PART I

Why Cybersecurity Needs Better Measurements for Risk

CHAPTER 1

The One Patch Most Needed in Cybersecurity

There is nothing more deceptive than an obvious fact.

—Sherlock Holmes
The Bascombe Valley Mystery[1]

In the days after September 11, 2001, increased security meant overhauled screening at the airport, no-fly lists, air marshals, and attacking terrorist training camps. But just 12 years later, the FBI was emphasizing the emergence of a very different concern: the "cyber-based threat." In 2013, FBI director James B. Comey, testifying before the Senate Committee on Homeland Security and Governmental Affairs, stated the following:

> . . .we anticipate that in the future, resources devoted to cyber-based threats will equal or even eclipse the resources devoted to non-cyber based terrorist threats.
>
> —FBI director James B. Comey, November 14, 2013[2]

This is a shift in priorities we cannot overstate. How many organizations in 2001, preparing for what they perceived as the key threats at the time, would have even imagined that cyber threats would have not only equaled but exceeded more conventional terrorist threats? Yet as we write this book, it is accepted as our new "new normal."

Admittedly, those outside of the world of cybersecurity may think the FBI is sowing seeds of Fear, Uncertainty, and Doubt (FUD) to some political end. But it would seem that there are plenty of sources of FUD, so why pick cyber threats in particular? Of course, to cybersecurity experts this is a non-epiphany. We are under attack and it will certainly get worse before it gets better.

Yet resources are limited. Therefore, the cybersecurity professional must effectively determine a kind of "return on risk mitigation." Whether or not such a return is explicitly calculated, we must evaluate whether a given defense strategy is a better use of resources than another. In short, we have to measure and monetize risk and risk reduction. What we need is a "how to" book for professionals in charge of allocating limited resources to addressing ever-increasing cyber threats, and leveraging those resources for optimum risk reduction. This includes methods for:

- How to measure risk assessment methods themselves.
- How to measure reduction in risk from a given defense, control, mitigation, or strategy (using some of the better-performing methods as identified in the first bullet).
- How to continuously and measurably improve on the implemented methods, using more advanced methods that the reader may employ as he or she feels ready.

Let's be explicit about what this book isn't. This is not a technical security book—if you're looking for a book on "ethical hacking," then you have certainly come to the wrong place. There will be no discussions about how to execute stack overflows, defeat encryption algorithms, or execute SQL injections. If and when we do discuss such things, it's only in the context of understanding them as parameters in a risk model.

But don't be disappointed if you're a technical person. We will certainly be getting into some analytic nitty-gritty as it applies to security. This is from the perspective of an analyst or leader trying to make better bets in relation to possible future losses. For now, let's review the scale of the challenge we are dealing with and how we deal with it currently, then outline a direction for the improvements laid out in the rest of the book.

The Global Attack Surface

Nation-states, organized crime, hacktivist entities, and insider threats want our secrets, our money, and our intellectual property, and some want our complete demise. Sound dramatic? If we understand the FBI correctly, they expect to spend as much or more on protecting us from cyber threats than from those who would turn airplanes, cars, pressure cookers, and even people into bombs. And if you are reading this book, you probably already accept the gravity of the situation. But we should at least spend some time emphasizing this point if for no other reason than to help those who already agree with this point make the case to others.

The Global Information Security Workforce Study (GISWS)—a survey conducted in 2015 of more than 14,000 security professionals, including 1,800 federal employees—showed we are not just taking a beating, we are backpedaling:

> When we consider the amount of effort dedicated over the past two years to furthering the security readiness of federal systems and the nation's overall security posture, our hope was to see an obvious step forward. The data shows that, in fact, we have taken a step back.
>
> —(ISC)^2 on the announcement of the GISWS, 2015[3]

Indeed, other sources of data support this dire conclusion. The UK insurance market, Lloyd's of London, estimated that cyberattacks cost businesses $400 billion globally per year.[4] In 2014, one billion records were compromised. This caused *Forbes* magazine to refer to 2014 as "The Year of the Data Breach."[5,6] Unfortunately, identifying 2014 as the year of the data breach may still prove to be premature. It could easily get worse.

