INTRUSION DETECTION SYSTEMS (IDSS)

Information security intrusion detection systems (IDSs) were first commercially available in the late 1990s. An IDS works like a burglar alarm in that it detects a violation of its configuration (analogous to an opened or broken window) and activates an alarm. This alarm can be audible and/or visual (producing noise and lights, respectively), or it can be silent (taking the form of an e-mail message or pager alert). With almost all IDSs, system administrators can choose the configuration of the various alerts and the associated alarm levels for each type of alert. Many IDSs enable administrators to configure the systems to notify them directly of trouble via e-mail or pagers. The systems can also be configured-again like a burglar alarm-to notify an external security service organization of a "break-in." The configurations that enable IDSs to provide such customized levels of detection and response are quite complex.

IDS Terminology

In order to understand IDS operational behavior, you must first become familiar with some terminology that is unique to the field of IDSs. The following is a compilation of relevant IDS-related terms and definitions that were drawn from the marketing literature of a well-known information security company, TruSecure, but are representative across the industry:

Alert or Alarm: An indication that a system has just been attacked and/or continues to be under

attack. IDSs create alerts of alarms to notify administrators that an attack is or was or occurring and may have been successful. Alerts and alarms may take
the form of audible signals, e-mail messages, pager notifications, pop-up windows, or log entries (these are merely written, i.e., they do not involve taking any action).

False Attack Stimulus: An event that triggers alarms and causes a false positive when no actual attacks are in progress. Testing scenarios that evaluate the configuration of IDSs may use false attack stimuli to determine if the IDSs can distinguish between these stimuli and real attacks.

False Negative: The failure of an IDS system to react to an actual attack event. Of all failures, this is the most grievous, for the very purpose of an IDS is to detect attacks.

False Positive: An alarm or alert that indicates that an attack is in progress or that an attack has successfully occurred when in fact there was no such attack. A false positive alert can sometimes be produced when an IDS mistakes normal system operations/activity for an attack. False positives tend to make users insensitive to alarms, and will reduce their quickness and degree of reaction to actual intrusion events through the process of desensitization to alarms and alerts. This can make users less inclined, and therefore slow, to react when an actual intrusion occurs.

Noise: The ongoing activity from alarm events that are accurate and noteworthy but not necessarily significant as potentially successful attacks. Unsuccessful attacks are the most common source of noise in IDSs, and some of these may not even be attacks at all, but rather employees or the other users of the local network simply experimenting with scanning and enumeration tools without any intent to do harm. The issue faced regarding noise is that most of the intrusion events detected are not malicious and have no significant chance of causing a loss.

Site Policy: The rules and configuration guidelines governing the implementation and operation of IDSs within the organization.

Site Policy Awareness: An IDSs ability to dynamically modify its site policies in reaction or response to environmental activity. A so-called Smart ID can adapt its reaction activities based on both guidance learned over the time from the administrator and circumstances present in the local environment. Using a device of this nature, the IDs administrator acquires logs of events that fit a specific profile instead of being alerted about minor changes, such as when a file is changed or a user login fails. Another advantage of using a Smart IDS is that the IDS knows when it does not need to alert the administrator-this would be the case when an attack using a known and documented exploit is made against systems that the IDS knows are patched against

that specific kind of attack. When the IDS can accept multiple response profiles based on changing attack scenarios and environmental values, it can be made much more useful.

True Attack Stimulus: An event that triggers alarms and causes an IDS to react as if a real attack is in progress. The event may be an actual attack, in which an attacker is at work on a system compromise attempt, or it may be a drill, in which security personnel are using hacker tools to conduct tests of a network segment.

Confidence Value: A value associated with an IDS's ability to detect and identify an attack correctly. The confidence value an organization places in the IDS is based on experience and past performance measurements. The confidence value, which is a type of fuzzy logic, provides an additional piece of information to assist the administrator in determining whether an attack alert is indicating that an actual attack in progress, or whether the IDS is reacting to false attack stimuli and creating a false positive. For example, if a system deemed capable of reporting a denial-of-service attack with 90% confidence sends an alert, there is a high probability that an actual attack is occurring.

Alarm Filtering: The process of classifying the attack alerts that an IDS produces in order to distinguish/sort false positives from actual attacks more efficiently. Once an \IDS has been installed and configured, the administrator can set up alarm filtering by first running the system for a while to track what types of false positives it generates and then adjusting the classification of certain alarms. For example, the administrator may set the IDS to discard certain alarms that he or she knows are produced by false attack stimuli or normal network operations. Alarm filters are similar to packet filters in that they can filter items by their source or destination IP addresses, but they have the additional capability of being able to filter by operating systems, confidence values, alarm type, or alarm severity.

Alarm Clustering : A consolidation of almost identical alarms into a single higher-level alarm. This consolidation will reduce the total number of alarms generated, thereby reducing administrative overhead, and will also indicate a relationship between the individual alarm elements.

Alarm Compaction: Alarm clustering that is based on frequency, similarity in attack signature, similarity in attack target, or other similarities. Like the previous form of alarm clustering, this will reduce the total number of alarms generated, thereby reducing administrative overhead, and will also indicate a relationship between the individual alarm elements when they have specific similar attributes.