Intrusion Detection and Prevention

Outlines:

- Intrusion
- Types of Intrusion
- Intrusion Detection Models
- Intrusion Prevention Models

By: Arash Habibi Lashkari
July - 2010

Definition

- Intrusion Detection
  - Intrusion detection is a technique of detecting unauthorized access to a computer system or a computer network.
  - An intrusion into a system is an attempt by an outsider to the system to illegally gain access to the system. Intrusion prevention, on the other hand, is the art of preventing an unauthorized access of a system's resources.
  - The two processes are related in a sense that while intrusion detection passively detects system intrusions, intrusion prevention actively filters network traffic to prevent intrusion attempts.

- Intrusion
  - An *intrusion* is a deliberate unauthorized attempt, successful or not, to break into, access, manipulate, or misuse some valuable property and where the misuse may result into or render the property unreliable or unusable.
  - The person who intrudes is an *intruder*.

- There are six types of intrusions:
  - Attempted break-ins, which are detected by a typical behavior profiles or violations of security constraints. An intrusion detection system for this type is called anomaly-based IDS.
  - Masquerade attacks, which are detected by a typical behavior profiles or violations of security constraints. These intrusions are also detected using anomaly-based IDS.
  - Penetrations of the security control system, which are detected by monitoring for specific patterns of activity.
  - Leakage, which is detected by atypical use of system resources.
  - Denial of service, which is detected by atypical use of system resources.
  - Malicious use, which is detected by atypical behavior profiles, violations of security constraints, or use of special privileges.

Intrusion Detection Systems (IDSs)

- An *intrusion detection system (IDS)* is a system used to detect unauthorized intrusions into computer systems and networks. Intrusion detection as a technology is not new, it has been used for generations to defend valuable resources.
- These are three models of intrusion detection mechanisms: *anomaly-based* detection, *signature-based* detection, and *hybrid* detection.

- Anomaly Detection –
  - Anomaly based systems are “learning” systems in a sense that they work by continuously creating “norms” of activities. These norms are then later used to detect anomalies that might indicate an intrusion.
  - Anomaly detection compares observed activity against expected normal usage profiles “leaned”. The profiles may be developed for users, groups of users, applications, or system resource usage.

- Misuse Detection -
  - The misuse detection concept assumes that each intrusive activity is represent able by a unique pattern or a *signature* so that slight variations of the same activity produce a new signature and therefore can also be detected.
  - Misuse detection systems, are therefore, commonly known as *signature systems.* They work by looking for a specific signature on a system. Identification engines perform well by monitoring these patterns of known misuse of system resources.
- Hybrid Detection -
  - Because of the difficulties with both the anomaly-based and signature-based detections, a hybrid model is being developed. Much research is now focusing on this hybrid model.

Types of Intrusion Detection Systems

- Intrusion detection systems are classified based on their monitoring scope. There are: network-based intrusion detection and host-based detections.
- Network-Based Intrusion Detection Systems (NIDSs)
  - NIDSs have the whole network as the monitoring scope. They monitor the traffic on the network to detect intrusions. They are responsible for detecting anomalous, inappropriate, or other data that may be considered unauthorized and harmful occurring on a network. There are striking differences between NIDS and firewalls.

- Host-Based Intrusion Detection Systems (HIDS)
  - Recent studies have shown that the problem of organization information misuse is not confirmed only to the “bad” outsiders but the problem is more rampart within organizations. To tackle this problem, security experts have turned to inspection of systems within an organization network. This local inspection of systems is called *host-based intrusion detection systems* (HIDS).
  - Host-based intrusion detection is the technique of detecting malicious activities on a single computer.
  - A host-based intrusion detection system, is therefore, deployed on a single target computer and it uses software that monitors operating system specific logs including system, event, and security logs on Windows systems and syslog in Unix environments to monitor sudden changes in these logs.
  - When a change is detected in any of these files, the HIDS compares the new log entry with its configured attack signatures to see if there is a match. If a match is detected then this signals the presence of an illegitimate activity.

- The Hybrid Intrusion Detection System
  - Both NIDS and HIDS are each patrolling its own area of the network for unwanted and illegal network traffic. They, however, complement each other. Both bring to the security of the network their own strengths and weaknesses that nicely complement and augment the security of the network.
  - Hybrids are new and need a great deal of support to gain on their two cousins. However, their success will depend to a great extent on how well the interface receives and distributes the incidents and integrates the reporting structure between the different types of sensors in the HIDS and NIDS spheres. Also the interface should be able to smartly and intelligently gather and report data from the network or systems being monitored.

The Changing Nature of IDS Tools

- Recent studies have shown that the majority of system intrusion actually come from insiders. So newer IDS tools are focusing on this issue and are being built to counter systems intrusion, new attack patterns are being developed to take this human behavior unpredictability into account.
- To keep abreast of all these changes, ID systems are changing constantly.
- The primary focus of ID systems has been on a network as a unit where they collect network packet data by watching network packet traffic and then analyzing it based on network protocol patterns “norms,” “normal” network traffic signatures, and network traffic anomalies built in the rule base. But since networks are getting larger, traffic heavier, and local networks more splintered, it is becoming more and more difficult for the ID system to “see” all traffic on a switched network such as an Ethernet. This is leading to new designs of IDS.

Other Types of Intrusion Detection Systems

- Although NIDS and HIDS and their hybrids are the most widely used tools in network intrusion detection, there are others that are less used but more targeting and, therefore, more specialized.
- Because many of these tools are so specialized, many are still not considered as being intrusion detection systems but rather intrusion detection add-ons or tools.

