

Security System Design for Productivity: An Introduction

J Wolfgang Goerlich
Written December 2009

Abstract

This paper discusses measuring Information Security (InfoSec) productivity using the $P > D + R$ formula. It then delves into techniques for improving detection time and accuracy. Conveying information using both visualization and sonification techniques are considered.

Creative Commons Copyright and Use Notice

You are free: to Share -- to copy, distribute, display, and perform the work; to Remix -- to make derivative works; under the following conditions:



Creative Commons Attribution-ShareAlike 2.5
<http://creativecommons.org/licenses/by-sa/2.5/legalcode>



Attribution. You must attribute the work in the manner specified by the author or licensor.



Share Alike. If you alter, transform, or build upon this work, you may distribute the resulting work only under a license identical to this one.

For any reuse or distribution, the reader must make clear to others the license terms of this work. Any of these conditions can be waived if the reader obtains permission from the author.

Information in this document is subject to change without notice. Companies, names, and other information used in examples herein are fictitious unless otherwise noted. This document is for informational purposes only. The author does not make any warranties, express or implied, in this document.

Microsoft Exchange, .Net, SharePoint, SQL Server, Terminal Services, and Windows are registered trademarks of Microsoft Corporation.

Citrix WinFrame, MetaFrame, Presentation Server, NFuse and Web Interface are registered trademarks of Citrix Corporation.

All other trade names referred to are the Servicemark, Trademark, or Registered Trademark of their respective manufacturers.

Table of Contents

Security System Design for Productivity: An Introduction	1
Abstract.....	1
Creative Commons Copyright and Use Notice	2
Table of Contents	3
Security System Design for Productivity: An Introduction	4
Introduction.....	4
Productivity Defined.....	4
Cognitive Load.....	6
Preparing the Data.....	7
Designing the System.....	7
Next Steps.....	9
Conclusions.....	10
References.....	11

Security System Design for Productivity: An Introduction

Introduction

Centuries ago, it was said that all roads lead to Rome. Today it is more appropriate to say all paths cross the Internet. Improving the productivity of the security professionals improves the organization's security itself and, thus, improves functioning world-wide.

Why is that? It has to do with the nature of securing a computer system and the modus operandi of computer criminals. Physical world security tends to favor the defender. InfoSec, by contrast, favors the criminal. The protection mechanisms are not as intuitive as brick and mortar. One cannot easily see the holes and gaps in cyber defense. Further, the loss from a cyber crime is also not intuitive as spotting a broken down door or missing valuables. Cyber crimes often result in losses that are invisible to the lay person, and the criminals are able to come and go without leaving a trace. In the physical world, we used to talk about criminals committing the perfect crime. In the InfoSec realm, the talk is about defenders who must commit the perfect defense.

Productivity Defined

To improve upon the situation, we need to improve upon the productivity of the InfoSec professional. Yet how shall we measure productivity? An ideal security metric for its flexibility, simplicity, and broad applicability is that Protection must be greater than Detection plus Response (Schwartz, 1999). This can be stated as $P > D + R$. Effective security is created by extending the time protection mechanisms take to be breached, or reducing the time it takes for an InfoSec professional to detect and respond to an attempted breach. The manual intervention is primarily in detection and response. The measure of productivity, then, is in how quickly the InfoSec pro completes these tasks.

Effective security begins with protection. Protection mechanisms include those built-in by the hardware and software manufacturers, and those designed and configured by the InfoSec pro. The goal of protection mechanisms is to increase the time and effort necessary to gain access to sensitive computer resources. The longer it takes, the more likely it is that the attacker may give up or the InfoSec pro may thwart the attack. Likewise, the more effort it takes, the more likely it is the attacker may make a mistake or lack the skill set to compromise the system. Protection is a measure of time that a system can withstand attack.

There are non-tangible protection mechanisms such as the actions taken by the InfoSec pro in preparation for an attack. These include establishing policies and providing user awareness training, incident response planning, response kits (also called jump bags), and so on. These actions blur the line somewhat between protection, detection, and response. Preparation improves the fundamentals and cuts detection and response time. Yet protection is about setting the pieces on the board before the game commences, and hence preparation belongs in this phase.

