CSE 410/565 Computer Security
Spring 2022

Lecture 19: Intrusion Detection

Department of Computer Science and Engineering
University at Buffalo

Lecture Outline

- Intruders
- Intrusion detection
  - host-based
  - network-based
  - hybrid
  - attacks on intrusion detection systems

Intruders

- Different types of intruders
  - hackers
    - people who break into computers to gain status within hacking community
    - even benign intruders consume resources and must be stopped
  - criminal organizations
    - more determined attackers with a target goal (e.g., to gain access to sensitive or financial data)
    - often act quickly and with fewer mistakes
    - obscure use of stolen financial data to complicate investigation

Intruders

- Types of intruders
  - insiders
    - employees who misuse their privileges with or without malice
    - example: access to IRS data by employees, employees who take databases upon leaving an organization
- The goal is to defend against all of the above
- Often a strong barrier is built at the network perimeter
  - firewalls, packet filtering, stricter policies, intrusion detection
  - special precautions must be made to defend against internal threats

Intruders

- Often the following defenses are used to counter insider intrusion
  - enforce least privilege, permit access only to resources needed for the job
  - use authentication to access sensitive information
  - log accesses and other relevant information
  - upon job termination promptly revoke all privileges
  - when an employee with access to sensitive information leaves, can be useful to store information about their privileges and data for future references in case an accident happens

Intrusion Detection

- 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

Intrusion Detection Systems

- IDSs can be classified based on how they recognize suspicious activity
  - misuse detection (signature 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

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

- What we often want from an IDS
  - continuous operation
  - minimum human intervention
  - small overhead, ability to scale
  - ability to adapt to changes in user behavior and system characteristics over time
  - resistance to compromise (ability to monitor itself)
  - ability to be reconfigured on the fly, without restarting
- Often all of the above are extremely difficult to achieve simultaneously
  - e.g., ability to adapt in anomaly-based detection often has a higher human supervision cost

