Intrusion Detection Systems

Intrusion Detection Systems (IDS) are becoming integral parts of network monitoring. This section identifies some of the more common IDS methods in use, and it discusses the capabilities and weaknesses in them. IDS is a relatively new technology, and it shows a lot of promise in helping to detect network intrusions. Intrusion Detection (ID) is the process of monitoring events in a system or network to determine if an intrusion is occurring.

An *intrusion* is defined as any activity or action that attempts to undermine or compromise the confidentiality, integrity, or availability of resources. Firewalls, as you may recall, were designed to prevent access to resources by an attacker. Intrusion Detection Systems report and monitor these activities.

Several key terms are necessary to explain the technology and facilitate the discussion in this section:

Activity An *activity* is an element of a data source that is of interest to the operator. This could include a specific occurrence of a type of activity that is suspicious. An example of this might be a TCP connection request that occurs repeatedly from the same IP address.

Administrator The *administrator* is the person responsible for setting the security policy for an organization. He is responsible for making decisions about the deployment and configuration of the IDS. The administrator should make decisions regarding alarm levels, historical logging, and session monitoring capabilities. He is also responsible for determining the appropriate responses to attacks and ensuring that those responses are carried out.

Alert An *alert* is a message from the analyzer indicating that an event of interest has occurred. This alert would contain information about the activity, as well as specifics of the occurrence. An alert may be generated when an excessive amount of ICMP traffic is occurring or when repeated logon attempts are failing. A certain level of traffic is normal for a network. Alerts occur when activities of a certain type exceed a preset threshold. For instance, you would not want to generate an alert every time someone from outside your network pings a server using the PING program. However, if the pings seemed more frequent, or exceeded a predetermined threshold, you would want to generate an alert.

Analyzer The *analyzer* is the component or process that analyzes the data collected by the sensor. The analyzer is looking for suspicious activity. Analyzers work by monitoring events and determining whether unusual activities are occurring, or they can use a rules-based process that is established when the IDS is configured.

Data Source The *data source* is the raw information that the IDS uses to detect suspicious activity. The data source may include audit files, systems logs, or the network traffic as it occurs.

Event An *event* is an occurrence in a data source that indicates that a suspicious activity has occurred. An event may generate an alert. Events are logged for future reference. Events will also typically trigger a notification that something unusual may be happening in the network. An IDS might begin logging events if the volume of inbound e-mail connections suddenly spikes. This event might be an indication that someone is probing your network. The event might trigger an alert if a deviation from normal network traffic patterns occurs or if an activity threshold has been crossed.

Manager The *manager* is the component or process the operator uses to manage the IDS. The IDS console is a manager. Configuration changes in the IDS are made by communicating with the IDS manager.

Notification *Notification* is the process or method by which the IDS manager makes the operator aware of an alert. This might include a graphic display highlighting the traffic or an e-mail sent to the administrative staff of the network.

Operator The *operator* is the person primarily responsible for the IDS.

Sensor A *sensor* is the IDS component that collects data from the data source and passes it to the analyzer for analysis. A sensor can be a device driver on a system, or it can be an actual black box that is connected to the network and reports to the IDS. The important thing to remember is that the sensor is a primary data collection point for the IDS.

The Intrusion Detection System, as you can see, has many different components and processes that work together to provide a real-time picture of your network traffic. Figure 4.6 shows the various components and processes working together to provide an IDS. It is important to remember that data can come from many different sources and must be analyzed to determine what is actually occurring. As you can see, the IDS is not a traffic-blocking device. It is a traffic-auditing device.