Because all protection mechanisms have a finite lifespan, attackers will succeed if given enough time. A report on by Verizon (2008) on over 500 incidents paints a chilling picture. A typical attack took hours (36%) to days (28%) to be successful. A typical detection took months (63%). In fact, a majority of the attacks in the survey (82%) were never detected at all. The failure was not on the part of the data, it was on the part of the detection. "Evidence of events leading up to 82 percent of data breaches was available to the organization prior to actual compromise."

Detection is a must. The impact of an attack is exponential to the time the attacker has on the system. The longer the time window between compromise and detection, the higher exposure and thus the higher the cost. Richard Bejtlich (2008) has coined this relationship *intrusion debt*. "The longer the intrusion remains active, I would argue, the more debt one builds." The impact can thus be significantly reduced by speeding up the detection time. Detection mechanisms include sensors, data feeds, correlation and signature engines, and notification devices. The InfoSec pro, once notified, will then have to investigate further.

The result of detection is response. The first step is determining if it is an event or incident. Events are system anomalies while incidents are actual attacks. The response to events may be as simple as logging that they occurred, or as in-depth as performing system diagnostics and repair. Incident response, which is more formal as the results may involve law enforcement, includes containment, eradication, and recovery. Containment aims to stop the spread of the attack. In severe cases, it may include completely removing the affected system(s) from the network. Once contained, all compromised code or malware must be eradicated and cleaned off. This may include a complete rebuild of the computer system(s) involved. Once cleaned or rebuilt, the data can then be restored and operations recovered.

Thus the lifecycle of an attack is architecting and configuring for security, preparation, identification, containment, eradication, and recovery. The lessons learned is then captured in a review session and fed back into the security architecture and configuration. The choke point is identification, specifically, detection time. This single factor has the greatest determination of success. All other activities depend upon detection. Yet detection can be very difficult.

Attackers have many places to hide their activities or obfuscate their attacks as normal network traffic. Moreover, many normal network activities can appear like attacks. When an intrusion detection systems that miss attacks in the day-to-day noise, a false negative occurs. When an intrusion detection system incorrectly identifies normal traffic as an attack, a false positive occurs. False negatives increase the detection time and thereby directly impact productivity. False positives, on the other hand, waste time in unwarranted response and can train personnel to ignore alerts. In fact, it can be shown that the false positive rate is the limiting factor for intrusion detection performance (Axelsson, 2000). The ideal is minimizing both false positives and false negatives.

The performance metric for security is $P > D + R$. Protection (P) is the time the system can withstand an attack. Detection (D) is the time it takes for the system and the personnel to recognize such an attack. Once detected, time is spent determining if this is a false positive (event) or an actual intrusion (incident). Response (R) is the time it takes to contain and eradicate the attack. This paper will primarily focus on techniques for improving detection, with a secondary view of how the techniques support response.

Cognitive Load

The first set of techniques we will explore aim to optimize the InfoSec professional's mental resources. Detection requires attention to detail and patience, which in turn requires a calm mind and a deliberative manner. Many days may go past without an attack occurring. Some attacks may turn out to be false positives. Thus the mental condition must be maintained in the face of potential boredom or loss of focus. Security incidents, by contrast, are rare times of high adrenaline and high stress. The mental condition here must be maintained so costly mistakes can be avoided. Depending upon the skill of the attacker, the InfoSec pro will need a fair degree of skill and luck to protect the system.

As Seneca of Rome put it, "luck is what happens when preparation meets opportunity." An InfoSec pro prepares for detection by studying the current literature available on threats and attack methodologies. The professional prepares, too, by using detection tools during penetration tests and vulnerabilities assessments. These tests and assessments mimic the activities of actual attackers. There should be step-by-step procedures for classifying events versus incidents and for activating the appropriate response. Such training and preparation places the InfoSec pro in a ready state.

