1

Cybersecurity

Understanding Vulnerability

The hack of Ashley Madison's website put an indelible human face on the vulnerabilities of cybersecurity in a way few other incidents could. Ashley Madison's entire business model of infidelity depended on discretion and anonymity to protect users from embarrassing disclosure. Public exposure of Madison's member names, email addresses, and other personal information has served as a case study illustrating the chasm between the expectation of security and real-life measures taken to ensure data are not easily exposed.

Ashley Madison did use a bcrypt algorithm in PHP, according to Wire.com. Madison's use of encrypted passwords surpassed the security of other recently hacked websites even though the hackers were able to crack the hash to discover the account holder's real password. Adding insult to injury, members were charged $19 for additional security to delete all personal information from the site; however, Madison failed to delete the data completely.

The Impact Team, the name used by the hackers, targeted Ashley Madison on moral grounds, implicating the firm's business model for facilitating adultery. The Impact Team allegedly disclosed Ashley Madison created fake accounts using female bots to engage male customers, artificially boosting overall membership and growth numbers prior to a planned initial public offering. The precise methods used to hack Madison are not clear; however, by deconstructing the hack we can see the most likely weaknesses.

Very large amounts of data were publicly released, suggesting a breach of administrator access to the database or an insider accomplice acted as a whistleblower through a Gray Hat hack. Enterprise database infrastructure is a common cause of an overwhelming number of hacker attacks. The hackers raised the bar by posting Madison account data including personal identifiable information on a public site, allowing the media to implicate public figures, employees of law enforcement and government agencies, and clergy for extra measure.

One of the most frequently exploited vulnerabilities involves access to a database directly from the Internet or website using a cross-site scripting (XSS) attack. XSS is a computer security vulnerability found in web applications. XSS allows hackers to download computer code called "script" into customer-facing websites, bypassing access controls. XSS carried out on websites accounted for roughly 84% of all security vulnerabilities documented by Symantec as of 2007.* Other minor control failures, such as human error, may also be factors contributing to the severity of a hack. The exact details of the attack, like for most attacks, are not exhaustive, though nevertheless symptomatic of a broader narrative of failures at Madison that existed long before the actual exposure of data occurred. Ashley Madison was heading into a storm of its own making long before The Impact Team exposed the firm.

According to McKinsey Research, 2010 was a transitional year for private equity investments. Private equity fund assets in the United States and Canada declined almost 90% from 2007 to 2009, from a peak of $506 billion to $64 billion. By 2010, markets began to stabilize, with stock prices rising sharply and oil climbing from $35 a barrel to more than $100. In addition, the Federal Reserves' Quantitative Easing (QE) program had been in place for two years, providing low-cost liquidity to institutional investors. Basically, there was lots of cash available for promising new ventures and Avid Life Media, the parent firm of Ashley Madison, wanted its share of cash to raise capital to expand.

In January 2010, Avid Life Media's CEO Noel Biderman managed a profitable portfolio of media assets at a time when Canada's private equity business was very active. Mr. Biderman actively sought $60 million in venture capital money to finance the acquisition of a much larger firm, Moxey Media, with the promise of an exit for institutional investors through a reverse takeover of an existing shell company on the Canadian exchange. Biderman managed two similar

* http://eval.symantec.com/mktginfo/enterprise/white_papers/b-whitepaper_exec_summary_internet_security_threat_report_xiii_04-2008.en-us.pdf

sites, CougarLife.com, catering to well-off older women seeking relationships with younger men, and EstablishedMen.com, which aimed to connect young women in search of wealthy men willing to subsidize their lives in exchange for intimate relations.

Morality, or to be more exact, the facilitation of immoral behavior, proved to be too high a hurdle for private equity investors in Canada. Biderman's Canadian bankers simply did not want to be associated with a "sinful venture" no matter how lucrative it might become. Fast-forward five years to spring 2015, when Avid Life Media made efforts to line up investors for an IPO in Europe. It was rumored the reason Mr. Biderman picked London to launch an IPO was because "Europe is the only region where we have a real chance of doing an IPO" presumably because of its more liberal attitudes toward adultery. Meetings with wealthy investors included claims of 36 million users, strong financials, and female-to-male ratios of 50:50 as part of the sales pitch to investors.

In April 2015, Avid Life Media prematurely informed Bloomberg of its plan to raise $200 million in private equity with a $1 billion valuation. *Fortune* magazine and, later Reuters, picked up the story but learned later that investors had pressured the firm to improve liquidity before funding commitments were finalized. Avid Life's growth rates were difficult to reconcile, posting earnings before interest, taxes, depreciation, and amortization (EBITDA) of $8 million in 2009 with $30 million in sales and by 2014 sales of $115 million reported to Bloomberg, an almost fourfold increase.

By July 2015, news broke that The Impact Team had publicly posted Ashley Madison's customer records and personal details and the fictitious account claims. Throughout the month of August 2015, The Impact Team posted several data dumps on the open web prompting public recrimination and eventually the resignation of CEO Noel Biderman, scuttling the planned IPO.

