Multi-client targeted attacks – the "missing link" of cyber catastrophe models

A CYBER INDUSTRY STEERING GROUP WHITE PAPER
BY AXIS, GALLAGHER, MOODY'S, AND MUNICH RE.

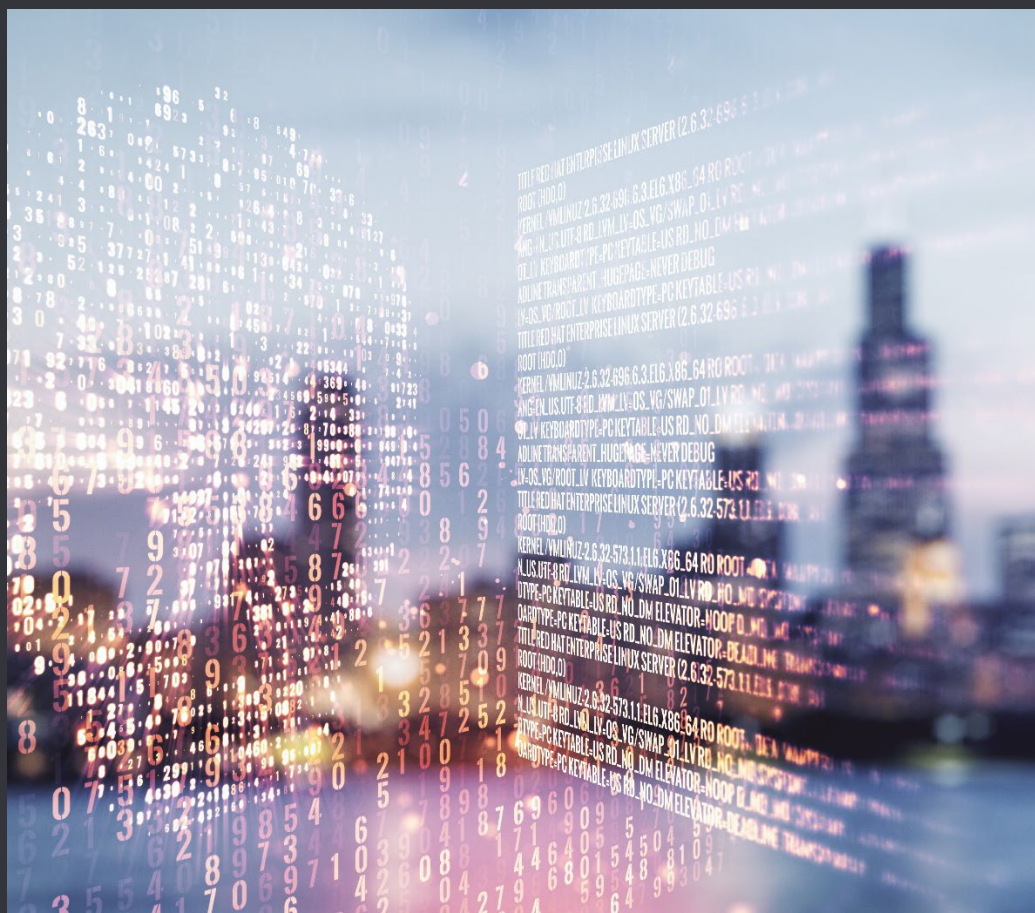

The information contained in this white paper is derived from public sources and general market knowledge. No company-specific price, modelling, customer, or any other information was used to inform this white paper or its conclusions. This white paper is for informational purposes only.

INTRODUCTION

At the core of (re)insurer business lies the ability to understand and quantify potential losses. By quantifying risks assumed and returns earned, (re)insurers can better understand potential impacts, including to solvency (such as bankruptcy risk) and develop strategies to mitigate these risks, improving stability. These decisions to optimize, hedge, and other strategic decisions manifest at a business unit or product level (such as a cyber perspective), as well as at a wider enterprise level.

Inaccurate or incomplete awareness, measurements, and parameterizations of this risk mean (re)insurers have an incomplete view of the risk and may be more likely to make suboptimal business decisions.

(Re)insurers typically parameterize these risks using mathematical models to represent individual risk losses and catastrophe models to represent correlated and systemic risks. However, there is a weakness in how this approach tends to manifest: Due to the intangible correlations in cyber risk, neither the catastrophe and accumulation models nor the mathematical models adequately account for all high-severity, low-footprint events, meaning events that cause a series of large losses but only to a small number of companies with one or more common features. These types of events can happen via many single targeted attacks or from outages and attacks on and via technology and service providers with relatively small customer penetration (or only impacting a small subset of their customers). The latter types are typically covered within cyber catastrophe models; as such, we focus on the first type and refer to them as multi-client targeted attacks (MCTAs).

