

The slide features a blue vertical bar on the left. The top section has a light gray background with the author's name and title on the left, and the date on the right. Below this is a dark blue banner with a background image of a person in a server room, the Counterpane logo, and the slogan 'INTELLIGENT ALERT. INSTANT RESPONSE. IMMEDIATE DEFENSE.' The main title is centered in a large, dark blue font. At the bottom, there is a white box with the online availability URL.

Bruce Schneier
CTO, Counterpane Internet Security

June 2002

INTELLIGENT ALERT. INSTANT RESPONSE. IMMEDIATE DEFENSE.

Counterpane[™]
Internet Security

Fixing Network Security by Hacking the Business Climate

This presentation is available online at:
<http://www.counterpane.com/presentation4.pdf>

Talk Description:

Network security has long been considered an engineering problem, and companies try to solve it by applying technologies. This approach is failing; the technologies are failing and the problem is worsening. What we need are security processes, such as detection, response, and deterrence. However, the only way to get corporate management to adequately address security is to change the risk-management equation. This can be achieved by enforcing liabilities, and giving corporate management the means to reduce or insure against those liabilities. It's only after we do all of these things will the Internet be a safe and secure place.

This presentation is available on-line at <http://www.counterpane.com/presentation4.pdf>

About the Author:

Internationally renowned security technologist and author Bruce Schneier is the Founder and the Chief Technical Officer of Counterpane Internet Security, Inc., the world leader in Managed Security Monitoring. Counterpane provides security monitoring services to Fortune 2000 companies world-wide. He is the author of six books on security and cryptography, including the security best seller, "Secrets & Lies: Digital Security in a Networked World." His first book, "Applied Cryptography," has sold over 150,000 copies world-wide, and is the definitive work in the field. Schneier designed the Blowfish and Twofish encryption algorithms, and writes the influential "Crypto-Gram" monthly newsletter. He is a frequent lecturer on computer security and cryptography.


Bruce Schneier's biography is available on-line at <http://www.counterpane.com/schneier.html>

About Counterpane Internet Security, Inc.

Counterpane Internet Security, Inc. is the innovator and acknowledged leader in providing Managed Security Monitoring (MSM) services. MSM combines people and technology to safeguard businesses. Working from a network of technically sophisticated Secure Operations Centers (SOCs) and using progressive analysis tools, Counterpane has built the most advanced analysis, correlation, detection, and diagnosis technology, comprising of a Sentry monitoring probe on the customer's network and the Socrates knowledge base inside the SOCs. Using this technology, Counterpane's expert Security Analysts are able to detect security incidents-both external intrusions and insider attacks-in real time, and tailor immediate, effective responses for its customers. It has partnered with leading security companies, consulting organizations, and VARs to deliver MSM services world-wide. Counterpane is headquartered in Sunnyvale, CA, and has two operational SOCs: one in Mountain View, CA, and the other in Chantilly, VA.

More information about Counterpane is available on-line at <http://www.counterpane.com/>

Computer Security is Critical



- Security is one of the fundamental building blocks of the Internet
 - Everything we do on the Internet requires some security
 - The limits of security will become one of the limits of the Internet
- Attacks make it riskier to do business on-line
 - Attacks increase the cost of doing business
 - Attacks make companies hesitant to be on-line
- But...the rewards of being on the Internet far outweigh the risks

2

INTELLIGENT ALERT. INSTANT RESPONSE. IMMEDIATE DEFENSE.

The Importance of Security


When I began working in computer security, the only interest was from the military and a few scattered privacy advocates. The Internet has changed all that. The promise of the Internet is to be a mirror of society. Everything we do in the real world, we want to do on the Internet: conduct private conversations, keep personal papers, sign letters and contracts, speak anonymously, rely on the integrity of information, gamble, vote, publish digital documents. All of these things require security. Computer security is a fundamental enabling technology of the Internet; it's what transforms the Internet from an academic curiosity into a serious business tool. The limits of security are the limits of the Internet. And no business or person is without these security needs.

The risks are real. Everyone talks about the direct risks: theft of trade secrets, customer information, money. People also talk about the productivity losses due to computer security problems. What's the loss to a company if its e-mail goes down for two days? Or if ten people have to scramble to clean up after a particularly nasty intrusion? I've seen figures as high as \$10 billion quoted for worldwide losses due to the ILOVEYOU virus; most of that is due to these productivity losses.

More important are the indirect risks: loss of customers, damage to brand, loss of goodwill. Regardless of how the million-credit-card-number theft at Egghead.com turned out, some percentage of customers decided to shop elsewhere. When CD Universe suffered a credit card theft in early 2000, it cost them dearly in their war for market share against Amazon.com and CDNow. In the aftermath of the Microsoft attack in October 2000, the company spent much more money and effort containing the public relations problem than fixing the security problem. The public perception that their source code was untainted was much more important than any effects of the actual attack.

And more indirect risks are coming. European countries have strict privacy laws; companies can be held liable if they do not take steps to protect the privacy of their customers. The U.S. has similar laws in particular industries-banking and healthcare-and there are bills in Congress to protect privacy more generally. We have not yet seen shareholder lawsuits against companies that failed to adequately secure their networks and suffered the consequences, but they're coming. Can company officers be held personally liable if they fail to provide for network security? The courts will be deciding this question in the next few years.

As risky as the Internet is, companies have no choice but to be there. The lures of new markets, new customers, new revenue sources, and new business models are just so great that companies will flock to the Internet regardless of the risks. There is no alternative. This, more than anything else, is why computer security is so important.


Counterpane™
Internet Security

Computer Security is Failing

- By every measure, the Internet is becoming less secure all the time
 - More attacks, more damage, more losses
- Security products are failing
 - If firewalls are so great, why aren't they keeping attackers out?
 - If anti-virus products work so well, why do viruses wreak havoc so often?
- More users means more systems means more problems
- Despite advances in security research and development, the problems get worse every year

3

INTELLIGENT ALERT. INSTANT RESPONSE. IMMEDIATE DEFENSE.

CSI's Computer Crime and Security Survey

For the past six years, the Computer Security Institute has conducted an annual computer crime survey. In 2001, 64% of respondents reported "unauthorized use of computer systems" in the last year. 25% said that they had no such unauthorized uses, and 11% said that they didn't know. The number of incidents was all over the map, and the number of insider versus outsider incidents was roughly equal. 70% of respondents reported their Internet connection as a frequent point of attack (this has been steadily rising over the six years), 18% reported remote dial-in as a frequent point of attack (this has been declining), and 31% reported internal systems as a frequent point of attack (also declining).

The types of attack range from telecommunications fraud to laptop theft to sabotage. 40% experienced a system penetration, 36% a denial-of-service attack. 26% reported theft of proprietary information, and 12% financial fraud. 18% reported sabotage. 23% had their Web sites hacked (another 27% didn't know), and over half of those had their Web sites hacked ten or more times (90% of the Web site hacks resulted in vandalism, and 13% included theft of transaction information).

What's interesting is that all of these attacks occurred despite the wide deployment of security technologies: 95% have firewalls, 61% an IDS, 90% access control of some sort, 42% digital IDs, etc.

The financial consequences are staggering. Only 196 respondents would quantify their losses, and those totaled \$378 million. From under 200 companies! In one year! This is a big deal.

To get a copy of this survey, visit http://www.gocsi.com/prelea_000321.htm


The Honeynet Project

The Honeynet Project measures actual computer attacks on the Internet. According to their most recent results, a random computer on the Internet is scanned dozens of times a day. The life expectancy of a default installation of Red Hat 6.2 server, or the time before someone successfully hacks it, is less than 72 hours. A common home user setup, with Windows 98 and file sharing enabled, was hacked five times in four days. Systems are subjected to NetBIOS scans an average of 17 times a day. And the fastest time for a server being hacked: 15 minutes after plugging it into the network.

My essay on the Honeynet Project: <http://www.counterpane.com/crypto-gram-0106.html#1>

The Honeynet Project homepage: <http://project.honey.net.org/>

Old Problems Don't Go Away



- We still don't know how to reliably write vulnerability-free code
 - Buffer overflows have been a security problem for decades
- Internet attacks are as old as the Internet
 - Hackers have been bypassing firewalls since they were invented
- Authentication problems have nothing to do with the technology
- Nothing has been invented that effectively deals with insider abuse

INTELLIGENT ALERT. INSTANT RESPONSE. IMMEDIATE DEFENSE.4

How to Think About Security

If security has a silly season, we're in it. After September 11, every two-bit peddler of security technology crawled out of the woodwork with new claims about how his product can make us all safe again. Every misguided and defeated government security initiative was dragged out of the closet, dusted off, and presented as the savior of our way of life. More and more, the general public is being asked to make security decisions, weigh security tradeoffs, and accept more intrusive security.

Unfortunately, the general public has no idea how to do this.

But we in computer security do. We've been doing it for years; we do it all the time. And I think we can teach everyone else to do it, too. What follows is my foolproof, five-step, security analysis. Use it to judge any security measure.

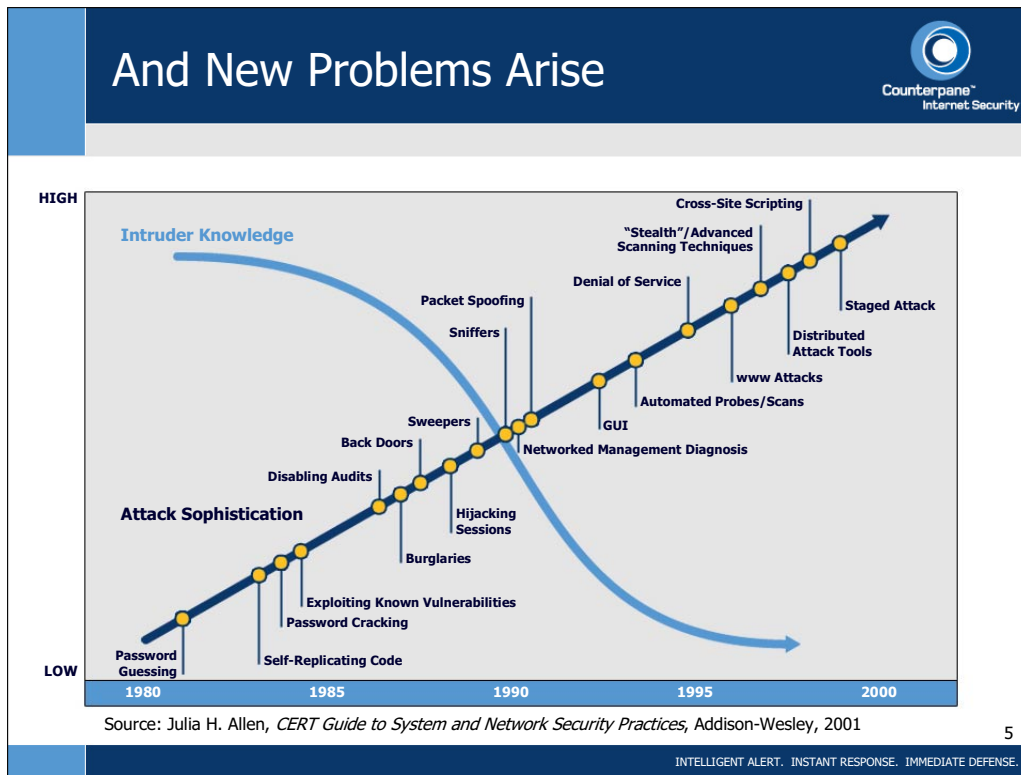
Step one: What problem does the security measure solve? You'd think this would be an easy one, but so many security initiatives are presented without any clear statement of the problem. National ID cards are a purported solution without any clear problem. Increased net surveillance has been presented as a vital security requirement, but without any explanation as to why. (I see the problem not as one of not having enough information, but of not being able to analyze and interpret the information already available.)

Step two: How well does the security measure solve the problem? Too often analyses jump from the problem statement to a theoretical solution, without any analysis as to how well current technology actually solves the problem. The companies that are pushing automatic face recognition software for airports and other public places spend all their time talking about the promises of a perfect system, while skipping the fact that existing systems work so poorly as to be useless. Enforcing a no-fly zone around a nuclear reactor only makes sense if you assume a hijacker will honor the zone, or if it is large enough to allow reaction to a hijacker who doesn't.