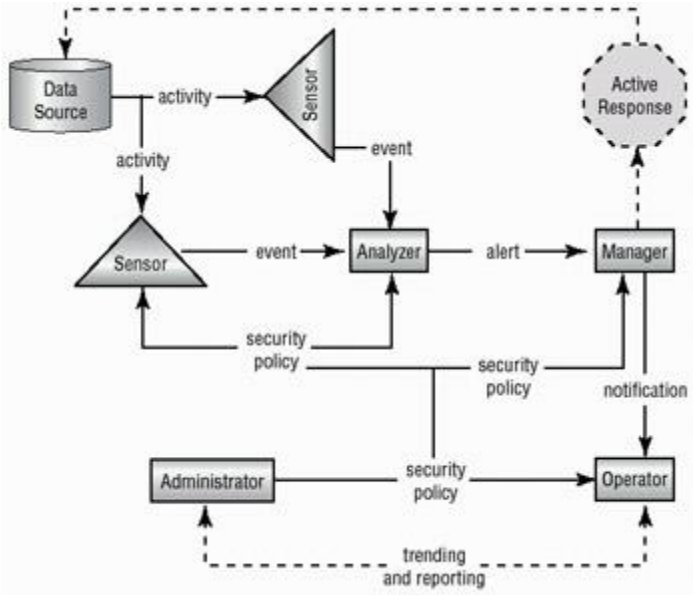

Figure 4.6: The components of an IDS working together to provide network monitoring

Two primary approaches are used in Intrusion Detection Systems: *Misuse- Detection IDS (MD-IDS)* and *Anomaly-Detection IDS (AD-IDS)*. An MD-IDS is primarily focused on evaluating attacks based on attack signatures and audit trails. Attack signatures describe a generally established method of attacking a system. For example, a TCP flood attack begins with a large number of incomplete TCP sessions. If the MD-IDS knows what a TCP flood attack looks like, it can make an appropriate report or response to thwart the attack. Figure 4.7 illustrates an MD-IDS in action. Notice that this IDS system uses an extensive database to determine the signature of the traffic. This process resembles an antivirus software process.

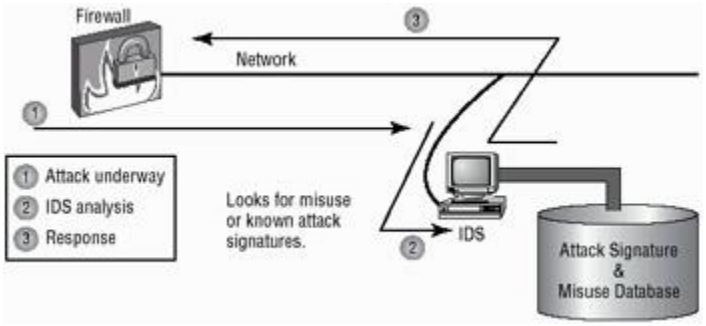

Figure 4.7: An MD-IDS in action

Intrusion Detection Systems are primarily focused toward reporting events or network traffic that deviate from historical work activity or network traffic patterns. For this reporting to be effective, administrators should develop a baseline or history of typical network traffic. This baseline activity provides a stable, long-term perspective on network activity. An example of this might be a report generated when a higher-than-normal level of ICMP responses are received in a specified time period. Such activity may indicate the beginning of an ICMP flood attack. The system may also report when a user who does not normally dial in at night requests administrative access to the system. Figure 4.8 demonstrates an AD-IDS tracking and reporting excessive traffic in a network. The AD-IDS process frequently uses artificial intelligence or expert systems technologies to learn what normal traffic for a network is.

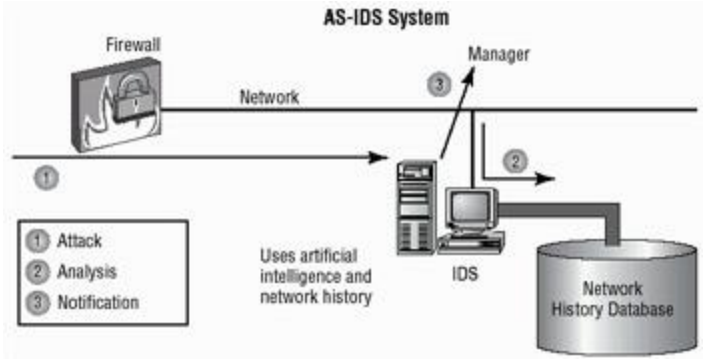

Figure 4.8: AD-IDS using expert system technology to evaluate risks

MD-IDS and AD-IDS are merging in most commercial systems. They provide the best opportunity to detect and thwart attacks and unauthorized access. Unlike a firewall, the IDS exists to detect and report unusual occurrences in a network, not block them.

The next sections discuss network-based and host-based implementations of IDS and the capabilities that they provide.

Network-Based IDS

A *Network-based IDS (N-IDS)* approach to IDS attaches the system to a point in the network where it can monitor and report on all network traffic. This can be in front of or behind the firewall, as shown in Figure 4.9. Placing the N-IDS in front of the firewall provides monitoring of all network traffic going into the network. This approach allows a huge amount of data to be processed, and it allows you to see all of the traffic that is coming into the network. Putting the N-IDS behind the firewall only allows you to see the traffic that penetrates the firewall. Although this approach reduces the amount of data processed, it does not allow you to see all of the attacks that might be developing.

