KOVRR
Cyber Risk Modeling

Cyber Catastrophes Explained

Towards a definition for the insurance industry

SEPTEMBER 2020

Executive Summary

In this paper, we will introduce the definition of a cyber catastrophe as used in Kovrr's modeling methodology. We hope this will contribute to the ongoing debate within the insurance industry, provide clarity around what we view as the most costly scenarios for the industry, and ultimately fulfill our role by enabling efficient risk transfer using this definition in our models.

When developing our definition we sought to understand the specific vulnerabilities, tools, and techniques leveraged in a cyber event. We looked at a range of cyber incidents that occurred in the past and identified the principal characteristics that led to large accumulations of loss. In doing this, we realized that we needed to address the misalignment of the term "event" by cybersecurity and insurance professionals.

To overcome this challenge, we have analyzed the way the cybersecurity industry classifies attacks and have developed a novel methodology of grouping companies affected by the same attack as a way of defining what constitutes a cyber event for insurance purposes.

The resulting definition of cyber catastrophe is:

- An infrequent cyber event that causes severe loss, injury, or property damage to a large population of cyber exposures.
- A cyber event that starts with a disruption in either a service provider or a technology, and unfolds by replicating this disruption whenever possible.

A cyber catastrophe related to technology failure can be broken down into three main phases: expansion, remission, and transition. The development of a cyberattack through these three phases can typically span from one to six months. The timeframe discussed can be a good basis for an hours clause for cyber.

One of the benefits of using this definition of cyber catastrophe is that it avoids having to take attribution into consideration. It should also be noted that this definition applies to all types of policies, regardless of whether cyber coverage is affirmative or silent.

The damage caused by service providers and technology events can be described by implementing an impact based modeling approach, which links the effect of a cyber event with its financial implications. Keeping the impact at the center of the model allows it to focus on the most important aspect of cyber insurance, the effect of the cyberattack on insureds.

Introduction

The systemic nature of cyber risk poses a significant threat to the insurance industry, with the potential for one event to accumulate losses from large portions of one or more portfolios. For years the industry has been tackling the challenge of managing cyber risk, but arguably there is still some inconsistency in the basic definitions of "cyber event" and "cyber catastrophe."

For natural hazards, clear definitions have been designed to guide the industry in building models that can accurately assess and model frequency, severity, and financial losses caused by natural disasters. Many of these modeling techniques have also been applied to modeling cyber catastrophes, however, the definition of cyber catastrophe is more complex. These complexities introduce new parameters that must be considered when classifying cyber events and determining the standard for labeling a cyberattack as a catastrophe.

The main challenge in modeling insurance losses caused by cyber catastrophes is that there is almost no prior experience; this is due to both the rarity of mass-scale cyber events and because cyber insurance is a relatively new product. Therefore, assessing the cost of a cyber catastrophe requires looking beyond insurance claims and researching specific past cyber events which might have caused large accumulations - assuming insurance had been in force at the time.

Defining a catastrophe

The American Academy of Actuaries defines catastrophes as: "Infrequent events that cause severe loss, injury, or property damage to a large population of exposures. While the term is most often associated with natural events (e.g., earthquakes, floods or hurricanes), it can also be used when there is concentrated or widespread damage from man-made disasters (e.g., fires, explosion, pollution, or nuclear fallout)".[1]

The breakdown of this definition sets the following parameters for catastrophic events:

- Low frequency
- High severity
- A large population of affected exposures

While these characteristics are relevant, their lack of any specificity as to the type of peril poses a challenge to their cyber application. The problems arise mainly from the heterogeneous nature of cyber incidents and the difficulty of threat actor attribution.

The variety of causes of loss pertaining to cyber is diverse and can encompass a range of events. Events can range from a severe service provider outage that occurs with no known intent, to a swift attack campaign exploiting a vulnerability in a common software leading to data theft. The service provider outage could affect tens of thousands of businesses that rely on that specific service provider, and the campaign could in a very short time span affect hundreds, if not thousands of systems implemented in millions of businesses, which could lead to a large accumulation.

From a cyber perspective, these two examples are seemingly worlds apart, however from an insurance perspective, both events are low frequency, high severity, and affect a large population.

Attribution is another challenge, one that manifests itself in two different ways. On the one hand, there is the problem of establishing whether two claims from companies in different countries and industries can be identified as arising from the same event.

