

Technical Briefing Paper

Title: Considerations for selecting and implementing an Intrusion Detection System

Synopsis: This paper presents the organisational and technical considerations that need to be addressed prior to implementing IDS.

Why choose an IDS?

Intrusion detection systems (IDS) should be viewed as an additional tool in the continuing job of maintaining the security of a corporate system. IDS are complementary to the use of firewalls and effective security policy within an organisation.

IDS will respond to threats such as internal attacks, external attacks using legitimate routes and exploiting allowed rules in firewalls that other host and network security products are unable to counter.

The main capabilities that IDS offer are:

- They increase the overall security of your environment.
- They allow monitoring of the network traffic inside your firewalls.
- They allow examination the contents of information packets transferred into and across the network, thus detecting for example, "buffer overflow" types of attack.
- They can detect changes to files and directories on client and server machines within the network.
- They can detect irregular access times. For example a user may only be expected on the network in office hours, an IDS would detect the unauthorised use of the network by that user out of hours.
- IDS can also detect "below-the-noise" type attacks that slowly scan a network perhaps over a period of weeks or even months.

Selection considerations

There are a large variety of intrusion detection systems available, suitable for almost any circumstance. They range from freeware versions to commercial systems costing many thousands of pounds. Some are designed to monitor whole networks, whilst others are deployed on a per-machine basis. All these systems have their particular strengths and weaknesses and are generally best suited for a particular circumstance.

Types of IDS

One of the main distinctions between types of Intrusion Detection Systems is whether they are network based or host-based systems. These perform quite different functions but complementary and it is not unusual to see both types of system deployed.

Network-based systems examine the individual packets flowing through a network. Unlike firewalls, which typically only look at IP addresses, ports and ICMP types, network based intrusion detection systems ("NIDS") are able to understand all the different flags and options that can exist within a network packet. A NIDS is therefore able to detect maliciously crafted packets that are designed to be overlooked by a firewall's relatively simplistic filtering rules. Hackers often craft such traffic in order to "map out" a network, as a form of pre-attack reconnaissance.

Whilst network-based IDS look at all the traffic flowing by on your network, **host-based intrusion detection systems** are concerned with what is happening on each individual computer or "host". They are able to detect such things as repeated failed access attempts or changes to critical system files.

An **application based IDS** is a host-based system that is particular to a type of service (e.g. an IDS built particular for a web-server or mail server).

A **target based IDS system** is built to check the integrity of a particular system and its onboard software including operating system software. Target based systems are often called file-integrity assessments since they use check-sum based software to determine whether a system has been tampered with.

Misuse and anomalous activity detection

A second level of categorisation of intrusion detection systems is between those based on the detection of misuse and those based on the detection of anomalous use.

Misuse detection within a network-based intrusion detection system typically involves checking for illegal types of network traffic, for instance combinations of options within a network packet that should never legitimately occur. Misuse detection by a host-based intrusion detection system would include attempts by a user to execute programs for which they have no legitimate need.

Detection of **anomalous activity** relies on the system knowing what is "regular" network traffic, and thus what isn't. Anomalous traffic to a host-based intrusion detection system might be interactive accesses outside of normal office hours. An example of anomalous traffic on a network-based IDS would be repeated attempted access by one remote machine to many diverse services on one or more of your internal systems, all in quick succession. This is indicative of someone conducting a "port scan" of your systems.

Many modern systems use a combination of both misuse and anomalous detection engines.

Passive and Reactive

A further method of categorising intrusion detection systems is by their passive or reactive nature.

Passive systems simply detect the potential security breach, log the information and raise an alert. This system relies on a well-trusted incident response routine being in place to take appropriate action on alerts. This is much better than a pre-programmed response using an inflexible set of rules, however can mean unseen costs for engineer call outs to false alarms.

Reactive systems are designed to respond to the illegal activity, for example by logging off a user or by reprogramming a screening router to disallow network traffic from a suspected hostile source. Whilst a reactive system might seem like an ideal solution there are serious drawbacks to such systems as they can potentially take actions that could shut down business critical services.

Implementation Considerations

What are the information assets to be protected?

The first step towards determining the most suitable balance of intrusion detection assets on a network is to understand the underlying network upon which the IDS will operate.

