CYBERHOUND

IPS/IDS ESSENTIALS

WHAT ARE INTRUSION DETECTION & PREVENTION SYSTEMS?

A technical whitepaper on network security.

INTRODUCTION

Intrusion Detection Systems (IDS) provide visibility to IT security personnel that an unauthorised event may have taken place. IDS systems sit off to one side of the traffic path (port mirroring, receiving Netflow data, etc.) and therefore do not actively block suspect packets.

Intrusion Prevention Systems (IPS) control traffic flowing into, or out of, a network or host, looking for known exploits, or unusual activity, and then take action on that traffic - typically by dropping the packets. By their very nature, an Intrusion Prevention System is also an Intrusion Detection System, however it may not provide the same level of visibility as a dedicated IDS.

In fact, the term "intrusion" can be misleading as these systems are looking for more than active intrusion attempts. The image of a hacker in a dark room leaning over their Kali laptop comes to mind when terms such as intrusion detection and prevention come up.

However, whilst IDS/IPS can and do help with active hacking and exploit attempts, they also scan for much more, such as data loss prevention (DLP), which can greatly assist, given today's heightened awareness of data breach laws.

With this in mind, it is important to identify the functionality required. For example, if providing visibility and logs into a centralised Security Information and Event Monitoring (SIEM) solution is the primary goal, then the complexity of deploying an IPS would not be required. If, however, controlling the traffic passing between two or more networks is the goal, then an IPS sitting inline to be able to take action is necessary.

On the surface, the role of a firewall looks very similar to that of an IDS/IPS. In reality, the role a firewall performs is quite different to that of an IDS/IPS. Traditional, and even "next generation" firewalls, work around specific rules to dictate what traffic is allowed into, or out of a network.

This can be based on source/destination IP and port traffic, or doing a deeper dive into the packet to determine the application that is generating that traffic. It's important to note that a firewall itself does not typically do that deep packet inspection to determine if the traffic is part of an exploit attempt, or is malicious in intent. However many modern firewall systems have embedded IDS/IPS modules to do just that, which forms the backbone of a Unified Threat Management (UTM) system.

IPS

IPS are more complex to deploy as you need to form a channel that all the data must flow through in order for it to do its job, meaning the network topology may need to be altered to accommodate the IPS appliance. Some firewalls can have IPS modules embedded within them (UTMs), which can make it easier to deploy, whereas some IPS are dedicated appliances, so will need to slot into position. Questions around changes to the logical network flow to be able to inspect the correct traffic, bandwidth requirements, and high availability all need to be answered, just to scope out the network deployment of an IPS (other factors such as logging and alerting of events and actions taken also form a significant portion of the conversation).

IDS

Unlike IPS, IDS is not required to be inline of the data path. Most Intrusion Detection Systems can receive traffic courtesy of port mirroring, or netflow data from switches and routers. Whilst it's not the role of an IDS to actively control that data, it must have some effective way of alerting relevant staff when it flags an event. Systems are becoming smarter, with the advent of more advanced reporting SIEM systems.

SIEM

IPS/IDS logs may be shipped to a third party SIEM solution for log analysis, event correlation and data analytics to better assist in identifying threats and determining the required course of action. This forms a crucial part of an organisation's overall risk management strategy. Many Managed Security Providers (MSP's) offer Security Operations Centre (SOC) services to better manage risk to the network and an organisation's data.

NETWORK & HOST

Intrusion detection and prevention systems typically come in two different flavours: network and host based.

- A Network IDS/IPS inspects packets flowing through a network (egress and ingress), to identify suspect activity.
- A host-based IDS/IPS runs on individual servers and workstations, and looks for any exploit attempts on that particular host.

The visibility of a network IDS/IPS system is typically greater than that of a host-based system, however care must be taken to ensure that the network IPS does not form a bottleneck and impede traffic flow.

What can (network) IDS/IPS do? Within a Next Generation Firewall, IPS is often the first line of defence against malicious traffic, cyber criminals and zero-day security vulnerabilities. Utilised early in the packet flows, the IPS engine will scan traffic and match against an IPS signature ruleset. Upon detection, the next generation packets are dropped by the firewall.

DETECTION METHODS

Detecting intrusion attempts is typically done in a couple of different ways (signature or anomalous) with each method having their pros and cons.

Signature detection looks within packet sequences for specific patterns or strings. For example, checking known patterns within an exploit payload can detect attacks that are attempting to exploit a particular buffer overflow vulnerability. This method is usually very quick, and assuming the signature database is updated constantly, produces few false positives (legitimate traffic being classified as malicious). Assuming the signature map is configured for the site in question, it can be used within high bandwidth traffic flows quite effectively. Virtually every IDS/IPS on the market today employs signature detection to some degree.

There are some limitations to signature detection systems. Primary amongst those is the need to know what to look for. An exploit for a buffer overflow usually always needs to be released in the wild before signatures can be created for IDS/IPS to use. This increases the time between an exploit being in the hands of attackers, to IDS/IPS vendors producing a valid signature.

