CHAPTER 1

The One Patch Most Needed in Cybersecurity

Everything's fine today, that is our illusion.

—Voltaire, 1759[1]

In a single year, cyberattacks resulted in one billion records compromised and financial losses of $400 billion.[2] This led *Forbes Magazine* to declare it "The Year of the Mega Data Breach."[3,4] Soon afterward, the head of the largest insurer, Lloyd's of London—a marketplace for a collection of insurance companies, reinsurance companies, and other types of financial backers—said cybersecurity is the "biggest, most systemic risk" he has seen in his 42 years in insurance.[5] The year of that article was 2014. A lot has happened since then.

By multiple measures, cybersecurity risk has been increasing every year since 2014. For example, records breached in 2021 were, according to one source, 22 times as high as 2014.[6] It hasn't peaked. We will only become more dependent on—and more vulnerable to—the technologies that drive our prosperity. We can try to reduce these risks, but resources are limited. To management, these risks may seem abstract, especially if they haven't experienced these losses directly. And yet we need to convince management, in their language, that these issues require their attention and a significant budget. Once we have that, we can try to identify and address higher priority risks first.

The title of this book is self-explanatory. We will talk about how we can measure risks in cybersecurity and why it is important to change how we currently do it. For now, we will just make the case that there is a reason to be worried—both about the threats to cybersecurity and the adequacy of methods to assess them.

Insurance: A Canary in the Coal Mine

One of the authors of this book, Richard Seiersen, was a chief information security officer (CISO) who is now working as a chief risk officer (CRO) for the cyber insurance firm Resilience. These two viewpoints provide a useful perspective on cybersecurity risk. Insurance is at its core a "put your money where your mouth is" business. When insurers make bad bets, it catches up to them. However, to be competitive they can't just charge whatever they want. They have a strong incentive to gather a lot of data, do the math, and work out what makes a good bet in the risks they cover. That doesn't mean they are always right. It is a bet, after all. But they've done their homework, and their analysis is usually better than what most firms can muster.

Richard would point out that insurance companies reveal their concerns about a risk when they increase their premiums, tighten their underwriting requirements, or quit selling a type of coverage altogether. They are a sort of canary in the coal mine of risks. What has been happening in the field of cyber insurance is a type of leading indicator CISOs should pay attention to.

According to the National Association of Insurance Commissioners (NAIC), in the years from 2017 to 2021, total premiums collected have increased by 45%, and the share of those premiums paid out for claims has more than doubled in that same period.[7] This means total cyber insurance claims paid has more than tripled in that same period. Note that claims just cover *some* losses. They exclude the retention (what retail consumer insurance would call a deductible), anything over the limit covered by the insurer, and exclusions such as acts of war.

If claims are completely independent of each other, then there is some expected variation from year to year, just like the total from rolling 100 dice will give you a slightly different answer from rolling another 100 dice. Insurance companies plan for that in how they manage their reserves for paying claims. However, the amount of change the NAIC observed is far beyond what could be explained as a random fluke. There are common, underlying trends driving all claims to be more frequent and more costly. This is the "systemic risk" mentioned earlier by the head of Lloyd's. In addition to claims being larger and more frequent, there is a risk that many claims could all happen at the same time making it difficult or impossible for an insurer to cover them. Furthermore, certain recent developments have made systemic risks an even bigger problem for insurers.

A key legal battle created a new systemic risk for insurers, which forced them to rewrite policies or, in some cases, get out of cyber insurance. In January 2022, Chubb, the largest cyber insurance provider, lost a case over whether it should cover $1.4 billion in losses claimed

by the pharmaceutical giant Merck.[8] Merck was hit by malicious code known as "NotPetya," which encrypted the data on thousands of Merck's computers.

Because the source of the attack was six Russians with ties to Russian intelligence agencies, Chubb argued that this was an act of war, which is typically excluded in property insurance. But the court ruled the policy excluded only physical warfare, not cyber war. Other insurers took notice and responded by writing much more stringent underwriting requirements. In August 2022, Lloyd's of London advised all cyber insurers selling through its platform to stop selling coverage for cyberattacks that are sponsored by government agencies.[9]

The NotPetya malware, which attacked Merck, was based on some previous code known as Petya. While Petya was used in ransomware, the attack on Merck did not demand ransom. This code, created by Russia to attack Ukrainian systems, was simply for the purpose of destruction. While both destructive and financially harmful for those effected, it could have been much worse.

A different hack on the company SolarWinds shows how widely one piece of malware can be spread. SolarWinds develops system performance monitoring software. One of its software packages, the Orion network management system, is used by over 30,000 public and private institutions. In 2020, it was revealed that an update for Orion, which SolarWinds sent to its customers, contained some malicious code. While the attack was widespread, it seems that companies (especially insurance) dodged a bullet. The SolarWinds hack was a major attack by any standard, but the motive appears to be access to classified government data more than exploiting individuals.

The original target of NotPetya, on the other hand, was a single Ukrainian software company but the malicious code leaked out to numerous Ukrainian entities such as the National Bank of Ukraine—and spread across the globe. It led to billions of dollars of impact—much of which was collateral damage. And still it was not as widespread as the malware that hit SolarWinds. If one attack were as widespread as SolarWinds and as destructive as NotPetya, we would have had a completely different story.