For some portfolios, especially those focused on particular niches, MCTAs are arguably more destructive than the events typically represented in a catastrophe model. For larger, more diverse portfolios, these events are more likely to occur around the transition between individual risk losses and catastrophe events (meaning lower return period/higher frequency regions of the exceedance probability curve that are often represent a risk to earnings) but could still have a material impact on capital for severe losses in the tail.

The lack of understanding about MCTAs and their parameterization represents a notable gap in the completeness of cyber risk measurement frameworks, increasing the likelihood of incomplete decision-making regarding this risk.

BACKGROUND

(Re)insurers typically think of risk portfolios as being driven by two idealized yet fundamental mechanisms of losses:

(i) Individual, uncorrelated risk losses, which would typically be parameterized by two different distributions:

(a) Attritional — Small-value losses, which have a high frequency and are therefore highly predictable with low volatility

(b) Large — Large-value losses, which have a low frequency of occurrence and high volatility

(ii) Correlated catastrophe and accumulation losses

To accurately weigh risk, (re)insurers must measure both of these mechanisms comprehensively and completely. For classes like property insurance, individual, uncorrelated losses would represent events such as individual building fires and theft. Furthermore, the attritional and large loss models should be well parameterized using historical data, whereas the correlated catastrophe and accumulation losses would be parameterized using an independent catastrophe model, which itself is typically built using a combination of physical modeling of the peril and historical calibration.

Cyber catastrophe models and the industry's awareness of systemic cyber risk are currently focused on two main perils: widespread malware incidents, such as the historic Not-Petya and WannaCry events, and outages of systemically important IT services and technologies, namely cloud outages. These events are by far the most probable drivers of large, capitally depleting events at the insurance industry level.

These events typically impact many companies with low-to-moderate-severity losses for every impacted risk. This then aggregates to large industry losses.

Catastrophe models do not account for every conceivable pairing of footprint and severity; most notably for this paper, they do not include all permutations where a very small footprint coincides with extremely high severity for impacted firms.

There is potential for these small-footprint, high-severity events to have a total dollar magnitude comparable to that of smaller malware or cloud events; as mentioned in the introduction, we call these events multi-client targeted attacks (MCTAs).

For larger diverse portfolios, the impact of these "missing" MCTA events is likely to manifest around lower return period/higher frequency regions of the exceedance probability curve, that typically represent risk to earnings, where portfolio optimization and other strategic decisions are taken, and where regulators, notably Lloyd's, are applying more scrutiny.

For more focused (regional, single industry, or similar) portfolios, these MCTA events will likely have the greatest impact and impact the higher return periods/lower frequency regions of the exceedance probability curve.

THE MULTI-CLIENT TARGETED ATTACK

From a modeling perspective, there are different triggers for MCTAs. These factors are themselves both correlated and entwined, meaning many of them are likely to be at play within a given "event":

Technological:
The most obvious triggers for MCTAs come from vulnerabilities in software or hardware used by insureds. This does not have to be a widely used software or hardware but might also be an industry-, size-, or country-specific piece of technology.

One example of a technological MCTA event is the Snowflake incident in 2024. Snowflake is a service provider used by many companies for storing and analyzing data, and in 2024, many of the organization's customers were breached because they had a specific configuration of Snowflake. Known affected companies included AT&T, Ticketmaster, and Santander Bank.

Technique:
Phishing attacks — in which individuals are tricked into providing passwords, payment information, or personal data to attackers — are a typical access vector. This can lead to business email compromise and the transfer of funds outside the company in question. Although this technique is often viewed as unrelated between organizations, attackers may select

companies with shared characteristics such as processes or business models. This approach helps them sound credible and improves the likelihood of success. These events may be difficult to distinguish from increased attritional activity.

Ideological:

Although those attacked might not share a common vulnerability, there may be something else motivating hackers to attack them.

Examples include attackers with an anti-capitalism background (such as those attacking banks and financial institutions) or who are motivated by an environmental agenda (like attacking car manufacturers or oil and gas companies).

There are also examples across industries in the world of large sports events (such as the Olympic Games and FIFA World Cup) where companies that sponsor these events might be attacked.