Step three: What other security problems does the measure cause? Security is a complex and inter-related system; change one thing and the effects ripple. If the government bans strong cryptography, or mandates back-doors, the resultant weaker systems will be easier for the bad guys to attack. National ID cards require a centralized infrastructure that is vulnerable to abuse. In fact, the rise of identity theft can be linked to the increased use of electronic identity. Make identities harder to steal through increased security measures, and that will only make the fewer stolen identities more valuable and easier to use.

Step four: What are the costs of the security measure? Costs are not just financial, they're social as well. We can improve security by banning commercial aircraft. We can make it harder for criminals to outrun police by mandating 40 mph speed maximums in automobiles. But these things cost society too much. A national ID card would be enormously expensive. The new rules allowing police to detain illegal aliens indefinitely without due process cost us dearly in liberty, as does much of the PATRIOT Act. We don't allow torture (officially, at least). Why not? Sometimes a security measure, even though it may be effective, is not worth the costs.

(Continued on next page.)



How to Think About Security (cont.)


Step five: Given the answers to steps two through four, is the security measure worth the costs? This is the easy step, but far too often no one bothers. It's not enough for a security measure to be effective. We don't have infinite resources. We don't have infinite patience. As a society, we need to do the things that make the most sense, that are the most effective use of our security dollar.

Some security measures pass these tests. Increasing security around dams, reservoirs, and other infrastructure points is a good idea. Not storing railcars full of hazardous chemicals in the middle of cities should have been mandated years ago. New building evacuation plans are smart, too. These are all good uses of our limited resources to improve security.

This five-step process works for any security measure, past, present, or future

- 1) What problem does it solve?
- 2) How well does it solve the problem?
- 3) What new problems does it add?
- 4) What are the economic and social costs?
- 5) Given the above, is it worth the costs?

When you start using it, you'd be surprised how ineffectual most security is these days. For example, only two of the airline security measures put in place since September 11 have any real value: reinforcing the cockpit door, and convincing passengers to fight back. Everything else falls somewhere between marginally improving security and a placebo.

Why?
Counterpane™
Internet Security

Complexity

- Complexity is the enemy of security
 - More complexity = less security
- The Internet is getting more complex faster than our ability to secure it
 - More features and increasing code size
 - Increasing extensibility
 - More interactions and dependencies
 - Increasing connectivity
- We are not willing to sacrifice features for security
- The monoculture of the Internet doesn't help either

6

INTELLIGENT ALERT. INSTANT RESPONSE. IMMEDIATE DEFENSE.

Complexity and Security

The future of digital systems is complexity, and complexity is the worst enemy of security.

Digital technology has been an unending series of innovations, unintended consequences, and surprises, and there's no reason to believe that will stop anytime soon. But one thing has held constant through it all, and it's that digital systems have gotten more complicated. The Internet is probably the most complex machine mankind has ever built, and it's not getting any simpler anytime soon.

As a consumer, I think this complexity is great. There are more choices, more options, more things I can do. As a security professional, I think it's terrifying. As cyberspace continues to get more complex, it will continue to get less secure. This is true for several reasons.

The first reason is the number of security bugs. All software contains bugs, and some of these bugs will affect security.

The second reason is the extensibility of complex systems. Complex systems are necessarily modular; code arrives when needed and does not all come pre-packaged and pre-installed. But extensibility means increased security flaws, because security often fails where modules interact. For example, last year's epidemic of macro viruses showed that Microsoft Word and Excel need to be secure. Java applets need to be secure for any possible use, not only for the uses they are intended. Malicious e-mail attachments can tunnel through firewalls. Convenience features in Microsoft Outlook can compromise security.


The third reason is the increased testing requirements for complex systems. The only reasonable way to test the security of a system is to perform security evaluations on it. However, the more complex the system is, the more security-related errors it will have. More possible interactions creates more work during the security evaluation. Checking every possible configuration is effectively impossible. Thus the difficulty of performing security evaluations also grows very rapidly with increasing complexity. The combination of additional (potential) weaknesses and a more difficult security analysis unavoidably results in insecure systems.

The fourth reason is that the more complex a system is, the harder it is to understand. There are all sorts of vulnerability points—human-computer interface, system interactions—that become much larger when you can't keep the entire system in your head.

The fifth reason is the difficulty of analysis. The more complex a system is, the harder it is to do analysis. Everything is more complicated: specification, design, implementation, and use. And everything is relevant to security analysis.

(Continued on next page.)

Technology Won't "Solve" Security



- Security is fundamentally a people problem, not a technology problem
 - Look around
- If you want to achieve real security, don't focus on the technologies
 - Look at the businesses
 - Look at the business motivations
 - Look at the business costs

7

INTELLIGENT ALERT. INSTANT RESPONSE. IMMEDIATE DEFENSE.

Complexity and Security (cont.)

A more complex system loses on all fronts. It contains more weaknesses to start with, its modularity exacerbates those weaknesses, it's harder to test, it's harder to understand, and it's harder to analyze. And it gets worse: the security of the overall system is limited by the security of its weakest link. Any single weakness can destroy the security of the entire system.

Systems are becoming more complex more quickly. Microsoft Windows is a poster child for this trend. Windows 3.1, released in 1992, had 3 million lines of code; Windows 95 has 15 million; Windows 98 has 18 million; and Windows 2000 has between 35 million and 60 million lines of code, depending on whom you believe. (In comparison, Linux, even with the addition of X Windows and Apache, is still under 5 million lines of code.)


The networks of the future, necessarily more complex, will be less secure. The technology industry is driven by demand for features, for options, for speed. There are no standards for quality or security, and there is no liability for insecure software. Hence, there is no economic incentive to create high quality. Instead, there is an economic incentive to create the lowest quality the market will bear. Unless customers demand higher quality and better security, this won't change.

The only reasonable response is to deal with this truism. Recognize that the digital world will be one of ever-expanding features and options, ever-faster product releases, ever-increasing complexity, and ever-decreasing security. Put processes in place to deal with this insecurity. Stop relying on the technology to save you.

I have long championed network monitoring and human intervention as a way to achieve security in complex systems. Monitoring provides constant feedback as to the efficacy of security, and human intervention provides a resilience that no automatic software product can match. The Internet isn't going to get any simpler in the near future; we need to modify our security practices to keep pace.

Reference: *Normal Accidents: Living with High-Risk Technologies*, by Charles Perrow, Basic Books, NY, 1999

Security: Two Models



**Threat Avoidance:
The Military Model**

- Security is an absolute
 - Figure out what the threats are, and avoid them
 - Either you're secure or you're not
- Follows a computer engineering mentality
 - Find and solve them
- Security becomes a barrier to business
 - Remember the crypto wars?

**Risk Management:
The Business Model**

- Security is relative:
 - Many risks and solutions
 - Things fail all the time
- Variety of options:
 - Accept the risk
 - Mitigate the risk with technology
 - Mitigate the risk with procedures
 - Transfer the risk

8

INTELLIGENT ALERT. INSTANT RESPONSE. IMMEDIATE DEFENSE.

Security and Risk Management

Ask any network administrator what he needs security for, and he can describe the threats: Web site defacements, corruption and loss of data due to network penetrations, denial-of-service attacks, viruses and Trojans. The list seems endless, and an endless series of news stories proves that the threats are real.

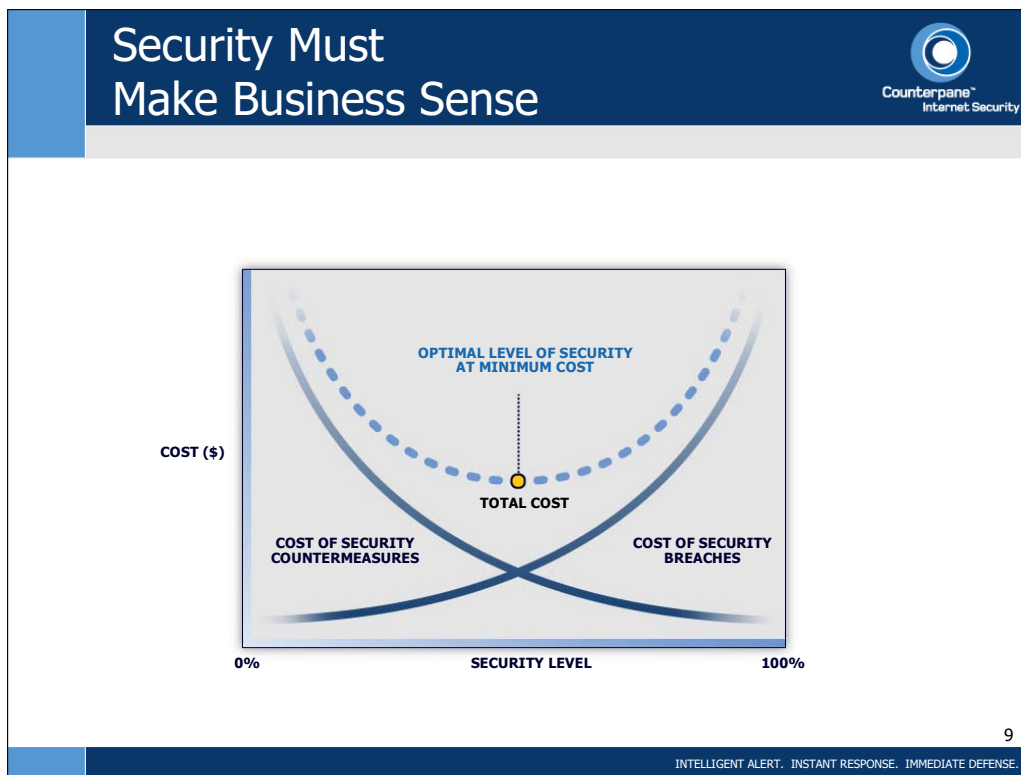
Ask that same network administrator how security technologies help, and he'll discuss avoiding the threats. This is the traditional paradigm of computer security, born out of a computer science mentality: figure out what the threats are, and build technologies to avoid them. The conceit is that technologies can somehow "solve" computer security, and the end result is a security program that becomes an expense and a barrier to business. How many times has a security officer said: "You can't do that; it would be insecure"?

This paradigm is wrong. Security is a people problem, not a technology problem. There is no computer security product-or even suite of products-that acts as magical security dust, imbuing a network with the property of "secure." It can't be done. And it's not the way business works.

Businesses manage risks. They manage all sorts of risks; network security is just another one. And there are many different ways to manage risks. The ones you choose in a particular situation depend on the details of that situation. And failures happen regularly; many businesses manage their risks improperly, pay for their mistakes, and then soldier on. Businesses are remarkably resilient.

To take a concrete example, consider a physical store and the risk of shoplifting. Most grocery stores accept the risk as a cost of doing business. Clothing stores might put tags on all their garments and sensors at the doorways; they mitigate the risk with a technology. A jewelry store might mitigate the risk through procedures: all merchandise stays locked up, customers are not allowed to handle anything unattended, etc. And that same jewelry store will carry theft insurance, another risk management tool.

More security isn't always better. You could improve the security of a bank by strip-searching everyone who walks through the front door. But if you did this, you would have no business. Studies show that most shoplifting at department stores occurs in dressing rooms. You could improve security by removing the dressing rooms, but the losses in sales would more than make up for the decrease in shoplifting. What all of these businesses are looking for is adequate security at a reasonable cost. This is what we need on the Internet as well-security that allows a company to offer new services, to expand into new markets, and to attract and retain new customers. And the particular computer security solutions they choose depend on who they are and what they are doing.



9

INTELLIGENT ALERT. INSTANT RESPONSE. IMMEDIATE DEFENSE.

The Business Case for Security Monitoring

Five years ago a firewall was all you needed for security on the Internet. Back then, no one had ever heard of denial-of-service attacks shutting down Web servers, let alone common gateway interface scripting flaws and the latest vulnerabilities in Microsoft Outlook Express. But in the wake of recent years came intrusion detection systems, public-key infrastructure, smart cards and biometrics. New networking services, wireless devices and the latest products regularly turn network security upside down. It's no wonder CIOs can't keep up.

What's amazing is that no one else can either. Computer security is a 40-year-old discipline; every year there's new research, new technologies, new products, even new laws. And every year things get worse.

