Cyber Black Swans

Gaining visibility into tail events when managing cyber insurance portfolios

JULY 2020

www.kovrr.com

In March 2011, a powerful earthquake hit off the coast of Tōhoku, Japan, generating a devastating tsunami that overwhelmed all flood defenses. Up until then, scientists did not expect an earthquake in that region beyond magnitude eight but this specific event exceeded all accepted scientific predictions and expectations with a magnitude nine. The event was unanticipated, caused major financial impact, and called upon scientists to review their understanding of subduction zones. Events like this have come to be known as black swans.

Cyber is a relatively new peril in the insurance landscape; companies have limited experience in underwriting and modeling the risk, and the risk itself has evolved in line with the advances of technology. Moreover, cyber insurance is still a developing market: scope of coverage is not very consistent, and policy terms are evolving rapidly. Against this backdrop, the industry is still interrogating itself about what a cyber black swan might look like, and how much it would cost.

Black swans were first discussed by Nassim Nicholas Taleb in his 2001 book *Fooled by Randomness*, which aptly concerned financial events. His definition was based on three main characteristics: unexpected; causing a major impact; and most importantly, explainable, even though only in hindsight. Black swans are particularly undesirable events in the financial sector. Actuaries and exposure managers aim to avoid black swans, or to put it another way avoid unexpected volatility of losses. To be prepared for this kind of occurrence is key not only for an insurance company's survival but also for its success.

Insurance professionals need to be as proficient at understanding cyber risk as they are with other types of risk. The need stems mainly from three forces at play. Firstly, the risk already resides in insurance companies' books in a non-affirmative form, for example claims from cyber events could affect property and casualty policies. Secondly, cyber insurance buyers are becoming more sophisticated and demanding coverage fit for their risk management needs, including limits commensurate with the potential loss. Lastly, since economies with high insurance penetration recover more quickly after a catastrophe, insurance companies have an important role to play in enhancing resilience to large cyber events in the economies where they operate.

