



Global Information Assurance Certification Paper

Copyright SANS Institute
Author Retains Full Rights

This paper is taken from the GIAC directory of certified professionals. Reposting is not permitted without express written permission.

Interested in learning more?

Check out the list of upcoming events offering
"Security Essentials (Security 401)"
at <http://www.giac.org/registration/gsec>

Dynamic Network Protection

Abstract

Most dialogues of network protection focus on the software and procedures used to fortify networks: firewalls, biometrics, access controls, and encryption. This paper presents an outline of dynamic protection mechanisms assisting an administrator in verifying and maintaining the protection of a network.

It discusses why there has been a need for such software and how protection mechanisms have been attacked. The report also describes the software available in this field, with specific emphasis on Intrusion Detection software.

Introduction

Computer networks have been becoming fundamental to the operation of modern organizations. As the dependency on networks increases, the need to control networked assets becomes increasingly critical. At the same time, networks have been becoming ever more important – in terms of their operation, the assets they offer, and the information they contain. In this way, they become not only more important to an entity itself – they also become an attractive target for unfriendly parties.

The concepts of defending assets have not been new; protection of physical assets has been a soundly developed part of any organizational arrangement. With the uncontrolled growth of inter networks, the logical assets of an entity have been increasingly exposed. It has been now possible for an intruder to penetrate a network, appropriate or vandalize a company's most important assets, and leave – all without leaving any physical track.

A wide assortment of protection mechanisms have been developed, aimed at safeguarding the logical assets of an entity: access controls, firewall technologies, encryption and cryptographic authentication, biometrics and the like. These measures have one common factor in that they attempt to prevent unauthenticated access to assets. What has been missing has been a responsive element – the protection guards, monitoring and alarm elements present in physical protection structures.

Dynamic Network Protection has been comprised of a number of procedures that address this shortcoming. The aim has been not only to reduce the number of successful abuses of a network, but also to give an early warning of abuses in progress. Lastly, the objective has been to ensure that misuse of the network does not go unnoticed.

A wide assortment of protection software and mechanisms have been currently available. This begs the question: Why has been there a need for dynamic protection?

In order to answer this, let us consider the 2002 CSI/FBI Computer Crime and Protection Survey [1].

The survey (dated April 2002) was conducted over 503 US companies. These companies had an assortment of protection structures in place.

In spite of these measures, 90% of these companies reported experiencing unauthenticated use of their computer networks. Twenty one percent did not know if their networks had been abused. Forty percent reported outside penetration of their networks. This survey has also challenged the conventional wisdom that the threat from inside the entity has been far greater than the threat from outside the entity.

Many of the organizations had been unable to quantify their losses due to intrusions – for the 223 organizations that had been, the total losses exceeded \$ 455 million. Clearly, in spite the presence of protection mechanisms (with the vast majority of organizations having access controls and firewalls in place), exploitation of networks continue – sometimes without the entity even being aware of the breach.

As a specific case, consider an organizational Web location. As an entity's most visible Internet network, these have long been favored point of assault. With the development of electronic commerce and the increasing use of the Internet, an entity's Web location has been developing a significant commercial value. By the same token, assaults on these locations could do significant harm to an entity – in loss of revenue, loss of customer confidence and damage to information networks.

A good illustration of the risks involved has been the “Solar Sunrise” assaults on US government locations [2]. During this series of assaults, a wide assortment of web locations had been defaced or disabled – including such locations as the FBI, the US Army main Web location, a number of government divisions, the US Information Agency, and the US Senate.

Clearly, conventional, static protection mechanisms such as firewalls have been incapable of offering complete defense. Dynamic Protection mechanisms such as Intrusion Detection must have a place in any secure network.

Dynamic protection mechanisms

Dynamic network protection, as described in this document, encompasses networking software and networks that allow network administrators to observe, inspect and improve the protection of their networks. Many conventional protection mechanisms have been effective in enforcing protection in a network, but lack the responsiveness necessary to maintain protection on an ongoing basis.

