



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: Network, Endpoint, and Cloud (Security 401)"
at <http://www.giac.org/registration/gsec>

The Power of Symantec ManHunt™
German Rincon
June 22, 2003
SANS GSEC Practical Assignment v. 1.4b

Abstract:

This paper is intended to give a wide range of information about how just implementing a Network Intrusion Detection System (NIDS), in a company concerned about its security is not enough, and how Symantec ManHunt™ a Network Intrusion Detection System (NIDS), can transcend the concept of this kind of security systems by having different detection mechanisms and features related to its main detection method which is protocol anomaly detection (PAD).

Introduction:

Since we have all taken advantage of Internet and its benefits being connected to the world, many of us don't realize how important network security could be. Firewalls deployed at the perimeter worked fairly well when there was limited interaction between internal and external networks, the internal users were trusted and the value of the network-available assets was limited.

Companies have purchased Firewalls as means to safeguard their Intranet. Unfortunately, Firewalls only offer a degree of perimeter access protection, but do not guarantee elimination of attack. In a dynamic environment, it is not unusual to discover miss-configured Firewalls or those with poorly defined policies. Furthermore, certain connections guarantee some degree of access (e.g. HTTP / FTP). Lastly, Firewalls do not prevent the use of modems as means to enter or leave a network.

However, there have been considerable changes in recent years. Network-aware applications and interactions between networks have greatly increased, and while access is being granted on a greater scale to these business-critical functions, attackers and their tools have become more sophisticated.

The nature of attacks on the Internet has changed radically since the design and implementation of traditional Intrusion Detection Systems (IDS). When these Intrusion Detection Systems (IDS), were developed, virtually all attacks across the network were intrusions, or break-ins. Conventional intrusions have relied on "stealth" to hide the malicious activity, as well as the identity of the intruder, until the damage has been done. The problem has always been to keep up with the latest attack methods in order to recognize that an attack is taking place; and the rate of new and modified attacks is ever increasing.

The concepts:

Intrusion detection products enable network administrators to develop proactive strategies to stop hackers or unauthorized users with malicious intent from misusing systems.

When we talk about Intrusion Detection Systems (IDS), management automatically assumed it is THE solution to all network, organization and social problems. Most people deal with this technology like it is a monolithic solution. This is not a good way to consider any security technology; it does not work like that. The majority fails to recognize that IDS' initial design and function is to protect the organization's vital information from an outsider.

As Mr. Dirk Lehmann at sans.org says:

“Sometimes, a distinction is made between misuse and intrusion detection. The term intrusion is used to describe attacks from the outside, whereas, misuse is used to describe an attack that originates from the internal network. However, most people don't draw such distinctions.”¹

Types of Intrusion Detection Systems:

Intrusion Detection Systems (IDS) detect network intruders and perform other important tasks. “Intrusion Detection Systems (IDS), can track user's activity from entry to exit, guard against known types of attack, detect network policy violations, and keep tabs on normal network activity, making abnormal behavior easier to spot”². Intrusion detection systems (IDS) are divided into three main categories: Host-based (HIDS), Network-based (NIDS) and Decoy-Based IDS.

Host-based Intrusion Detection Systems (HIDS):

There are two Host-based Intrusion detections systems: application specific and operating system-specific. In both types, an agent generally runs on the server being monitored, and analyzes log files, access records, and application log files.

Anomaly detection modules, which are based on statistical comparisons to normal patterns, are typically used on Host-based systems. In the case of operating system monitors, abnormal sessions, such as unsuccessful logins which are compared to a behavioral model of normal usage using criteria, such as time of access and the number and types of files created and accessed.

¹ Lehmann, Dirk. “What is ID?”

² Randall, Neil. “Intrusion Detection Systems: Who is in Here?”

One example of what these systems do is, it takes a snap shot of your existing system files and matches it to the previous snap shot. If the critical system files were modified or deleted, the alert is sent to the administrator to investigate.