Once ready, the intrusion detection system must interrupt professional when an alert occurs. On its face, this is simple enough. In reality, it is anything but. Optimizing mental resources means protecting attention and awareness. Any interruption that fails to do this becomes a disruption. Such disruptions extend the time on the task and create what is known as a resumption lag. This lag is the few minutes of disorientation a person experiences after being interrupted. The time of such disruptions adds to the detection time, thus decreasing overall security effectiveness.

There are many ways to reduce the interruption overhead. One method is to interrupt on task boundaries (Adamczyk, 2004). The InfoSec professional is then interrupted only after completing the task at hand. The mental workload can be used to determine when the person is mentally available (Iqbal, 2005). Such methods allow for gradations wherein some alerts wait for task boundaries and others do not. For example, urgent alerts on key systems with a high confidence level could be configured to interrupt regardless of where the person was in a given task sequence. Another alternative is to pool interruptions into related categories. The interrupt occurs less often and carries more information. This is ideal for situations wherein several alarms may trigger based upon one underlying cause. Once the InfoSec pro is working on the cause, subsequent events are queued so as to not impact the professional's workflow.

The workflow must be maintained for the professional to give it their full, undivided attention. Multitasking is as costly as interruptions because it is, in fact, self-induced interruptions. While the impact on the InfoSec professional has not been studied, the impact has been studied on information workers in general. Basex, an IT research firm, looked at a thousand such workers and found they lost 2.1 hours per day due to multi-tasking (Wallis, 2006). Where does the time go? Other researchers have found that the culprit is resumption lag which takes, on average, 25 minutes to get back onto task (Thompson, 2005). Intrusion detection systems should focus on the task at hand, encouraging workflow and discouraging multitasking.

Click-thru or drill-down interfaces support this objective. Marty (2008) refers to it as the information seeking mantra: "Overview first, zoom and filter, then details on-demand." These present the summary information on a flat UI. The end-user can then click on an

item to view different aspects, or drill deeper to view the details. Such interfaces provide information gradation and allow users to step through a series of stages to provide ever more granular data. Gradation helps in identifying false positives because corresponding or conflicting information can be found. Gradation also helps later, during response, when digital forensics is required because both the summary and the source data are available.

Preparing the Data

Information gradation enables the user to jump from analysis to source data, from summary to specifics. The data itself comes from a broad range of courses; including applications, network services, operating systems, proxies, traffic flows, packet captures, firewalls, intrusion detection systems and passive network analysis (Marty, 2008). Sensors placed on the computers themselves (host-based) and on network connections (network-based). The raw data itself can be logged and stored: network packet captures, network transactions, application and operating system logs, system performance metrics. The statistical summary data is also logged and stored: packets per second, transactions per second, warning and error counts, performance over time. Analysis data, based on the raw data, is the third category of information that is created and stored: signature-based detection, anomaly-based analysis, and advanced filtering.

As an aside, let us look at the intersection of signature-based and anomaly detection. A signature-based model relies upon a pattern-matching engine that matches a known attack to the current system state. Detection speed is the primary advantage. The disadvantages are potential false positives due to normal traffic matching the signature and a higher rate of false negatives due to novel attack methods. By contrast, anomaly-based analysis treats every anomaly from normal traffic patterns as a potential attack. Anomaly-based analysis has a higher rate of false positives and a significantly lower rate of false negatives. Ideal intrusion detection systems combine both, which has been shown to more than double the detection rate with a corresponding 3% increase in false positives (Hwang, 2007).

Once the data is logged, data mining and analysis tools are used to convert it for output. Take, for example, a worm malware infection report. The statistics on traffic between infected hosts is extrapolated into a data set. Host-based information about system performance and data access is extrapolated into another data set. The first can be used to determine the extent and spread of the infection. The second provide a view into the information exposure. Both can then be returned to the InfoSec professional to be used in containment and eradication. One must remember, however, that data only becomes information at the boundary of man and machine.