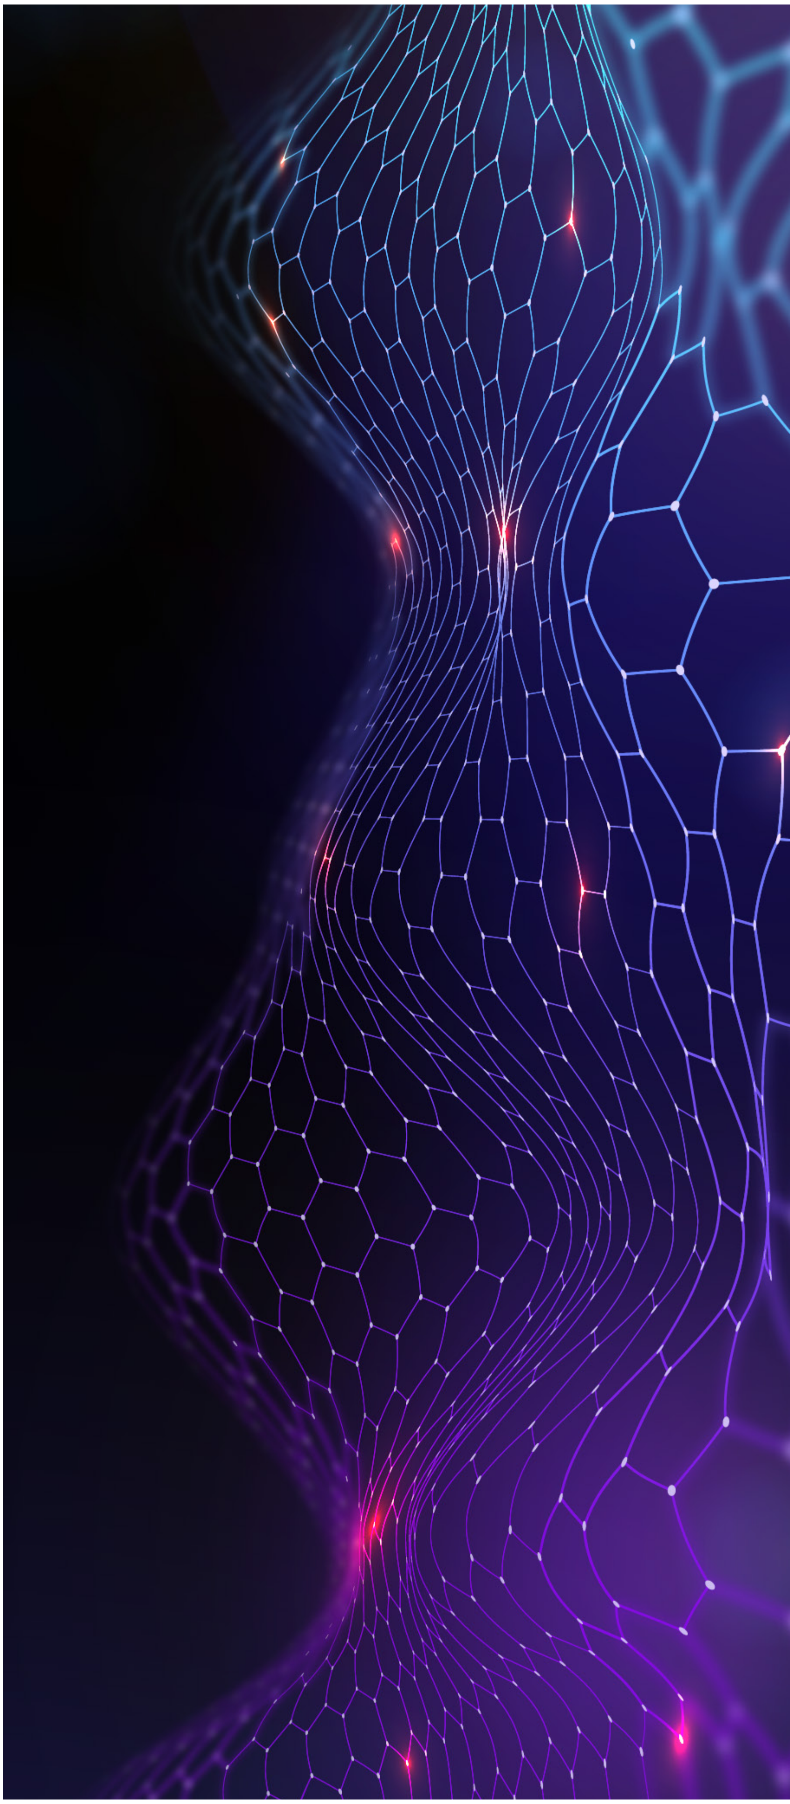

The Footprint of a Cyber Event

An effective solution for managing cyber risk allows practitioners to identify drivers of loss—risks in the portfolio that are most likely to contribute to an event. Solutions need to properly capture the correlation within a portfolio, in order to distinguish which risks will be affected, and to what extent those risks will incur serious financial loss. For natural hazards, correlation is determined by geographic proximity. For example, in an earthquake, the most affected properties will be the ones closest to the epicenter. In cyber, geographic proximity is not enough because events propagate through computer connections.

To better illustrate the problem, let's consider a major bug in a very popular technology. For example, the type of vulnerability that might allow remote code execution, that is the ability for a malicious threat actor to take control of a server or any other endpoint. Millions of businesses, all around the world, are potentially at risk. A campaign exploiting this type of vulnerability will start with the specific aim of maximizing the return for the threat actors involved, meaning an initial target will be identified based on the industry sector and country the attack is most likely to succeed in. All these factors can be modeled, using a combination of game theory and cyber security knowledge—however, pinpointing exactly which company will be targeted first is a challenge.

Often in such cases, several companies are targeted as starting points for the cyber event. Each of these initial targets will be exploited to target several others, in a chain reaction generating the same spread as a pandemic. How fast and wide this chain reaction will go is determined by how many business partners of each affected company rely on the same technology, and by the mechanism of transmission—described by the type of attack vector and the requirements for human interaction. Just as in a pandemic, the speed and reach of the attack can be summarized with a virality factor representing the number of companies each infected company is able to reach. Stronger virality is produced by more popular technology and less requirements for human interactions—for example not needing people to click on links.

