



SANS Institute

Information Security Reading Room

We're Lost, But We're Making Good Time!

Benjamin Grubin

Copyright SANS Institute 2021. Author Retains Full Rights.

This paper is from the SANS Institute Reading Room site. Reposting is not permitted without express written permission.

We're Lost, But We're Making Good Time!

The Failure of Vulnerability Research

Benjamin P. Grubin, CISSP
18 September 2001
Practical for GSEC 1.2f

Preface

Throughout history, information has been protected by one form of security or another. Time and time again, we have stood by to watch the latest techniques compromised by a few dedicated, intelligent, and resourceful people purely as an intellectual pursuit. Whether referred to as black hat or white, the fact remains that scores of people on both sides of the electronic battlefield spend countless hours determining ways to undermine the security of fragile technological systems.

While a small segment of the information security professionals spend their time attempting to design secure architectures for their businesses, the more significant majority of participants in the infosec world tirelessly find new and previously unknown methods to compromise even the toughest and most leading edge security solutions. We as security professionals are facing a problem far worse than poor software or ignorant system administrators. We are facing ourselves.

In early 1994 Internet Security Systems took the corporate security world by storm with their Internet Security Scanner product. Along with several competitors such as Wheelgroup (now Cisco: <http://www.cisco.com>) and Secure Networks, Inc. (now NAI: <http://www.nai.com>), they began providing tools to the business community that allowed them to “institutionalize” their information security. In the form of a vulnerability scanner, businesses were provided a means to detect, catalogue, and respond to a growing number of threats to their enterprise in the form of software bugs. While several primitive methods existed previously to perform “scans” of network systems and report suspected trouble spots, this time businesses began to truly *believe* in vulnerability scanning as a best practice.

Why shouldn't they? Indeed, the vacuum hole that existed before with regards to security best practices was ripe for a product to fill it. For the first time, companies were offered not just a product but a new service, patterned on the successful anti-virus products: vendors would constantly research new security issues and incorporate the detection into their products. To businesses, it seemed folly not to participate in such a program, and protect themselves with the best product possible. The best product, of course, was the one that could identify the most problems. This attitude was pervasive, and influenced the development of modern intrusion detection technology along the same lines. If a product could scan for vulnerabilities, it could be made to detect attacks on them in progress. Now there are two products to sell!

Indeed, many other security technologies have been developed in many different areas, but the impact vulnerability scanning and intrusion detection technologies have made on the information security profession is inarguably huge. Now, because these products make possible a form of metrics by which to judge one's security posture, a company can feel a measure of safety. My intention with this paper is to show this safety as a farcical creation of the information security industry.

The Problem

If you look at how and why these early products were developed, you will see many parallels with the anti-virus business. Many widely-known and easily exploitable security issues existed, and the Internet community seemed unable or unwilling to address them. For many, the problem was simply knowing that the vulnerability existed. This is the problem that vulnerability scanners solved. They codified these vulnerabilities and allowed an unsophisticated user to detect whether he or she was at risk. The early tools that accomplished this, tools such as SATAN[1] and ISS[2], were widely distributed. Companies saw these products and began asking for enhancements to detect certain problems in their own networks. But the real problem truly arose when vulnerability scanners became a commercial enterprise.

As a result of the commercialization of vulnerability scanners, security companies that sold them needed a way to distinguish themselves from the competition. The only way they saw to accomplish this was to have more "checks" than the other guy. Unfortunately, they were working in a (at the time) finite environment. Only so many known security problems existed for them to include in their products. When a new security problem was found, the vendors would scramble to include a check for it in the latest product. Of course the real competition here was who could get the check out *faster*, since once a vulnerability was released, it was impossible for only one of the companies to keep it secret for long. They had to find a new method to distinguish themselves from the competition. What better way than to find the people who were releasing the "exploits", and employ them to make their work proprietary.

Of course this led to an interesting side effect. Suddenly, what was once an intellectual pursuit was now an economic one. People were being *paid* to research and publish new security vulnerabilities, and of course, produce a method to check for them in their respective company's product. Now the game wasn't only who had the most checks, but who had the best vulnerability research team. Full disclosure was the name of this game, and it became the new measure of security companies in general: how much exploitable information could they release to the public?

Previous to this ominous step, the information security industry was a passive, or reactive participant in the environment. Now they were active, and fiercely competitive. This competitive drive led to heavy recruitment of security hobbyists who were then encouraged to find the most disastrous, piercing vulnerabilities in the hottest commercial products.

According to the MITRE CVE[3], out of over 3000 documented vulnerabilities, over 70% of them were initially researched and discovered by an information security product vendor. It was here that the paths of the anti-virus community and the information security community diverged. Imagine if the anti-virus product vendors began writing and distributing virii, forcing companies to constantly buy new versions of the software to keep up. That's not competition, that's extortion—but that same type of activity in the information security industry is called research.