- System Integrity Verifiers (SIVs)
  - SIVs monitor critical files in a system, such as system files, to find whether an intruder has changed them. They can also detect other system components' data; for example, they detect when a normal user somehow acquires root/administrator level privileges. In addition, they also monitor system registries in order to find well known signatures.
- Log File Monitors (LFM)
  - LFMs first create a record of log files generated by network services. Then they monitor this record, just like NIDS, looking for system trends, tendencies, and patterns in the log files that would suggest an intruder is attacking.

- Honeypots
  - A *honeypot* is a system designed to look like something that an intruder can hack. They are built for many purposes but the overriding one is to deceive attackers and learn about their tools and methods.
  - Honeypots are also add-on/tools that are not strictly sniffer-based intrusion detection systems like HIDS and NIDS. However, they are good deception systems that protect the network in much the same way as HIDS and NIDS.
  - Since the goal for a honeypot is to deceive intruders and learn from them without compromising the security of the network, then it is important to find a strategic place for the honeypot. In the DMZ for those networks with DMZs or behind the network firewall if the private network does not have a DMZ.

Response to System Intrusion

- A good intrusion detection system alert should produce a corresponding response.
- A good response must consist of pre-planned defensive measures that include an incident response team and ways to collect IDS logs for future use and for evidence when needed.

- Incident Response Team
  - An *incident response team* (IRT) is a primary and centralized group of dedicated people charged with the responsibility of being the first contact team whenever an incidence occurs. An IRT must have the following responsibilities:
    - keeping up-to-date with the latest threats and incidents,
    - being the main point of contact for incident reporting,
    - notifying others whenever an incident occurs,
    - assessing the damage and impact of every incident,
    - finding out how to avoid exploitation of the same vulnerability, and
    - recovering from the incident.

- IDS Logs as Evidence
  - IDS logs can be kept as a way to protect the organization in case of legal proceedings. If sensors to monitor the internal network are to be deployed, verify that there is a published policy explicitly stating that use of the network is consent to monitoring.

Challenges to Intrusion Detection Systems

- There is an exciting future and challenges for IDS as the marriage between it and artificial intelligence takes hold
- Although there are also IDS challenges in many areas including in the deployment of IDSes in switched environments.
- Deploying IDS in Switched Environments
  - Network-based IDS sensors must be deployed in areas where they can "see" network traffic packets. However, in switched networks this is not possible because by their very nature, sensors in switched networks are shielded from most of the network traffic. Sensors are allowed to "see" traffic only from specified components of the network.
  - One way to handle this situation has traditionally been to attach a network sensor to a mirror port on the switch. But port mirroring, in addition to putting an overhead on the port, gets unworkable when there is an increase in traffic on that port because overloading one port with traffic from other ports may cause the port to bulk and miss some traffic.

- Other issues still limiting IDS technology are:
  - False alarms. Though the tools have come a long way, and are slowly gaining acceptance as they gain widespread use, they still produce a significant number of both false positives and negatives,
  - The technology is not yet ready to handle a large-scale attack. Because of its very nature it has to literally scan every packet, every contact point, and every traffic pattern in the network. For larger networks and in a large-scale attack, it is not possible that the technology can be relied on to keep working with acceptable quality and grace.
  - Unless there is a breakthrough today, the technology in its current state cannot handle very fast and large quantities of traffic efficiently.
  - Probably the biggest challenge is the IDS's perceived and sometimes exaggerated capabilities. The technology, while good, is not the cure of all computer network ills that it is pumped up to be. It is just like any other good security tool.

Implementing an Intrusion Detection System

- An effective IDS does not stand alone. It must be supported by a number of other systems. Among the things to consider, in addition to the IDS, in setting up a good IDS for the company network are:
  - *Operating Systems.* A good operating system that has logging and auditing features. Most of the modern operating systems including Windows, Unix, and other variants of Unix have these features. These features can be used to monitor security critical resources.
  - *Services*. All applications on servers such as Web servers, e-mail servers, and databases should include logging/auditing features as well.
  - *Firewalls*. A good firewall should have some network intrusion detection capabilities.
  - *Network management platform*. Whenever network management services such as OpenView are used, make sure that they do have tools to help in setting up alerts on suspicious activity.

Intrusion Prevention Systems (IPSs)

- Although IDS have been one of the cornerstones of network security, they have covered only one component of the total network security picture since they have been and they are a passive component which only detects and reports without preventing.
- A promising new model of intrusion is developing and picking up momentum. It is the *intrusion prevention system* (IPS) which, is to prevent attacks.
- Like their counterparts the IDS, IPS fall into two categories: network-based and host-based.

Network-Based Intrusion Prevention Systems (NIPS)

- Because NIDSs are passively detecting intrusions into the network without preventing them from entering the networks, many organizations in recent times have been bundling up IDS and firewalls to create a model that can detect and then prevent.
- The bundle works as follows.
  - The IDS fronts the network with a firewall behind it. On the detection of an attack, the IDS then goes into the prevention mode by altering the firewall access control rules on the firewall. The action may result in the attack being blocked based on all the access control regimes administered by the firewall.
  - The IDS can also affect prevention through the TCP resets; TCP utilizes the RST (reset) bit in the TCP header for resetting a TCP connection, usually sent as a response request to a non-existent connection. But this kind of bundling is both expensive and complex, especially to an untrained security team. It suffers from *latency* – the time it takes for the IDS to either modify the firewall rules or issue a TCP reset command. This period of time is critical in the success of an attack.

Host-Based Intrusion Prevention Systems (HIPSs)

- Most HIPSs work by *sand-boxing*, a process of restricting the definition of acceptable behavior rules used on HIPSs. HIPS prevention occurs at the agent residing at the host. The agent intercept system calls or system messages by utilizing dynamic linked libraries (dll) substitution.
- The substitution is accomplished by injecting existing system dlls with vendor stub dlls that perform the interception.

Questions