This is generally achieved from discussions with system owners and administrators, supported with up-to-date network architecture diagrams. Gathering network information on systems may form part of the implementation task.

It is best practice prior to the installation of IDS to ensure that the network is as secure as possible. This may involve network security health checking with host vulnerability scans to ensure that there is no unauthorised hardware/software on the network.

Ensuring appropriate physical coverage of the IDS

Generally, IDS is targeted at critical infrastructure assets such as servers and domain controllers that provide essential services, hence an assessment of network resources is often required to determine what is vital to the organisation. The system owner and administrators generate the "critical asset list".

As well as targeting particular aspects of intrusion detection towards critical assets, it is also important to ensure that IDS physically covers all appropriate networks (for example parts of the network may be cryptographically separated using VPN technology, other parts may be separated by switch technology). Assessments of those network domains, IT systems, and hosts found on respective networks should be made to ensure that all aspects are covered.

What are the perceived risks?

Each network has its own critical functions, and supplies different business needs. For example, an internal network may supply an organisation's corporate intranet, whereas there may be a smaller isolated network that

supplies web services (e.g. http, https, and ftp). However, it is more likely that a firewall could connect internet to an internal network. Each network will be faced with a range of threats, some more potent, some held at bay with other functionality (e.g. firewalls) but not exclusively shut out for reasons from attacker skill to flaky software configurations. Threat examples could include:

- Closed internal network that may have malicious insiders installing "sniffer" software to detect other users passwords,
- Opened to the internet may have "script-kiddies" running scanning software against particular web servers looking for unpatched software holes, that can subsequently be exploited so that yet another site is defaced.

A detailed threat assessment should be made with support from system owners and administrators based on the network's operational environment (including aspects such as Internet access, levels of required employee trust, operating systems and network systems used, and hardware utilised).

A threat assessment is relatively ineffective alone, since it does not help focus IDS resources particularly well. Hence a business impact analysis is made that looks at particular failures in confidentiality, integrity, availability, non-repudiation and authentication. The business impact analysis investigates how each particular failure adversely changes business process, and suggests countermeasures that, in this case, include intrusion detection.

The threat assessment and business impact analysis would be done by interview with key stakeholders including system users, administrators and owners.

Integrating with existing security infrastructure

It is important to understand what existing security is in place on the network. For example, there may be policies regarding the installation of new machines and limitations on software installed, there may be traffic analysis already being carried out, there may be a configured firewall, there may even be network virus software and web/e-mail monitoring software installed on servers.

It is important to co-ordinate and manage the installation of an IDS so that its effect is maximised, so that it doesn't clash (e.g. provides constant false positives caused by administrative port scanning), and so that the IDS policy of use fits well with the overall security policy.

Integrating with existing security policy

It is usual for an organisation to have a detailed security policy. However it is likely that this is a high level document, detailing the user "do's and don'ts" on the internal network and accessing the Internet. In order to successfully select and install an IDS system it is necessary to understand how IDS will support this policy. To implement IDS requires policy to extend down into high levels of technical detail – such that particular events may be blocked or recorded.

The information developed in this detailed technical policy is important since it may dictate to some extent the nature and way that the IDS solution is implemented. Conversely, it may be necessary at this point to document and propose technical level policy and procedure. For example, defining and documenting network protocol use and supported traffic flows.

In some cases it may even be necessary to propose appropriate changes to security policy to maximise the effectiveness of any suggested IDS solution.

The way forward

What are the steps to take towards selection and implementation of IDS?

- Mapping functional capability to requirements
- Product evaluation and selection
- Product customisation
- Incident response

There are a variety of products on the market, together with some extremely capable freeware offerings. Different IDS products offer different functionality. Table 1 (overpage) shows the range of different approaches to IDS and the security roles they cover.

Product evaluation ultimately resulting in IDS selection. Many IDS manufacturers (Appendix A provides a small sample list of some commercial and freeware IDS) offer evaluation periods where the system can be tested free. It is highly advisable to take this offer up, even if it does mean putting some manpower effort in to determine whether the product is suitable – it is better to do this than to have an unworkable solution sat doing nothing after a year.