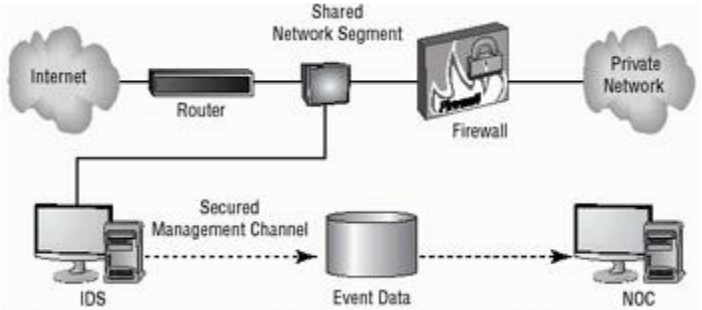

Figure 4.9: N-IDS placement in a network determines what data will be analyzed.

The N-IDS can be attached to a switch or a hub, or it can be attached to a tap. Figure 4.10 illustrates a connection to the network using a hub. Many hubs and switches provide a monitoring port for troubleshooting and diagnostic purposes. This port may function in a manner similar to a tap. The advantage of the tap approach is that the IDS is the only device that will be using the tap.

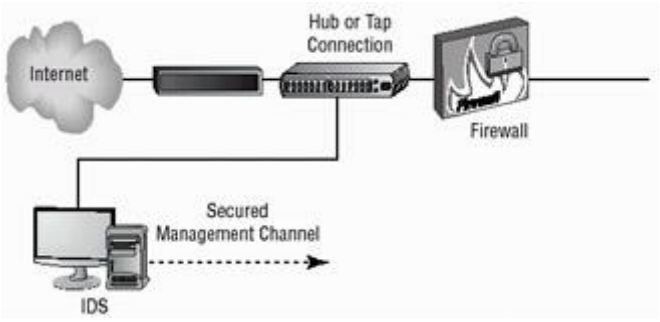

Figure 4.10: A hub being used to attach the N-IDS to the network

In either case, the IDS will monitor and evaluate all of the traffic to which it has access. Two basic types of responses can be formulated at the network level: passive and active. They are briefly explained in the following section.

Real World Scenario: These Network Audit Files Are Killing Me

You are the network administrator of a relatively busy network. Your company has gone through a couple of cutbacks, and your staffing is limited. You want to make sure that your network stays as secure as you can make it. What can you do to ease your workload?

You should consider two primary possibilities to protect your network: either install an IDS or reduce the logging levels of your network audit files.

You might be able to reduce the amount of logged traffic in your audit files by changing the settings on what you audit. However, changing audit rules would prevent you from seeing what is happening on your network because most events would not be logged.

Installing an IDS would allow you to establish rules that would provide a higher level of automation than you could achieve by reviewing audit files. Your best solution might be to convince your company to invest in an IDS. An IDS could send you an e-mail or alert you when an event has been detected.

Passive Response

A *passive response* is the most common type of response to many intrusions. In general, passive responses are the easiest to develop and implement. Passive response strategies include:

Logging The *logging* response is the recording of an event has occurred and under what circumstances it occurred. Logging functions should provide enough information to administrators about the nature of the attack to determine what has happened and to assist in evaluating the threat. This information can then be used to devise methods to counter the threat.

Notification The notification response is the communication of event- related information to the appropriate personnel when an event has occurred. This would include relaying any relevant data about the event to help evaluate the situation. If the IDS is manned full time, messages can be displayed on the manager's console to indicate that the situation is occurring.

Shunning Shunning or ignoring an attack is a very common response. This might be the case if your IDS notices an Internet Information Services (IIS) attack occurring on a system that is running another web-hosting service, such as Apache. The attack will not work because Apache does not respond in the same way that IIS does, so why pay attention to it? In a busy network, many different types of attacks can occur simultaneously. If you are not worried about an attack succeeding, why waste energy or time investigating it or notifying someone of the attack? The IDS can merely make a note of it in a log and move on to other more pressing business.

