



NISCC Technical note 10/03

Issued 19 November 2003

Deployment Guidance for Intrusion Detection Systems

Key Points

- Factors to take into account when considering deploying Intrusion Detection Systems (IDS)
- Best use of the available sensor types
- Configuration and tuning of IDS
- Maintenance issues
- Implications for the resources required to investigate and act on the outputs from IDS

**National Infrastructure
Security Co-Ordination Centre**
PO Box 832
London, SW1P 1BG

Tel: 020 7821 1330 Ext 4511
Fax: 020 7821 1686
Email: enquiries@niscc.gov.uk
Web: www.niscc.gov.uk

NISCC Technical Note 10/03 : Deployment Guidance for Intrusion Detection Systems

Aim

1. This NISCC Technical Note supplements NISCC Technical Note 09/03. Its purpose is to assist people who are thinking about implementing, upgrading or procuring an Intrusion Detection System (IDS) (be it self-managed or part of a managed service), by providing the basic knowledge for them to enter into discussions with vendors and managed service providers.
2. This note discusses issues applicable to all IDS installations, and does not deal with particular products, services or technologies.
3. A comprehensive glossary explaining all the terminology used is available at <http://www.securityfocus.com/infocus/1728>

Summary

4. This note discusses:
 - a. the need to understand what needs to be protected, and from whom, so that the system meets the security requirement and is also cost effective;
 - b. the pros and cons of network and host-based sensors, deployed in various configurations, to reduce costs and maximise coverage of critical business services;
 - c. the integration of the IDS network with the existing infrastructure, by the use of in-band or out-of-band reporting, and how this affects the overall security stance of the organisation;
 - d. how to tune an IDS system to tailor it to the environment, the pitfalls of tuning and the possible consequences;
 - e. how to configure an IDS to its environment to make it more efficient and produce fewer false positives;
 - f. the use of Intrusion Prevention Systems to mitigate some of the risks posed by Internet attackers, and the possible side-effects;
 - g. the need to keep the IDS and associated rules and signatures up to date and
 - h. the need to invest the necessary human resources in the IDS installation to ensure that alerts are received and acted upon promptly.

Introduction

5. First ask yourself what you are trying to protect, and secondly, from whom.
 - a. To answer the first question, try to ignore the physical aspects of the network and consider the business processes that are important to you, such as payroll, public services and enterprise data storage. Assign a nominal level of sensitivity: which would be worse - a web page defacement or a compromise of your accounting systems? Which could you most easily recover from? Could there be a significant impact on share value and business image as a result of a web page defacement?
 - b. To answer the second question, you have to consider who is likely to attack those systems. Is an employee or an external attacker most likely to be interested in the payroll system? Or in your Internet presence?
6. Keeping these thoughts in-mind during the planning of an IDS installation will almost certainly lead to a more suitable and more secure deployment.

Sensor Placement

7. Your IDS budget will only go so far. Each sensor that you deploy will have not only a procurement and a support cost, but an associated staffing cost. There are strong arguments in favour of a few well placed, well managed IDS sensors in preference to a large number of poorly managed sensors. Should budget constraints prevent complete system coverage, consider deploying a small number of sensors to the most critical areas of your network and have a small number of sensors that can be moved around the other important areas on a monthly/quarterly basis. It won't give you the full coverage, but will at least cover the entire network in sections over time.
8. Having identified the services your network offers it should now be a simpler task to identify the priority of deploying IDS in certain areas of the physical network. The choice of sensor, and its positioning, will depend on a number of factors that will differ with each network.

Network Sensors

9. Network sensors are good for monitoring a large number of hosts with a single sensor. When networks were primarily built with hubs it was possible to monitor traffic to, from and between all the hosts on the hub with a single sensor. With the advent of switched infrastructures this became more difficult. If considering this type of sensor deployment, it is important to identify the capabilities of the switches used in your organisation. Many have some form of 'port spanning' that can turn one of the switch ports into a 'sniffing' port, but few can guarantee that every packet seen by the switch will be seen on this port, especially under high load. The resulting data aggregation may also overload the sensor and result in packets being dropped.
10. Another option is to deploy the network sensors on the connections between network segments, for example between the workstation switches and the server switches or between the De-Militarised Zone (DMZ) and the internal network. A potential weakness of this configuration is that communications between hosts on the same network switch will not be detected by the sensor. However, installation may be simplified. The chosen option should depend on the assessments made earlier for sensitive services and potential attackers.

