Managed Security Monitoring: Closing the Window of Exposure

Bruce Schneier

Introduction

The Internet is critical to business. Companies have no choice but to connect their internal networks to the rest of the world—to link with customers, suppliers, partners, and their own employees. But with that connection comes new threats: malicious hackers, criminals, industrial spies. These attackers regularly steal corporate assets and intellectual property, cause service breaks and system failures, mar corporate brands, and scare customers.

This document discusses how to deal with Internet security as a business problem: how to handle it in the same way as the rest of the corporate threats. Unless companies can successfully deal with those threats, they will never be able to unlock the full potential of the Internet. Companies that cannot adequately protect their customers and shareholders will fail, and those that can will succeed.

Traditional approaches to computer security have failed. Despite decades of research, and hundreds of available security products, the Internet has steadily become more dangerous, not less. In this paper I argue that the historical security model of *threat avoidance* is flawed, and that it should be abandoned in favor of a more businesslike *risk management* model. Traditional security products—largely preventive in nature—embody the threat avoidance paradigm: either they successfully repel attackers, or they fail. The unfortunate reality is that every security product ever sold has, on occasion, failed.

A security solution based on risk management encompasses several strategies. First, some risk is accepted as a cost of doing business. Second, some risk is reduced through technical and/or procedural means. And third, some risk is transferred, through contracts or insurance.

Most people concentrate on the second approach and attempt to solve the risk through the purchase of security equipment. However, technical risk reduction cannot be achieved this way; newly discovered attacks, proliferation of attack tools, and flaws in the products themselves all result in a private network becoming vulnerable at random (and increasingly frequent) intervals for random amounts of time. The only way to stay ahead of these vulnerabilities is through detection and response: active network monitoring. And the most cost-effective way to do that is through a *Managed Security Monitoring* service. The reality of today's Internet will make Managed Security Monitoring a mandatory security requirement.

Security and Risk Management

Ask any network administrator what he needs security for, and he'll describe the threats: Web site defacements, corruption and loss of data due to network penetrations, denial of service attacks, viruses, loss of good name and reputation. The list seems endless, and an endless slew of press articles prove that the threats are real.

Counterpane
Internet Security

Most computer security is sold as a prophylactic: encryption prevents eavesdropping, firewalls prevent unauthorized network access, PKI prevents impersonations. To the world at large, this is a strange marketing strategy. A door lock is never sold with the slogan: "This lock prevents burglaries." But computer-security products are sold that way all the time.

There exists no computer-security product—or even a suite of products—that acts as magical security dust, imbuing a network with the property of "secure." Security products are risk management tools, some more effective than others, that reduce the risk of financial loss due to network attacks. These tools should be deployed when the savings due to risk reduction are worth the investment in the tool. Otherwise, it is cheaper to accept or insure the risk than it is to deploy the tool.

For example, it makes no sense to purchase a $10,000 safe to secure a $1000 diamond. Even if you could buy a $500 safe, a $300 insurance policy would be a smarter purchase. But if you could buy a $100 safe *and* a $100 insurance policy that requires the safe, that would be the most cost-effective solution of all.

This is important. To a CEO, what is important is risk management. The CEO doesn't care if risk is reduced through technical means, operational procedures, or insurance. It's all the same. Blindly adding technologies to avoid the threats is not smart business; carefully adding technologies that provide for cost-effective risk reduction is.

The Window of Exposure

A company's computer network could be likened to a building, and the windows and doors to the Internet access points. Continuing this analogy, strong door and window locks could help keep out intruders, and office-door locks and locked filing cabinets could help prevent "insider" attacks. Of course these preventive security measures are not enough, and a well-protected building also has alarms: alarms on the doors and windows, and maybe motion sensors and pressure plates in critical areas inside.

The Internet is much more complicated than a building, and constantly changing. Every day there are new vulnerabilities discovered, new attack tools written, and new legitimate services offered. Whenever a new way to attack a house—or a network—is discovered, there exists a *window of exposure* until that attack method is prevented.