Active Response

An *active response* involves taking an action based upon an attack or threat. The goal of an active response would be to take the quickest action possible to reduce the potential impact of an event. This type of response requires clear plans for how to deal with an event, clear policies, and intelligence in the IDS in order to be successful. An active response would include one of the responses briefly described here:

Terminating Connections, Processes, or Sessions If a flood attack is detected, the IDS can cause the subsystem, such as TCP, to force resets to all of the sessions that are underway. This will free up resources and allow TCP to continue to operate normally. Of course, all valid TCP sessions will be closed and will need to be reestablished—but at least this will be possible, and it may not have a great effect on the end users. The IDS will evaluate the events and determine the best way to handle them. Figure 4.11 illustrates TCP being directed to issue RST or reset commands from the IDS to close all open sessions to TCP. This type of mechanism can also terminate user sessions, or stop and restart any process that appears to be operating abnormally.

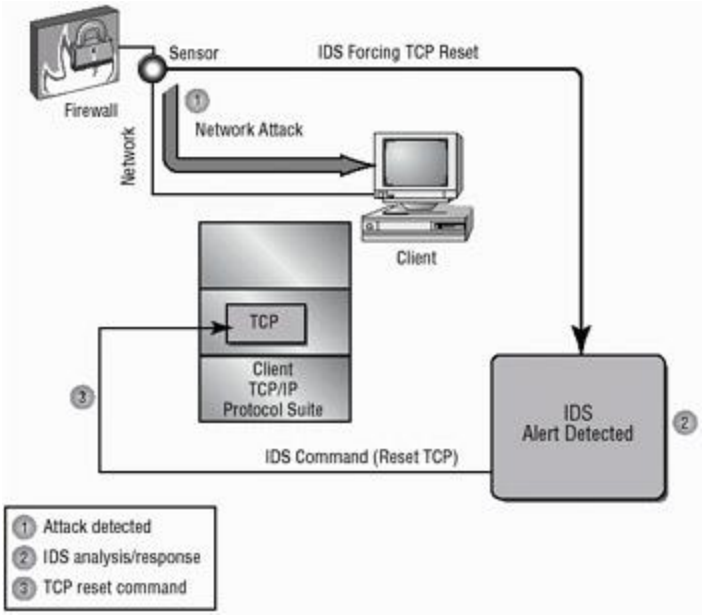

Figure 4.11: IDS instructing TCP to reset all connections

Network Configuration Changes If a certain IP address is determined to cause repeated attacks on the network, the IDS can instruct a border router or firewall to reject any requests or traffic from that particular IP address. This configuration change can remain in effect permanently or for a specified period. Figure 4.12 illustrates the IDS instructing the firewall to close Port 80 for 60 seconds to terminate an IIS attack.

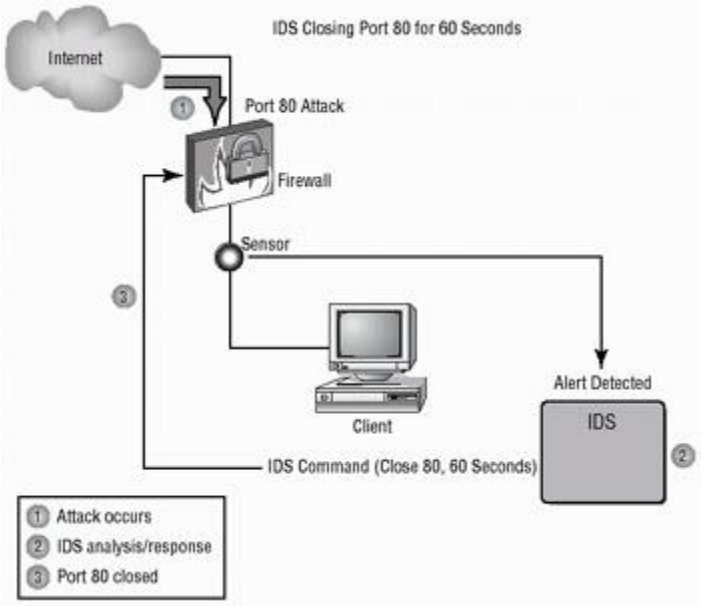

Figure 4.12: IDS instructing the firewall to close Port 80 for 60 seconds to thwart an IIS attack

If the IDS determines that a particular socket or port is being attacked, it can instruct the firewall to block that port for a specified amount of time. This will effectively eliminate the attack, but it might also inadvertently cause a self-imposed DoS situation to occur by eliminating legitimate traffic. This is especially true for Port 80 (HTTP or web) traffic.