The change in act-of-war exclusions combined with an apparent increase in frequency of exactly that kind of event adds to a growing list of systemic risks. 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. We said this in the first edition of this book in 2016, and it is just as true now if not more. If this results in multiple major claims in a short period of time, this may be a bigger burden than insurers can realistically cover. What worries the insurance companies is that even the biggest attacks seen so far aren't as big as they could have been.

The Global Attack Surface

As we mentioned above, insurance companies have the option of limiting their risk by changing policy language or simply not selling insurance to you if they feel your risk is too great. They can simply choose not to participate in that risk. You don't have that option. Whatever risks your insurers won't cover are still carried by you. Nation-states, organized crime, hacktivist entities, and insider threats want our secrets, our money, and our intellectual property, and some want our complete demise. That's not just being dramatic. If you are reading this book, you probably already accept the gravity of the situation.

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 total of all kinds of 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 as a customer of a major retailer, 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 not always with the knowledge or direct permission of all interested parties like you.

Various definitions of attack surface consider the ways into and out of a system, the defenses of that system, and sometimes the value of data in that system.[10,11] Some definitions 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. As we saw in the 2013 data breach of the major retailer Target, every possible connection matters. That exploit came from a vendor with access to their HVAC systems.[12]

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, 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 10 from 2001 to 2022 (half a billion to

5 billion).[13] 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).[14] Also, on a global scale, it acts as an important multiplier on the following dimensions.

- *Each person does more online.* We were already doing a lot online, and the pandemic accelerated it. In 2020, the first year of the pandemic, e-commerce increased by 43%.[15] The videoconference became the new normal for meeting during the pandemic. There is also a different kind of "pandemic" in the growing number of connected devices per person. The Internet of Things (IoT) constitutes another potential way for an individual to use the Internet even without their active involvement—attacks on IoT devices tripled in the first half of 2019.[16]
- *Vulnerabilities increase.* A natural consequence of the previous two factors is the number of ways such uses can be exploited increases. According to the Mitre CVE database, total known vulnerabilities grew at a rate of under 8% per year from 2005 to 2015 and then grew at more than 20% per year from 2016 to 2021. And there are more sophisticated actors looking for more ways to find and exploit those vulnerabilities.
- *The possibility of a major breach "cascade."* The NotPetya attack showed how destructive an attack could be for the organizations it hit, and SolarWinds showed how widespread an attack could be. But even in 2013, breaches such as the one at Target mentioned earlier show how routine and mundane connections between organizations can be an attack vector. Organizations such as 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. The kinds 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 be the fastest growing risk.

Our naive and obvious hypothesis? Attack surface and breach are correlated. 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, comments such as the Lloyd's of London insurers' are not yet so dated and cannot be dismissed as alarmist. Even with the giant breaches at Target, Anthem, and Sony behind us, we believe we haven't seen "The Big One" yet.

The Cyber Threat Response

In order to respond to competition and pandemics, organizations have had to be even more aggressive about adopting online solutions. The public wants to shop online, order meals online, track deliveries online, and more. Businesses and schools have had to allow for more remote work forcing more meetings through services such as Zoom, Webex, and Teams.

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.

Are organizations at least paying attention to and prioritizing these risks? There are several surveys that can answer this. Nearly every major consulting firm produces annual reports on risks, based on surveys of their executive-level clients. These reports are meant to cover all risks, not just cybersecurity, but cybersecurity has had a growing presence in these reports. Their methods of selecting executives, questions the surveys ask, and analysis of results vary, but they tell a similar story: cybersecurity gets a growing share of attention. In some cases, cybersecurity rates more concern among leadership than any other risk.

McKinsey, the largest and oldest of the major consulting firms, produces annual reports on all types of risks facing an organization. In the 2010 McKinsey report on enterprise risks, cybersecurity was not mentioned once. In 2016, cybersecurity was mentioned but received the least coverage in terms of mentions than any other risk. In the 2021 "McKinsey on Risk" report, cybersecurity risk was mentioned more often than all other risks including finance, regulation, geopolitics, competition, or even COVID.

PWC's 25th Annual Global CEO Survey, published in 2022, asked CEOs which types of threats to growth concerned them most. The number - one most cited threat was cybersecurity, with 49% of CEOs offering that response. This ranked ahead of geopolitical conflict, macroeconomic volatility, and health risks in the year of Russia's invasion of Ukraine and after two years of COVID.

Generally, executive-level attention on cybersecurity risks has increased, and attention is followed by resources. The boardroom is beginning to ask questions such as "Will we be breached?" or "Are we better than this other

recently breached company in our industry?" 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 in the 1990s, but for most of the time since then growth in CISOs was slow. *CFO Magazine* acknowledged that hiring a CISO as recently as 2008 would have been considered "superfluous."[17] By the time the first edition of this book was written (2016) the CISO role was becoming much more common. Now, nearly all Fortune 500 companies have a CISO or a similar VP, SVP, or C-level person dedicated to cybersecurity.[18]

Firms have also been showing a willingness, perhaps more slowly than cybersecurity professionals would like, to allocate serious resources to solve this problem. Cybersecurity spending per employee more than quadrupled from 2012 to 2020 to $2,691 (even adjusted for inflation).[19] According to cyberseek.org, in August 2022, the US labor market reached a milestone of *one million total workers in cybersecurity*.