Rarely is a totally new technique for picking door locks invented, one that renders existing lock technology obsolete. Imagine for a moment it has. At the point of invention, there exists a window of exposure for all buildings that have these sorts of locks. As long as no one knows about the lockpicking technique, the window is small. As criminals learn about the technique, the window grows in size. If the technique is published and every criminal learns about it, the window is very wide. At this point, there is nothing anyone can do about the problem; the locks are vulnerable. Only after a lock manufacturer designs and markets a lock that is resistant to this technique can people start to install the new locks. The window closes slowly but, since some buildings will never get these new locks, never completely.

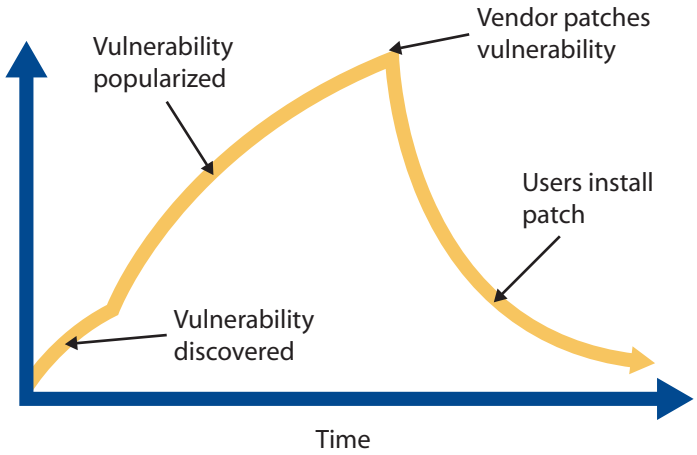

Figure 1

This is what happens daily on the Internet. And because the Internet is much more dynamic and unstable than a building, the repercussions are much worse. Someone discovers a new attack methodology that renders some networks vulnerable to attack. The exposure grows as more people learn about this vulnerability. Sometimes the window of exposure grows very slowly: there are attacks that are known by a few academics and no one else. Sometimes the window grows very quickly: some hacker writes an exploit that takes advantage of the vulnerability and distributes it free on the Internet. Sometimes the software vendor patches the vulnerable software quickly, and sometimes the vendor takes months or years. And some network administrators install patches quickly and religiously, while others never do.

To take just one example: In July 1998, a bug report was published about Microsoft's Internet Information Server. No one knows how long—if at all—it was used to attack systems before then. Microsoft issued a patch fixing the vulnerability shortly after it was published. In July 1999, Microsoft issued a second warning of the vulnerability and the need to install the patch. Even so, in January 2000 the vulnerability was used to steal credit card numbers from several high-profile Web sites.

But that's just one isolated vulnerability out of the dozens that were discovered in networked products that week, and are discovered every week. There isn't a single window of exposure, but rather the superposition of many windows of exposure. The result is a constant state of exposure in corporate networks: there is *always* a way for a determined hacker to break into your corporate network. As an example, look at *eWeek* magazine's "Openhack" project, where they built a test network and offered prizes for successful attacks. As it said in an article about the project, "eWeek Labs' Openhack.com e-business site was built from the ground up with security in mind, and was co-designed and co-maintained by security company Guardent Inc. Yet Openhack was cracked—by two different people in less than one month." [1] And these are defenders who invested a lot of effort to do it right.

Reducing the Window of Exposure

The obvious defense is to make these windows of exposure as small as possible. We have two options. We can try to limit the number of people who know about the attacks, thereby reducing the window in space. And we can try to increase the speed at which vendors patch software to eliminate vulnerabilities (and how fast those patches are installed), thereby reducing the window in time. Both are being tried today.

Reducing the Window in Space

Some security experts advocate limiting the amount of vulnerability information available to the public. The idea is that the less attackers know about attack methodologies, and the harder it is for them to get their hands on attack tools, the safer networks become.

