

CERT-In Security Guideline CISG-2003-06

CERT-In

Indian Computer Emergency Response Team

Handling Computer Security Incidents

IDS – Intrusion Detection System

**Department of Information Technology
Ministry of Communications and Information Technology
Government of India**

Contents

1.	Introduction	2
2.	Purpose & Scope	2
3.	What is Intrusion detection system?	3
4.	IDS vs Firewall	3
5.	Need for IDS	4
6.	Components of an IDS	4
7.	How they work?	4
8.	Characteristics of a good IDS	5
9.	Types of IDS	5
9.1	Architecture	5
9.1.1	<i>Host-Target Co-location</i>	5
9.1.2	<i>Host-Target Separation</i>	5
9.2	Control Strategy	6
9.2.1	<i>Centralized</i>	6
9.2.2	<i>Distributed</i>	6
9.3	Timing	6
9.3.1	<i>Interval-Based (Batch Mode)</i>	6
9.3.2	<i>Real-Time(Continuous)</i>	6
9.4	Information Sources	6
9.4.1	<i>Network-Based IDSs</i>	6
9.4.2	<i>Host-Based IDSs</i>	6
9.4.3	<i>Application-Based IDSs</i>	6
9.5	IDS Analysis	7
9.5.1	<i>Misuse Detection</i>	7
9.5.2	<i>Anomaly Detection</i>	8
9.6	Response Options	8
9.6.1	<i>Active responses</i>	8
9.6.2	<i>Passive responses</i>	8
10.	Tools that complement IDSs	8
10.1	<i>Vulnerability Assessment Tools</i>	8
10.2	<i>File Integrity Checkers</i>	8
10.3	<i>Honey Pots & Padded Cells</i>	8
11.	Selection considerations	9
11.1	<i>Functional requirements</i>	9
11.2	<i>Performance requirements</i>	11
11.3	Guidance parameters	11
12.	Deployment considerations	13
12.1	<i>Deployment strategy</i>	13
12.2	<i>Deployment location</i>	13
12.3	<i>Alarm strategy</i>	14
13.	Limitations of IDSs	14
14.	Types of Computer Attacks Commonly Detected by IDSs	15
15.	Thumb rules	16
16.	Conclusion	17
17.	References	17
	Appendix : <i>Merits & Demerits of different types of IDSs</i>	18

1. Introduction

There has been steep escalation in the number of security incidents reported in the first quarter 2003, compared with incidents reported in all of 2000. This state of affairs is attributed by ever increasing number of vulnerabilities and attacks. Therefore, it is very important that the security mechanisms of a system are designed so as to prevent unauthorized access to system resources and data. Completely preventing breaches of security appear, at present, unrealistic. However, organizations can try to detect these intrusion attempts so that action may be taken before it is too late.

In the 1980s, most intruders were experts, with high levels of expertise and individually developed methods for breaking into systems. They rarely used automated tools and exploit scripts. Today, anyone can attack Internet sites using readily available intrusion tools & exploit scripts that capitalize on widely known vulnerabilities and damaging intrusions can occur in a matter of seconds. Intruders hide their presence by installing modified versions of system monitoring and administration commands and by erasing their tracks in audit and log files. In the 1980s and early 1990s, denial-of-service attacks were infrequent and not considered serious. Today, successful denial- of-service attacks can put e-commerce based organizations such as online stockbrokers and retail sites out of business [Ref:1].

Intrusion detection is the process of monitoring the events occurring in a computer system or network and analyzing them for signs of intrusions, defined as attempts to compromise the confidentiality, integrity, availability, or to bypass the security mechanisms of a computer or network. Intrusions are caused by attackers accessing the systems from the Internet, authorized users of the systems who attempt to gain additional privileges for which they are not authorized, and authorized users who misuse the privileges given them.

2. Purpose & Scope

This document is intended as a primer in intrusion detection for understanding what security goals intrusion detection mechanisms serve, what types of IDS exists, their merits & demerits, how to select and deploy intrusion detection systems for specific system and network environments, how to manage the output of intrusion detection systems, and how to integrate intrusion detection functions with the rest of the organizational security infrastructure.

Each security protection serves to address a particular security threat to the system. Furthermore, each security protection has weak and strong points. Only by combining them and by looking at the security in depth organizations can protect from a realistic range of security attacks. Firewalls serve as barrier mechanisms, barring entry to some kinds of network traffic and allowing others, based on a firewall policy. IDSs serve as monitoring mechanisms, watching activities, and making decisions about whether the observed events are suspicious. They can spot attackers circumventing firewalls and report them to system administrators, who can take steps to prevent damage.

3. What is Intrusion detection system?

Intrusion detection systems (IDSs) are software or hardware systems that automate the process of monitoring the events occurring in a computer system or network, analyzing them for signs of security problems. As network attacks have increased in number and severity over the past few years, intrusion detection systems have become a necessary addition to the security infrastructure of most organizations.