1. https://www.actuary.org/sites/default/files/files/catmonograph_june01.4.pdf/catmonograph_june01.4.pdf

For example, how can this be applied to two claims for ransomware? One possible way is through attribution, but the key here is that **attribution is not the only parameter we can use to establish a relationship between the two attacks. In fact, the specific vulnerabilities and exposures[2] exploited in the attacks, and the tools and techniques employed for such exploits, are significantly more relevant.**

In our white paper on silent cyber[3] we defined a campaign as follows:

"An attack, or series of attacks, leveraging specific vulnerabilities, tools and techniques."

The concept of a "campaign" allows for an established commonality between two separate and seemingly distant claims, without necessarily answering the question of attribution.

On the other hand, war exclusions may require a sophisticated way to identify not only malicious intent but belligerent intent as well. A recent report by The Geneva Association[4] highlights the need to clarify policy wordings around the concepts of war and terrorism. The report introduces the concept of "hostile cyber activity" (HCA) to include State-sponsored cyberattacks, and makes a distinction among different categories of cybercrime based on motivation and expectation.

Additionally, there is a technical challenge in pinpointing the exact identity of the threat actor behind an attack, as attackers use multiple sophisticated techniques to mask their identity.

	Cyber Crime	Cyber Terror	HCA
Motivation	Money	Money / Intimidation	Intimidation
Expectation	Value	Value / Fear	Fear

These types of classifications allow for establishing the type of attack without necessarily answering the question of attribution.

In 2017, NotPetya, a wiper attack disguised as ransomware, disrupted businesses around the world and became the first cyber event to be officially classified as a cyber catastrophe. It is widely believed that NotPetya originated as a Russian military cyberattack on Ukraine, causing collateral damage in the shape of $10 billion in global losses and, according to PCS estimates, $3 billion in insured losses.

Setting aside the attribution, however, a closer analysis of the payoff from the attack shows a complete disregard for any tangible value it might have generated. Based on this trait alone, NotPetya could be classified as hostile cyber activity, without any reference to Russia or Ukraine. The concept of campaign and the classification of cyberattacks according to their expectation are powerful tools in helping to define a cyber catastrophe.

2. Here we use the terms 'vulnerability' and 'exposure' in the sense understood by the cyber security community, which differs from insurance practitioner terminology. Please refer to https://www.cvedetails.com and to https://nvd.nist.gov/vuln-metrics/cvss for more information.
3. https://www.kovrr.com/resource/cyber-risk-from-peril-to-product
4. https://www.genevaassociation.org/research-topics/cyber/CTCW-common-language

Defining a cyber event for insurance

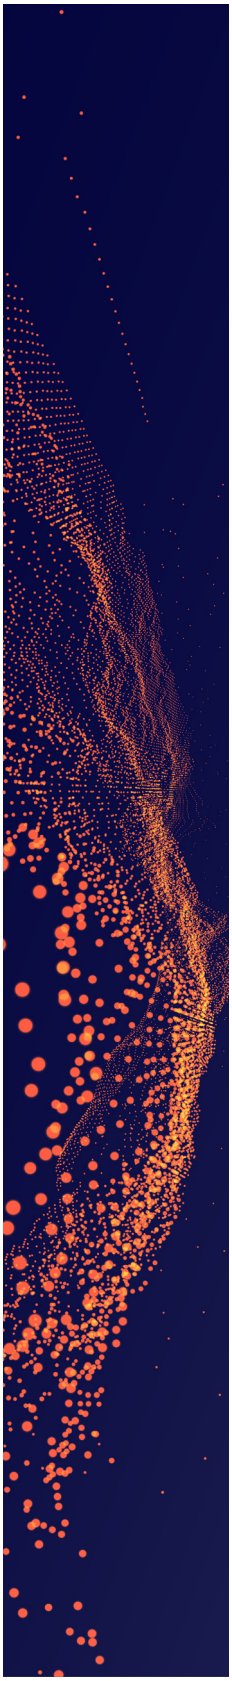

Before further exploring cyber catastrophes, it is worth taking a step back and looking at the definition of a cyber event. The term "cyber event" has a different meaning for a cybersecurity expert and a cyber insurance professional. In order to properly model cyber events for insurance it has to be understood that the same classification as the cybersecurity industry cannot be used. We will discuss below the differences between the definitions, and provide an example.

