Understanding how IDS works

 **www.techrepublic.com**/article/solutionbase-understanding-how-an-intrusion-detection-system-ids-works/

In a perfect world, the only people who are accessing your network are people you know about and trust. You want vendors and clients to be able to access your network. You also want remote users to be able to access the network at night.

But what about that guy sitting alone in his bedroom hyped up on caffeine and sugar with nothing better to do but to see what he can do to your network? How do you tell the difference between legitimate and illegitimate network access?

That's where an intrusion detection system comes in. When used in concert with an intrusion prevention system, you can detect and stop hackers before they get anywhere close to important data on your network. Here's how they work.

Why won't a firewall work?

A network intrusion is an attack that comes from outside the LAN, usually across the Internet. Intrusion detection systems and intrusion prevention systems (IDS/IDP) are important elements in any network security plan.

Some organizations depend on their firewalls for IDS/IDP, and major firewall vendors are building some intrusion detection and prevention features into their products. However, a dedicated IDS/IDP device can detect a much broader range of malicious activities than those built into most firewalls.

The line between firewall and IDS technology is becoming more blurred, however. Traditionally, a difference was that an IDS could "understand" the contents of packet headers, such as flags and options, rather than just looking at IP addresses and ports. Firewalls have gotten smarter, though, and application layer filtering firewalls can perform the same sort of deep inspection as an IDS. Add-on products, often made by IDS vendors, can make these firewalls perform even more efficiently in the IDS role.

Technically speaking, an IDS performs the same function for your network that a security alarm system performs for your home or office building. That is, it monitors for intrusions and lets some one know when one occurs, but doesn't really do anything to prevent the intruder from entering. You probably wouldn't consider implementing an alarm system on a building without also coupling it with preventative measures such as strong locks and perhaps security guards.

Likewise, an IDS needs to be combined with mechanisms to prevent intrusions as well as detecting intrusion attempts. That's why most products today that are called IDS are really combination IDS/IPS.

What's the difference between IDS and VS?

You may be wondering how an IDS or IPS differs from a vulnerability scanner (VS), since they seem to perform similar functions. A vulnerability scanner works more like "preventative medicine" – you scan the network with VS software to determine whether any vulnerabilities exist that can be exploited by hackers and attackers, and patch the holes it finds before an attack actually occurs.

Unfortunately, many attacks are based on protocol or application characteristics that can't be "fixed" without impairing the operation of your production network. An IDS spots attack attempts in their early stages and lets you know what's happening, and an IPS employs countermeasures to stop an attack in progress.

There's a place on your network for both vulnerability scanners and intrusion detection/prevention systems; the two can work together to keep your resources more secure.

Intrusion detection approaches

IDS vendors implement their products in different ways and there are consequently several ways to categorize intrusion detection systems. The first is based on the scope of the IDS's monitoring; that is, whether it is installed on and uses data from a single host computer, or is a network-based product that monitors traffic on the network as a whole, as well as analyzes data from individual computers.

Another difference in implementation has to do with how the vendor markets the system, either as a software product or as an integrated hardware device (appliance).

Host-based intrusion detection

A host-based IDS is one in which the software is installed on a single system and the data from that system is used to detect intrusions. Because the host-based IDS protects the server "at the source," it can more intensely protect that specific computer. The host-based system usually examines log files on the computer to search for attack signatures. Important system files and executables may also be checked periodically for unexpected changes. A host based system will also monitor ports and trigger an alert if certain ports are accessed.

Network-based intrusion detection

A network-based IDS monitors data from network traffic as well as data from one or more host computers to detect intrusions. A network-based IDS analyzes data packets sent over the network, and generally uses a "promiscuous" network adapter (one that is capable of reading all of the packets sent over the network, rather than just those packets addressed to it). The network-based IDS examines packet headers, which are generally not seen by the host-based IDS. This allows the detection of Denial of Service (DoS) and other types of attacks that may not be detected by a host-based IDS.

IDS software

IDS software can be installed on a regular PC running a standard network operating system, and has the same advantages as a software firewall in comparison to a firewall appliance:

- You have more control over the selection of the hardware.
- If you already have available hardware, you may pay less than for an appliance.
- There are freeware software IDS and IPS products, although they tend to be more rudimentary than the commercial products.
- You can upgrade the hardware more easily (add or upgrade the processor, add more memory, etc.).