In a nutshell, intrusion detection systems do exactly as the name suggests: they detect possible intrusions. More specifically, IDS tools aim to detect computer attacks and/or computer misuse, and to alert the proper individuals upon detection. An IDS installed on a network provides much the same purpose as a burglar alarm system installed in a house. Through various methods, both detect when an intruder/attacker/burglar is present, and both subsequently issue some type of warning or alert.

Successful IDSs can recognize both intrusions and denial-of-service activities and invoke countermeasures against them in real time. To realize this potential, more accurate detection and reduced false-alarm rates are needed. The strength of an IDS lies in its ability to perform well the following functions:

- Monitoring and analysis of system events and user behaviors
- Testing the security states of system configurations
- Baseline the security state of a system, then tracking any changes to that baseline
- Recognizing patterns of system events that correspond to known attacks
- Recognizing patterns of activity that statistically vary from normal activity
- Managing operating system audit and logging mechanisms and the data they generate
- Alerting appropriate staff by appropriate means when attacks are detected.
- Measuring enforcement of security policies encoded in the analysis engine
- Providing default information security policies
- Allowing non-security experts to perform important security monitoring functions.

4. IDS vs Firewall

Although IDSs may be used in conjunction with firewalls, which aim to regulate and control the flow of information into and out of a network, the two security tools should not be considered the same thing. Firewalls can be thought of as a fence or a security guard placed in front of a house. They protect a network and attempt to prevent intrusions, while IDS tools detect whether or not the network is under attack or has, in fact, been breached. IDS tools thus form an integral part of a thorough and complete security system. They don't fully guarantee security, but when used with security policy, vulnerability assessments, data encryption, user authentication, access control, and firewalls, they can greatly enhance network safety.

A detailed knowledge of what a hacker can do and what should and shouldn't be allowed through the firewall is required before embarking on the configuration adventure, and a slip of the mouse is all it takes to open up a hole big enough for average hacker to drive the proverbial bus through. The problem is, a badly configured firewall can be worse than no firewall at all, since it will engender a false sense of security. To protect an organisation completely, therefore, it is necessary to provide a second line of

defense, and in order to achieve this, an entire category of mechanism exists in the form of Intrusion Detection Systems (IDS).

5. Need for IDS

IDSs are an integral and necessary element of a complete information security infrastructure performing as the logical complement to network firewalls. Simply put, IDS tools allow for complete supervision of systems, regardless of the action being taken, such that information will always exist to determine the nature of the security incident and its source.

Intrusion detection allows organizations to protect their systems from the threats that come with increasing network connectivity and reliance on information systems. The most compelling reasons to acquire and use IDSs are[Ref: 2]:

- To prevent problem behaviors by increasing the perceived risk of discovery and punishment for those who would attack or otherwise abuse the system,
- To detect attacks and other security violations that are not prevented by other security measures
- To detect and deal with the preambles to attacks (commonly experienced as network probes and other “doorknob rattling” activities)
- To document the existing threat to an organization
- To act as quality control for security design and administration, especially of large and complex enterprises
- To provide useful information about intrusions that do take place, allowing improved diagnosis, recovery, and correction of causative factors.

6. Components of an IDS

The three fundamental functional components of any IDS are Information source, Analysis and Response.

- *Information Sources* – the different sources of event information used to determine whether an intrusion has taken place. These sources can be drawn from different levels of the system, with network, host, and application monitoring most common.
- *Analysis* – the part of intrusion detection systems that actually organizes and makes sense of the events derived from the information sources, deciding when those events indicate that intrusions are occurring or have already taken place. The most common analysis approaches are *misuse detection* and *anomaly detection*.
- *Response* – the set of actions that the system takes once it detects intrusions. These are typically grouped into active and passive measures, with active measures involving some automated intervention on the part of the system, and passive measures involving reporting IDS findings to humans, who are then expected to take action based on those reports.

7. How they work?

Intrusion detection systems are made up of three functional components, information sources, analysis, and response. The system obtains event information from one or more information sources, performs a pre-configured analysis of the event data,

and then generates specified responses, ranging from reports to active intervention when intrusions are detected.

Intrusion detection systems serve three essential security functions: they monitor, detect, and respond to unauthorized activity by insiders and outsider intrusion. Intrusion detection systems use policies to define certain events that, if detected will issue an alert. In other words, if a particular event is considered to constitute a security incident, an alert will be issued if that event is detected. Certain intrusion detection systems have the capability of sending out alerts, so that the administrator of the IDS will receive a notification of a possible security incident in the form of a SMS, page, email, or SNMP trap. Many intrusion detection systems not only recognize a particular incident and issue an appropriate alert, they also respond automatically to the event. Such a response might include logging off a user, disabling a user account, and launching of scripts.