It's not about the technology.

Network security is an arms race, where the attackers have all the advantages. First, potential intruders are in what military strategists call "the position of the interior": the defender has to defend against every possible attack, while the attacker only has to find one weakness. Second, the immense complexity of modern networks makes them impossible to properly secure. And third, skilled attackers can encapsulate their attacks in automatic programs, allowing people with no skill to use them.

The way forward is not more products, but better processes. We have to stop looking for the magic preventive technology that will avoid the threats, and embrace processes that will let us manage the risks. And that doesn't mean more prevention; it means detection and response.

On the Internet this translates to constant monitoring of your network. In October 2000, Microsoft discovered that an attacker penetrated its corporate network weeks earlier, doing untold damage. Administrators discovered this breach when they noticed 20 new accounts being created on a server. Then they went back through their audit records and pieced together how the attacker got in and what he did. If someone had been monitoring those audit records—from the firewalls, servers and routers—in real time, the attacker could have been detected and repelled at the point of entry.

Monitoring also means vigilance; attacks come from all over and at all hours. It means that experts need to continuously monitor with the tools and expertise at hand to figure out what is happening. Throwing an intrusion detection system onto a network and handing a system administrator a pager isn't monitoring, any more than giving a bucket to the guy at the other end of a fire alarm replaces a fire department.

Prevention systems are never perfect. No bank ever says: "Our safe is so good, we don't need an alarm system." No museum ever says: "Our door and window locks are so good, we don't need night watchmen." Detection and response are how we get security in the real world, and it's the only way we can possibly get security on the Internet. CIOs must invest in monitoring services if they are to maintain security in a networked world.

C
Counterpane™
Internet Security

Current Risk Management Thinking: Enterprise Customers

- Costs of Improving Security
 - Significant expense
 - Reduced functionality
 - Annoyed users
- Costs of Ignoring Security
 - Bad press
 - Angry customers
 - Regulatory pressure
- Result: Do what everyone else does, and no more

10
INTELLIGENT ALERT. INSTANT RESPONSE. IMMEDIATE DEFENSE.

Full Disclosure

Microsoft is leading the charge to restrict the free flow of computer security vulnerabilities. In November Scott Culp, manager of the security response center at Microsoft, published an essay describing the current practice of publishing security vulnerabilities to be "information anarchy." He claimed that we'd all be a lot safer if researchers would keep details about vulnerabilities to themselves, and stop arming hackers with offensive tools. Last week, at Microsoft's Trusted Computing Forum, Culp announced a new coalition to put these ideas into practice.

This is the classic "bug secrecy vs. full disclosure" debate. I've written about it previously in Crypto-Gram; others have written about it as well. It's a complicated issue with subtle implications all over computer security, and it's one worth discussing again.

The Window of Exposure

I coined a term called the "Window of Exposure" to explain the evolution of a security vulnerability over time. A vulnerability is a bug; it's a programming mistake made by a programmer during the product's development and not caught during testing. It's an opening that someone can abuse to break into the computer or do something normally prohibited.

Assume there's a vulnerability in a product and no one knows about it. There is little danger, because no one knows to exploit the vulnerability. This vulnerability can lie undiscovered for a short time -- Windows XP vulnerabilities were discovered before the product was released -- or for years. Eventually, someone discovers the vulnerability. Maybe it's a good guy who tells the developer. Maybe it's a bad guy who exploits the vulnerability to break into systems. Maybe it's a guy who tells no one, and then someone else discovers it a few months later. In any case, once someone knows about the vulnerability, the danger increases.

Eventually, news of the vulnerability spreads. Maybe it spreads amongst the security community. Maybe it spreads amongst the hacker underground. The danger increases as more people learn about the vulnerability. At some point, the vulnerability is announced. Maybe it's announced on Bugtraq or another vulnerability Web site. Maybe it's announced by the security researcher in a press release, or by CERT, or by the software developer. Maybe it's announced on a hacker bulletin board. But once it's announced, the danger increases even more because more people know about it.

Then, someone writes an exploit: an automatic tool that exercises the vulnerability. This is an inflection point, and one that doesn't have a real-world analog for two reasons. One, software has the ability to separate skill from ability. Once a tool is written, anyone can exploit the vulnerability, regardless of his skill or understanding. And two, this tool can be distributed widely for zero cost, thereby giving everybody who wants it the ability. This is where "script kiddies" come into play: people who use automatic attack tools to break into systems. Once a tool is written, the danger increases by orders of magnitude.

Then, the software developer issues a patch. The danger decreases, but not as much as we'd like to think. A great many computers

(Continued on next page.)

C
Counterpane™
Internet Security

Current Risk Management Thinking: Software Vendors

- Costs of Improving Security
 - Fewer features
 - Significant expense
 - Delayed product offerings
 - Annoyed users
- Costs of Ignoring Security
 - Bad press
 - Users switching to competitors
- Result: Talk big about security, but do as little as possible

11

INTELLIGENT ALERT. INSTANT RESPONSE. IMMEDIATE DEFENSE.

Full Disclosure (cont.)

on the Internet don't have their patches up to date; there are many examples of systems being broken into using vulnerabilities that should have been patched. I don't fault the sysadmins for this; there are just too many patches, and many of them are sloppily written and poorly tested. So while the danger decreases, it never gets back down to zero.

You can think of this as a graph of danger versus time, and the Window of Exposure as the area under the graph. The goal is to make this area as small as possible. In other words, we want there to be as little danger as possible over the life cycle of the software and the particular vulnerability. Proponents of bug secrecy and proponents of full disclosure simply have different ideas for achieving that.

History of Full Disclosure

During the early years of computers and networks, bug secrecy was the norm. When users and researchers found vulnerabilities in a software product, they would quietly alert the vendor. In theory, the vendor would then fix the vulnerability. After CERT was founded in 1988, it became a clearinghouse for vulnerabilities. People would send newly discovered vulnerabilities to CERT. CERT would then verify them, alert the vendors, and publish the details (and the fix) once the fix was available.

The problem with this system is that the vendors didn't have any motivation to fix vulnerabilities. CERT wouldn't publish until there was a fix, so there was no urgency. It was easier to keep the vulnerabilities secret. There were incidents of vendors threatening researchers if they made their findings public, and smear campaigns against researchers who announced the existence of vulnerabilities (even if they omitted details). And so many vulnerabilities remained unfixed for years.


The full disclosure movement was born out of frustration with this process. Once a vulnerability is published, public pressures give vendors a strong incentive to fix the problem quickly. For the most part, this has worked. Today, many researchers publish vulnerabilities they discover on mailing lists such as Bugtraq. The press writes about the vulnerabilities in the computer magazines. The vendors scramble to patch these vulnerabilities as soon as they are publicized, so they can write their own press releases about how quickly and thoroughly they fixed things. The full disclosure movement is improving Internet security.

At the same time, hackers use these mailing lists to learn about vulnerabilities and write exploits. Sometimes the researchers themselves write demonstration exploits. Sometimes others do. These exploits are used to break into vulnerable computers and networks, and greatly decrease Internet security. In his essay, Culp points to Code Red, Li0n, Sadmin, Ramen, and Nimda as examples of malicious code written after researchers demonstrated how particular vulnerabilities worked.

Those against the full-disclosure movement argue that publishing vulnerability details does more harm than good by arming the criminal hackers with tools they can use to break into systems. Security is much better served, they counter, by keeping the exact details of vulnerabilities secret.

(Continued on next page.)

Why Don't Large Organizations Install Patches?



- Costs of Installing a Patch
 - Fixed cost to evaluate patch, assess integration issues
 - Variable cost to install patch across network
- Costs of Not Patching Network
 - Security risk
- Result: The smaller the network, the more likely the patch gets installed.

12

INTELLIGENT ALERT. INSTANT RESPONSE. IMMEDIATE DEFENSE.

Full Disclosure (cont.)

Full-disclosure proponents counter that this assumes that the researcher who publicizes the vulnerability is always the first one to discover it, which simply isn't true. Sometimes vulnerabilities have been known by attackers (sometimes passed about quietly in the hacker underground) for months or years before the vendor ever found out. The sooner a vulnerability is publicized and fixed, the better it is for everyone, they say. And returning to bug secrecy would only bring back vendor denial and inaction.

That's the debate in a nutshell: Is the benefit of publicizing an attack worth the increased threat of the enemy learning about it? Should we reduce the Window of Exposure by trying to limit knowledge of the vulnerability, or by publishing the vulnerability to force vendors to fix it as quickly as possible?

What we've learned during the past eight or so years is that full disclosure helps much more than it hurts. Since full disclosure has become the norm, the computer industry has transformed itself from a group of companies that ignores security and belittles vulnerabilities into one that fixes vulnerabilities as quickly as possible. A few companies are even going further, and taking security seriously enough to attempt to build quality software from the beginning: to fix vulnerabilities before the product is released. And far fewer problems are showing up first in the hacker underground, attacking people with absolutely no warning. It used to be that vulnerability information was only available to a select few: security researchers and hackers who were connected enough in their respective communities. Now it is available to everyone.

This democratization is important. If a known vulnerability exists and you don't know about it, then you're making security decisions with substandard data. Word will eventually get out -- the Window of Exposure will grow -- but you have no control, or knowledge, of when or how. All you can do is hope that the bad guys don't find out before the good guys fix the problem. Full disclosure means that everyone gets the information at the same time, and everyone can act on it.

And detailed information is required. If a researcher just publishes vague statements about the vulnerability, then the vendor can claim that it's not real. If the researcher publishes scientific details without example code, then the vendor can claim that it's just theoretical. The only way to make vendors sit up and take notice is to publish details: both in human- and computer-readable form. (Microsoft is guilty of both of these practices, using their PR machine to deny and belittle vulnerabilities until they are demonstrated with actual code.) And demonstration code is the only way to verify that a vendor's vulnerability patch actually patched the vulnerability.

This free information flow, of both description and proof-of-concept code, is also vital for security research. Research and development in computer security has blossomed in the past decade, and much of that can be attributed to the full-disclosure movement. The ability to publish research findings -- both good and bad -- leads to better security for everyone. Without publication, the security

(Continued on next page.)

C
Counterpane™
Internet Security

Why did Firewalls Succeed?

- It's not because they're effective
 - Most firewalls are configured to be ineffective
 - E-mail encryption is effective, too
- It's because auditors started requiring firewalls
 - Cost of Adding a Firewall
 - Money
 - User annoyance
 - Cost of not Adding a Firewall
 - Failing an audit
 - Not doing what the rest of the industry is doing

13

INTELLIGENT ALERT. INSTANT RESPONSE. IMMEDIATE DEFENSE.

Full Disclosure (cont.)

community can't learn from each other's mistakes. Everyone must operate with blinders on, making the same mistakes over and over. Full disclosure is essential if we are to continue to improve the security of our computers and networks.

Bug Secrecy Example

You can see the problems with bug secrecy in the digital-rights-management industry. The DMCA has enshrined the bug secrecy paradigm into law; in most cases it is illegal to publish vulnerabilities or automatic hacking tools against copy-protection schemes. Researchers are harassed, and pressured against distributing their work. Security vulnerabilities are kept secret. And the result is a plethora of insecure systems, their owners blustering behind the law hoping that no one finds out how bad they really are.

The result is that users can't make intelligent decisions on security. Here's one example: A few months ago, security researcher Niels Ferguson found a security flaw in Intel's HDCP Digital Video Encryption System, but withheld publication out of fear of being prosecuted under the DMCA. Intel's reaction was reminiscent of the pre-full-disclosure days: they dismissed the break as "theoretical" and maintained that the system was still secure. Imagine you're thinking about buying Intel's system. What do you do? You have no real information, so you have to trust either Ferguson or Intel.

Here's another: In November, a release of the Linux kernel came without the customary detailed information about the OS's security. The developers cited fear of the DMCA as a reason why those details were withheld. Imagine you're evaluating operating systems: Do you feel more or less confident about the security the Linux kernel version 2.2, now that you have no details?

Full Disclosure and Responsibility

Culp has a point when he talks about responsibility. (Of course, Scott is avoiding "mea Culpa.") The goal here is to improve security, not to arm people who break into computers and networks. Automatic hacking tools with easy point-and-click interfaces, ready made for script kiddies, cause a lot of damage to organizations and their networks. There are such things as responsible and irresponsible disclosure. It's not always easy to tell the difference, but I have some guidelines.