Network-Based Intrusion Detection Systems (NIDS):

Network-based intrusion detection systems have the benefit of potentially analyzing all layers of the network communication. These tools can reside on their own servers; this way they can eliminate performance hits on the application server(s). They can also use a rule base to describe common attack techniques.

Decoy-Based IDS:

Decoy systems or honeypots as they are commonly known, provide additional level of security within the network infrastructure. They are considered systems “set and forget” IDS sensors because they are set as systems where the whole purpose is to capture unauthorized activity. This means any packet entering or leaving a deception system is suspect by nature.

Methods of Intrusion Detection:

There are several technologies to detect malicious activity. The three of the most widely distributed are:

Signature-Based Detection:

Most of the IDS products on the market are based on a system that examines every packet on the network traffic and compare it to every signature, for a specific pattern of attacks. This means that for every exploit, the IDS vendor must create a signature for that attack in order to detect it, meaning that the attack must be known.

As network speed increases, the resource that the IDS sensor uses to look at every packet decreases causing some packets to be discarded and therefore allowing attacks to slip by undetected. In addition, to this, it is also important to consider the amount of time it takes to the IDS vendor to identify the attack, create a signature and release an update.

Behavior-Anomaly Detection:

A less common method of intrusion detection is the ability to detect statistical anomalies. Examples include detection of excessive use, detection of use at unusual hours and detection of changes in system calls made by user processes.

Taking into consideration that anomalies can be detected without having to understand the underlying cause, even legitimate use of the system could trigger anomalies and therefore lead to a very high number of false positives.

Protocol-Anomaly Detection (PAD):

“One of the key differences between anomaly detection and other forms of detection is that, rather of defining what is not allowed or bad, it defines what is allowed and good”³. This detection is performed at the application protocol layer. It focuses on the structure and content of the communications and is designed to analyze a protocol and requires defining a model of such protocol normal usage. The model can be defined as the rule of normal usage for it. Any use of the protocol outside of this model can be considered as an anomaly. When protocol rules are modeled directly in the sensors, it is easy to identify traffic that violates rules such as unexpected data and extra or invalid characters.

In a network based Intrusion Detection System (IDS), anomaly filters would disassemble the data packets for each network protocol, and check if they are built in compliance with the protocol standards, as described in RFCs (Request for Comments), or equivalents. However, protocols are seldom implemented according to their standards, and anomaly filters should thus be designed in a flexible way, in order to fit not only to the official usage of a protocol, but mainly to a given model of a protocol's 'normal' usage. Building such a model requires analyzing common protocol implementations in practice, in order to define the limits of what is officially and unofficially the standard for this protocol.

Protocol anomaly detection systems are theoretically faster than signature based Intrusion Detection Systems (IDS). They can potentially identify most of the attacks, including new and unknown ones (Zero-Day Attacks), without requiring attack-dependent knowledge. While longer to develop than signature filters, they however will not require regular updates. Another thing to have in mind when implementing this type of systems is reading the alerts, which is a difficult task, since they do not provide clear information about the nature of the threat. They should therefore be monitored by experienced personal.

Looking at the playground from this point of view, it really leads us to one common question and that is, what should be better, a Protocol Anomaly Detection (PAD) Intrusion Detection System (IDS), or a signature-based Intrusion Detection System (IDS)?

Well, just as an antivirus depends on its virus definitions in order to catch viruses in a file, Signature-Based Intrusion Detection Systems (IDS), try to match the packet payloads to a database of “signatures” to detect an attack. When the Signature-Based IDS, discovers a match, then an event is reported. This is very interesting since taking into consideration the dependability of these kind of Intrusion Detection Systems (IDS), and its ability to update the signature database we could think that they could only detect known attacks. All this without taking in mind the process of comparing every single event on the wire to a signature database and the latencies involved in those communications, really limit the amount of traffic that can be monitored.

