

Intrusion Detection Systems

Intrusion Detection Systems (IDS) can be a useful way to monitor networks and critical computers for signs of unusual activity. They can provide early warning of security and other problems, allowing incidents to be dealt with quickly and their impact reduced. However an IDS that is not properly configured and maintained is likely to generate many spurious alerts, wasting staff time and possibly missing important new signs.

Tuning an IDS to provide the appropriate level of alerts for a particular network or host is likely to take some time. IDS look for unusual signs but they cannot determine whether activity is hostile or merely a change resulting from an unforeseen use of the network or system. Each alert will normally need to be checked by a system or network manager. An alert caused by normal activity is called a 'false positive', while hostile activity that does not generate an alert is a 'false negative'. All IDS are likely to generate both types of error: tuning the system to reduce false negatives will almost always increase the number of false positives and vice versa. The appropriate sensitivity is a matter of local policy and should be governed by the risk to the systems being monitored and the resources that are available to respond to alerts.

Network-based IDS, in particular, will normally detect *attempts* to breach security rather than whether the attempt was successful or not. Identifying successful intrusions is likely to require prompt investigation by a skilled human. If this is not done then the IDS's warning is likely to be wasted.

How Do They Work?

Different IDS packages are designed to monitor activity on hosts and on networks. Packages are named here only as examples of their type. Host-based IDS such as Tripwire, AIDE and Samhain take a snapshot of the files on a computer and then generate alerts if there are unexpected changes to the permissions, ownership or content of critical files. These can, for example, detect tampering with password files, system programs or security configurations. Host-based IDS are particularly useful on critical servers.

Network-based IDS such as Snort® examine network packets and compare them against rules designed to identify particular types of unusual activity. Rules may match packet headers or content. Like anti-virus programs, network-based IDS must have their rulesets updated regularly to look for new problems. Network-based IDS often run on dedicated systems as large rulesets can require high-performance hardware. Such systems may also need special network connections as on a switched network the IDS will receive only packets addressed to itself and not those for other systems. Many switches can be configured to copy all packets to one port (known as a 'spanning port' or 'port mirroring') and this may be the best way to connect an IDS. Alternatively the IDS function may be incorporated into a firewall or router, though these tend to be less flexible than stand-alone systems.

A third type of IDS collects information about network flows, rather than individual packets. Flow information is written to a database that can be queried either by a human investigator or by regular jobs to generate alerts for suspicious patterns of flows. Packages such as Argus and NFSen can be used to detect network scanning, unofficial servers and many other types of problem. Flow data can be collected from a spanning port or generated by suitable routers. Graphical presentations of flow data or other historical information from IDS can be useful in identifying trends and changes.

What Can They Do?

Intrusion Detection Systems are usually configured to send an alert to a human being when they detect suspicious activity. Many types of alerts can be used, from entries in logfiles to e-mail or text messages sent by SMS.

It is also possible to configure an IDS to act automatically to try to reduce the impact of the suspicious activity, making it an Intrusion Prevention System (IPS). A network-based IDS might, for example, dynamically alter a firewall configuration to block packets from a particular IP address for a few minutes, or a host based IPS such as Microsoft®'s URLScan or Apache™'s ModSecurity may use its knowledge of the application layer to detect and block suspicious requests to the web server it protects.

There are a number of potential problems with automated reaction so it is important to balance these risks against the danger of allowing the activity to continue:

- False positives will cause innocent, and possibly vital, activity to be blocked;
- An attacker may forge packets to appear to come from a critical system, such as a DNS server, thus causing you to deny service to yourself;
- The link between the IDS and the firewall or other system implementing the preventive action must be secure. An attacker must not be able to interfere with the link so as to turn off the protection provided by the firewall.

Some IDS can be configured to send packets to the source of the suspicious traffic. This option should be used, if at all, with extreme care. Even in a genuine attack the machine making an intrusion attempt is likely to be an innocent third party and any response that could be interpreted as an attack could well be a criminal offence.

References

Host-based IDS:

Tripwire: <http://sourceforge.net/projects/tripwire/>
<http://www.tripwire.com/>

AIDE: <http://sourceforge.net/projects/aide>

SAMHAIN: <http://la-samhna.de/samhain/>

Network IDS:

Snort®: <http://www.snort.org/> (Many output and alert options in 'Contrib' section)

Argus: <http://www.qosient.com/argus/>

NFSen: <http://nfsen.sourceforge.net/>

Application level IPS:

Microsoft® URLScan: <http://www.microsoft.com/technet/security/prodtech/IIS.msp>

Apache™ ModSecurity: <http://www.modsecurity.org/>

Training



JANET runs a training course on Using Logfiles for Security, which covers the use of IDS. More details are available at:

<http://www.ja.net/services/training/courses/Logfiles.html>