First, I am opposed to attacks that primarily sow fear. Publishing vulnerabilities that there's no real evidence for is bad. Publishing vulnerabilities that are more smoke than fire is bad. Publishing vulnerabilities in critical systems that cannot be easily fixed and whose exploitation will cause serious harm (e.g., the air traffic control system) is bad.

Second, I believe in giving the vendor advance notice. CERT took this to an extreme, sometimes giving the vendor years to fix the problem. I'd like to see the researcher tell the vendor that he will publish the vulnerability in a few weeks, and then stick to that promise. Currently CERT gives vendors 45 days, but will disclose vulnerability information immediately for paid subscribers. Microsoft

(Continued on next page.)

C
Counterpane™
Internet Security

Current Risk Management Strategies

- Disclaim liability
 - UCITA
- Transfer liability
 - Shrink-wrap software licenses
- Invoke a "best practice" standard
- Results:
 - Security is opaque: hard to assess and analyze
 - There's no real incentive to improve security

14
INTELLIGENT ALERT. INSTANT RESPONSE. IMMEDIATE DEFENSE.

Full Disclosure (cont.)

proposes a 30-day secrecy period. While this is a good idea in theory, creating a special insider group of people "in the know" has its own set of problems.

Third, I agree with Culp that it is irresponsible, and possibly criminal, to distribute easy-to-use exploits. Reverse engineering security systems, discovering vulnerabilities, writing research papers about them, and even writing demonstration code, benefits research; it makes us smarter at designing secure systems. Distributing exploits just make us more vulnerable. I'd like to get my hands on the people who write virus creation kits, for example. They've got a lot to answer for.

This is not clear-cut: there are tools that do both good and bad, and sometimes the difference is merely marketing. Dan Farmer was vilified for writing SATAN; today, vulnerability assessment tools are viable security administration products. Remote administration tools look a lot like Back Orifice (although less feature-rich). L0phtCrack is a hacker tool to break weak passwords as a prelude to an attack, but LC 3.0 is sold as a network administration tool to test for weak passwords. And the program that Dmitry Sklyarov was arrested for writing has legitimate uses. In fact, most tools have both good and bad uses, and when in doubt I believe it is better to get the information in the hands of people who need it, even if it means that the bad guys get it too.

One thing to pay attention to is the agenda of the researcher. Publishing a security vulnerability is often a publicity play; the researcher is looking to get his own name in the newspaper by successfully bagging his prey. The publicizer often has his own agenda: he's a security consultant, or an employee of a company that offers security products or services. I am a little tired of companies that publish vulnerabilities in order to push their own product or service. Although, of course, a non-altruistic motive does not mean that the information is bad.

I like the "be part of the solution, not part of the problem" metric. Researching security is part of the solution. Convincing vendors to fix problems is part of the solution. Sowing fear is part of the problem. Handing attack tools to clueless teenagers is part of the problem.


The Inevitability of Security Vulnerabilities

None of this would be an issue if software were engineered properly in the first place. A security vulnerability is a programming mistake: either an out- and-out mistake like a buffer overflow, which should have been caught and prevented, or an opening introduced by a lack of understanding the interactions in a complex piece of code. If there were no security vulnerabilities, there would be no problem. It's poor software quality that causes this mess in the first place.

While this is true -- software vendors uniformly produce shoddy software -- the sheer complexity of modern software and networks means that vulnerabilities, lots of vulnerabilities, are inevitable. They're in every major software package. Each time Microsoft releases an OS it crows about how extensive the testing was and how secure it is, and every time it contains more security vulnerabilities than the previous OS. I don't believe this trend will reverse itself anytime soon.

(Continued on next page.)

Fixing the Problem



- This is a business problem, not a technological problem
- If you want to change the risk-management analysis, change the costs
 - Security must affect the bottom line in an obvious way
- In order to improve security, the CEO needs to care

15

INTELLIGENT ALERT. INSTANT RESPONSE. IMMEDIATE DEFENSE.

Full Disclosure (cont.)

Vendors don't take security seriously because there is no market incentive for them to, and no adverse effects when they don't. I have long argued that software vendors should not be exempt from the product liability laws that govern the rest of commerce. When this happens, vendors will do more than pay lip service to security vulnerabilities: they will fix them as quickly as possible. But until then, full disclosure is the only way we have to motivate vendors to act responsibly.

Microsoft's motives in promoting bug secrecy are obvious: it's a whole lot easier to squelch security information than it is to fix problems, or design products securely in the first place. Microsoft's steady stream of public security vulnerabilities has led many people to question the security of their future products. And with analysts like Gartner advising people to abandon Microsoft IIS because of all its insecurities, giving customers less security information would be good for business.

Bug secrecy is a viable solution only if software vendors are followers of W. Edwards Deming's quality management principles. The longer a bug remains unfixed, the bigger a problem it is. And because the number of systems on the Internet is constantly growing, the longer a security vulnerability remains unfixed, the larger the window of exposure. If companies believe this and then act accordingly, then there is a powerful argument for secrecy.

However, history shows this isn't the case. Read Scott Culp's essay; he did not say: "Hey guys, if you have a bug, send it to me and I'll make sure it gets fixed pronto." What he did was to rail against the publication of vulnerabilities, and ask researchers to keep details under their hats. Otherwise, he threatened, "vendors will have no choice but to find other ways to protect their customers." That's the attitude that makes full disclosure the only viable way to reduce the window of vulnerability.

In his essay, Culp compares the practice of publishing vulnerabilities to shouting "Fire" in a crowded movie theater. What he forgets is that there actually is a fire; the vulnerabilities exist regardless. Blaming the person who disclosed the vulnerability is like imprisoning the person who first saw the flames. Disclosure does not create security vulnerabilities; programmers create them, and they remain until other programmers find and remove them. Everyone makes mistakes; they are natural events in the sense that they inevitably happen. But that's no excuse for pretending that they are caused by forces out of our control, and mitigated when we get around to it.



Four Steps Towards Security

- Step 1: Enforce Liabilities
- Step 2: Allow Parties to Transfer Liabilities
- Step 3: Provide Mechanisms to Reduce Risk
- Step 4: Rational Prosecution Leads to Deterrence

16

INTELLIGENT ALERT. INSTANT RESPONSE. IMMEDIATE DEFENSE.

Military History and Network Security

In warfare, the defender's military advantage comes from two broad strengths: the ability to quickly react to an attack, and the ability to control the terrain.

The first strength is probably the most important; a defender can more quickly shift forces to resupply existing forces, shore up defense where it is needed, and counterattack. Here we see the same themes from elsewhere in this booklet: how detection and response are critical, the need for trained experts to quickly analyze and react to attacks, and the importance of vigilance. I've built Counterpane's MSM service around these very principles, precisely because it can dramatically shift the balance from attacker to defender.

The defender's second strength also gives him a strong advantage. He has better knowledge of the terrain: where the good hiding places are, where the mountain passes are, how to sneak through the caves. He can modify the terrain: building castles or SAM batteries, digging trenches or tunnels, erecting guard towers or pillboxes. And he can choose the terrain on which to stand and defend: behind the stone wall, atop the hill, on the far side of the bridge, in the dense jungle. The defender can use terrain to his maximum advantage; the attacker is stuck with whatever terrain he is forced to traverse.

On the Internet, this second advantage is one that network defenders seldom take advantage of: knowledge of the network. The network administrator knows exactly how his network is built (or, at least, he should), what it is supposed to do, and how it is supposed to do it. Any attacker except a knowledgeable insider has no choice but to stumble around, trying this and that, trying to figure out what's where and who's connected to whom. And it's about time we exploited this advantage.

Think about burglar alarms. The reason they work is that the attacker doesn't know they're there. He might successfully bypass a door lock, or sneak in through a second-story window, but he doesn't know that there is a pressure plate under this particular rug, or an electric eye across this particular doorway. MacGyver-like antics aside, any burglar wandering through a building wired with alarms is guaranteed to trip something sooner or later.

Traditional computer security has been static: install a firewall, configure a PKI, add access-control measures, and you're done. Real security is dynamic. The defense has to be continuously vigilant, always ready for the attack. The defense has to be able to detect attacks quickly, before serious damage is done. And the defense has to be able to respond to attacks effectively, repelling the attacker and restoring order.

This kind of defense is possible in computer networks. It starts with effective sensors: firewalls, well-audited servers and routers, intrusion-detection products, network burglar alarms. But it also includes people: trained security experts that can quickly separate the false alarms from the real attacks, and who know how to respond. It includes an MSM service. This is security through process. This is security that recognizes that human intelligence is vital for a strong defense, and that automatic software programs just don't cut it.

Step 1: Enforce Liabilities



- The purpose of liability is to hold someone accountable
- Today, there are no real consequences for having bad security
- Enforcing liabilities will change that
 - Software manufacturers liable for defective products
 - Companies liable for mishandling their customers' data
- Liability will increase the transparency of security

17

INTELLIGENT ALERT. INSTANT RESPONSE. IMMEDIATE DEFENSE.

Liability Enforcement Options



- Industry-defined standards
 - Not legally binding
 - More rhetoric than results
 - Difficult to enforce
- Federal regulation
 - Difficult to implement and enforce
 - Would face strong lobbying opposition from vendors
- Lawsuits
 - Takes years to settle down
 - Despite problems with licensing agreements, few have succeeded
 - The most likely option

18

INTELLIGENT ALERT. INSTANT RESPONSE. IMMEDIATE DEFENSE.

Liability and Security

Today, computer security is at a crossroads. It's failing, regularly, and with increasingly serious results. I believe it will improve eventually. In the near term, the consequences of insecurity will get worse before they get better. And when they get better, the improvement will be slow and will be met with considerable resistance. The engine of this improvement will be liability -- holding software manufacturers accountable for the security and, more generally, the quality of their products -- and the timetable for improvement depends wholly on how quickly security liability permeates cyberspace.

Network security is not a problem that technology can solve. Security has a technological component, but businesses approach security as they do any other business risk: in terms of risk management. Organizations optimize their activities to minimize their cost/risk ratio, and understanding those motivations is key to understanding computer security today.

For example, most organizations don't spend a lot of money on network security. Why? Because the costs are significant: time, expense, reduced functionality, frustrated end users. On the other hand, the costs of ignoring security and getting hacked are small: the possibility of bad press and angry customers, maybe some network downtime, none of which is permanent. And there's some regulatory pressure, from audits or lawsuits, that add additional costs. The result: a smart organization does what everyone else does, and no more.

The same economic reasoning explains why software vendors don't spend a lot of effort securing their products. The costs of adding good security are significant -- large expenses, reduced functionality, delayed product releases, annoyed users -- while the costs of ignoring security are minor: occasional bad press, and maybe some users switching to competitors' products. Any smart software vendor will talk big about security, but do as little as possible.


Think about why firewalls succeeded in the marketplace. It's not because they're effective; most firewalls are installed so poorly as not to be effective, and there are many more effective security products that have never seen widespread deployment. Firewalls are ubiquitous because auditors started demanding firewalls. This changed the cost equation for businesses. The cost of adding a firewall was expense and user annoyance, but the cost of not having a firewall was failing an audit. And even worse, a company without a firewall could be accused of not following industry best practices in a lawsuit. The result: everyone has a firewall, whether it does any good or not.

Network security is a business problem, and the only way to fix it is to concentrate on the business motivations. We need to change the costs; security needs to affect an organization's bottom line in an obvious way. In order to improve computer security, the CEO must care. In order for the CEO to care, it must affect the stock price and the shareholders.

I have a three-step program towards improving computer and network security. None of the steps have anything to do with the technology; they all have to do with businesses, economics, and people.

(Continued on next page.)

Problems with Liability Enforcement



- The complexity of the problem makes assigning liabilities difficult
- The international nature of the Internet makes uniform liabilities difficult
- Free software could be choked to death, since individual programmers are unlikely to have sufficient resources to accept liability
- Liability enforcement has the potential to stifle still-nascent software and computer industry

19

INTELLIGENT ALERT. INSTANT RESPONSE. IMMEDIATE DEFENSE.