In recent years, a number of protection software have been developed that may best be classified under this heading: while these software often have no direct effect in preventing misuse, they allow administrators to improve the overall protection of their networks.

Examples include:

- Intrusion Detection Systems (IDS) – Intrusion Detection Systems monitor the state of a network, attempting to recognize and report improper behavior. These networks defend a network in much the same way as protection cameras defend buildings: by letting protection staff keep an eye on what has been going on.
- Network Protection Scanners – Network scanners inspect a network or host network, looking for known vulnerabilities. The best known example has been the Satan tool – it scans nodes and connected networks for a specific series of vulnerabilities, reports any found, and suggests solutions.
- Network Integrity Checkers – Many of the ways in which networks have been attacked involve changes to the host's software and data. Integrity checkers compare the contents of a network to a known safe state – allowing administrators to know exactly what has been changed.
- Honeytrap networks – If an IDS has been a protection camera, this has been an alarm; networks whose sole purpose has been to be attacked. By closely monitoring these networks, network administrators may observe intruders in action – allowing them to repair, learn and strengthen protection against future assaults.
- Special purpose software – Specific software have been developed to address protection vulnerabilities present in networks. While not as generally applicable as those listed above, still deserve a place in every administrator's toolkit. Examples include: network scanning tool, password crackers and sniffers.

In a world where protection mechanisms had been infallible, none of these networks would be necessary. In fact, none of these networks may, in itself, prevent an assault from succeeding.

The Limitations of Static Protection

The static ways such as firewalls have been effective in ensuring the protection of any network. Even in realistic environments, static protection mechanisms have been capable of significantly improving the protection of networked assets. In spite of the wide assortment of protection mechanisms available, intrusions continue to occur.

Based on this fact, a number of limitations in static protection mechanisms may be identified.

1) The defense offered by these mechanisms has been limited in scope. While these mechanisms may be effective in the context in which they have been applied, they do not offer universal defense. For example, firewalls, while being effective against external assault, offer no defense against internal exploitation. The same type of argument applies to other mechanisms: authentication has been weak to trust networks, where the authentication mechanisms have been bypassed. Encryption only defends information while in an encrypted form.

2) The protection mechanisms themselves have been sensitive to technical and implementation issues. Such networks may become weak due to theoretical advances (such as the DES encryption standard, which may no longer be considered completely secure [2]).

3) Protection mechanisms must be correctly applied in order to be effective. Many of the protection mechanisms available have been very intricate (both in arrangement and in application), and a single mistake may be enough to nullify the efficacy of the network. An example of this has been the use of dial-in lines allowing direct access to a trusted network. No matter how good the firewall blocking to that network has been, it is still defenseless.

4) Static protection mechanisms, by their very nature, have been prone to silent failure. Often, the first sign that your protection has failed comes when it has been far too late (such as when an entire server has been wiped clean – an effective way for an intruder to erase a history of his actions). Even when a network's protection has not yet been penetrated, that may lead to a mistaken sense of protection. In general, these mechanisms also may not recognize when they have been under assault – at best, an assault has been logged as a series of failed transactions.

5) Associated with the previous point has been the issue of remedial information. Once a failure has been identified, it may be difficult or impossible to track the cause of that failure. Information on the identity of the intruder may allow the effects of an intrusion to be alleviated – but none of the mechanisms described offer any such capabilities. The audit information collected by some software, while being useable, does not have sufficient detail to allow this type of diagnostics.

6) Lastly, the protection mechanisms may themselves be subject to assault. Authentication servers may be corrupted, firewalls crashed or circumvented, and cryptographic distribution channels may be compromised. In many cases it has been a simple exercise to disable network by attacking its underlying infrastructure. A good illustration of this has been a number of software that have been freely available, aimed at allowing users to get around the restrictions applied by protection mechanisms – anonymous proxies and the like.

The fundamental issue with static protection mechanisms has been that they have been essentially passive (analogical to a medieval fort around the city without archers defending the walls). While this may be sufficient for a degree of protection, it does not hold up in the world of modern networks, where administrators have been often overworked, do not have the necessary skills, and where the assaults on networks have been intensifying.