³ Hernacky, Brian “Why companies should PAD their networks”

Based on this kind of methods of Intrusion Detection system (IDS), lets concentrate in one of the major features of Symantec ManHunt, which is exactly its ability to detect what Symantec calls “zero-day” attacks, with this kind of method. Besides, these filters are not able to detect the few attacks that cannot be considered as protocol anomalies. It is advised to use them in collaboration with signature filters.

Symantec ManHunt:

First of all, lets take a look at how Symantec Manhunt works in the network, its components, and the way it should be implemented.

Within a network, multiple Symantec ManHunt nodes (a system running Symantec ManHunt), can work together as a Symantec ManHunt cluster and share event data. A Symantec ManHunt cluster can be formed from up to 100 Symantec ManHunt nodes across multiple network segments within multiple network locations.

By default, the first Symantec ManHunt installation is designated as a primary master node, and all other ManHunt nodes within the cluster are designated as slave nodes.


Symantec ManHunt analyzes traffic on the network by using event data from multiple sources like high performance sensors, which allow Symantec ManHunt to monitor many ports. The sensors use switch port analyzers (SPAN), to listen to network flows that are directly attached to the sensors by copying all of a particular port’s incoming or outgoing traffic to another port. This enables sensors to monitor 100% of the traffic on the ports they are monitoring. In order to understand this concept of SPAN, lets take a look at this brief example:

Switch Port Analyzers (SPAN):¹⁰


The SPAN feature was introduced on switches because of a fundamental difference they have with hubs. When a hub receives a packet on one port, it will send out a copy of that packet on all ports except on the one where it was received. After a switch boots up, it will start to build up a Layer 2 forwarding table based upon the source MAC address of the different packets received. Once this forwarding table has been built, the switch forwards traffic destined for a MAC address directly to the corresponding port.

¹⁰ “Configuring the catalyst port analyzer (SPAN)”

On a switch, after host B's MAC address is learned, unicast traffic from A to B is only forwarded to B's port, and therefore not seen by the sniffer:



An extra feature is needed that will artificially copy unicast packets sent by Host A to the sniffer port:



In this above diagram, the sniffer is attached to a port that is configured to receive a copy of every single packet that is sent by host A. This port is called a SPAN port.

Another source to highlight from where Symantec ManHunt analyzes traffic on the network is FlowChaser sensors. FlowChaser sensors receive information about network flows from various devices (Cisco Router, the sniffer, etc.) and the FlowChaser Database in Symantec ManHunt stores the data to accelerate the TrackBack attack response. The TrackBack function is designed to track a data stream to the source within the cluster, or, if the source is outside the cluster, to its entry point into the cluster (Figure 1.2). It does this by using its sensor resources to gather information from switches and routers, systematically looking for the data stream with matching characteristics. Symantec ManHunt uses its knowledge of the network topology to make choices as to which devices to interrogate about the attack stream.




Figure 1.2 TrackBack Function.⁷

In order for Symantec ManHunt node to monitor a network, it requires information regarding the topology of its ManHunt cluster, that is, the topology of the network or portion of the distributed network in which it resides and from which it can gather information. Each ManHunt node also requires relevant data regarding connections to other networks, whether they are autonomous systems or other portions within a distributed network. From the administration console, there is a way to enter data about the network topology.

To build a network topology, there must be access to the network, and determine which devices will be monitored. Once the network map has been created, nodes can be added to the administration console. The following are brief descriptions of the types of nodes that can be added to the administration console:⁷

- Location Nodes—A location node represents any physical or logical grouping of network segments.
- Symantec ManHunt Nodes—A Symantec ManHunt node represents the software that is running on a single machine. By default, the first ManHunt node installed in a cluster propagates database and configuration changes to all other nodes within the cluster.
- Switches, hubs—Switches are further categorized as those that support SMON (switch monitoring), those that do not support SMON, and those that ManHunt are configured to work essentially like hubs (non-steerable switches).
- Copy Ports—A copy port provides a link between a switch or hub and a ManHunt device so that the ManHunt device can listen to traffic on the switch or hub. When connected to a switch, Symantec ManHunt sensors use switch port analyzers (SPAN) to listen to network flows that are directly attached to the sensor by copying all of a particular port's

⁷ "Symantec ManHunt 2.2." Administration Guide.

incoming or outgoing traffic to another port. This enables it to monitor the traffic without slowing the traffic.

