


5 strategies for proactively embracing failure

by Steve Shillingford



As information security professionals, we constantly ask, “Are we doing enough?” To which, the answer is usually a resounding “No.” So, we embark on an often endless cycle of product and process evaluation, purchase and implementation only to end up plagued by our initial insecurity - that we’re still not doing enough, that we’re still not secure.

This paranoia is driven by a multitude of factors. Beyond the desire to succeed in our professional roles, consider the influence of highly publicized breaches, the endless succession of "next generation" security toolsets, the barrage of threats including the next "zero-day" exploit, and the evolving government regulations meant to ensure information security in the first place. Collectively, these factors breed an industry-wide fear of catastrophic system failure. Naturally, we are inclined to embody this fear by building systems aimed solely at preventing it.

This logic is flawed. The preventative security solutions that we employ today only protect against “known threats” - or those that have already been identified by our existing info security systems. Meanwhile, attackers continue to persevere, relentlessly. Admitting susceptibility to these security loopholes, or “unknown threats”, that facilitate failure, may prove more useful than focusing so much on known

threats. Viewed in a broader light of environmental adaption in complex systems, we begin to recognize that planning for failure is not only important, it's fundamental in addressing our objective of comprehensive security.

Armed with the handful of strategic initiatives outlined here, IT security professionals can begin accounting for the inevitability of failure, improving their overall security posture.

Understanding failure

Before we begin planning for failure, we must learn to accept that information security is an imperfect entity; that, despite our best efforts, any defense measure we employ will fail; and that if this failure is unavoidable, it must be factored into the information security process. For the same reasons banks use video surveillance while simultaneously deploying prevention measures (guards, alarms, etc.), IT organizations need to embrace the notion that

prevention is not a credible stand-alone measure. Educating ourselves, our teams, and our departments about the various theories and studies that lend credibility to this contrarian thinking is imperative. There are many complex systems outside IT where failure preparation is standard practice. Why should the network be viewed any differently?

We should first consider the factors that drive innovation in the information security products landscape. The rules that govern economics and natural selection help eliminate the inferior, unaffordable and ineffectual offerings. We are able to select from a range of best-of-breed defensive solutions that, upon deployment, instill a sense of reasonable confidence in our information security systems.

This confidence sets the stage for a less-than-desirable consequence. The more fit a security product is perceived to be, the more likely it is to recklessly reassure us that we are secure. And while we're caught up in a fleeting sense of security, the same market dynamics that drive product innovation are fueling the evolution of pervasive and agile threats.

To overcome their adversaries, attackers have become increasingly covert, both in the means through which they're infiltrating our systems and their intended end result. Typically, an information security system's ability to detect and protect against these attacks depends on deterministic strategies, where products are configured to address only known threats or events. If attacker's operations are unknown, and therefore go undetected, preventive countermeasures cannot be adapted to thwart their attacks.

In an attempt to counter such deficits, hybrid products have appeared - those that include deterministic and heuristic strategies, such as behavioral- or anomaly-based methods. Currently, the time it takes these products to help us accurately identify and resolve malicious network activity is insufficient in containing the damage caused an attack. But when these hybrid technologies become pervasive - as the competitive product ecosystem suggests they will - attackers will adapt and evade, becoming simultaneously able to impersonate "normal" behavior and remain relatively undetectable.

Van Valen's Red Queen hypothesis helps explain information security product developers' and hackers' tendency to one-up each other. It suggests the balance between competing species evolves dynamically - a state where adaptive improvement is always possible for both species so they continually evolve in relationship to one another and keep up with the evolutionary improvement of their counterparts. In the context of information security, product vendors and attackers continually compete for survival, each incrementally trumping each other's more advantageous attributes without driving their competition into extinction. So as quickly as a system can be updated to protect against an identified threat, an unknown, more adaptable threat can compromise the systems' effectiveness.

In this way, active countermeasures to known threats only provide the illusion of control. Bruce Schneier explains this well in his book *Beyond Fear*, when he outlines "security theater." According to Schneier, security theater describes countermeasure solutions that provide the feeling of improved security while doing little or nothing to actually ensure safety. This is not to say that a firewall doesn't, in fact, protect against the known threats for which it's configured against; instead, it draws attention to our tendency as information security professionals to be blinded with confidence in our defensive efforts and ignore the potential vulnerabilities of our current systems to unknown threats.

Industry reports lend additional credibility to the insidious and pervasive nature of network attackers and can often provide clues to what isn't working on an industry-wide level. A June 2008 security survey conducted by InformationWeek reported that while 95 percent of the organizations surveyed had security budgets that were the same or increased from 2007, 66 percent of them suspected their vulnerability to breaches to be the same or worse as they were in 2007. The same survey participants suggested that firewalls, antivirus tools, encryption and VPNs were only effective about two-thirds of the time, providing ample opportunity for successful attacks.

Or take, for example, the recent findings about a non-dictionary attack on the popular wireless encryption method Wi-Fi Protected

Access (WPA), which were presented at the PacSec Tokyo 2008 conference by academic researchers Erik Tews and Martin Beck. In their paper, entitled "Practical Attacks against WEP and WPA," they report finding a hole in part of 802.11i that forms the basis of WPA. Leveraging this weakness, they were able to break the temporary Key Integrity Protocol (TKIP) in under 15 minutes.