However after an initial run of successful attacks, the campaign will become known to the developers of the technology, who will devise a patch to fix the vulnerability. As users around the world begin implementing the patch, the campaign loses momentum. From an initial potential reach of millions of businesses, in the end the campaign may have reached no more than a few thousands, a success rate running to single digit percentage points at most, but better measured as a rate per mille (x successes per thousand tries). In comparison, an earthquake is likely to cause some sort of damage to the majority of properties within its reach. With numbers such as these, determining the footprint of a cyber event, and identifying which companies were affected, is problematic.

Why a Cyber Catastrophe Model?

Currently, insurers are cautious about growing their cyber book because of the dynamic nature of the risk, the relative lack of loss experience and the uncertainty of how a large event might unfold. Given the complexity of the threat from cyber, actuarial techniques are not fully applicable, and many insurers have gone back to basics by managing the risk using deterministic scenarios. These allow exposure managers and insurance executives to better understand the mechanisms of loss, grasp which lines of business might be affected and develop specific affirmative coverages.

Obtaining a full view of risk using deterministic scenarios alone can be challenging. This is because deterministic scenarios have three major limitations:

1. Deterministic scenarios fail to properly capture the correlation within a portfolio
2. Highly prescriptive scenarios fail to allow for more complex types of attacks (one campaign can have both ransomware and data theft traits)
3. The frequency of each deterministic scenario is not a modeled parameter, yet in cyber frequency is the most important parameter as each event is man made

The combined effect of these three major limitations is that deterministic scenarios can be misleading when trying to avoid surprises. Employing scenarios exclusively, leaves too much room for events to unfold in totally unexpected ways.

The aim of any exposure management tool is to enable an estimate of what is commonly known as value at risk (VaR). VaR methods have been used in the financial markets since the 1980s and were introduced to a wider audience when J.P. Morgan published their methodology in the 1990s. In insurance, where the single biggest threat to the balance sheet is a catastrophe event, VaR is usually estimated using catastrophe models. The output is an exceedance probability curve which is then plugged into the capital model, typically developed by actuaries using a Dynamic Financial Analysis (DFA) tool. Although a catastrophe model is not the only tool available for exposure managers, and is not the only tool capable of estimating VaR, it is most commonly used because it provides a very flexible framework for portfolio management.

There are four components in a catastrophe model:

- stochastic event catalog
- hazard module
- vulnerability module
- financial module

The stochastic event catalog is generated using a simulation method and is intended to include all possible events that can occur. It might seem an impossible task to account for every possible event, and in some ways it is. What is most important is to account for every possible impact. For example, there could be a never ending list of reasons for a service provider to experience a one week outage, but in reality the only thing that matters is that a one week outage and its impact is included in the catalog.

The frequency of each event is one of the parameters of the model, and due to the fact that events are generated stochastically, their features and traits are usually more complex than those that have actually occurred. For instance, an earthquake model might allow for very high magnitude events even where there is no historical record of large earthquakes, because the possibility of higher magnitudes for the region is scientifically possible.

Crucially, the one thing catastrophe models are best at capturing is the correlation within a portfolio. By modeling the footprint of each event in the stochastic catalog, a catastrophe model allows exposure managers to identify the drivers of loss, that is those risks contributing the most to the VaR.

The Shape of a Cyber EP Curve

In order to best determine the footprint of an event, it is important to have the correct view of the hazard. Kovrr's modeling efforts are focused on the causes of a cyber event: service providers and technologies. The key observation in our methodology is that every cyber catastrophe starts with a disruption in either a service provider or a technology and unfolds by replicating this disruption whenever possible. Kovrr's view of the hazard is also informed by the service providers and technologies each entity in a portfolio relies upon.