Geopolitical:

Other ideological events might be triggered by geopolitical events if companies are attacked for supporting one or another side in a conflict, as we have seen in the war between Ukraine and Russia, or simply doing business in a country caught up in a conflict. Like in the case of "technique" trigger, these events may be difficult to distinguish from increased attritional activity.

Cyber insurance

Cyber insurance has the distinctive potential to introduce correlation to itself. Threat actors realize that those who purchase cyber insurance are likely to be more inclined to pay ransomware and extortion demands. Therefore, threat actors may try to obtain lists of companies that buy insurance by attacking companies in the insurance value chain, with the correlating factor potentially being insured by a particular carrier or utilizing a particular broker.

Success amplification

This isn't necessarily a correlation mechanism in itself but it makes correlation more likely. Here we are referring to success itself. Once threat actors carry out a successful attack against an organization, they are likely to try other similar attacks. A good example of this would be the series of attacks against UK institutions Co-op, Harrods, and Marks & Spencer in 2025.

Reward amplification

Criminal threat actors are usually, as the name suggests, after money, either directly (stealing or extorting money) or indirectly (stealing resources or data that is then sold). Their targeting typically considers the return on investment of the incursion and how they can best monetize their attacks. Like in the some of the other trigger types, these events may be difficult to distinguish from increased attritional activity.

This reward amplification leads to much of the industry-focused clustering we see in cyberattacks. For example, healthcare records containing protected health information are typically seen as the most valuable personal data asset on the black market, and its price will vary over time due to the data's availability and usefulness. This has led to clear waves of attacks focused on healthcare institutions that hold large amounts of this data.

Other industries have experienced similar waves of ransomware attacks, in which attackers stop certain critical business processes, apply maximum leverage to the victim, and urge them to pay the ransom as soon as possible.

NOTABLE CHALLENGES

When modeling all the aforementioned events, there are clear pain points in addition to the usual difficulties in cyber accumulation modeling:

Probability and permutations:

These events are tied to specific vulnerabilities, industries, sizes, countries, and other characteristics, sometimes in combination, meaning the number of plausible permutations is very large; the more granularly we model cyber risk, the more permutations arise.

Typically, catastrophe models use predefined event sets to represent these combinations. Continuing with this approach would require an extremely large event set, which is difficult to parameterize and time-consuming to run.

Sampling bias or the believability of bad luck:

As previously mentioned, these MCTAs could have a focus on individual industries — partially due to threat actors discovering software or process-driven vulnerabilities — with individual industry verticals. However, some industry verticals can be particularly small and only contain a very limited number of companies but be highly correlated due to technology and vulnerability bottlenecks.

For example, if an insurance portfolio is centered on a particular industry vertical, the MCTA concept implies the potential for an event that causes many of the risks within it to be impacted at one time and therefore cause a significant monetary loss. However, the insurer may not accept these as being reasonable, in part because there is inevitably no or minimal prior experience of this. This is not necessarily a challenge in the modeling of the risk but rather in having the risk treated as being real.

CONCLUSIONS

Cyber catastrophe models are key tools for (re)insurers to measure and understand the degree of cyber insurance risk they have assumed. These tools are still evolving and cannot yet be considered exhaustive, which means gaps remain in their ability to capture all relevant risk drivers. This itself should not be too surprising considering even in natural catastrophe risk, the range of perils being modeled continues to increase. That is, despite natural catastrophe modeling being a roughly 30-year-old discipline, flood and wildfire modeling are a much more recent introduction. The clear difference here is that most natural catastrophe models describe individual perils (with the exception of some perils such as severe convective storm, which encompasses tornado, hail, and straight-line wind), whereas cyber models attempt to model the insurance product and therefore the entire peril range.

This paper highlights the importance of addressing gaps in current risk measurement frameworks, particularly the omission of MCTAs. These events can arise from myriad correlating factors, such as shared technological characteristics, exploitation techniques, and attacker motivations. These factors (or the strength of the correlation) can be amplified both by the attacker's success in a particular incursion and by the reward they gain from the attack.

Depending on portfolio composition, MCTAs may have a substantial impact on modeled results since they are typically underrepresented at lower return periods represent a potential substantial impact on niche portfolios at both shorter and further-out return periods. They can also potentially represent a less substantial impact for large, diversified portfolios.