Ashley Madison's ambitions for growth at all costs was at least as important a contributing factor as the weak security controls and fake member counts. Most of the media stories focused on the salacious understory of adultery but the true cause may have been far simpler. Ashley Madison's story is not an isolated incident by any measure. "Biderman's Dilemma" illustrates how poor decision making contributes to security weakness in complex ways that are not always apparent. The lesson of Madison's story illustrates the importance of looking beyond a single cyber event to understand the dynamics within a firm that contribute to the root cause of a breach. Madison's hack was indicative of a failure of decision making. The hack was a symptom of a larger problem: management's inability to grow the firm without the financial resources needed for expansion. For better or worse, security is a tradeoff between risks and opportunities to grow the business. How one makes a choice between the options presented may depend on what is valued more as opposed to an analysis of the risks.

According to Dell's Threat Report, more than 1.7 trillion intrusion prevention system (IPS) attacks were blocked in 2014 versus 2.16 trillion in 2015, a 73% increase representing a tripling since 2013.* The main function of IPSs is to identify malicious activity, log information about this activity, attempt to block/stop it, and report it. With more than 88 trillion attacks on application traffic observed by just one vendor, it's clear that the magnitude and velocity of attacks represent the weaponization of cyber threats as a growing phenomenon.

Even more disturbing is the recent discovery of a massive security gap exploited by hackers that had previously been rated an improbable risk factor. In September 2015, "Security researchers uncovered clandestine attacks across three continents on the routers that direct traffic around the Internet, potentially allowing suspected cyberspies to harvest vast amounts of data while going undetected for at least one year," as reported in a Cisco router breach.†

According to FireEye, a security research firm, a hacker in a highly sophisticated attack used malicious code, dubbed "SYNful Knock," to take over routers used by Cisco Systems, one of the world's top suppliers spanning three continents and used to direct traffic around the Internet. "Routers are attractive to hackers because they operate outside the perimeter of firewalls, anti-virus, behavioral detection software and other security tools that organizations use to safeguard data traffic. Until now, they were considered vulnerable to sustained

* http://www.computerweekly.com/news/4500273520/Encrypted-traffic-security-analysis-a-top-priority-for-2016-says-Dell-Security

† http://www.reuters.com/article/cybersecurity-routers-cisco-systems-upda-idUSL5N11L0VM20150915

denial-of-service attacks using barrages of millions of packets of data, but not outright takeover."*

"This finding represents the ultimate spying tool, the ultimate espionage tool, the ultimate cybercrime tool," according to the CEO of the cyber research firm. "If you own the routers (seize control) you own all of the data of the companies and government organizations that sit behind the routers." Apparently, these were legacy routers that are no longer sold by Cisco but were being maintained for existing customers. Although Cisco claimed no responsibility for the vulnerability, the firm speculated that hackers either gained access to the targeted customers' administrative credentials or acquired physical access directly to the routers.

The full extent of this massive breach is still unfolding as of the writing of this book. Yet when you look at the narrative used to explain the hack, it is clear that the expectation of security was based on untested speculation or not challenged for validation. No doubt, many new lessons will emerge from the SYNful Knock hack, and one of these may be insight into more effective ways to communicate news of far-reaching breaches in security. A "soft notice" of the SYNful Knock breach was used by Cisco to known customers impacted based on an initial assessment of the breach; however, what about clients who may be impacted but may be harder to detect? How widely should a cyberattack be communicated outside of impacted systems? The lack of robust reporting of cyberattacks, along with the stigma associated with media, regulatory, and shareholder scrutiny, provides hackers with a head start to continue attacking others in an industry or modify the attack to strike in a new way in the near future.

SYNful Knock and similar attacks represent a residual threat that lingers beyond the initial breach. The severity and frequency of propagation of additional infection or the return of hackers to soft targets is high in large attacks. Research points out that firms are reluctant to report attacks or have delayed reporting due to fear of litigation or the need for more time to investigate the root cause of the breach thoroughly, impacting timely defensive response and dissemination of critical details helpful to others similarly exposed to the threat.

* http://uk.reuters.com/article/us-cybersecurity-routers-cisco-systems-idUKKCN0RF0N420150915

A national "clearinghouse of cyberattack data" is needed as part of a self-regulatory system to improve the response time for events with large-scale impact. The creation of a national clearinghouse of cyberattack data should be given "Safe Harbor" status for reporting fully and completely in stages during and after the event updates on attack characteristics. A "clearinghouse" facilitates the creation of a single repository that ensures data quality through a standardized reporting regime. In addition, the establishment of safe harbor provisions is critical for minimizing adverse litigation: a standardized stochastic database of sufficient size provides a credible source for projecting trends and developing useful patterns for security response. A "clearinghouse" also establishes criteria for sharing data with law enforcement and the larger community in anonymity while investigations continue protected by safe harbor from defending lawsuits, allowing firms to conduct more thorough analysis of the root cause and thereby improving reporting accuracy. The need for a "self-regulatory" association is addressed later in this book but the purpose is to build a collaborative cybersecurity community and leverage thinking from a broad range of disciplines and standards organizations.