In fact, the founder and head of XL Catlin, the largest insurer in Lloyd's of London, said cybersecurity is the "biggest, most systemic risk" he has seen in his 42 years in insurance.[7] Potential weaknesses in widely used software; interdependent network access between companies, vendors, and clients; and the possibility of large coordinated attacks can affect much more than even one big company like Anthem, Target, or Sony. XL Catlin believes it is possible that there could be a simultaneous impact on multiple major organizations affecting the entire economy. They feel that if there are multiple major claims in a short period of time, this is a bigger burden than insurers can realistically cover.

What is causing such a dramatic rise in breach and the anticipation of even more breaches? It is called attack surface. "Attack surface" is usually defined as the kind of total of all exposures of an information system. It exposes value to untrusted sources. You don't need to be a security professional to get this. Your home, your bank account, your family, and your identity all have an attack surface. If you received identity theft protection as a federal employee, or a customer of Home Depot, Target, Anthem, or Neiman Marcus, then you received that courtesy of an attack surface. These companies put the digital you within reach of criminals. Directly or indirectly, the Internet facilitated this. This evolution happened quickly and without the knowledge or direct permission of all interested parties (organizations, employees, customers, or citizens).

Various definitions of the phrase consider the ways into and out of a system, the defenses of that system, and sometimes the value of data in that system.[8,9] Some definitions of attack surface refer to the attack surface of a system and some refer to the attack surface of a network, but either might be too narrow even for a given firm. We might also define an "Enterprise Attack Surface" that not only consists of all systems and networks in that organization but also the exposure of third parties. This includes everyone in the enterprise "ecosystem" including major customers, vendors, and perhaps government agencies. (Recall that in the case of the Target breach, the exploit came from an HVAC vendor.)

Perhaps the total attack surface that concerns all citizens, consumers, and governments is a kind of "global attack surface": the total set of cybersecurity exposures—across all systems, networks, and organizations—we all face just by shopping with a credit card, browsing online, receiving medical benefits, or even just being employed. This global attack surface is a macro-level phenomenon driven by at least four macro-level causes of growth: increasing users worldwide, variety of users worldwide, growth in discovered and exploited vulnerabilities per person per use, and organizations more networked with each other resulting in "cascade failure" risks.

- *The increasing number of persons on the Internet.* Internet users worldwide grew by a factor of 6 from 2001 to 2014 (half a billion to 3 billion). It may not be obvious that the number of users is a dimension in some attack surfaces, but some measures of attack surface also include the value of a target, which would be partly a function of number of users (e.g., gaining access to more personal records)[10] Also, on a global scale, it acts as an important multiplier on the following dimensions.
- *The number of uses per person for online resources.* The varied uses of the Internet, total time spent on the Internet, use of credit cards, and various services that require the storage of personal data-automated transactions are growing. Per person. Worldwide. For example, since 2001 the number of websites alone has grown at a rate five times faster than the number of users—a billion total by 2014. Connected devices constitute another potential way for an individual to use the Internet even without their active involvement. One forecast regarding the "Internet of Things" (IoT) was made by Gartner, Inc: "4.9 billion connected things will be in use in 2015, up 30 percent from 2014, and will reach 25 billion by 2020."[11] A key concern here is the lack of consistent security in designs. The National Security Telecommunications Advisory Committee determined that "there is a small—and rapidly closing—window to ensure that the IoT is adopted in a way that maximizes security and minimizes risk. If the country fails to do so, it will be coping with the consequences for generations."[12]

- *Vulnerabilities increase.* A natural consequence of the previous two factors is the number of ways such uses can be exploited increases. This is due to the increase in systems and devices with potential vulnerabilities, even if vulnerabilities per system or device do not increase. At least the number of *discovered* vulnerabilities will increase partly because the number of people actively seeking and exploiting vulnerabilities increases. And more of those will be from well-organized and well-funded teams of individuals working for national sponsors.
- *The possibility of a major breach "cascade."* More large organizations are finding efficiencies from being more connected. The fact that Target was breached through a vendor raises the possibility of the same attack affecting multiple organizations. Organizations like Target have many vendors, several of which in turn have multiple large corporate and government clients. Mapping this cyber-ecosystem of connections would be almost impossible, since it would certainly require all these organizations to divulge sensitive information. So the kind of publicly available metrics we have for the previous three factors in this list do not exist for this one. But we suspect most large organizations could just be one or two degrees of separation from each other.