8. Characteristics of a good IDS

An intrusion detection system should address the following issues, regardless of what mechanism it is based on:

- It must run continually without human supervision. The system must be reliable enough to allow it to run in the background of the system being observed. However, it should not be a "black box". That is, its internal workings should be examinable from outside.
- It must be fault tolerant in the sense that it must survive a system crash and not have its knowledge-base rebuilt at restart.
- On a similar note to above, it must resist subversion. The system can monitor itself to ensure that it has not been subverted.
- It must impose minimal overhead on the system. A system that slows a computer to a crawl will simply not be used.
- It must observe deviations from normal behavior.
- It must be easily tailored to the system in question. Every system has a different usage pattern, and the defense mechanism should adapt easily to these patterns.
- It must cope with changing system behavior over time as new applications are being added. The system profile will change over time, and the IDS must be able to adapt.
- It must be difficult to fool.

9. Types of IDS

IDSs may be described according to their architecture, control strategy, system monitoring approaches, analysis strategy and the timing of information sources and analysis[Ref:2].

9.1 Architecture

IDSs can be categorized according to their architecture - how the functional components are arranged with respect to each other. The primary architectural components are the Host, the system on which the IDS software runs, and the Target, the system that the IDS is monitoring for problems. This gives rise to the following:

9.1.1 *Host-Target Co-location*: In early days of IDSs, most IDSs ran on the systems they protected. This was due to the fact that most systems were mainframe systems, and the cost of computers made a separate IDS system a costly extravagance. This presented a problem from a security point of view, as any attacker that successfully attacked the target system could simply disable the IDS as an integral portion of the attack.

9.1.2 *Host-Target Separation*: With the advent of workstations and personal computers, most IDS architects moved towards running the IDS control and analysis systems on a separate system, hence separating the IDS host and target systems. This improved the security of the IDS as this made it much easier to hide the existence of the IDS from attackers.

9.2 **Control Strategy**

IDSs can also be categorized according to the Control Strategy, which describes how the elements of IDS is controlled, and how the input and output of the IDS is managed.

9.2.1 *Centralized*: Under centralized control strategies, all monitoring, detection and reporting is controlled directly from a central location

9.2.2. *Fully Distributed*: Monitoring and detection is done using an agent-based

9.2.2 *Partially Distributed*: Monitoring and detection is controlled from a local control node, with hierarchical reporting to one or more central location(s).

9.2 **Timing**

Timing refers to the elapsed time between the events that are monitored and the analysis of those events.

9.3.1 *Interval-Based (Batch Mode)*: In interval-based IDSs, the information flow from monitoring points to analysis engines is not continuous. In effect, the information is handled in a fashion similar to “store and forward” communications schemes. Many early host-based IDSs used this timing scheme, as they relied on operating system audit trails, which were generated as files. Intervalbased IDSs are precluded from performing active responses.

9.3.2 *Real-Time(Continuous)*: Real-time IDSs operate on continuous information feeds from information sources. This is the predominant timing scheme for networkbased IDSs, which gather information from network traffic streams.

9.4 **Information Sources**

The most common way to classify IDSs is to group them by information source. Some IDSs analyze network packets, captured from network backbones or LAN segments, to find attackers. Other IDSs analyze information sources generated by the operating system or application software for signs of intrusion.

9.4.1 *Network-Based IDSs*

These IDSs detect attacks by capturing and analyzing network packets. Listening on a network segment or switch, one network-based IDS can monitor the network traffic affecting multiple hosts that are connected to the network segment, thereby protecting those hosts. Network-based IDSs often consist of a set of single-purpose sensors or hosts placed at various points in a network. These units monitor network traffic, performing local analysis of that traffic and reporting attacks to a central management console. As the sensors are limited to running the IDS, they can be more easily secured against attack. Many of these sensors are designed to run in “stealth” mode, in order to

make it more difficult for an attacker to determine their presence and location. The majority of commercial intrusion detection systems are networkbased.

9.4.2 Host-Based IDSs

Host-based IDSs operate on information collected from within an individual computer system. This vantage point allows hostbased IDSs to analyze activities with great reliability and precision, determining exactly which processes and users are involved in a particular attack on the operating system. Furthermore, unlike networkbased IDSs, host-based IDSs can “see” the outcome of an attempted attack, as they can directly access and monitor the data files and system processes usually targeted by attacks.

Host-based IDSs normally utilize information sources of two types, operating system audit trails, and system logs. Operating system audit trails are usually generated at the innermost (kernel) level of the operating system, and are therefore more detailed and better protected than system logs. However, system logs are much less obtuse and much smaller than audit trails, and are furthermore far easier to comprehend. Some host-based IDSs are designed to support a centralized IDS management and reporting infrastructure that can allow a single management console to track many hosts. Others generate messages in formats that are compatible with network management systems.

9.4.3 Application-Based IDSs

Application-based IDSs are a special subset of host-based IDSs that analyze the events transpiring within a software application. The most common information sources used by application-based IDSs are the application’s transaction log files. The ability to interface with the application directly, with significant domain or application-specific knowledge included in the analysis engine, allows application-based IDSs to detect suspicious behavior due to authorized users exceeding their authorization. This is because such problems are more likely to appear in the interaction between the user, the data, and the application.

