

EBOOK

— The Needed —

BREA KTHR OUGH

in **Cybersecurity**

Why—despite continued investment in cyberdefenses—catastrophic data breaches are more prevalent than ever and how control systems alongside new technologies driven by network metadata can turn this dismal state around.



About the authors



Ricardo Villadiego
Founder & CEO

Ricardo Villadiego (RV) is a seasoned entrepreneur and visionary technology leader focused on cybersecurity. His last 20 years have been spent in the quest of solving some of the most prevalent cybersecurity challenges organizations face. RV founded Easy Solutions, a global organization focused on the prevention and detection of electronic fraud. Subsequently, RV led the cybersecurity business unit at Cyxtera Technologies, where he developed a long-term vision and execution plan. His passion for technology and cybersecurity have triggered yet another venture, and he created Lumu Technologies with a clear objective: help organizations detect compromises at speed.



Fernando Cuervo
Director of Product Growth

Lumu Technologies' Director of Product Growth Fernando Cuervo has more than 10 years of experience in cybersecurity and product management, having served in different roles such as SOC Manager, IT Manager, Product Owner, and others. Fernando has also presented cybersecurity talks internationally and actively shares his professional views on blogs and news websites. He attained a degree in systems engineering from the Central University of Bogotá and holds further accredited certifications in ITIL, Scrum Product Ownership, and Strategic Management.



Claudio Deiro
Senior Technical Product Architect

Claudio graduated from the Polytechnic University of Turin with a degree in electronic engineering. Describing himself as 'unapologetically an engineer', his career led him towards the cybersecurity industry, where he served as Technical Product Manager and Technical Product Architect until he joined Lumu Technologies as Senior Technical Product Architect. Claudio is also the co-founder of Fundación S, a non-profit organization that helps Colombians to develop the necessary skills and aptitudes necessary to succeed in the software industry.

Table of contents

04 Introduction

05 Why We Need a Breakthrough

- 6 The State of Cybersecurity
- 8 How Did We Get Here?
- 9 The Needed Breakthrough

12 Limitations in Cybersecurity Testing

- 13 The Genesis of Testing
- 14 An Example of Testing Done Well
- 16 Why Traditional Testing is Hardly Enough
- 18 The Evolution of Security Testing

21 The Path to Continuous Compromise Assessment

- 22 The Country of the Blind
- 22 Cybersecurity as a Control System
- 23 Compromise as a Disease
- 23 A Viable Path to ContinuousCompromise Assessment

29 Improving Decision Making in Cybersecurity

- 31 How People Make Decisions
- 32 Complexity in CybersecurityDecision Making
- 33 The Prevalence of Poor Decision Making in Cybersecurity
- 35 The Psychology of Error: Biases in Our Perception of Security
- 36 Overcoming Biases
- 37 The Cost of Insecurity
- 39 Better Data Makes Better Decisions

40 If Not Now, Then When?

- 41 References

Introduction

I've committed the last 20 years of my life to solving some of the most pressing cybersecurity challenges faced by enterprises. In the last 5 years or so, one question has been at the forefront of my mind: Why are we not continuously and intentionally looking for compromises? Why is it that the one thing that truly matters to every cybersecurity defense is the one thing that we are not measuring?

Cybersecurity is complex. Success in complex scenarios relies on the system's ability to regulate and recover from disturbances. The word 'cybernetic' stems from the Greek word 'kybernetes' meaning 'steersman'. As a boat veers to one side, the steersman automatically corrects the rudder, maintaining a steady course. Cybersecurity might be more complex than a simple boat flowing down a river, but how could we have lost this ability?

In practice, self-regulation is done via closed-loop systems or "error-controlled systems" defining error as the state of compromise for a particular cybersecurity incident. The faster the industry moves towards developing the necessary cybersecurity capabilities that help an organization assess its continuous status of compromise, the faster that cyber-resilience will be achieved. With small but deliberate changes to the cyber-security architecture, the disparity between the cyber incident and the detection of the breach can be dramatically shortened.

This is the breakthrough my colleagues and I passionately believe we need to achieve.

— Ricardo Villadiego


Why We Need a Breakthrough

Part 01

The State of Cybersecurity

The cybersecurity industry is sizzling. Organizations continue to commit increasing budgets to their cybersecurity efforts. Between 2015 and 2019, enterprises deployed an astonishing \$670 Billion, according to a Forbes cybersecurity roundup¹.

Yet, during the same period, the number of security breaches increased exponentially, and the amount of exposed data resulted in a crisis of global scale. According to the U.S. Identity Theft Resource Center², the number of breaches grew from 1 257 in 2018, an already frightening number, to 1473 in 2019.



Number of U.S. data breach

Year	Banking/ Credit/ Financial	Business	Educational	Gov/military	Medical	Total
2013	35	194	54	60	271	614
2014	38	263	57	91	332	781
2015	71	312	58	63	275	779
2016	51	497	997	72	373	1090
2017	134	907	128	79	384	1632
2018	135	575	78	100	369	1257
2019	108	644	113	83	525	1473

The same study reflects that the problem is widely spread across all industries.

It is untrue that organizations that adhere to the most stringent regulatory standards—such as in the banking sector—perform better than others that are less regulated, or that industries that invest heavily in cybersecurity are less breached. It is probably time to accept that investment does not necessarily translate to protection.

For years, we have been conditioned to define success in terms of the investment of time and money. In cybersecurity, this well-proven formula is not producing the results we should expect to see from an industry that has such a high level of investment.

That universal formula (Success = Time + Money) has worked in most aspects of life, from sports to sending a man to the moon. That same formula is producing impressive results in the health field. In August of 2019, the WHO and the National Institute of Allergies and Infectious Diseases announced a cure for Ebola, and significant progress has been made towards an HIV vaccine, a disease that meant death about 20 years ago.

However, the cybersecurity outlook is disappointing, and the cyber-war may be lost. Another indication of this is an iconic brand like Capital One announcing being a victim of a massive breach. How we arrived to this point deserves exploration.

How Did We Get Here?

There are four main drivers that led the industry to its current state of compromise and uncertainty:

- a. **Ever-evolving threats** generate an infinite number of vulnerabilities that enterprises must attempt to defend. Cybersecurity technologies continue to be mostly reactive which leads to a vicious “cyber cycle” of attackers scanning networks, developing exploits and attacking systems, with defenders detecting attacks, analyzing exploits and patching such systems.
- b. **Unlimited capital** flowing into the industry is fueling defense vendors that fall into the “detect, then mitigate” approach. The results are technologies in-market that are not ready for primetime, inherently unstable and becoming obsolete as soon as deployments are completed without ever testing if they delivered on their promise.
- c. As a result of a.) and b.), cyber-defense architectures have **grown in complexity**, stacking an avalanche of vendors that neglected management and monitoring capabilities, hence adding little incremental protection to the system. The complexity and cost associated with it create a false sense of security, especially at higher levels in the organization.
- d. Today's society is psychologically wired to find instant gratification. This notion is translated to problem-solving as the pursuit of **the magical solution** (“the silver bullet”). This behavior and the inability to embrace the idea of being breached led practitioners and decision-makers to accept the current framework of innovation in cybersecurity: detect then mitigate.

To make matters worse, many of today's cybersecurity solutions and architectures work as an open-loop system at their core. This means systems do not take into account the redeeming features of closed-loop systems, in which the ideal output (in this case the state of no compromise) is measured continuously to make sure that changes are applied to the system (the cybersecurity architecture).

It's impossible to obtain different results doing more of the same. In order to break the cyber cycle, cybersecurity needs to make a fundamental shift towards applying control theory to continuously measure the value of reference. For a given organization, it must be “no compromise.” Any deviation from the reference value should be promptly identified and mitigated by adjusting the cyber defense architecture.

