Bright Hub – How IDS works

www.brighthub.com/computing/smb-security/articles/65416.aspx

Trying to find out how an intrusion detection systems works? It would take a book to completely explain this; nevertheless, this article will provide an introductory answer to the question.

- **Introduction**

  There are many different kinds of intrusion detection systems (IDS). There are physical IDS (i.e. for buildings or rooms), and there are digital type of IDS used to provide an alarm system for computer networks or computers themselves. This article will concern itself with computer network or computer intrusion detection systems.

  In this article we will introduce the reader to the concepts of how intrusion detection systems work. To do this, I will cover the following:

  - Answer the question: What is an Intrusion Detection System?
  - Enumerate and explain the various IDS components
  - Explain how the components come together as a system

  When done, this article as a whole shall answer the question "How does an intrusion detection systems work?"

- **What is an Intrusion Detection System?**

  Before we tackle the question of how intrusion detection systems work, we need to understand what it is. The easiest way to explain this is by making an analogy with physical alarm systems-e.g. the type you might have in your house.

  A house alarm system, when armed, will sound when one or more of its entry sensors is tripped. There are different kinds of sensors: contact or motion. So if you forget to lock your door or someone breaks open your door while your alarm is armed, your home alarm will sound.

  Computer network or computer intrusion detection systems basically provide the same function. In this case, your computer or your network is the one using an alarm system. So when a computer hacker is trying to find ways into your network or your computer, it can provide you with an alert (via e-mail or some means through its own management software) that something or someone is attempting to break in or has already broken in.

  The intruder from this perspective can be a live malicious hacker, or it could be an Internet worm or malware intended to exploit your computer's vulnerabilities in order to spread itself.

  There are two basic types of IDS: host-based and network-based. Host-based IDS are designed to provide IDS functions for the protection of the host where it resides. In most

cases, a host-based IDS also runs as a service and/or and application within the host it is protecting. Network-based IDS are designed to watch network traffic and alert of any malicious activities it sees on the network port it is monitoring. Network-based IDS typically run on a separate host or dedicated appliance designed to perform this function.

Combining these two solutions to protect your computers can be part of a layered network and computer protection plan.

- ## Components of Intrusion Detection Systems

  Network or computer intrusion detection systems all have these basic components:

  - Sensor
    - Activity or packet capture engine
    - Behavioral or signature detection engine
  - Backend
    - Event recording database
    - Alerting engine
  - Frontend
    - User interface
    - Command & control

  These components could be physically separate or could all reside on the same host. There could be one or more sensors, and in larger setups, one or more backends.

  ### Sensor

  The sensor is the primary component for detecting hacking activities on the computer or on the network. It has a packet capture and activity capture engine to help it get access to activities efficiently and quickly. Most IDS have a signature database which it uses to determine the presence of an event, and more advanced IDS have behavioral activity detection to determine malicious behavior. The good thing about the latter is that it allows the sensor to detect what is called a "zero day attack"- compared to signature-based detection which can only detect activities that have already been seen in the wild before.

  ### Backend

  The backend is where actual alerting and recording happens. This allows the sensor to focus on the function of detection for efficiency and speed. The backend collects all events detected by the sensors, and is the component that does the function of alerting. Alerting can come in the form of:

  - Log. Log to the database.
  - E-mail. The alert can be sent to one or more receipients.

- SNMP trap. There are applications out there that can collect SNMP traps of various kinds. The backend can send an SNMP trap to a SNMP trap collection and viewer (e.g. HP Openview).
- Block. Some advance IDS have the ability to cause a connection block (i.e. cause a connection reset--TCP reset--between the hacker's computer and the target)
- Display. The alert can be sent to a console that shows the various events that the sensor is detecting.

Aside from providing the repository and the alerting mechanism, the backend provides the IDS setup and configuration storage.

Frontend

The frontend is the IDS's direct user interface. From the frontend, the user can do the following:

- View events that the sensor has detected
- Setup IDS configuration
- Update signature database and behavioral detection engine
- Update sensor and other parts of the IDS

How they work

The components of an IDS work together as a whole to provide an early warning or post intrusion alerting system. In explaining how things work, we will use a network-based IDS as a point of reference.

On network-based IDS, the sensor will typically run as a separate host, and the backend and frontend will generally run on another host. In a bigger environment, there will be more than one sensor, and the backend and frontend are on a separate host.

Sensors Detect and Report

Sensors typically have a capture network interface and a management network interface. Sensors tap into the network so that it can listen to the various communications within its reach. Fifteen years ago, network hubs were common, and as such connecting the network interface of a sensor to any port would give it access to all traffic that goes to any ports of that hub. When network switches became common, ports used by sensors required configuration to span other ports (i.e. to copy traffic from ports of interest).

As the sensor sees traffic from the network, its capture engine quickly passes it into a buffer, necessary to help it to keep up with the load, and if the buffer isn't large enough or if the engine isn't fast enough, it could drop packets. The detection engine then goes through the buffer and performs network protocol analysis. In performing protocol analysis, it may sometimes have to find multiple packets in order to properly compose a complete higher level protocol message, or in order to see if certain event thresholds are

exceeded. Signature-based event detection happens here. As soon as it detects an event, it reports it to the backend.

Behavioral malicious events require more processing and thus will require the sensor detect certain activities that are fundamental to malicious network attacks. As soon as the sensor detects a malicious behavioral event, it sends this event to the backend. Note that behavioral event detection is more complex and may require the detection of multiple primitive events and event count thresholds before triggering.

Backend Collect and Alert

The IDS backend is the hub of the IDS. All events reported by one or more sensors end up in the event repository database. It is in the backend where it determines how an event is handled. Most IDS allow the user to configure how it responds to events. Critical events typically are setup to send email, SNMP traps, displayed, and in more advanced IDS cause a block. The least significant events may simply be logged.

Frontend - Command and Control

An IDS would be difficult to use without a frontend. The frontend is where a user can setup, configure, and update the entire IDS.

As events collect in the backend, the frontend when opened shows events that have been configured to display (with its appropriate color coding or display indication). It is here where the user can manage detected events. IDS typically can get very event-noisy, meaning that it will report way too many events so that a user can get complacent and thus start to ignore reported events.

In order for an IDS to be useful, it should be fine-tuned to report only events of real significance and use. It is through the console that the user is able to fine tune the IDS detection and response, and if done correctly, will provide the user an early warning system from any intrusion. And if an intrusion actually occurs and is somehow missed, the IDS can report on actual intrusion events.

- ## Summary

How do intrusion detection systems work? By enumerating the various components of an IDS and explaining how each component functions in relation to the others, we've answered this question. Note that it can take a book to help answer this question, and without the right technical background, this article may not make too much sense. However, if you have basic computer and network knowledge, the explanation in this article should suffice.