9.5 IDS Analysis

There are two primary approaches to analyzing events to detect attacks: misuse detection and anomaly detection. Misuse detection, in which the analysis targets something known to be “bad”, is the technique used by most commercial systems. Anomaly detection, in which the analysis looks for abnormal patterns of activity, has been, and continues to be, the subject of a great deal of research. Anomaly detection is used in limited form by a number of IDSs. There are strengths and weaknesses associated with each approach, and it appears that the most effective IDSs use mostly misuse detection methods with a smattering of anomaly detection components.

9.5.1 Misuse Detection

Misuse detectors analyze system activity, looking for events or sets of events that match a predefined pattern of events that describe a known attack. As the patterns corresponding to known attacks are called *signatures*, misuse detection is sometimes called “signature-based detection.” The most common form of misuse detection used in commercial products specifies each pattern of events corresponding to an attack as a separate signature. However, there are more sophisticated approaches to doing misuse detection (called “state-based” analysis techniques) that can leverage a single signature to detect groups of attacks.

9.5.2 **Anomaly Detection**

Anomaly detectors identify abnormal unusual behavior (anomalies) on a host or network. They function on the assumption that attacks are different from “normal” (legitimate) activity and can therefore be detected by systems that identify these differences. Anomaly detectors construct profiles representing normal behavior of users, hosts, or network connections. These profiles are constructed from historical data collected over a period of normal operation. The detectors then collect event data and use a variety of measures to determine when monitored activity deviates from the norm. The measures and techniques used in anomaly detection include threshold detection, statistical measures (both parametric & non-parametric), rule-based measures. Additionally, measures like neural networks, genetic algorithms, and immune system models are also being used. Current commercial IDSs mostly exploit the first two measures.

9.6 **Response Options**

Once IDSs have obtained event information and analyzed it to find symptoms of attacks, they generate responses. Some of these responses involve reporting results and findings to a pre-specified location. Others involve more active automated responses. Though researchers are tempted to underrate the importance of good response functions in IDSs, they are actually very important. Commercial IDSs support a wide range of response options, often categorized as active responses, passive responses, or some mixture of the two.

9.6.1 **Active Responses**

Active IDS responses are automated actions taken when certain types of intrusions are detected. There are three categories of active responses.

- **Collect additional information** - about a suspected attack and to decide whether action should be taken.
- **Change the Environment** - halt an attack in progress and then block subsequent access by the attacker. (eg. inject TCP reset packets, reconfiguring routers and firewalls)
- **Take Action Against the Intruder** – the “strike-back” option

9.6.2 **Passive Responses**

Passive IDS responses provide information to system users, relying on humans to take subsequent action based on that information. Many commercial IDSs rely solely on passive responses.

- **Alarms and Notifications** - like onscreen alert or popup window. displayed on the IDS console
- **SNMP Traps and Plug-ins**- that post alarms and alerts to central network management consoles

10. **Tools that complement IDSs**

There are several tools that complement an IDS and they can be used to enhance an organization’s intrusion detection capability. These tools are often labeled as intrusion detection products by vendors since they perform similar functions. Some of the most common tools are:

10.1 Vulnerability Assessment Tools - test to determine whether a network or host is vulnerable to known attacks and are essentially batch mode misuse detectors that

operate on reliably generated “snapshot” of the security state of a system at a particular time & results of specified test routines.

10.2 File Integrity Checkers -utilize message digest or other cryptographic checksums for critical files and objects, comparing them to reference values, and flagging differences or changes. They are extremely valuable to those conducting a forensic examination of systems that have been attacked, as they allow quick and reliable diagnosis of the footprint of an attack

10.3 Honey Pots & Padded Cells -

Honey Pots are decoy systems that are designed to lure a potential attacker away from critical systems and are re designed to divert an attacker from accessing critical systems, collect information about the attacker’s activity, and encourage the attacker to stay on the system long enough for administrators to respond. They are filled with fabricated information designed to appear valuable but that a legitimate user of the system wouldn’t access. *Padded cells* Rather than trying to attract attackers with tempting data, a padded cell operates in tandem with a traditional IDS and when the IDS detects attackers, it seamlessly transfers then to a special padded cell host. Once the attackers are in the padded cell, they are contained within a simulated environment where they can cause no harm.