The Needed Breakthrough

The global cybersecurity crisis is a problem that must be solved, or significantly improved in the short term. VCs are rightfully continuing to deploy capital in the industry. With a threat landscape that can and will evolve infinitely, addressing the problem is ever more critical.

The focus must be placed in building cybersecurity capabilities that fundamentally disrupt the current state of cybersecurity. We need to rethink our security paradigm from the long-standing one of trying to keep adversaries out of our networks. Organizations have to assume that cybercriminals are already inside. This is known in government circles as “Assumption of Breach.” Deborah Hayden of the NSA’s Information Assurance Directorate has said as much back in December, 2010³.

The industry lacks a factual process that provides certainty around cyber incidents, which is one of the two drivers for making the right cybersecurity decisions. At Lumu, we call this process **Continuous Compromise Assessment**.

To better understand this concept, it is necessary to first revise the Cyber Kill Chain⁴, which is a model for the identification and prevention of cyber intrusions activity. The model identifies what adversaries must complete in order to achieve the objective. The following is a simplified graph of the process.



A closer look at the different stages among the multiple variations of the Cyber Kill Chain unveils the common denominator that enables adversaries evil intent: **network access**. Network traffic is ground zero for illuminating threats. Almost all threats must first be downloaded and then communicate back to its C&C to provide any value to attackers.

The ability to collect network traffic to illuminate threats may be the **feedback loop** that many cybersecurity and academic researchers have been envisioning for over a decade. Even with the advances in bandwidth and storage, collecting network traffic for a large organization might be cost-prohibitive. The problem now evolves into how to collect signals of network traffic in a way that accurately represents the summary of the “conversations” within an organization.

In his book *Secrets and Lies*, Bruce Schneier formulates “that often the patterns

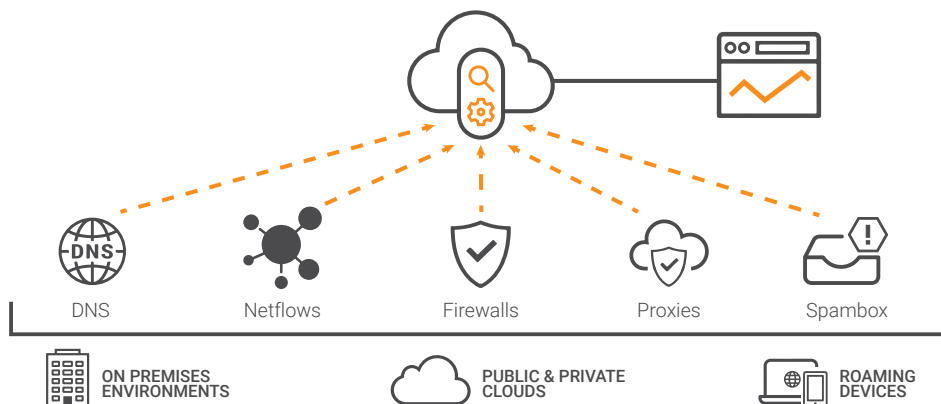
of communications are just as important as the content of the communication.” For example, the simple fact that Alice telephones a known terrorist every week is more important than the details of their conversation. Putting this together with the steps associated with the Cyber Kill Chain, we can quickly realize that the process of compromising a device and a network will make that device and network behave differently. Here are a few steps to illustrate the process:

- » The end-user who the adversary is targeting will point his or her device to a new **host**.
- » If the attack is successful, the device will attempt to connect with the adversary’s infrastructure (**C&C**) seeking instructions and/or exfiltrating information.
- » In more sophisticated attacks, the adversary will need to escalate privileges and, in order to do so, the compromised device will attempt communications with adjacent devices and/or high-value targets within the now compromised organization. This is a clear sign of **lateral movement**.
- » As the adversary conquers new victims, more devices will attempt to connect with the adversary’s infrastructure.

Further analysis of the described steps among many others facilitated the key elements of metadata, from required network traffic to an accurate representation the summary of conversations within an organization, as described in the following table:

Network Metadata	Why it Matters
DNS Queries	Collecting DNS Queries provides context into the attempt of connections from the organization’s devices towards adversarial infrastructure.
Network Flows	Among other malicious behavior, network flows provide insights into an organization’s devices that are controlled by the adversaries and attempt to move laterally.
Access Logs of Perimeter Proxies or Firewalls	In cases where the attacks avoid domain resolution, the traces of adversarial contact will lie in the access log of firewalls or proxies, depending on the organization’s network configuration.
Spambox	Email is the preferred method by attackers to deliver exploits to the organization’s end-users ⁶ . Analyzing the organization’s spambox provides insights into the type of attacks an organization is receiving, but more importantly if end-users are accessing such attacks and the organization is at a high risk of compromise.

Signaling traffic in this form instead of doing a full packet capture is optimal, as it represents only a tiny fraction of the total network traffic. Yet it's still possible to identify the compromise level of an organization.



Specific techniques have been developed to facilitate the data collection process while minimizing friction in the multiple environments that define a network nowadays.

The remaining problem to solve is how to make it a continuous process. Collecting and processing these signals for a specific timespan is feasible, but it is challenging. Organizations can quickly become disenchanted due to the level of complexity in data collection and processing, even using tools that promise to handle at least some of these key signals, like SIEMs or network flow collectors.

To solve this last piece, a reliable, accurate and continuous process is required from collection to Illumination as shown in the following image.



Only once the continuous process is implemented can we say that the feedback loop has been built and this can be considered the breakthrough for cybersecurity in modern days. A continuous compromise assessment process will not only simplify the decision-making process for managers and practitioners but will also entirely change the dynamics of the cybersecurity ecosystem and the cyber cycle of attackers versus defenders.

Limitations in Cybersecurity Testing

Part 02

The Genesis of Testing

Our ability to form a question that can later become a hypothesis has been an important enabler of human evolution. Parting from a moment of curiosity, humans long to understand what occurs when certain conditions are met. As such, simple and complex theories have been tested repeatedly throughout history, becoming the beginning of the scientific method as we know it today.

Rome was not built in a day, and neither are the methodologies used to test hypotheses. This has taken us several hundreds of years to get to where we are today. After a rather muted period in scientific advancement during the Dark Ages and Renaissance periods, the world experienced a time of incredible discoveries. Many European scholars became exposed: Aristotle, the greatest thinker of psychology, politics and ethics, prompted the well-known *Dialogue Concerning the Two Chief World Systems*⁵. Ptolemy, astronomer and geographer known for being the first to consider the Earth as the center of the universe; Euclid, known for the way space, time and shapes are conceived. Arguably, one of the most influential personas of the period was Francis Bacon, a lawyer and philosopher who was the first to formalize the concept of a true scientific method. The outcomes of his research were heavily influenced by the great minds of Nicolaus Copernicus (1473-1543) and Galileo Galilei (1564-1642) who invented the telescope, and used it to study the sun and planets.

The accomplishments of the period were the necessary setup for a radical revolution in science. The work of Isaac Newton (1642-1727) in mathematics, integral and differential calculus, and astronomy resulted in the definition of laws of motion. Newton's work marked the beginning of a new world enabled by the proper use of modern science. The success of the scientific method was further demonstrated by the creation of the cell theory⁶, made possible by the invention of the microscope by Antoni Van Leeuwenhoek (1632-1723). This discovery served as a milestone for science in general, as it exposed the hidden world that exists beyond the limits of human vision. As such, Matthias Jakob Schleiden (1804-1881) and Zoologist Theodor Schwann (1810-1882) concluded that both all plants and all animals are composed of cells. In 1858, Rudolf Virchow (1821-1902) expanded the work of Schleiden and Schwann by proposing that all living cells must rise from pre-existing cells. These emblematic discoveries were followed by a handful of scientists that further explained how the universe functioned, including Pasteur, Einstein or Hawking.