11. Network taps can be employed for this purpose or a simple hub can be installed with a feed to the sensor. Some more recent network IDS systems are run in-line and can be used to block potential attacks from traversing from one network segment to another.

12. You need to decide on the performance you need from your sensors. Most vendors now advertise sensors capable of monitoring traffic at speeds over 100Mbit/s; but this will be in ideal conditions. The best approach is to ask for a demonstration of the sensor running on your network at the normal daily peak, and remember that your bandwidth utilisation may go up.

13. Network sensors are not all good news. They will generally be blind to traffic that is encrypted (e.g. SSH or HTTPS) and they can easily be provoked into producing lots of false positives. There are tools that produce packets containing the very things that IDSs look for, for example, '*IDS wakeup*' is a suite of tools for testing IDSs. However, it can also be used to produce a huge number of alerts, which may reduce the performance of the IDS and could hide a real attack.

14. If your network has any connections to external organisations it is an excellent idea to monitor this connection with a network sensor. This not only protects your network from potential attacks from the other organisation's network, but will also inform you of outbound attacks. Before installing a network IDS in an environment that may contain data from other organisations, it is essential to gain consent and check that the appropriate agreements have been made with the data owners.

15. A network sensor placed outside the firewall will give an indication of the threat being mitigated by the firewall but it will do little else. Do not allow the number of alerts it produces to cloud the real issue, which is security inside the firewall. In the last year there have been a number of vulnerabilities discovered in IDS sensors. If the management network for the sensor bypasses the firewall, this could allow an attacker to gain access to your critical networks. If there is a strong case for having a sensor outside the firewall, it should be isolated from the internal IDS network.

Host Sensors

16. Host sensors work by monitoring the connections to and from the host, and the logs generated and system calls made by running programs. Host sensors can see what is happening on the host and are therefore more likely to identify attacks (including those committed locally) and are generally less likely to give false alarms. However, a single host sensor can only monitor a single host, becoming costly for large numbers of hosts, and sensors may not be available for all the architectures used in your organisation.

17. Due to the resource overheads, the host sensor may adversely affect services running on the host, although this is becoming less of a problem as sensors become more efficient. For workstations, simple personal firewalls will often be a more cost-effective substitute, especially if these can be managed centrally.

Simple Network Example

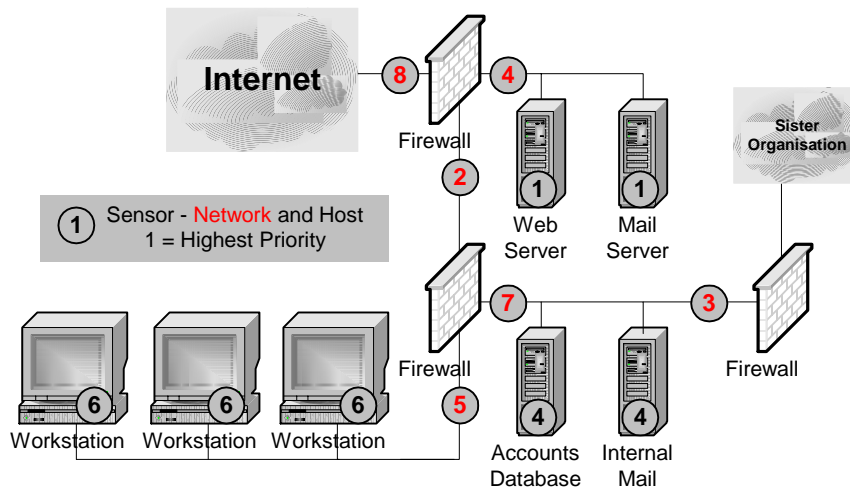


Figure 1. Simple network showing priorities for deploying sensors

18. Fig 1 shows a simple network of a type likely to be used by a small organisation offering some web services. The diagram shows host and network sensors placed at all applicable points.