It seems reasonable that of these four trends the earlier trends magnify the latter trends. If so, the risk of the major breach "cascade" event could grow faster than the growth rate of the first couple of trends.

Our naïve, and obvious, hypothesis? Attack surface and breach are correlated. If this holds true, then we haven't seen anything yet. We are heading into a historic growth in attack surface, and hence breach, which will eclipse what has been seen to date. Given all this, the FBI director's comments and the statements of Lloyd's of London insurers cannot be dismissed as alarmist. Even with the giant breaches like Target, Anthem, and Sony behind us, we believe we haven't seen "The Big One" yet.

The Cyber Threat Response

It's a bit of a catch-22 in that success in business is highly correlated with exposure. Banking, buying, getting medical attention, and even being employed is predicated on exposure. You need to expose data to transact business, and if you want to do more business, that means more attack surface. When you are exposed, you can be seen and affected in unexpected and malicious ways. In defense, cybersecurity professionals try to "harden" systems—that is, removing all nonessentials, including programs, users, data, privileges, and vulnerabilities. Hardening shrinks, but does not eliminate, attack surface. Yet even this partial reduction in attack surface requires significant resources, and the trends show that the resource requirements will grow.

Generally, executive-level attention on cybersecurity risks has increased, and attention is followed by resources. The boardroom is beginning to ask questions like "Will we be breached?" or "Are we better than Sony?" or "Did we spend enough on the right risks?" Asking these questions eventually brings some to hire a chief information security officer (CISO). The first Fortune 100 CISO role emerged more than 20 years ago, but for most of that time growth in CISOs was slow. *CFO Magazine* acknowledged that hiring a CISO as recently as 2008 would have been considered "superfluous."[13] In fact, large companies are still in the process of hiring their first CISOs, many just after they suffer major breaches. By the time this book was written, Target finally hired their first CISO,[14] and JPMorgan did likewise after their breach.[15]

In addition to merely asking these questions and creating a management-level role for information security, corporations have been showing a willingness, perhaps more slowly than cybersecurity professionals would like, to allocate serious resources to this problem:

- Just after the 9/11 attacks the annual cybersecurity market in the United States was $4.1 billion.[16] By 2015 the information technology budget of the United States Defense Department had grown to $36.7 billion.[17]
- This does not include $1.4 billion in startup investments for new cybersecurity-related firms.[18]
- Cybersecurity budgets have grown at about twice the rate of IT budgets overall.[19]

So what do organizations do with this new executive visibility and inflow of money to cybersecurity? Mostly, they seek out vulnerabilities, detect attacks, and eliminate compromises. Of course, the size of the attack surface and the sheer volume of vulnerabilities, attacks, and compromises means organizations must make tough choices; not everything gets fixed, stopped, recovered, and so forth. There will need to be some form of acceptable (tolerable) losses. What risks are acceptable is often not documented, and when they are, they are stated in soft, unquantified terms that cannot be used clearly in a calculation to determine if a given expenditure is justified or not.

On the vulnerability side of the equation, this has led to what is called "vulnerability management." An extension on the attack side is "security event management," which can generalize to "security management." More recently there is "threat intelligence" and the emerging phrase "threat management."

While all are within the tactical security solution spaces, the management portion attempts to rank-order what to do next. So how do organizations conduct security management? How do they prioritize the allocation of significant, but limited, resources for an expanding list of vulnerabilities? In other words, how do they make cybersecurity decisions to allocate limited resources in a fight against such uncertain and growing risks?