- Interfaces—All interfaces attached to devices defined in the topology tree.
- External Sensors—Symantec ManHunt has the ability to accept input from other external sensors (Smart Agents).

Symantec ManHunt Detection Mechanisms:

Symantec ManHunt includes several detection mechanisms including protocol anomaly detection (PAD), Stateful Signatures, custom signatures, DoS Detection, IP traffic rate monitoring, IDS evasion detection, and IP fragment reassembly, to detect attacks (Figure 1.0). Based on this introduction, let's take a look at how each one of these mechanisms work.

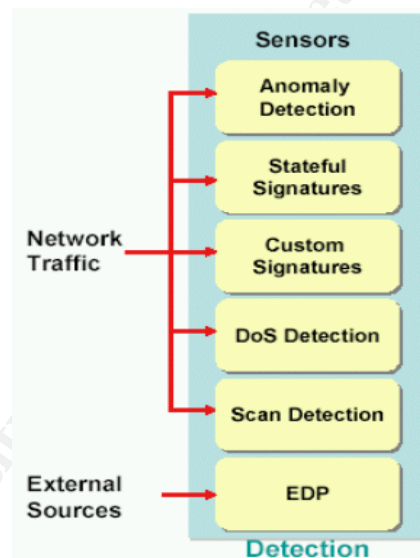


Figure 1.0 Symantec ManHunt Detection Mechanisms.⁵

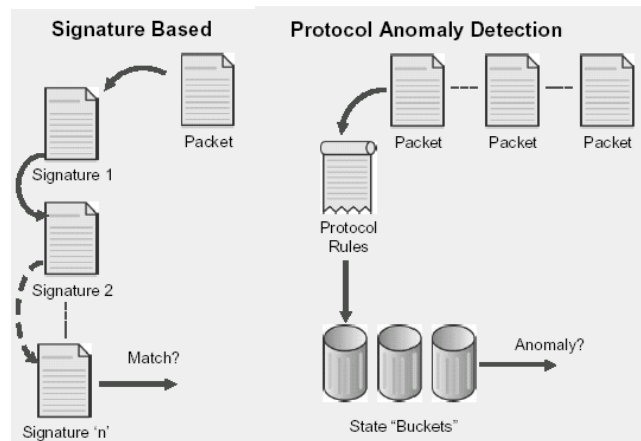
Protocol Anomaly Detection (PAD) with Symantec ManHunt:

One of the most important characteristics of Symantec ManHunt is that it does not rely on a signature database only to identify network threats. Taking into consideration the process of comparing every network event to a signature database, and the latencies involved in those communications the severely limit the volume of traffic that can be monitored. Symantec Manhunt's primary approach to identify threats on the network is through protocol anomaly detection (PAD), this is done by comparing data to the protocol model on each state and

⁵ Hill, Steve "Symantec ManHunt 2.2"

looking for activities outside of the normal behavior for that protocol. It also, understands and is able to gauge the severity of the anomaly. By this, Symantec Manhunt is able to detect novel attacks and not just known attacks.

Anomaly Detection Intrusion Detection Systems (IDS) examine entire flows and identifies violations in the protocol. By doing this, Anomaly Detection Intrusion Detection Systems (IDS), can detect not only known, but also unknown attacks.



Protocol Anomaly Detection vs. Signature-Based Detection

Because most of the attacks (buffer Overflows, authentication exploits), exploit application layer protocols, the protocol anomaly detection method in Symantec ManHunt looks for the anomalies in the requests and the responses in this kind of traffic. It detects when a protocol stream does not match the protocol definition and flags the flow as being unusual. For example, if an http request results in a shell prompt, then the protocol has been violated and Symantec ManHunt flags the event. If an attacker attempts a buffer overflow by sending 8000 bytes across when only 512 are allowed, the protocol has been violated and Symantec ManHunt flags the flow.