A legal framework is evolving; however, more is needed. A sense of urgency is felt for addressing eCommerce and security issues across borders and boundaries that didn't exist when current law was written. A Bloomberg article describing the frustration technology firms feel in dealing with a legal system that is challenged to keep pace with advancements on the web quoted a comment by Larry Page at a Google developer's conference in 2013: "The law can't be right if it's 50 years old. Like, it's before the internet."* No one expects Congress to act any time soon, leaving firms to depend on a patchwork of court precedents to wade through a number of operational cases. Examples include the following: How are classifications (employees or contractors) for Gig-economy workers selected? What protections are provided under existing copyright laws? What jurisdictional powers does the United States have over data stored in cloud servers across international boundaries? And can France expand Europe's "right to be forgotten" worldwide? These cases touch on a very small number

* http://www.bloomberg.com/news/articles/2016-06-23/the-right-to-be-forgotten-and-other-cyberlaw-cases-go-to-court

of important considerations, many not yet raised, including a lack of guidance on security and knock-on liabilities in the event of a breach.

The sophistication of SYNful Knock has been attributed to advanced nation-states, such as China or Russia, yet these assumptions may prove inaccurate as well. Hackers have become adroit at covering their tracks to their true identity and source of origin, making complex assumptions based on incomplete data inadequate for accurate attribution. What we do know is that several countries, including the United States, are formally developing cyber talent in specialized educational programs from the high school level through college and university. Experts are now well aware of or strongly suspect that nation-states have used cyberattacks to steal intellectual property and monitor certain assets deemed critical in counter-surveillance exercises. Attempts have been made to develop rules of engagement between nations in tactics and strategies dealing with cyber espionage.

A common theme in security weakness points to system and infrastructure complexity as firms layer policy and security infrastructure in a labyrinth designed to create what many call the "M & M" defense, hard on the outside and soft in the center. To explain the need for more security, and thus increased complexity, an equally complex taxonomy has been developed to help laypersons and senior management understand why these resources are needed. The M & M defense is one analogy used but many others have also cropped up. The analogies depict anecdotal solutions with no real analytical or quantifiable justification for these investments. Some of the more interesting analogies include "Cyber Pearl Harbor," "Brakes on a Racing Car," "Holistic Security," "Fortress Security," "Looking around Corners," and the list goes on. In fact, one recent report listed 32 examples of colorful analogies used by security professionals to describe their cyber programs but not one of them explains how it reduces or mitigates cyber risk.

The language of risk or, more succinctly, the lack of insightful communications about risk, creates unnecessary complexity in security response. Too often fear about the uncertainty of a risk, or worse, a false sense of security, creates contradictions in security that lead to poor outcomes. Cybersecurity is not alone in its imprecision in communicating risk concepts. However, the language of risk is a key indicator of the maturity of a cybersecurity program, with security complexity an outcome that inevitably leads to system failures. One appropriate analogy is, "If you don't have a planned destination any road will lead there." Senior management should expect to know and understand exactly which risks will be reduced, and residual risks remain in a cybersecurity program. Further, in making selections about security defense strategy, security analysts must distinguish between which recommendations are assumptions and which data represent facts. Industry benchmarks against standards or "best practice" within an industry is insufficient for developing assurance. Assurance can be derived only for a robust quantitative and qualitative analysis of credible data about cyber risks, and that takes time to develop. What we do know today is that complexity is the enemy of good cybersecurity or any risk management program; therefore, strategies to streamline complexity and make security intuitive will be more effective.

In recent years, organizations have thrown massive resources at a moving target. As soon as the threat vector changes, security defenses are rendered inadequate subject to new vulnerabilities or more sophisticated breach behavior. Security professionals are aware that a patchwork of defensive strategies is not sustainable, but implementing an enterprise solution is still elusive. AlgoSec, a securities research firm, conducted a 2012 study of the complexity of network security with more than 100 IT professionals from its global database. The findings are a very small sample and should not be extrapolated broadly, however, the results are consistent with a common understanding of network complexity.

More than half of the respondents stated that network security complexity had actually contributed to cyberattacks. Instead of streamlining security policies as threats change, new measures are "bolted onto" existing protocols, creating more complexity and resulting in human error and inconsistency in execution from too many policies to manage. Adding to the level of complexity, security professionals with vendor-specific skillsets are required to support multiple vendor systems, adding costly redundancies and inefficient manual processes.

"This is interesting considering that 95% of organizations use network security devices from multiple vendors. Even as more policies, vendors and devices have been added to increasingly complex environments, an estimated 75% of organizations still manually manage network security."* "Automation and consolidation are two valid ways to simplify network security policy management and reduce the risk of misconfiguration," according to AlgoSec.* Simplicity requires a more precise vision for cybersecurity and an understanding of the obstacles that lead to better outcomes. New approaches are needed.

