


Effective Intrusion Detection



With careful configuration and management, intrusion detection systems can make a valuable contribution to IT infrastructure security

This white paper examines some of the key issues that organisations need to consider as part of implementing an effective intrusion detection solution:

- Performance considerations for network- and host-based detection
- Scalability of intrusion detection solutions
- Deployment and integration of different vendor's products
- Tuning and optimisation of sensors
- Integration with incident handling strategies.

Intrusion Detection Systems (IDS) are fast becoming an indispensable component of many organisation's security strategies.

The rapid growth in their use is largely based on the premise that any operational IT system is likely to be the subject of some form of attack, or reconnaissance-type activity, during its lifetime. Certainly, in the case of Internet-facing systems, many are likely to be probed, scanned and interrogated within the first few hours of being connected to a network.

It's important, also, to consider the decreasing lifespan of malicious attacks. A shrewd and determined attacker can easily scan a system for vulnerabilities, gain access to it, modify or remove information and then eradicate any trace of their actions in a matter of minutes. It really doesn't take long.

When attacks do succeed, consider how they're recognised. A defaced web site is easy to spot. Unfortunately, history shows that it's normally a customer or prospect that reports the intrusion first. More difficult to identify is theft of client information. Here, it's often an irate supplier or credit card organisation that acquaints you with the news.

Clearly, then, any measure that can help to identify attempts to compromise IT systems and alert operations staff to the fact must be of potential benefit. This is what an Intrusion Detection System can help achieve.

But designing, deploying and managing an IDS solution is more than just an investment in technology. As with many other aspects of information security, a

combination of product, people and process will be required that can:

- Monitor networks and host IT systems for evidence of an attack that's either underway or in preparation. It's important to monitor at both levels simply because some network-based attacks will be undetectable by servers. Conversely, network-level detection won't always identify attacks that are concealed, say, in encrypted network traffic.
- Be able to create and maintain intrusion detection policies and deploy these rapidly across local and remote networks to individual intrusion detection sensors.
- Analyse, consolidate and correlate data from a wide variety of different sources such as network and host detectors themselves, system event or error logs and then present this vast quantity of information in such a way that operations staff can derive value from it.

Keeping pace with the traffic

One of the first areas to consider in designing an intrusion detection solution is data analysis – the process of examining network traffic or host activity for signs of an attack. A number of issues need to be assessed including the bandwidth of a network itself.

Consider, for a moment, a fast Ethernet LAN segment carrying data at 100Mbps. If servers are connected to it via a switch – rather than using a traditional hub-based architecture – it's not uncommon to find IDS sensors that will struggle to intercept the high, aggregated traffic volumes that can easily be generated where multiple systems are communicating at hundred megabit per second rates.

This property is called throughput and is vitally important when introducing intrusion detection into a busy intranet environment. It's clearly unacceptable to deploy products which will drop packets during peak periods and thereby fail to identify potential attacks. Look for IDS solutions, therefore, that can support sustained throughputs of several hundred Mbps or, alternatively, which

are capable of connecting directly onto network switch backplanes.

Throughput is also an issue when considering host-based intrusion detection. Bear in mind that levels of activity on servers may increase significantly during an attack and it's vital to confirm that an IDS sensor will keep pace with anticipated peak traffic levels. It's also important, of course, to establish that the performance of a monitored system will not be impacted by the intrusion detection processes running on the server itself.

A similar problem can arise where the combined monitoring traffic from large numbers of network sensors begin to represent an unacceptable proportion of total network activity. For these type of implementations, identify IDS solutions where sensors can be cascaded together and which are capable of consolidating monitoring traffic in order to reduce the overhead introduced to a network.

Scalable solutions

In practice, it's not just large companies that will be faced with deployments of significant numbers of sensors. Today's eCommerce environments frequently comprise multiple subnets with multiple access points. There are likely to be additional back-end connections into existing infrastructure or third party links too. And, of course, host systems themselves have to be considered.

Any IDS solution, then, must afford adequate scalability. This is measured not only in the number of separate sensors or agents that can be supported, but also in the level of work each agent – and the IDS solution as a whole – needs to accomplish in order to correctly identify and analyse what is happening across a total IT infrastructure.

Operations staff need to be presented with a clear picture of suspicious activities that have been identified anywhere across their domain. Having to wade through vast quantities of individual system logs, traffic reports or other sources of activity information is not a viable option.

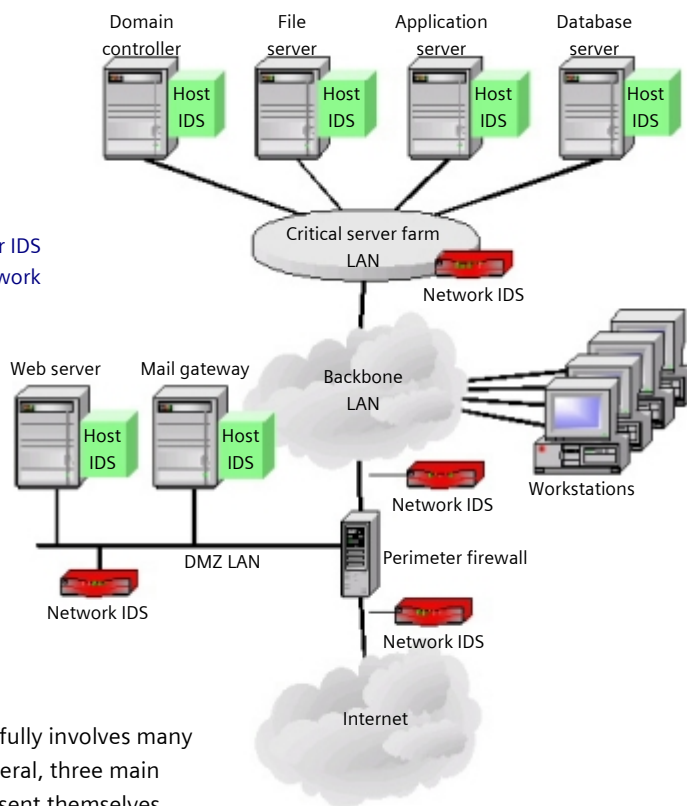
Furthermore, a centralised reporting capability needs to be backed up with supporting information to enable the