Host-Based Intrusion Detection

- 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 Intrusion Detection

- 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 Intrusion Detection

- 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

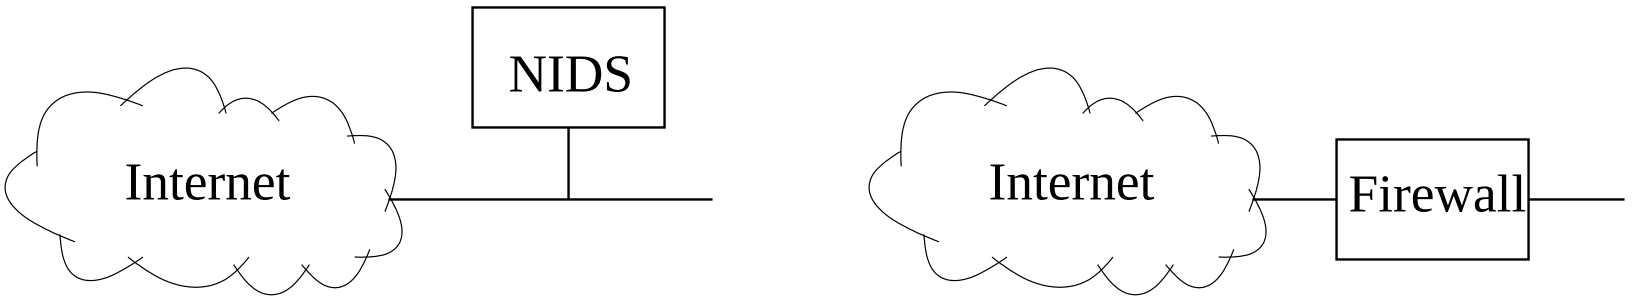

Network-Based Intrusion Detection

- Where NIDS is positioned matters

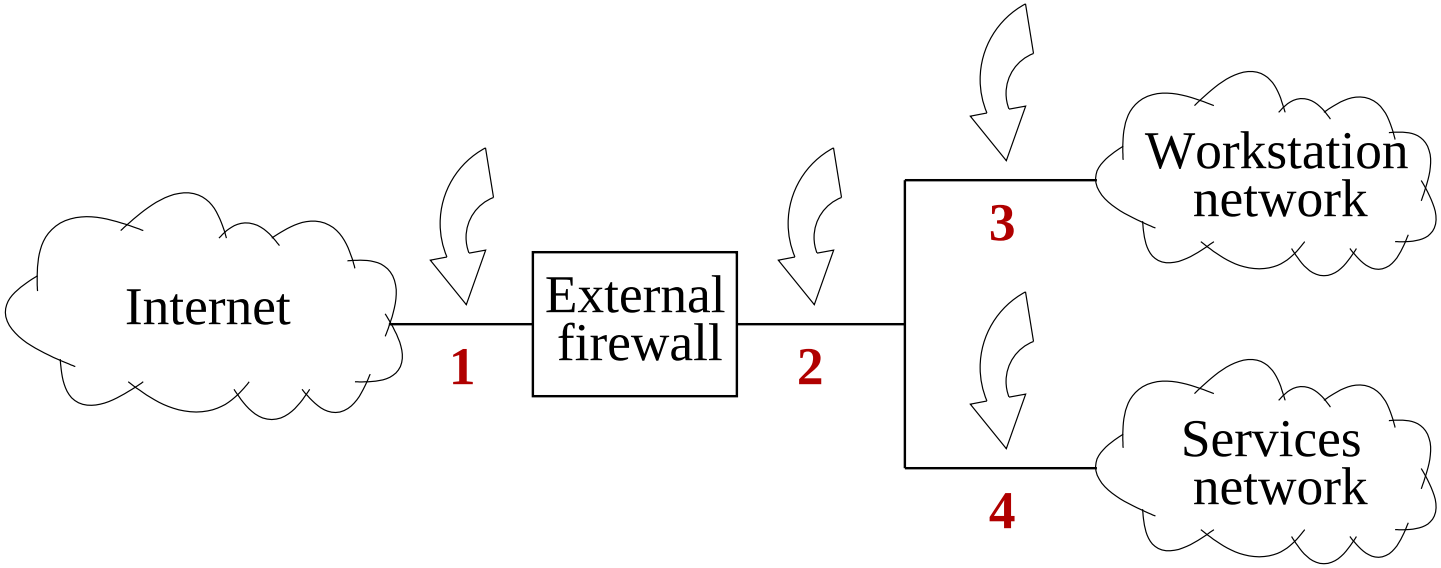

  - point 1: complete picture of traffic, lots of data
  - point 2: can recognize problems with firewall, see outgoing attacks
  - points 3 and 4: increased visibility of attacks on the local network, can see internal attacks

Network-Based Intrusion Detection

- 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 Intrusion Detection

- 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

Network-Based Intrusion Detection

- Challenges in running NIDS
  - necessity to handle large volume of traffic
  - ability to correctly maintain the state of each machine on the network
  - ability to withstand attacks on NIDS itself
- Attacks on NIDS
  - algorithmic complexity attacks
  - evasion attacks
  - stealthy port scanning

Attacks on NIDS

- Algorithmic complexity attacks
  - DoS attacks are already serious for denying service, but can be more severe as a component of an attack
  - DoS attack on IDS enables other attacks to remain undetected
- Example: complexity attack on hash table
  - on average, a hash table has $O(n)$ overhead to insert n elements
  - in the worst case, it may have $O(n^2)$ overhead to insert n elements
  - Perl implementation for 90 thousand inserts
    - random: < 2 sec
    - worst case: > 6500 sec

Attacks on NIDS

- Complexity attack against Bro
  - Bro used simple XOR to "hash" values for hash table
    - easy to find collisions
  - for example, Bro port scanning detector keeps a hash table of destination IP addresses
    - keep the list of destination IP addresses for each (source IP, destination port)
  - using source IP spoofing one can exploit this structure to perform DoS attack

Performance	Attack	Random
Total CPU time	44.5 min	0.85 min
Hash table time	43.78 min	0.02 min