Liability and Security (cont.)

Step one: enforce liabilities. This is essential. Today there are no real consequences for having bad security, or having low-quality software of any kind. In fact, the marketplace rewards low quality. More precisely, it rewards early releases at the expense of almost all quality. If we expect CEOs to spend significant resources on security -- especially the security of their customers -- they must be liable for mishandling their customers' data. If we expect software vendors to reduce features, lengthen development cycles, and invest in secure software development processes, they must be liable for security vulnerabilities in their products.

Legislatures could impose liability on the computer industry, by forcing software manufacturers to live with the same product liability laws that affect other industries. If software manufacturers produced a defective product, they would be liable for damages. Even without this, courts could start imposing liability-like penalties on software manufacturers and users. This is starting to happen. A U.S. judge forced the Department of Interior to take its network offline, because it couldn't guarantee the safety of American Indian data it was entrusted with. Several cases have resulted in penalties against companies who used customer data in violation of their privacy promises, or who collected that data using misrepresentation or fraud. And judges have issued restraining orders against companies with insecure networks that are used as conduits for attacks against others.

However it happens, liability changes everything. Currently, there is no reason for a software company not to offer more features, more complexity. Liability forces software companies to think twice before changing something. Liability forces companies to protect the data they're entrusted with.

Step two: allow parties to transfer liabilities. This will happen automatically, because this is what insurance companies do. The insurance industry turns variable-cost risks into fixed expenses. They're going to move into cyber-insurance in a big way. And when they do, they're going to drive the computer security industry...just like they drive the security industry in the brick-and-mortar world.


A company doesn't buy security for its warehouse--strong locks, window bars, or an alarm system--because it makes it feel safe. It buys that security because its insurance rates go down. The same thing will hold true for computer security. Once enough policies are being written, insurance companies will start charging different premiums for different levels of security. Even without legislated liability, the CEO will start noticing how his insurance rates change. And once the CEO starts buying security products based on his insurance premiums, the insurance industry will wield enormous power in the marketplace. They will determine which security products are ubiquitous, and which are ignored. And since the insurance companies pay for the actual liability, they have a great incentive to be rational about risk analysis and the effectiveness of security products.

And software companies will take notice, and will increase security in order to make the insurance for their products affordable.

Step three: provide mechanisms to reduce risk. This will happen automatically, and be entirely market driven, because it's what the insurance industry wants. Moreover, they want it done in standard models that they can build policies around. They're going to look to

(Continued on next page.)

Privacy Liability Litigation



- Disclosure of customer data in violation of privacy promises:
 - US Bankcorp litigation, Bank of America litigation, Amazon.com litigation, FTC vs. Toysmart.com
- Obtaining personal information by misrepresentation or fraud:
 - FTC vs. ReverseAuction.com, Valan vs. General Motors
- Tracking or monitoring Internet users without permission or disclosure:
 - Steart vs. Yahoo, DoubleClick litigation, AOL litigation, Toys 'R' Us litigation

20

INTELLIGENT ALERT. INSTANT RESPONSE. IMMEDIATE DEFENSE.

Liability and Security (cont.)

security processes: processes of secure software development before systems are released, and processes of protection, detection, and response for corporate networks and systems. And more and more, they're going to look towards outsourced services.


The insurance industry prefers security outsourcing, because they can write policies around those services. It's much easier to design insurance around a standard set of security services delivered by an outside vendor than it is to customize a policy for each individual network.

Actually, this isn't a three-step program. It's a one-step program with two inevitable consequences. Enforce liability, and everything else will flow from it. It has to.

Much of Internet security is a common: an area used by a community as a whole. Like all commons, keeping it working benefits everyone, but any individual can benefit from exploiting it. (Think of the criminal justice system in the real world.) In our society we protect our commons -- our environment, healthy working conditions, safe food and drug practices, lawful streets, sound accounting practices -- by legislating those goods and by making companies liable for taking undue advantage of those commons. This kind of thinking is what gives us bridges that don't collapse, clean air and water, and sanitary restaurants. We don't live in a "buyer beware" society; we hold companies liable for taking advantage of buyers.

There's no reason to treat software any differently from other products. Today Firestone can produce a tire with a single systemic flaw and they're liable, but Microsoft can produce an operating system with multiple systemic flaws discovered per week and not be liable. This makes no sense, and it's the primary reason security is so bad today.

Liability Enforcement: 2001



- A U.S. Court forced the Department of Interior offline, because they couldn't adequately protect Native American data
- Judges are issuing restraining orders against companies who are launching pads for attacks against third parties
- Microsoft's "responsible disclosure" movement is an attempt to turn attention from the entity responsible for the vulnerability in the first place

21INTELLIGENT ALERT. INSTANT RESPONSE. IMMEDIATE DEFENSE.

Judges Punish Bad Security

The first involves the U.S. Department of Interior. There's an ongoing litigation between Native Americans and the U.S. Government regarding mishandling of funds. After seeing for himself how insecure the Department's computers were, and that it was possible for someone to alter records and divert funds, a U.S. District Judge ordered the department to disconnect its computers from the Internet until its network is secured.

The second involves a couple of Web hosting companies. One day, C.I. Host was hit with a denial-of-service attack. They traced at least part of the attack to companies hosted by Exodus Communications. C.I. Host filed an injunction against Exodus, alleging that they committed or allowed a third party to commit a DOS attack. A Texas judge issued a temporary restraining order against three of Exodus's customers, forcing them to disconnect from the Internet until they could prove that the vulnerabilities leading to the DOS attack had been fixed.

I like this kind of stuff. It forces responsibility. It tells companies that if they can't make their networks secure, they have no business being on the Internet. It may be Draconian, but it gets the message across.


On the Internet, as on any connected system, security has a ripple effect. Your security depends on the actions of others, often of others you can't control. This is the moral of the widely reported distributed denial-of-service attacks in February 2000: the security of the computers at eBay, Amazon, Yahoo, and CNN.com depended on the security of the computers at the University of California at Santa Barbara. If Eli Lilly has bad computer security, then your identity as a Prozac user may be compromised. If Microsoft can't keep your Passport data secure, then your online identify can be compromised. It's hard enough making your own computers secure; now you're expected to police the security of everyone else's networks.

This is where the legal system can step in. I like to see companies told that they have no business putting the security of others at risk. If a company's computers are so insecure that hackers routinely break in and use them as a launching pad for further attacks, get them off the Internet. If a company can't secure the personal information it is entrusted with, why should it be allowed to have that information? If a company produces a software product that compromises the security of thousands of users, maybe they should be prohibited from selling it.

I know there are more instances of this happening. I've seen it, and some of my colleagues have too. Counterpane acquired two customers recently, both of whom needed us to improve their network's security within hours, in response to this sort of legal threat. We came in and installed our monitoring service, and they were able to convince a judge that they should not be turned off. I see this as a trend that will increase, as attacked companies look around for someone to share fault with.

This kind of thing certainly won't solve our computer security problems, but at least it will remind companies that they can't dodge responsibility forever. The Internet is a vast commons, and the actions of one affect the security of us all.

Step 2: Allow Parties to Transfer Liabilities



- Insurance spreads liability risk among a group
- Insurance converts a variable liability into a fixed liability
- Insurance is the CEO's primary risk-management tool
- In the real world, insurance drives the security business

22

INTELLIGENT ALERT. INSTANT RESPONSE. IMMEDIATE DEFENSE.

Insurance and the Future of Network Security

Eventually, the insurance industry will subsume the computer security industry. Not that insurance companies will start marketing security products, but rather that the kind of firewall you use -- along with the kind of authentication scheme you use, the kind of operating system you use, and the kind of network monitoring scheme you use -- will be strongly influenced by the constraints of insurance.

Consider security, and safety, in the real world. Businesses don't install building alarms because it makes them feel safer; they do it because they get a reduction in their insurance rates. Building-owners don't install sprinkler systems out of affection for their tenants, but because building codes and insurance policies demand it. Deciding what kind of theft and fire prevention equipment to install are risk management decisions, and the risk taker of last resort is the insurance industry.

This is sometimes hard for computer techies to understand, because the security industry has trained them to expect technology to solve their problems. Remember when all you needed was a firewall, and then you were safe? Remember when it was an intrusion detection product? Or a PKI? I think the current wisdom is that all you need is biometrics, or maybe smart cards.

The real world doesn't work this way. Businesses achieve security through insurance. They take the risks they are not willing to accept themselves, bundle them up, and pay someone else to make them go away. If a warehouse is insured properly, the owner really doesn't care if it burns down or not. If he does care, he's underinsured. Similarly, if a network is insured properly, the owner won't care whether it is hacked or not.

This is worth repeating: a properly insured network is immune to the effects of hacking. Concerned about denial-of-service attacks? Get bandwidth interruption insurance. Concerned about data corruption? Get data integrity insurance. (I'm making these policy names up, here.) Concerned about negative publicity due to a widely publicized network attack? Get a rider on your good name insurance that covers that sort of event. The insurance industry isn't offering all of these policies yet, but it is coming.

When I talk about this future at conferences, a common objection I hear is that premium calculation is impossible. Again, this is a technical mentality talking. Sure, insurance companies like well-understood risk profiles and carefully calculated premiums. But they also insure satellite launches and the palate of wine critic Robert Parker. If an insurance company can protect Tylenol against some lunatic putting a poisoned bottle on a supermarket shelf, anti-hacking insurance will be a snap.

(Continued on next page.)

Insurance Drives Security



- What the insurance industry needs are standard risk models
- What the insurance industry needs are standard protection profiles
- What the insurance industry needs is more security, not empty press releases

23

INTELLIGENT ALERT. INSTANT RESPONSE. IMMEDIATE DEFENSE.

Insurance and the Future of Network Security (cont.)

Imagine the future.... Every business has network security insurance, just as every business has insurance against fire, theft, and any other reasonable threat. To do otherwise would be to behave recklessly and be open to lawsuits. Details of network security become check boxes when it comes time to calculate the premium. Do you have a firewall? Which brand? Your rate may be one price if you have this brand, and a different price if you have another brand. Do you have a service monitoring your network? If you do, your rate goes down this much.


This process changes everything. What will happen when the CFO looks at his premium and realizes that it will go down 50% if he gets rid of all his insecure Windows operating systems and replaces them with a secure version of Linux? The choice of which operating system to use will no longer be 100% technical. Microsoft, and other companies with shoddy security, will start losing sales because companies don't want to pay the insurance premiums. In this vision of the future, how secure a product is becomes a real, measurable, feature that companies are willing to pay for...because it saves them money in the long run.

Other systems will be affected, too. Online merchants and brick-and-mortar merchants will have different insurance premiums, because the risks are different. Businesses can add authentication mechanisms -- public-key certificates, biometrics, smart cards -- and either save or lose money depending on their effectiveness. Computer security "snake-oil" peddlers who make outlandish claims and sell ridiculous products will find no buyers as long as the insurance industry doesn't recognize their value. In fact, the whole point of buying a security product or hiring a security service will not be based on threat avoidance; it will be based on risk management.

And it will be about time. Sooner or later, the insurance industry will sell everyone anti-hacking policies. It will be unthinkable not to have one. And then we'll start seeing good security rewarded in the marketplace.

24

Insurance Creates a Marketplace for Security



- Vendors could accept liabilities from their customers
- Customers could accept liabilities from the vendors
- Insecurity could be bought and sold
 - Just like some pollution rights are
- Insurance would provide real incentive to increase security

INTELLIGENT ALERT. INSTANT RESPONSE. IMMEDIATE DEFENSE.

Richard Clarke on 9/11's Lessons


Richard Clarke, the U.S. government's cyber-security czar, often talks about six lessons from September 11th. His talks start out well, but fall apart when he states that businesses should, out of the goodness of their hearts and concern for their way of life, produce and use more secure products. And his lessons don't just apply to combating cyber-terrorism; they apply to everyday network security.