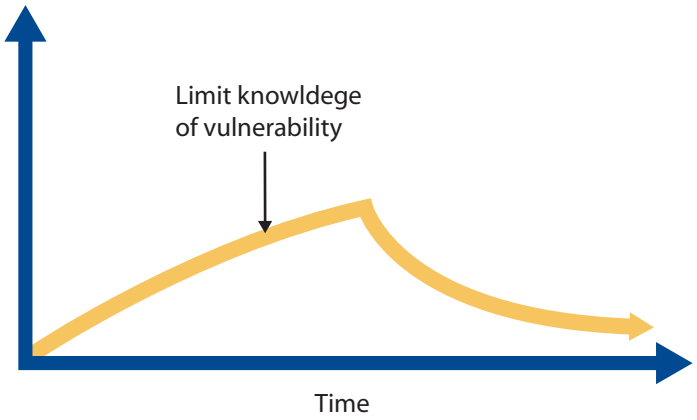

Figure 2

[1] *"Openhack: Lessons Learned," eWeek, 7 Aug 2000, p. 1.*

This would work in theory, but unfortunately it is impossible to enforce in practice. There is a continuous stream of research in security vulnerabilities, and most of these result in public announcements. Hackers write new attack exploits all the time, and they quickly end up in the hands of malicious attackers. There have been some isolated incidences of a researcher deliberately not publishing a vulnerability he discovered, but public dissemination of vulnerability information is the norm…because it is the best way to improve security.

Reducing the Window in Time

The majority of efforts to reduce the window of exposure centers around the time axis. The window remains open until the vendor patches the vulnerability *and* the network administrator installs the patches. Ideally, the vendor will distribute the patch before any exploits are written.

This also works a lot better in theory than in practice. There are many instances of security conscious vendors publishing patches in a timely fashion. But there are just as many examples of security vendors ignoring problems, and of network administrators not bothering to install existing patches. A series of credit card thefts in early 2000 was facilitated by a vulnerability in Microsoft IIS that was discovered, and a patch was released for, a year and a half earlier.

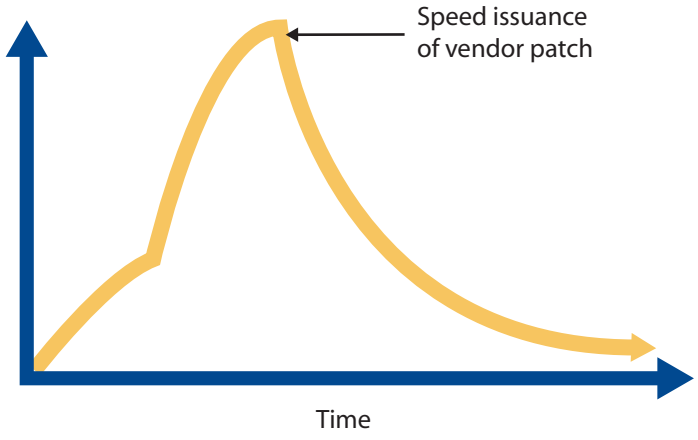

Figure 3

Closing the Window

The problem is that for the most part, the size and shape of the window of exposure is not under the control of network administrators: it's under the control of hackers and software vendors. All the network administrators can do is install patches when they're available.

The lockpicking example was a naive one, because it assumed that the door lock was the only thing standing between an attacker and access. This is often not true; secure buildings not only have door and window locks, but alarm systems. And it is an alarm system that is the key to closing the window of exposure. Not reducing it, but closing it.

Good security is based not only on preventive countermeasures—locks in this example—but also on detection and response. Preventive countermeasures provide defense in two ways: they provide a barrier to overcome, and they force the attacker to spend time overcoming the barrier. Good security is based on prevention, detection, and response. An effective alarm system coupled with a fast response system can repel the attacker before he has done his damage, and hopefully before he has even overcome the preventive barrier.

To a building with an alarm system, a new lockpicking technique is no longer a critical security concern. Of course the security manager should upgrade his locks to the new resistant technology, but until he can, he has his alarm system to fall back on. Even if an attacker successfully picked the lock, he would still be detected and he would still be repelled (or caught).

The same kind of thinking can apply to computer networks. Attackers can be detected inside corporate networks, regardless of which vulnerability they used to enter. New vulnerabilities can be detected *before* the security products are patched to resist them. In many cases, the window of exposure can be closed.

Of course there are always caveats. Attackers will attempt to bypass the alarm system. This is always possible, but made difficult because alarms provide defense in depth. It is not enough to bypass the door lock and the door alarm; an attacker has to bypass all internal alarms. On a computer network, alarms can be in every server, every router, every network software package. Bypassing them all is truly a Herculean task.

And some attacks don't trip alarms. A trusted employee can steal from a company without setting off any alarms. That same employee could disable the alarm system. Still, an alarm system makes it much harder to penetrate a building undetected. And in any attack that takes time, an alarm system gives the police the ability to respond in time to stop a crime. This is why banks have alarm systems in addition to blast-proof vaults.