The ability to enumerate all service providers and technologies upon which a company relies is comparable to being able to geocode a property at coordinate level. It's desirable, but not always possible. Our analysis shows that companies of similar size, in the same country and within the same industry tend to rely upon similar service providers and technologies. This observation allows us to map the hazard even when available data on the underlying companies is limited, in a very similar way to the process of geocoding in property. Location, industry and size are the minimal set of data points required to assess cyber accumulations—CRA-Zones™. These three core pieces of information allow for the hazard to be mapped but full enumeration of service providers and technologies ensures the most accurate mapping. Additional data points can also be used to reach intermediate levels of accuracy.

To return to the example of a campaign exploiting a major bug in a very popular technology: coding the hazard in the fashion described above makes it possible to enumerate all risks in the portfolio relying on this specific technology, and thus identify potential targets. After this initial step there remains the challenging task of identifying the risks most likely to suffer a loss, usually a relatively small subset of all the possible companies that could have been affected.

In this example, there are two distributions:

1. The probability of one company (among those relying on the technology) of being attacked and incurring a loss.

This is analyzed by combining hazard parameters, for example, the security posture of the company with event parameters, such as the country and industry where the campaign started.

2. The success rate of the campaign - that is the probability of the campaign affecting many companies.

The second distribution is best described as an event parameter. Instead of trying to pin-point each initial target, assumptions about where the campaign started are coded in terms of CRA-Zones™, and the propagation of the attack is then modeled according to the interconnectedness of different CRA-Zones™.

Determining which risks will incur substantial financial damage and which will not requires the combination of these two distributions. Fortunately, this is a well-known mathematical problem, solved by defining a copula.

Correlation in Cyber Portfolios

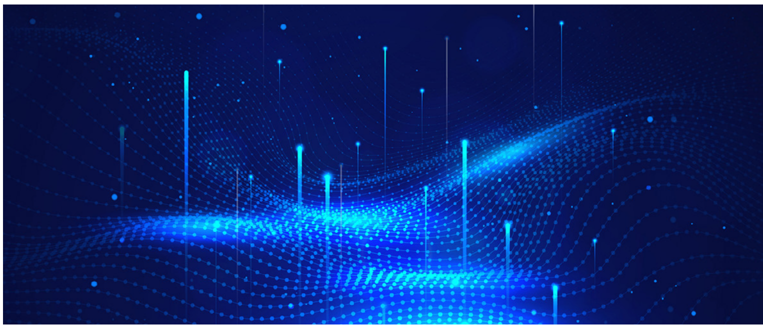

In earthquake modeling, the level of ground-shaking for each property in the portfolio is a function of the epicenter and the magnitude. Using these two parameters, one can calculate the amount of energy released underneath each property, and thus determine the damage suffered by each building. There might however be local variations around this damage, the classic example is two identical buildings standing next to each other, only one of which is destroyed by an earthquake. These types of situations are often used to justify the need for producing probabilistic output by event; but a model that ignores this type of variation in event outcome can still be considered acceptable in natural catastrophe modeling.

A cyber event, however, does not have a complete description of the intensity of impact for each company, even when looking at two companies of the same size, in the same country, relying on exactly the same technology and service providers. In these types of situations, there is still a material probability that at least one of the two companies will not suffer any loss from the event. Therefore, the need for probabilistic output by event is not an option here, it is a necessity. A model that ignores this type of variation in the outcome from a cyber event is inaccurate.

In order to define a probability distribution around a loss outcome, a copula can be used to link characteristics of a single risk (for instance good risk management practices) with characteristics of an event (for example the propensity of a campaign to become viral). In simple terms, a copula is a function describing a way of putting together these two pieces of information. Formally, it is defined as a multivariate distribution where each variable represents one of the risks among those relying on the technology. To illustrate what this means let's consider the case of a portfolio consisting of two risks. The probability of each suffering a loss is known, as is the chance they will be impacted at the same time. A copula is a function that couples (hence the name) the two separate distributions using information about the correlation.