Certainly a lot of expert intuition is involved, as there always is in management. But for more systematic approaches, the vast majority of organizations concerned with cybersecurity will resort to some sort of "scoring" method that ultimately plots risks on a "matrix." This is true for both very tactical level issues and strategic, aggregated risks. For example, an application with multiple vulnerabilities could have all of them aggregated into one score. Using similar methods at another scale, groups of applications can then be aggregated into a portfolio and plotted with other portfolios. The aggregation process is typically some form of invented mathematics unfamiliar to actuaries, statisticians, and mathematicians.

In one widely used approach, "likelihood" and "impact" will be rated subjectively, perhaps on a 1 to 5 scale, and those two values will be used to plot a particular risk on a matrix (variously called a "risk matrix," "heat map," "risk map," etc.). The matrix—similar to the one shown in Figure 1.1—is then often further divided into sections of low, medium, and high risk. Events with high likelihood and high impact would be in the upper-right "high risk" corner, while those with low likelihood and low impact would be in the opposite "low risk" corner. The idea is that the higher the score, the more important something is and the sooner you should address it. You may intuitively think such an approach is reasonable, and if you thought so you would be in good company.

			Impact				
			Negligible	**Minor**	**Moderate**	**Critical**	**Catastrophic**
			1	2	3	4	5
Likelihood	**Frequent**	5	Medium	Medium	**High**	**High**	**High**
	Likely	4	Medium	Medium	Medium	**High**	**High**
	Occasional	3	Low	Medium	Medium	Medium	**High**
	Seldom	2	Low	Low	Medium	Medium	Medium
	Improbable	1	Low	Low	Low	Medium	Medium

FIGURE 1.1 The familiar risk matrix (a.k.a. heat map or risk map)

Various versions of scores and risk maps are endorsed and promoted by several major organizations, standards, and frameworks such as the National Institute of Standards and Technology (NIST), the International Standards Organization (ISO), MITRE.org, and the Open Web Application Security Project (OWASP), among others. Most organizations with a cybersecurity function claim at least one of these as part of their framework for assessing risk. In fact, most major software organizations like Oracle, Microsoft, and Adobe rate their vulnerabilities using a NIST-supported scoring system called the "Common Vulnerability Scoring System" (CVSS). Also, many security solutions also include CVSS ratings, be it for vulnerability and/or attack related. While the control recommendations made by many of these frameworks are good, it's how we are guided to prioritize risk management on an enterprise scale that is amplifying risk.

Literally hundreds of security vendors and even standards bodies have come to adopt some form of scoring system. Indeed, scoring approaches and risk matrices are at the core of the security industry's risk management approaches.

In all cases, they are based on the idea that such methods are of some sufficient benefit. That is, they are assumed to be at least an improvement over not using such a method. As one of the standards organizations has put it, rating risk this way is adequate:

> Once the tester has identified a potential risk and wants to figure out how serious it is, the first step is to estimate the *likelihood*. At the highest level, this is a rough measure of how likely this particular vulnerability is to be uncovered and exploited by an attacker. It is not necessary to be over-precise in this estimate. Generally, identifying whether the likelihood is low, medium, or high is sufficient.
>
> —OWASP[20] (emphasis added)

Does this last phrase, stating "low, medium, or high is sufficient," need to be taken on faith? Considering the critical nature of the decisions such methods will guide, we argue that it should not. This is a testable hypothesis and it actually *has been* tested in many different ways. The growing trends of cybersecurity attacks alone indicate it might be high time to try something else.

So let's be clear about our position on current methods: *They are a failure. They do not work.* A thorough investigation of the research on these methods and decision-making methods in general indicates the following (all of this will be discussed in detail in Chapters 4 and 5):

- There is no evidence that the types of scoring and risk matrix methods widely used in cybersecurity improve judgment.
- On the contrary, there is evidence these methods add noise and error to the judgment process. One researcher—Tony Cox—goes as far as to say they can be "worse than random." (Cox's research and many others will be detailed in Chapter 5.)
- Any appearance of "working" is probably a type of "analysis placebo." That is, a method may make you feel better even though the activity provides no measurable improvement in estimating risks (or even adds error).
- There is overwhelming evidence in published research that quantitative, probabilistic methods are effective.
- Fortunately, most cybersecurity experts seem willing and able to adopt better quantitative solutions. But common misconceptions held by some—including misconceptions about basic statistics—create some obstacles for adopting better methods.