Another example of Protocol Anomaly Detection (PAD), detected by Symantec ManHunt is the imapex2 exploit, which is a common IMAP 4 buffer overflow protocol. In a buffer overflow, the attacker sends more data than the acceptable in the protocol standard. The victim allocates buffer space based on the attackers request. "When more data is received, the buffer is overrun, after this it is possible to grant remote access or process termination".⁴


Example (Imapex2 – A common IMAP4 exploit):

```
S:* OK recourse.com IMAP4rev1 v12.264 server ready
C:* AUTHENTICATE {67}
S:+ Ready for argument
```

⁴ Hanson, Jeffrey P. "Microsoft Outlook / Outlook Express..."

attacker never completes the process by failing to send the ACK, leaving the session in an invalid state. To formulate this into an actual attack, the attacker repeats this process until the victim runs out of buffer space and the process aborts. This is an example of a Denial of Service (DoS) attack.

Example of a Syn Flood attack:



Based on this, Symantec ManHunt has the ability to detect this type of attacks by sorting the SYN packets into queues.

Scan Detection:

Symantec ManHunt has the ability to detect if more than certain connections to adjacent ports are attempted within a limited amount of time, a scan has probably been initiated. For example if Symantec ManHunt detects a variety of attempts to connect to ports 80 HTTP, 23 telnet, 25 SMTP, 111 RPC and port 22 SSH on one victims machine, or, several systems within few seconds, it will flag the event as an IP sweep.

Stateful Signature Detection:

Higher levels application logic attacks cannot be detected with state models because they do not actually violate the protocol. For this kind of attacks, Symantec ManHunt uses a signature database. However, Symantec ManHunt does not compare every exploit to every signature in the database.

One example of this kind of detection is an attacker a CGI exploit to gain access to host A, a Web server. Symantec ManHunt detects that the packet was an HTTP packet and then performs signature matching within the HTTP signature file to detect the attack.

Custom Signature Detection:

Symantec ManHunt has the ability to add customized signatures on what Symantec calls Hybrid mode. Hybrid mode allows Symantec ManHunt to run custom signature detection on incoming data in addition to protocol anomaly detection.

This mechanism is useful to write custom rules for specific company security policies; for example, a company could write a custom signature to detect employees who are using file sharing or instant messenger programs. “However when Hybrid Mode is activated, ManHunt sensor performance may be slightly degraded due to the processing-intensive nature of signature detection”.⁷

The signature file uses a subset of Snort 1.8 signature language. However, some options behave different when signatures are compiled in Symantec ManHunt different than compiled with Snort.

IDS Evasion Attack Detection:

Taking into consideration that attackers usually use a ‘smoke screen’ to split attacks to pass traditional Intrusion Detection Systems (IDS), For example, an attacker may launch a SYN flood, and, while the IDS is trying to keep up with the SYN packets, send a CGI exploit to gain access to the server. The way Symantec ManHunt detects these types of attacks is by detecting a hash for each event, which includes the events’ contents and the destination IP. Then the hashes are sorted into queues by type.

Event Dispatch Protocol (EDP):

In addition to the entire network traffic detection methods, Symantec ManHunt uses an EDP (Event Dispatch Protocol) Proxy to receive event data from “*external sources*” called Smart Agents, and correlate that data with all other events. This allows Symantec ManHunt to correlate events not only with some kinds of firewalls (Checkpoint and NetScreen), but even from others Intrusion Detection Systems (IDS), like: Cisco IDS, Snort, Tripwire, Okena Stowatch, Enterasys Dragon IDS and, ISS RealSecure. Since Symantec ManHunt is monitoring event information across the network, it can identify a threat on one part of the network and gather information from additional devices to monitor the network without having to inspect all traffic.