The *Real* Problem

The crux of the issue lies not in technology, but instead how security is perceived. The reality of the security business is that information security is a human problem, not a technological one. Unfortunately, the human problem is also one much more difficult to address. Only now are primitive technologies entering the market that can use heuristics and expert systems to detect *inappropriate* activity. This type of activity is not triggered by an exploit, but instead by a computer being able to recognize that a *person* is doing something that is inconsistent with what he is expected to be doing.

The real problem now lies in the business community, which has been duped by so-called security professionals into thinking that a constant war against buggy software is not only rational, but effective. In reality, it is neither. It is simply not possible to write software or design hardware that is free of problems. If a human problem exists, which according to the Computer Crime and Security Survey[4] it does in great numbers, you must attack that problem.

While the CSI survey has reported in recent years that external attacks have exceeded internal attacks, I believe these results are inappropriately weighted. Businesses believe the statistic reported to them by devices such as network and host intrusion detection systems, designed to catch the types of vulnerabilities most common on the internet. In reality, internal attacks rarely take the form of a software bug, but instead appear as normal activity to an intrusion detection device. It is simply someone abusing their privileges to do something they should not. This type of attack is the most insidious and without a doubt the most damaging, and the most difficult to detect, leading to the skewed “external attack” figure.

As a result, companies lean on the existing security technologies to protect themselves. Constantly in a cycle of responding to the latest security threat, businesses wage war against the for-profit vulnerability research teams they are in fact paying to protect them. They are unable to focus on solving the human problem, because the magnitude of the vulnerability problem is so large and utterly unsolvable—made that way by the effect of security companies *encouraging* the discovery of new security flaws even outside themselves by the promise of employment and status in the industry to those that can find them.

The Human Problem

The war against software bugs that compromise security will always be a nagging issue, but to rely on it as a means to combat the information security problem as a whole is ineffective at best. Instead, as security professionals, we must focus on how to solve the *human* problem technologically.

People are difficult creatures to predict. Utilizing misuse and anomaly detection[5], we are just becoming able to attack the root of the problem—identifying inappropriate behavior and either terminating or reporting it. Credit card and telephone companies have invested millions into this type of technology[6] to determine potential cases of abuse. While they also attempt to provide a level of technical security around the use of devices such as credit cards and telephones, the real damage prevention occurs in their proprietary fraud abuse systems. These systems utilize previous spending habits of legitimate consumers to constantly compare activity to an established norm. We in the information technology field are just beginning to have access to commercial technologies containing expert systems to accomplish this same task on an enterprise IT level.

In a traditional physical security setting, human security guards combined with electronic alerting systems are the most widely implemented. These security systems combined with access control, authentication, and authorization systems such as ID cards and electronic, auditable locking systems, allow a human to be alerted if a person is doing something other than they should be. The difference between this world and the world of network and information security is that there are no set of physical rules that are standard. In the physical world, we are ruled by the laws of physics. One cannot go through a door without opening it. The rules of the physical world are much easier to describe to something like a security system. Anything you can't describe with those types of rules must be identified by a human being. This too is easier in the physical world, since a human can interact with the world and be a part of it. The same cannot be said for the miles of cables carrying electronic data in a typical enterprise. These cables are opaque, you can't stick your head in and look around. So we must design tools to do that for us, and out of the vast amounts of data, cull just that which might be considered interesting. This is where the difficulty lies.

Inside that wire, in the world of data, there are no physical rules. The rules are determined by each application passing data over that wire, each written by a different person (or persons), and none conforming to a universal set of laws such as nature. So the rules we use to describe how things *should* work inside that tangle of cables cannot be static. Instead, the rules must be adaptive and intelligent, learning how things normally operate, and ensuring that anything outside those boundaries can be examined by a person in depth. Only once the rules have been explained properly (“Joe can access application Y from 9 to 5.”) the same principles as physical security systems can begin to apply. Joe can access his application during the day without being flagged, but if he does anything *else* it will trigger an alarm. It is this “else” condition where traditional security

systems based around vulnerabilities fail us. It is not economical to constantly update the else condition, and to attempt to do so invites the type of trouble we are now experiencing due to the corporate vulnerability research phenomenon.

Products such as Silent Runner[7] provide a means to distribute “security guards” across your network. These security guards can not only utilize existing deployed security infrastructure such as intrusion detection, but use anomaly detection to analyze patterns and distinguish inappropriate behavior. Instead of “Code Red worm detected on PC”, the type of information generated by intrusion detection, you can instead get information such as “The payroll PC’s are trying to talk to the Internet more than normal”. This kind of information is not only far more useful in detecting real trouble and reducing false positives, but is more generalized. It is not tied to a specific bug, worm, or exploit, but instead relies on patterns to see that something is wrong. Nobody has to constantly update the software with new issues, because anything unusual (such as someone trying to run an exploit) is much easier to see as an anomaly than through a group of vulnerability signatures. The value of this approach is clear.

