EXPERT INSIGHT

Cybersecurity Threats, Malware Trends, and Strategies

Mitigate exploits, malware, phishing, and other social engineering attacks

Tim Rains

Packt

Cybersecurity Threats, Malware Trends, and Strategies

Mitigate exploits, malware, phishing, and other social engineering attacks

Tim Rains

Packt>
BIRMINGHAM - MUMBAI

Cybersecurity Threats, Malware Trends, and Strategies

Producer: Tushar Gupta
Acquisition Editor – Peer reviews: Divya Mudaliar
Content Development Editor: James Robinson-Prior
Technical Editor: Karan Sonawane
Project Editor: Janice Gonsalves
Copy Editor: Safis Editing
Proofreader: Safis Editing
Indexer: Rekha Nair
Presentation Designer: Sandip Tadge

First published: May 2020

Production reference: 1280520

Published by Packt Publishing Ltd.
Livery Place
35 Livery Street
Birmingham B3 2PB, UK.

ISBN 978-1-80020-601-4

www.packt.com

packt.com

Subscribe to our online digital library for full access to over 7,000 books and videos, as well as industry leading tools to help you plan your personal development and advance your career. For more information, please visit our website.

Why subscribe?

- Spend less time learning and more time coding with practical eBooks and Videos from over 4,000 industry professionals
- Learn better with Skill Plans built especially for you
- Get a free eBook or video every month
- Fully searchable for easy access to vital information
- Copy and paste, print, and bookmark content

Did you know that Packt offers eBook versions of every book published, with PDF and ePub files available? You can upgrade to the eBook version at www.Packt.com and as a print book customer, you are entitled to a discount on the eBook copy. Get in touch with us at customercare@packtpub.com for more details.

At www.Packt.com, you can also read a collection of free technical articles, sign up for a range of free newsletters, and receive exclusive discounts and offers on Packt books and eBooks.

Contributors

About the author

Tim Rains worked at Microsoft for the better part of two decades where he held a number of roles, including Global Chief Security Advisor, Director of Security, Identity and Enterprise Mobility, Director of Trustworthy Computing, and was a founding technical leader of Microsoft's customer facing Security Incident Response team.

Currently, Tim works at Amazon Web Services as Regional Leader, Security and Compliance Business Acceleration for Europe, the Middle East and Africa.

Tim lives with his wife Brenda and their two sons Tristan and Liam in London, England.

I'd like to thank my wife Brenda for her encouragement, assistance, and patience, without which this book would not have been written. Yuri Diogenes, thank you for your encouragement. Thank you Dr. ir Johannes Drooghaag for being our Technical Reviewer and to the entire team at Packt.

About the reviewer

Dr. ir Johannes Drooghaag is an Executive Consultant, Trainer, Author and Thought Leader for Cybersecurity, Sustainability and Privacy. After 30 years in leading roles at international corporations, Dr. ir Johannes Drooghaag founded his consultancy under his own name, which is based on the *Spearhead Management Model* he developed.

In his work, he focuses on building bridges between human potential and technological capabilities. This means teaching highly skilled technical leaders the soft skills of leadership and the mindset of leaders and growing the technical savviness of business leaders to be able to make the right decisions about risks and mitigations.

Dr. ir Johannes Drooghaag advocates for kids to be involved in Cybersecurity and Privacy Awareness as early as possible with free programs like *Internet Safety for Kids (and their Parents)*. His workshop *Cyber Security for Road Warriors (and Couch Potatoes)* focuses on frequent travelers and remote workers and has made thousands of people aware of the cyber risks surrounding them and the steps they can take to avoid becoming a victim.

Table of Contents

Preface

Imagine you are in a submarine, submerged miles below the surface surrounded by dark, freezing water. The hull of the submarine is under constant immense pressure from all directions. A single mistake in the design, construction or operation of the submarine spells disaster for it and its entire crew.

This is analogous to the challenge **Chief Information Security Officers** (**CISOs**) and their teams face today. Their organizations are surrounded on the internet by attackers that are constantly probing for ways to penetrate and compromise their organization's IT infrastructure. The people in their organizations receive wave after wave of social engineering attacks designed to trick them into making poor trust decisions that will undermine the controls that their security teams have implemented. The specters of ransomware and data breaches continue to haunt CISOs, **Chief Information Officers** (**CIOs**) and **Chief Technology Officers** (**CTOs**) of the most sophisticated organizations in the world.

After conducting hundreds of incident response investigations and publishing thousands of pages of threat intelligence, I have had the opportunity to learn from and advise literally thousands of businesses and public sector organizations all over the world. I wrote this book to share some of the insights and lessons I've learned during this extraordinary journey.