A cybersecurity expert classifies an event as a cyber campaign that can last anywhere between days to years.[5,6] In this context, a campaign is usually defined in terms of threat actors, attack methods, tools, and payloads. Some campaigns might include several attack methods. An attack method can be repurposed by a different threat actor which will lead to a separate campaign classification. In the WannaCry and NotPetya events both ransomwares used EternalBlue (CVE-2017-0144) as part of their exploitation process.

A cyber insurance expert defines an event in terms of policy wordings and losses, with the concept of an event being developed in parallel with the understanding of how losses can accumulate. For the cyber insurance expert, a cyberattack is just one of many possible causes of loss, because an insured loss can occur with no malicious actor or intent behind it. It is also important to note that the threat actor identity is less of a concern for the insurance expert, possibly excluding the case of a nation-state sponsored attack (which is still subject to debate on coverage). The main parameters considered by insurance experts are the frequency and severity of different losses arising from cyber.

The different viewpoints of the two experts might generate completely different interpretations of the same situation. Let's take, for example, the Emotet cybercrime operation. Emotet was first identified in 2014, and began as a banking trojan with the goal of stealing banking credentials from its victims. The malware continued to evolve over the years, introducing sophisticated techniques to steal data and evade detection by security software. Over time, additional abilities were added, e.g., the ability to transform the victim's machine to a spam propagation server, or data extortion capabilities, and hijacking the victims' data for ransom.[7]

From a cybersecurity perspective, the development of Emotet is classified as one ongoing Emotet campaign, which evolves and adds capabilities as part of the natural advancement of threat actors. From a cyber insurance perspective, an Emotet victim that suffers from asset abuse (the spam propagation repurpose of its machines) is very different from one that will suffer a ransomware extortion attempt.

Besides its shape-shifting nature, another important aspect of Emotet is the amount of time the campaign has survived. The campaign has been ongoing for six years. For cybersecurity experts, it is very convenient to refer to Emotet as one event, because the name alone carries a lot of meaning in terms of attack methods, tools and payloads, but for cyber insurance experts, the time element poses a problem. For insurance purposes, it is conceivable that losses arising from Emotet would not be regarded as one event, even in cases where they are completely homogeneous.

5. https://www.zdnet.com/article/click-fraud-zeroaccess-botnet-rises-from-the-ashes/
6. https://www.clearskysec.com/wp-content/uploads/2020/06/CryptoCore_Group.pdf
7. https://www.malwarebytes.com/emotet/

Defining a cyber catastrophe

A cyber catastrophe is an infrequent cyber event that causes severe loss, injury or property damage to two or more, but typically a large population of cyber exposures.

To fulfill the latter requirement, the following must also happen:

1. **A large population of exposures needs to be in harm's way.**
   An important observation is that for this to happen, a disruption to an important common system or process must occur. The importance of such a system or process will differ, however there must be some common ground, meaning all potentially at risk entities will rely on that system or process for business. From this rather simple observation, we can then classify catastrophic cyber events into two main categories: one category relates to service providers and the other relates to technologies.
2. **The campaign must have the potential to spread rapidly and uncontrollably.**
   Here we use the term campaign in a fairly loose sense, meaning not only a cyberattack but also the sort of human error or omission potentially leading to a cyber incident. Examples of such errors and omissions are misconfigurations, bugs in newly deployed code, or even electricity blackouts.

Below is a list of examples of events that have been studied during the development of Kovrr's Impact Based Modeling Framework and therefore contributed to our definition of cyber catastrophe. The list marks events which have caused significant accumulations:

Key Historical Cyber Catastrophes

Year	Event	Attack Type	Impact
1999	Melissa	Mass Mailing Office Macro Virus	Email servers at corporations and government agencies worldwide became overloaded, and some had to be shut down entirely, including at Microsoft. The attack affected the availability of the email servers.
2000	ILOVEYOU	VBS Script Malware	Disrupted the operations of businesses and government agencies including Ford, Merrill Lynch, the Pentagon and the British Parliament. The attack affected the confidentiality, integrity, and availability of the machines.
2001	Code Red	Computer Worm	Attacked web servers around the world and caused defacement and Denial of Service. The attack affected the availability and integrity of the machines.
2003	SQL Slammer	Computer Worm	Caused a denial of service on internet hosts and dramatically slowed general internet traffic. The attack affected the availability of the servers.
2008	Conficker	Computer Botnet	The worm attacked Windows machines slowing them down and disrupting their work and was present in systems owned by The Armed Forces of Germany. The United Kingdom Ministry of Defense, The French Navy, hospitals and more. The attack affected the availability and integrity of the machines.
2016	Dyn DNS Provider Outage	DDoS	The 2016 Dyn cyberattack was a series of distributed denial-of-service attacks targeting systems operated by Domain Name System (DNS) provider Dyn. The attack caused major internet platforms and services to be unavailable to large swathes of users in Europe and North America. The attack affected the availability of the DNS service.
2017	WannaCry	Ransomware	The attack targeted Windows machines. The attack encrypted hundreds of thousands of computers in more than 150 countries. The attack demanded cryptocurrency in ransom to unlock the files. The attack affected the availability and confidentiality of the machines.
2017	NotPetya	Malware/ Ransomware	The attack targeted Windows machines encrypting data. NotPetya heavily affected supply chain logistics companies such as the shipping giant Maersk, postal company FedEx, and the Port of Rotterdam. The attack affected the availability and confidentiality of the machines.
2018	AWS Cloud Disruption Event	N/A	Parts of Amazon Web Services' US-East-1 region experienced approximately half an hour of downtime. Some customers' instances and data could not be restored because the hardware running them experienced complete failure. The attack affected the availability of the service.
2018	Microsoft Office 365 Outage in EU, Asia, US	N/A	Users from organizations all over the world, including the UK parliament, were unable to login to their email accounts or anything else hosted on Office 365, Microsoft's cloud computing service, for more than 15 hours. The attack affected the availability of the service.

A service provider event is either an outage or a degradation, a disruption at the source (e.g., email provider, cloud provider, DNS provider) causing the service provided to be temporarily unavailable or unreliable (it can also affect the data a client is storing on the service). A technology event is one that affects products and technologies that an insured entity relies upon to do business (e.g., databases, web servers, IoT, networking devices). The impact of both types of events is described by how they affect the confidentiality, integrity or availability of services or data for the insured entity.

The characteristics of a service provider event are comparable to natural catastrophes, mostly because they have a defined time frame which allows for a more straightforward impact analysis. It is important to know which services an entity relies on and the geolocation of each service's associated data center. Additionally, further insights regarding the service type and availability are needed to estimate if the service will suffer from an outage or a disruption. This can be compared to needing to know if a property lies in proximity of the epicenter of an earthquake.

A catastrophe related to a technology carries a layer of complexity beyond a standard natural catastrophe event. That complexity is due to two main characteristics unique to cyber:

- A campaign is rarely confined to a geographical space.
- It can also be unbounded by time.

We can make several useful observations about the spread of an attack:

- Several entities relying on a particular technology does not mean they will all suffer a loss. In fact only a fraction will be affected. This is discussed at greater length in our white paper on cyber black swans.[8]
- For a cyberattack on a specific technology to affect a large population of exposures, there must be some sort of automated propagation mechanism (one of the key elements that affects the virality factor).
- Not all possible attack vectors are equally likely. For example, attacks that rely on 1-click or 0-click exploits (i.e. requiring minimal or no user interaction) are more likely, while attacks based solely on social engineering are less likely, as these require user interaction.

To tackle the deeper question of how time affects a campaign, let's consider an example. Conficker was discovered in November of 2008. The malware was propagated by using a remote code execution (RCE) vulnerability in Windows Server Service (CVE-2008-4250), enabling an attacker to take full control of a windows machine remotely.[9] There was no human interaction required in order to trigger this vulnerability. This enabled Conficker to infect over 9 million devices in a time span of six months. This is considered a very high virality factor. The economic damage attributed to the attack was approximately 9.1 billion USD.[10]