Designing the System

The study of human-computer interaction systems (HCI) deals with this translation of data to information. The subset of HCI that deals with security is known as HCI-S. HCI-S focuses exclusively on how to improve security through interface design. The six main elements of HCI-S are: visibility of system status; aesthetic and minimalist design; end-user satisfaction; feature availability; learnability and intuitiveness; and trust (Muñoz, 2007).

A minimalist design returns to the idea of reducing interruptions and multitasking. It also facilitates satisfaction as well-designed interruptions have been shown to reduce the end-user's irritation and increase respect for the system (Iqbal, 2005). Feature availability and learnability speak to the way a system conforms to how the InfoSec pro works. It is mentally ergonomic, with all the knobs and buttons where the person expects them. Thus the HCI-S criteria, when properly applied, create an environment that makes the InfoSec professional more productive in finding and identifying data points.

The next step is to represent the data points in visual ways. This step has three goals. The first is to, again, convey information in a fashion that is non-intrusive and provides information gradation. The second goal is to make the system visually appealing; contributing to aesthetics and end-user satisfaction. The third and final goal is to reveal patterns in the information that might otherwise go unnoticed. For example, it can be shown that false positive detection improves when the source of an attack is shown on a three-dimensional globe (Marty, 2008). Since most actual attacks come from specific countries, such a display readily shows the InfoSec pro the country of origin and thus speeds up detecting false positives.

Visual cues can come from a notification system. Such a system serves to alert the InfoSec pro when specific actions occur. For example, a processor utilization threshold is exceeded on a specific host. Another example is when an attack signature is noted in the traffic flow. These specific notifications are then sent to the person for immediate response. Again, the delivery method must minimize the impact of the interruption.

One way to minimize this interruption is by slow growth displays. The threshold alert itself has a threshold. It must occur so many times in a given time period to alert the user. The initial alert is nominal, perhaps a small icon in the task tray or on a ticker window. Succeeding alerts are larger as the problem continues. Or the alerts roll-up into a colored icon on the screen. Many of today's email monitors, such as Google's Gmail icon, work this way. The icon changes color when messages are present. The end-user can then retrieve more details (e.g., the subject line of the emails in the inbox) by activating a command on the icon. Another example is Scope which coordinates several information notification feeds into one window (Van Dantzich, 2002).

Immediate and slow growth notification schemes are both problem-centric. They provide an alert only on exceptions; yet do not provide information on the ongoing state. Only by understanding the ongoing state can a professional anticipate problems and proactively respond before the state is such that the thresholds are exceeded. Thus state graphs are also needed for the HCI-S display. Peripheral displays, ones that show continuous information in the periphery of the user's environment, are ideal for real-time state graphs.

To put all of these pieces into perspective, let us take network traffic on an Ethernet port as an example. Threshold alerts will be triggered when errors exceed a certain number and when traffic exceeds a certain throughput. These alerts are displayed automatically using a combination of careful interrupting and slow growth notifications. Line graphs that display the ongoing throughput provides real-time state information. Information gradation is achieved when the InfoSec pro can click into the stream, pull out specific network conversations, and view these conversations in a protocol analyzer such as Wireshark.

Another method for displaying information is to use audio in place of or in conjunction with visual communication (Muñoz, 2007). Conveying data over nonverbal audio systems is known as sonification. Several sonification systems for network and intrusion

monitoring have been studied, including Peep (Gilfix, 2000), NeMoS (Malandrino, 2003), and JListen (Gopinath, 2004).

It is important to stress that sonification is merely representation of data over audio channels. The exact same data analysis is performed for both visualization and sonification. The information that is conveyed differs, of course, as any representation hides some details while illuminating others. Sonification has been shown to have a positive effect in identifying false positives (Malandrino, 2003).

Immediate notification in sonification revolves around the use of sound iconography. For example, NeMoS rang a tubular bell when Snort logged a denial of service attack. Slow growth methods also exist by increasing one or more of the following: alert volume; alert tone; alert frequency. It should be noted, however, that using tone or, similarly, using different musical notes, can fail as it requires the InfoSec professional to have a fair degree of musical skill and understanding.