Curiosity is the single most important element that has moved science forward. The drive to understand what lies beyond the human eye. The most important discoveries of the world have only been possible through testing a wide array of hypotheses. As such, testing is the biggest enabler of the modern world. The luxuries of life, including

electricity, transportation, manufacturing, and the Internet itself are possible by relentless testing. But, what is the process that allows us to systematically take a simple question to a viable conclusion. The most simple scientific method consists of six basic steps:

1. **Make an observation:** The majority of scientific inquiry starts with an observation that piques curiosity.
2. **Ask a question:** The purpose of the question is to narrow the focus of the inquiry, to identify the problem in specific terms.
3. **Form a Hypothesis:** Suggest a possible answer in the form of a Hypothesis. A hypothesis is stated as an "if-then" statement
4. **Conduct an experiment:** Set up to test a specific hypothesis which must be controlled.
5. **Analyze Data:** Collect quantitative and qualitative data. On that information you can find evidence to support or reject the hypothesis.
6. **Iterate:** Do it again with the new information.

There are several real world applications of testing that go beyond academic objectives. The world's most powerful armies are a perfect example, who have made use of it in preparation for war. Military groups have continuously needed to test their strategies to simulate, evaluate, perfect their defenses, and confirm these were designed and executed effectively.

The need for testing has stood the test of time from generation to generation. However, not all testing is created equal. Testing has the power to build industries, and take them from small to dominant empires. Unfortunately, testing also has the power of sabotaging industries when it is not performed diligently.

An Example of Testing Done Well


Now that we have settled on the criticality of testing, it is important to understand why testing is an ongoing quest for perfection proximity. Testing will point out errors in controllable and uncontrollable variables. Instead of pursuing perfection, our duty is to reach an acceptable threshold of error. Even with these acceptable thresholds, the different elements involved may very well fail. The airline industry is a good example of how precise testing can make an industry mostly predictable, and ultimately, successful. Because the effects of mechanical or process errors can be catastrophic,

this industry is heavily focused on testing. Therefore, this has made flying the safest method of transportation.

According to The Economist, the world saw an important decrease in the number of airplane accidents after 1972, an increase in the 80s, and a steady drop after the 90s, reaching a perfect year in 2017 with zero accidents or fatalities. Recognizing that airplane travel has become more accessible, this is an important accomplishment for the industry.

Safer skies

Global passenger flights, number of accidents and fatalities



The reasons for the decrease are numbered below:

- » Aviation accidents not only cost money, but also cost lives. This means, there is a sense of urgency and high pressure to solve any given problem.
- » An airplane is a rather complex machine. However, the aviation industry is mostly streamlined. There are a few large vendors that manufacture aircrafts, which ensure that all the parts and procedures are standardized.

- » All technology is developed with a clear purpose. After an accident, there are mandatory detailed investigations that show the root cause. This leads to the development of technology that helps ensure the same accident does not happen again. This is why the art of flying has been nearly perfected. The industry itself has pioneered an open and collaborative format to show incidents and accidents and receive support from its community.

The entire industry is constantly trying to improve and learn from mistakes that lead to accidents. A key component to their success is their open community, where the interested parties transparently share detailed insights, accept their own errors and take actions for continuous improvement.

Aviation is the perfect example where there is a clear motivation and urgency to solve any and all problems. This is an industry that has developed the tools and methodologies to measure and continuously lower their margin of error, demonstrating the benefits of testing when done correctly. Several other industries must learn from the airplane industry. Cybersecurity is no exception.

Why Traditional Testing is Hardly Enough

Security testing today has two big branches: penetration testing and vulnerability assessment. The first one tries to break the enterprise's defenses, and the second one shows an organization's level of exposure indirectly, demonstrating the known vulnerabilities. They focus on testing for risks outside the corporate network. How organizations organize and execute them is highly dependable on their resources. As such, these vary greatly across verticals. There are serious effects to the lack of standardized methodologies, processes, industry collaboration, and transparency across the industry

Nonetheless, there are some popular methodologies largely used by the industry:

Test Basis	Test Type
Whitebox testing: Full information about the target is shared with the testers.	Vulnerability identification in software: Must give feedback to developers on coding practices
Blackbox testing: No information is shared with the testers about the target.	Scenario to identify vulnerabilities: The tester explores particular scenarios to find whether it leads to a vulnerability in your infrastructure.
	Scenario to test detection and response: The goal here is to measure the detection and response capabilities of the organization.

Penetration testing and vulnerability assessment's primary goal was to test networks, which is only part of the problem. Alone, they are insufficient, being preventive techniques. Prevention is rendered useless once compromise takes place. These tools have fallen short of expectations for the following reasons:

1. False Hypothesis: As we saw testing a hypothesis is the foundation of scientific method, but it is important to select the correct hypothesis to test. The hypothesis of cybersecurity in general is to test defenses assuming that we are secure, but what happens if the adversary is already inside?
2. Incomplete: Traditional security testing is incomplete by nature because it only tests defenses and finds vulnerabilities (outside), neglecting the true state of compromise (inside).
3. Limited View: Traditional security testing was designed to show an image of vulnerabilities of critical assets on a specific date, but systems, configurations and threats change on a daily basis.
4. Relies on the weakest link: The industry assumes that the attacker gets inside the networks exploiting vulnerabilities but the truth is that it is easy to send an email to compromise an organization. We rely on people to not be in a compromising position which is unrealistic.

The purpose of penetration testing is to simulate whether an attacker can pass an organization's defenses. These tests are not conclusive. Often, pentesters have different abilities when compared to the attacker as well as less time to perform each test. In the case of vulnerability assessments, the single purpose is to measure the exposure of a company. These tests are performed based on hypothetical scenarios of potential exploitation vectors from the attacker. This means, organizations can learn about the potential of attacks instead of confirmed attacks.

An additional shortcoming of pentests is their focus exclusively on the critical assets. As an industry, we have a misconception that all attacks occur on servers and databases. The truth is that most attacks start with malicious email targeting the employees, meant to compromise a device and move laterally until higher value assets are found. Lastly, organizations often still rely on legacy systems difficult or impossible to upgrade. Such systems may be exposing them to a wide range of vulnerabilities, that penetration testing may detect. However, it may not be possible to take action due to the legacy nature of the systems.

The reasons exposed above are not meant to discourage penetration testing and vulnerability assessments. However, the industry as a whole has set unrealistic expectations on these tools, that exceed what they were designed to do. They are certainly not where testing should end. In fact, the way the industry executes testing is useful, but hardly enough.

The greatest evidence that we are missing a part on our security strategy is that there are breaches making headlines weekly, even in companies that devote a lot of resources and comply with pentesting and vulnerability assessment regulatory requirements.

The Evolution of Security Testing

Testing has a critical role in the evolution of today’s digitally-driven, ever-connected world. The life humans carry out is only possible because of the perfection of the testing discipline. Everything that humans see, touch and experience, requires an enormous amount of testing before it becomes a reality. Because the more one improves the testing practice, the better the outcome, testing has the power of building up an entire industry. Unfortunately, the opposite is also true.

Cybersecurity is yet to perfect the art of testing for the greater good of the industry. We have already analyzed why testing in cyber is severely flawed but it is worth pondering on the fact that organizations rarely find vulnerabilities that they are not purposely looking for and determined to find. It is likely that compromises are not found because they are not actively being looked for. This fact certainly merits the question: **Why are organizations not diligently and systematically determined to identify compromises?**

Recent incidents demonstrate that adversaries have remained inside of enterprise networks for long periods of time, going absolutely undetected, even after the execution of multiple pentests and vulnerability assessments.

