

Zen and the Art of Intrusion Detection

David Beesley of consultancy Network Defence explains how you and your networks can achieve inner calm in the face of Internet security risks



If a tree falls in a forest with no-one to hear it, does it make a sound? So goes a typical zen-like philosophical question. While it's a thought-provoking question, what has it got to do with Intrusion Detection Systems (IDS)?

The answer is simple – if you're not there to watch the tree fall, do you need to know whether it fell or not? The same principle applies with IDS. There's a forest of threats to your network out there, so how do you establish where the threats really lie?

Do you set up your IDS to watch every tree, and risk getting lost in the detail? Or do you narrow the focus to only the trees that you think represent the biggest threats, and risk missing the source of an attack?

There's no easy answer. What's more, IDS are known for producing many false positives. They can miss new attacks and let them go undetected. They need regular care and maintenance. But IT's plate is usually full with issues like desktop software management, backups, anti-virus issues, virtual private networks, firewalls, spam filtering and more.

So IDS can be something of a handful for the IT department, demanding a delicate balancing of priorities if a hard-pressed team is to manage an IDS effectively. It's no wonder that because of these issues, outsourcing IDS activity to a trusted third party is a popular choice, allowing IT staff to concentrate on its core activities.

But how do you ensure that the third party delivers the right managed service to suit your needs? Which managed IDS configuration will suit your network best? Should you be doing internal detection of potential attacks, or perimeter detection? Here's a guide to identifying your network's needs, the questions you should ask of a thirdparty IDS provider – and to the shortest route to inner network calm.

Configuration matters

There are a number of possible configurations for a managed intrusion detection service. The correct configuration will depend on your security policy – and should be agreed in consultation with the managed IDS provider.

Usually, IDS are configured either for internal detection or for perimeter protection. For internal detection the normal practice is to place sensors to protect key servers within the organisation. These sensors will connect back to a centralised database.

These sensors will have two network interface cards. One card will be placed in promiscuous mode and will be gathering data. This card will not have a fixed IP address and the network cable will be configured so that data cannot be transmitted from this interface back to the network, so the sensor is in 'stealth' mode. The second interface will be configured with an IP address and ideally will be connected to a separate VLAN or subnet. This will send data from the sensor back to the database.

For perimeter protection sensors can be placed inside or outside the perimeter firewall. If placed outside, the sensor will see all attacks that are inbound to the organisation, including any that are stopped by the firewall. This position gives an overall threat level for the organisation, and will use an encrypted tunnel is used to communicate with the management network.

The management network includes a database, which each sensor writes alerts to, and a GUI front end which allows managed service staff to view alerts in real time and respond accordingly.

In these configurations, no change is made to your firewall set-up and no communication is allowed into your network from the sensor. From a risk assessment perspective, if an attacker compromised the sensor, they would have access to all data passing through the untrusted network, so sensitive traffic should be encrypted while passing across the internet, or sent via an alternative method.

Internal affairs

However, if sensors are placed inside the firewall, the IDS will only see attacks that have passed through the firewall. This reduces the overall number of attacks seen, but it also prevents viewing of threats that have been blocked by the firewall.

This solution is intrinsically safer as it means that both the customer's firewall and the sensor's firewall need to be defeated before the attacker can attempt to compromise the sensor.

If you have a switched internal network, then sensor placement is crucial. To avoid overwhelming sensors on larger networks, look for a solution involving multiple sensors, with each protecting a single server or group of servers all passing data back to a central database.

Questions, questions

When it comes to choosing a third-party IDS provider, it's important to ask the right questions to ensure you get the service that matches your needs.

The first question to ask is what IDS solution does the provider use? Is it open source or proprietary? The best answer is an open-source solution, such as Snort. The key advantage of widely-used open-source software is that identification of and updates for new attacks are written extremely quickly. Usually a Snort signature for a major new alert is available faster than updates for commercial solutions – which means you get earlier protection against new threats as they arise.

Second, can the third party manage the IDS on a 24x7 basis, if required? Unfortunately, hackers rarely work conventional office hours, so round-the clock cover should be available to you.

Third, how does the provider link to the sensors on your network? It should be done via an encrypted tunnel, so that the data cannot be accessed and used.

Fourth, what service level options are available to you from the service provider? How fast can they update the managed IDS? What speed of response can they offer in reaction to differing threat levels? A proactive provider should sit down with you and assess the various threats, the risk levels, and recommend appropriate responses.

Finally, what reporting can the provider offer? You should get quick, pro-active

alerts in the event of a major issue, backed by regular reports which outline the overall attack profile on your company's network.

All of these points should be enshrined in a comprehensive service level agreement that determines the activities that are carried out by the managed service organisation, and the level of reporting and alerting that takes place.

With the right approach to a managed IDS service, you can save the costs of buying your own IDS solution, and spare yourself the ongoing stresses of configuring and managing the solution – leading to inner calm in both mind and budget.