1. "We have enemies." Everyone does. Companies have competitors. People have others who don't like them. Some enemies target us by name, others simply want to rob someone and don't care whom. Too many organizations justify their inattention to security by saying: "Who would want to attack us?" That just doesn't make sense.
2. "Don't underestimate them." Don't. Whether it is a DVD pirate living in a country with no copyright laws, or a hacker kid who spends days trying to break into networks, cyberspace attackers have proven to be better funded, smarter, and more tenacious than people estimate. If you assume your enemies won't be able to figure out your defenses and bypass them, you're not paying attention.
3. "They will use our technology against us." This is especially true in cyberspace. Almost all attacks involve using the very network being attacked. Maybe it's a vulnerability in the software; maybe it's a feature that should never have been created. Hacking is judo: using network software to do things it was never intended to do.
4. "They will attack the seams of our technology." As bad as most cryptography is out there, it's almost always easier to break a system by some other method. Attacks on the seams -- the places where different technologies come together -- are more fruitful. Think of the FBI reading PGP-encrypted mail by installing a keyboard sniffer, or people who bypass copy-protection controls by mimicking them rather than breaking them. This lesson is obvious to anyone who has broken security software.
5. "Our technology is surprisingly interdependent." That's certainly clear. We've seen vulnerabilities in IIS affect all sorts of systems. We've seen malicious code use features of Microsoft Word and Outlook to spread. A single SNMP vulnerability affects hundreds of products. Interdependence is how the Internet works. It's also how it fails.
6. "The only way to solve this problem is for government and industry to work together." This is more subjective, but I agree with it. I don't think that industry can do it alone, mostly because they have no incentive to do it. I don't think that government can do it alone, because they don't have the capability. Clarke seems to think that it's government's job to provide some funding, high-level coordination, and general cheerleading. I think it's government's job to provide a financial incentive to business. If you want to fix network security, hack the business model.

Clarke spends a lot of time visiting companies and talking to them about security. I'm sure the CEOs give him a warm reception, and tell him that they'll make their stuff more secure. But what happens a month later, when budgets are tight and a release date looms? Will the CEO remember his promise to Clarke, or will he listen to the demands of the market? If the government really wants the CEO to care, it's going to have to make security a market force.

I don't see any other way.

Hacking Insurance will Become Ubiquitous



- Someday, soon, computer-security insurance will be as common as fire insurance
- This will give insurance companies enormous clout in the security marketplace
- They will be able to drive products and services
- They will want standard products and services, so they can better write policies

- The insurance industry is offering policies anyway, but it is unclear what they're worth
- The insurance industry will continue to offer policies

25

INTELLIGENT ALERT. INSTANT RESPONSE. IMMEDIATE DEFENSE.

Prevention, Detection, and Response

Most computer security is sold as a prophylactic: encryption prevents eavesdropping, firewalls prevent unauthorized network access, PKI prevents impersonation. To the world at large, this is a strange marketing strategy. A door lock is never sold with the slogan: "This lock prevents burglaries." No one ever asks to purchase "a device that will prevent murder." But computer security products are sold that way all the time. Companies regularly try to buy "a device that prevents hacking." This is no more possible than an anti-murder device.

When you buy a safe, it comes with a rating. 30TL-30 minutes, tools. 60TRTL-60 minutes, torch and tools. What this means is that a professional safecracker, with safecracking tools and an oxyacetylene torch, will break open the safe in an hour. If an alarm doesn't sound and guards don't come running within that hour, the safe is worthless. The safe buys you time; you have to spend it wisely.


Real-world security includes prevention, detection, and response. If the prevention mechanisms were perfect, you wouldn't need detection and response. But no prevention mechanism is perfect. This is especially true for computer networks. All software products have security bugs, most network devices are misconfigured, and users make all sorts of mistakes. Without detection and response, the prevention mechanisms only have limited value. They're fragile. And detection and response are not only more cost effective, but also more effective, than piling on more prevention.

On the Internet, this translates to monitoring. When a forensics team investigates an intrusion, they comb through the audit logs of the network's routers, servers, firewalls, etc. Using those logs, they piece together the attacker's actions: how he got in, what he did, what he stole. If someone can monitor those audit logs in real time, he could figure out what the attacker IS DOING. And if he can respond fast enough, he can repel the attacker before he does real damage.

That's real security. It doesn't matter how the attacker gets in, or what he is doing. If there are enough motion sensors, electric eyes, and pressure plates in your house, you'll catch the burglar regardless of how he got in. If you are monitoring your network carefully enough, you'll catch a hacker regardless of what vulnerability he exploited to gain access. And if you can respond quickly and effectively, you can repel the attacker before he does any damage. Good detection and response can make up for imperfect prevention.

And prevention systems are never perfect. No bank ever says: "Our safe is so good, we don't need an alarm system." No museum ever says: "Our door and window locks are so good, we don't need night watchmen." Detection and response are how we get security in the real world, and they're the only way we can possibly get security on the Internet. CIOs must invest in network monitoring services if they are to properly manage the risks associated with their network infrastructure.

Step 3: Provide Mechanisms to Reduce Risk



- There are two places to reduce risk:
 - Before software is released
 - After software is released
- We need to do both
- Techniques and processes to improve software quality
 - *Building Secure Software*, by Viega and McGraw
- Protection, detection, and response to improve network security

26

INTELLIGENT ALERT. INSTANT RESPONSE. IMMEDIATE DEFENSE.

Monitoring Network Security

Network monitoring implies several things. It implies a series of sensors in and around the network. Luckily, these are already in place. Every firewall produces a continuous stream of audit messages. So does every router and server. IDSs send messages when they notice something. Every other security product generates alarms in some way.

But these sensors by themselves do not offer security. You have to assume that the attacker is in full possession of the specifications for these sensors, is well aware of their deficiencies, and has tailored his attack accordingly. He may even have passwords that let him masquerade as a legitimate user. Only another human has a chance of detecting some anomalous behavior that gives him away.

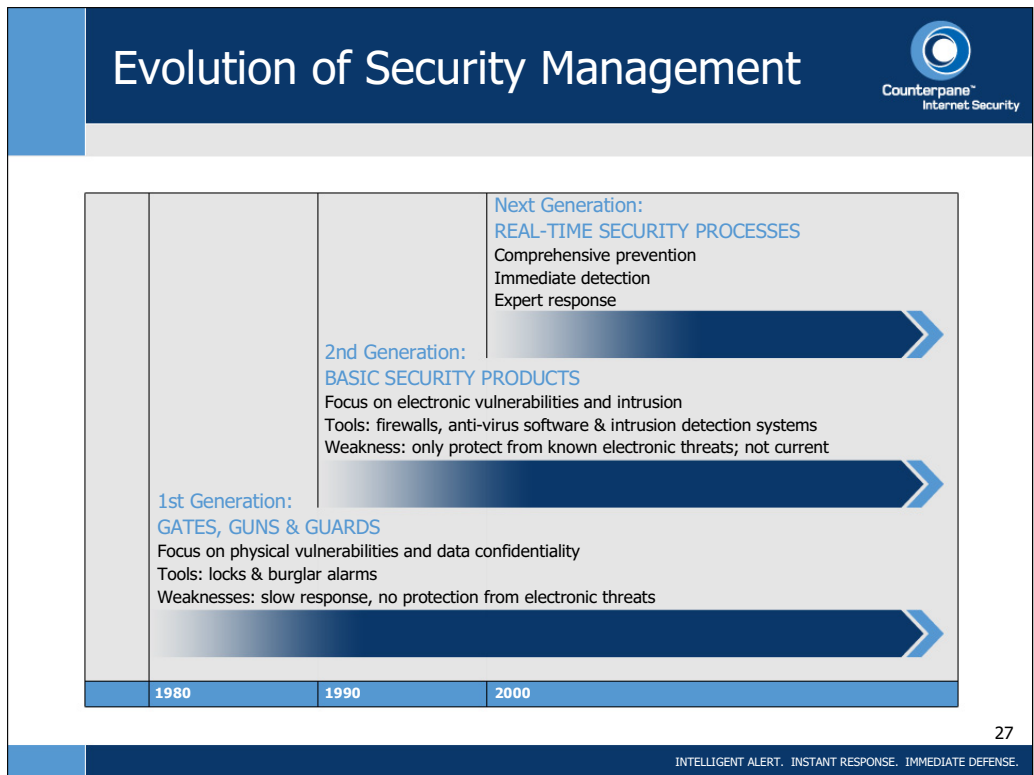
The first step is intelligent alert. Network attacks can be much more subtle than a broken window. Much depends on context. Software can filter the tens of megabytes of audit information a medium-sized network can generate in a day, but software is too easy for an attacker to fool. Intelligent alert requires people. People to analyze what the software finds suspicious. People to delve deeper into suspicious events, determining what is really going on. People to separate false alarms from real attacks. People who understand context.

By itself, an alert is only marginally useful. More important is to know how to respond. This is the second step of good network monitoring. Every attack has a response. It could be as simple as shutting off a particular IP address to repel an attacker. It could be as drastic as taking a corporate network off the Internet. Again, people are the key. Software can only provide generic information; real understanding requires experts.

And finally, the response must be integrated with the business needs of the organization. Security engineers only see half the information. They understand attacks and their security significance, but they don't understand the business ramifications. A large e-business might keep its Web site up and running even if it is being attacked; preventing the loss of revenue may be more important than the site's immediate security. On the other hand, a law firm may have the exact opposite response; the sanctity of its customers' data might be more important than having its Web site available.

This is detection and response as applied to computer networks. Network devices produce megabytes of audit information daily. Automatic search tools sift through those megabytes, looking for telltale signs of attacks. Expert analysts examine those telltales, understanding what they mean and determining how to respond. And the owner of the network—the organization—makes security decisions based on ongoing business concerns.

To make network monitoring work, people are needed every step of the way. Software is just too easy to fool. Software doesn't think, doesn't question, doesn't adapt. Without people, computer security software is just a static defense. Marry software with experts, and you have a whole different level of security.



Monitoring and Resilient Security

During the course of the year 2000, several groups of Eastern European hackers broke into at least 40 companies' Web sites, stole credit card numbers, and in some cases tried to extort money from their victims. The network vulnerabilities exploited by these criminals were known, and patches that closed them were available-but none of the companies had installed them. In January 2001, the Ramen worm targeted known vulnerabilities in several versions of Red Hat Linux. None of the thousands of infected systems had their patches up to date. In October 2000, Microsoft was molested by unknown hackers who wandered unchallenged through their network, accessing intellectual property, for weeks or months. According to reports, the attackers would not have been able to break in if Microsoft's patches had been up to date. The series of high-profile credit card thefts in January 2000, including the CD Universe incident, were also the result of uninstalled patches. A patch issued eighteen months previously would have protected these companies.

What's going on here? Isn't anyone installing security patches anymore? Doesn't anyone care?


What's going on is that there are just too damn many patches. It's simply impossible to keep up. I get weekly summaries of new vulnerabilities and patches. One alert service listed 19 new patches in a variety of products in the first week of March 2001. That was an average week. Some of the listings affected my network, and many of them did not. Microsoft Outlook alone had over a dozen security patches in the year 2000. I don't know how the average user can possibly install them all; he'd never get anything else done.

Security based on patches is inherently fragile. Any large network is going to have hundreds of vulnerabilities. If there's a vulnerability in your system, you can be attacked successfully and there's nothing you can do about it. Even if you manage to install every patch you know about, what about the vulnerabilities that haven't been patched yet? (That same alert service listed 10 new vulnerabilities for which there are no patches.) Or the vulnerabilities discovered but not reported yet? Or the ones still undiscovered?

Good security is resilient. It's resilient to user errors. It's resilient to network changes. And it's resilient to administrators not installing every patch. Managed Security Monitoring is an important part of that resilience. Monitoring makes a network less dependent on keeping patches up to date; it's a process that provides security even in the face of ever-present vulnerabilities, uninstalled patches, and imperfect products.

In a perfect world, systems would rarely need security patches. The few patches they did need would automatically download, be easy to install, and always work. But we don't live in a perfect world. Network administrators are busy people, and networks are constantly changing. Vigilant monitoring is by no means a panacea, but it is a much more realistic way of providing resilient security.

Security Monitoring



- NOT device monitoring, but another way of thinking about security
- Real-time detection can catch intruders in process
- Real-time response can repel intruders before they do lasting damage
- Good monitoring is resilient
 - Works even if the threat environment changes
 - Makes individual vulnerabilities less relevant
- Good monitoring provides a feedback loop
 - How else do you know how your security is?
- Good monitoring requires human experts
 - Vigilant, adaptive, relentless

28

INTELLIGENT ALERT. INSTANT RESPONSE. IMMEDIATE DEFENSE.