The views and opinions expressed in this book are my own and not those of my past or present employers.

Who this book is for?

This book is for CISOs, aspiring CISOs, senior managers in the office of the CISO, CIOs, CTOs and other roles who have meaningful responsibility for the cybersecurity of their organizations.

What this book covers

Chapter 1, *Ingredients for a Successful Cybersecurity Strategy*, provides a detailed look at the ingredients that are necessary for a successful cybersecurity program.

Chapter 2, *Using Vulnerability Trends to Reduce Risk and Costs*, provides a unique
20-year view of vulnerabilities, using vulnerability disclosure data from the National Vulnerability Database. This will help the reader more accurately evaluate the efficacy of cybersecurity strategies discussed in later chapters.

Chapter 3, *The Evolution of the Threat Landscape – Malware*, provides a unique data-driven perspective of how malware has evolved around the world over a 10 year period. This helps the reader understand the types of malware threats they face and which malware threats are most, and least, prevalent.

Chapter 4, *Internet-Based Threats*, examines some of the way's attackers have been using the internet and how these methods have evolved over time. This chapter dives into phishing attacks, drive-by download attacks and malware hosting sites.

Chapter 5, *Cybersecurity Strategies*, discusses the major cybersecurity strategies employed in the industry for the past 20 years or so. This chapter introduces the Cybersecurity Fundamentals Scoring System, which enables the reader to estimate an efficacy score for any cybersecurity strategy.

Chapter 6, *Strategy Implementation*, provides an example of how one of the best cybersecurity strategies identified can be implemented. This chapter illustrates how an Attack-Centric Strategy, namely the Intrusion Kill Chain, can be implemented.

Chapter 7, *Measuring Performance and Effectiveness*, looks at the challenge that CISOs and security teams have always had and how to measure the effectiveness of their cybersecurity program. This chapter examines how to measure the performance and effectiveness of a cybersecurity strategy.

Chapter 8, *The Cloud – A Modern Approach to Security and Compliance*, provides an overview of how the cloud is a great cybersecurity talent amplifier. This chapter looks at how the cloud can mitigate the ways enterprises typically get compromised. Additionally, this chapter dives into how security teams can use encryption and key management to protect data in the cloud.

To get the most out of this book

- You'll already understand basic **Information Technology** (**IT**) concepts and have some experience using, implementing, and/or operating IT systems and applications.
- Experience managing enterprise IT and/or cybersecurity teams will be helpful, but is not strictly required.
- You'll bring a healthy appetite to learn about some of the aspects of cybersecurity that you might not have been exposed to in the past.

Conventions used

The following conventions are used in the book:

A block of code is set as follows:

```
{
    "eventVersion": "1.05",
    "userIdentity": {
        "type": "AssumedRole",
        "principalId": "Example:user123",
        "arn": "arn:aws:sts::Example:assumed-role/Admin/user123",
        "accountId": "Example-ID",
}
```

Bold: Indicates a new term or an important word.

Download the color images

We also provide a PDF file that has color images of the screenshots/diagrams used in this book. You can download it here: `https://static.packt-cdn.com/downloads/9781800206014_ColorImages.pdf`.

Get in touch

Feedback from our readers is always welcome.

General feedback: Email `feedback@packtpub.com`, and mention the book's title in the subject of your message. If you have questions about any aspect of this book, please email us at `questions@packtpub.com`.

Errata: Although we have taken every care to ensure the accuracy of our content, mistakes do happen. If you have found a mistake in this book we would be grateful if you would report this to us. Please visit, `http://www.packtpub.com/submit-errata`, selecting your book, clicking on the Errata Submission Form link, and entering the details.

Piracy: If you come across any illegal copies of our works in any form on the Internet, we would be grateful if you would provide us with the location address or website name. Please contact us at `copyright@packtpub.com` with a link to the material.

If you are interested in becoming an author: If there is a topic that you have expertise in and you are interested in either writing or contributing to a book, please visit .

Reviews

Please leave a review. Once you have read and used this book, why not leave a review on the site that you purchased it from? Potential readers can then see and use your unbiased opinion to make purchase decisions, we at Packt can understand what you think about our products, and our authors can see your feedback on their book. Thank you!

For more information about Packt, please visit `packtpub.com`.

1

Ingredients for a Successful Cybersecurity Strategy

There's no doubt that enterprises today, more than ever, need effective cybersecurity strategies. However a sound strategy is not in and of itself a guarantee of success. There are several ingredients that are necessary for a cybersecurity program to be successful. This chapter will describe what a cybersecurity strategy looks like and each of the necessary ingredients for success in detail.