How cybersecurity assesses risk, and how it determines how much it reduces risk, are the basis for determining where cybersecurity needs to prioritize the use of resources. And if this method is broken—or even just leaves room for significant improvement—then that is the highest-priority problem for cybersecurity to tackle! Clearly, putting cybersecurity risk-assessment and decision-making methods on a solid foundation will affect everything else cybersecurity does. If risk assessment itself is a weakness, then fixing risk assessment is the most important "patch" a cybersecurity professional can implement.

A Proposal for Cybersecurity Risk Management

In this book, we will propose a different direction for cybersecurity. Every proposed solution will ultimately be guided by the title of this book. That is, we are solving problems by describing how to measure cybersecurity risk—*anything* in cybersecurity risk. These measurements will be a tool in the solutions proposed but also reveal how these solutions were selected in the first place. So let us propose that we adopt a new quantitative approach to cybersecurity, built upon the following principles:

- *It is possible to greatly improve on the existing methods.* Many aspects of existing methods have been measured and found wanting. This is not acceptable for the scale of the problems faced in cybersecurity.

- *Cybersecurity can use the same quantitative language of risk analysis used in other problems.* As we will see, there are plenty of fields with massive risk, minimal data, and profoundly chaotic actors that are regularly modeled using traditional mathematical methods. We don't need to reinvent terminology or methods from other fields that also have challenging risk analysis problems.
- *Methods exist that have already been measured to be an improvement over expert intuition.* This improvement exists even when methods are based, as are the current methods, on only the subjective judgment of cybersecurity experts.
- *These improved methods are entirely feasible.* We know this because it has already been done. One or both of the authors have had direct experience with using every method described in this book in real-world corporate environments. The methods are currently used by cybersecurity analysts with a variety of backgrounds.
- *You can improve further on these models with empirical data.* You have more data available than you think from a variety of existing and newly emerging sources. Even when data is scarce, mathematical methods with limited data can still be an improvement on subjective judgment alone. Even the risk analysis methods themselves can be measured and tracked to make continuous improvements.

The book is separated into three parts that will make each of these points in multiple ways. Part I will introduce a simple quantitative method that requires little more effort than the current scoring methods, but uses techniques that have shown a measurable improvement in judgment. It will then discuss how to measure the measurement methods themselves. In other words, we will try to answer the question "How do we know it works?" regarding different methods for assessing cybersecurity. The last chapter of Part I will address common objections to quantitative methods, detail the research against scoring methods, and discuss misconceptions and misunderstandings that keep some from adopting better methods.

Part II will move from the "why" we use the methods we use and focus on how to add further improvements to the simple model described in Part I. We will talk about how to add useful details to the simple model, how to refine the ability of cybersecurity experts to assess uncertainties, and how to improve a model with empirical data (even when data seems limited).

Part III will take a step back to the bigger picture of how these methods can be rolled out to the enterprise, how new threats may emerge, and how evolving tools and methods can further improve the measurement of cybersecurity risks. We will try to describe a call to action for the cybersecurity industry as a whole.

But first, our next chapter will build a foundation for how we should understand the term “measurement.” That may seem simple and obvious, but misunderstandings about that term and the methods required to execute it are behind at least some of the resistance to applying measurement to cybersecurity.