To understand better how network complexity contributes to vulnerability I will borrow a concept used by John Doyle, the John G. Braun Professor of Control and Dynamical Systems, Electrical Engineering, and BioEngineering at the California Institute of Technology. Doyle explored the nature of complex systems by looking at the engineering design of the Internet. "One line of research portrays the Internet as 'scale-free' (SF) with a 'hub-like' core structure that makes the network simultaneously robust to random losses of nodes yet fragile to targeted attacks on the highly connected nodes or 'hubs.' The resulting error tolerance with attack vulnerability has been proposed as a previously overlooked 'Achilles' heel' of the Internet."†

Doyle's findings were a surprising discovery and have become more evident with the growth of cyberattacks more broadly. "Unfortunately, the Internet's strong robustness and adaptability coexists with an equally extreme fragility to components 'failing on,' particularly by malicious exploitation or hijacking of the very mechanisms that confer its robustness properties at higher levels in the protocol stack. Worms, viruses, spam, and denial-of-service attacks remain familiar examples. This RYF tradeoff is a critical aspect of the Internet, and much research is devoted to enhancing these protocols in the face of new challenges."‡

Doyle introduced the concept of the "Robust Yet Fragile" (RYF) paradigm to explain the five components of network design used to build a robust system. Each design component is built on the concept of adding robustness to networks to handle today's evolving business needs. *Reliability* is robustness to component failures. *Efficiency* is robustness to resource scarcity. *Scalability* is robustness to changes in the size and complexity of the system as a whole. *Modularity* is

* http://www.algosec.com/en/resources/examining_the_dangers_of_complexity_in_network_security_environments

† http://www.pnas.org/content/102/41/14497.full, PNAS 2005 102 (41) 14497–14502; published ahead of print October 4, 2005, doi:10.1073/pnas.0501426102

‡ http://www.maths.adelaide.edu.au/matthew.roughan/Papers/PNAS_2005.pdf

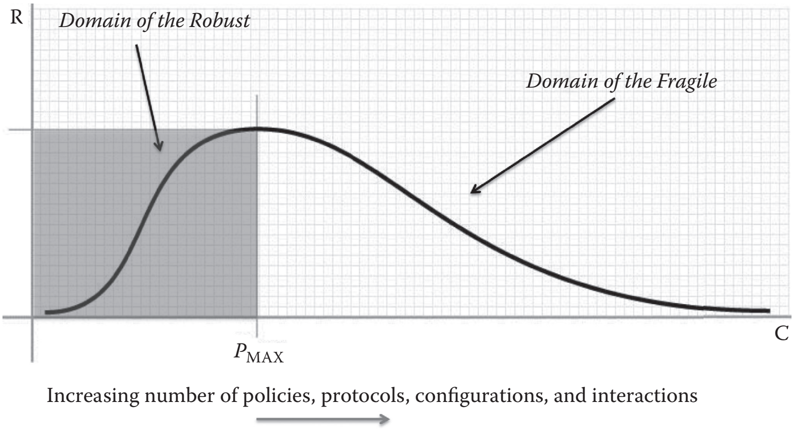

Figure 1.1 Robustness versus complexity: systems view.

robustness to structure component rearrangements. *Evolvability* is robustness of lineages to changes on long timescales. The Robust Yet Fragile concept is represented in Figure 1.1.*

The graph describes the optimal point of robust network design. Like all systems of equilibrium, the point at which robust network design leads to unnecessary complexity is a paradox faced by security professionals and systems architects. Systems such as the Internet are robust for a single point of failure yet fragile to a targeted attack. As networks bolt on more stuff to build scale, the weight of all that stuff becomes more risky. The cost of the tools that were designed to make business competitive and efficient has begun to exceed the benefit as an indirect result of vulnerabilities in scale. The security paradox is represented as the rising cost of marginal security at the point of fragility. As systems become more complex, the point of diminishing returns appears in the frequency and severity of incidents requiring remediation.

The long-tail risk of the "Domain of the Fragile" in the graph demonstrates increasing uncertainty and the likelihood of losses exceeding expectations. Ironically, as organizations build scale more resources are outsourced to vendors and third-party providers, creating the unintended effect of extending fragility beyond the full control of the

* http://www.maoz.com/~dmm/talks/I2_member_meeting_2013.pdf

organization. In other words, cloud computing, wireless devices, and other tools designed to reduce costs and streamline infrastructure may be lengthening the Domain of the Fragile in excess of the short-term benefits of business convenience. This observation does not mean that third-party vendors should not be used, but what it does suggest is an understanding of the incremental risk exposure outsourcing adds to infrastructure. The point is that as a firm moves further along the RYF curve, cost savings are not the only consideration. The art and science of measuring network complexity is still evolving, as are the standards and security tools used by vendors to address these exposures. Organizations need a framework for measuring and defining robustness and early warnings of increased fragility.

More than 60% of technology experts predicted that between 2016 and 2025 a major cyberattack would occur resulting in "significant loss of life or property losses/damage/theft in the tens of billions of dollars."* Others believe the threats are hype by software vendors to promote anxiety to justify new products and services. The idea of a "Cyber Pearl Harbor" is frequently attributed to former defense secretary Leon Panetta. "In a speech at the Intrepid Sea, Air and Space Museum in New York, Mr. Panetta painted a dire picture of how such an attack on the United States might unfold. He said he was reacting to increasing aggressiveness and technological advances by the nation's adversaries, which officials identified as China, Russia, Iran and militant groups."† Secretary Panetta was not the first to use the concept, which dates as far back as 1991 when Win Schwartu introduced the possibility in testimony to Congress. Although nothing of this magnitude has happened to date, it is not out of the realm of possibility.