To address this, modeling approaches that incorporate these overlooked scenarios would promote a more complete and resilient view of cyber risk. This in turn would enable (re)insurers to make better-informed decisions and mitigate risks effectively.

The Insurance Steering Group is actively collaborating across the cyber insurance value chain and continues to invest time and resources into identifying and proposing potential approaches to these issues. This is the first in a series of papers describing problems like this, with forthcoming papers more deeply contextualizing, qualifying, and quantifying these and other problems.

ACKNOWLEDGEMENTS

This white paper was developed by AXIS, Gallagher, Moody's, and Munich Re as part of the Cyber Industry Steering Group initiative. The authors wish to thank the broader ISG membership for their thoughtful review, constructive feedback, and insights shared during steering group meetings, which have materially strengthened this work. Although this paper reflects the perspectives of the named authors, it has been informed by the collective expertise of the ISG community.

What is the Cyber Industry Steering Group?

The Cyber Industry Steering Group is an industry initiative launched by Moody's, alongside cyber market participants, to help facilitate the growth of the global cyber insurance market.

It brings together leading members of the insurance community — including (re)insurers, brokers, technology partners, and other stakeholders — to advance understanding of the cyber threat landscape, build resilience, improve analytics and modeling, and attract new sources of capital to meet increasing customer needs.

This group believes that by advancing the industry's understanding and quantification of cyber risk, this will help to attract capacity and support the market's growth.

This group is conducting important primary research, developing improved cyber risk assessment tools, and leading a series of market education initiatives.

Find out more about the Cyber Industry Steering Group here.

The ISG is a cyber industry steering group comprised of stakeholders collaborating with the collective goal of supporting the sustainable growth of the cyber insurance market. The ISG members did not establish any legal partnership, joint venture or similar for the purposes of the ISG or producing this paper and no ISG member is constituted the agent of another or otherwise authorised to act on another ISG member's behalf. Without limiting the foregoing, no ISG member shall have any responsibility for any act or omission of any other ISG member.

Authors

For the purposes of this white paper, any and all reference to the authors shall mean:

AXIS
AXIS Capital Holdings Limited is a publicly listed global specialty insurer and reinsurer organized under the laws of Bermuda, with its principal executive offices at **29 Richmond Road, 3rd Floor, Pembroke HM 08, Bermuda**. **AXIS Capital Holdings Limited** is subject to supervision by the Bermuda Monetary Authority, which acts as the Group Supervisor for the AXIS group of companies. AXIS conducts insurance and reinsurance business through various licensed subsidiaries in multiple jurisdictions and, in certain countries, including the United States, AXIS insurance and reinsurance entities may operate subject to applicable regulatory restrictions.

Stephen Gibson
Stephen.gibson@axiscapital.com
+44 74 6871 8228

Moody's
Moody's Analytics is a provider of data, analytics, research, software, and advisory solutions for risk management and decision making and operates as a business unit of **Moody's Corporation**, a company organized under the laws of the **United States**. Moody's Analytics operates through licensed and authorized entities in multiple jurisdictions worldwide and provides analytics and advisory services to financial institutions, insurers, corporates, and public sector market participants. Moody's Analytics maintains its principal offices in the United States at **7 World Trade Center, 250 Greenwich Street, New York, NY 10007**.

Christopher Vos
christopher.vos@moodys.com
Matthew Harrison
matthew.harrison@moodys.com

Gallagher
Gallagher Re is the global reinsurance broking and advisory business of Arthur J. Gallagher & Co., a company organized under the laws of the **United States**. Gallagher Re operates through licensed and authorized brokerage entities in multiple jurisdictions worldwide and provides reinsurance placement, advisory, and analytics services to insurance and reinsurance market participants. **Gallagher Re** maintains its principal offices in the **United Kingdom at The Walbrook Building, 25 Walbrook, London EC4N 8AW.**

Simon Heather
simon_heather@gallagherre.com
+44 (0)7793047511

Munich Re
Münchener Rückversicherungs-Gesellschaft Aktiengesellschaft in München ("**Munich Re**," Commercial Register Munich Number: HRB 42039) is supervised by the German Bundesanstalt für Finanzdienstleistungsaufsicht (BaFin). **Munich Re** is a reinsurance company organized under the laws of **Germany** with its Registered Office at Koeniginstrasse 107, 80802 Munich. In some countries, including in the United States, **Munich Re** holds the status of an unauthorized reinsurer.

Dr. Stephan Brunner
Sbrunner@munichre.com
+49 89 38914799