Reporting:

In addition to these mechanisms of detection, one of the strong key points of Symantec ManHunt from my point of view is the ability to organize data by incidents from its sensors, which are made up of events. For example, if an attacker probes port 21 on an ftp server a new incident is created containing a port sweep event. If that same attacker, five minutes later tries an anonymous login because he found the port open, the event will be added to the same incident, if it is still active. The user can define how long an incident remains active. As long as an incident is active ManHunt will continue to add events to it. After an incident is closed, a new event will cause a new incident to be created.

⁷ “Symantec ManHunt 2.2.” Administration guide.

Conclusion:

To be effective, a network security solution must be made up of several layers that address the various types of threats faced by today's networks. Intrusion Detection Systems (IDS) will not pick up every attack, no matter what kind of system the company has deployed. However, remember that signature-based IDS do not detect new attacks and since protocol anomaly systems can detect many new attacks like Code Red, Code Red II, and Nimda, corporations should, at minimum, be able to strengthen their defenses at the gates to their networks.

From my point of view, Symantec ManHunt is a well-built detection system that will rely on multiple detection mechanisms (as seen), each one of them covering some portion of the threat space. Organizations concerned that the protocol anomaly detection system may not detect certain threats should consider a solution that provides additional forms of detection in complement to anomaly detection and as we could see it, Symantec ManHunt counts on several detection mechanisms that can mitigate most of those concerns.

References:

- [1] Lehmann, Dirk. "What is ID?" Intrusion Detection FAQ. 16 April 2003.
URL: http://www.sans.org/resources/idfaq/what_is_id.php (3 June, 2003)
- [2] Randall, Neil. "Intrusion Detection Systems: Who's in Here?" 5 December 2000.
URL: <http://www.pcmag.com/article2/0,,110115,00.asp?kc=PCAV10209KTX1K0100360>
(3 Jun. 2000)
- [3] Hernacky, Brian. "Why companies should PAD their networks" 23 January 2003.
URL: <http://www.computerworld.com/securitytopics/security/story/0,10801,77813,00.html>
(4 June, 2003)
- [4] Hanson, Jeffrey P. "Microsoft Outlook / Outlook Express GMT Field Buffer Overflow Vulnerability" 10-15 December. 2000
URL: http://www.giac.org/practical/Jeffrey_Hanson_GCIH.doc (10 June, 2003)
- [5] Hill, Steve. "Symantec ManHunt 2.2." Technology Overview. 19 March 2003.
URL: http://www.wizardsecurity.net/manhunt_summit.ppt (9 June, 2003)
- [6] ZyXEL Communications Corporation. 2000
URL: http://www.zyxel.com/support/supportnote/zywall10/faq/fw_faq.htm#8 (8 June, 2003)
- [7] Symantec. "Symantec Manhunt 2.2." Administration Guide. 2002
URL: ftp://ftp.symantec.com/public/english_us_canada/products/manhunt/2.2/manuals/manhunt_admin.pdf (2 June, 2003)
- [8] Lemonnier, Erwan. "Protocol Anomaly Detection in Network-Based IDS" 28 June. 2001

URL: http://erwan.lemonnier.free.fr/exjobb/report/protocol_anomaly_detection.pdf (7 June, 2003)

[9] Symantec. "Symantec Manhunt 2.2." Installation Guide. 2002
URL: ftp://ftp.symantec.com/public/english_us_canada/products/manhunt/2.2/manuals/manhunt_intall.pdf (2 June, 2003)

[10] "Configuring the catalyst switch port analyzer (SPAN)." 14 May. 2003
URL: <http://www.cisco.com/warp/public/473/41.html#prereq> (17 June, 2003)

[11] Symantec. "Defining Protocol Anomaly Detection." 2003.
URL: <https://enterprisesecurity.symantec.com/content/displaypdf.cfm?SSL=YES&EID=0&PDFID=346> (8 June, 2003)

[12] Das, Kumar. "Protocol Anomaly Detection for Network-based Intrusion Detection." (13 August. 2001)
URL: <http://www.sans.org/rr/papers/30/349.pdf> (7 June, 2003)

© SANS Institute 2003, Author retains full rights.