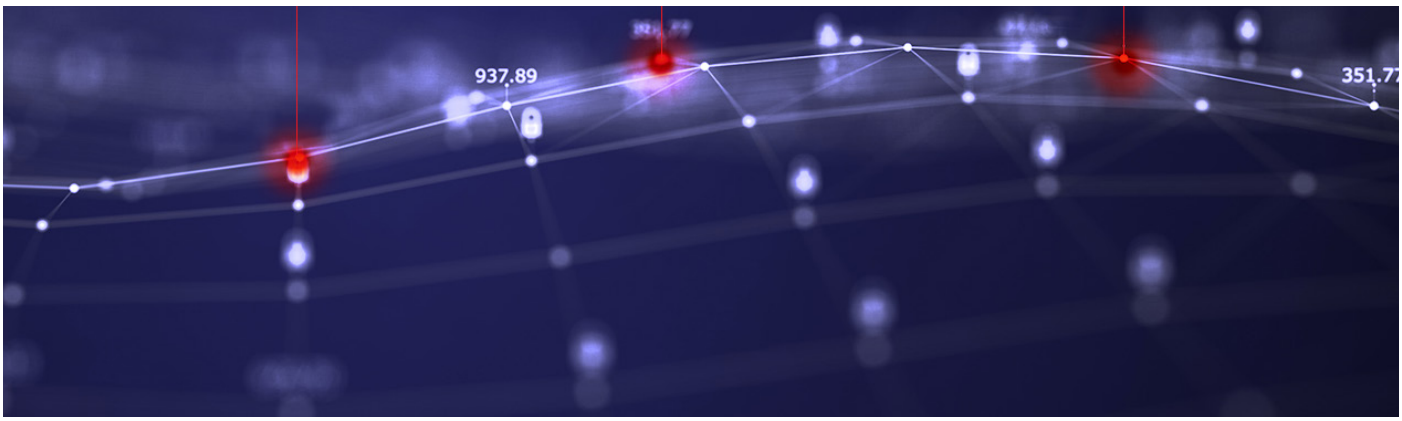

8. https://www.kovrr.com/resource/cyber-black-swans
9. https://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2008-4250
10. https://www.zdnet.com/article/confickers-estimated-economic-cost-9-1-billion/

Conficker's activity continued after the initial infection phase. In 2013, five years after the attack began, Conficker had more than 1 million identified infections. In the following years, Conficker still managed to achieve at least 100,000 infections every year.[9] Although it never became an insurance event, Conficker had all the elements for being one. If the cyber insurance market had been more developed at the time, there would have been insured losses from Conficker. However, the event begs the question, would they have all been accumulated under the same catastrophe code? It is fairly clear that the answer would be no. However, understanding where to draw the line in regard to time is more complicated.

Having classified cyber catastrophes in two categories related to service providers and technologies, it is important to note that these types of events are not mutually exclusive: a technology event might affect a service provider causing a catastrophic outage or a large scale privacy event.

The effect of an hours clause on cyber

When drafting reinsurance contracts, significant emphasis is placed on the definition of an event, and in drawing the boundaries of what losses cedants are allowed to accumulate within one event. Usually the definition of an event, or even more precisely of a catastrophe, is accompanied by what is commonly referred to as an hours clause. This is a section of the wording that limits the time period during which claims from the same event may be accumulated. The time period is usually measured in consecutive hours - 72 (for 3 days) or 168 (for a week). For example, a standard approach to assessing damage due to an earthquake involves 72 hours. Damage caused by aftershocks more than 72 hours after the first quake cannot be accumulated and must be settled separately.

An hours clause can allow what could be thought of as one event to become two, or for two events to become one. One example of the latter are storms Lothar and Martin which hit Europe in 1999 within one day of each other, and thus were referred to as one storm by many reinsurance companies for recovering purposes. Another example would be the Camp Fire, a wildfire that raged in California in 2018 and became the deadliest, most destructive and most expensive wildfire in the state's history in terms of insured losses. The fire started on November 8th, 2018 and was fully contained only after seventeen days on November 25th, 2018.[10] With most hours clauses set at 7 days, many cedants were able to claim multiple times on their reinsurance contracts, effectively splitting Camp Fire into two events.

For cyber there is currently no standard definition of a catastrophic event or hours clause. However, to parameterize a timeframe for modelling purposes these terms must be defined. After looking at several campaigns, the main traits of technology related cyber catastrophes became apparent and can be summarized using Conficker as the main example.

Firstly, viral campaigns' infection rates are most successful for months, not days. In fact, using a few simplifying assumptions to deal with the presence of multiple strains, in our analysis each Conficker outbreak reached a peak during a period of six months. The main reason for this is the obvious interaction between infection and patching rates, the balance of which takes time to develop.

Secondly, viral campaigns cannot be too successful without exposing the perpetrators. In this respect, Conficker is an extreme example of a viral cyber campaign, one that possibly represents an upper bound of sorts. In fact, considering

11. https://www.trendmicro.com/en_us/research/17/l/conficker-downad-9-years-examining-impact-legacy-systems.html
12. https://www.fire.ca.gov/incidents/2018/11/8/camp-fire/

that the attack enabled full control of a windows machine remotely, Conficker never reached its full potential as a cyber weapon. Armed with a different payload, such as ransomware, its impact might have been much worse. The reasons why it never did are subject to speculation, but experts believe its high profile status and the scrutiny it attracted, prevented Conficker's controllers from taking full advantage of it.