For example, if there is a 1% probability of having a loss of 1M for each risk, a copula allows us to answer the question "What is the probability of having a combined loss of 2M?" To keep it simple, assume the distribution of losses for both risks to be a Bernoulli trial - either there is a loss of 1M (with probability 1%) or there is no loss (with probability 99%). If we believed the two risks were perfectly correlated, that is they suffered losses always at the same time, the answer would be 1%, since either they both had a loss of 1M or they both had no loss. If we assume the two risks to be completely independent the answer would be $(1\%)^2 = 0.01\%$ since the joint probability of independent events is computed by multiplication (or because the joint distribution would now be a binomial). Using a copula allows us to determine where in the range of 0.01% and 1% the combined probability falls.

Drivers of Loss

Value at Risk (VaR) describes the probability of facing a loss of a certain amount or more within a specified time frame. For example, a 1% VaR of 1M means there is a 1% probability the portfolio will suffer a loss of 1M or more over the next year. Tail Value at Risk (TVaR) is a metric used to answer a related question: assuming the business will suffer a loss of at least 1M, what is the most likely amount of such loss? This is often described in terms of return periods, for instance, a 1-in-100 year loss. TVaR is computed as the mean above a threshold. In our example, it is the mean across all modeled losses above 1M. TVaR is an additive metric and therefore best suited to rank the contributions to an overall loss.

A catastrophe model is the best tool for the identification of drivers of loss and TVaR is a common way to portray them. Drivers of loss are sections of a portfolio, or in some cases single risks, identified as being the most likely significant contributors to a catastrophe. It should be noted that different elements of the portfolio are likely to be drivers of loss depending on the risk metric used. For example, a facultative contract with a very high attachment point compared to the rest of the portfolio, is unlikely to be a main contributor to the annual average loss but may well be the main driver of loss in the tail. The ability to identify and compare drivers of loss at different return periods is essential for exposure management. Kovrr's platform provides visibility on the main drivers of loss using additive metrics including Annual Average Loss and TVaR.

In the quest to avoid surprises, it is not enough to have a catalog including all possible black swans, however, it is important to be able to identify which portions of the portfolio are the most vulnerable. VaR and TVaR metrics directly emanate from the correlation assumptions in the model and the shape of the resulting loss distribution, it is therefore essential these assumptions are well understood in order for users to make the most of them. This is yet another clear difference with natural hazards, where the correlation is completely described by the geography of the peril. For exposure managers working in cyber, Kovrr's model offers full transparency on the main drivers of loss and the assumptions behind their identification.