Outsourced Managed Security Monitoring

The key to a successful detection and response system is vigilance: attacks can happen at any time of the day and any day of the year. While it is possible for companies to build detection and response services for their own networks, it's rarely cost-effective. Staffing for security expertise 24 hours a day and 365 days a year requires six full-time employees; more, if you include supervisors and backup personnel with more specialized skills. Even if an organization could find the budget for all of these people, it would be very difficult to hire them in today's job market.


Retaining them would be even harder. Security monitoring is inherently bursty: six weeks of boredom followed by eight hours of panic, then seven weeks of boredom followed by six hours of panic. Attacks against a single organization don't happen often enough to keep a team of this caliber engaged and interested.

In the real world, this kind of expertise is always outsourced. It's the only cost-effective way to satisfy the requirements. I may only need a doctor twice in the coming year, but when I need one I may need him immediately. I may need specialists. Out of a hundred possible specialties, I may need two of them-and I have no idea beforehand which ones. I would never consider hiring a team of doctors to wait around until I happen to get sick. I outsource my medical needs to my clinic, my emergency room, my hospital. Similarly, a network needs to outsource its security monitoring to an MSM service.

Aside from the aggregation of expertise, an MSM service has other economies of scale. It can more easily hire and train its personnel, simply because it needs more of them. And it can build an infrastructure to support them. Vigilant monitoring means keeping up to date on new vulnerabilities, new hacker tools, new security products, and new software releases. An MSM service can spread these costs among all of its customers.


An MSM provider also has a much broader view of the Internet. It can learn from attacks against one customer, and use that knowledge to protect all of its customers. And, from its point of view, attacks are frequent. There's a reason you don't have your own fire department, even if you can afford one. When the fire department comes to your house, you want it to have practiced on the rest of the neighborhood. To an MSM company, network attacks are everyday occurrences; its experts know exactly how to respond to any given attack, because in all likelihood they have already seen the same attack many times before.

In the real world, security is always outsourced. Every building hires another company to put guards in its lobby. Every bank hires another company to drive its money around town. Security is important, complex, and distasteful; it is smarter to outsource than to do it yourself.


 Counterpane™
 Internet Security

Risk-Reduction Means Standardized Security

- Most businesses cannot afford bespoke security
- Standardization leads to quantized risk reduction
- Outsourcing is the only way to make security scale
- Outsourcing is what the insurance industry will want



29
INTELLIGENT ALERT. INSTANT RESPONSE. IMMEDIATE DEFENSE.

Outsourcing Security

If the decision to outsource network security is a difficult one, the decisions of what to outsource and where seem impossible. The stakes are high. On the one hand, the promises of outsourced security seem so attractive: the potential to significantly increase your network's security without hiring half a dozen people or spending a fortune is impossible to ignore. On the other hand, there are the stories of managed security companies going out of business, and bad experiences with outsourcing other areas of IT. It's no wonder that paralysis is the most common reaction to the whole thing.

I believe that network security will continue to be outsourced, because there's no other way to deal with the shortage of skilled computer security experts, the increasing requirements for businesses to open their networks, and the ever-more-dangerous threat environment. For the Internet to succeed as a business tool, security has to scale. Outsourcing is how it will do that.

Over the past few years, we've seen many different companies offering different capabilities under the general category of "managed security services." The field is so confusing that even the industry analysts can't agree on how to categorize them. This company offers to manage your firewall. That company offers periodic vulnerability scans. Another offers to manage your security policy, or monitor your network, or install your IDS, or host your computers. Some of these businesses make sense, and some of them don't. Some will survive, and some of them won't. Knowing which is which is the first step.


What to Outsource

You won't outsource everything, because some things just don't outsource well. Either they're too close to your business, or they're too expensive for an outsourcing company to deliver efficiently, or they simply don't scale well. Knowing the difference is important.

Medical care is a prime example of outsourcing that we can use for comparison. Everyone outsources healthcare; we don't act as our own doctor. More to the point, no one hires a private personal doctor. And we all know what aspects of medical care we like: the ambulance picks up in seconds and rushes us to the hospital, a team of medical experts spares no expense in running tests to figure out what's wrong and in doing whatever it takes to cure us, someone else paying the bill. And we all know what aspects we don't like: ill-equipped and ill-staffed hospitals, HMOs telling us that we can't have that particular test or that a specialist isn't warranted in this case, getting stuck with the bill. When I imagine the wonders of healthcare in the future, I think of automatic monitoring systems that watch our every heartbeat and automatically alert doctors if there's a problem. When I imagine future healthcare horrors, I think of decisions about our health made by accountants and being forced to accept the decisions of others.


The aspects of outsourced healthcare we like involve immediate access to experts. Any medical emergency requires experts, and the faster they can pay attention to us the better off we'll be. The aspects of outsourced healthcare we don't like involve management. Our healthcare is our responsibility, and we don't want someone else making life and death decisions about us.

Outsourcing Security



Counterpane™
Internet Security

- In a world of liabilities, “best practices” become important
- Outsourcing is a way to standardize best practices
- Insurance companies can tie policies to these best practices
- Outsourcing levels the playing field



30

INTELLIGENT ALERT. INSTANT RESPONSE. IMMEDIATE DEFENSE.

Outsourcing Security (cont.)

Network security is no different. Outsource expert assistance: vulnerability scanning, monitoring, consulting, forensics. Don't outsource management.

This truism has been borne out in the industry. Salinas Network Services was the largest firewall management company. For a price, Salinas would manage your firewall. Earlier this year, it disappeared. There just wasn't a business in managing firewalls for other companies. The companies demanded too much individual attention for the money they were willing to pay. Firewall management is just too core to a business. They had no choice but to treat their Salinas contacts as employees.

Pilot Network Services offered secure network management. It's business was to host your computers securely, manage all security devices, test your applications before putting them up on the network...effectively becoming your security management group. They're gone now too - same problem.

Some consulting companies are doing well and some are not. This is more a function of is the quality of the service they offer. Consulting is, and always will be, a profitable business. Outsourcing occasional requirements for expertise transcends any single area.

Outsourced security companies that are doing well are the ones that offer well-defined services that companies need. For example:

- consulting companies like @Stake and Foundstone for expert advice and assistance: strategic security consulting, penetration testing, forensics, etc.
- security VARs for product installation and configuration
- TruSecure for certification and expert assistance
- Counterpane for network security monitoring


In all of these cases, the company buying the outsourced services retains control of its own security. This is important for the company purchasing the services, but it is also important for the vendor. By not demanding a management role, the security companies can offer a useful, effective, and scalable service.

Arguments for Outsourcing

The primary argument for outsourcing is financial: a company can get the security expertise it needs much more cheaply by hiring someone else to provide it. Take monitoring, for example. The key to successful security monitoring is vigilance: attacks can happen at any time of the day and any day of the year. While it is possible for companies to build detection and response services for their own networks, it's rarely cost-effective.

Network Security:
What to Outsource

- Outsource emergency response, but not management
 - Policy development
 - Implementation and installation
 - Security monitoring
 - Vulnerability testing
 - Expert consulting
 - Forensics



31

INTELLIGENT ALERT. INSTANT RESPONSE. IMMEDIATE DEFENSE.

Outsourcing Security (cont.)

Staffing for security expertise 24 hours a day and 365 days a year requires five full-time employees-more, if you include supervisors and backup personnel with specialized skills. Even if an organization could find the budget for all of these people, it would be very difficult to hire them in today's job market.

Retaining them would be even harder. Security monitoring is inherently erratic: six weeks of boredom followed by eight hours of panic, then seven weeks of boredom followed by six hours of panic. Attacks against a single organization don't happen often enough to keep a team of this caliber engaged and interested.

This is why outsourcing is the only cost-effective way to satisfy the requirements. Think about healthcare again. I may only need a doctor twice in the coming year, but when I need one I may need him immediately, and I may need specialists. Out of a hundred possible specialties, I may need two of them -- and I have no idea beforehand which ones. I would never consider hiring a team of doctors to wait around until I happen to get sick. I outsource my medical needs to my clinic, my emergency room, my hospital. Similarly, a network will outsource its security monitoring.

Aside from the aggregation of expertise, an outsource monitoring service has other economies of scale. It can more easily hire and train its personnel, simply because it needs more employees. And it can build an infrastructure to support them. Vigilant monitoring means keeping up to date on new vulnerabilities, new hacker tools, new security products, and new software releases. Outsourced security companies can spread these costs among all of their customers.

An outsource company also has a much broader view of the Internet. It can learn from attacks against one customer, and use that knowledge to protect all of its customers. And, from its point of view, attacks are frequent. No matter how wealthy you are, you do not hire a doctor to sit in your living room waiting for you to get sick. You get better medical care from a doctor that sees patient after patient, learning from each one. To an outsource security company, network attacks are everyday occurrences; its experts know exactly how to respond to any given attack, because in all likelihood they have already seen it many times before.

How to Choose an Outsourcer

This is difficult, because it's hard to tell the difference between good computer security and bad computer security. But by the same token, it's hard to tell the difference between good medical care and bad medical care. If we're not health experts ourselves, we can sometimes be led astray by bad doctors that appear to be good. So, how do you choose a doctor? Or a hospital? I choose one by asking around, getting recommendations, and going with the best I can find. Medical care involves trust; I need to be able to trust my doctor.

Step 4: Rational Prosecution and Education will Lead to Deterrence

- There is no magic security device that prevents murder
- Detection and response work after the fact

- Why are do we all feel safe in this room?

- We feel safe because we live in a lawful society
- We feel safe because criminals are prosecuted
- We feel safe because we deter crime
- We feel safe because we educate our populace

32
INTELLIGENT ALERT. INSTANT RESPONSE. IMMEDIATE DEFENSE.

Outsourcing Security (cont.)

Security outsourcing is no different; you should to choose a company you trust. To determine which one, talk with others in your industry or ask analysts. Go with the industry leader. In both security and medical care, you don't use a little-known maverick unless you're desperate.

Watch companies that have conflicts of interest. Some outsourcers offer security management and monitoring. This worries me. If the outsourcer finds a security problem with my network, will the company tell me or try to fix it quietly? Companies that both sell and manage security products have the same conflict of interest. Consulting companies that offer periodic vulnerability scans, or network monitoring, have a different conflict of interest: they see the managed services as a way to sell consulting services. There's a reason companies hire outside auditors: it keeps everyone honest. I believe that outsourcers that offer combined management/monitoring services will be among the next to disappear. And if a company decides to outsource its security-device management, it is essential that it outsource its monitoring to a different company.

In any outsourcing decision that involves an ongoing relationship, the financial health of the outsourcer is critical. The last thing you want is to embark on a long-term medical treatment plan, only to have the hospital go out of business in the middle. Companies that entrusted their security management to Salinas and Pilot were left stranded when those companies went out of business. Companies that choose the wrong security consulting group will have the same problem. Look for companies that: 1) are leaders in their fields, 2) do one thing well, not those that try to do everything, and 3) have a history.


The Future of Outsourcing

Modern society is built around specialization; more tasks are outsourced today than ever before. We outsource fire and police services, government (that's what a representative democracy is), and food preparation. In general, we outsource things that have one of three characteristics: it's complex, important, or distasteful. In business we outsource tax preparation, payroll, and cleaning services. Outsourcing security is nothing new: all buildings hire another company to put guards in their lobbies, and every bank hires another company to drive its money around town.

Computer security is all three: complex, important, and distasteful. Its distastefulness comes from the difficulty, the drudgery, and the 3:00 a.m. alarms. Its complexity comes out of the intricacies of modern networks, the rate at which threats change and attacks improve, and the ever-evolving network services. Its importance comes from this fact of business today: companies have no choice but to open up their networks to the Internet.

Doctors and hospitals are the only way to get adequate medical care. Similarly, outsourcing is the only way to get adequate security on today's networks.

Problems with Prosecution



- It can be difficult to backtrack attacks
- It can be difficult to prove who the attacker is
 - Separation between the "attacker" and the "attacker's computer"
- It can be difficult to prosecute
 - Complicated and technical evidence
 - General fears lead to irrational prosecution

- International nature of the Internet compounds all of these problems

33INTELLIGENT ALERT. INSTANT RESPONSE. IMMEDIATE DEFENSE.

