection Systems

- Who is likely intruder?
  - May be outsider who got thru firewall
  - May be evil insider
- What do intruders do?
  - Launch well-known attacks
  - Launch variations on well-known attacks
  - Launch new or un-known attacks
  - Use a system to attack other systems
  - Etc.

Intrusions

- An IDS is designed to detect automated attacks and threats, including the following:
  - **Port scans:** information gathering intended to determine which ports on a host are open for TCP connections
  - **Denial-of-service attacks:** network attacks meant to overwhelm a host and shut out legitimate accesses
  - **Malware attacks:** replicating malicious software attacks, such as Trojan horses, computer worms, viruses, etc.
  - **ARP spoofing:** an attempt to redirect IP traffic in a local-area network
  - **DNS cache poisoning:** a pharming attack directed at changing a host's DNS cache to create a falsified domain-name/IP-address association

Intrusion Detection Systems

- Intrusion detection is not perfect, two types of errors are
  - false positives: legitimate behavior of authorized users is classified as an intrusion
  - false negatives: an intrusion is not recognized as suspicious activity
- False negatives result in higher losses than false positives
  - thus a higher rate of false positives is normally tolerated than the rate of false negatives
  - if an error rate is very high, warnings tend to get ignored
  - proper tuning of the system is important
- The earlier intrusion is detected, the better
  - it is easier to recover while the damage is small

Intrusion Detection Systems (IDS)

- Intrusion detection system (IDS) is a security service that monitors and analyzes system events
- IDS classification
  - host-based IDS
    - monitors events and characteristics of a single host for suspicious activity
  - network-based IDS
    - monitors data on the network for traces of suspicious activity
    - often a single monitor scans data sent to/from many machines on the network
  - hybrid IDS
    - combines information gathered from hosts and network

Host-based IDS

- A host-based IDS runs on a single host
  - it is best positioned to evaluate the state of the machine
- It can monitor events and activity such as
  - login and session activity
    - frequency and location, time since last login, failed login attempts
    - events of security importance can include break-in into a dead account, logins from unusual locations or unusual hours, password guessing, etc.
  - program execution activity
    - monitored activity can include execution denials, resource utilization and execution frequency

Host-based IDS

- Monitored events and activity
  - file access activity
    - record frequency of different types of file access, denial of access
    - look for abnormal usage patterns, suspicious activity such as copying system programs or opening devices directly
  - some combination of the above
    - e.g., users who login after hours often access the same files they used earlier
- If a host-based IDS runs on each host, information from different machines can be collected and managed at a central facility
  - the central manager receives aggregate information and distributes updates to all machines running the IDS

Network-based IDS

- A network-based IDS monitors traffic corresponding to many machines on a network
  - often such a monitor is passive
    - NIDS receives a copy of the traffic
  - a **firewall**, on the other hand, performs active filtering
    - all traffic goes directly through it
  - active filtering adds overhead and normally needs to be minimized