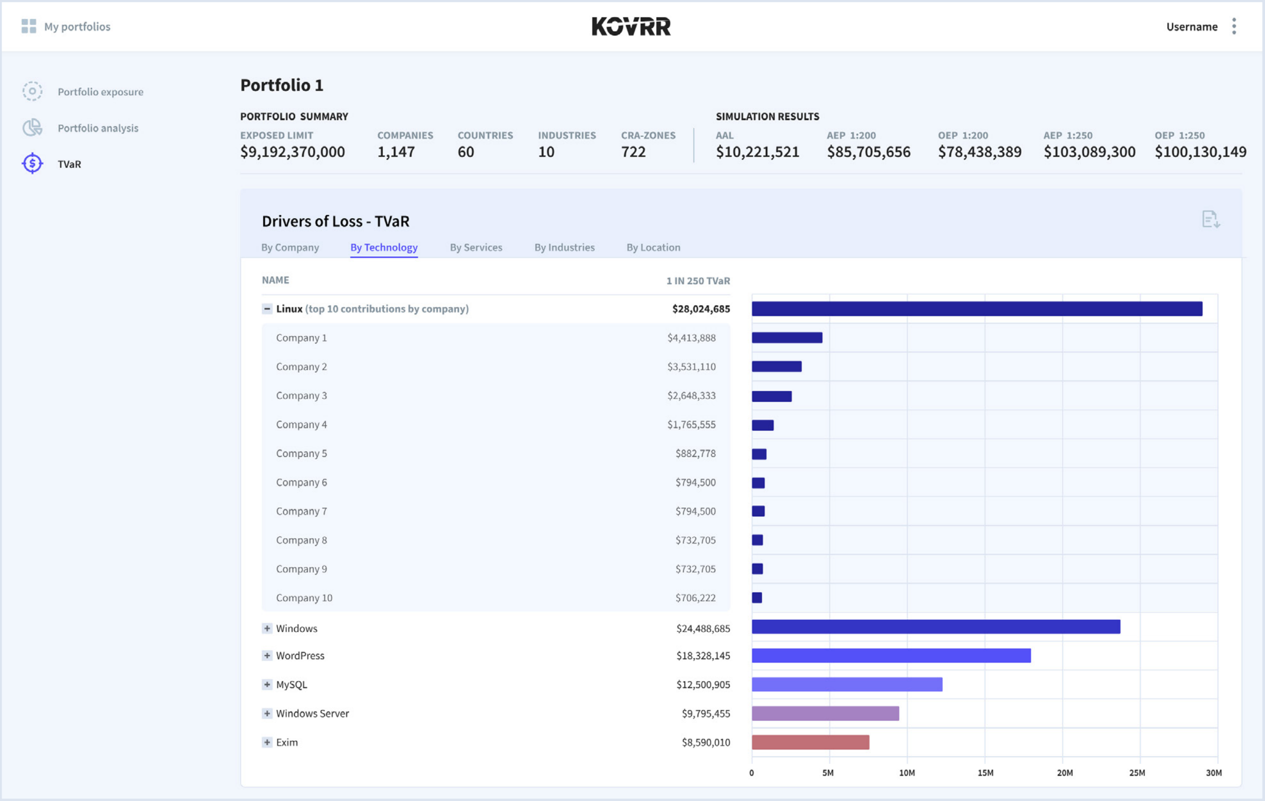

Portfolio 1

PORTFOLIO SUMMARY					SIMULATION RESULTS				
EXPOSED LIMIT	COMPANIES	COUNTRIES	INDUSTRIES	CRA-ZONES	AAL	AEP 1:200	OEP 1:200	AEP 1:250	OEP 1:250
$9,192,370,000	1,147	60	10	722	$10,221,521	$85,705,656	$78,438,389	$103,089,300	$100,130,149

Drivers of Loss - TVaR

By Company | By Technology | By Services | By Industries | By Location

NAME	1 IN 250 TVaR
− **Linux** (top 10 contributions by company)	**$28,024,685**
Company 1	$4,413,888
Company 2	$3,531,110
Company 3	$2,648,333
Company 4	$1,765,555
Company 5	$882,778
Company 6	$794,500
Company 7	$794,500
Company 8	$732,705
Company 9	$732,705
Company 10	$706,222
+ Windows	$24,488,685
+ WordPress	$18,328,145
+ MySQL	$12,500,905
+ Windows Server	$9,795,455
+ Exim	$8,590,010

Kovrr's dashboard reflecting drivers of loss by technology

These metrics are often a good way for users to build confidence in a model. Underwriters and exposure managers have an existing understanding of their portfolio and are looking for tools that provide additional insights. The identification of drivers of loss is a good test for any model because it can confirm pre-existing knowledge while providing additional information on cyber risk.

Conclusion

Cyber is characterized by interdependencies, an evolving nature of threats and high uncertainty with respect to both exposure and hazard data. By taking all these aspects into account, Kovrr's model is best placed to recognize where the potential for extreme events truly lies in a portfolio, providing users with a tool capable of not only identifying those black swans but also managing their impact.

Kovrr's solution enables insurance professionals to make informed decisions based on the ability to drill down from portfolio-level exposures to their respective single risks' drivers of loss and to adjust their underwriting guidelines accordingly. Exposure managers and insurance executives use this information to make decisions about risk transfer, with the goal of reducing volatility and thus avoiding black swans. Identifying drivers of loss is essential for quantifying reinsurance needs, making decisions around diversification in specific markets, and ultimately validating premium adequacy. In turn, Kovrr offers reinsurers the ability to deeply understand their cedants' portfolios, gaining insights which enable them to target support where it is most needed and ultimately deploy capital more efficiently. Going forward the challenge will no longer be finding black swans but managing a flock of them.

KOVRR
Cyber Risk Modeling

The Author

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

Kovrr's Shalom Bublil, Naomi Weisz, Amir Kessler, and Amir Shur 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