Nevertheless, warning about a risk is very different from taking effective actions to prevent or mitigate the risk. I will spend time later in the book to review the research and development various groups have started for next generation security. Secretary Panetta's points should be reframed into a question: How must the cybersecurity

* http://www.defenseone.com/threats/2014/10/cyber-attack-will-cause-significant-loss-life-2025-experts-predict/97688/

† http://www.nytimes.com/2012/10/12/world/panetta-warns-of-dire-threat-of-cyberattack.html?_r=0

community respond most effectively to an adversary with advanced skills operating in near anonymity using increasingly powerful tools? The solutions may depend on how each of these characteristics is viewed collectively by industry, the military, and government. Intelligence is needed to answer these questions most effectively, as well as enhanced processes for analyzing the insights derived from the data.

A recent report by the U.S. Department of Justice demonstrates the broad impact attributed to cyberattacks. In April 2015, a cyberattack of federal computer systems exposed up to 22.1 million U.S. government personnel, or 7% of the population in this one event alone.* Cybersecurity is cited as one of the Department of Justice's highest priorities; however, the FBI was granted a budget of only $314 million in 2014, and 82 out of 134 open jobs for computer scientists have been left unfilled under the Justice Department's Next Generation Cyber Initiative launched in 2012.

Congress, it appears, has not been persuaded to commit fully to anything more than a piecemeal approach to cybersecurity. So far, the cyberwar has been fought in small skirmishes but without better intelligence it may be hard to see beyond the horizon. Recent hacks of government agencies and national voter records resemble behaviors associated with exploratory reconnaissance missions but we may never fully understand the purpose of these hacks or what damage, if any, was done.

Cyber threat intelligence researchers have developed surveillance systems to monitor activity in the dark web and networks like Tor to thwart hackers before a breach is launched. Instead of building ever more elaborate security processes, in-house security researchers have developed proprietary "spider intercepts" to crawl the dark web for nefarious behavior that might lead to a cyberattack on its customers.

For example, a security research firm recently uncovered an unauthorized Twitter account created to appear as part of a legitimate bank's customer service department. The Twitter account suddenly began to offer customer assistance but was immediately thwarted by a spider intercept designed by the firm hired by the real bank as a

* http://www.reuters.com/article/2015/07/30/us-usa-fbi-cyberattack-idUSKCN0Q428220150730

proactive defense strategy to detect fraud.* A diverse growth industry of cyber vendor software has exploded in the last five years in response to demand from firms of all types seeking help to defend against an onslaught of attacks. As organizations consider their options for offensive and defensive strategies, vendor selection should be incorporated into a framework designed to simplify cybersecurity.

No matter which camp you fall into, tradeoffs between complexity and security are inevitable using the tools and knowledge available today. The question is, How does an organization assume the right level of complex layers while balancing appropriate security for its business model? Don't expect me to answer the question! The answer is different for each organization. Steve Jobs developed a template for thinking about enterprise models when he created Apple's ecosystem of devices. Jobs described the solution as being at the "Crossroads of Technology and Liberal Arts." For Jobs, this undoubtedly referred to how humans interact with technology.

To paraphrase Jobs, "Technology alone is not enough. Technology married with the liberal arts, [technology] married with [the] humanities is what yield's us the result that makes our heart sing." Jobs further described a post-PC ecosystem that must be easier to use than PCs, more intuitive and integrated. Simplicity is the genius behind Apple's success. Jobs' insightful vision reimagined how humans interact with technology. Apple's ecosystem changed the possibilities for how people interact in their personal lives and increasingly in business life as well. Jobs' model of simplicity may also have applications in cybersecurity.

The National Institute of Standards and Technology (NIST) restated [a similar] observation in a 2013 White House announcement of a [new] Framework to Improve Critical Infrastructure Cybersecurity. "Consequently, we believe that the strategy and tactics we use as defenders must necessarily focus on operational loss minimization."† Said a different way, the NIST recognizes that focusing on too many objects is less effective. However, will a defensive

* http://www.bloomberg.com/news/articles/2015-04-22/in-the-dark-corners-of-the-web-a-spider-intercepts-hackers

† https://www.whitehouse.gov/the-press-office/2013/02/12/executive-order-improving-critical-infrastructure-cybersecurity

strategy alone be enough to slow the loss of critical data? Cybersecurity also needs good offensive weapons with the capability to recognize, respond, and execute defensive strategies without the assistance of humans.

As Steve Jobs eloquently described the problem, "Technology alone is not enough!" Technology designed to anticipate human behavior and internal threats is part of new research being explored in collaboration with government, research universities, and private industry. Sounds simple: build networks with the ability to perform security that we can trust to free humans to focus on the things that add value! How does trust get built into networked information systems and what are the challenges facing an organization that wants to redesign trustworthiness into a cybersecurity defense strategy?

One of the big challenges in cognitive hacks revolves around the issue of trustworthiness ("integrity"). Trustworthiness, or Cyber Trust, focuses on developing systems that are "more predictable, more accountable and less vulnerable to attack."* Developing machines that learn and recognize patterns requires a completely new ecosystem. How does a machine become "smart"? Smart systems require situational awareness that a threat exists and have the ability to select a corrective action that is appropriate for a specific attack.