nature, source, and potential implications of an attack to be quickly established. If a suspect connection has been detected, for example, it may be useful to examine the information that's been transmitted over that connection. If files containing sensitive or confidential information have been accessed, it's clearly useful to identify the transgressors. Similarly, where the perimeter of a network is being repeatedly probed, discovering the source of this traffic can be invaluable for subsequent investigative or evidential purposes.

type of sensors required – and their location – can certainly be established.

The second issue relates to the actual IDS products to be used. Although a combination of different supplier's systems may work, a fully integrated solution encompassing network- and host-based sensors from a single vendor will generally be advantageous. When a server comes under attack, for instance, being able to easily integrate evidence from network sensors with that from host agents will inevitably help to corroborate the best course of corrective action to take.

Possible locations for IDS sensors in a typical network



One supplier or two

Deploying IDS successfully involves many challenges but, in general, three main issues are likely to present themselves during this phase of a project. The first is how, and where, to deploy individual sensors. Should they be implemented at the perimeter only or internally as well. Are they required inside or outside a firewall. Is it necessary to add them to all servers or just critical or externally-facing ones.

Unfortunately, there are no easy answers to these questions. Each organisation's IT environment and business strategy needs to be individually assessed in order to correctly identify the threats faced and the risks involved. From this assessment, though, the number and

Bear in mind, too, that different vendor's technology will often work in dissimilar ways and one will often be more suited to working with an organisation's existing investment in technology than another. Sensors which are software-based, for example, may demand an operating environment that's incompatible with a company's IT strategy and be unfamiliar to staff. Selecting appliance-based IDS solutions can address this type of concern.

Thirdly, it's important to assess the ease with which a chosen IDS solution will

integrate with existing network or systems management suites. When attacks are detected, speed of response is of the essence and seamless integration with established problem reporting and escalation procedures will be vital in minimising the risk to normal operations.

Delivering the benefit

Simply implementing IDS 'out of the box' will rarely achieve the promised benefits. In fact, without correct tuning and optimisation, an IDS solution will frequently introduce more problems than it has been designed to resolve. A common reason for this relates to false positives and false negatives.

A false positive occurs when an IDS triggers an alarm in response to a sequence of events it believes represent an attack but which actually isn't. A common example is a user who has forgotten their password and who is trying several possible candidates in an attempt to log on. The IDS may detect this as an attempt to guess a password – as indeed it is – and initiate a chain of events which would lead to operations and security staff becoming involved to investigate the perceived incident.

Conversely, a false negative occurs when an IDS misses an attack. Citing a more technical example, a common early sign of an attack is a port scan. This involves an attacker making a large number of connection attempts to establish which network services a system, such as a web server, will respond to. It's a widespread technique that's used to identify potential vulnerabilities in systems.

Most IDS solutions will track this type of activity by either looking for a large number of connections occurring during a set time interval or a series of connections over a range of numeric service ports. A shrewd attacker, however, might slow down their connection attempts to avoid breaching the IDS threshold or, alternatively, randomise the order of the ports scanned in order to disguise their activity.

Time spent benchmarking an IT environment and measuring the levels of activity that can trigger IDS alarms is

highly recommended. These findings can then be analysed and reflected in the configuration of specific sensors.

Adopting these types of actions will help to minimise false positives as well as maximising the accuracy and performance of an IDS alerting function. It's essential that operations staff develop faith in new security initiatives and a continual stream of false alarms will very quickly result in lost confidence and a wasted IDS investment.

A planned approach to management

It should be evident by now that intrusion detection is not just a product - it's a process. Ongoing management is central to achieving the proposed benefits from an IDS solution. For many organisations, though, running and managing intrusion detection systems 24 hours a day isn't practical and a combination of operational coverage during office hours, with some form of automated response during overnight and weekend periods, may be an acceptable compromise.

Automated response should be used with care, however, and may not be appropriate for every situation. Automatic actions may often involve closing suspect connections or blocking access from certain sources. Both of these activities could, of course, result in denying service to legitimate users. In a small environment where round the clock manning is not possible, however, they can often represent an acceptable first level of defensive response.

Others organisations, whilst being able to justify 7x24 operations, may lack the necessary levels of expertise to correctly interpret and respond to every form of attack. Using the services of a Managed Security Provider (MSP) can often provide a workable and cost-effective solution in this scenario.

In all cases, though, the information provided by an IDS solution must be used in conjunction with an agreed incident handling strategy. Operations staff – whether within the organisation itself or employed by an MSP – will then clearly understand the actions to take in response to specific alerts from an IDS. It may be appropriate, for example, to leave systems

running to gather further evidence of an attack. In other cases, closing systems down to prevent loss or compromising of data may be a safer alternative.

All such activities should be defined in an incident handling strategy together with procedures for further investigation, incident escalation and, importantly, measures for preventing – or minimising the impact – of similar, future attacks.

Finally, as well as incident management, IDS solutions should also include documented procedures for updating sensor attack signatures, implementing policy changes or modifying sensor configurations. Clearly, in a large, dispersed environment, selecting IDS products where these functions can be completed rapidly from a central point, will provide not only a more cost effective solution but a far more secure one also.