An Internet security service based on detection and response—an Internet alarm system—is a fundamentally different way to manage the risk of network attack. It is so effective that the insurance company Lloyd's of London based a new anti-hacking insurance policy around Counterpane's Managed Security Monitoring (MSM) service. A customer of Counterpane can get anti-hacking insurance from Lloyd's. There's no security assessment required. There are no specific security products required. There are no complicated rules about installing vendor patches or adhering to specific policies. If you are a Counterpane customer, you can purchase the insurance. Lloyd's will even cover attacks that make use of vulnerabilities for which patches exist, even if the customer didn't install the patch. Why is Lloyd's offering this insurance policy without complicated assessments, procedures, and caveats? Because Counterpane's service actually closes the window of exposure for networks it monitors; it doesn't just reduce it. Choosing one product over another doesn't make much of a risk management difference, but adding the Counterpane service does.

Detection and Response in Computer Networks

The details are considerably more complicated. Detection and response is easy with a building alarm, but considerably more complicated with a computer network. Knowing how a network is being attacked and by what sort of adversary, or even knowing *that* a network is under attack, can be difficult to ascertain. Determining an appropriate response can be a complicated and subtle decision process. And knowing how to effectively respond requires considerable expertise and experience.

Counterpane Internet Security, Inc. is in the business of Managed Security Monitoring: network security detection and response. The service works much like a traditional physical alarm company, with the addition of vigilant human response, enabling it to work in the complicated arena of the Internet.

Central to this service are security analysts. These highly trained security professionals are the front line defenders. They interpret the alarms from the customer networks, determine the appropriate response, and contact the customer. Assisting these analysts is the SOCRATES information processing system, which includes detailed information about attacks and attack tools, possible responses, and customer network configurations and operations. Feeding SOCRATES are the Counterpane Sentries, the alarm system residing on the customer network. The Sentry monitors all critical parts of the customer networks—the routers, the servers, as well as network security devices like firewalls and IDSs.

Unlike traditional building alarm systems, new attacks and new attack tools appear every day. Counterpane's Network Intelligence group is tasked with staying on top of these developments. This group monitors a variety of information channels—vendor patches and security bulletins, open-source Web sites and newsgroups, and underground hacker bulletin boards and chat areas—to make sure that SOCRATES always has the most current information and the Sentries are always able to detect the broadest possible attacks.

Outsourcing Detection and Response

The key to a successful detection and response system is vigilance: attacks can happen at all times of the day and all days of the year. While it is possible for companies to build a detection and response service for their own

networks, it is rarely cost-effective for them to do so. Staffing for security expertise 24 hours a day and 365 days a year requires five full-time employees, and more if you include supervisors and backup personnel for critical tasks. Even if an organization could find the budget for these people, it would be very difficult to find them on the job market. Retaining them would be even harder; attacks against a single organization don't happen often enough to keep a team of this caliber engaged and interested. Staffing an intelligence organization, necessary to keep the system up to date, is just as difficult. A Managed Security Monitoring provider can spread these costs among all of its customers, making the type of detection and response necessary to reduce risks attainable and affordable.

Aside from the aggregation of expertise, a single outsourcing organization has other economies of scale as well. It has a much larger network visibility. It can learn from attacks against one customer, and use that knowledge to protect all of its customers. And attacks would be commonplace. To a team just protecting its own company, an attack would be a rare event. To a Managed Security Monitoring company, attacks would be everyday occurrences; the experts would know exactly how to respond to any given attack, because in all likelihood they would 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. Whenever you see something with those characteristics, it is smarter to outsource than to do it yourself.

Conclusion

Security is a process, not a product. Traditional preventive security products go a long way to securing computer networks, but they can never close the window of exposure. All existing networks are vulnerable to attack. Looking at the problem as one of risk management, detection and response are far more effective security tools than prevention can ever be. And Managed Security Monitoring is the most cost-effective way, as well as *the* most effective way, to reduce the risk of financial losses due to network attacks.

Counterpane
Internet Security

3031 Tisch Way, 100 Plaza East
San Jose, CA. 95128

www.counterpane.com

P: 408.260.7500
F: 408.556.0889