Throughout this chapter, we'll cover the following topics:

- Defining the term *cybersecurity strategy*
- Common ways organizations become compromised, and how the mitigation of these are fundamental to effective cybersecurity
- Understanding the difference between an attacker's motivation and their tactics
- Additional guidance on formulating a successful cybersecurity strategy

Let's begin with a fundamental question that we'll need to answer before discussing cybersecurity strategy in any detail: what do we actually mean when we say "cybersecurity strategy"?

What is a cybersecurity strategy?

Organizations that have a super-strong security culture, essentially have cybersecurity baked into them. For everyone else, there's strategy. In my experience, the terms "strategy" and "tactics" are poorly understood in the business world. One person's strategy is another person's tactics. I once worked with a Corporate Vice President who would tell me that I was talking about tactics when I was explaining our strategy. Throughout my career, I've been in meetings where people have talked past each other because one person is discussing strategies and the other is discussing tactics.

Additionally, security and compliance professionals sometimes use the term "strategy" when they are referring to frameworks, models, or standards. There are lots of these in the industry and many organizations use them. For example, ISO standards, NIST standards, OWASP Top 10, CIS Benchmarks, STRIDE, risk management frameworks, SOC 2, PCI, HIPAA, the Cloud Security Alliance Cloud Controls Matrix, the AWS Cloud Adoption Framework Security Perspective, AWS Well-Architected Security Pillar, and many more. All of these can be helpful tools for organizations seeking to improve their security postures, comply with regulations, and demonstrate that they meet industry standards.

I'm not proposing a new dictionary definition of the term "strategy," but I do want to explain what I mean when I'm discussing cybersecurity strategies in this book. In my view, there are at least two critical inputs to a cybersecurity strategy:

1. Each organization's high-value assets
2. The specific requirements, threats, and risks that apply to each organization, informed by the industry they are in, the place(s) in the world where they do business, and the people associated with each organization

High Value Assets (**HVAs**) are also known as "crown jewels." There are many definitions for these terms. But when I use them, I mean the organization will fail or be severely disrupted if the asset's confidentiality, integrity, or availability is compromised. HVAs are rarely the computers that the organization's information workers use. Yet I've seen so many organizations focus on the security of desktop systems as if they were HVAs.Given the importance of HVAs, it would be easy to focus on them to the exclusion of lower-value assets. But keep in mind that attackers often use lower-value assets as an entry point to attack HVAs. For example, those old development and test environments that were never decommissioned properly, typically, aren't HVAs. But they are often found to be a source of compromise.

One of the first things a CISO needs to do when they get the job is to identify the organization's HVAs. This might be more challenging than it sounds as the crown jewels might not be obvious to people that don't possess expertise specifically related to the business they are supporting. Interviewing members of the C-suite and members of the board of directors can help to identify assets that would truly cause the business to fail or be severely disrupted.

Working backward from the organization's objectives can also help identify its HVAs. As CISOs do this analysis, they should be prepared for some nuances that weren't initially obvious. For example, could the business still meet its objectives without power, water, heating, air conditioning, and life-safety systems? Depending on the business and the type of building(s) it uses, if elevators aren't available, is there any point letting employees and customers through the front door? Customers might be willing to walk up a few flights of stairs, but would they be willing to walk up 40 flights of stairs if that was necessary? Probably not.

If this disruption was sustained for days, weeks, or months, how long could the business survive? Where are the control systems for these functions? And when was the last time the security posture of these systems was assessed? Identifying an organization's HVAs doesn't mean that CISOs can ignore everything else. Understanding which assets are truly HVAs and which aren't helps CISOs prioritize their limited resources and focus on avoiding extinction events for the organization.

Once the CISO has identified their organization's crown jewels, the next step is to ensure that the C-suite and board of directors understand and agree with that list. This clarity will be very helpful when the time comes to request more resources or different resources than the organization has leveraged in the past. When the organization needs to make hard decisions about reductions in resources, clarity around HVAs will help make risk-based decisions. The time and effort spent getting the senior stakeholder community on the same page will make the CISO's life easier moving forward.

The second critical input to a cybersecurity strategy is the specific requirements, threats, and risks that apply to the organization, informed by the industry they are in, the place(s) in the world where they do business, and the people associated with it. This input helps further scope the requirements of the cybersecurity program. For example, the industry and/or location where they do business might have regulatory compliance requirements that they need to observe, or they could face stiff fines or get their business license revoked. Keep in mind that most organizations can't identify all possible threats and risks to them. That would require omniscience and is a natural limitation of a risk-based approach.