19. For this small organisation connected to the Internet the greatest threat is considered to come from external attackers and Internet worms. The company is extremely worried about web page defacements and its public image. The reasoning behind the assigned priorities is given in Table 1.

Priority	Sensor	Reasoning
1	Host	The high Internet threat, combined with the concern for public image, make the Internet facing servers a priority. A network sensor could be deployed, but a host based sensor is more likely to detect attack in encrypted traffic such as HTTPS.
2	Network	To protect the remaining network, a network sensor is deployed to monitor all traffic entering the internal networks from the Internet.
3	Network	The only remaining connection out of the network is then monitored with a further network sensor. This sensor will not only identify incoming attacks, but will also alert if a network is being used to attack the sister organisation.
4	Host	At this stage there may be no way of detecting internal activity. As the internal servers contain the most valuable information, host sensors are deployed to protect them. In many organisations the priority of internal servers may be even greater than that of Internet facing servers.
5	Network	A network sensor is deployed to monitor traffic to/from the workstations. This is the most low cost way of protecting the workstations and may also detect any attempts to launch attacks from them against other parts of the network.
6	Host	The workstations themselves are then protected with host sensors offering a far greater level of security.
7	Network	This network sensor is almost redundant as all traffic visible to it should either be seen by sensor 2 or 5.
8	Network	The final sensor to be deployed is a network sensor outside the firewall. The organisation decided that it did not have the resources or skills required to constantly monitor the alerts from this sensor, but considered that it might be useful for identifying portscanning activity or problems such as DoS (although this information can generally be obtained by analysis of firewall logs).

Sensor priority reasoning

Table 1.

20. Each organisation will have a different view of the nature of the threat, and the priorities will change accordingly.

IDS Communication

21. Connection of the IDS sensors to the management console can be done either by utilising the existing network infrastructure (in-band) or by installing a secondary network for IDS communications (out-of-band).

In Band

22. The obvious advantage of in-band communications is cost. For network sensors the network that is being monitored is also used to send alerts and receive updates. For host sensors the network connection of the host computer is used. If communications are to be sent between networks that are separated by firewalls, additional rules will need to be created to allow the passage of the relevant traffic.

23. One downside of in-band communications is the potential for an attacker to use the IDS to disrupt the network. For instance, if an attacker was to send a large number of connections to a public web server that contains a known exploit, for every connection made, the sensors monitoring on the route of the network packets would create an alarm. This multiplication in network traffic can lead to a reduction in available bandwidth, and potentially even a denial of service.

Out of Band

24. For out-of-band communications, a separate IDS network must be created that connects all the sensors to the management console. In general, network sensors will be connected to the network being monitored by taps and the sensor will have a second network interface connected to the IDS network, thus completely isolating the IDS from the production network.

25. This architecture is best used in critical networks, where it is important that the IDS should have no effect on the performance of the network. As discussed previously, a badly implemented out-of-band network could be an additional security risk. If the sensors are not connected to the network with one-way taps, an attacker could compromise a sensor and gain access to the IDS network. In this scenario, it may be possible to bypass internal firewalls and gain access to critical servers. Making sure the network sensors are hardened on the IDS network interface can also reduce this risk.

Tuning

26. The secret of a good IDS installation is careful tuning. This is the process of re-classifying as legitimate the events that create false positives. Out of the box, most IDSs will be programmed with signatures that will cause alerts when they see traffic that may be considered legitimate in your organisation. A good example of this is the ICMP echo reply/request commonly used by the tool 'ping' for network discovery. Used internally, 'ping' can be extremely useful for finding network problems, however, it may also be used by an attacker to map your internal network from afar. Tuning the IDS sensors so that ICMP echo messages between internal machines are considered benign will reduce the number of false positives while retaining the ability to alert when Internet sourced messages are seen in the internal network.

27. Another form of tuning is where an IDS alerts because an exploit was attempted that is not relevant for an architecture (e.g. an IIS buffer overflow against an Apache web server) or one that requires a previously patched vulnerability to be present (Code Red etc).

28. In both cases it is quite simple to remove the rule from all the sensors and thus remove the problem. However, a more sensible approach is to carry out tuning on a sensor by sensor basis. Tuning the sensor that sees the web server traffic to log rather than alert when it sees the Apache exploit or Code Red will preserve the ability to identify the same attacks elsewhere in the network.