IDS appliances

IDS appliances come as "turn key" packages with the software already installed, often on a proprietary operating system. Advantages of hardware-based IDS include:

- IDS appliances offer "plug and play" setup and functionality with an easier learning curve.
- The proprietary OS may be less vulnerable to exploitation and attack.
- The hardware can be optimized for high performance.
- Cost of an appliance may be less than the combined cost of IDS software, the PC operating system and the hardware on which to run it.

How IDS works

IDS systems can use different methods for detecting suspected intrusions. The two most common broad categories are by pattern matching and detection of statistical anomalies.

Pattern matching

Pattern matching is used to detect known attacks by their "signatures," or the specific actions that they perform. It is also known as signature-based IDS or misuse detection. The IDS looks for traffic and behavior that matches the patterns of known attacks. The effectiveness is dependent on the signature database, which must be kept up to date.

Pattern matching is analogous to identifying a criminal who committed a particular crime by finding his fingerprint at the scene. Fingerprint analysis is a type of pattern matching.

The biggest problem with pattern matching is that it fails to catch new attacks for which the software doesn't have a defined signature in its database.

Statistical anomaly

Anomaly-based detection watches for deviations from normal usage patterns. This requires first establishing a baseline profile to determine what the norm is, then monitoring for actions that are outside of those normal parameters. This allows you to catch new intrusions or attacks that don't yet have a known signature.

Anomaly detection is analogous to a police officer who walks or drives a particular beat every day and knows what is "normal" for that area. When he sees something that's out of the ordinary, it creates reasonable suspicion that criminal activity may be going on, even though he may not know exactly what crime is being committed or who is responsible.

There are several different anomaly detection methods, including:

- Metric model
- Neural network
- Machine learning classification

A problem with anomaly-based IDS is the higher incidence of false positives, because behavior that is unusual will be flagged as a possible attack even if it's not.

Where the IDS fits in your security plan

Your edge or front-end firewall is the first line of defense in protecting your network against intruders, and it will likely have its own intrusion detection capability, although it may detect and prevent only a limited number of known attacks/intrusions. A network-based IDS is often placed between the edge firewall and a back-end firewall that protects the internal network from the publicly accessible network in between, called the DMZ or perimeter network or screened subnet.

Placing the IDS in this location allows it to do its job on all traffic that gets through the edge firewall and provides an extra layer of protection for the DMZ, which is the most vulnerable part of your network since it contains your public servers such as Internet-accessible Web servers, DNS servers, front-end mail servers, etc.

Putting the IDS in front of the edge firewall would result in a greater load on the IDS, since it would respond to many scans, probes and attack attempts that could otherwise be filtered out by the firewall. Also, the huge number of alerts might lead to a "IDS who cried wolf" situation in which administrators would start ignoring the alerts when many of them don't lead to real attacks.

You may also want to place an IDS behind the back-end firewall to detect intrusions on the internal LAN.

A multi-layered approach

The best security is afforded by using one than one IDS (for example, an IDS in the DMZ and another on the internal network) and by using both network and host-based IDS. Host-based IDS can be installed on critical servers for multi-layered protection.

Incident response

The detection of intrusions is only the first step in making an organization more secure and protecting against intruders. The real key is what happens after the intrusion is detected: your incident response plan.

To be effective, response must be as immediate as possible. That's why your IDS needs to include notification features and you need to set them up so that the alerts get to the proper people as quickly as possible after an intrusion is detected.

Your incident response team should practice incident response procedures before the "real thing" occurs, so that each member knows his/her role and acts instinctively to accomplish the team's goals, which include:

- Preventing damage (or further damage)
- Tracking/identifying the intruder
- Preserving evidence in case the incident leads to criminal prosecution and/or civil litigation

Popular IDS/IDP

There are many IDS solutions on the market today, as well as free/open source IDS tools that you can download. The best solution for your organization depends on your network's size, security needs, existing security infrastructure, budget and IT department structure and workload.