11. Selection considerations

The wide array of intrusion detection products available today addresses a range of organizational security goals and considerations. Given this range of products and features, the process of selecting products that represent the best fit for your organization’s needs is, at times, difficult.[Ref: 2, 5]

11.1 Functional requirements

- As the network-computing environment increases in complexity, so do the functional requirements of IDSs. Common functional requirements of an IDS being deployed in current or near-term operational computing environments include the following:
 - The IDS must continuously monitor and report intrusions.
 - The IDS must supply enough information to repair the system, determine the extent of damage, and establish responsibility for the intrusion.
 - The IDS should be modular and configurable as each host and network segment will require their own tests and these tests will need to be continuously upgraded and eventually replaced with new tests.
 - Since the IDS is assigned the critical role of monitoring the security state of the network, the IDS itself is a primary target of attack. The IDS must be able to operate in a hostile computing environment and exhibit a high degree of fault-tolerance and allow for graceful degradation.
 - The IDS should be adaptive to network topology and configuration changes as computing elements are dynamically added and removed from the network.
 - Anomaly detection systems should have a very low false alarm rate. Given the projected increase in network connectivity and traffic, simply decreasing the percentage of overall false alarms may not be sufficient as their absolute number may continue to rise.

- The IDS should be able to learn from past experiences and improve its detection capabilities over time. A self-tuning IDS will be able to learning from false alarms with the guidance of system administrators and eventually on its own.
- The IDS should be able to be easily and frequently updated with attack signatures as new security advisories and security patches become available and new vulnerabilities and attacks are discovered.
- Decision support tools will be necessary to help system administrators respond to various attacks. The IDS will be required not only to detect anomalous events, but also to take automated corrective action.
- The IDS should be able to perform data fusion and be able to process information from multiple and distributed data sources such as firewalls, routers, and switches. As real-time detection demands push networked-based solutions to re-programmable hardware devices that can download new capabilities, the IDS will need to be able to communicate with the hardware-based devices.
- Data reduction tools will be necessary to help the IDS process the information gathered from data fusion techniques. Data mining tools will be helpful in running statistical analysis tools on archived data in support of anomaly detection techniques.
- The IDS should be capable of providing an automated response to suspicious activity. Rapid changes in network conditions and limited network administration expertise make it difficult for system administrators to diagnose problems and take corrective action to minimize the damage that intruders can cause.
- The ability to detect and react to distributed and coordinated attacks will become necessary. Coordinated attacks against a network will be able to marshal greater forces and launch many more and varied attacks against a single target. These attacks can be permutations of known attacks, be rapidly evolving, and be launched at little cost to the attackers.
- Distributing the computational load and the diagnostic capabilities to agents scattered throughout the network adds a level of fault-tolerance, but it is often necessary for the system administrator to have control over the IDS from a central location.
- The IDS should be able to work with other Commercial Off-the-Shelf (COTS) security tools, as no vendor toolset is likely to excel in or to provide complete coverage of the detection, diagnosis, and response responsibilities. The IDS framework should be able to integrate various data reduction, forensic, host-based, and network-based security tools. Interoperability and conformance to standards will further increase the value of the IDS.
- IDS data often requires additional analysis to assess any damage to the network after an intrusion has been detected. Although the anomalous event was the first detected, it may not be the first attempt to gain unauthorized access to the network. Post event analysis will be needed to identify compromised machines before the network can be restored to a safe condition.
- The IDS itself must also be designed with security in mind. For example, the IDS must be able to authenticate the administrator, audit administrator actions, mutually authenticate IDS devices, protect the IDS data, and not create additional vulnerabilities.

11.2 Performance requirements

An IDS that is functionally correct, but that detects attacks too slowly is of little use. Thus one must enumerate several performance requirements for IDSs. The IDS performance requirements include:

- To the extent possible, anomalous events or breaches in security should be detected in real-time and reported immediately to minimize the damage to the network and the loss or corruption of confidential information.
- The IDS must not place undue burden or interfere with the normal operations for which the systems were bought and deployed to begin with. This requirement makes it necessary for the agents to be cognizant of the consumption of network resources for which they are competing. There is a tradeoff between additional levels of security monitoring and the performance penalty to be paid by other applications.
- The IDS must be scalable. As new computing devices are added to the network, the IDS must be able to handle the additional computational and communication load.

11.3 Guidance parameters

The following parameters may be used as guidance when preparing a specification for acquiring an intrusion detection product[Ref. 1].

- Technical and Policy Considerations
 - system environment
 - technical specifications of systems environment
 - technical specifications of current security protections
 - goals of the organization
 - system environment and management culture in the organization
 - security goals and objectives
 - protecting from threat originating outside organization
 - protecting from insider attack
 - use of the output of the IDS to determine new needs
 - use of the IDS to maintain managerial control (non-security related) over network usage
 - existence of security policy
 - structure of the policy
 - general job descriptions of system users
 - existence of reasonable use policies or other management provisions
 - defined processes for dealing with specific policy violations
- Organizational Requirements and Constraints
 - requirements that are levied from outside the organization
 - requirements for public access to information on organization's systems
 - other security-specific requirements levied by law
 - internal audit requirements for security best practices or due diligence
 - system subject to accreditation
 - requirements for law enforcement investigation and resolution of security incidents