Test Basis	Test Type
Citrix	10 Years
Marriott	4 Years
Yahoo	Several Breaches. Months
Equifax	6 Months

The most overlooked component of cybersecurity testing is the hypothesis that networks are compromised. The focus is placed on the erroneous assumption that organizations are secured, and no compromise exists. There is a famous quote by Steve Denn that depicts the described situation: “One can never make the same mistake twice, because the second time, it is not a mistake, it is a choice.” As an industry, the mistaken assumption should no longer be considered an error, but instead an action that is taken deliberately.

As we adjust the standard scientific method to fit the needs of the cybersecurity industry, the steps below must be followed:

1. **Make an observation:** The cybersecurity industry is underperforming. Data breaches continue to grow in scale and sophistication, despite investment surges.
2. **Ask a question:** Why do breaches continue to happen?
3. **Form a Hypothesis:** If we evolve the security testing methods then the breaches will decrease.
4. **Conduct an experiment:** Analyze network data to find compromises. Check if this information provides additional value when compared to traditional security testing.
5. **Analyze Data:** Analyze how those findings improve or not an organization's stance against risk and whether it is improving its cyber-resilience.
6. **Iterate:** Develop a "rinse and repeat" culture. Repeat the process with additional information

It is imperative to urge organizations to assume their networks are compromised, and work tirelessly to prove otherwise. **This is the most critical hypothesis that security practitioners must be continuously testing.** The most critical step is to admit that organizations must challenge the status quo and evolve their thinking if different results are desired.

Here are some points to start this evolution:

1. **Mindset change.** There is a need to internalize the assumption of being compromised and focus on proving that it is not the case. Keep in mind that traditional security testing is needed but not enough. If organizations want to be cyber-resilient, **measuring compromise** is not negotiable.
2. **Unlock the value of the organization's own metadata.** Organizations are sitting on a gold mine: their own network data. Pro Tip: DNS queries are possibly the most valuable information for compromise detection and only a few companies are using it.
3. **Engage in Continuous Compromise Assessment.** Once the data is unlocked, the visibility and intelligence can be extremely valuable. A continuous compromise assessment program can ensure that compromise is detected dynamically and in real-time. Measure effectiveness. Start with a baseline and take deliberate actions to eliminate the compromise, and increase your cyber-resilience.

The single purpose of most of the cyber defense strategies is to avoid being compromised. Yet, this is useless if a compromise happens, and the function of detecting and measuring compromise is absolutely neglected. Continuous compromise detection has become a necessity. Organizations that can unlock what hides under their own data will become empowered to perfect their defense strategies. For this reason, the feedback loop between defenses implemented and compromise detection must be closed. The diagram below explains it in detail:

Compromise detection complements existing testing and vulnerability tools, and helps companies evolve and perfect their own testing practice. Today, the cybersecurity industry faces a solid opportunity: to arm organizations with the right knowledge on compromised levels through the implementation of tried and true testing methodologies.

The Path to Continuous Compromise Assessment

Part 03


The Country of the Blind

An important contributor to the cybersecurity problem outlined above is the limited visibility on the state of compromise of our networks. As such, organizations need to diligently find ways to develop enough visibility to understand what goes on within their networks. There is an emphasis on building ever-better walls to signal a boundary not to be crossed, and to give time to guns to intervene. There is a missing element, one that it is perceived as obvious: eyes. Eyes to document the infractions and guide the guns.

In contemporary cybersecurity we have an abundance of walls, but we are almost blind and we have almost no guns. While the scarceness of means to efficaciously retaliate is an issue beyond the reach of technical solutions alone, involving international treaties, nation states goodwill, and the like, we—as cybersecurity professionals—have the ability, and the duty, to provide our customers the tools to timely identify compromises and breaches.

Cybersecurity as a Control System

An alternative point of view is to look at cybersecurity's problem as a control system [4].



For the system to work—that is to say to keep the compromise level within acceptable limits and prevent serious damage—a precondition is the existence of sufficiently precise and timely feedback. Nowadays the feedback is often the message of a security researcher that notifies the victim that its data was found on the dark web. It might be a uranium centrifuge unexpectedly failing: too late. Alternatively, it may be a stream of thousands of alerts each day triggered by heuristic rules: too much.

Compromise as a Disease

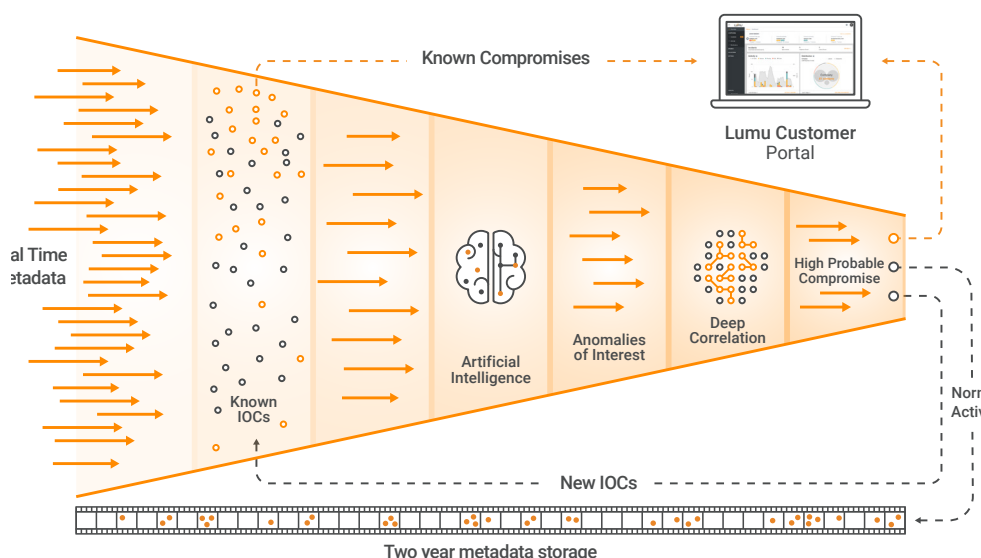
To use a healthcare metaphor, compromise is like an infection. Firewalls and EDRs—the current staples of cybersecurity—are preventative measures. Now, “no amount of prevention will help you when prevention fails”⁷, and therefore you become sick. To fight back the infection we need diagnostic tools and antibiotics.

In cyber the cure can be easy—format a machine, change access credentials, strengthen firewall rules, run cleanup tools—to diagnose the infection is not. Our machines and our networks are a mess, much like our bodies. Much like their real life counterparts, computer viruses have learned to mutate, making the traditional signature-based approach to detection reactive and very much dependent on the agility of the provider. Thousands of processes, on thousands of machines, interchange thousands of packets every second. Finding the malicious threat is like finding the proverbial needle in the haystack. Behind a single IP address in the cloud there can be hundreds of applications. Some of them may be malicious, or infected themselves. But how do we find out?

There is simply no way for the unaided human eye to make sense of all this noise.

This is also the reason why a whitelist approach—only permitting what is known to be safe—won’t work outside of very special cases. The whitelist would eventually outgrow any management capacity, and what at one moment is known to be safe may become infected, and infectious, a moment later.