Trust may seem an innocuous defensive strategy but it is core to basic cybersecurity. However, building trust into Networked Information Systems (NIS) is harder than you might expect. Experts have known for some time that networked information systems are not trustworthy and the technology needed to make them trustworthy has not been available.† The Defense Advanced Research Projects Agency (DARPA) commissioned a study in 1999 to "look beyond policy, procedures and vulnerabilities to a richer set of solutions only new science and technology could provide." The study committee was convened by the Computer Science and Telecommunications Board (CSTB) to assess information systems trustworthiness and the prospects for new technology to increase trustworthiness. If this study were conducted today the issue of trustworthiness would have to be expanded to include a wider range of technologies associated with

* https://people.eecs.berkeley.edu/~tygar/papers/Trust_in_Cyberspace.pdf

† http://www.nap.edu/catalog/6161/trust-in-cyberspace

cognitive hacking not contemplated in 1999; however, the focus is nonetheless instructive in understanding the core challenge of building trust into networked information systems. The study examined "the many dimensions of trustworthiness (e.g., correctness, security, reliability, safety, survivability), the state of the practice, and the available technology and science base."*

Trustworthiness is defined as an expectation that "the system does what is required despite environmental disruptions, human user and operator errors, and attacks by hostile parties. Further, there is an assumption that design and implementation errors must be avoided, eliminated, or somehow tolerated."* Research concludes that the reality of designing a completely trustworthy network is impractical to build. Security professionals therefore must develop strategies for dealing with building trustworthiness into NIS. The challenge of building trust into systems revolves around several critical factors that must be accounted for by security professionals after the fact. Trustworthiness is costly to design and requires advanced skills to implement in configurations that might suit a large number of customers who seek to customize security in different ways.

Observations in the study point to a dilemma between market demands and increased security functionality. "The market has responded best in dimensions, such as reliability, that are easy for consumers (and producers) to evaluate, as compared to other dimensions, such as security, which address exposures that are difficult to quantify or even fully articulate."* The market has favored purchasing commercial off-the-shelf solutions over custom solutions that are more costly and take longer to implement.

To grow faster, solution providers rush to capture market share delivering products to market without trustworthiness functionality because the market has not shown an interest. Solution providers have also been reticent to add functionality that makes configuration and implementation harder for end users. Research in the study suggests that we may be years away from developing trust into NIS at a price point that the market would bear. It is ironic, however, that the industry is willing to spend billions of dollars on cybersecurity after installing NIS without the level of trustworthiness needed to

* https://people.eecs.berkeley.edu/~tygar/papers/Trust_in_Cyberspace.pdf

prevent or partially mitigate the risks. This is the contradiction in how humans evaluate risks and make tradeoffs in security that appear to be rational on the one hand but look irrational on further analysis. The challenge of building NIS with the appropriate level of trust was also evaluated by the study and found that a path might be possible but would require external forces to drive designers to reconsider delivery of trustworthy systems. Either customer demand changes, requiring NIS providers to redesign systems with robust security, or regulatory sentiment changes as a result of an escalation in cyber risk that is deemed unacceptable. What design or engineering changes are required to build cost-effective NIS solutions?

Networked information systems are often large, complex structures that are designed to address the needs of specific organizations. Over time, as the needs of the firm grow through mergers and acquisitions, geographic expansion or obsolescence network complexity inevitably grows, contributing to diminished trustworthiness. Very little research has been conducted over a diverse population of networked information systems; therefore, little understanding exists for improving the design and engineering of these systems to keep up with changes. The root contributing factor that enables the success of hackers is a system of our own design.

We see this pattern repeated over and over again without learning the lesson that we are the designers of our own risks. The incremental costs of repairing our mistakes appear as incremental marginal costs, when in fact these costs, in aggregate, exceed the cost of mitigation in the first instance. Networked information systems are the plumbing that connects us in the eCommerce universe we now live in, and like those of the plumbing connecting households to municipal facilities, the costs of replacing lead pipes with less toxic ones are prohibitive. Going forward, is the alternative method then the use of "Smart" systems? Can we build new applications that account for the inherent lack of trustworthiness in NIS, reducing the need for manual processes or constant human intervention? The answer is yes and work has begun in the research of a new science in Intelligence and Security Informatics (ISI).

What is a "smart system" and how would these applications provide defense against cyberattacks? The role of situational awareness in cybersecurity has garnered a great deal of attention and is the subject of new research in smart systems using Intelligence and Security Informatics (ISI). ISI is defined as the development of advanced information technologies, systems, algorithms, and databases for international, national, and homeland security related applications, through an integrated technological, organizational, and policy-based approach (Mehrotra et al. 2006).* ISI represents a very large body of intensive research in smart applications to solve a diverse set of problems, including cognitive hacking.

Recently, Cybenko et al. (2002a,b) "defined cognitive hacking as an attack on a computer system directed at the mind of the user of the system, which, in order to succeed, had to influence the user's perceptions and behavior." "The National Science Foundation and the National Institute of Justice have recently called for new research in intelligence and security informatics to study semantic attacks and countermeasures."†

In addition to work in ISI security other related areas, research has been undertaken on deception detection in the fields of psychology; communications in the fields of forensic linguistics; and in literary and linguistic computing, in particular research on authorship attribution. This book borrows heavily from this research in Chapter 2 to explore what has been learned and ways in which cognition leads to vulnerability and potentially new approaches to understand and address security more efficiently. This work is timely, as the marginal cost of risk continues to rise, leading to disruptions in business requiring risk transfer strategies to mitigate cyber risk.