Notes

1. Sir Arthur Conan Doyle, “The Boscombe Valley Mystery,” *The Strand Magazine*, 1891.
2. Greg Miller, “FBI Director Warns of Cyberattacks; Other Security Chiefs Say Terrorism Threat Has Altered,” *Washington Post,* November 14, 2013, www.washingtonpost.com/world/national-security/fbi-director-warns-of-cyberattacks-other-security-chiefs-say-terrorism-threat-has-altered/2013/11/14/24f1b27a-4d53-11e3-9890-a1e0997fb0c0_story.html.
3. Dan Waddell, Director of Government Affairs, National Capital Regions of (ISC)^2 in an announcement of the Global Information Security Workforce Study (GISWS), www.isc2.org, May 14, 2015.
4. Stephen Gandel, “Lloyd’s CEO: Cyber Attacks Cost Companies $400 Billion Every Year,” Fortune.com, January 23, 2015, http://fortune.com/2015/01/23/cyber-attack-insurance-lloyds/.
5. Sue Poremba, “2014 Cyber Security News Was Dominated by the Sony Hack Scandal and Retail Data Breaches,” *Forbes Magazine*, December 31, 2014, www.forbes.com/sites/sungardas/2014/12/31/2014-cyber-security-news-was-dominated-by-the-sony-hack-scandal-and-retail-data-breaches/#1c79203e4910.
6. Kevin Haley, “The 2014 Internet Security Threat Report: Year Of The Mega Data Breach,” *Forbes Magazine*, July 24, 2014, www.forbes.com/sites/symantec/2014/07/24/the-2014-internet-security-threat-report-year-of-the-mega-data-breach/#724e90a01a98.
7. Matthew Heller, “Lloyd’s Insurer Says Cyber Risks Too Big to Cover,” CFO.com, February 6, 2015, ww2.cfo.com/risk-management/2015/02/lloyds-insurer-says-cyber-risks-big-cover/.
8. Jim Bird and Jim Manico, “Attack Surface Analysis Cheat Sheet.” OWASP.org. July 18, 2015, www.owasp.org/index.php/Attack_Surface_Analysis_Cheat_Sheet.
9. Stephen Northcutt, “The Attack Surface Problem.” SANS.edu. January 7, 2011, www.sans.edu/research/security-laboratory/article/did-attack-surface.
10. Pratyusa K. Manadhata and Jeannette M. Wing, “An Attack Surface Metric,” *IEEE Transactions on Software Engineering* 37, no. 3 (2010): 371–386.

11. Gartner, “Gartner Says 4.9 Billion Connected ‘Things’ Will Be in Use in 2015” (press release), November 11, 2014, www.gartner.com/newsroom/id/2905717.
12. The President’s National Security Telecommunications Advisory Committee, “NSTAC Report to the President on the Internet of Things,” November 19, 2014, www.dhs.gov/sites/default/files/publications/IoT%20Final%20Draft%20Report%2011-2014.pdf.
13. Alissa Ponchione, “CISOs: The CFOs of IT,” *CFO*, November 7, 2013, ww2.cfo.com/technology/2013/11/cisos-cfos/.
14. Matthew J. Schwartz, “Target Ignored Data Breach Alarms,” *Dark Reading* (blog), *InformationWeek,* March 14, 2014, www.darkreading.com/attacks-and-breaches/target-ignored-data-breach-alarms/d/d-id/1127712.
15. Elizabeth Weise, “Chief Information Security Officers Hard to Find—and Harder to Keep,” *USA Today*, December 3, 2014, www.usatoday.com/story/tech/2014/12/02/sony-hack-attack-chief-information-security-officer-philip-reitinger/19776929/.
16. Kelly Kavanagh, “North America Security Market Forecast: 2001–2006,” Gartner, October 9, 2002, www.bus.umich.edu/KresgePublic/Journals/Gartner/research/110400/110432/110432.html.
17. Sean Brodrick, “Why 2016 Will Be the Year of Cybersecurity,” *Energy & Resources Digest,* December 30, 2015, http://energyandresourcesdigest.com/invest-cybersecurity-2016-hack-cibr/.
18. Deborah Gage, “VCs Pour Money into Cybersecurity Startups,” *Wall Street Journal*, April 19, 2015, www.wsj.com/articles/vcs-pour-money-into-cybersecurity-startups-1429499474.
19. PWC, *Managing Cyber Risks in an Interconnected World: Key Findings from the Global State of Information Security Survey 2015,* September 30, 2014, www.pwc.be/en/news-publications/publications/2014/gsiss2015.html.
20. OWASP, “OWASP Risk Rating Methodology,” last modified September 3, 2015, www.owasp.org/index.php/OWASP_Risk_Rating_Methodology.