After publishing thousands of pages of threat intelligence when I worked at Microsoft (Microsoft Corporation, 2007-2016), I can tell you that there are global threats that have the potential to impact everyone, but there are also industry-specific threats and regional threats. Using credible threat intelligence to inform the strategy will help CISOs prioritize capabilities and controls, which is especially helpful if they don't have unlimited resources. Trying to protect everything as if it's of the same value to the organization is a recipe for failure. CISOs have to make trade-offs, and it's better if they do this knowing the specific threats that really apply to the industry and region of the world where they do business. This doesn't mean CISOs can ignore all other threats, but identifying the highest-risk threats to their organization's crown jewels will help them focus resources in the most important places.

I have dedicated three chapters in this book to help you understand the threat landscape and how it has evolved over the last 20 years. *Chapter 2*, *Using Vulnerability Trends to Reduce Risk and Costs*, dives deep into vulnerability management and will show you how vulnerability disclosures have trended over the past two decades. *Chapter 3*, *The Evolution of the Threat Landscape – Malware*, focuses on how malware has evolved over the last 20 years. *Chapter 4*, *Internet-Based Threats*, examines internet-based threats that every organization should seek to mitigate.

Without the two inputs I've described here, CISOs are left implementing "best practices" and industry standards that are based on someone else's threat model. Again, these can be helpful in moving organizations in the right direction, but they typically aren't based on the HVAs of individual organizations and the specific threats they care about. Using best practices and industry standards that aren't informed by these two inputs will make it more likely that there will be critical gaps.

At this point, you might be wondering what a cybersecurity strategy looks like. The following diagram represents a cybersecurity strategy. HVAs are central and are supported by the other parts of the strategy. The cybersecurity fundamentals include the foundational capabilities that support a successful security program, such as vulnerability management and identity management, among others.

Advanced cybersecurity capabilities are investments that organizations should make as they become very proficient at the fundamentals. If your organization isn't really good at the fundamentals, then don't bother investing in advanced cybersecurity capabilities, as attackers won't need to do anything "advanced" to successfully compromise the environment and subvert those advanced capabilities.

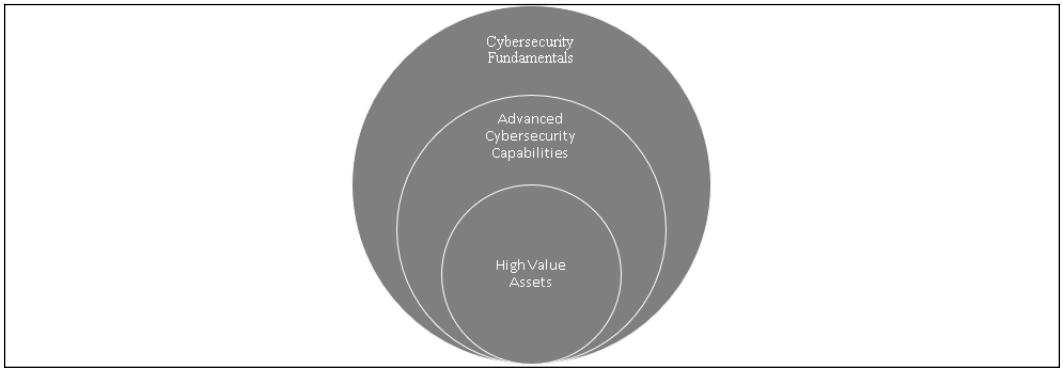

Figure 1.1: An illustrative example of a cybersecurity strategy

Now that we have a good idea of what cybersecurity strategy entails, let's examine what I consider to be a critical ingredient of cybersecurity strategies: the common ways that organizations are compromised.

How organizations get initially compromised and the cybersecurity fundamentals

The foundation of the strategy is what I call the "cybersecurity fundamentals." A solid foundation is required for a successful strategy. The cybersecurity fundamentals are based on the threat intelligence I mentioned earlier. After performing hundreds of incident response investigations and studying Microsoft's threat intelligence for over a decade, I can tell you with confidence that there are only five ways that organizations get *initially* compromised. After the initial compromise, there are many, many **tactics, techniques, and procedures** (**TTPs**) that attackers can use to move laterally, steal credentials, compromise infrastructure, remain persistent, steal information, and destroy data and infrastructure. Some of these have been around for decades and some are new and novel.

The five ways that organizations get initially compromised are what I call the "cybersecurity usual suspects":

1. Unpatched vulnerabilities
2. Security misconfigurations
3. Weak, leaked, and stolen credentials
4. Social engineering
5. Insider threats

The cybersecurity fundamentals are the part of the strategy that focuses on mitigating the cybersecurity usual suspects. Let's look at each one of these in more detail, starting with the exploitation of unpatched vulnerabilities.

