

Background

This NISCC technical note is intended as an introduction to Intrusion Detection Systems (IDSs), but also explores possible security features of IDSs and their current state of implementation. It is vendor-independent and does not discuss features available on particular IDSs. The intention of the paper is to increase technical awareness about IDSs and to enable potential purchasers in determining the security features of IDSs most appropriate to their networks. The note is aimed at managers and security officers in the organisation who wish to inform their decisions on the choice of IDSs that available, and the impact of those choices. A checklist is provided at the end of this note.

Summary

1. Given the options available, it is clear that there are a large number of possible types of IDSs, some more useful than others at detecting intrusions from particular classes of attackers. Each combination of options will have its own advantages and disadvantages. Currently most commercial IDSs are signature based because the accuracy of anomaly detection systems is not sufficiently high given the additional complexities associated with implementing anomaly detection. Statistical inference techniques in misuse detection are not yet widely deployed in commercial systems.
2. A current problem with many network based and network node IDSs is the fact that they do not capture (sufficient) packet data in order to corroborate an alert. One solution would be to implement a circular buffer in memory that is written to disk when an alert is issued by the IDS.
3. Nevertheless, despite the lack of maturity of intrusion detection and difficulties with corroboration, there a number of commercial and open source intrusion detection systems currently available, as well as a smaller number of anomaly detection systems. Provided that it is accepted that the IDS will not capture all intrusions, that it may interpret legitimate traffic as an attack, and that it will require attention (training or tuning) to maintain its usefulness, then IDSs can be an important part of an organisation's computer security strategy.
4. A useful measure for testing the effectiveness of the IDS is the percentage of attacks it identifies and misidentifies from traffic that contains a known number of attacks, such as a scan from a vulnerability scanner. Massachusetts Institute of Technology (MIT) have collated a sample based on real network data and have performed a bake off of IDSs, eg [1].

Introduction

1. This note is not intended to be exhaustive in its coverage. The following URLs link to papers that provide alternative or more detailed coverage:

- http://www.sans.org/newlook/resources/IDFAQ/ID_FAQ.htm
- <http://csrc.nist.gov/publications/nistpubs/800-31/sp800-31.pdf>
- <http://online.securityfocus.com/infocus/1524>
- <http://online.securityfocus.com/infocus/1534>
- <http://online.securityfocus.com/infocus/1544>
- <http://online.securityfocus.com/infocus/1553>
- <http://online.securityfocus.com/infocus/1569>

2. It is possible to define an intrusion as "any set of actions that compromise the integrity, confidentiality or availability" [1] of a resource. This definition captures the notion of an intrusion as being a deliberate security violation, although it is arguable that unintentional unauthorised access would also count as an intrusion. In the context of this note, intrusion detection is the detection of intrusions against computer systems.

3. An IDS is like burglar alarm for computer systems and networks. The function of an IDS is to detect intrusions and alert system administrators. It is not the purpose of an IDS to stop attacks: that is the function of a firewall. Instead IDSs provide a check that the firewall is enforcing the organisational security policy, detect internal misuse and correlate attacks over a number of different sources and locations. It should be stressed that intrusion detection is currently not a mature technology. There is an ongoing research effort in the academic community into intrusion detection. In practice this means that the burglar alarm may alert when there is no intrusion and sometimes it will not alert when there is an intrusion.

4. The scope of this note is restricted to intrusion of detection. Alternative approaches are limiting the exposure of the system to attack and limiting the impact of attacks. Building dependable systems and network infrastructures are the subject of ongoing research [2]. This research will not be discussed further in this note.

5. This note aims to help with understanding the types of intrusion detection systems that have so far been implemented in commercial or research systems. There are three sorts of classification, namely by:

- Techniques used
- Domain of applicability
- Type of attacker

Techniques Used

6. Intrusion detection systems come in two broad categories: those that search for predefined patterns and those search for unusual or anomalous patterns. The former type is commonly called *misuse detection*, while the latter type is known as *anomaly detection*. (It should be noted that "misuse" is often associated with abuse of privilege by authorised users, but "misuse detection" has a more

Network based intrusion detection

- Host based intrusion detection
- Network node intrusion detection

Network Based Intrusion Detection

18. Network based intrusion detection is detection of attacks in network services, such as SMTP (email) or HTTP (web). Attacks of this kind are common, especially given the internet and extranet connectivity of many organisations.