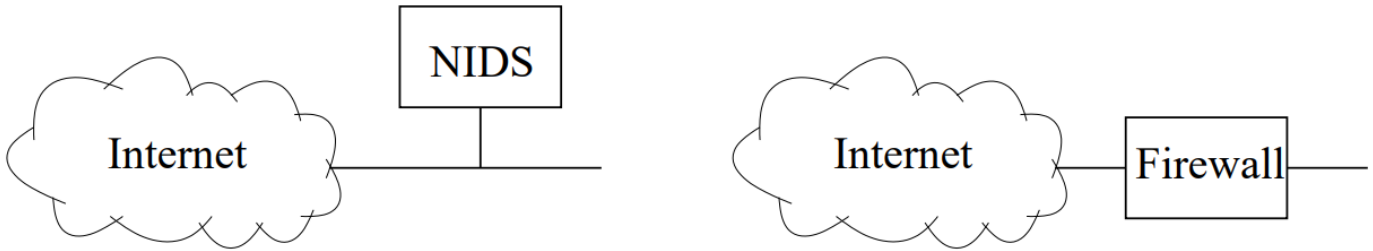

Network-based IDS

- Where NIDS is positioned matters

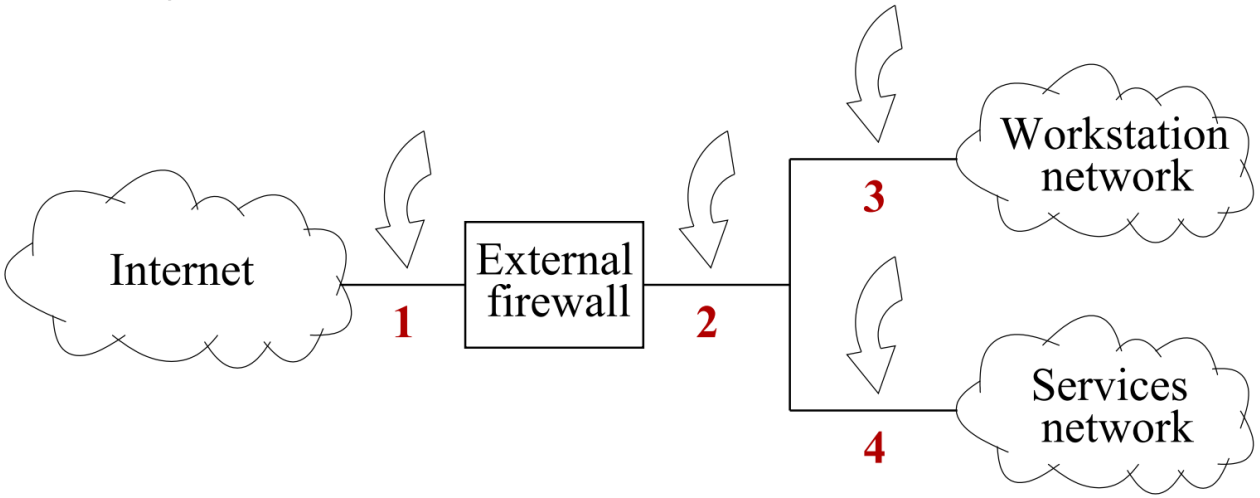

  - point 1: complete picture of traffic, lots of data
  - point 2: can recognize problems with firewall, see outgoing attacks
  - points 3 and 4: increased visibility of attacks on the local network, can see internal attacks

Network-based IDS

- A NIDS is often stateful and performs deep packet inspection
  - full stream reassembly
  - analysis at network, transport and/or application layers
    - **network layer**: IP, ICMP protocols, illegal header values, spoofed addresses
    - **transport layer**: analysis of TCP and UDP headers, detection of unusual packet fragmentation, floods, scans
    - **application layer**: understanding of DHCP, DNS, HTTP, Network File System (NSF), remote login and many other protocols; detection of buffer overflow attacks, malware propagation, etc.
  - detection of DoS attacks, scanning, malware (worms)

Network-based IDS

- Example systems
  - Snort
    - can be host-based or network-based
    - can monitor traffic inline (supports intrusion prevention) or passively
    - intrusion detection/prevention is rule-based
  - Bro
    - provides passive monitoring of network traffic
    - suitable for high-speed high-volume detection
  - commercial appliances

Intrusion Detection Systems

- IDSs can be classified based on how they recognize suspicious activity
  - misuse detection (signature or heuristic based)
    - define what constitutes an intrusion attempt through a set of rules
    - e.g., specific patterns in network traffic, a combination of events
    - can detect only known/encoded intrusion attempts
  - anomaly detection
    - train the system on clean data to understand behavior of legitimate users
    - use it to monitor real data and detect anomalous behavior
    - advantages: more flexible, can detect unknown misuses
    - disadvantages: higher error rate, difficult to tune

Signature or Heuristic Detection

- Signature approaches
  - Match a large collection of known patterns of malicious data against data stored on a system or in transit over a network
  - The signatures need to be large enough to minimize the false alarm rate while still detecting a sufficiently large fraction of malicious data
  - Widely used in anti-virus products, network traffic scanning proxies, and NIDS
- Rule-based heuristic identification
  - Involves the use of rules for identifying known penetrations or penetrations that would exploit known weaknesses
  - Rules that identify suspicious behavior can also be defined, even when the behavior is within the bounds of established patterns of usage
  - Typically rules used are specific
  - SNORT is an example of a rule-based NIDS

Signature Detection Example

- Failed login attempts may indicate password cracking attack
- IDS could use the rule "N failed login attempts in M seconds" as **signature**
- If N or more failed login attempts in M seconds, IDS warns of attack
- Note that the warning is specific
  - Admin knows what attack is suspected
  - Admin can verify attack (or false alarm)

Signature Detection

- But if attacker knows the signature, he can try $N-1$ logins every M seconds!
- In this case, signature detection slows the attacker, but might not stop him

Signature Detection

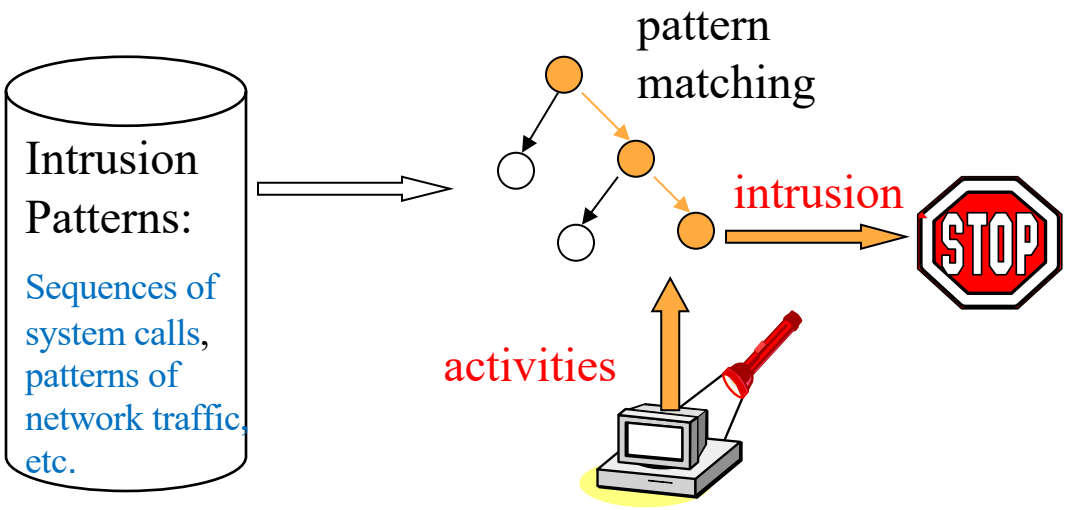

Example: ***if*** (traffic contains "x90+de[^\r\n]{30}") ***then*** "attack detected"

Advantage: Mostly accurate. But problems?