If an exploit changes (*exploit.a -> exploit.b*), then an additional signature for *exploit.b* is almost always required, therefore constant signature updates are critical to the effective working of an IDS/IPS.

Additionally, frameworks like Metasploit make it relatively easy to obscure the exploit payload to trick simple signature-based IDS/IPS by employing NOP (no operation) generators and payload encoder methodologies.

An **anomaly-based** IDS/IPS looks at the traffic and decides if it's normal or not. If not, then it alerts and potentially takes action. To be able to do this, the system needs to go through a learning stage, where it establishes a baseline of normal behaviour. After this training period, it will then look for any activity that falls outside the established baseline. There may well be multiple training periods to take into account events like end of month and end of financial year activity, and quiet network periods (between Christmas and New Year, for example).

These systems are also getting "smarter" based around the fact they are geared towards artificial intelligence, and how systems like neural networks develop over time.

An example of anomalous detection is when a machine is infected with a worm. That machine would then start scanning all other hosts around it to determine if any are vulnerable. The signature for the worm may not be in the database (indeed, the worm may not even be known yet), however the behaviour of that host would typically fall outside the baseline, thus prompting an alert and action. This provides good zero-day security, with caveats.

Anomaly-based IDS/IPS typically cannot detect malicious activity if that activity falls within the network baseline, or adheres to protocol standards.

They are also prone to more false positives than signature-based IDS/IPS - organisations are rarely, if ever, static. What an IDS/IPS would pick up as anomalous could simply be increased customer traffic, or larger than expected file transfers.

Significantly, creating rules for anomaly-based IDS/IPS are more complex. These systems must know about all the protocols they would expect to see and how those protocols should behave, otherwise they can't tell if the traffic is normal or malicious. For well established traffic types such as HTTP/S, SMTP, and DNS, this should pose no problem. However, for custom applications that have their own communication protocols, that just happen to use TCP port 443, it can be difficult to incorporate that into the IDS/IPS (this is especially so if the communications protocol is proprietary and closed-source).

CONCLUSION

Intrusion Prevention provides an important defence layer of both known and unknown threats allowing automated intervention, protection and analysis of the network and the users within. IDS/IPS is a complex process, branching over a number of different areas within an organisation.

As such, these systems need to be carefully evaluated - a task that should be done in consultation with security experts, and vendors with a long, proven-track record within your sector. It should be implemented as a well-defined and integrated layer of your overall threat management strategy. Additionally, unless you have a dedicated security team, and significant budget, stand-alone IDS/IPS systems can be too complex and expensive to implement.

Most network IDS/IPS services form part of a Unified Threat Management system (UTM) that also incorporates firewall functionality, and usually web scanning and reporting. The reporting on such systems tie the various roles of a UTM together, which helps deliver a holistic overview of the threat landscape to an organisation, better arming the IT team to manage and mitigate threats as they appear.

ABOUT CYBERHOUND

CyberHound is an innovative provider of cybersecurity and internet solutions to schools. CyberHound has invested millions of dollars in developing a unique K12 solution for Australian schools. This investment increases every year to ensure ongoing innovation.

CyberHound has developed one of the most advanced set of multi-layer features to deliver reliable cybersecurity for schools. These are all supported and updated by CyberHound's Managed Security Cloud Services - all delivered securely from the most secure data centres in Australia.

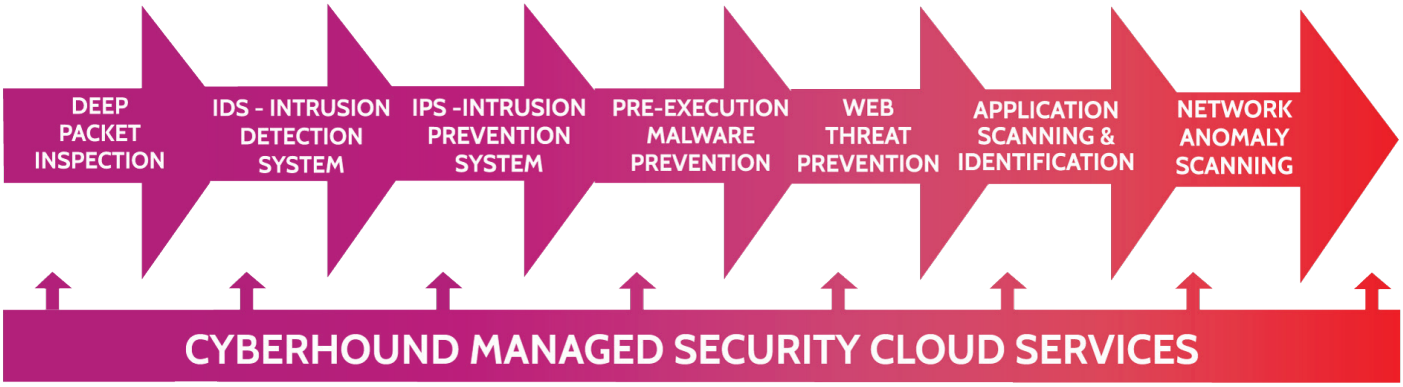

CYBERHOUND.COM