These findings carry implications for information security professionals in enterprises worldwide. Sure, we can upgrade to WPA2 if we haven't already, but how long until this encryption method is cracked?

Findings from both industry and academic research encourage us to more closely scrutinize our own security systems and processes. Combined with exposure to theories, such as Red Queen and Schneier's security theater, we begin to understand the ever-evolving nature of attackers and their ability to evade the security products we employ to detect and protect against them. As a group we should acknowledge the imperfect nature of our information security systems and processes. Equipped with this new perspective, we can more effectively address questions like those around securing wireless networks - we can begin accounting for inevitable failure as a fundamental tenet of design.

As a group we should acknowledge the imperfect nature of our information security systems and processes.

Risk mitigation

Evaluating security infrastructure in accordance with risk management theory provides a valuable framework with which to start accounting for system and process failure. As long as we are trying to protect assets, we must accept that some combination of existing threats (or attacks from which we are trying to protect our assets) and vulnerabilities (or the way in which an attacker prevails), can put those assets at risk. Identifying our organizations' assets, calculating their individual risk and employing a risk management model can help us determine our organizations' specific threshold for risk.

ISO International Standard ISO/IEC 15408-1:2005, also known as Common Criteria Part 1, offers a straightforward formula for calculating the relationship among variables, such as threats and vulnerabilities that account for an assets' quantitative risk (standards.iso.org/ittf/PubliclyAvailableStandards/). Similarly, the NIST Special Publication 800-30 provides a simple decision chart for determining an organization's acceptability of risk (csrc.nist.gov/publications/nistpubs/).

A recent adaptation of these basic risk mitigation theories—issued by the American National Standards Institute (ANSI) and the Internet Security Alliance in a guide called

"The Financial Impact of Cyber Risk: 50 Questions Every CFO Should Ask" (webstore.ansi.org/cybersecurity.aspx) - takes a more holistic approach. The guide suggests that organizations calculate network security risks for specific attacks or events by first asking questions of every department or group within the organization that might be affected. This comprehensive pooling of data seeks to ensure better accuracy in determining the organizations' potential risks, and the costs associated with them, because it involves everyone who might be affected by a security breach.

This equation, supported by the basic theories outlined in the Common Criteria Part 1 and NIST Special Publication 800-30, can help information security professionals in determining what risk management actions, if any, should be executed.

The formulas inform allocation of resources - essentially, helping us assess the type of protection we can afford in terms of time, money, energy and space consumption, human resources, tolerability and sustainability. They allow us to arrive at an acceptable level the cost associated with our organizations' specific risks while also directing us to where vulnerabilities persist.

Devalue data

We should not be surprised when we're faced with reports about the information security risks of Internet communication, such as VoIP, SMS-linked micro-blogs or social networks. We should anticipate them because, by nature, network information is highly vulnerable.

Consider the characteristics of posting data on the Internet. Technology makes instant communication simple. This communication can be private, and often private communication is centered around sensitive matters. Communication can also be public, and sometimes public communications can reveal too much information. As such, Internet communication is becoming increasingly transparent. Because we're inclined to capitalize on the simplicity it provides, what, as information security professionals, can we do to ensure that sensitive enterprise data remains relatively private?

We can turn to Red Team exercises—or security practice drills issued by the US government. An example is the confidential report that recently spurred an onslaught of “Potential for Terrorist Use of Twitter” stories in the media. However, we should proceed with caution when directing our attention to such exercises, as we don't want to adopt their alarmist perspective.

We must also avoid making it easy to hack a system. This may seem obvious, but it's a surprisingly common oversight. In the recent Sarah Palin hack, the hacker simply reset Palin's password using her birth date, ZIP code and information about where she met her spouse—all information available through a simple Google search. It seems someone would have thought to adjust the password settings on her personal email accounts or take them down entirely.

Any amount of time wasted in implementing a patch widens the window of time during which an organizations' data is vulnerable to a known threat, and system failure.

We must heed these warnings. More specifically, we should use communication modes other than the Internet when transmitting sensitive enterprise data. To help ensure that all employees take such precautions, not just those of us in IT, we can block users within our network from accessing non-corporate email, VoIP, micro-blogging and social network accounts. We can also provide warnings and education that deters them from using personal accounts to send company documents and information when working outside of the enterprise network.

Finally, and most importantly, we can try to devalue data whenever possible. We can use full-disk and database encryption so that when a loss or breach occurs, the thief finds the data inaccessible or, at least, very expensive. We can use unique passwords with the help of a password manager so that if one password is compromised, others aren't. We can use “one-time” data instances such as one-time passwords or one-time credit card numbers.

Accounting for known threats

We must ensure that failure doesn't occur because of a known issue. With risk mitigation theories, we can more accurately determine which information security product investments will lessen risks associated with known threats and events. Aggressively applying more- or less-comprehensive detection and prevention solutions based on these determinations is imperative.