19. Network based IDSs employ sensors that listen to the network segments of the network and report to a central management console which is typically used for analysis and reporting. Network sensors can be implemented on routers by means of additional cards. One sensor will be needed for each network segment if the packets are routed to the segments by a switch (unless the switch allows traffic on the same virtual local area network to be copied to a mirror Switch Port Analyser port).

20. A number of network based IDSs can block attackers' IP addresses via the network firewall. Blocking IP addresses can be effective if the IP address is involved in an attack. To implement IP blocking a low false positive rate is needed. Moreover, if IP blocking is used, attackers may be able to deny service to a legitimate IP address by spoofing hostile traffic with that IP address.

21. While network based IDSs are often effective at detecting network attacks, they cannot analyse encrypted packets and recent IDS's are prone to drop packets at rates above 70 MBits/sec. In addition, some network based IDSs do not capture all packets relevant to the attack, so that it is not possible to verify that the attack is not legitimate traffic.

Host Based Intrusion Detection

22. Host based intrusion detection analyses activity on a particular computer by analysing the operating system and application logs, generally in near real time.

23. Host based IDSs employ sensors on each computer on which intrusion detection is required, typically sending data to a central analysis and reporting console. It is advisable to install a host based IDS on each host that provides an essential service to the company or that is used to store sensitive data. Examples would be mail and web servers in the De-Militarised Zone (DMZ) and file and application servers on the internal network.

24. An interesting type of host based IDS is one which analyses the behaviour of an application (eg by forming a model of its system calls). Attacks can often be identified by unusual application behaviour, eg by a different set or different pattern of system calls.

25. Host based IDSs are very good at detecting intrusions on a particular host and for recording details of the intrusion. They do not in general contain information on the origin of the attack.

Network node intrusion detection

26. Network node intrusion detection is network intrusion detection that only analyses traffic to or from one particular host. The deployment of network node IDSs is therefore the same as for host based IDSs, but the attack types that the IDS identifies are related to network attack.

27. Because they report on network attacks relating to a host, network node IDSs can be combined with a host based IDS to provide greater information on the cause of the attack. Network node IDSs require a security policy that identifies security critical hosts on the network (eg web servers, domain controllers), and as such have a different emphasis to network based IDSs.

28. The ideal IDS would include both host based and network or network node intrusion detection.

Type of Attacker

29. The type of IDS is also determined by the type of attacker that an IDS needs to identify. Broadly there are two types of attacker:

- External attackers
- Legitimate system users

External Attackers

30. To identify attacks by external attackers, it is important to have adequate physical and personnel security measures in place and to have a firewall controlling all remote electronic access. The firewall should be configured to log all probes and unexpected traffic. A network based intrusion detection sensor could be placed on the external side of the firewall if the potential level of threat were important. It is more common, however, to place a network based intrusion detection sensor on the internal interfaces to the firewall (including any De-Militarised Zones (DMZs), where mail servers and web servers are typically located). The sensor would report detected intrusions to the management console, which would alert the system administrator.

31. It is also important to have host based operating systems sensors in place for any computers that are exposed to the external network. The reason for this is that if the network based intrusion detection sensors fail to detect an attack from the analysis of the network traffic, vulnerabilities in the network services offered could lead to arbitrary code being executed on the host offering the vulnerable service.

Legitimate System Users

32. Legitimate system users may misuse their privileges to gain access to data to which they are not permitted access, modify or corrupt data, cause the system to crash or otherwise violate the terms of the system security policy. It is much harder to counter the threat from system users than from external attackers. It is essential to have good procedures for personnel security so that the level of trust appropriate to the user can be identified. It is recommended that the system security policy permits monitoring of email, web and other network services offered by the organisation. From a technical perspective, one can then pass internal network traffic through content filtering programs in order to identify abuses of the system security policy.