A Viable Path to Continuous Compromise Assessment




Organizations often overlook the power of their own network metadata. Today, this is the most promising path for transformational improvement in the world of cybersecurity. This goldmine contains incredible potential, as long as it is used correctly. The process outlined below exemplifies how to best leverage network traffic:

1. For starters, metadata is collected in real time to later be contrasted against a large pool of known, certified IoCs, coming from the organization itself as well as private and public sources of curated detailed cyber threat intelligence. Alerts are generated upon the identification of matching data.
2. Once alerts are generated, organizations have a clear indication that prompt action is required.
3. All incoming metadata should then be put through artificial intelligence and heuristic inference engines that would allow us to understand anomalous behavior, in order to reduce false negative rates. For example, unusual traffic patterns generated by an asset coming in contact with points within the network that are out of the ordinary and/or with a certain frequency. What results from this filtering process is a list of anomalies of interest, which may represent compromise.
4. Anomalies of interest should be later put through a deep correlation process which consists of taking the traffic deemed as likely to be related with malicious actors and confirming its compromise nature. For example, cybercriminals often use the same group of IP addresses or a specific segment in the network, as well as the same domains in rotation. The deep correlation step generates only alerts of high-probability.
5. The last step in this process is to store the residual network metadata for a period of time and leverage emerging IoCs for further correlation and analysis. This step is critical because it will enable organizations to constantly improve their Continued Compromise Assessment process.

The implementation of Continuous Compromise Assessment can have a transformative effect on the cybersecurity industry. The most natural question at this stage is 'Why has this not been done before?' First, we did not know the impact of the interconnected world and the extent of malicious attacks. Knowing what we know now, Continuous Compromise Assessment will be an absolutely critical step in any enterprise's cybersecurity strategy. Secondly, this process is only possible as of very recently. About 10 years ago, carrying out this process would not have been effective or practical. There are a few reasons why, which are enumerated below:

1. Data Storage Cost & Computing Power

The cost of storing data has decreased 3,000 times in the last 20 years while computing power has increased 10,000 times since the year 2000⁸. These conditions have helped build the perfect scenario for the collection and administration of large volumes of metadata as well as the execution of deep learning capabilities.



2. The Cloud

There are good reasons to place processes and storage in a cloud computing environment. A critical factor for the success of a system of this kind is the time between when new intelligence is available and when such intelligence is incorporated into the system. With on-premise systems there will always be a delay in the distribution of such information.

On the other hand, with a cloud-based system, new intelligence will be available to all users as soon as it is available to the system.

The most important factor that makes a cloud deployment ideal for a system like this is the removal of all maintenance and management burdens on its users. The valuable time of skilled security professionals should no longer be spent on system maintenance, including monitoring disk space usage, or writing rules to catch the

latest infection. All these menial tasks are transferred to the cloud environment. Security professionals can concentrate on investigating and remediating incidents.

3. ML & AI Renaissance

Artificial intelligence was all the rage in the late '60s and early '70s. It then went out of fashion, for a decade (or four). But now, with immensely more powerful hardware and somehow curbed expectations, it is riding high again. Besides the fact that talking about AI is trending, the use of machine learning and anomaly detection for Continuous Compromise Assessment may bring real advantages

Today's machine learning algorithms are sufficiently well understood and can be supported by enough computing power to be successful in practical contexts.

Storage cost, computing power, combined with cloud, machine learning and AI, make this approach absolutely practical, effective and possible. As they say, the devil is in the details and a detailed look at how data is collected merits a discussion.

How to Leverage Your Metadata & Overcome Challenges Along the Way

Data collection: Do you really want to see it all? The ground truth is in the network data. Unlike logs, that can be tampered with or simply deleted, or EDRs (Endpoint Detection and Response), that have to play at the same level as the attacking software, there is no way for an attacker to interfere with packet capture and analysis. So a conscientious network administrator should capture and analyze everything. Right?

As the ones familiar with Betteridge's law of headlines⁹ have already guessed, the answer is no: it would cost too much. You would need roughly double the amount of bandwidth and computing power of the original network, simply to analyze all its data.

Fortunately, "traffic patterns reveal a lot about any organization and are much easier to collect than actual communication data"¹⁰. This means that a much more effective approach is possible. In this approach, what is collected and analyzed for the whole network is the metadata.

Please note that a partial approach, where only the critical systems are under control, would not suffice. Attackers inside your network would be able to move laterally and take control of less critical systems until they are in the position of reaching the resources they are after without raising suspicion. In case you are interested, reading the account, in her own words, of how Phineas Fisher hacked Hacking Team [9] can provide a good idea of the steps an attacker may take, from a peripheral firewall to an unsecured test database to a backup storage and finally—with some additional steps—to everything. It is therefore necessary to include in the analysis all network devices, including the ones considered less critical.

Metadata Collection and Consolidation

Collecting the metadata seems to be a rather trivial process. After all, most organizations most likely already generate and collect such metadata. However, there exist at least two problems with collecting metadata:

1. Metadata comes in different formats, making it difficult to collect, organize and consume.
2. Some organizations do not have processes in place for metadata collection.

The first problem can be overcome using de facto standards, like Cisco's NetFlow¹¹ and Elastic's Packetbeat [11]. We can address the second problem using a stack of existing software components that can be easily customized to fulfill the user's needs.

Metadata Analysis - The Trouble with Indicators of Compromise

When looking at an organization's network metadata a particular IP or domain is identified, it is easy to conclude that this network has been compromised. Well, not so fast. An approach to compromise detection based solely on Indicators of Compromise (IoCs) is bound to fail for at least three good reasons:

1. **Reactive approach:** IoCs should be identified, confirmed, and divulged. Looking for IoCs does not help the first victims of an attack, or the targets of customized attacks. An IoC-only approach will present a high false negative rate.
2. **Noise:** IoC lists suffer from a high noise level. Part of the reason is that they are often compiled in automated form. But the high reuse and sharing rates for network resources imply that just seeing an IP is most often not sufficient proof of a compromise. An IoC-only approach will present a high false positive rate.
3. **Lack of context:** Context is often necessary for the interpretation of the data. And without interpretation it is impossible to really understand what is going on and take the appropriate corrective actions.

Metadata Analysis - Anomaly Detection

The occurrence of a compromise will cause a change in the behavior of the network, possibly a very subtle change to escape detection. For example, a botnet agent will phone home to let the botmaster know a new bot is available. A worm will try to contact neighbors to infect as many machines as possible. A coin miner will contact the C&C to get new jobs and report results.

The emergence of anomalous behaviour gives us a chance to discover compromise, if we can appropriately learn how the network behaves and detect changes

Unfortunately, the network is ever-evolving. The behavior of a network can change for reasons as benign as the installation of new software, or a new version of already existing software. Or the deployment of a new web application. And possibly almost every user will need a unique combination of applications and will present a unique behavioral pattern.

Relying solely on anomaly detection would therefore generate an unreasonably high number of false alerts.

Metadata Analysis - Suspicious Behavior Detection

When analysis begins, a compromise could already be there, possibly in a fairly high number of machines. On the other hand, the changes introduced in the behaviour of the network can be so subtle, ever so gradual, that the anomaly detection system does not trigger.

Therefore, we need a set of models, heuristics and rules that can detect suspicious behavioral patterns.

Metadata Analysis - Deep Correlation

At this stage, we have a series of anomalies, such as a machine contacting a web server never seen before; and suspicious patterns, such as a series of machines constantly posting small amounts of data to an unknown external server. There is little we can say with this information. Maybe the website is little known, but showed up in the results of a particular search, and hosts relevant content. Maybe custom software is in use that is sending home some telemetry or diagnostic data.

Now, let's imagine that the little known web server is hosted by a known bulletproof hosting provider^{12 13}, and in other instances a visit to this host has been followed by malicious activity. Most would agree that it would be wise to take a look at that machine. If, instead, it turns out that the website was recently created by a reputable owner, the anomaly can probably be safely ignored. On the other hand, imagine that the same posting pattern is also observed in other unrelated users, and the receiving system is not known to be managed by a legitimate organization. Would you look into it?