Can't detect new attacks

Signature Detection

- Advantages of signature detection
  - Simple
  - Detect known attacks
  - Know which attack at time of detection
  - Efficient (if reasonable number of signatures)
- Disadvantages of signature detection
  - Signature files must be kept up to date
  - Number of signatures may become large
  - Can only detect known attacks
  - Variation on known attack may not be detected

Anomaly Detection

- A variety of classification approaches are used:
  - Statistical
    - Analysis of the observed behavior using univariate, multivariate, or time-series models of observed metrics
  - Knowledge based
    - Approaches use an expert system that classifies observed behavior according to a set of rules that model legitimate behavior
  - Machine-learning
    - Approaches automatically determine a suitable classification model from the training data using data mining techniques

Anomaly Detection

- Anomaly detection systems look for unusual or abnormal behaviors
- There are (at least) two challenges
  - What is normal for this system?
  - How “far” from normal is abnormal?
- Statistics is obviously required here!
  - The **mean** defines normal
  - The **variance** indicates how far abnormal lives from normal

Anomaly Detection

- Advantages
  - Chance of detecting unknown attacks
  - May be more efficient (since no signatures)
- Disadvantages
  - Reliability is unclear
    - High false positive/false negative
  - Anomaly detection indicates something unusual, but lack of specific info on possible attack!
  - Today, cannot be used alone
  - Must be used with a signature detection system

Honeypots

- A further component of intrusion detection technology is the honeypot
- Decoy systems designed to:
  - 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.
- Systems are filled with fabricated information that a legitimate user of the system wouldn't access
- Resources that have no production value
  - Therefore incoming communication is most likely a probe, scan, or attack
  - Initiated outbound communication suggests that the system has probably been compromised

Honeypots

Example of Honeypot Deployment

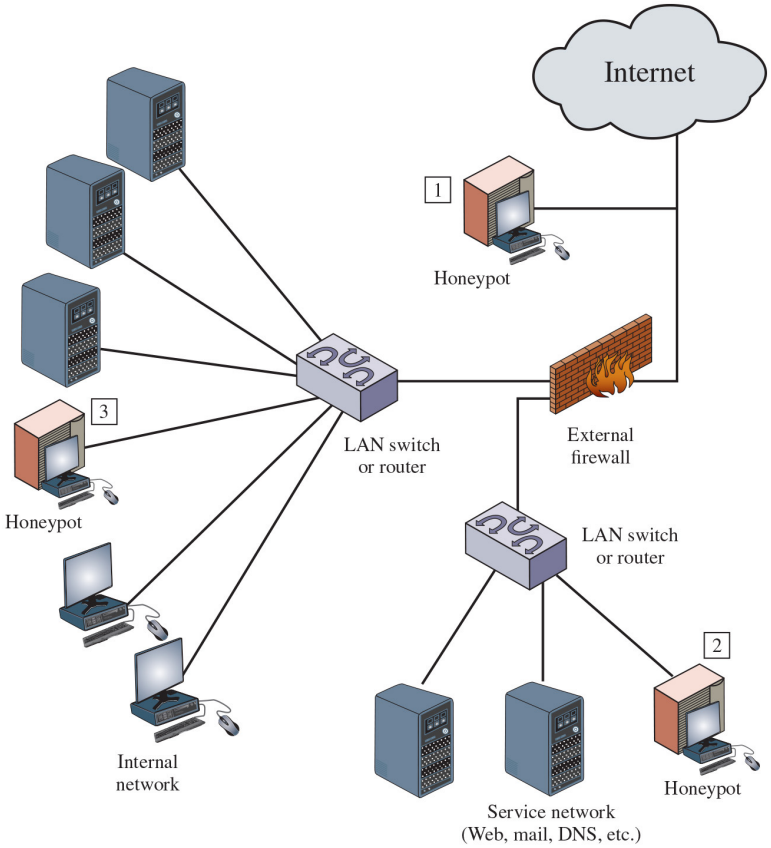

Summary

- Firewalls: first line of defense
- Intrusion detection systems
  - Based on deploy position:
    - host-based: best positioned to detect attacks on a machine
    - network-based: monitors traffic of the entire network
    - hybrid
  - Based on detection method:
    - signature-based: effective, but don’t recognize new attacks
    - anomaly-based: can find novel attacks, but often result in many false positives
- Effort must be applied to protect the IDS itself from attacks