Smaller companies on a tight budget may benefit from freeware solutions that provide basic IDS without additional cost. Larger organization with sufficient IT personnel may prefer to implement commercial IDS software, while those with overburdened IT departments may prefer the quick setup and low overhead of an appliance-based solution.

IDS freeware

The most popular freeware or open source IDS products include:

- Snort does real time analysis of IP packets and can search content and perform pattern matching to detect common attacks such as buffer overflows, SMB probes, port scans and CGI attacks. There are versions available for both Linux and Windows. There are plug-ins available that extend its detection and reporting capabilities.

Snort is often thought of as a protocol analyzer, and the line between such "sniffers" and IDS can be thin. Other protocol analyzers and monitors, such as Sunbelt Software's LanHound, can also perform some IDS functions.

GFI LANGuard S.I.M. is an "entry level" IDS that also provides forensic evidence gathering, Web site monitoring, and system recovery for Windows servers and workstations. Tripwire is a file integrity checker for UNIX/Linux that can be used for host-based intrusion detection (not intrusion prevention). It works by creating a baseline of a system's files and directories and identifying changes, based on a policy file you configure to define which attributes should be compared. The source code can be downloaded from http://sourceforge.net/projects/tripwire or you can get packages with preconfigured policy files from Linux vendors such as RedHat, Mandrake, SuSE and Debian. There is also a commercial enterprise version of Tripwire available for purchase at http://www.tripwire.com.

Commercial IDS software

Although the current trend is toward appliance-based intrusion detection systems, there are some well-regarded software IDS products available:

- **Internet Security Systems (ISS) RealSecure** is a network-based IDS that monitors TCP, UDP and ICMP traffic and is configured to look for attack patterns. It can implement countermeasures (IPS functionality) if an attack is detected. ISS has bundled its product and created add-ons for popular software-based firewalls such as Microsoft's Internet Security and Acceleration (ISA) server and Check Point's Firewall-1.
- **GFI LANGuard Security Event Log Monitor (S.E.L.M.)** is a host-based IDS that performs real-time monitoring of Windows systems' security event logs to detect anomalous events and send alerts.

IDS appliances

Many vendors make IDS appliances. Some of the most popular include:

- **IntruShield**, made by McAfee, is an enterprise-level intrusion detection and prevention appliance that comes in six models ranging from the 1200, which supports 100 Mbps throughput and provides 2 Ethernet ports, to the 4010, which supports up to 2 Gbps with 12 gigabit ports. It offers centralized management and provides real-time preventative measures to protect against attacks.
- **Cisco Secure Intrusion Detection System (Cisco IDS)** rack-mount appliances provide monitoring using stateful pattern recognition, protocol parsing, heuristic detection and anomaly detection. There are four different models in the IDS 4200 series, with performance ranging from 80 Mbps to 1000 Mbps.
- **Top Layer Attack Mitigator IPS** protects against content based attacks such as worms, Trojans, viruses and vulnerability exploits and rate-based attacks such as Denial of Service (DoS) and Distributed DoS, using deep packet inspection and stateful analysis engines, and reduces false positives with its advanced protocol validation modules.
- **Proventia IDS** appliances are based on the ISS software and come in models designed for aggregate network bandwidths of 200 to 1200 Mbps for up to 14 network segments.

Halting the barbarians at the gate

Intrusion detection is just as essential to your network as a burglar alarm system is to commercial buildings or homes where valuables are kept. A good IDS will also include IPS functionality; rather than just telling you someone is breaking into your network, it will do something about it. IDS/IPS products can be host- or network-based (and the two can be used in conjunction) and can be implemented via software installed on one of your network's servers or as a dedicated appliance. There are many reputable brands of IDS/IPS out there, and if your organization doesn't currently have an IDS solution in place, you should consider adding one to your existing security infrastructure.

Full Bio

Debra Littlejohn Shinder, MCSE, MVP is a technology consultant, trainer, and writer who has authored a number of books on computer operating systems, networking, and security. Deb is a tech editor, developmental editor, and contributor to over 20 additional books on subjects such as the Windows 2000 and Windows 2003 MCSE exams, CompTIA Security+ exam, and TruSecure's ICSA certification.