The examples attempt to show how correlating metadata with available intelligence enables filtering out a significant part of the detected events, leaving for human investigation only events that with high probability are compromises. It also allows us to enrich the provided information with context, so that human analysis is easier

Metadata Storage - The Importance of Memory

After analysis, the metadata is stored for a substantial amount of time. This, while generating non-negligible costs, allows for forensic analysis and reanalysis. The importance of forensic analysis is pretty obvious.

The idea with reanalysis is that the metadata will be scrutinized again as new intelligence or algorithms become available. Threats that unfortunately escaped the first round will therefore be discovered, hopefully before significant damage occurs.

It is of course vital that the data is stored correctly, so that the needed information could be efficiently retrieved. The storing technology should also allow for a high degree of flexibility, as it may not be obvious what exactly will be needed in the future. Fortunately, the technological advances in big data treatment make it possible to meet these requirements.

Improving Decision Making in Cybersecurity


Part 04

How People Make Decisions

Human minds, though amazing in their own right, are unable to adequately encompass the complexity of real-world systems, according to cognitive psychologist Herbert A. Simon's concept of Bounded Irrationality¹⁴. We frequently resort to reasoning shortcuts and other mental biases that lead us to adopt 'satisfying' solutions, rather than optimal ones.

Event-based Decision Making

According to renowned author, John D. Sterman¹⁵ "Where the world is dynamic, evolving, and interconnected, we tend to make decisions using mental models that are static, narrow, and reductionist." Sterman goes on to state that we tend to interpret experience through a series of open-loop events, where problems resulting from disparities between our goals and situation, require decisions that lead to results. In this paradigm, the decision-maker acts as a 'satisficer' and does not consider the inter-connectivity and feedback from real-world, dynamic systems.



We have evolved to make decisions in this manner in order to continue functioning under stringent limitations. Such constraints include having limited time to make decisions, too much information to process, not enough meaning from the information, and fallible memories for retaining it all. Cognitive psychologist Daniel Kahneman calls this thinking fast¹⁶—when we are not able to think slowly. While these "short cuts" may lead to cognitive biases—more on that later—they are crucial for decision-making efficiency.

The average cybersecurity team could make hundreds of these reactive decisions in a day. Every alert is an opportunity for a decision and there simply isn't time to think 'slowly' at every juncture. The trick is knowing if the decision calls for some slow thinking.

Proactive Investment Strategies

In *Managerial Perspectives on Risk and Risk Taking*¹⁷, James G. March and Zur Shapira state that “in conventional decision theory formulations, choice involves a trade-off between risk and expected return.” Therefore, rational decision-makers invest in cybersecurity when the investment will yield a positive return, or rather when the cost of the investment is less than the potentially catastrophic loss it prevents. Indeed, the greatest responsibility of any modern CISO considers just this: how to invest budgets and resources in a way that most effectively reduces breaches and their consequences.

In any cybersecurity architecture, investments must be made into prevention, detection, and response systems, all of which have an influence on the other. This means that the feedback loop needs to be closed in order to measure the effects of the changes to the system. Traditionally this has been difficult to achieve since organizations have not had the ability to measure compromise as an output of the system.

There is a time and place for each of these decision making paradigms. The latter type of decision is strategic. They consider having the right tools and that enough resources are available when they are needed. The former is eminently tactical, covering how to employ such tools and resources to minimize damage. The important part is understanding when each type of decision making is called for.

Complexity in Cybersecurity Decision Making

In a research paper¹⁸, issued by the Cybersecurity Interdisciplinary Systems Laboratory at MIT Sloan, researchers attempted to determine why poor decision making was so prevalent in cybersecurity. The researchers ran a cybersecurity simulation game that mimicked the complex systems—including prevention, detection, and response—needed in a modern enterprise’s cybersecurity program. Players had to choose how to invest in these processes, in order to protect against attacks and ultimately protect their enterprise’s bottom line. Two groups of players were invited to play the simulation game. One group consisted of cybersecurity professionals, the other of inexperienced players.

The study found that both groups struggled in making effective decisions, but over multiple iterations, both groups managed to improve their scores. There were two major sources of complexity that needed to be overcome:

Uncertainty Concerning Cyber Incidents

The study considers that uncertainty surrounding the cost of an incident hampers decision making. In cases where deterrents are successful, it can be difficult to measure the cost of a hypothetical cyber incident. Operators may also underestimate the frequency of attacks. The combination of these two factors creates an impression of the “expected cost” of insecurity that does not bear resemblance to reality. Even in cases where security operators have a good grasp on the expected cost of a breach, biases might cause operators to act irrationally—see the section *The Psychology of Error: Biases in Our Perception of Security* for more on this topic.

Delays in Complex Systems

The study looks at how investments in prevention, detection, or response can take time to have observable effects once implemented. Additionally, each investment needs time for implementation, and operators need time for training and overcoming learning curves. In a reactive decision-making paradigm, the development of cybersecurity capabilities only after the detection of an attack, means the organization’s information systems will not properly recover in time and will remain vulnerable. A closed-loop decision-making process fares better, but the delays in feedback would mean that constant adjustment and measurement of the system would be needed to reach an optimal state.

The MIT paper showed that exposure to such large-scale breach events and their management improves overall cybersecurity decision making. However, waiting for such events to occur is a very expensive way to learn how to deal with them.

The Prevalence of Poor Decision Making in Cybersecurity

Equifax - a Compendium of Errors

The poster-child of data breaches’ first example of poor decision making was a lack of preventative maintenance. Hackers made use of a widely-known vulnerability (that had been reported only 3 days earlier) in their complaints portal to gain initial access. If the vulnerability had only been promptly patched, there would not have been a breach.

The attackers’ second move—moving laterally while escalating privileges—was also made easier by a lack of preventative measures. If Equifax had chosen to invest in the proper segmentation of systems, the attack would have been more easily limited to their customer complaint platform.

The attackers were able to have access to Equifax's databases for 76 days¹⁹. At that time, they had reportedly not renewed an encryption license. Therefore, the encrypted personal information of approximately half of all Americans was able to pass through their HTTPS interception without being inspected. Only when the encryption had been updated—ten months late—did full network visibility resume, and was the attack detected.

Once the attack was discovered, Equifax's response showed terrible event-based reactive decision making. They delayed publicizing the breach for a month, when transparency in such events is the best policy. During that time little was done in terms of mitigating its effect on the American people, although several executives sold stock in the company—one being convicted for insider trading.

Capital One

In early 2019, an attacker exposed a vulnerability in Capital One's cloud integration in order to steal the credentials from over 100 million credit applications. The attacker executed a Server Side Request Forgery²⁰ to trick a misconfigured web application firewall into relaying information including current credentials. This type of vulnerability had been known for years, but required specialized knowledge related to Amazon Web Services' Identity and Access Management as well as EC2 to identify and fix. Ultimately, a lack of investment in these in-demand cybersecurity skills led to a vulnerability that could have easily been avoided.

Marriott - the Breach that Lasted 4 Years

On November 30th, 2018, Marriott Hotels announced a breach²¹ that had been detected on September 8th. The breach affected the network of a chain of hotels—Starwood—that Marriott had purchased in 2016. It soon became apparent that Starwood had been breached in 2014 and remained compromised for 4 years. The attack exposed over 500 million customer records including passwords and credit card details. The breach was typical of a phishing attack that installed a Remote Access Trojan and a password sniffer in order to gain access and administrator privileges.

The most worrying aspect of the Marriott breach is that the compromise was allowed to persist for 4 years. This reveals that a key cybersecurity rule was not followed: assume you are compromised and prove otherwise²². It also highlighted the importance of IT and security due diligence in the event of mergers and acquisitions. As the proprietor of Starwood, Marriott laid off most of their corporate staff, including IT and security staff. The new reservation system was not ready to manage the hundreds of newly acquired hotels, so the old understaffed and malware-maligned system was allowed to continue serving customers until the breach was discovered two years later.