Intrusion Detection

Intrusion Detection has as its primary aim the detection of exploitation of computer network. The ideal IDS is capable of detecting intrusive behavior in progress, notify protection staff of the issue, and be capable of taking action to minimize the risk posed by such exploitation.

A second, less obvious aim of IDS has been to collect data on network behavior, in order to facilitate recovery after intrusions, identify the origin of the assault, and serve as legal evidence in the case of a prosecution in the aftermath of an episode.

IDS goals may be broken down into the following specific points:

- IDS must be capable of accurately differentiating normal or adequate user behavior from potentially damaging actions.
- IDS must be capable of scaling across the large composite networks increasingly present in the real world.
- IDS must be capable of handling the intricate structures and interactions of modern networks, and must be capable of deployment across an assortment of network architectures.
- IDS must be capable of adapting in response to new assaults and usage patterns, with minimal human intervention.
- IDS must offer reports of assaults in real time, ideally as the intrusion has been in progress – allowing protection staff to take corrective action.
- IDS must cooperate with other protection mechanisms, increasing the overall protection of networks. IDS must be capable of detecting failures or assaults on other protection mechanisms, forming a second level of defense.
- IDS must be capable of responding to intrusive behavior by increasing its monitoring in the relevant sections, increasing the protection in relevant sections, or by excluding intrusive behavior.
- IDS must recognize abusive behavior in all sections of a network.
- IDS must defend itself against assaults, ensure the integrity of the greater network and audit information, and ensure that a compromised or unfriendly component may not adversely affect the operation of the network as a whole.
- IDS must continue to operate in the presence of network failures, unreliable transmission, high network loads, and denial of service assaults..
- IDS must generate audit information for network profiling and use in the recovery of intrusions. Specifically, IDS must generate information in a manner that would allow it to be admissible as evidence in a court.

Intrusion Detection networks have evolved from batch oriented structures to intricate, distributed real time networks of components. IDS basic general model has emerged, allowing for discrete components to be distinguished. [3]

A classic IDS arrangement consists of the following components:

Sensors: These components gather data for an IDS. Sensors take the form of monitoring processes on networked hosts (extracting information from the host event logs), or of dedicated network monitors connected to an observation point on a network node. From there, a network monitor inspects all visible network traffic. These networks also filter the event logs, generating summaries that are intended for IDS Monitors.

Monitors: These are the processing elements of an IDS. Monitors receive and interpret event summaries received from sensors. These event summaries get inspected for suspicious activity. The suspicion reports are forwarded to resolvers.

Resolvers: These elements are responsible for determining appropriate responses such as: notification of administrator, changing the behavior sensors and monitors or reconfiguration of firewalls.

Controllers: Configuration of components has been possible with controllers. These IDS components simplify network administration and allow administrative staff to rapidly reconfigure IDS components.

The division between IDS components has often been indistinct especially in newer systems.

Intrusion Detection Procedures

Intrusion Detection methodologies may be broken down into two major categories: Misuse Detection and Anomaly Detection.

Misuse Detection (MIDS) attempts to match actual behavior against known intrusive patterns. An assortment of procedures have been used to model and recognize assault patterns, such as expert networks [4], signature examination (used in [5] and [9]), Petri nets, and genetic algorithms. A common element between these procedures has been that they attempt to represent the fundamental nature of a known assault in such a way that variations on that assault may be distinguished from normal user behavior. Anything that has been not recognized as an assault has been accepted as legal behavior.

The dominant form of misuse detection used has been signature examination. A limitation of this methodology, and MIDS in general, has been that the signature set requires constant review as new assaults develop. In addition, as more assaults and assault variations become available, the number of signatures against which an event flow must be checked becomes larger.

Anomaly Detection (AIDS) attempts to model the usual behavior of users. Any action that does not correspond to expectations has been considered suspicious. The strength of this strategy lies in the ability to differentiate normal user behavior, abnormal user behavior, and intrusive behavior. Procedures used for constructing models include statistical measures [7], expert networks, neural networks, and user behavior profiling. Actual behavior gets compared to known patterns or expected behavior.