29. Over-tuning an IDS will reduce its effectiveness. It is imperative that no tuning is undertaken without full knowledge of what has been observed, what that means to the network and what mitigating steps have been taken to ensure there is no risk from such an attack. When tuning is the only option it should be as selective as possible and the reasons for doing so should be documented.

30. Tuning will be an ongoing activity. If a rule is initially too generic and produces too many false alarms it may be revised by the IDS vendor. It can be a good idea periodically to reinstate previous versions of revised rules and check that activity generating false alarms is still ongoing. When rule updates are made, the improved rules may not be activated if the previous version was excluded from the active rule set.

31. False alarms should be far less common on internal network segments, but once again careful tuning is required, especially with respect to Internet services such as web browsing.

Configuration

32. Most IDSs allow the user to enter substantial network information to indicate which servers are offering which services. This information can be used by the IDS to decide the severity of any alerts produced and also to improve performance by only looking for relevant activity. For example if the IDS is aware which servers offer HTTP services and on which port, the IDS needs only look at traffic to and from those servers/port for HTTP-based exploits.

33. Some more advanced IDSs can also use this information to identify anomalous activity such as a connection to TCP port 80 on a machine that is not declared as a web server.

Intrusion Prevention

34. Almost all current IDSs offer some form of active countermeasures even if they are not specifically called Intrusion Prevention Systems (IPS). The most common method is to attempt to break the communication between the attacker and the target by sending TCP reset packets to either or both of them. This method can be extremely useful, but has the limitation of not being able to deal sufficiently well with UDP traffic and not at all with single packet exploits such as the Slammer worm.

35. Some IDSs can be integrated with firewalls to block all traffic from the source of an attack.

36. With ever higher specification machines it has now become feasible to produce a device that can act as an in-depth firewall. These devices, often called In-Line IDS, will inspect every packet for dubious content before allowing them to pass. Whereas firewalls generally work at the transport layer of the OSI model, In-Line IDS can work right up to the application level.

37. As discussed earlier, any form of intrusion prevention must be implemented carefully so as not to allow attackers to utilise this functionality for their benefit. All IPSs have the ability to be used to deny service to legitimate users. By spoofing legitimate IP addresses, attackers can launch attacks that appear to come from innocent hosts. If the IPS system is configured to deny access to subsequent connections from that host, the attacker has effectively denied service to the host.

Signature updates

38. At present most IDSs are primarily signature or rule-based, meaning that if they see an event or sequence of events they produce an alert. For an IDS to be effective, these signatures and rules need to be updated regularly and it is essential that a method for doing so is considered. This can be a problem if the IDS vendor only supports Internet updates, but the IDS network has no connection to the Internet. This will be particularly relevant to networks that may contain sensitive material.

39. Anomaly-based IDSs may still require periodic updates, depending on the way that anomalies are detected, and the same issues apply.

Human Costs

40. The most serious consideration in any IDS installation should be that of who undertakes the monitoring of logs and alerts. If you have no intention of monitoring IDS alerts or taking action on them, you should secure your network in other ways.

41. As an IDS's primary role is to identify nefarious activity on the network it seems reasonable that the best people to monitor the system are members of the security team. In many organisations the securing of a network is undertaken by the system and network administrators. By having a different team monitor the IDS, you can add defence in depth to your networks. This also guarantees that the security team is getting their information first hand and promptly.

42. If it is not possible to arrange for the IDS system to be monitored by your own staff, for example if it is purchased as part of a managed service, it is essential to build monitoring requirements into the agreement with your service provider. This should specify who will be informed when alerts are generated and state timescales for reporting in and out of working hours. You should also require weekly or monthly reports that detail any events and any tuning that has been done.

Legal Issues

43. When deploying any IDS sensor it is essential to understand what data is going to be monitored and recorded. All users of the system must be asked to sign an agreement that defines appropriate usage of the system and gives permission for their communications to be monitored for lawful purposes.

44. If connections to other organisations exist, an agreement should be drawn up and signed by the appropriate persons at the other organisation that authorises the monitoring of all communications between the networks.