Marriot's response²³ to the breach caused further problems by using a wide range of email domains and websites, some of which lacked HTTPS certification. This led to a variety of phishing attacks imitating Marriott in the wake of the breach.

The Psychology of Error: Biases in Our Perception of Security

Security comes at a cost, whether it is in the form of a loss of money, convenience, or opportunities. For example, locking your front door means trading increased security at the cost of a minor inconvenience. We all have an instinctive understanding that a trade-off needs to be made. We have developed the ability to make these cost-benefit decisions quickly through cognitive biases: shortcuts that go around our limitations in time, memory, meaning, and dealing with excessive information. As Bruce Schneier says in his TEDtalk²⁴, "We are highly optimized for risk decisions that are endemic to small family groups in the East-African highlands in 100,000 BC." Our instincts inform our perception of security. Unfortunately, as the following examples illustrate, how we perceive security can differ greatly from its reality.

We Exaggerate Rare Risks

Many people fear flying even though it is safer than driving a car²⁵. This is because we tend to underestimate common risks. News stories of flaming airplane wreckage feature prominently in peoples' association of flying. Part of the problem is that it is precisely the rarity of these events that make them newsworthy. However, the more attention is devoted to these events in news headlines, the larger the risk seems to us. In fact, air travel has become progressively safer over the years.

Comparing the numbers of commercial air disasters with the numbers of data breaches reveals increasing security in air travel, and decreasing security for personal data. Yet 'having your data stolen' is not a fear that people hold, despite breaches becoming so commonplace that they rarely make the front page. Given the number of known records breached—and allowing for some unknown breaches—every one of us has had our private data breached multiple times.

The Unknown Is Feared More than the Familiar

We tend to trust people or things we know rather than those we do not know. System administrators do not patch known vulnerabilities for fear of introducing instability in their systems. Additionally, adopting new technologies is delayed in preference for more familiar legacy technologies. It is for this simple reason that phishers target users with emails that imitate trusted senders.

Personified Risks Are Given Priority Over Anonymous Risks

We struggle to accept risks when they are just abstractions. This is the reason why faceless attack groups—as well as hurricanes—are given names. It becomes more urgent when you know that Samurai Panda or APT4 is after you than some obscure Chinese officer.

We Underestimate Risks in Situations Where We Feel in Control

When we willingly adopt a risk posture, we tend to underestimate it. People feel in control when they have just deployed a new firewall, some magical virtualization technology, or even a visibility solution. A CISO may think they are in control because they just deployed the latest state-of-the-art EDR. This can lead to seriously underestimating adversaries and their ability to get around these measures.

We Misjudge Objects When We Have Poor Visibility

In behavioral psychology, it has been found that people with poor vision tend to think that objects are farther away than they really are. The same happens when security operators have poor visibility into the compromises in their network infrastructure. In these cases, it is assumed that the risk of compromise is more remote than it actually is.

Overcoming Biases

How can we align our perception of security with its reality? How do we know if the proper amount is being spent on security and spent effectively? It's important to realize that we are all susceptible to biases. However, the first true step towards achieving this is arming ourselves with the facts—and keeping these facts updated.

$$R_s = C(P_0 - P_s)$$

R_s: Return of investment of a given solution s

C: Cost of a breach for my organization

P₀: Probability of a breach in a given time frame, with the current posture

P_s: Probability of a breach in the same time frame, adopting the solution

The first fact that needs clarity is the cost of insecurity. A clear understanding of the cost of a breach forms one part of the equation that tells you if your security trade-off is balanced. This used to be a difficult number to quantify, but each year brings better reporting²⁶ that helps you understand the consequences for your industry, company size, and geographic region.

The second critical fact is your business' individual risk of a breach. Lumu's Continuous Compromise Assessment was developed to determine your organization's real-time factual level of compromise. The result of this process is a baseline for your cybersecurity architecture. This metric informs those big strategic decisions like "Are my security tools delivering on their promises?" and "Where do I need further investment?"

The Cost of Insecurity

As we have stated before, investment decisions are transactional. An investment has to be justified by its return. In cybersecurity, the return is the costs associated with the breach that is avoided by the investment. It has been noted that "difficulties in measuring the costs and benefits of information security investments cloud the vision of the rational decision-maker."²⁷ However, with each year better information regarding the cost and frequency of breaches becomes available through a range of reputable resources. It has become pivotal for cybersecurity operators to acquaint themselves with the real cost of insecurity in order to make an informed decision.

What Motivates Attackers?

Cybercrime is big business. A report by Atlas VPN²⁸ estimated that cybercrime generates \$1.5 trillion annually. The largest component—\$860 billion—of this total comes from illegal online trading. The selling of trade secrets and intellectual property theft accounts for another \$500 billion. Trading stolen data—anything from credit cards to birthdates—generates another \$160 billion. A further \$1.6 billion is made by selling crimeware or Crimeware-as-a-Service. While individual ransomware attacks provide great returns for threat actors and cause extensive damage, it 'only' accounts for \$1 billion of the total revenues of cybercrime.

State actors are driven by more than profit motive. These might conjure up images of strategic attacks like those we have seen carried out against nuclear centrifuges or election meddling. However, private citizens are also at risk. The Equifax breach that exposed the personal data of nearly half of all Americans were believed to have been carried out by Chinese spies for the purposes of espionage.

How Are Attackers Getting In?

It should be no surprise that as in previous years, the most common method of entry for breaches was hacking/intrusion²⁹. This category, accounting for 39% of all breaches, includes breaches through phishing, ransomware/malware, and skimming. The second-largest category, unauthorized access (37%) continued its growth trend from 2018, largely due to the increased prevalence of credential stuffing. The remaining 24% of compromises resulted from employee negligence, accidental exposure, data on the move, physical theft, and insider theft.

How Long Are They Avoiding Detection?

The average time to detect a compromise increased to 207 days in 2020³⁰. A further 73 days were required to contain these threats. Interestingly, these figures varied greatly depending on their region or industry. For example, German organizations required 160 days to identify and contain compromises, compared to 380 days in Brazil. Financial and banking organizations performed somewhat better than most, requiring 233 days while healthcare providers performed worst, requiring 329 days.

What Are They Getting Out?

The number of breaches increased in 2019 and so did the number of records exposed. In total, 870 million records were exposed, of which 165 million are considered to be 'sensitive records'. Financial institutions were attackers' main source of sensitive records, accounting for 101 million exposed records.

The Impact and Cost

Cybersecurity spending has increased by 44% since 2014, and yet we continue to see an increase in the number of breaches and records exposed. In 2019 the number of breaches increased by 17%. The impact of each breach also increased, especially in the USA, where the average cost of a breach amounted to \$8.64 million, more than double the global average.

From the data, it is clear that no industry is safe from breaches. Even the industries that were fastest to detect and contain compromises were still unacceptably slow. Industries subject to the most stringent regulations are failing to protect sensitive data. Complying with the minimum demands of regulators or comparing yourself with industries that are faring worse, is far from enough.

Despite the direct correlation between dwell time and ransomware attacks, the time required for compromise detection is only increasing. Threat actors are constantly evolving their tactics, techniques, and procedures to ensure better deliverability. There

needs to be a tactical and mindset change if strategists and operators are going to be able to turn around the hard reality our industry is up against.

Better Data Makes Better Decisions

Whether making quick tactical decisions or longer-term strategic ones, acting upon good information always aids the process. Let's look at some of the qualities this information needs.

Timely and Up-to-Date

Being able to make decisions quickly requires access to the newest information. Lacking information can lead to uncertainty and delays. Delays, in turn, can lead to growing doubts and more ineffective decision making.