The search for an appropriate balance between security and the cost of risk has reached a tipping point. Banks, insurance companies, and financial services firms initially absorbed the cost of security to protect customers and business relationships. But as the cost to defend against cyber risk has been rising rapidly there are signs many firms may begin to push back. The cost of liability is unsustainable for either insurers or small business to handle alone, prompting a shared approach to the risk of cybersecurity.‡

* http://www.security-informatics.com/about

† http://www.ists.dartmouth.edu/library/301.pdf

‡ http://www.marketwatch.com/story/do-you-need-enterprise-grade-cybersecurity-2015-09-21?dist=beforebell

Individuals are pretty well protected when it comes to fraudulent transfers from their bank accounts. Regulation E of the Electronic Fund Transfer Act requires banks to bear the burden in most circumstances. However, to the surprise of many small business owners, banks are not responsible for lost funds due to a cybersecurity breach.* Insurers are stepping in to offer insurance with a condition. Insureds may be required to participate in risk assessments, training, and computer system audits or to pay monthly for monitoring services. These shared risk models have pluses and minuses including the fact that an insurer's primary business is not cybersecurity. Yet, a model in which insurers share the risk with a small business in a bundled program may prove very attractive.

Financial services firms are also raising awareness with consumers about the need to have adequate security on home PCs and mobile devices. However, only 15% of broker dealers and 9% of investment advisers have policies in place that explain liability in the event of a cyber breach according to a Securities and Exchange Commission survey in February 2015.

The point here is that cyber risk is fast becoming an additional cost of doing business on the web. The implications are far reaching as the mobilization of the Internet expands to a variety of devices and spawns new industries. The Internet of Things (IoT) is a concept of connecting any device with an on and off switch to the Internet. This includes everything from cell phones, coffee makers, washing machines, headphones, lamps, wearable devices to almost anything else you can think of. This also applies to components of machinery such as a heating system in a building, car operating systems, or hospital medical devices used to monitor patient care. The rush to market without security in the IoT market raises the bar of trustworthiness to a magnitude few can imagine today. The lack of an agreement to build robust security means that hackers will be able to link billions of devices into an army of drones capable of launching more powerful attacks.

This constant rush forward to introduce new tech products has largely been unregulated, with little, if any, attention paid to security until consumer data are hacked or security breaches are made public

* http://www.npr.org/sections/alltechconsidered/2015/09/15/440252972/when-cyber-fraud-hits-businesses-banks-may-not-offer-protection

by tech firms. Regulators should become more active participants in setting security standards and expectations for data protection in new product development. So where is security today? What is the state of cybersecurity? To answer that question, I looked at data recently released by FireEye, one of the top cybersecurity research firms in the country, to get a sense of the current state of corporate defense combating cybersecurity. The results painted a dismal picture of cybersecurity in general.

FireEye produced its first Special Report, Cybersecurity's Maginot's Line, in 2014 and followed up with trends in 2015 ("Maginot Revisited") gathered from 1,600 network and email sensors installed in real-world corporate networks. Maginot's Line was named for a line of fortifications deployed in ways to slow or repel attack despite its strength and elaborate design, the line was unable to prevent an invasion by German troops who entered France via Belgium.

John Doyle's concept of RYF discussed earlier identified the same weakness in his Domain of the Fragile.* Caveat alert: The data are from a vendor's report and may not be statistically representative of security practice used more broadly. The findings are instructive just the same. The firms in the study had deployed layers and layers of fortress-like IT security measures around the enterprise in an attempt to prevent unauthorized access by threat actors. FireEye installed its sensors behind these existing layers of security to monitor the network of firms participating in the study, giving FireEye a unique perspective on the effectiveness of the "fortress" model of security. "Any threat observed by FireEye in the study had passed through all other security defenses."†

What FireEye discovered was a wakeup call! Attackers are bypassing conventional security measures almost at will! Even more disturbing is that security breaches are widespread across industries and geographic regions. "The new data reaffirms our [FireEye's] initial

* http://www.pnas.org/content/102/41/14497.full, PNAS 2005 102 (41) 14497–14502; published ahead of print October 4, 2005, doi:10.1073/pnas.0501426102

† https://www2.fireeye.com/WEB-2015RPTMaginotRevisited.html

findings. It shows attacks getting through multiple layers of conventional defense-in-depth tools in the vast majority of deployments."*

Before diving into the results, it's important to explain how FireEye conducted the study. FireEye examined data from 1,214 security deployments over two overlapping six-month test periods, comparing the change from the first test. The findings were conclusive, with a particular focus on advanced persistent threat (APT) actors.