Unpatched vulnerabilities

A vulnerability is a flaw in software or hardware design and/or the underlying programming code that allows an attacker to make the affected system do something that wasn't intended. The most severe vulnerabilities allow attackers to take complete control of the affected system, running arbitrary code of their choice. Less severe vulnerabilities lead to systems disclosing data in ways that weren't intended or denying service to legitimate users. In *Chapter 2*, *Using Vulnerability Trends to Reduce Risk and Costs*, I provide a deep dive into vulnerability management and some of the key vulnerability disclosure trends over the past 20 years. I'll save that in-depth discussion for the next chapter, but I will provide some more context here.

Attackers have been using vulnerabilities to compromise systems at scale since at least the days of Code Red and Nimda in 2001. In 2003, SQL Slammer and MSBlaster successfully disrupted the internet and compromised hundreds of thousands of systems worldwide by exploiting unpatched vulnerabilities in Microsoft Windows operating systems. In the years following these attacks, a cottage industry developed an ongoing effort to help enterprise organizations, those with the most complex environments, inventory their IT systems, identify vulnerabilities in them, deploy mitigations for vulnerabilities, and patch them. At the end of 2019, there were over 122,000 vulnerabilities disclosed in software and hardware products from across the industry, on record, in the National Vulnerability Database (National Vulnerability Database, n.d.). As you'll read in *Chapter 2*, *Using Vulnerability Trends to Reduce Risk and Costs*, the number of vulnerabilities disclosed across the industry surged between 2016 and 2020, reaching levels never seen before.

An economy has evolved around the supply and demand for vulnerabilities and exploits, with a varied list of participants including vendors, attackers, defenders, various commercial entities, governments, and others. The number of participants in this economy and their relative sophistication make it harder for organizations to protect themselves from the exploitation of vulnerabilities in their IT environment by pressurizing the associated risks. Using unpatched vulnerabilities are a mainstay of attackers' toolkits.

Organizations that are highly efficient and competent at vulnerability management make it much harder for attackers to successfully attack them.

A well-run vulnerability management program is a fundamental component and a critical requirement of a cybersecurity strategy. Without it, organizations' cybersecurity efforts will fail regardless of the other investments they make. It's important enough to reiterate this point. Unpatched vulnerabilities in operating systems, and the underlying platform components that advanced cybersecurity capabilities rely on, enable attackers to completely undermine the effectiveness of these investments. Failing to efficiently address ongoing vulnerability disclosures in the "trusted computing base" that your systems rely on renders it untrustworthy.

An accurate inventory of all IT assets is critical for a vulnerability management program. Organizations that can't perform accurate and timely inventories of all their IT assets, scan all IT assets for vulnerabilities, and efficiently mitigate and/or patch those vulnerabilities, shouldn't bother making other investments until this is addressed. If your organization falls into this category, please reread the preface section of this book and recall the submarine analogy I introduced. If the CISO and vulnerability management program managers rely on their organization's IT group or other internal partners to provide IT asset inventories, those inventories need to be complete – not just inventories of systems they want to comply with.

Assets that don't show up in inventories won't get scanned or patched and will become the weak link in the security chain you are trying to create. Very often, this is at odds with the uptime objectives that IT organizations are measured against, because patching vulnerabilities increases the number of system reboots and, subsequently, decreases uptime even if everything goes smoothly. My advice in scenarios where asset inventories are provided by parties other than the vulnerability management program itself is to trust but verify. Spend the extra effort and budget to continually check asset inventories against reality. This includes those official and unofficial development and test environments that have been responsible for so many breaches in the industry over the years.

If the sources of asset inventories resist this requirement or fail to provide accurate, timely inventories, this represents the type of risk that the board of directors should be informed of. Providing them with a view of the estimated percentage of total asset inventory currently not managed by your vulnerability management program should result in the sources of asset inventories reprioritizing their work and the disruption of a dangerous status quo. I'll discuss vulnerability management in more detail in *Chapter 2*, *Using Vulnerability Trends to Reduce Risk and Costs*, of this book. I'll also discuss vulnerability management in *Chapter 8*, *The Cloud – A Modern Approach to Security and Compliance*, on cloud computing.

The cloud can render the old-fashioned methods of inventorying, scanning, and patching security vulnerabilities obsolete.