Consistent Quality

Comprehensiveness should not come at the expense of quality. An example would be the prevalence of false alarms. Low-quality alerts cause alert fatigue and security operators to ignore alerts, as in the case of the boy who cried wolf. Alerts can only achieve certainty in response to known attacks with documented techniques and assets. Novel attacks will have to be represented by anomalies that require investigation. However, the investigative burden can be eased and alert fatigue lessened by improving the orchestration between alerts and investigating teams, and by providing contextual information.

Greater Visibility

As with poor eyesight, poor network visibility leads to errors in judgment. Greater network visibility helps to understand the main output of a cybersecurity system: its level of compromise. This level of compromise is crucial feedback information that can inform where additional investment is necessary in the system and tell you if investments are performing according to their promise.

Support Taking Action

Having too many options exacerbates delays in decision making. We frequently spend a lot of time trying to choose the best option. Paradoxically, it can be best to make a good choice, and then commit the resources that it needs to become a great choice in retrospect. However, to do so requires that the initial choice was made based on accurate intelligence and that the necessary resources are available for its follow up.

If Not Now, Then When?

Almost every report written in cybersecurity ends with a high sense of urgency. It is the nature of the industry we are in. This one is no different. There is no doubt that it is time for a breakthrough in cybersecurity. The current state is simply unsustainable. This is not meant to sell the famous fear. On the contrary, the solution is largely in the hands of practitioners. In addition, because a lot of missing pieces have recently fallen into place: cost of storage, computing power, cloud infrastructures and functional machine learning.

As an industry, it is our responsibility to respond to the democratization of cybercrime with the democratization of cybersecurity. The tools for executing sophisticated attacks are readily available. This also means that advanced and efficient technology is available to all who want them, and those who dare to put them to work.

The challenges facing cybersecurity might seem complex and daunting. However, success in complex scenarios lies in the system's ability to regulate from disturbances. With small but deliberate changes to the cyber-security architecture, the disparity between the cyber incident and the detection of the breach can be dramatically shortened.

Individual breaches and the systemic risk they represent can only be contained if we actively look for compromise, and make this a foundational component of our security testing frameworks and strategies. This fundamental shift is in our hands to execute, and there is no time to waste.

References

- 1 2020 Roundup Of Cybersecurity Forecasts And Market Estimates. <https://www.forbes.com/sites/louisacolumbus/2020/04/05/2020-roundup-of-cybersecurity-forecasts-and-market-estimates/?sh=7ec29f4c381d>
- 2 2019 Identity Theft Resource Center 2019 Data Breach Report. <https://www.idtheftcenter.org/2019-data-breaches/>
- 3 Assumption of Breach: The New Security Paradigm by Jeffrey Carr.
- 4 The Cyber Kill Chain, developed by Lockheed Martin. <https://www.lockheedmartin.com/en-us/capabilities/cyber/cyber-kill-chain.html>
- 5 Encyclopedia Britannica. Aristotle | Biography, Contributions, & Facts. https://es.wikipedia.org/wiki/Enciclopedia_Brit%C3%A1nica
- 6 How Stuff Works: Cell Theory. <https://science.howstuffworks.com/innovation/scientific-experiments/scientific-method4.html>
- 7 Security boulevard: NEW TECH: 'Network Traffic Analysis' gets to ground truth about data moving inside the perimeter. <https://securityboulevard.com/2019/04/new-tech-network-traffic-analysis-gets-to-ground-truth-ab>
- 8 Pradeep Menon: An Executive Primer to Deep Learning. <https://medium.com/@rpradeepmenon/an-executive-primer-to-deep-learning-80c1ece69b34>
- 9 Betteridge's Law of Headlines. <https://whatis.techtarget.com/definition/Betteridges-law-of-headlines>
- 10 Bruce Schneier (2000) - Secrets and Lies: Digital Security in a Networked World.
- 11 Cisco IOS NetFlow. <https://www.cisco.com/c/en/us/products/ios-nx-os-software/ios-netflow/index.html>
- 12 BlueAngelHost: Bulletproof Hosting. <http://www.blueangelhost.com/blog/bulletproof-hosting/>
- 13 Krebs On Security: Meet the World's Biggest 'Bulletproof' Hoster. <https://krebsonsecurity.com/2019/07/meet-the-worlds-biggest-bulletproof-host/>
- 14 Simon, Herbert (1957). A Behavioral Model of Rational Choice, in Models of Man, Social and Rational: Mathematical Essays on Rational Human Behavior in a Social Setting. New York: Wiley.
- 15 Sterman, John D. (2000). Business Dynamics: Systems thinking and modeling for a complex world. McGraw Hill. https://en.wikipedia.org/wiki/Business_dynamics
- 16 Kahneman, Daniel (2011). Thinking, Fast and Slow. London: Penguin Books. <https://archive.org/details/thinkingfastslow0000kahn/page/14>
- 17 March, James G.; Shapira, Zur (1987). Managerial Perspectives on Risk and Risk Taking. Management Science. <https://semanticscholar.org/paper/bed01b90c5f0b03cd60b20b8ace2ba56c1b2f942>
- 18 Jalali MS, Siegel M, Madnick S. Decision-Making and Biases in Cybersecurity Capability Development: Evidence from a Simulation Game Experiment. Journal of Strategic Information Systems 2018. <https://scholar.harvard.edu/jalali/publications/decision-making-and-biases-cybersecurity-capability-development-evidence>
- 19 Tech Crunch: Equifax breach was 'entirely preventable' had it used basic security measures, says House report. <https://techcrunch.com/2018/12/10/equifax-breach-preventable-house-oversight-report/>
- 20 Hackerone: Server-Side Request Forgery (SSRF). <https://www.hackerone.com/blog-How-To-Server-Side-Request-Forgery-SSRF>
- 21 CSO Online: Marriott data breach FAQ: How did it happen and what was the impact?. <https://www.csoonline.com/article/3441220/marriott-data-breach-faq-how-did-it-happen-and-what-was-the-impact.html>
- 22 Krebs on Security: What the Marriott Breach Says About Security. <https://krebsonsecurity.com/2018/12/what-the-marriott-breach-says-about-security/>
- 23 TechCrunch: Marriott's breach response is so bad, security experts are filling in the gaps — at their own expense. <https://techcrunch.com/2018/12/03/marriott-data-breach-response-risk-phishing/>
- 24 Schneier, B. (2010, October) The Security Mirage. https://www.ted.com/talks/bruce_schneier_the_security_mirage
- 25 Forbes, Ricardo Villadiego: Aiming For The Sky: Cybersecurity Lessons From The Airline Industry. https://www.ted.com/talks/bruce_schneier_the_security_mirage
- 26 Lumu Technologies: 2020 Compromise Flashcard, 2020 Ransomware Flashcard. <https://lumu.io/resources/2020-compromise-flashcard/> <https://lumu.io/resources/2020-ransomware-flashcard/>
- 27 Chai et Al, 2011: Firms' information security investment decisions: stock market evidence of investors' behavior.
- 28 AtlasVPN: Cybercrime annual revenue is 3 times bigger than Walmart's. <https://atlasvpn.com/blog/cybercrime-annual-revenue-is-3-times-bigger-than-walmarts>
- 29 ID Theft Centre: 2019 End-of-Year Data Breach Report. <https://www.idtheftcenter.org/2019-end-of-year-data-breach-report-download/>
- 30 IBM: Cost of a Data Breach Report 2020. <https://www.ibm.com/security/data-breach>



Illuminating threats
and adversaries

www.lumu.io

Lumu Technologies Inc. | 8350 NW 52nd Terrace Suite 301, Miami, FL 33166 | info@lumu.io | +1 (877) 909-5868