Attacks on NIDS

- NIDS evasion
  - attack might rely on the fact that NIDS is not the target host and might have incomplete picture
  - complete fragment reassembly is necessary to detect certain attacks
  - NIDS only has partial knowledge of what the host sees
    - Time-To-Live (TTL) expires before reaching the host
    - packets that exceed the maximum transmission unit (MTU) are dropped
  - ambiguities in TCP/IP (e.g., overlapping IP and TCP fragments)
    - different OSs implement the standard differently

Attacks on NIDS

- Small TTL attack

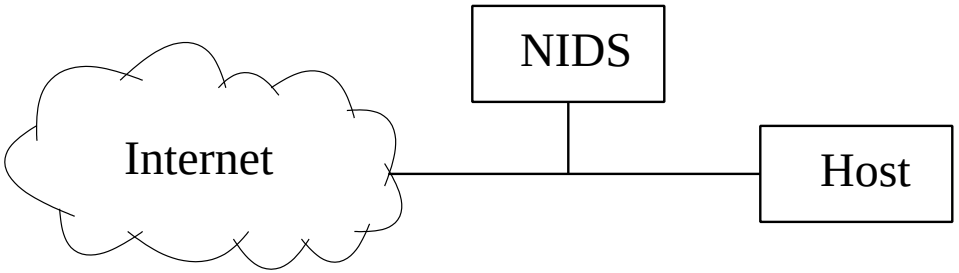

Attacker's data stream:

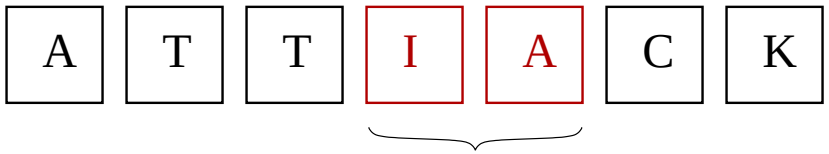

same TCP sequence number
"I" has short TTL

NIDS sees:

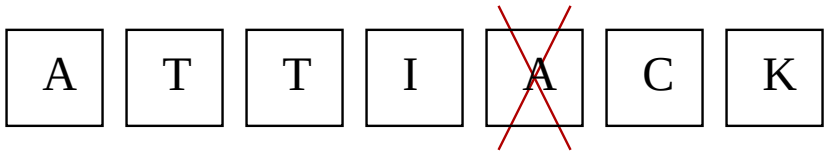

End host sees:

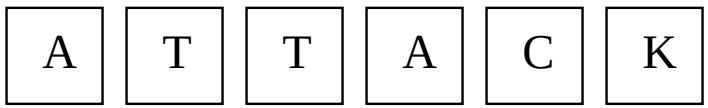

Attacks on NIDS

- Fragment overlap attack

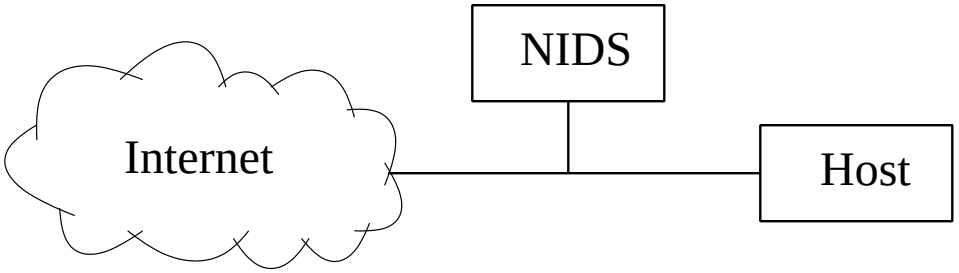

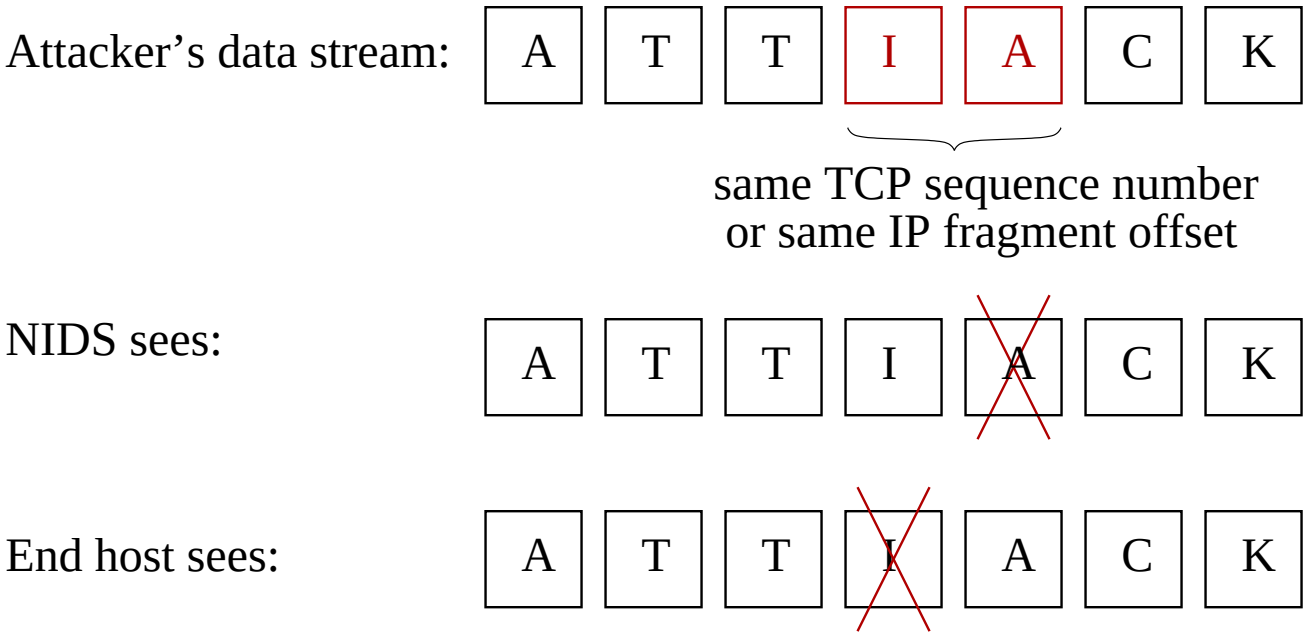

Attacks on NIDS

- How do we defend against such attacks?
  - solution: introduce traffic normalizer to avoid ambiguities

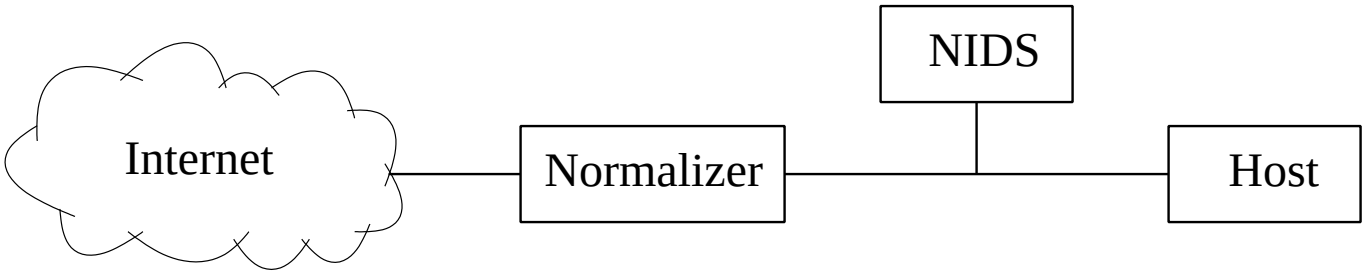

  - drop overlapping IP/TCP fragments
  - increase TTL in packets with low TTL
- But IDS evasion can still be possible
  - different interpretation of strings of characters at higher levels
  - e.g., A T T I DEL A C K

Intrusion Detection

- For more reliable detection, NIDSs can be placed at different points inside the network
  - one monitor for the entire network
  - a monitor inside each subnet
  - this results in a distributed IDS
- Hybrid IDSs can be most effective
  - run IDS both on hosts and network
  - combine the data for improved decision making

Conclusions

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