Of course, one challenge with the approach I just described is environments that have embraced **Bring Your Own Device** (**BYOD**) policies that allow information workers to use their personal mobile devices to access and process enterprise data. The underlying question is whether enterprise vulnerability management teams should inventory and manage personal devices? This debate is one reason why many security professionals originally dubbed BYOD as "Bring Your Own Disaster." Different organizations take different approaches when answering this question. Some organizations give employees corporate-owned and fully managed mobile devices, while others require personal devices to enroll in enterprise mobile device management programs. I've also seen a more passive management model, where users are required to have a access pin on their devices and aren't allowed to connect to their employers' networks if the latest mobile operating system version isn't installed on their devices. Some organizations use **Network Access Control** (**NAC**) or **Network Access Protection** (**NAP**) technologies to help enforce policies related to the health of systems connecting to their network. Minimizing the number of unpatched systems allowed to connect to enterprise networks is a best practice, but can be challenging to accomplish depending on corporate cultures and mobile device policies. Collecting data that helps security teams understand the risk that mobile devices pose to their environments is very helpful for a rationalized risk-based approach.

Next, we'll consider security misconfigurations. Like unpatched vulnerabilities, security misconfigurations can potentially enable attackers to take a range of actions on a system including disrupting its operation, stealing information, lowering security settings or disabling security features, seizing control of it, and using it to attack other systems.

Security misconfigurations

Security misconfigurations can be present in a system as the default setting, like a preset key or password that is the same on every system manufactured by a vendor. Security misconfigurations can also be introduced gradually as a system's configuration changes incrementally as it's managed over time.

After performing hundreds of incident response investigations while I was on the customer-facing incident response team at Microsoft, I can tell you that a significant percentage of systems get initially compromised through security misconfigurations.

This is especially true of internet-facing systems such as web servers, firewalls, and other systems found in enterprise **demilitarized zones** (**DMZs**). Once a misconfiguration enables an attacker to control a system in a DMZ or use it to send authenticated commands on the attacker's behalf (such as a server-side request forgery attack), the attacker aspires to use the system to gain access to other systems in the DMZ and ultimately get access to systems inside the internal firewall of the organization. This has been a common pattern in attackers' playbooks for 20 years or more.

Security misconfigurations have also plagued endpoint devices, such as PCs, smartphones, and **Internet of Things** (**IoT**) devices. The infrastructures that these endpoints connect to, such as wireless access points, are also frequently probed by attackers for common misconfigurations. Security misconfigurations have also been an issue in **industrial control systems** (**ICS**). For example, one scenario with ICS that has burned security teams in the past is "fall back to last known status," which can override more recent security configuration changes in favor of former, less secure settings. Hardcoded credentials and vulnerable default configurations have long haunted manufacturers of all sorts of software and hardware across the industry.

A well-run vulnerability management program typically includes identifying security misconfigurations as part of its scope. Many of the same vulnerability scanners and tools that are used to identify and patch security vulnerabilities are also capable of identifying security misconfigurations and providing guidance on how to address them. Again, organizations should forego big investments in advanced cybersecurity capabilities if they aren't already very proficient at identifying and mitigating security misconfigurations in their environment. There's no point in spending a bunch of money and effort looking for the **advanced persistent threat** (**APT**) in an environment if attackers can use decades-old lists of hardcoded passwords, which are available on the internet, to successfully compromise and move around the environment. Even if CISOs found such attackers in their IT environment, they would be powerless to exorcise them with unmanaged common security misconfigurations present.

Some of the biggest breaches in history were a result of an initial compromise through a combination of unpatched vulnerabilities and security misconfigurations. Both can be managed through a well-run vulnerability management program. This is a non-optional discipline in any cybersecurity strategy that should be resourced accordingly. Don't forget, you can't manage what you don't measure; complete, accurate, and timely IT asset inventories are critical for vulnerability management programs. Trust but verify asset inventories, always. It's worth keeping in mind that the cloud provides several advantages over the old on-premises IT world. I'll discuss this in detail in *Chapter 8*, *The Cloud – A Modern Approach to Security and Compliance*, in this book.

Security misconfigurations can be present by default with new hardware and software, or can creep in over time. Another ongoing threat that requires constant attention is that of compromised credentials. Organizations must constantly and proactively work to mitigate this threat vector.

Weak, leaked, and stolen credentials

Compromised IT environments due to weak, leaked, or stolen credentials are common. There are several ways that credentials get leaked and stolen, including social engineering such as phishing, malware that does keystroke logging or steals credentials from operating systems and browsers, and compromised systems that cache, store, and/or process credentials. Sometimes, developers put projects on publicly available code-sharing sites that have secrets such as keys and passwords forgotten in the code. Old development and test environments that are abandoned but still running will ultimately yield credentials to attackers after not being patched over time.