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 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 whether 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 information and event management." More recently there is "threat intelligence" as well as "threat management." While all are within the tactical security solution space, 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, surveys conducted by Hubbard Decision Research in 2016 found that about 80% of organizations concerned with cybersecurity will resort to some sort of "scoring" method. 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.

Just over 50% of respondents plot risks on a two-dimensional matrix. In this 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.

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). 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 such as Oracle, Microsoft, and Adobe rate their vulnerabilities using a NIST-supported scoring system called the "Common Vulnerability Scoring System" (CVSS). 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,

			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 (aka Heat Map or Risk Map)

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 including the risk matrix. 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 beneficial to some degree. 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. *(emphasis added)*
>
> —OWASP[20]

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 later chapters):

- 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 we will discuss more—Tony Cox—goes as far as to say they can be "worse than random."
- 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 an improvement over unaided expert intuition.

- The improvement over unaided expert intuition is measurable even when the inputs are partially or entirely subjective—as long as the inputs are *unambiguous* quantities from *calibrated* experts.
- 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.

It's not all bad news. We'll show later in this book that cyber risk quantification (CRQ) is catching on and is supported by a growing number of tools. As we will show in later chapters, even simple quantification methods will be better than the still widely used risk matrix.

A Proposal for Cybersecurity Risk Management

In this book, we will propose a different direction for cybersecurity. Every proposed solution will ultimately align with 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.* As we mentioned already, we know that some methods measurably outperform others, even when comparing two purely subjective methods. The approaches that are currently the most popular use methods that are among the worst performing. 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.
- *These improved methods are entirely feasible.* We know this because it has already been done. 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 in many different industries. Nearly all of the analysts we trained in these methods had no quantitative risk analysis background (i.e., statisticians, actuaries, etc.). This is not just theory.

- *Even these improved methods can be improved further—continuously.* 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.

With improved methods, the cybersecurity professional can effectively determine a kind of "return on control." We can evaluate whether a given defense strategy is a better use of resources than another. In short, we can measure and monetize risk and risk reduction. To get there, we just need a "how to" book for professionals in charge of allocating limited resources to addressing ever-increasing cyber threats and leveraging those resources for optimal risk reduction.

This book is separated into three parts. Part I will introduce a simple quantitative method that requires a 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.

To get started, 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.

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.

Notes

1. Voltaire, *Poème sur le désastre de Lisbonne* (Poem about the Lisbon disaster), 1759.
2. 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/.
3. Sue Poremba, "2014 Cyber Security News Was Dominated by the Sony Hack Scandal and Retail Data Breaches," *Forbes Magazine,* December 31, 2014.
4. Kevin Haley, "The 2014 Internet Security Threat Report: Year of the Mega Data Breach," *Forbes Magazine,* July 24, 2014.
5. Matthew Heller, "Lloyd's Insurer Says Cyber Risks Too Big to Cover," CFO.com, February 6, 2015.
6. Risk Based Security, "2021 Year End Data Breach QuickView Report," February 2022.
7. National Association of Insurance Commissioners, Memorandum to Property and Casualty Insurance Committee, Report on the Cybersecurity Insurance Market, October 20, 2021.
8. Andrea, Vittorio. "Merck's $1.4 Billion Insurance Win Splits Cyber from 'Act of War,'" *Bloomberg Law*, January 19, 2022.
9. Daphne, Zhang. "Lloyd's Cyber Insurance Tweaks Stir Coverage Restriction Concern," *Bloomberg Law*, August 26, 2022.
10. Stephen Northcutt, "The Attack Surface Problem," SANS.edu. January 7, 2011, www.sans.edu/research/security-laboratory/article/did-attack-surface.
11. Pratyusa K. Manadhata and Jeannette M. Wing, "An Attack Surface Metric," *IEEE Transactions on Software Engineering* 37, no. 3 (2010): 371–386.

12. 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.
13. DataReportal.com, *Digital 2022 April Global Statshot*, April 21, 2022.
14. Jim Bird and Jim Manico, "Attack Surface Analysis Cheat Sheet," OWASP .org, July 18, 2015.
15. Mayumi, Brewster. "Annual Retail Trade Survey Shows Impact of Online Shopping on Retail Sales During Covid-19 Pandemic," United States Census Bureau, April 27, 2022.
16. "Top Cybersecurity Statistics, Trends, and Facts," *CSO Online,* October 2021.
17. Alissa Ponchione, "CISOs: The CFOs of IT," *CFO*, November 7, 2013
18. "List of Fortune 500 Chief Information Security Officers," *Cybercrime Magazine*, 2022.
19. Deloitte, "Reshaping the Cybersecurity Landscape," July 2020.
20. OWASP, "OWASP Risk Rating Methodology," last modified September 3, 2015.