33. One useful technique is to assign profiles to users. There are two ways of doing this. One way would be to define security access groups or roles and to use a host based intrusion detection system to detect violations of the security access policy. Another approach would be to define activity profiles for users or types of users. For example, individual A may usually get in at 9:00am, log in to server X, perform some word processing and go home at 5:00pm, while individual B may come in at 10:00am, log into server Y, perform some software development and go home at 8:00pm. If A were to log on to server Y at 10:00pm, this activity would be doubly anomalous in terms of time access and server accessed. For profiling of this kind host based intrusion detection sensors and/or audit logs could also be used, for which a machine learning IDS would be suitable as it could be calibrated against each user profile.

34. Network based intrusion detection systems also have a role to play in detecting misuse by legitimate system users, especially in detecting denial of service attacks, but host based IDSs tend to be more useful because the individuals will already have users accounts on the computer system.

Protecting the IDS

35. IDSs themselves are often the target of attack because it is in attackers' interest to prevent detection of attacks. For this reason the operating systems on which the IDS run, and the operating system of the data collection unit and the management console, should be hardened as far as possible. Logging of events should be piped via a point to point connection to a logging server. Furthermore, in the case of a network based IDS, it is often advisable to deploy a passive tap that listens on the network without producing network traffic of its own, thus being detectable by network monitoring software.

Checklist

36. A checklist of the issues raised in this note are listed below.

Is the main threat from outside or in? [Emphasis on network or host based IDS respectively]
Do you offer network services outside your organisations? [Emphasis on network based IDS]
Do you have a switched network? [Emphasis on network node IDS]
Is the IDS signature based?
Does the vendor supply regular signature updates? [If signature based IDS]
Can you write your own signatures? [If signature based IDS]
Does the IDS perform misuse detection or anomaly detection?
How well does the IDS perform on a well-known sample of test data?
How much tuning does the IDS require?
Can you tune out persistent false alarms? [If misuse detection IDS]
How much training does the IDS require? [If anomaly detection IDS]
How much packet information does the IDS store when it identifies an event as an incident?
Does the IDS offer host and network based intrusion detection?
Can the IDS be integrated with firewalls and routers to provide a centrally managed solution?
Can the IDS correlate across hosts? [If host based IDS]
Does the IDS block intrusions?

Notes and References

[1] Quoted from R. Heady, G. Luger, A. Maccabe and M. Servilla "The architecture of a network level intrusion detection system" Technical Report, Computer Science Department, University of New Mexico, August 1990.

[2] See, for example, the work of Carnegie Mellon University, http://www.cert.org/nav/index_purple.html, in the US and the MAFTIA project, <http://www.research.newcastle.ec.org/maftia/index.html>, in Europe .

[3] See <http://www.ciac.org/bulletins/I-117.shtml>

[4] J. Kuri, G. Navarro, L. Mé, L. Heye "A Pattern Matching Based Filter for Audit Reduction and Fast Reduction of Potential Intrusions" in *Lecture Notes in Computer Science: Recent Advances in Intrusion Detection* 1907, eds. H. Debar, L. Mé and S. Felix Wu 2000, 17-27.

[5] A. Valdes and K. Skinner "Adaptive, Model-Based Monitoring for Cyber Attack Detection" in *Lecture Notes in Computer Science: Recent Advances in Intrusion Detection* 1907, eds. H. Debar, L. Mé and S. Felix Wu 2000, 80-92.

[6] R. Lippman, J.W. Haines, D.J Fried, J. Korba and K. Das "Analysis and Results of the 1999 DARPA Off-line Intrusion Detection Evaluation" in *Lecture Notes in Computer Science: Recent Advances in Intrusion Detection* 1907, eds. H. Debar, L. Mé and S. Felix Wu 2000, 162-181.