Deception A *deception* active response fools the attacker into thinking the attack is succeeding while monitoring the activity and potentially redirecting the attacker to a system that is designed to be broken. This allows the operator or administrator to gather data about how the attack is unfolding and what techniques are being used in the attack. This process is referred to as *sending them to the honey pot*, and it is described later in this chapter. Figure 4.13 illustrates a honey pot where a deception has been successful.

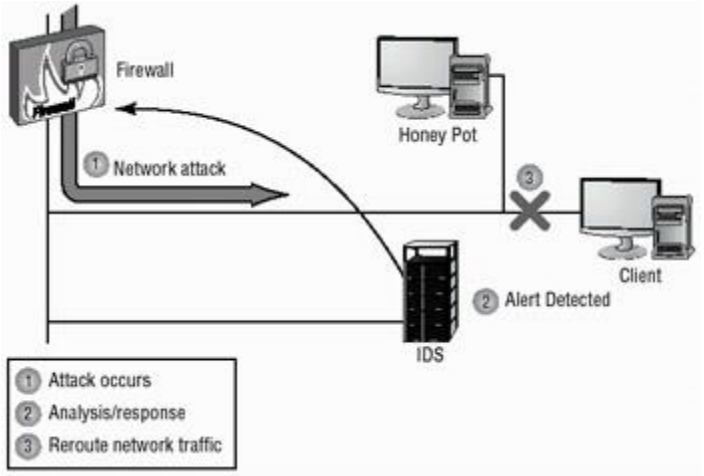

Figure 4.13: A network honey pot deceives an attacker and gathers intelligence.

The advantage of this type of response is that all activities will be watched and recorded for analysis when the attack is completed. This is a very difficult scenario to set up, and it is dangerous to allow a hacker to proceed into your network, even if you are monitoring the events.

This approach is frequently used when an active investigation is underway by law enforcement and they are gathering evidence to ensure a successful prosecution of the attacker.

Host-Based IDS

A *Host-based IDS (H-IDS)* is designed to run as software on a host computer system. These systems typically run as a service or as a background process on the computer system. H-IDS systems will examine the machine logs, systems events, and applications interactions. H-IDS systems do not normally monitor incoming network traffic to the host. H-IDS systems are very popular on servers that use encrypted channels or channels to other servers. Figure 4.14 illustrates an H-IDS installed on a server. Notice that the H-IDS interacts with the logon audit and kernel audit files. The kernel audit files are used for process and application interfaces.

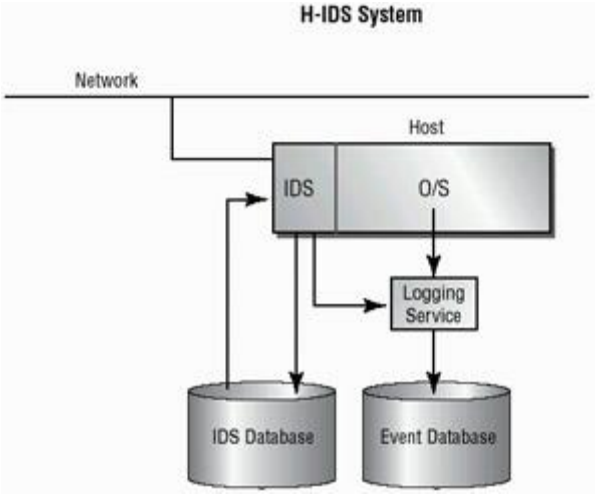

Figure 4.14: A Host-based IDS interacting with the operating system

Two major problems with H-IDS are not easily overcome. The first problem involves a compromise of the system. If the system is compromised, the log files that the IDS uses to report to may become corrupt or inaccurate. This may make fault determination difficult or the system unreliable.

The second major problem with H-IDS is that it must be deployed on each system that needs it. This can create a headache for administrative and support staff.

One of the major things that H-IDS provides is the potential to keep checksums on files. They can be used to inform system administrators that files have been altered by an attack. This makes recovery easier because determining where tampered has occurred is easier.

Note Host-based IDS systems typically respond in a passive manner to an incident. An active response would theoretically be similar to those provided by a network-based IDS system.