Massive lists of stolen and leaked credentials have been discovered on the internet over the years. In addition to these lists, the availability of high-performance computing clusters and GPU-based password cracking tools have rendered passwords, by themselves, ineffective to protect resources and accounts. Once passwords have been leaked or stolen, they can be potentially leveraged for unauthorized access to systems, in "reuse" attacks and for privilege escalation. The usefulness of passwords, by themselves, to protect enterprise resources has long passed. Subsequently, using **multi-factor authentication** (**MFA**) is a requirement for enterprises and consumers alike. Using MFA can mitigate stolen and leaked credentials in many, but not all, scenarios. Using MFA, even if attackers possess a valid username and password for an account, they won't get access to the account if attackers don't also possess the other factors required for authentication. Other factors that can be used for authentication include digital certificates, one-time passwords and pins generated on dedicated hardware or a smartphone app, a call to a preregistered landline or mobile phone, and more.

MFA isn't a silver bullet for weak, leaked, or stolen passwords, but it's super helpful in many scenarios. There have been some successful attacks on some MFA methods. For example, SIM-swapping attacks to intercept pin codes sent to preregister mobile phones via SMS. Another real limitation of MFA is that it isn't ubiquitous in enterprise IT environments. Organizations with decades of legacy applications that use old-fashioned authentication and authorization methods are less likely to fully mitigate the risk with MFA. Even if the latest systems and cloud-based services require MFA, chances are there are more legacy applications that cannot utilize it easily.

A picture of an iceberg comes to mind here. Several CISOs that I've talked to have experienced this limitation firsthand during penetration tests that exposed the limitations of MFA in their environments. Still, MFA should be widely adopted as it successfully mitigates many attack scenarios where weak, leaked, and stolen passwords are involved. It should be required for new systems being adopted and the risks posed by the old systems without it should be carefully considered and mitigated where possible. There are several vendors that specialize in such mitigations.

When an on-premises enterprise environment is initially compromised, attackers use leaked or stolen credentials to perform reconnaissance and to look for other credentials that have been cached in the environment. They are especially on the lookout for administrator credentials that could give them unlimited access to resources in the compromised environment. Typically, within seconds of the initial compromise, attackers try to access the victim organization's user account directory service, such as Microsoft **Active Directory** (**AD**), to dump all the credentials in the directory. The more credentials they can use to move and stay persistent, the harder it will be to expel them from the environment – they can persist indefinitely. Attackers will try to steal user account databases. If attackers successfully get all the credentials from their directory service, then recovery really is aspirational.

Once attackers have stolen hashed credentials, the weakest of these credentials can be cracked in offline attacks in a matter of hours. The longer, uncommon, and truly complex passwords will get cracked last. There have been raging debates for decades about the efficacy of passwords versus passphrases, as well as appropriate character lengths, character sets, password lockout policies, password expiration policies, and the like. Guidance for passwords has changed over the years as threats and risks have changed and new data has become available. Some of the people I worked with on Microsoft's Identity Protection team published password guidance based on the data from 10 million credential attacks per day that they see on their enterprise and consumer identity systems. "Microsoft Password Guidance" (Hicock, 2016) is recommended reading.

When credentials are leaked or stolen from an organization, it doesn't take attackers long to run them through scripts that try to log in to financial institutions, e-commerce sites, social networking sites, and other sites in the hopes that the credentials were reused somewhere. Reusing passwords across accounts is a terrible practice. Simply put, credentials that provide access to more than one account have a higher ROI for attackers than those that don't. Sets of compromised credentials that can provide access to corporate resources and information, as well as social networks that can also serve as a rich source of information and potential victims, are valuable.

Using unique passwords for every account and using MFA everywhere can mitigate this risk. If you have too many accounts to assign unique passwords to, then use a password vault to make life easier. There are numerous commercially available products for consumers and enterprises.

Identity has always been the hardest part of cybersecurity. Identity governance and management deserves its own book. I offer a very incomplete list of recommendations to help manage the risk of weak, leaked, and stolen credentials:

- MFA can be very effective – use it everywhere you can. Microsoft published a great blog post about the effectiveness of MFA called "Your Pa$$word Doesn't Matter" (Weinert, 2019) that is recommend reading.
- You should know if your organization is leaking credentials and how old those leaked credentials are. Using a service that collects leaked and stolen credentials, and looks for your organization's credentials being sold and traded online, can give you a little peace of mind that you aren't missing something obvious. Getting some idea as to the age of these credentials can help decide if password resets are necessary and the number of people potentially impacted.
- Privileged Access Management solutions can detect pass-the-hash, pass-the-ticket, and Golden Ticket attacks, as well as attackers' lateral movement and reconnaissance in your infrastructure:
    - Many of these solutions also offer password vaulting, credential brokering, and specialized analytics. Some of these solutions can be noisy and prone to false positives, but still, they can help you to manage and detect weak, leaked, and stolen credentials.
- In cloud-based environments, identity and access management (IAM) controls are the most powerful controls you have. Taking advantage of all the power that IAM controls offer can help you to protect and detect resources in the cloud. But this is one control set area that can proliferate into an unmanageable mess quickly. Extra thoughtful planning around your organization's IAM strategy will pay huge security dividends.