- organization's resource constraints
 - budget for acquisition and life cycle support of intrusion detection hardware, software, and infrastructure
- existence of sufficient number of existing staff to monitor an intrusion detection system full time
- authority to instigate changes based on the findings of an intrusion detection system
- IDS Product Features and Quality
 - product sufficiently scalable for your environment
 - the product been tested
 - tested against functional requirements
 - product been tested against attack
 - user level of expertise targeted by the product
 - designed to evolve as the organization grows
 - adapt to growth in user expertise
 - adapt to growth and change of the organization's systems infrastructure
 - adapt to growth and change of the security threat environment
 - support provisions for the product
 - commitments for product installation and configuration support
 - commitments for ongoing product support
 - subscriptions to signature updates included
 - How often are subscriptions updated
 - How quickly after a new attack is made public will the vendor ship a new signature
 - Are software updates included
 - How quickly will software updates and patches be issued after a problem is reported to the vendor
 - Are technical support services included? What is the cost
 - What are the contact provisions for contacting technical support (email, telephone, online chat, web-based reporting)
 - Are there any guarantees associated with the IDS
 - training resources does the vendor provide as part of the product
 - additional training resources are available from the vendor and at what cost
- Reporting and Archiving Capabilities
 - capabilities to generate routine reports and other detailed information documents
 - reports of system events and intrusions detected over a particular reporting period
 - statistics or logs generated by the IDS in formats suitable for inclusion in database systems or for use in report generating
- Failsafe considerations for IDS responses
 - existence of failsafe features that are meant to protect the IDS from being circumvented or defeated by an attacker.

- provision to provide silent, reliable monitoring of attackers without broadcasting alarms and alerts in plaintext over the monitored network
- use of encrypted tunnels or other cryptographic measures to hide and authenticate IDS communications.

12. Deployment considerations

12.1 Deployment Strategy

Intrusion detection technology is a necessary addition to every organization's computer network security infrastructure. However, given the deficiencies of today's intrusion detection products, and the limited security skill level of many system administrators, an effective IDS deployment requires careful planning, preparation, prototyping, testing, and specialized training.

It is necessary to perform a thorough requirements analysis, careful selection of the intrusion detection strategy and solution that is compatible with the organization's network infrastructure, policies, and resource level.

Organizations should consider a staged deployment of IDSs to allow personnel to gain experience and to ascertain how many monitoring and maintenance resources they will require. The resource requirements for each type of IDS vary widely, depending on the organization and systems environment. IDSs require significant preparation and ongoing human interaction. Organizations must have appropriate security policies, plans, and procedures in place so that personnel know how to handle the many and varied alarms IDSs produce.

It is recommend to consider a combination of network-based IDSs and hostbased IDSs to protect an enterprise-wide network. Start with network-based IDSs as they are usually the simplest to install and maintain. Next, protect critical servers with host-based IDSs. Utilize vulnerability analysis products on a regular schedule to test IDSs and other security mechanisms for proper function and configuration. Honey pots and related technologies should be used conservatively and only by organizations with a highly skilled technical staff that are willing to experiment with leading-edge technology. Furthermore, such techniques should be used only after seeking guidance from legal counsel.

12.2 Deployment Location

One question that arises when deploying network-based IDSs is where to locate the system sensors. There are many options for placing a network-based IDS with different advantages associated with each location:

- Behind each external firewall, in the network DMZ
 - Sees attacks, originating from the outside world, that penetrate the network's perimeter defenses.
 - Highlights problems with the network firewall policy or performance
 - Sees attacks that might target the web server or ftp server, which commonly reside in this DMZ
 - Even if the incoming attack is not recognized, the IDS can sometimes recognize the outgoing traffic that results from the compromised server

- Outside an external firewall
 - Documents number of attacks originating on the Internet that target the network.
 - Documents types of attacks originating on the Internet that target the network
- On major network backbones
 - Monitors a large amount of a network's traffic, thus increasing the possibility of spotting attacks.
 - Detects unauthorized activity by authorized users within the organization's security perimeter.
- On critical subnets
 - Detects attacks targeting critical systems and resources.
 - Allows focusing of limited resources to the network assets considered of greatest value.

12.3 Alarm strategy

When deploying IDSs, the questions of which IDS alarm features to use and when are important issues. Most IDSs come with configurable alarm features, which allow a wide variety of alarm options, including SMS, email, paging, network management protocol traps, and even automated blocking of attack sources. Although these features may be appealing, it is important to be conservative about using them until you have a stable IDS installation and some sense of the behavior of the IDS within your environment. Some experts recommend not activating IDS alarms for as long as several months after installation. In cases where the alarm and response features include automated response to attacks, specifically those that allow the IDS to direct the firewall to block traffic from the ostensible sources of the attacks, be extremely careful that attackers do not abuse this feature to deny access to legitimate users.