Scarce use of AIDS approach is caused by the following:

- Network overhead involved in maintaining and checking intricate behavioral models.
- Overhead involved in maintaining profiles for every user and process involved.
- Difficulty categorizing valid changes in user behavior.
- Issues modeling intricate networks accurately.
- Generation of large numbers of false positives.
- The ability for intruders to trick IDS to ignore intrusive behavior.

Monitor processing patterns

The ideal IDS would be capable of detecting all assaults in real time, and offer historical summaries. In practice, IDS networks often break down into real time or batch categories.

Real time networks suffer from performance issues (inspection of large amounts of information in real time). The ability of real time detection to observe and respond to intrusions in progress is of a great value. Most products appear to fall within this group. In addition, real time networks require a sensor and its monitor to reside on the same host, due to communication overhead.

Batch networks collect event data at defined intervals. Batch inspections allow more intricate examination and do not suffer from many of the performance issues inherent in real time processing. Since this technique places a delay between the intrusion and detection, it has been most appropriate to low threat environments. This style of review places a lower processing load on sensor modules, and allows storage overhead (which may be significant) to be central. Lastly, the availability of historical information surrounding an intrusion may greatly simplify the repair and strengthening of protection vulnerabilities.

Intrusion Detection and Dynamic Protection mechanisms offer a number of benefits to an entity:

- Intrusion Detection networks may offer protection for other protection mechanisms. In many cases, an assault will target protection mechanisms directly. Intrusion Detection networks may trigger alerts, allowing the issue to be repaired and minimize damages to the system.
- Intrusion Detection networks allow network administrators to form a clearer view of what the true protection state of their networks has been. Audit trails and network logs often contain important information, but have generally been in a format that have been unusable to all but the most expert of users.

- Intrusion Detection networks have been designed to extract information useful in tracking intrusions.
- In addition to recognizing the origin of exploitation, IDS may often identify the exact nature of that exploitation. This allows for mitigation of the effects of such exploitation, and to update procedures and defenses to prevent future recurrences. Intrusions commonly include the modification of system files to facilitate future access and to erase logs of the intrusion.
- There have been a number of complications in using computer generated logs in legal proceedings. Must the need arise to prosecute an intruder, the data held in IDS logs may be more likely to offer adequate evidence – particularly if the IDS was designed with this aim in mind.
- Intrusion Detection software may be able to recognize network failures. Many assaults have been based on creating illegal input to networks.
- When combined with network protection scanners and similar software, IDS may identify protection issues in networks before they become dangerous. For example, finding out that a firewall has been weak to a specific assault while configuring protection allows for early preventative action.
- IDS networks may help to identify which assaults have been used against your networks, and what network assets have been being targeted. This allows network administrators to boost protection where it has been needed, instead of where it may be needed.
- Every month, new vulnerabilities are discovered. Detection software comes with extensive libraries of intrusion signatures. This relieves the network administrators of the responsibility of keeping track of what new assaults might be implemented against them.
- Keeping track of the protection of a network has been an intricate task. IDS products have implanted knowledge on network protection, which allows less specialized administrative staff to maintain network defenses successfully.
- In order to use Dynamic Protection software successfully, the organizational protection strategy must be soundly developed. By offering detailed information on the protection status and behavior of a network, IDS may help in establishing a comprehensive protection strategy.
- In addition, many Dynamic Protection software include recommendations giving guidance in formulating and refining a Protection Strategy.

Clearly, Intrusion Detection networks offer a number of advantages in terms of network protection and management. However, IDS does not offer a complete solution to network protection. In specific, there have been a number of limitations and issues that restrict the usefulness of current IDSs:

- An IDS may not stop ongoing intrusions. While an IDS may be capable of detecting an intrusion while it has been occurring, it has been essentially a reporting tool – it may not disconnect abusive connections. Many current IDS claim the capability of blocking intrusions, but these capabilities generally depend on manipulating other protection mechanisms already in place (for example, having a firewall block a specific IP address)
- An IDS may not track intrusive behavior in environments with poor authentication and identification structures. If it is possible for a user to gain anonymous access IDS might be capable of isolating the intrusive behavior, but may not track back the intrusion. In addition, many intrusions consist of separate steps. IDS may be unable to correlate these steps if they do not have a common origin. In addition, current IDS networks suffer from false positive results thus adversely impacting legitimate users. Pattern based ways of recognizing users has not yet sufficiently matured to offer a solution.
- IDSs have been created to collect information on intrusions and attempt to track such behavior to its origin. However, due to the current nature of networking protocols and networks, the best an IDS may generally do has been to track an intrusion to its point of entry into the defended network. In the same way, IDS will attempt to identify the nature of an attack. However, it will often be impossible for an automatic network to fully comprehend the nature of an assault. Therefore, while an IDS has been an invaluable tool diagnosing an assault, human expert will generally be required for incident handling.
- In order to fully defend an entity, an IDS must be aware of the protection strategy of that entity. Every IDS has a specific mechanism for categorizing adequate and inadequate behavior, initially based on a general, baseline methodology. Unless an IDS has been specifically configured to recognize specific actions as intrusive it will not flag those actions. For example, browsing through other users' files may be against an entity's protection strategy, but it will not generally trigger an IDS response.
- Intruders have been very aware of the presence of IDS capabilities on a network, and will often directly assault such networks. IDS may not operate correctly if the information it receives has been manipulated. Must an intruder succeed in disabling an IDS sensor, the network will, at best, retain records up to the loss of contact. A more dangerous scenario has been where an intruder takes over and impersonates a sensor: no alert will be generated from losing contact, and an intruder may then feed random information to the monitor.
- IDSs generally depend on monitoring all traffic on a network segment, or all of the event logs. With the current increasing use of network bandwidth, it has been becoming impossible for any machine to properly monitor a network link under heavy load. This implies that some parts of an assault may be missed. A similar issue has been the increasing use of switching technology in networks – where an IDS sensor would have to be implanted into the switch hardware in order to ensure that it may filter all traffic. One possible solution to this has been to place IDS sensors on natural network bottlenecks.

- In order to recognize assaults, IDS has to model the effect of an intrusion on the networks it has been defending. Specifically, since different networks respond differently to the same intrusion, it becomes impossible for IDS to accurately predict the effect of any given sequence.
- New assault techniques have been continually discovered. Current IDS networks have limited capabilities for detecting assaults that differ significantly from previously known assaults. AIDS may have some success in detecting such assaults, but IDS software must be updated and maintained continually to ensure their coverage remains intact.
- Lastly, current IDS software suffers from scaling issues. Modern networks have been continually becoming larger, and assaults have been emerging that make use of distributed origins, and attacking wider groups of targets. Where the behavior observed by a local sensor might not reflect intrusive behavior, the global picture may be entirely different. As an example, consider a password cracking attack. IDS would likely notice a large number of failed authentication attempts with a common IP address. However, an intruder could spread the probes across a wide range of machines and networks. At any given single location, this would not be observed as intrusive behavior (a small number of failed authentication attempts has been generally adequate), but viewed across the entire network, this must be recognized as an assault.

Contemporary Intrusion Detection Networks

This section describes a number of software and procedures currently used in intrusion detection. Due to space constraints, only some of these softwares have been briefly described.

Manual review procedures

Full scale IDS networks may not be appropriate for many of smaller systems such as home networks. There have been a number of procedures available for adding detection capabilities to such networks.

The first way has been a specialized form of misuse detection. On a classic network, there have been a number of services that will not be in use (IMAP (143) or HTTP (80)). Any attempts to connect to such services would be considered suspicious.

Consider a network with the IMAP (143), SMB (139), and HTTP (80) ports. By connecting a dummy service to these ports, it appears to an intruder as if this has been a valid port. Connections to these ports would trigger a script which emails details of that connection to a protection officer.

Another way makes use of the log files and audit information already being gathered on the host. It essentially boils down to a host based anomaly detection network where any event not explicitly filtered has been reported.