Products like Silent Runner combat the human problem, using technological means that are currently in a fledgling state of development, but I feel it is important to highlight it as an example of what truly useful technologies are becoming available in this constant war.

Eliminating Commercial Vulnerability Research

I believe the difficult but necessary next step to refocusing our efforts on the true problem is to eliminate sponsored vulnerability research on third-party products. The true responsibility for finding and correcting software problems lies with the vendors and users of products. It cannot be the responsibility of the security product vendor to advise their clients to switch to a newer revision of another vendor’s software. This has a multitude of problems such as compatibility issues, lack of vendor support, and cost of upgrading. Instead, the vendor must work with the customer to resolve security issues, and then announce and upgrade customers as necessary.

The one lost function of a vulnerability research group is the ability to keep tabs on what new vulnerabilities and exploits have been developed by freelancers (or hackers). This is much better accomplished by a team dedicated to doing just that. Such teams exist in companies like Security Focus (<http://www.securityfocus.org>), who provide a database to subscribers listing current known vulnerabilities in products. These databases could in turn be used by security product vendors as a baseline for scanning signatures, when applicable. Then security product vendors would be free to compete at a more meaningful level: against the usability, stability, and overall quality of the competitors’ products. The inherent conflict of interest is removed.

Conclusion

This is a complex problem to solve, but the appropriate research is underway, and many products are emerging on the market. The important point I have tried to communicate is that the existing metrics preached by consultants and vendors alike is misleading and unreliable. We must focus our attention on developing and implementing viable solutions to the human security problem, and not on technological gimmicks that offer empty promises.

Technical stopgaps, especially those based on vulnerability research, can never be effective in preventing real security issues. Software will always have bugs. Instead, enterprise IT architectures and appropriate access control, authentication, and authorization procedures must be improved. This can only be accomplished by having a skilled security professional involved in every stage of planning, implementing, and maintaining an IT infrastructure. These measures are the only way to understand what *appropriate use* really means in the context of a business—and the only way to understand how to protect the business from real threats. In time, this understanding of appropriate use will be learned and understood by capable systems.

While waiting for these future security systems to become mainstream, the best preparation is understanding, organizationally, what appropriate use means and defining the rules and policies around it. Only then will you be able to implement these rules in the security systems of the future. Focusing on this now instead of vulnerability detection and correction, is the only way to ensure you will be ahead of the curve instead of behind when new technology is introduced. In the long run, the cost of this is nothing compared to the band-aid solutions and the incredible cost of recovering from a security incident.

Sources

[1] SATAN, System Administrator's Tool for Analyzing Networks

URL: <http://www.cs.ruu.nl/cert-uu/satan.html>

[2] Internet Security Scanner, Internet Security Systems

URL: <http://www.iss.net>

[3] Common Vulnerabilities and Exposures (CVE) Database, MITRE Corporation

URL: <http://cve.mitre.org/cve/>

[4] Computer Crime and Security Survey, Computer Security Institute / Federal Bureau of Investigation

URL: <http://www.gocsi.com/prelea/000321.html>

[5] Sundaram, Aurobindo "An Introduction to Intrusion Detection" (23 Jan 01)

URL: <http://www.acm.org/crossroads/xrds2-4/intrus.html>

[6] Hodgson, Jeffrey "Telephone Fraud: What to Look For in a Fraud Detection System"

URL: <http://www.beckcomputers.com/botfiles/botlookfor.html>

[7] Silent Runner, Raytheon

URL: <http://www.silentranner.com>

References

Forristal, Jeff, Shipley, Greg “Vulnerability Assessment Scanners”
Network Computing, January 2001

URL: <http://www.networkcomputing.com/1201/1201f1b1.html>

eEye Digital Security, Vulnerability Research

URL: <http://www.eeye.com/html/Research/Advisories/index.html>

Internet Security Systems, X-Force

URL: <http://xforce.iss.net/>

W. Lee, S. J. Stolfo, and P. K. Chan. “Learning patterns from Unix process execution traces for intrusion detection.”

AI Approaches to Fraud Detection and Risk Management AAAI Press (July 1997)

Avolio, Frederick “A Multi-Dimensional Approach to Internet Security” (May 1998)

URL: <http://www.avolio.com/MultiDimensional.html>

Ghosh, Anup, Schwartzbard, Aaron & Schatz, Michael “Learning Program Behavior Profiles for Intrusion Detection”

URL:

http://www.usenix.org/publications/library/proceedings/detection99/full_papers/ghosh/ghosh.html/

© SANS Institute 2001, Author retains full rights