13. Limitations of IDSs

For the effective deployment of an IDS, it is important to understand the limitations of it. IDS cannot perform the following functions:

- Compensating for weak or missing security mechanisms in the protection infrastructure. Such mechanisms include firewalls, identification and authentication, link encryption, access control mechanisms, and virus detection and eradication.
- Instantaneously detecting, reporting, and responding to an attack, when there is a heavy network or processing load.
- Detecting newly published attacks or variants of existing attacks.
- Effectively responding to attacks launched by sophisticated attackers
- Automatically investigating attacks without human intervention.
- Resisting attacks that are intended to defeat or circumvent them
- Compensating for problems with the fidelity of information sources
- Dealing effectively with switched networks.

14. Types of Computer Attacks Commonly Detected by IDSs

Three types of computer attacks are most commonly reported by IDSs: system scanning, denial of service (DOS), and system penetration. These attacks can be launched locally, on the attacked machine, or remotely, using a network to access the target. An IDS operator must understand the differences between these types of attacks, as each requires a different set of responses.

14.1 Scanning Attacks

A scanning attack occurs when an attacker probes a target network or system by sending different kinds of packets. Using the responses received from the target, the attacker can learn many of the system's characteristics and vulnerabilities. Scanning attacks do not penetrate or otherwise compromise systems. Various tools used to perform these activities include: network mappers, port mappers, network scanners, port scanners, or vulnerability scanners.

Scanning attacks may yield:

- The topology of a target network
- The types of network traffic allowed through a firewall
- The active hosts on the network
- The operating systems those hosts are running
- The server software they are running
- The software version numbers for all detected software

14.2 Denial of Service Attacks

Denial Of Service (DOS) attacks attempt to slow or shut down targeted network systems or services. In certain Internet communities, DOS attacks are common.. While often used for such trivial purposes, DOS attacks can also be used to shut down major organizations. In well publicized incidents, DOS attacks were charged with causing major losses to electronic commerce operations, whose customers were unable to access them to make purchases. There are two main types of DOS attacks: flaw exploitation and flooding. It is important for an IDS operator to understand the difference between them.

14.3. Flaw exploitation DOS Attacks

Flaw exploitation attacks exploit a flaw in the target system's software in order to cause a processing failure or to cause it to exhaust system resources such as the 'ping of death' attack. This attack involves sending an unexpectedly large ping packet to certain Windows systems. The target system could not handle this abnormal packet, and a system crash resulted. With respect to resource exhaustion attacks, the resources targeted include CPU time, memory, disk space, space in a special buffer, or network bandwidth. In many cases, simply patching the software can circumvent this type of DOS attack.

14.4 Flooding DOS Attacks

Flooding attacks simply send a system or system component more information than it can handle. In cases where the attacker cannot send a system sufficient information to overwhelm its processing capacity, the attacker may nonetheless be able to monopolize the network connection to the target, thereby denying anyone else use of the resource.

With these attacks, there is no flaw in the target system that can be patched. While there are few general solutions to stop flooding attacks, there are several technical modifications that can be made by a target to mitigate such an attack.

The term “distributed DOS” (DDOS) is a subset of DOS attacks. DDOS attacks are simply flooding DOS attacks where the attacker uses multiple computers to launch the attack. These attacking computers are centrally controlled by the attacker’s computer and thus act as a single immense attack system. An attacker cannot usually bring down a major ecommerce site by flooding it with network packets from a single host. However, if an attacker gains control of 20,000 hosts and subverts them to run an attack under his direction, then the attacker has a formidable capability to successfully attack the fastest of systems, bringing it to a halt.

14.5 Penetration Attacks

Penetration attacks involve the unauthorized acquisition and/or alteration of system privileges, resources, or data. These violate integrity and control of a system. A penetration attack can gain control of a system by exploiting a variety of software flaws.

While penetration attacks vary tremendously in details and impact, the most common types are:

User to Root: A local user on a host gains complete control of the target host

Remote to User: An attacker on the network gains access to a user account on the target host

Remote to Root: An attacker on the network gains complete control of the target host

Remote Disk Read: An attacker on the network gains the ability to read private data files on the target host without the authorization of the owner

Remote Disk Write: An attacker on the network gains the ability to write to private data files on the target host without the authorization of the owner

15. Thumb rules

- Ensure that one (or more) network-based detectors are monitoring each network segment by installing a detector on the segment itself or on a segment boundary device that has the ability to inspect all packets on the subnet.
- Identify all the servers that the organization’s security policy deems critical to the enterprise. Deploy a host-based intrusion detection mechanism on each of them.
- Within each critical server, identify all critical network applications and deploy an application-based intrusion detection system on each one.
- Obtain a current set of attack signatures from IDS vendors and install accordingly. Use a configuration management tool to track the signature file information on all systems.
- Based on the list of trusted and distrusted entities that have been compiled, generate a policy for host- and application-based intrusion detection systems to ensure that all unauthorized access attempts are logged and dealt with appropriately.