Additionally, as the information security ecosystem evolves, more known threats are revealed. These threats are often brought to our attention by the security products' vendors in the form of a patch or signature file. Though it may seem obvious, staying abreast of these updates within our existing infrastructures is of equivocal, if not greater, importance to investing in new products or upgrades. Any amount of time wasted in implementing a patch widens the window of time during which an organizations' data is vulnerable to a known threat, and system failure.

What happens when a known threat infiltrates our systems during this window of vulnerability? Or, if it attacks before the patch itself is issued, when the threat is still unknown?

Though vendors may tout the idea of “zero-day threat detection,” more often than not evidence suggests the contrary. Rather, it points to an undefined time period before a patch was issued and implemented when vulnerable systems were successfully attacked and exploited.

The controversy surrounding the Microsoft MS08-067 emergency patch is an example of this. The patch was issued on October 23, 2008 to remedy the Windows RPC exploit. Yet, Trojans capitalizing on the flaw were identified the day following its release. Further analysis of these strains suggested that they may have been in circulation before the patch was issued, perhaps as early as September 29. The concept of “zero day” goes out the window, but the potential exploitations or events that occurred because of the vulnerability remain.

Incorporating an incident response plan into our information security practices and processes provides us with the ability to better identify the cause and extent of a breach.

Negative day threat detection and network forensics

Incorporating an incident response plan into our information security practices and processes provides us with the ability to better identify the cause and extent of a breach. With a plan in place, we can account for the fallibility of patches and defensive solutions as well as the pervasive nature of system threats and vulnerabilities.

A well-executed incident response plan incorporates a number of variables. Above all, it must contain the direct damage caused by an attack. It must provide the tools for a methodical and timely response, curbing the indirect damage, such as negative publicity, reduced customer confidence, or legal repercussions. If set up properly, it can also identify and resolve the root causes of an incident so repeat occurrences can be avoided. The hallmark of such a plan is network forensics technology - or more specifically, traffic capture, regeneration and search solutions.

Capitalizing on the advancements in data storage, which increase space at lower costs, these solutions record all data crossing a network and store it for later recall and analysis. This complete record of network traffic provides context to alerts or events. Once we identify a threat, we can navigate through traffic history and search evidence surrounding the actual event, not just superficial metadata such as log files and header information. We can use this evidence to view and replay, with

full fidelity, the events that predated classification of the threat.

In the case of vendor-issued patches, network forensics technologies offer “negative day threat detection.” That is, the patch serves as an incident or notification of a previously unknown threat. And we can go back, even weeks prior to the issuance of patch, and use the published threat patterns to search for instances of the offensive malware that might have crossed the network since the first reports of the incident.

But alert mechanisms that rely on pre-defined signatures, patterns or data or those that are identified by security vendors or researchers are hardly infallible. We can also leverage capture technology for surveillance - a process of continuously capturing and monitoring traffic for detection of any atypical activity or anomaly. Specifically in high-risk or vulnerable areas of a network, monitoring traffic records can help us proactively distinguish between legitimate alerts and false positives. They can help us uncover previously undetected, or unknown, breaches.

When prevention fails, detection is key. Network forensics tools equip us better in efficiently realizing known and unknown breaches. We can more effectively stem further loss or future loss of sensitive data and update existing controls to avoid repeat attacks. These tools provide necessary fortitude to any effective incident response plan and help us account for failure.

Complex systems reside in a state of equilibrium where events have individual and aggregate impacts. For example, why is it difficult to immunize against certain viruses? Because in many cases, these viruses evolve and evade the cocktails of drugs that seek to prevent them from successfully attacking healthy cells. This characteristic holds true for any complex system with multiple inputs and outputs.

Another system of moderate-to-sufficient complexity that's worth examining is security in a bank. A bank has a diverse collection of defenses to protect against robbery, including a vault, time-release locks, bulletproof glass and security guards. But it also employs a security measure that accounts for the failure of those defensive solutions: surveillance cameras. If the defensive measures fail to detect and prevent a robbery until after the money and robbers are long gone, authorities would

not turn to the security guard for eye-witness testimony. They'd rely more heavily on the forensic record of evidence provided by the cameras.

Why then, as information security professionals, do we think our organizations are any different than the virus or the bank? We should know better than to believe we're safe from network failure of some undetermined variety and magnitude. We must take into consideration education about the pervasiveness of threats, the specific risks that our organizations face in the wake of these threats, the importance of devaluing data and the role played by network forensics technologies. Once we realize the impediments to adopting these strategies are non-existent, we can move to implement them throughout our information security systems and processes - ultimately accounting for network failure.

Steve Shillingford has more than 15 years of experience in sales, operations and management in technology companies. He joined Solera Networks (www.soleranetworks.com) in early 2007 from Oracle Corporation, where he was responsible for some of the largest deals in the company during his tenure, all in the Rocky Mountain region. Steve was named top salesperson within Oracle in 2005 as a result of this success. Prior to joining Oracle in 2000, Steve had held several sales and operational management positions at Novell over the preceding seven years. Steve holds a B.S. with honors in Psychology from Brigham Young University.