These procedures, while corresponding to many of the classic mechanisms used in heavier IDS software, may markedly increase the effectiveness of the protection already in place on a network – which has been exactly the aim of intrusion detection.

SHADOW/CIDER [8]

The CIDER (Cooperative Intrusion Detection Evaluation and Response) toolkit has been a series of public domain software, aimed at automating the information gathering for intrusion detection networks. The SHADOW (SANS's Heuristic Examination network for Defensive Online Warfare) network has been constructed from these freely available components. As a result, setting up a Shadow intrusion detection network involves minimal cost and user expertise, while the network has been easily customizable.

Structurally, the Shadow network consists of a number of scripts layered on top of commonly available Unix software. The Shadow network makes use of a series of distributed Sensors, collecting traffic obtained on their local network segments. These sensors consist of Unix machines running *tcpdump*.

Under this methodology, an intruder would be unable to gain much information on the IDS capabilities of the network from a compromised sensor, and the individual sensors require minimal resources.

Network Flight Recorder (NFR)

NFR (available from <http://www.nfr.net>) has been one of the most discussed new IDS networks. It combines content based monitoring and a filtering mechanism.

Referring to the previous dialogue of IDS procedures, NFR has been a pure network monitor – with all the advantages and issues that entails.

NFR has been a real time detection network, processing packets during capture, whereas Shadow has been essentially an offline network.

Black ICE

Black ICE is a host based network IDS, running on a MS Windows platform. Even though it was aimed at the consumer market, it has a number of interesting features (from a technical point of view). These include an extremely simple installation, route tracing of intruders, and the ability to block intrusive connections [6].

Black ICE has been a hybrid between current IDS methodologies. By being host based, it has been capable of successfully defending that single host with a minimal performance cost.

Conclusion

Hopefully, this report has given the reader an overview of the field of protection software that do more than just keeping intruders out.

One issue that would appear to have been left behind in the development of this field has been the application of small scale, simple networks for detection. A wide assortment of powerful distributed networks have been available, but have too much administrative load and complexity for many small networks.

There have been IDS networks that promise lightweight sensors – trading off administrative complexity for power. At the other end of the scale, a number of powerful single point detection networks have been developed which have fallen out of fashion in recent years. There remains a definite split between networks that offer powerful detection, and networks that have been configurable.

The most powerful form of detection has been one that knows and understands the network it has been defending. Nevertheless, protection of computer networks may not exclusively rely on static defenses (firewalls). Increased number of trained protection guards must be available on a nonstop basis to defend the network as soon as it has been attacked.

© SANS Institute 2003, Author retains full rights.

Bibliography

[1] Computer Protection Institute "2002 CSI/FBI Computer Crime and Protection Survey", April 7, 2002.

<http://www.gocsi.com/press/20020407.html>

[2] RSA Laboratories "Frequently Asked Questions about Today's Cryptography".

<ftp://ftp.rsasecurity.com/pub/labsfaq/faq4pdf.zip>

[3] Herve Debar, Marc Dacier, Andreas Wespi "Towards a taxonomy of intrusiondetection networks", Computer Networks 31 (1999) 805822.

[4] Ulf Lindqvist, Phillip A. Porras, "Detecting Computer and Network Misuse Through the ProductionBased Expert Network Toolset (PBEST)", Proceedings of the 1999 IEEE Symposium on Protection and Privacy.

<http://www.sdl.sri.com/emerald/pbestsp99cr.pdf>

[5] Cisco Networks "NetRanger Intrusion Detection Network Technical Overview".

http://www.cisco.com/warp/public/cc/pd/sqsw/sqidsz/tech/ntran_tc.htm

[6] ISS, "Black ICE Defender User's Guide version 2.9".

<http://documents.iss.net/literature/Black ICE /BI Defender 29 User Guide.pdf>

[7] James P. Anderson, "Computer Protection Threat Monitoring and Surveillance". Technical Report, James P. Anderson Co., Fort Washington, PA, April 1980.

[8] SHADOW Homepage

<http://www.nswc.navy.mil/ISSEC/CID/>

© SANS Institute 2003. Author retains full rights.