Product customisation is a process that passes from installation and implementation into the normal operation phase of the life of the IDS. This is because IDS takes time to “know” what normal means on your network. For example high amounts of network management may be misinterpreted as information gathering, network audits may be interpreted as network attacks. User applications may make network requests may be interpreted as malicious. On the other hand, it may be necessary to add new signatures since the default installation of the IDS system may not provide operational coverage of particular application level processes that are deemed business critical – for example unauthorised use of banking connections that use bespoke banking network software.

Incident response functions within the organisation should be augmented to include reaction to alerts from IDS. If there is no incident response functionality within an organisation careful consideration should be made as to the reasoning for owning an IDS since it is only as good as the responses made to the alerts generated. However the installation of IDS is an ideal time to consider broader incident response, perhaps encompassing disaster recovery.

Incident response does not necessarily mean calling in the troops every time an alarm turns red. Many systems allow automated response to certain

threat signatures (e.g. IP address blocking), others may not be time critical and can be dealt with inside office hours – however the IDS system should be finely tuned so that network operations will only respond to time critical threats and attacks.

System properties	Nature of System	IDS type				Intrusion analysis method		
		Application based	Host based	Target based	Network based	Real-time	Signature assessment	Integrity assessment
Security Functionality	Description							
<i>Confidentiality</i>	File/system resource access		✓				✓	
	Policy violation	✓	✓	✓	✓		✓	✓
	Password strength	✓	✓				✓	
<i>Integrity</i>	Malicious software insertion		✓	✓		✓	✓	✓
	Malicious software detection			✓			✓	✓
	Network service attacks				✓		✓	
	Web server-based attacks	✓		✓			✓	
<i>Availability</i>	DoS				✓		✓	
	Firewall failure/ misconfiguration	✓			✓		✓	
	Attacks over encrypted networks	✓	✓				✓	
	Unusual activity detection	✓	✓	✓	✓	✓		
<i>Miscellaneous</i>	System/network error detection		✓				✓	
	Liability exposure for own IT attacking others	✓	✓	✓	✓	✓	✓	✓
	Post-hoc damage analysis	✓	✓	✓	✓	✓	✓	✓
Capabilities	Smaller management overhead				✓			
	Packet-based analysis			✓	✓			
	Prevention of evidence removal			✓	✓			
	Speed of detection			✓	✓			
	Malicious intent detection				✓			
	Complements other forms of security				✓			
	Operating system independence				✓			
	Verification of attack in progress	✓	✓					
	System specific activity	✓	✓					
	Coverage benefits	✓	✓					
	Monitoring capability at key points	✓	✓					
	No need for addition hardware	✓	✓	✓				

Table 1: A table of IDS functionality versus IDS type and analytic behaviour (greyed out=not applicable)

Appendix A: Example IDS Products

Some product examples are given below. (The omission or inclusion of a product on this list should not be seen as any form of [lack of] endorsement for that product.)

- Cisco Secure Intrusion Detection System – <http://www.cisco.com/warp/public/44/jump/secure.shtml>

This product is a network based IDS. The Cisco Secure IDS is a real-time intrusion detection security system that includes the Cisco Secure IDS Sensors, security appliances that act as "sniffers," and the Cisco Secure IDS Director, a centralized management console.

- Symantec/Axent Intruder Alert (ITA) & NetProwler- <http://enterprisesecurity.symantec.com/content/productlink.cfm>

ITA is a host based IDS which complements and integrates with NetProwler a real-time network IDS. Both systems can be managed from a single console.

- ISS RealSecure – http://www.iss.net/securing_e-business/security_products/intrusion_detection/

This product has a modular architecture which includes both network and host-based intrusion detection capabilities. RealSecure has a management console which is used to manage all its IDS plug-in modules

- SNORT- <http://www.snort.org>

SNORT is a freeware network based IDS which is available for a large number of platforms. Being freeware there are also a large number of complimentary freeware contributions available.

- Tripwire – <http://www.tripwire.com>

Tripwire is a filesystem integrity checker, a utility that compares properties of designated files and directories against information stored in a previously generated database. It has additional complimentary products Tripwire HQ Manager and HQ Connector. HQ Manager is a software console that allows multiple installations of Tripwire to be managed and HQ connector provides the communication between the two.

Tripwire software has several applications: intrusion detection, software verification, policy compliance, forensics, damage assessment and recovery, and auditing.