Ongoing state can be conveyed using an ecology of sounds to represent a variety of data feeds. Sometimes, as with Peep, the sounds are literally from ecology such as woodlands or jungle sounds. Much like crickets telling the temperature by their chirps, synthetic crickets can convey the measure of activity thru digital chirps. Alternatively, music can be synthesized around the metrics and data points provided by the monitoring. Both natural and musical sound-scapes allow the InfoSec pro to build an intuition around the network performance. The InfoSec pro's mind then becomes the signature and anomaly detection engine, quickly alerting the consciousness when something does not sound quite right.

Next Steps

The system should also be flexible and customizable. An InfoSec pro may, for example, schedule an audio alert to play at a certain time. Or the pro may change priority on certain items, adding and removing as he sees fit. Scope, for example, allowed people to modify the position of items and groupings whenever desired (Van Dantzich, 2002). A memo notepad capability is also important to facilitate notes thru out the security, preparation, identification, containment, eradication, and recovery cycle.

Enabling forensics is also of importance. While this paper focuses primarily on improving detection, we should keep in mind that some detected results will be actual incidents of attack, and some of these will require involving law enforcement and potentially legal proceedings. Enabling forensics is done in two ways: by providing all of the information within the system, not just the summary or statistics but also the original source data, and by enabling the data to be extracted in line with the chain of custody.

The system should also foster communication by means of visualization and sonification. Such communication, of course, occurs between the system and the InfoSec professional. In the case of forensics, this is communication between the InfoSec professional and law enforcement. During normal operations, this is communication between professional, coworkers, and management.

Conclusions

All paths cross the Internet. Security of Internet applications and Internet-connected systems is of utmost importance in today's business world. Given security can be defined and measured using the $P > D + R$ formula, effort should be spent in considering improving Protection (P), Detection (D), and Response (R). Much has already been written on protection mechanisms. Thus this paper sought to review Detection methods and how intrusion detection systems can be designed to increase the InfoSec professional's effectiveness and support Response activities. These methods include data gathering, visualization, and sonification techniques, and presenting in methods that minimize interruptions and multitasking.

References

Adamczyk, P. B. (2004). If Not Now, When? The Effects of Interruption at Different Moments Within Task Execution. *CHI-2004*, (pp. 271-278). Vienna, Austria.

Axelsson, S. (2000). The base-rate fallacy and the difficulty of intrusion detection. *Transactions on Information Systems Security* , 186-205.

Gilfix, M. C. (2000). Peep (the network auralizer): Monitoring your network with sound. *LISA 2000 14th Systems Administration Conference*, (pp. 109-117).

Gopinath, M. (2004). Auralization of intrusion detection system using Jlisten. *Thesis, Birla Institute of Technology and Science, India* .

Hwang, K. (2007). Hybrid intrusion detection with weighted signature generation over anomalous Internet episodes. *IEEE Transactions on Dependable and Secure Computing, Vol. 4 Issue 1* , 41-55.

Iqbal, S. B. (2005). Investigating the Effectiveness of Mental Workload as a Predictor of Opportune Moments for Interruption. *CHI-2005*, (pp. 1489-1492). Vienna, Austria.

Malandrino, D. M. (2003). Network Monitoring with Sound. *Proceedings of the 2003 International Conference on Auditory Display*, (pp. 251-253). Boston, MA, USA.

Marty, R. (2008). *Applied Security Visualization*. Addison-Wesley Professional.

Muñoz, J. M. (2007). Integration of Auditive and Visual Feedback in the Design of Interfaces for Security Applications. *CLIHC 2007*. Rio de Janeiro, Brazil.

Schwartau, W. (1999). *Time Based Security*. Interpact Press.

Thompson, C. (2005, October 16). Meet the Life Hackers. *New York Times* .

Van Dantzich, M. R. (2002). Scope: Providing awareness of multiple notifications at a glance. *In Proceedings of the 6th International Working Conference on Advanced Visual Interfaces*.

Wallis, C. S. (2006, January 16). Help! I've Lost My Focus. *Time* .