APT attacks are not your run of the mill hacks. APT actors may receive direction and support from a national government and, as the name implies, are the most tenacious users of a wide range of tactics and tools in the pursuit of their attack. APT malware is also very stealthy, allowing actors to cloak their actions and, in many cases, their identity. The presence of APT malware does not mean it is being directed by an APT actor but its presence demonstrates the sophistication of the attacker. APT malware is identified with the subtype, "APT," such as BACKDOOR.APT. GH0STRAT. Now for the results.

Brief Summary of Results

"Ninety-six percent of systems across multiple industry types were breached and twenty-seven percent of the breaches involved malware."* The following data represent percentage breaches by industry verticals participating in the FireEye study: 100% in Legal, 30% in Retail, 29% in Auto & Transportation, 28% in Entertainment & Media, 37% in Healthcare & Pharmaceuticals, 30% in Services & Consulting, and 32% in High Tech.

Attacks are increasingly focused on compromising systems through the use of advanced malware. Figure 1.2 gives a breakdown: in short, 96 out of 100 attacks were successful even with layers and layers of security in place! Security defenses were ineffective but not for the reason one would think. Instead, hackers simply found more effective ways to bypass the defenses that were in place. One commonality among all industries is the "attack vector," meaning hackers, at least in this study, have concentrated their efforts on two parts of the fortified infrastructure to deliver their malware.

* https://www2.fireeye.com/WEB-2015RPTMaginotRevisited.html

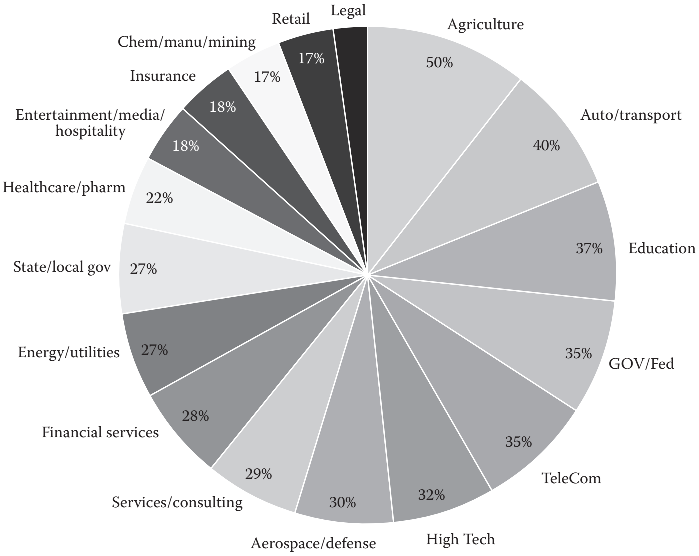

Figure 1.2 Percent of individual industries with advanced malware concentration.

Email and web traffic were cited as the most common and effective ways used to deliver attacks. Social media and malvertisement must now be included as high growth delivery channels, suggesting the problem is spreading. Either the user's web browser was compromised or the attack involved the use of APT malware to trick email users into opening an infected document. Why try to defeat security defenses when the hacker can so easily trick you into providing access to the enterprise? The truth is that the most successful cyberattacks use simple approaches. Cognitive hacks are effective—using a variety of media—with low-cost tactics, yielding tremendous results. In some cases, a form of "crowd-sourcing" for malware has evolved in the deep web, allowing hackers to create more sophisticated versions of successful tools for new attacks rendering defenses useless.

Motivated enemies have exploited human behavior since before the Trojan War to defeat the defenses of its adversaries. It seems not much has changed except the tools used to execute the means to the end. We (humans) are the weak link in Maginot's Line! Recognizing the root cause of the problem is the first step in finding new solutions. Fortunately, a great deal of pioneering work is being conducted to expand our understanding of the role "situational awareness" or cognition contributes to cyber vulnerabilities.

Returning to the final insights in the study, several industry types experienced a higher concentration of breaches. The following industry types were breached 100% of the time: Agriculture, Auto/Transportation, Education, Healthcare/Pharm, and Retail. Ninety percent of all industries in the study experienced one or more breaches except Aerospace and Defense, which recorded a breach 76% of the time. While these results indicate a high level of failure, they demonstrate these industries have either hardened security or attackers have simply been less successful for reasons not identified in the study.

These findings may not be extrapolated uniformly as a benchmark but are informative nonetheless. The FireEye test can serve as a proxy for thinking about cybersecurity and is instructive in evaluating assumptions about security in general. So far, we have taken a very broad brush to explain a nuanced problem, ignoring for a moment that the details help to paint a more complete picture. We will get to the data as we continue to track the digital footprint of cyberattacks. Let's now turn to the subjects of cognition, machine learning, artificial intelligence, and new research in trustworthiness in cyberspace.

References

Cybenko, G., Giani, A., and Thompson, P., "Cognitive Hacking and the Value of Information," Workshop on Economics and Information Security, May 16–17, 2002, Berkeley, California, 24, 2002a.

Cybenko, G., Giani, A., and Thompson, P., "Cognitive Hacking: A Battle for the Mind," *IEEE Computer*, 35(8), 2002b, 50–56.

Mehrotra, S., Zeng, D. D., and Chen, H. (Eds.), IEEE International Conference on Intelligence and Security Informatics, ISI 2006, San Diego, CA, USA, May 23–24, 2006. http://www.springer.com/us/book/9783540344780.