Counterpane Case Study: The Regence Group

The Regence Group, a major healthcare insurer in the Pacific Northwest, is an affiliation of Blue Cross/Blue Shield insurance companies in Washington, Oregon, Utah, and Idaho. A year and a half ago, the company decided to merge its individual IT security organizations into a single department and to consolidate corporate information security under one person. David MacLeod, Ph.D., CISSP, became The Regence Group's Chief Information Security Officer.

"The first thing I did is look at what we needed to do in order to have an effective information security program. I divided my tasks into two categories: what we should do internally and what we should outsource.

"Managed Security Monitoring rose to the top as something we needed to outsource. I didn't have the staff to man a monitoring station 24x7. I didn't have staff who knew enough about the diverse platforms throughout Regence Group. I didn't have people who could correlate and identify real threats, and weed out anomalies and scans."

Regence Chooses Counterpane

The Regence Group was attracted to Counterpane because they were the only company who offered the kind of comprehensive monitoring that the diverse network of The Regence Group required. "Lots of companies offered to monitor our firewalls and IDSs," said MacLeod, "but only Counterpane would monitor our entire network."

MacLeod was also impressed with Counterpane's endorsement from Lloyd's of London and the way the insurer viewed Counterpane's service offering. "It was clear that the predominant provider in this space was Counterpane."

"Installation was a breeze," recalls MacLeod. "We were up and running within days. And installation continues to be easy. Whenever we have more devices that need monitoring, we just point their audit logs at the Sentry. We call Counterpane and tell them about the new devices, and we're monitored. And if we forget to call, Counterpane calls us as soon as they see the anomalous log data.


"In fact," said MacLeod, "one weekend a corporate Web server went down. Counterpane noticed the problem and called us, before we even knew about it."

Counterpane Provides Enormous Value

To Regence, Counterpane's value can be clearly demonstrated by the number of security incidents they had to deal with. MacLeod compiled statistics to present to his Board of Directors. In the fourth quarter of 2002, Regence's IDSs logged a total of 145,645 incidents. Of those, 12,875 were high-level alerts--serious attempts to compromise the Regence Group network. Another 12,300 were medium-level alerts, and the other 120,000 were low-level alerts. Counterpane's monitoring and filtering of events reduced that number to just over 200 incidents to which Regence had to respond.

(Continued on next page.)

Lawless and the Internet



- The Internet is a lawless society
 - Like a society of warlords
- The rich can afford private security
 - The rest of us do without
- This is why so many are afraid to talk about being attacked

- We need to turn the Internet into a lawful society
- Education is a major part of this
 - We need laws that can be explained

34

INTELLIGENT ALERT. INSTANT RESPONSE. IMMEDIATE DEFENSE.

Regence Case Study (cont.)

MacLeod explained. "I don't have the staff to handle 140,000 alerts. I don't even have the staff to handle 12,000 alerts. What Counterpane does is boil all that noise down to 200 real incidents that my team needs to respond to. Because Counterpane gives us only the information we need when we need it, we can concentrate on the attacks that matter. The value of that kind of service is enormous."

Counterpane's fast response also reduced the effects of those 200 attacks. "The FBI recently reported that roughly 1% of all Internet attacks were successful," said MacLeod. "So if you look at our 140,000 incidents, we should have seen 1400 successful attacks in the fourth quarter. Even if you use the number of high-level alerts--12,000--that should still translate to 120 successful attacks. Because Counterpane was monitoring, and because they were able to notify us so that we could respond to attacks as they were happening, we had only one serious breach during that time period. And that was Nimda."

Even Nimda, which devastated so many other companies, was swiftly contained within Regence's network because of Counterpane. "We got off easy," explains MacLeod. "Even before the Nimda storm crossed our network, Counterpane was on the phone with us explaining what was happening. We started shutting doors and updating our systems immediately, and as a result Nimda compromised only 10% of desktops and 5% of servers. I credit Counterpane."


Counterpane's Global View

Regence entered into a two-year monitoring contract with Counterpane, with the expectation that they would use those two years to develop their own internal monitoring capabilities. Then, they would bring the function back in-house. But the first year proved that Regence could not replicate Counterpane's service and that continued outsourcing of Managed Security Monitoring to Counterpane remains the financially prudent approach.

MacLeod said: "When I look back at Code Red, Code Red II, and Nimda, what stands out in my mind is Counterpane's ability to respond before we were attacked. They regularly call us and tell us how we should respond, based on their experience with customers in earlier time zones."

"We could not possibly replicate Counterpane's service ourselves," continued MacLeod. "We couldn't staff it. And even if we could, we would not get the benefits of Counterpane's global view. They watch security incidents throughout the globe, and we benefit from that."

"Today, we can't do our jobs effectively without Counterpane," said MacLeod.



Conclusions

- The risks will always be with us
 - Just like the real world
- Security products will not solve the problems of Internet security
 - Any more than they solve security problems in the real world
- The best we can do is manage the risk
 - Just like we do in the real world
- The company that manages its risk better will be more profitable

35

INTELLIGENT ALERT. INSTANT RESPONSE. IMMEDIATE DEFENSE.

Counterpane Case Study: Womble Carlyle

Womble Carlyle Sandridge and Rice, PLLC is a 450-attorney law firm, with nine offices in the U.S. According to CIO Sean Scott, security is not only important for the firm's lawyers and clients, but is essential for the future success of the firm.

"As a law firm that serves high tech clients and prominent businesses, keeping up-to-date with technology-particularly security-is not an option, it's a necessity. Womble Carlyle strives to honor client privacy through a secure network. This is one of Womble Carlyle's competitive advantages."

Monitoring as a Solution

Womble Carlyle renewed research into security products last year. In addition to intensifying the sophistication of the firm's security, Womble Carlyle wanted a security audit. Several securities firms recommended Womble Carlyle contact Counterpane about security monitoring, as a way to provide long-term business security.

"By the time we called Counterpane in for a meeting, we were well-educated about the issues and were ready to do business," says Scott. "We realized that these guys knew what they're talking about, more so than some of the other security vendors we talked with. Scott added that he appreciated Counterpane's security expertise, honesty and candor.

Womble Chooses Counterpane

"We met with Vigilinx, RipTech, ISS, Symantec, Guardent, and a several other companies. Scott says. "Counterpane met our needs best. They offered competitive pricing and were vendor neutral." For Womble Carlyle, a vendor-neutral solution was critical.

"Talking to the other vendors gave us a good reality check and starting point," Scott added. "But Counterpane was the company we chose to do business with." Womble Carlyle was impressed not only with Counterpane's technical ability and its monitoring service, but by the fact that they did not also engage in consulting or sell hardware.


"Because Counterpane only does monitoring, we felt confident that they had our best interests in mind," Scott says. "They didn't try to sell us any hardware, software or consulting services. Others wanted to make money by selling us additional hardware and software."

Counterpane: The First Month

Installing the Counterpane service was painless. Scott recalls: "The installation was done in five hours, at most. I was shocked that something like that happened so quickly, with no ill affect on my network." Counterpane started monitoring Womble Carlyle immediately.

(Continued on next page.)

Tweak the Risk Equation Until the CEO Cares



- Today, security is an afterthought:
 - Schedule, performance, and features are what's important
 - "When the deadline approaches, good design is ignored"
 - If a mistake doesn't have immediate consequences, then it's not a mistake
- Liability enforcement changes this
- Until security is monetized, the CEO won't care
- Give the CEO tools to manage his risk

36

INTELLIGENT ALERT. INSTANT RESPONSE. IMMEDIATE DEFENSE.

Womble Case Study (cont.)

Initially, Counterpane analysts called Womble Carlyle often about security events. "My security engineers got a number of calls from Counterpane in the first week and a half," says Scott. "We would apply patches and fix problems in our network."

During the first month, Counterpane and Womble Carlyle worked together to get a more accurate picture of the Womble Carlyle network: to figure out what worked, what didn't, what was good, and what was bad. Security at Womble Carlyle soon improved. Scott says: "IT still get e-mails from Counterpane. But we don't get any more phone calls at night."

Counterpane Helps Make Security Easy

Counterpane is currently monitoring 150 devices on the Womble Carlyle network, including every Internet portal. Every time Womble Carlyle adds a new office, Counterpane monitoring is added. "Counterpane is as integral in our network as Cisco routers," adds Scott.

The most significant benefit that Womble Carlyle can point to is that Counterpane makes network security easier. "Counterpane gives IT the information we need to figure out what's important from a security perspective. Before Counterpane, we didn't always have the most comprehensive information," says Scott. Network security operates much more efficiently now.

Counterpane also gives Scott piece of mind. "If we were to be threatened by some breach, I know we would be alerted by Counterpane and advised on how to react. Womble Carlyle's IT professionals can now focus more of their time on providing attorneys with technical service in other areas.

End Results

Counterpane monitoring gives Scott and his security staff time to improve Womble Carlyle's security. Because they're spending their time more efficiently, they can concentrate on strategically adding security products and policies internally. "There are still fires to put out, but we put them out quickly and effectively," says Scott.

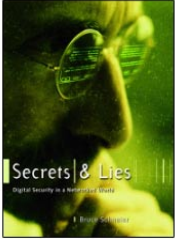
"If we didn't have Counterpane monitoring us, we would be going through the process of adding security internally, while at the same time we would have to monitor the workings inside our network, including improper activities."

Scott concluded: "If you don't keep up with security, you're finished. We all live in the age of data wreckage. People try to raid corporate e-mail, steal client lists, disrupt network operations. We are a law firm committed to protecting our data. To best serve our clients and the firm at large, IT must protect data security. Counterpane's an integral part of that security."

More case studies can be found at:

<http://www.counterpane.com/experiences.html>


Two Useful Resources
from Bruce Schneier



Secrets and Lies:
Digital Security in a Networked World

by Bruce Schneier

John Wiley & Sons, 2000
www.counterpane.com/sandl.html



Crypto-Gram:
Free Monthly Security Newsletter

by Bruce Schneier

www.counterpane.com/crypto-gram.html

37

INTELLIGENT ALERT. INSTANT RESPONSE. IMMEDIATE DEFENSE.

Crypto-Gram, by Bruce Schneier

Crypto-Gram is a free monthly e-mail newsletter that provides news, summaries, analyses, insights, and commentaries on computer security and cryptography. In addition to the regular news reports regarding Internet security and cryptography, here are some of the subjects that appeared in previous issues of Crypto-Gram:

- The Security Patch Treadmill
- Hard-Drive-Embedded Copy Protection
- A Cyber UL?
- Voting and Technology
- Why Digital Signatures Are Not Signatures
- AES
- Full Disclosure and the Window of Exposure
- Bluetooth
- Full Disclosure and the CIA
- Microsoft SOAP
- Computer Security: Will We Ever Learn?
- ILOVEYOU Virus
- UCITA
- Software Complexity and Security
- Publicizing Vulnerabilities
- Insurance and the Future of Network Security
- Internet Voting vs. Large-Value e-Commerce
- Code Signing in Microsoft Windows
- Digital Safe-Deposit Boxes
- Microsoft Hack (the Company, not a Product)
- Semantic Attacks: The Third Wave of Network Attacks
- PGP Vulnerability
- Microsoft Vulnerabilities, Publicity, and Virus-Based Fixes
- Security Risks of Unicode
- The Data Encryption Standard (DES)
- Trusted Client Software
- Microsoft Active Setup “Backdoor”
- Kerberos and Windows 2000
- Distributed Denial-of-Service Attacks
- Cookies

To subscribe to Crypto-Gram, visit <http://www.counterpane.com/crypto-gram.html>

Or send a blank message to crypto-gram-subscribe@chaparraltree.com

Back issues of Crypto-Gram are available at <http://www.counterpane.com>

Counterpane™ is a registered trademark of Counterpane.
© Copyright Counterpane Internet Security. All rights reserved.



INTELLIGENT ALERT. INSTANT RESPONSE. IMMEDIATE DEFENSE.

Counterpane™
Internet Security

Counterpane Internet Security, Inc.

19050 Pruneridge Ave.
Cupertino, CA 95014

1.888.710.8175

www.counterpane.com

Notes: