

[WP-003]

Whitepaper: Hiding an Intrusion Detection System (IDS)

A Theoretical Discussion on How to Play “Hide ‘N Go Peek”

Version 040305

March 2004

Author: Bob Radvanovsky, rsradvan@unixworks.com

(A special thanks goes to those listed for being my “sounding board” on this project.)

Copyright © 2004 Bob Radvanovsky. All rights reserved.



Limited Liability Statement

In no event shall the author(s) be liable for any found errors contained herein or for any direct, indirect, special, incidental or consequential damages (including lost profit or lost data) whether based on warranty, contract, tort, or any other legal theory in connection with the furnishing, performance, or use of this material.

The information contained within this document may be subject to change without any notice. No trademark, copyright, or patent licenses are expressly or implicitly granted (herein) with this whitepaper.

Documentation pertaining to any security-related technical or proprietary information, its data and all information provided and contained within this document is considered proprietary in nature and subject to copyright protection and is intended solely for use by its owner. Additionally, this documentation is solely for the purpose of discussing managed and timed proxy servers that are heterogeneous to any networked environment, and are not dependent upon any specified architecture, hardware platform or its software.

The name "LINUX" is a registered trademark of Linus Torvalds.

The name "UNIX" is a registered trademark of The Open Group.
[ref: <http://www.opengroup.org/legal.htm#trademarks>]

The name "Sourcefire" and "Snort" are registered trademarks of Sourcefire, Inc.
[patent pending] [ref: <http://www.snort.org>]

All other product names mentioned herein are trademarks or registered trademarks of their respective owners. NOTE: Any names not outlined or mentioned above are fictional in nature; as such, any relation to any name or trademark (if any) is purely coincidental.

Introduction

This document is an abstract notion in that it may be possible to (quite literally) "hide" an intrusion detection system on the secured-side of any given network. Without going into detailed lecture specific to monitoring both external and internal network traffic, intrusion detection systems are seeing a reintroduction into the commercial network as viable network tool. Reasons for IDS environments may be partially due to recent events following that tragic day of 9/11, or because of surmounting cyber threats through increasing Internet connectivity and its growing co-dependence. Nonetheless, IDS is on the rebound, mostly now as it has mutated into an even more useful tool through the combination of some self-intelligence and firewall prevention through another newly introduced technology called "intrusion prevention systems" (or "IPS"). Product manufacturers of this technology have several variations for their name, but essentially, it is actually two technologies combined: an IDS console (and its sensors), and a firewall. Thus, the term "IPS" to many who've worked with and been around since the first inception of IDS technologies, think/feel that "IPS" is just another spin-doctored/marketing job to push more product that may not necessarily work in today's heterogeneous networked corporate environments.



As IPS has been noticed by corporate executives, IDS environments are being revisited, too. With common network attack signatures being recorded and analyzed faster than ever, IDS environments are gaining popularity again, but this time, not necessarily with the larger corporate environment. It is with the small or medium sized company that does not have that large or deep budget to support their enterprise, unlike what most large corporate networks would require. Many industries recently, such as the healthcare and financial industries, have recently come under scrutiny as they will soon be required to comply to very complex and strict privacy laws recently passed by the United States Congress, requiring companies within the supply/service chain to compliance in some form or another. Thus, small or medium sized businesses that work with or support these two industries are starting to see IDS as one possible tool for their compliance use.

IDS in the Smaller Corporations

Purchasing (perhaps) one IDS console/server configuration would consume the entire year's budget for many small companies; thus, this document will emphasize the utilization of an Open Source solution: SNORT. SNORT remains a freely available product from SourceFire, Inc. Several applications have been developed over the years that work in conjunction with SNORT, thus providing additional capabilities for a better, more effective management of an IDS environment. SNORT has been cumbersome and difficult to work with in recent versions, but has made great strides in its management capabilities, mostly through several third-party products: SNORTCENTER and ACID. Utilizing a MySQL database back-end as the primary storage mechanism, SNORTCENTER and ACID combined, provide a very effective front-end using SNORT. Recently though, a group of developers have taken away the difficulty of installing and configuring the entire environment and created a much-awaited package distribution. It's called SENTINIX, and is built using GNU/Linux (aka Debian), which includes several packages: SNORT, SNORTCENTER, ACID, PHP, Apache, MySQL, Postfix, Mailscanner, SpamAssassin, Nessus, Cacti – with more added in the future.

As with any IDS environment, data gathered and stored on the IDS console server is vital to the business, and depending upon the company's data retention policy, will determine the level of criticality and what the business must do to ensure the safety of the data. Obviously, retention of this data may be necessary for forensic management as part of an investigation, or would be needed for problem determination, etc. Nonetheless, the safety of the IDS console server is important.

Since the IDS server is statically assigned an IP address, an idea came to mind. Since over 70% of all network-based attacks originate from internal sources, wouldn't it make sense to protect the IDS server from the inside? Aside from use of protecting the IDS with a firewall (which would now make it an IDS-hybrid firewall, or IPS), what if the IDS had some additional intelligence added to itself? One of the more common methods to signify a precursor to a possible network-based attack scans for open ports either for specific addresses, subnets, or all addresses for the entire network. The entire idea behind having a IDS is the preventative measure of passively monitoring such attacks, even ones as common as the port scan.



Blackbird IDS?

Stealth is the primary concern for the IDS environment. Stealth of the sensors is necessary, but (using the analogy of mining gold) not only is the gold mine important, but the transportation of the raw gold ore, the processing center, and finally, the gold vault. It is that vault that contains the items that so many want, similarly the same way that an intruder would know that his/her footsteps were heard walking into someplace where they didn't belong. Thus, the intruder would want to remove any proof of their trespassing, making any evidence of their attempts or network reconnaissance as valuable as gold itself.

So what am I pushing here? Quite simply, the notion that perhaps it is feasible to introduce yet-another-stealth method by utilizing DHCP or some pooled configuration of IP addresses to effectively protect the IDS server when it comes under heavy attack. This becomes useful when an internal attack against a critical piece of the network infrastructure cannot easily migrate to another IP address easily. Or can it?

A close friend of mine and myself attended a sales demonstration of another newly introduced commercial-grade firewall/IDS/VPN/all-in-one security product from one of the network security manufacturing vendor's presentations recently. Despite efforts to block external-going-into-internal network traffic, there appeared to be very little preventing internal traffic from attacking this server. The two of us looked surprised at each other, and both of us muttered the same thing: hack attack from the inside. How would this work? What measure or level would be necessary to ensure that this does not happen again? Corporate executives continue to think and operate at a level that it is necessary to protect and maintain perimeter defenses of the corporate environment at any and all costs. What these executives continue to fail with is that, although the business goals and directives continue to remain steady and consistent, the perspective towards newer and emerging technologies continues to be extremely lacking. A prime example is the emergence of the wireless access point (if you will, a wireless "router") into the corporate network environments at the local, departmental, or even sub-departmental levels. This poses serious risk and, depending upon the location of the access point, could cause serious detrimental and perhaps fatal, consequences through such an insecure introduction into any given corporate network. This lack of perspective at a departmental (or perhaps even at a sub-departmental) level demonstrates the ill-preparedness of most large corporations today, and how well they can not only manage and maintain their environments, but provide effective and immediate resolution in a critical scenario (as outlined above).



With this resulting mindset, this permeates and propagates throughout the corporate social and political structure. This ensures an almost guaranteed success rate for failure in the event of a security catastrophic event, such as that soon-to-be-internal breach from within that corporate network environment that didn't pay attention closely to their departmental networks, such as through the injection of malicious network attacks via the not-so-widely-known access point that was interconnected by one of the local network or systems administrators. As the corporate environment grows, so does the risk to greater exposure of either continued or added risk of the environment. Risk assessments in today's networked environments are becoming standard for companies that support critical infrastructure (financial, healthcare, transportation, food production/processing, utilities [electric, gas, water, sewage], municipalities, etc.); however, companies want to find more effective means of not only monitoring, but automating the monitoring process down to a simplified graph or report that corporate executives could easily digest.

But just like any other well-intentioned plans, there are always flaws, despite what may come from it. The measured environment would be "smart enough" (if you can program an IDS to "think" that it is under attack) to have some form of self-preservation, yet continue to operate, collect, sort, and maintain the IDS data within its "gold vault".

How the Process Works

The process is quite simple: if the IDS server receives an excessive number of network packets that appear to be malicious or have intent of subterfuging the environment, the IDS shuts down its network links to the internal/secured-side of the network, whilst continuing maintaining its connectivity with its sensor units, either remotely or locally.

After a period of time when the attacks have either subsided or migrated to some other location, the IDS server will re-establish its server connectivity to the internal/secured-side of the network. If, after reintroducing itself back into the corporate network environment, the attack of subterfuging resumes, the IDS shuts down its network links again to the internal/secured-side of the network, issues or reassigns a new IP addressed number, either within the same subnet, or some other network location, either through a DHCP server for a "special IP address pool" that is assigned by the DHCP server, or a random "address pool" internally within the IDS server, re-establishing its network link again, but in a different location.

After the IDS network links become active, a notification is sent to either to the IDS administrator or administration group, usually via pager or internal email, in a plaintext cryptic message that states that the "for a good time, call me at 215-80" (the message is obvious that the 2 numbers are the last 2 octets from an IP address) message – or something similar to that effect. In many cases, if a more plain-as-daylight, yet obscured message is sent, one at a time, to everyone concerned, it will only appear as a violation had occurred for those individuals in which the email was sent, rather than a notification that the IDS server was just under attack, and now resides in a different virtual location of the corporate network.

It's that simple!!!



All of this technology can be easily integrated with already existing technologies, thus the cost of implementing such an endeavor is much lower than introducing something that would require a much more significant implementation. As we promote and embrace the “Open Source Initiative”, this too, can be easily implemented using already developed products that could simply be patched using several applications currently available (with permission from their authors, of course).

For this whitepaper, since SNORT was mentioned, it would have multiple network links to several locations, preferably secured. If a malicious or subterfuged attack commences on only 1 link, then risk is minimal, and the IDS server observes network traffic based upon the intent of the attacker. If it continues, then within the IDS server, separate from the SNORT environment, another application would monitor for port scanning activities, et. al. Once a determination is made that a malicious or subterfuged event was occurring, this would then cause rise to alarm and trigger the necessary course of action, thus protecting the IDS server.

Another analogy would be the turtle and how it meets its threats in the animal kingdom. When the turtle comes under attack by either a predator or threatening scenario, it pulls its head and legs into its hardened shell. It waits until the threat leaves. Try as much as the predator or threatening animal may want to the turtle, but the likelihood that the turtle will be harmed is slim to none. Once the threat diminishes, the turtle’s head and legs reappear again. If threat comes back, either from the same predator or threatening animal, or even from a different one, the analogy is that the turtle would utilize one of its other openings to peer its head and legs through there, thus while the threat exists elsewhere, the turtle can safely peer its head and legs through that safer opening. Although it is a strange thought to think that a turtle could have that large or spansive area to peer its head and legs elsewhere, it’s that idea that counts, right? Nonetheless, it is that very concept that playing a form of “shell game” with the would-be attacker into thinking that the IDS server’s links would even remotely show up on the same network now has introduced a guessing probability into the arena, thus this randomness factorization helps reduce the risk perhaps, provided that the IDS does not reappear onto the same subnet that is currently under attack!

If the would-be attacker were to accurately guess where the IDS would show up next, this would pose a serious risk to the IDS server; however, given the fact that many larger enterprise corporate network environments, which consist of tens of thousands of addresses, would make the “hide ‘n go peek” that much more challenging to the attacker(s).

In closing, existing technologies could be elaborately combined together to create a single-point distribution, thus permitting the game of “hide ‘n go peek” with would-be attackers. These tools would consist of some of the more common utilities (such as DHCP), but would be part of something of a much grander/larger scale; in this case, the DHCP server is one of several utilities that would be applied towards this newer way of stealthy thinking and operations. Utilizing an aggregation of commonly applied technologies would be cost-effective in its implementation, provided that the corporate executives would be willing to agree to such a mechanism.

Because of this way of thinking, we can expect “stealth IDS” environments in the future.