I will discuss identity a little more in *Chapter 5*, *Cybersecurity Strategies* of this book.

An important aspect of protecting credentials involves educating information workers within an organization to be aware of social engineering attacks in which attackers may attempt to steal credentials through methods such as phishing. This is not the only way in which social engineering is used to compromise systems, however. We'll cover social engineering in a little more detail next.

Social engineering

Of the cybersecurity usual suspects, social engineering is the most widely used method. Simply put, social engineering is tricking users into making poor trust decisions. Examples of poor trust decisions include lowering the security posture of a system by changing its settings without understanding the possible outcomes of doing so or installing malware on a system. Attackers rely on the naivety of their victims in social engineering attacks.

The volume of social engineering attacks is orders of magnitudes larger than other types of attacks. For example, the volume of email phishing attacks Microsoft reported for July 2019 was 0.85% of the more than 470 billion email messages that flowed through Office 365 that month (Microsoft Corporation, n.d.). That's 4 billion phishing emails that all relied on social engineering, detected in a single month. Similarly, Trojans, a category of malware that relies on social engineering to be successful, has been the most prevalent category of malware in the world continuously for the last decade. I'll discuss this category of malware and many others, in detail, in *Chapter 3*, *The Evolution of the Threat Landscape – Malware*.

Given the massive volume of social engineering attacks, and their historical record of success, mitigating these attacks really isn't optional for enterprises. A fundamental component of an enterprise cybersecurity strategy is a mitigation strategy for social engineering attacks. Put another way, not including social engineering attacks in your cybersecurity strategy would mean ignoring the top way that organizations get initially compromised by volume.

Social engineering attacks are typically perpetrated by attackers external to organizations, to which users must be prepared through appropriate education and training. Another challenging threat to defend against is one from within. The final potential route of compromise, which we'll discuss next, is that of the insider threat.

Insider threats

When discussing insider threats with CISOs and security teams, I find it useful to break them down into three different categories, listed here from most likely to least likely:

1. Users and administrators that make mistakes or poor trust decisions that lead to bad security outcomes.
2. The lone wolf insider or a very small group of individuals that use their privileged access to steal information or otherwise negatively impact the confidentiality, integrity, or availability of the organization's information technology and/or data.

3. The mass conspiracy where multiple insiders work together to overcome the separation of duties that distributes the span of security control. I've found that enterprises typically bring this category up in discussions about risks in managed service provider environments and the cloud.

Mitigating insider threats is an important aspect of cybersecurity and is something that should be fundamental to any enterprise-wide strategy. Enforcing meaningful separation of duties and embracing the principle of least privilege are helpful, as are monitoring and auditing.

I became a big fan of deception technology after seeing how it can be used to mitigate insider threats. There are a few different approaches to deception technology, but the basic concept is to present attackers with a system, potentially with publicly known vulnerabilities or common security misconfigurations that, when interacted with, alerts defenders to the presence of attackers. This approach can help alert defenders to the presence of external attackers and insider threats. I've heard some security professionals refer to it as a "canary in the coal mine" for IT environments. Implementing deception technology with as few people involved as possible and keeping the program confidential can be helpful in exposing at least two of the three categories of insider threats that I have outlined.

Those are the five ways organizations get initially compromised. Defending against these five vectors of attack is fundamental to effective cybersecurity.

Focus on the cybersecurity fundamentals

To have a successful cybersecurity program, organizations need to get very good at continuously mitigating all five of these types of threats. This competency forms the foundation of a sound cybersecurity strategy. Other cybersecurity-related investments will potentially have diminishing returns if the foundation of the strategy is not solid.

After an attacker uses one or more of these five ways to initially compromise an organization, then they might employ a plethora of novel and advanced TTPs. Organizations that focus on the cybersecurity fundamentals make it much harder for attackers to be successful; that is, by focusing on the inside 85% of the bell curve below which the cybersecurity fundamentals sit, instead of the activities in the outlying 7.5% on either end of the curve, security teams will be much more successful. Unfortunately, the allure of hunting advanced persistent threats can take resources away from the less sexy, but critical, work in the middle of the curve.