- Update the policy created above with the list of Access Control Lists that have previously compiled. Ensure that anomaly detection policies reflect the network load histograms collected for each network segment.
- Harden and secure all intrusion detectors.
- Establish a policy for rotating logs, a copy of which should always be written to remote, removable media.
- Establish a policy for monitoring these systems. Use an SNMP agent and device-generated SNMP traps, when available.
- Review all policies and signature files on a regular basis.

16. Conclusion

Whenever an organization connects its network to the Internet, it opens up a whole can of worms regarding security. As the network grows, it will play host to numerous bugs and security loop holes. Organizations recognize the value of a good security policy to define what is allowed (and what is not allowed) in terms of network and Internet access and deploy a number of tools like a well-configured firewall to enforce that security policy. A badly configured firewall can be worse than no firewall at all, since it will engender a false sense of security. To protect an organization completely, therefore, it is necessary to provide a second line of defense, in the form of IDS.

17. References

1. Defending Yourself: The Role of Intrusion Detection Systems; John McHugh, Alan Christie, and Julia Allen; IEEE software.
2. Special Publication on Intrusion Detection Systems; <http://www.nist.gov>
3. Emerging Technology: Deploying an Effective Intrusion Detection System <http://www.networkmagazine.com/shared/article/showArticle.jhtml?articleId=8702894>
4. SANS IDS FAQ <http://www.sans.org>
5. CERT Guide to Systems & Network Practices; Julia H. Allen; Addison Wesley
6. An Overview of Issues in Testing Intrusion Detection System; <http://www.nist.gov>

Appendix

Merits & De-merits of different types of IDSs

	Merits	Demerits
N-IDS	<ul style="list-style-type: none"> • can monitor a large network. • little impact upon an existing network as they are passive devices • can be made very secure against attack and even made invisible to many attackers. 	<ul style="list-style-type: none"> • difficulty in processing all packets in a large/busy network • may fail to recognize an attack launched during periods of high traffic. • Many of the advantages of network-based IDSs don't apply to more modern switch-based networks. • cannot analyze encrypted information. • cannot tell whether or not an attack was successful; they can only discern that an attack was initiated. • problems dealing with networkbased attacks that involve fragmenting packets.
H-IDS	<ul style="list-style-type: none"> • ability to monitor events local to a host, can detect attacks that cannot be seen by a network-based IDS. • can often operate in an environment in which network traffic is encrypted, when the host-based information sources are generated before data is encrypted and/or after the data is decrypted at the destination host • unaffected by switched networks. • When Host-based IDSs operate on OS audit trails, they can help detect Trojan Horse or other attacks that involve software integrity breaches. These appear as inconsistencies in process execution. 	<ul style="list-style-type: none"> • harder to manage, as information must be configured and managed for every host monitored. • Since at least the information sources (and sometimes part of the analysis engines) for host-based IDSs reside on the host targeted by attacks, the IDS may be attacked and disabled as part of the attack. • not well suited for detecting network scans or other such surveillance that targets an entire network, because the IDS only sees those network packets received by its host. • can be disabled by certain denial-of-service attacks. • When host-based IDSs use operating system audit trails as an information source, the amount of information can be immense, requiring additional local storage on the system. • use the computing resources of the hosts they are monitoring, therefore inflicting a performance cost on the monitored systems.

A-IDS	<ul style="list-style-type: none"> • can monitor the interaction between user and application, which often allows them to trace unauthorized activity to individual users. • can often work in encrypted environments, since they interface with the application at transaction endpoints, where information is presented to users in unencrypted form. 	<ul style="list-style-type: none"> • may be more vulnerable than host-based IDSs to attacks as the applications logs are not as well-protected as the operating system audit trails used for host-based IDSs. • cannot detect Trojan Horse or other such software tampering attacks. • advisable to use an Application-based IDS in combination with Host-based and/or Network-based IDSs.
MD	<ul style="list-style-type: none"> • very effective at detecting attacks without generating an overwhelming number of false alarms. • can quickly and reliably diagnose the use of a specific attack tool or technique. • allows system managers to track security problems on their systems, initiating incident handling procedures. 	<ul style="list-style-type: none"> • can only detect those attacks they know about – therefore they must be constantly updated with signatures of new attacks. • are designed to use tightly defined signatures that prevent them from detecting variants of common attacks. State-based misuse detectors can overcome this limitation, but are not commonly used in commercial IDSs.
AD	<ul style="list-style-type: none"> • detect unusual behavior and thus have the ability to detect symptoms of attacks without specific knowledge of details. • can produce information that can in turn be used to define signatures for misuse detectors. 	<ul style="list-style-type: none"> • produce a large number of false alarms due to the unpredictable behaviors of users and networks. • often require extensive “training sets” of system event records in order to characterize normal behavior patterns.

N-IDS: Network IDS
H-IDS: Host IDS
A-IDS: Application IDS
MD: Mis-use Detection
AD: Anomaly Detection