The two observations above allow for a definition of the boundaries of a technology related cyber catastrophe, at least for modeling purposes. In its initial phase, the attack develops rather quickly in order to maximize whatever value its impact on the target allows. There is a second phase where the attack is still successful, however loses momentum, and in the third phase it transitions from a viral campaign into the sort of cyber incident that can happen every year. The development of the attack through these three phases typically spans a time horizon of one month, but can be as long as six months. The time frame discussed can be a good basis for an hours clause for cyber.

Scoping the timeframe of an event is a critical part of the modeling process. By adding this parameter, the appropriate definitions to a cyber campaign can be applied. This allows a modeler to categorize the high impact phases, e.g., of the Conficker cyber campaign as a catastrophic event, and address additional losses as attritional losses.

Impact Based Modeling for Cyber

Kovrr has developed a unique modeling approach based on the key observation that every cyber catastrophe starts with a failure in either a service provider or a technology. Setting aside the possibility for technology events to cause disruption to service providers, cyber catastrophes can be described as events that unfold over a confined period of time during which homogeneous losses accumulate.

While service provider events are easier to define and model, technology events present unique characteristics that call for additional elements to be defined and parameterized:

- Commonality of specific vulnerabilities, tools and techniques leveraged in campaigns.
- A measurement of the virality of a campaign in terms of the popularity of the technology and if human interaction is required for propagation.
- The impact of a campaign and how it affects the confidentiality, integrity or availability of services or data for the target entity.

Based on these three elements, and considering that attribution is not part of the Kovrr framework, the following example fits the definition of cyber catastrophe: several losses arising from ransomware attacks on different city governments in the US. It is conceivable that behind each loss there is a different threat actor, but the methods used are always the same.

The timeline of a cyber catastrophe event is a challenging topic, however, when technology related they can be divided into three phases spanning a period of up to six months: expansion, remission, and transition.

Despite the differing nature of both, the impact of service provider and technology events can both be described using similar methodology. The description is based on how the event affects the confidentiality, integrity or availability of services and data, keeping in mind that a combination of all three is also possible. Keeping the impact at the center of the model has two major benefits.

The first benefit is it allows the model to include the various steps of cyberattacks such as initial infection, persistence and propagation. A second benefit is it allows for the model to focus on the most important aspect of cyber insurance - the effect of the cyberattack on insureds. The next step is to link the effect with cost components, which are the basic building blocks of the loss, and describe each element of the inflicted damage in terms of a monetary amount. Cost components are also connected to different coverages, enabling the flexibility to manage both silent cyber exposure and affirmative cyber exposure, based on the specific product wording and its exclusions.

Kovrr's methodology was developed in order to align cybersecurity expertise with interests of the insurance industry. It aims to solve issues related to definition inconsistencies in the market related to both coverage and events. While the debate on wording continues, Kovrr's approach offers clients and prospects a well-defined framework on which to build their view of risk. We continue to listen and take an active role in the debate, firmly believing that our role in the insurance space is to enable efficient risk transfer for cyber.

KOVRR
Cyber Risk Modeling

The Authors

Marco Lo Giudice, PhD is Head of Pricing Models Development at Kovrr. He has worked in the catastrophe modeling and exposure management fields for thirteen years. Most recently, he served as the Local Head of Pricing at Tokio Millennium Re in the company's UK branch.

Avi Bashan is CTO of Kovrr and leads engineering and research efforts. He has worked in the cyber industry for fifteen years and started his career in the IDF intelligence technology unit. Previously he worked at Lacoon Mobile Security and most recently Check Point Software Technologies. Avi is a lecturer at Bar Ilan University's Business School and holds a B.Med.Sc from the Hebrew University of Jerusalem.

Kovrr's Naomi Weisz and CyDelta's Visesh Gosrani also contributed to this report.

About Kovrr

Kovrr's cyber risk modeling platform delivers global (re)insurers transparent, data driven insights into their affirmative and non-affirmative cyber risk exposures. The Kovrr platform is designed to help underwriters, exposure managers and catastrophe modelers understand, financially quantify and manage cyber risk by utilizing AI-powered risk models.

To learn more please contact the Kovrr team: contact@kovrr.com